Mantis — Manifesto

Mantis Biotech

July 2026

CONTENTS

Mantis — Manifesto

Mantis is building the infrastructure layer for the world's most accurate digital twins of human beings.

The observation

The average American now spends roughly 40% of their waking life using a computer. Compounded over an adult lifetime, that is nearly two decades of continuous, recorded interaction with machines. By 2025, the average connected person will generate nearly 5,000 discrete digital interactions per day, and humanity as a whole will produce around 400 million terabytes of new data every 24 hours.

For the first time in history, a person's behavior is being continuously and passively documented at high resolution. Every one of those interactions is a measurement of a person's state. When fed into the right models, measurements of the state become forecasts of the trajectory — a working simulation of what a unique individual will do next. Mantis intends to build those models.

The theory

Human behavior is difficult to predict on a fully encompassing level because it is considered a chaotic system. A chaotic system is one where two nearly-identical starting states pull apart over time: an initial uncertainty δ_0 grows roughly as $\delta(t) = \delta_0 \cdot e^{\wedge}(\lambda t)$, where λ is the Lyapunov exponent. The larger λ is, the faster small errors blow up. Its reciprocal, $1/\lambda$, is the Lyapunov time: the horizon past which any forecast decays into noise no better than a baseline guess. Earth's atmosphere has a Lyapunov time of roughly one to two weeks, which is precisely why a weather forecast is useless past that window.

The core theory of the company is that the right algorithms for processing human behavioral data produce Lyapunov times long enough to yield meaningful, valuable predictions. Raw behavior is chaotic; the correct coarse-grainings of it are not. Currently we separate human behavior into three regimes:

1. **Stable aggregates** — trends, rates, and baselines that average over the chaos beneath them.
2. **Stable routines** — habit- and constraint-anchored patterns.
3. **Micro-chaos** — the thin sliver of genuine unpredictability that no one can forecast and no diagnostic depends on.

The questions worth answering live almost entirely in the first two regimes. Our moat is the ability to discover and define these algorithms: identifying which coarse-grained features of a person's raw, chaotic data are stable, and forecasting them to the exact horizon at which a decision still has to be made. The full regime breakdown, a worked numerical example, and the peer-reviewed evidence quantifying each regime are in the Appendix.

The same statistical regularities that let markets clear, insurers price risk, and epidemiologists model outbreaks in aggregate persist at the level of the single person's own aggregates.

Why nobody has built it

Not for lack of data or ML talent. Massive incumbents like Google, Apple, and Meta each hold a deep slice of this record and employ the best machine learning teams on earth. Structural barriers stop them:

Silos. Each company sees only the behavior of its own products. A digital twin requires the integration of a person's data across devices, platforms, and physiological systems. No incumbent can acquire its competitors' slice.

Incentive misalignment. Every incumbent that holds this data monetizes it through a specific, narrow product: Google sells search and ads, Apple sells hardware, Meta sells attention. A holistic predictive model of the individual runs orthogonal to these direct incentives. Building it would mean diverting the crown-jewel data asset away from the justification for the company's valuation, toward a product with enormous reputational exposure. No public-company executive gets rewarded for that trade. The model that would serve the person rather than the platform sits permanently outside every incumbent's mandate.

The opportunity isn't too big for incumbents. It's structurally incompatible with any company whose business is a single downstream use of the data.

What Mantis is

Mantis is the neutral infrastructure layer for behavioral modeling. It plays the same structural role OAuth plays for identity, Plaid plays for banking, and Stripe plays for payments: a broker that succeeds precisely because it doesn't compete with either side.

That neutrality serves as a strong moat, because it aligns our incentives with both the individual and the company.

As a result, Mantis is also data-agnostic: our pipelines ingest and normalize any stream a person grants us (device telemetry, wearable physiology, digital activity, clinical records) and our models turn that unified record into a living, predictive twin of that individual's high-inertia cognitive and physical behavior. Companies then request scoped, revocable access to predictions from that twin, the way an app requests scoped access through OAuth.

Why healthcare first

We begin in medicine because it is where the value of individual-level prediction is highest, the permission to hold sensitive data is most clearly defined, and the demand is proven. Healthcare also imposes the strictest bar for privacy, validation, and accountability, so it's a fun place to get started.

On a more personal note, there is a 0% chance that I, as a founder, will ever do anything with my life other than try to improve the human experience. If you want to figure out a better way to predict purchasing behavior or advertise to consumers, go bark up another tree.

The horizon

Every consequential system humans build eventually gets a digital twin. The last and most valuable system without one is the human being. Mantis is that layer. We allow computers to understand and predict individual human behavior accurately, at the horizons that matter, in the service of the person being predicted.

