

HEX [PI] and HEX6-BLE: Research Overview for Academic Review

Oct 5, 2026 · Anna Maria Sarkisyan · Evrica Technologies Inc.

Public version. HEX [PI] is a probabilistic signal-recognition method with measured out-of-sample results; HEX6-BLE is a proposed sensing server and optical mesh that has not yet been built. This overview states what has been measured, what is hypothesis, and where collaboration is sought. Implementation detail is available separately under a non-disclosure agreement.

Summary

HEX [PI] (software, version 7.0) estimates the probability of a yes-or-no event in a time series from the series' own history. It first states the null: what the series' own record already implies. It then learns, one outcome at a time, whether the place of the present value among its recent surroundings changes those odds. It reports each probability next to the null, is scored only on what it adds to the null, and says so when a read has no verified edge. It reports no confidence score.

In a one-pass replay (280,252 calls on 43 series, every call made before its outcome, every series started from nothing), HEX [PI] added information on whether the next day's move is larger than usual. On 17 market series from 2016 to 2026 the gain over the null was +0.0043 bits per call (95% interval 0.0030 to 0.0056), with 57.9% of calls right against 57.2% expected from the null alone. On ten currencies never used in development it was +0.0082 bits per call (0.0060 to 0.0106). On daily maximum temperature at three stations never used before, it called next-day direction right 61.0% of the time against 54.6% expected. The one market read with no verified edge is next-day direction (51.4% against 51.2%, interval spanning zero), consistent with the weak-form efficient market hypothesis.

The market edge is small, and the 17 market series were used during development, so that run is a backtest. The currencies and temperature stations were fixed by a rule written before any run and are the independent evidence. Live calls, recorded from the release of version 7.0, are the third test.

HEX6-BLE (hardware) is a proposed low-cost sensing and control server for homes, greenhouses and farms. It reads physical inputs the way the software reads any series, through optical channels and environmental sensors, and includes one experimental

element, a fungal memristor. Servers link into a mesh over light and sign their readings to a tamper-evident ledger. No HEX6-BLE unit has been built.

This overview is for academic review. It separates what has been measured from what is proposed, and lists the questions on which Evrica Technologies Inc. seeks institutional collaboration.

Status of claims

Claim	Class	Evidence
On next-day move size in markets, HEX [PI] adds information beyond the series' own record	Measured	One-pass replay, 17 series, 46,478 calls: +0.0043 bits per call, 95% interval 0.0030 to 0.0056
The move-size result holds on series never used in development	Measured	10 currencies, 26,540 calls: +0.0082 bits per call, interval 0.0060 to 0.0106
On next-day market direction HEX [PI] has no verified edge, and reports that	Measured	46,478 calls: 0.0000 bits per call, interval -0.0007 to 0.0006
HEX [PI] adds information on next-day direction of daily maximum temperature	Measured	3 stations never used before, 29,408 days: 61.0% right against 54.6% expected; +0.047 bits per call
Learning within a series adds to the result	Measured	Against a copy frozen after its first 500 calls: +0.0036 bits per call, interval 0.0025 to 0.0048
Version 7.0 improves on version 6.5 on market move size	Measured	Same calls: +0.0037 bits per call, interval 0.0025 to 0.0050
The replay result holds on live calls	Not yet measured	Every call is recorded before its outcome from the release of 7.0
HEX [PI] beats a standard volatility model (e.g. GARCH)	Not tested	Open question

Claim	Class	Evidence
HEX [PI] 7.0 beats a general-purpose language model on the same calls	Not tested	Open question
Dried shiitake mycelium behaves as a memristor	Published by others	LaRocco et al., PLOS ONE, 2025
A fungal memristor adds forecasting signal in a HEX6-BLE server	Hypothesis	Test T6
HEX6-BLE scheduling lowers a building's grid energy and peak demand	Hypothesis	Test T5
Servers hold synchronisation over a short optical link	Design target	Test T4

Only the rows marked Measured are results of this work. The other rows should not be cited as findings.

HEX [PI]: how it works

This section describes the method in principle. Equations, parameter values and source code are in the confidential technical documentation.

