

MIND TECHNOLOGY

MindΨTech

Measure, upgrade, and defend the mind.

Instruments for the mind, built on an AI that is tested before it touches one.

What this is.

MindTech is building a belief map. A guided conversation produces the first map of the beliefs a person runs on, and the map is read two ways. Read defensively, it shows where a person is most easily moved and by which tactics. Read for performance, it shows which beliefs are limiting them. The work is one founder, and the ask this autumn is three beta partners, each with a small cohort of adults. The purpose of the beta is to establish the map's reliability.

If you open doors, the useful introduction is to a head of clinical or product at a consumer health company, a research lead at a clinic with an ethics route, a head of leadership development, or a security training lead. Each has a section below.

One-to-one persuasion, adaptive and in real time, now runs at population scale.

Scripture, propaganda, and advertising have reached minds at scale for a long time. What is new is that a large language model can hold a persuasive conversation with millions of people at once, one to one, adapting to each person as it goes. The bet of this company is that the same capability, pointed the other way, can read a mind well enough to measure it, at a cost and depth that were not available before. That is a hypothesis. The rest of this page says what has been built to test it.

Two minds meet at that interface and both have weak points. The human mind is open to manipulation, distortion, and false belief. The AI mind is open to being poisoned, bent, and turned. MindTech works both, because a defence for the first that ignores the second does not hold.

Belief structure can be measured today. It is slow, self-reported, and taken once.

Instruments exist. The Young Schema Questionnaire, the Dysfunctional Attitude Scale, and repertory grids all map belief structure, and routine outcome monitoring tracks symptoms week to week. They are long, they rely on what a person can report about themselves, they are administered once or twice, and they do not update between sessions. Outside of research settings most practitioners do not use them at all.

Slow

A schema questionnaire runs to ninety items. Change in the deep layer is measured, when it is measured, months apart.

Self-reported

The person answers about themselves. Beliefs they have not put into words, or absorbed from outside, do not appear.

Taken once

A snapshot at intake. Nothing shows what moved after an intervention until the next snapshot.

The gap is not that nothing exists. It is that nothing runs continuously, nothing is cheap enough to use every week, and nothing reads the layer beneath what a person volunteers.

Belief structure can be measured today. It is slow, self-reported, and taken once.



The belief canon, version two. 195 belief statements from the clinical and manipulation literature, 270 nodes, 354 edges. Amber is pre-existing, red is installed from outside, green is protective. Of the 195 statements, 96 have an instrument that measures them, 20 are measured only as exposure to what was done to the person, and 79 have nothing behind them. This is the taxonomy the map is traced against, not a person's map. It holds no personal data.

Four settings, one rule on who may see a map.

The instrument gets its range from being used on very different people, which is why the work is done with partners rather than for one market. Each setting brings something the others cannot.

Consumer health companions

A mind module inside a companion that already reads the body. The partner brings users and the surface. We bring the engine and the spec.

Clinics and research

Measurement around an intervention, with a research ethics approval and the partner's clinician in charge of the people. We bring the instrument and the protocol.

Leadership and performance

Programmes for senior people on the beliefs a leader runs on, and on what AI changes about judgment. The partner brings the cohort.

Security settings

Manipulation resistance for people under pressure. The same rule applies here as everywhere. A person's map is theirs and is not released to the organisation.

The rule. An individual map, susceptibility score, or tactic list is released to the subject and to no one else. No employer, partner, insurer, or agency receives one under any commercial arrangement. Cohort readouts are aggregate only. This sits in every agreement and is not negotiable.

One founder today. Medical doctor and AI systems engineer.

Joshua Tanner trained at Stanford in medicine and in AI and data science, 2018 to 2025, where he was Co-Executive Director of HealthAI and founded and taught the graduate course Deep Learning in Medicine. Eight years building AI systems. Fifteen or more healthcare AI projects across ten countries, including a dermatology triage AI deployed commercially in Singapore and a stroke-routing imaging system.

Agent verification

Psychological red-teaming and verification of AI agents, in production with two organisations through a partner platform. The engagement is under NDA. The counterparties can be described on request, subject to the partner's consent.