Appendix — Supporting theory and evidence

A. The three regimes of human behavior

Human behavior is chaotic at the micro level while being remarkably stable at the level that matters for prediction. This is the same principle that lets thermodynamics work despite the chaos of individual molecules, and lets climate be forecast for decades while weather dies at two weeks: the trajectory of a single particle is unforecastable, but the temperature of the gas is rock-solid. Chaos at the bottom averages into stability at the top. Behavior separates accordingly:

Stable aggregates — the target. The coarse-grained statistics of a person's behavior: their weekly average resting heart rate and how it's trending, their typical sleep architecture, the rate at which they miss medications, the slow drift of their cognitive baseline over months. These are the sums, means, and trends over many underlying events — and precisely because they average over the chaos beneath them, they have long Lyapunov times and very high predictability. This is where diagnostics live. To flag an infection, detect cognitive decline, or forecast decompensation, you do not need to know what a person will do at 3:47 p.m. Tuesday. You need the trend, and the trend is stable.

Stable routines — also reachable. Coarse individual behaviors anchored by habit and constraint: where you sleep, your commute, your spending categories. This is the regime the *Science* mobility work measured, topping out around 93% predictability.

Micro-chaos — the sliver beyond today's models. The exact word you'll type next, one specific impulse purchase, a single momentary mood spike. Lyapunov times of minutes; genuinely unforecastable with current methods. But note what this region actually contains: individual, instantaneous, one-off events — almost none of which are what a diagnostic or a care decision depends on. The unpredictable part of a human being is real, but for the questions worth asking, it is a thin and largely irrelevant sliver. The signal we want survives the averaging; the noise we can't yet predict washes out in it.

B. A worked example of the actionable window

Suppose an individual's day-to-day resting heart rate is noisy and effectively unpredictable at the single-day level, but its 7-day rolling average drifts slowly. Start from a 10% error on that average with an uncertainty that doubles roughly weekly, and the forecast stays under 50% error until 2^(t/7) crosses 5 — about 16 days. That is a two-week actionable window on the aggregate, extracted from data that looks like noise if you stare at any single day of it. The barrier isn't the phenomenon — it's knowing which function of the data to model, and at what horizon.

C. Empirical evidence

The theoretical ceiling on human predictability has been quantified, repeatedly, in the peer-reviewed record.

Movement. In a landmark 2010 *Science* study, Song, Qu, Blumm, and Barabási analyzed the anonymized mobility trajectories of millions of mobile-phone users and found a 93% theoretical upper bound on the predictability of an individual's location — and, strikingly, that this figure barely varied across people, independent of how far they routinely traveled. The intuition that frequent travelers are less predictable turned out to be wrong; nearly everyone is about equally forecastable. Human movement, one of the most basic behavioral signals, is roughly nine-tenths determined.

Personality and preference. In a 2013 *PNAS* study, Kosinski, Stillwell, and Graepel showed that ordinary Facebook Likes — nothing more — could accurately predict sexual orientation, ethnicity, religious and political views, intelligence, personality, substance use, and even whether a person's parents had separated. A 2015 follow-up in *PNAS* went further: a model using a person's Likes judged their personality more accurately than their coworkers, friends, roommates, and family, and with enough data (around 300 Likes) rivaled the accuracy of their own spouse. Machines already read us better than the people closest to us do.

Physiology. A single consumer smartwatch takes on the order of 250,000 physiological measurements per day, and that stream carries early warning of illness. In a 2020 *Nature Biomedical Engineering* study from Snyder's lab at Stanford, 81% of COVID-19 cases showed detectable physiological deviations — in resting heart rate, sleep, and activity — at or before symptom onset, some as much as nine days early. The signal that you are getting sick exists in your data before it exists in your awareness.

Sources

- U.S. average daily screen time ≈ 7 hours (~40% of waking hours): DataReportal / Statista (2025–2026)
- ~4,909 digital interactions per person per day by 2025: IDC Data Age projection
- ~402.7 million TB of new data created daily worldwide (~147–180 ZB/year): IDC / Statista (2024–2025)
- Average U.S. broadband household ≈ 641 GB/month: OpenVault Broadband Insights
- 93% theoretical upper bound on individual mobility predictability: Song, Qu, Blumm & Barabási, “Limits of Predictability in Human Mobility,” *Science* 327, 1018–1021 (2010)
- Facebook Likes predicting sensitive traits: Kosinski, Stillwell & Graepel, “Private traits and attributes are predictable from digital records of human behavior,” *PNAS* 110, 5802–5805 (2013)
- Computer models out-judging friends and family on personality: Youyou, Kosinski & Stillwell, “Computer-based personality judgments are more accurate than those made by humans,” *PNAS* (2015)
- Smartwatch presymptomatic illness detection: Mishra et al. (Snyder Lab, Stanford), “Pre-symptomatic detection of COVID-19 from smartwatch data,” *Nature Biomedical Engineering* (2020)