1. **Events.** For each series the engine estimates two yes-or-no events for the next step: whether the change is positive, and whether its size is larger than usual for that series.
2. **The null.** Before anything else the engine states the series' usual odds for each event: how often it has happened, tracked over several lengths of memory, each weighted by how well it has predicted. For direction it also counts where the present value ranks among all earlier values. The null never makes a call. It is the reference every call is scored against.
3. **Position.** The present value is ranked among its own recent values at several window lengths, and so is the size of the latest change. Ranks use no units, so the same code runs on any ordered series.
4. **Learning.** The engine learns, one outcome at a time, how much each rank shifts the odds away from the null on that series. Every such weight starts at zero, so on a new series the engine starts exactly at the null.

5. **Reversible influence.** The call is a weighted average of the null and two learned models. A model gains influence by giving more probability to what happened, and influence can return quickly after the series changes.
6. **Strength.** A call has a side and a strength. The strength shown is restated from how often past calls of that strength were right. The side is never changed.
7. **Strict ordering.** Nothing is learned from an outcome until that outcome is known. No call uses information from its own future.
8. **Verification.** Each call is scored by the information it added to the null. A read is presented as having an edge only when that gain, over its whole record, is clear of zero.

The output for each event is a probability, the null beside it, the number of scored steps behind it, and whether the read has a verified edge. A language model turns these numbers into prose for the user and does not alter them.

Each step follows established practice. Scoring every call before its outcome is prequential evaluation (Dawid, 1984) with a proper scoring rule (Gneiting and Raftery, 2007). The weighted averages are Bayesian mixtures (Vovk, 1990) with fixed-share updating (Herbster and Warmuth, 1998). The strength step is isotonic regression, with current theory in Dimitriadis et al. (2021) on stable reliability diagrams and Allen et al. (2025) on carrying in-sample calibration into out-of-sample guarantees.

Evaluation

Protocol

- **Data:** 43 series, 31 of them daily. 17 market series (January 2016 to September 2026) were used during development. 10 currencies and 3 temperature stations had never been used before. One further temperature series and 12 horse-racing series were used in earlier experiments.
- **One pass:** each series starts from nothing. At each step the engine sees the outcome, scores its last call, learns, and makes the next call. No setting is fitted on later data.
- **Written first:** the engine, the rule for which unused series to include, and seven predictions were fixed in writing before the run. One amendment followed the first run, to how exact ties are ranked. Everything was rerun and no prediction was changed.
- **Measure:** information gain over the null, in bits per call. This is the base-2 logarithm of the probability HEX [PI] gave to what happened, minus the same for the null. Beside it are the share of calls right and the share right had the null alone made every call

("expected").

- **Uncertainty:** block bootstrap over 20-day blocks, with every series' calls on those days kept together; 95% intervals. For the undated racing series the block is the jockey. A read is verified when its interval is above zero.
- **Checks:** rerunning the engine on truncated histories at 258 sampled points reproduced the recorded calls exactly. The product code matches the reference implementation to rounding error.

Main result, next-step calls

Series	Event	Calls	Right	Expected	Bits per call over the null (95% interval)	Status
Markets, 17 series, 2016–2026	Move size	46,478	57.9%	57.2%	+0.0043 (0.0030 to 0.0056)	Verified
Markets, 17 series, 2016–2026	Direction	46,478	51.4%	51.2%	0.0000 (–0.0007 to 0.0006)	No verified edge
Currencies, 10 never used before	Move size	26,540	62.1%	61.2%	+0.0082 (0.0060 to 0.0106)	Verified
Currencies, 10 never used before	Direction	26,540	54.2%	54.0%	+0.0011 (–0.0002 to 0.0025)	No verified edge
Temperature, 3 stations never used before	Direction	29,408	61.0%	54.6%	+0.0466 (0.0430 to 0.0499)	Verified
Temperature, 3 stations never used before	Move size	29,408	63.1%	61.6%	+0.0186 (0.0158 to 0.0212)	Verified
Temperature, Chicago	Direction	24,678	62.9%	50.7%	+0.0644 (0.0598 to 0.0689)	Verified
Temperature, Chicago	Move size	24,678	55.1%	54.1%	+0.0048 (0.0035 to 0.0061)	Verified
Horse racing, 12 jockeys	Direction	13,022	74.8%	75.0%	–0.0017 (–0.0020 to –0.0014)	No verified edge
Horse racing, 12 jockeys	Move size	13,022	52.9%	49.7%	+0.0122 (0.0001 to 0.0296)	Verified, lower bound at zero

On move size the engine adds information in markets and on currencies it had never seen. On temperature it adds more, on both events. On market direction it adds nothing. On racing direction it is slightly behind: the null alone is right three times in four there, and what the engine adds costs a little.