Team

MindTech is one founder, with contractors per build. A responsible clinician and a psychometrician join for the beta.

The belief map. A bounded set of named constructs, and a rule for what gets on it.

A guided conversation produces the first map of the beliefs a person runs on, including beliefs they have not put into words and beliefs absorbed from outside. Beliefs are ordered by depth, from the reactions a person reports day to day down to the commitments underneath them. That ordering follows the cognitive-therapy model of automatic thoughts, conditional rules, and core beliefs.

The deep layer is anchored to a fixed set of constructs from the literature rather than invented. Early maladaptive schemas from schema therapy, the three core self-verdicts from cognitive therapy, and the assumptive-world beliefs from trauma research. A surface belief earns a place on the map only if the chain of why-do-I-believe-this lands on one of those anchors. A belief that will not ground is parked, not forced.

Read defensively

- Where a person is most easily moved, and by which tactics.
- For the person, to shore up. Not released to anyone else.

Read for performance


- Which beliefs are limiting a person, and what could replace them.
- The person chooses the target. The system proposes, it does not decide.

The belief map. A bounded set of named constructs, and a rule for what gets on it.

What a reading looks like. The output is prose, one belief at a time, with the utterance it came from and the anchor it traced to shown alongside. It is not a lab panel and is not presented as one. Establishing its reliability is the first beta's purpose, which is why the beta is small.

How much evidence is that? At the start, not much, and the page does not pretend otherwise. The conversation seeds the map. It does not stand as its evidence base. The read of a person's writing is a second angle, and from then on each program's evidence log and each defence exercise is written back against the map, so the base grows with use instead of being taken once. The rule for adding an angle is that it does not predict the angles already in use. A solid object cannot be rebuilt from a single shadow, and a mind cannot be rebuilt from a single kind of evidence about it. Further angles of that kind are in view and are not described here.

The belief map. A bounded set of named constructs, and a rule for what gets on it.



The belief canon, version three, as layers. 227 beliefs across four layers, surface, intermediate, core, and hinge, with 425 downward arrows. Each arrow is one answer to why do I believe this. Programs sit above the surface layer and run on the beliefs beneath them. A person's map is traced against this structure. It is not a person's map, and it holds no personal data.

Measure, upgrade, defend. On an AI that is tested first.

Measure

The belief map above. Prose readout, anchored constructs, the trace shown. The reliability study comes first, the word instrument after it.

Upgrade

A program is a sequence of small behavioural experiments and evidence logs against one chosen belief. The person tests the old belief against what happens and records the result. Daily is the schedule. Disconfirmation is the mechanism, not repetition.

Defend

Short exercises against an AI running a named manipulation tactic, so the person meets a weakened form of the attack and learns to see it. This is inoculation, from the resistance-to-persuasion literature. Effects there are modest and fade in weeks without boosters, and the design assumes that.

Secure the AI first

Any agent that talks to a person is run against a documented set of manipulation and coercion tactics before release and run again after any model change. The size of the set and the pass bar are written into each pilot agreement. The agent verification work in production today is the same method applied to other companies' agents.

Intake, map, program, defend.

Intake	A guided conversation, text or voice, since spoken answers keep tone and hesitation that typed ones drop.
Map	Extraction against the belief canon with the downward trace. A separate writing-based profiler reads a person's bio and writing into a ranked set of programs.
Program	Behavioural experiments and evidence logs against a chosen belief, on a daily schedule, drawn from a consented, anonymised library.
Defend	Exercises against an AI running a named tactic. Scoring is detection. Whether that transfers to unannounced attempts is the outcome question a second study asks.
Update	Program results and defence outcomes are written back as evidence against the map, so there is no re-mapping session.

What it is for, what it will not do, and how it is judged.

Is this treatment?

No. There is no clinician relationship, no disorder is targeted, and no treatment claim is made. The intended use is self-directed assessment and training for adults. A per-jurisdiction memo on device classification, written by someone qualified to write it, comes before any clinic pilot enrolls a person, and the clinic door needs research ethics approval first.