The gain was positive on 12 of the 17 market series and 8 of the 10 new currencies, and at all three new temperature stations (direction right 63.1%, 64.5% and 60.5%). Status in the table is per row of the replay. The product keeps one status per read, starting from the market result, so it presents direction as having no verified edge on every series until its live record is decisive.

Learning, experience and the previous version

- **Learning adds to the result.** Against a copy of itself frozen after its first 500 calls on a series, the learning engine gained +0.0036 bits per call on market move size (0.0025 to 0.0048).
- **Experience matters.** On market move size the engine is behind the null over a series' first 100 calls (−0.011 bits per call), level by about 500, and ahead by +0.006 after 1,000.
- **Against version 6.5.** The previous engine reduced recent observations to a small set of discrete states. It runs beside 7.0 on every call as a comparator. On the same calls 7.0 gained +0.0037 bits per call on market move size (0.0025 to 0.0050) and was level on market direction.
- **Predictions.** Of the seven predictions written before the run, six held. The seventh, that restating strength would not lower information gain, came out neither way (−0.0001, interval −0.0010 to 0.0007).

Reading the result

The 17-market run is a development backtest: those series were used while designing every version. The currencies and temperature stations are the independent evidence. Two of the ten currencies are unusual, one held in a narrow band and one in crisis; both were kept because the inclusion rule written beforehand kept them.

In markets the move-size effect is volatility clustering, a long-established effect (Engle, 1982; Bollerslev, 1986). The result shows that the engine captures it beyond the series' own running rate, not that it captures it better than a volatility model.

The current benchmarks for this effect are realised-volatility models. Kilic (2025, revised 2026) finds that regime-switching HAR and ARFIMA models are hard to beat and that machine-learning methods do not consistently surpass them; Christensen et al. (2026) find machine learning ahead of HAR, more so at longer horizons. HEX [PI] has not been tested against either family, and that comparison is the first open question below.

The temperature direction result is most likely ordinary mean reversion, which a linear or persistence model would also capture. Version 7.0 has not been compared with one.

HEX6-BLE server and mesh: concept

This section describes a design. No unit has been built or tested, and component-level detail is in the confidential technical documentation.

The server

HEX6-BLE is a palm-sized server built from standard, low-cost parts: a single-board computer, a microcontroller, environmental sensors (temperature, humidity, light, soil moisture), current clamps for electricity metering, low-power laser diodes with photodiodes, and a secure element that signs readings. It senses a building's environment and energy use, runs HEX [PI] locally, and schedules flexible loads such as a battery or heat pump through the manufacturers' own interfaces.

The design principle is that physical inputs are read the way the software reads any series: each optical channel and each sensor is reduced to where its present reading ranks among its own recent readings, so one method serves both data and the physical environment.

The experimental element: a fungal memristor

A memristor is a component whose resistance depends on the history of current through it. LaRocco et al. (2025) reported that dehydrated shiitake mycelium shows memristive behaviour, switching at up to 5,850 Hz with about 90% accuracy when tested as volatile memory. The samples in that work are dried, not living tissue, and the result has not yet been independently replicated for this use.

The wider field is reviewed by Adamatzky et al. (2022), who describe fungal memristors, oscillators and sensors built from mycelium. Chatzipaschalis et al. (2025) model mycelium's electrical behaviour and reproduce it in FPGA and memristor circuits, which offers a simulation route for tests T1 to T3 before live samples are used.

This element is experimental, and the reason for including it is specific. HEX [PI] keeps tallies that fade at fixed rates. A memristor's state also fades after it is written. If the fungal element's fading is stable and repeatable, it could act as a physical record of the same history and be read as one more input. That is a hypothesis. Tests T1 to T3 establish whether the element behaves as required, and test T6 whether it adds any forecasting signal. If it does not, the element is removed and the server works without it. No other function depends on it.

The mesh

Each server links to its neighbours over short, low-power optical channels and exchanges timing and state. The mesh carries information, not power. A server that loses its links keeps working alone. Whether sharing evidence between neighbouring servers improves local forecasts is untested.

Optical wireless links for sensor networks are an established research area; Celik et al. (2023) survey the designs, trade-offs and commercial systems.

The ledger

Each server signs its readings with a hardware-held key. At intervals it publishes a single fingerprint (a Merkle root) of its recent readings to a public blockchain, while the raw readings stay on the server. Anyone holding a reading can later prove it was recorded at that time and has not been changed, without household data being exposed. Signing proves a record is unchanged; it does not prove the sensor told the truth.

A close precedent is Vasheghani Farahani and Treiblmaier (2025), who built and measured a blockchain-secured solar energy logger on low-power edge devices and report small operating overhead for a verifiable audit trail.

Safety

Lasers are Class 2 or lower and aimed along enclosed or above-head paths. Current clamps are installed by a licensed electrician. The building's own protection devices remain in charge of electrical safety.

Experimental programme

Six tests, each with a control and a pass criterion fixed before data is collected. Protocols are to be registered publicly before each test begins, and null results reported.

Test	Question	Control	Pass criterion
T1	Is the fungal element a memristor?	Plain resistor of matched resistance	Pinched hysteresis loop through the origin; none in the control

Test	Question	Control	Pass criterion
T2	Can light-driven pulses write a readable state?	Inert carbon disc in place of the element	Patterns recovered from the state above chance; control at chance
T3	Is the element's fading stable?	Repeat at a second temperature and humidity	Retention curve repeatable within a stated tolerance
T4	Does the optical link hold synchronisation?	Link blocked; ambient light varied	Error rate and clock agreement within stated targets
T5	Does scheduling save energy?	Randomised alternation with a fixed timer; weather as covariate	Lower grid energy and peak demand, interval excluding zero
T6	Does the fungal element add signal?	Shuffled copy of the element's channel	Information gain over the null rises with the real channel and not with the shuffled copy

T1 to T4 are bench tests. T5 and T6 need one instrumented building for at least eight weeks. T6 is the decisive test for the experimental element and T5 for the energy claim; either can fail independently of the other.

For scale on T5: in a living-lab residential building with battery storage and a heat pump, Le et al. (2023) report a 6% reduction in energy cost from optimised scheduling under uncertainty.

Eight weeks is enough to detect only a large effect in T5 or T6. The minimum detectable effect is to be computed from baseline variance and stated in each registration.

Limitations and open questions

Limitations

- **Scope of the market result.** The measured edge in markets is on move size. Next-day market direction has no verified edge.

- **The size effect is known.** That large moves follow large moves is volatility clustering. HEX [PI] has not yet been compared with a GARCH or realised-volatility baseline, so the result shows the engine captures the effect beyond the series' own record, not that it captures it better.
- **Modest effect.** In markets the gain is 0.004 bits per call and 0.7 points of share right. Its economic value after costs has not been assessed.
- **Development data.** The 17 market series were used while designing every version. The independent evidence is the ten currencies and three temperature stations.
- **Correlated series.** The series are not independent, so the effective sample is smaller than the number of calls. The intervals allow for same-day dependence by keeping every series' calls on a block of days together.
- **Slow start.** On a new series the engine loses information over its first hundred calls and is level with the null by about 500.
- **No live record yet.** History cannot prove the future. Live calls are recorded from the release of version 7.0.
- **Not repeated on 7.0.** A 70-call comparison with a language model and a comparison with a persistence model on temperature were run on earlier versions. Neither has been repeated, and neither is reported here.
- **Single implementation.** All results come from the authors' own code and have not been independently replicated.
- **Hardware is unbuilt.** Cost, link performance and energy savings are targets.
- **One study behind the fungal element.** Variation between samples, ageing, and sensitivity to humidity and temperature are unknown for this use. A resemblance between the engine's fading tallies and a memristor's fading state does not mean the element carries useful information.

Questions for collaborators

1. Does HEX [PI] add skill on move size beyond standard volatility models on the same days? And how does it compare with general-purpose language models given the same information, on several hundred pre-registered calls and more than one model?
2. Can the rule that declares an edge verified be replaced by a test that stays valid at every stopping time, so that the claim does not depend on when the record is read?
3. What is the retention behaviour of dehydrated shiitake mycelium under controlled humidity and temperature, and how much does it vary between samples?

4. Does sharing evidence between neighbouring servers improve local solar and load forecasts?
5. What saving in grid energy and peak demand does scheduling on calibrated probabilities achieve against a rule-based controller, in a randomised design?
6. For regenerative farming, which soil and microclimate signals are predictable at the hourly scale, and does acting on them change irrigation or yield?

Access, rights and references

Access to full detail

The complete technical documentation (equations, parameter values, component specification), the source code and the raw result files are available to academic reviewers under a non-disclosure agreement with Evrica Technologies Inc. Terms for publication, authorship and intellectual property are set by written agreement with each institution.