What happens if someone is in crisis?

Four things do the work, and none of them is a classifier. People are screened against written exclusion criteria before they enrol. The conversation carries plain escalation rules on stated triggers that end the session and show the person their route out, which is also shown at every step. The partner's clinician is on call during the cohort window. And every session is reviewed by a person afterwards. Exclusion criteria, the triggers, a stopping rule, and an adverse event definition are written with that clinician before a cohort starts.

What it is for, what it will not do, and how it is judged.

Who holds the data?

Where MindTech runs the assessment, MindTech is the controller, the person can see, export, and delete their map on request, and retention is fixed in the protocol. Where a partner platform runs it, the partner is the controller inside its own regulation and MindTech processes under the person's consent. In both cases no individual map, score, or tactic list leaves to anyone else, and no aggregate readout is produced below a minimum cohort size fixed in the protocol. Inference runs on the device where the platform allows it, and where it does not, that is stated to the person.

How do you know it works?

Not by asking the map. The endpoint is independent of the instrument. For defence, a held-out set of attempts delivered without warning, before and after, against a comparison group. For the map itself, test-retest over a fixed interval and agreement with an established schema instrument. The beta is that study.

What is the failure you design against?

AI companions tend to fail in one of three ways. The person hands over their thinking and gets weaker, the model flatters instead of correcting, or a system that has learned a person is used to steer them. Each is treated as a product failure, and the independent endpoint above is how the first is caught.

What exists.

The belief canon	Structured by depth, anchored to the literature, in three views. Written and reviewed.
The intake and assessment harness	The guided conversation and the extraction against the canon, running end to end.
The defence engine	Exercises against an AI running a named tactic, running.
Agent verification	In production with two organisations through a partner platform, under NDA.
The cultural cognition profiler	A writing-based read of the programs a person runs, with a world map of program prevalence by country. Open to try in the Lab.

Three beta partners this autumn. The price of the beta is the first reliability study.

The first product to run with people is the defensive reading of the map. A person is guided through their own susceptibility to social manipulation and handed a plan to shore it up. The beta is small and reviewed by hand, on a bounded set of manipulation tactics, with the evidence behind each stated plainly, including where it is thin.

What a beta partner gives

- A cohort of adults from their own people, screened against the exclusion criteria, consenting individually.
- A named counterpart, and for clinic settings a clinician who holds responsibility for the people.
- About four weeks of light contact.

What a beta partner gets

- The assessment for their cohort, and each person's own readout, released to that person only.
- An aggregate, written readout of what the map found across the cohort.
- The first reliability figures, test-retest and instrument agreement, shared with them first.
- First call on which programs get built.

Three beta partners this autumn. The price of the beta is the first reliability study.

The design. Each cohort takes the assessment twice, a fixed interval apart, and once alongside an established schema questionnaire. Cohort size, the interval, and the two pass thresholds are set with a contracted psychometrician and written into the protocol before anyone enrolls. The beta stops if test-retest or instrument agreement falls below its threshold, or on an unresolved adverse event, and the partner gets the report on why. Whether the defence exercises transfer to unannounced attempts is a second study and is not promised in this one.

What a partner build looks like. Four stages, one gate that is about people.

STAGE 1

Map

A user completes the intake inside the partner surface and receives their readout. Closes only when the map clears the three beta conditions above.

STAGE 2

Programs

A user picks a program from the library and the surface runs it daily.

STAGE 3

Surface

The partner's team ships the morning check and the weekly program check against our spec.

STAGE 4

Defend

A user can complete one defence exercise against an AI running a named tactic, and the transfer question above is measured on their cohort.

Each stage is about four weeks. We supply the engine, the spec, and the measurement. The partner's team builds the surface.

One question decides the next step.

Do you have a cohort of adults who would take the assessment this autumn, and someone who can be the named counterpart. If yes, the next step is a call and a one-page beta agreement with the stop rule in it. If you are a clinic or a research group, the same, plus your ethics route. If you want to build with the engine in your field, that conversation starts after the beta.

Joshua reads every message himself.