Contact: Anna Maria Sarkisyan, Founder and CEO, Evrica Technologies Inc.,
anna@evrica.tech

Copyright and rights

© 2026 Evrica Technologies Inc. All rights reserved.

HEX [PI] and HEX6-BLE are names used by Evrica Technologies Inc., a Delaware corporation. This overview may be shared for academic review. It may not be adapted or republished without written permission. Aspects of the methods described are the subject of a patent application. Third-party works cited remain the property of their owners.

This document describes research. It is not an offer of securities, investment advice or a performance guarantee.

References

1. Adamatzky, A., Ayres, P., Beasley, A. E., Chiolerio, A., Dehshibi, M. M., Gandia, A., et al. (2022). Fungal electronics. *Biosystems*, 212, 104588.
[doi:10.1016/j.biosystems.2021.104588](https://doi.org/10.1016/j.biosystems.2021.104588)
2. Allen, S., Gavrilopoulos, G., Henzi, A., Kleger, G.-R., and Ziegel, J. (2025). In-sample calibration yields conformal calibration guarantees. [arXiv:2503.03841](https://arxiv.org/abs/2503.03841)

3. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
4. Celik, A., Romdhane, I., Kaddoum, G., and Eltawil, A. M. (2023). A top-down survey on optical wireless communications for the Internet of Things. *IEEE Communications Surveys and Tutorials*, 25, 1–45. [doi:10.1109/COMST.2022.3220504](https://doi.org/10.1109/COMST.2022.3220504)
5. Chatzipaschalis, I., Tompris, I., Kleitsiotis, G., Chatzinikolaou, T. P., Fyrigos, I.-A., Ayres, P., Rubio, A., Adamatzky, A., and Sirakoulis, G. (2025). Mycelium electrical behavior: from morphological modeling to hardware implementation. Conference paper. [doi:10.5281/zenodo.17638776](https://doi.org/10.5281/zenodo.17638776)
6. Christensen, K., Siggaard, M., and Veliyev, B. (2026). A machine learning approach to volatility forecasting. [arXiv:2601.13014](https://arxiv.org/abs/2601.13014)
7. Chua, L. O. (1971). Memristor: the missing circuit element. *IEEE Transactions on Circuit Theory*, 18(5), 507–519.
8. Dawid, A. P. (1984). Statistical theory: the prequential approach. *Journal of the Royal Statistical Society, Series A*, 147(2), 278–292.
9. Dimitriadis, T., Gneiting, T., and Jordan, A. I. (2021). Stable reliability diagrams for probabilistic classifiers. *Proceedings of the National Academy of Sciences*, 118(8), e2016191118.
10. Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
11. Gneiting, T., and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378.
12. Herbster, M., and Warmuth, M. K. (1998). Tracking the best expert. *Machine Learning*, 32(2), 151–178.
13. Kilic, R. (2025, revised September 2026). Linear and nonlinear econometric models versus machine-learning models: evidence from realized-volatility forecasting. Finance and Economics Discussion Series 2025-061r1, Board of Governors of the Federal Reserve System. [doi:10.17016/FEDS.2025.061r1](https://doi.org/10.17016/FEDS.2025.061r1)
14. LaRocco, J., Tahmina, Q., Petreaca, R., Simonis, J., and Hill, J. (2025). Sustainable memristors from shiitake mycelium for high-frequency bioelectronics. *PLOS ONE*, 20(10). [doi:10.1371/journal.pone.0328965](https://doi.org/10.1371/journal.pone.0328965)
15. Le, A. T., Mazidi, M., and Steen, D. (2023). A novel energy management system for optimal energy and flexibility scheduling of residential buildings: a case study in HSB Living Lab. 2023 *IEEE Asia Meeting on Environment and Electrical Engineering*. [doi:10.1109/EEE-AM58328.2023.10394920](https://doi.org/10.1109/EEE-AM58328.2023.10394920)

16. Merkle, R. C. (1988). A digital signature based on a conventional encryption function. *Advances in Cryptology: CRYPTO '87*, 369–378.
17. Vasheghani Farahani, J., and Treiblmaier, H. (2025). A sustainability assessment of a blockchain-secured solar energy logger for edge IoT environments. *Sustainability*, 17(17), 8063. [doi:10.3390/su17178063](https://doi.org/10.3390/su17178063)
18. Vovk, V. G. (1990). Aggregating strategies. *Proceedings of the Third Annual Workshop on Computational Learning Theory*, 371–383.