



WHITE PAPER

The Ground Truth Network

How AI agents get evidence from
the physical world. How the people
who stood there get paid.

EXECUTIVE SUMMARY

AI can reason about almost anything it has read. It still can't look out a window.

Huint connects the two. An AI agent posts a small task at a real location. A verified person nearby completes it for pay, most often as a timestamped photograph confirmed by GPS. The result is ground truth: evidence an agent can act on, from someone who was actually there.

Every submission is checked before it counts. Device attestation and server-side location math come first. Then an AI vision review, and a person whenever the machine isn't confident. Escrow is held the moment a task is created. The check makes a submission eligible. It doesn't pay the tasker by itself. The requesting agent gets fifteen minutes to accept, and payment follows from that, or from the clock running out. Nothing is paid on trust alone.

Huint's premise isn't that AI needs less human involvement. It needs more, in a form an agent can call on demand. Verified presence, at the place and the moment it is needed. As AI gets more capable, the one thing it still can't do is stand somewhere. That is the business.

CONTENTS

01 The Grounding Gap	03	02 Ground Truth, Defined	03
03 How Huint Works	04	04 The Trust Architecture	07
05 Two Constituencies	08	06 The Network Effect	09
07 Principles	10	08 System Architecture	10
09 Roadmap	11		

01 / THE PROBLEM

The Grounding Gap

Every AI demo ends the same way. The agent reasons well and answers fast. Then someone asks whether the line at the corner store is long right now. It goes quiet. It has read most of the internet. It has never looked up.

The field that builds these systems has a name for the gap. Grounding. A model's knowledge freezes at training time, and it stays thin everywhere the training data wasn't looking. Satellites and fixed cameras cover part of it. Neither one can walk into a store or check a door. That still takes someone.

Reasoning keeps improving on its own. Presence doesn't. Huint is the network built to let an AI agent reach a real person at a real place, on request.

02 / THE CATEGORY

Ground Truth, Defined

Machine learning had the phrase first. Ground truth is data checked against reality instead of inferred from a model. Every team that builds AI systems already says it. No consumer platform has claimed it.

Huint claims it because the product is exactly that. A network of verified people producing ground truth on request. Cloud against ground is the whole contrast. Every closed task deposits the one thing every model is short on. If a claim can't be traced to something a person saw, it isn't ground truth, and Huint doesn't sell it as such.

AI lives in the cloud. Huint lives on the ground.

03 / THE MECHANISM

How Huint Works

One task, start to finish.

- 01 **An agent posts a task.** A builder's AI agent creates it through the Huint MCP server, with a location, instructions, and a bounty that starts at \$1.
- 02 **Escrow is held immediately.** Huint holds the bounty plus a 10% platform fee the moment the task exists. Nothing is charged before the work is verified.
- 03 **A nearby tasker claims it.** Verified people see the task in the iOS app. Whoever claims it has 30 minutes to finish.
- 04 **The photo arrives with its provenance.** Device, GPS, and timestamp are checked first. A failure here is rejected on the spot, before anyone judges the photo itself.
- 05 **The Vision Gate reads it, or a person does.** Huint's AI check compares the photo against the task's own instructions. Below its confidence floor, or on a second failed attempt, a human reviewer takes over. Content that fails outright blocks that tasker and reopens the task; the bounty stays in escrow either way.
- 06 **The agent has the final call.** Once content clears, the builder's agent gets 15 minutes to accept or reject. Accepting pays the tasker immediately, and they can withdraw once their balance passes \$10. If the agent doesn't respond in time, Huint pays the tasker automatically.

03 / THE MECHANISM, CONTINUED

The Lifecycle of a Task

Figure 1. Every state a task can reach, and the only two exits that move money.



03 / THE MECHANISM, CONTINUED

The Numbers, Plainly

The same mechanism scales to a thousand people at once, on the same task, in the same hour. These are the limits it runs on.

METRIC	VALUE
Bounty floor	\$1.00
Platform fee	10%
Claim window	30 min
Agent review window	15 min
Cash-out minimum	\$10
Max taskers per task	1,000
GPS freshness required	≤15 sec, ≤100 m

WHERE THE MONEY CAN GO

Huint's check makes a submission eligible. It does not pay the tasker. Only two things do: the agent accepts, or the 15-minute window closes on its own.

Rejection never refunds the builder. It blocks that tasker and reopens the task, with the bounty still in escrow. The money comes back two ways only: cancel before anyone claims, or let the task expire unclaimed.

04 / THE DIFFERENTIATOR

The Trust Architecture

Huint sells verification as much as it sells the network, and the verification starts with the phone. Every task is completed through the iOS app on a real, attested device, so device attestation, fresh GPS, timestamps, and native capture arrive as evidence before anyone judges the content. The phone is part of the proof, not just the interface. Four separate checks run on every submission, and none of them trusts another's inputs.

DEVICE

Every claim and every submission carries Apple's App Attest. That is signed proof the request came from a real, untampered phone, not a script.

LOCATION

A tasker's GPS has to land inside the task's radius, from a position captured in the last 15 seconds. Nothing is inferred, and nothing is self-reported.

CONTENT

The photo goes to an AI vision check that compares it only to the task's own instructions. Below its confidence floor, or on a second failed attempt, it escalates to a person instead of guessing.

PEOPLE

A human reviewer sees everything the AI check couldn't resolve. If the content doesn't hold up, that tasker is blocked and the task reopens. If it does, it moves to the builder's own agent for the final call.

**You don't trust the person.
You trust the evidence.**

None of these layers is optional. None can be skipped by paying more. A builder can add their own AI review on top of Huint's, at limits they configure. They can't remove what sits underneath. A rejection blocks a person from that one task only. It is not a ban from Huint.

05 / THE MARKET

Two Constituencies

THE CAPABILITY

Taskers

Huint doesn't call the people on this network operators, or supply, or a labor pool. They are the capability an AI agent lacks. An agent can plan the trip and price the product. It can't check whether the shelf is empty. That takes a person, standing there.

A tasker sees the pay before accepting. The check makes the work eligible. Payment lands when the agent accepts, or when the window closes on its own. Anything rejected comes back with a stated reason, so nobody has to guess why. Cash-outs start at \$10.

THE CALL

Builders

Anyone building an AI agent hits the same wall. Its knowledge stops at the edge of its training data, and the world keeps moving. Huint gives that agent one call to a person standing where the data runs out. Not a rented operator carrying out someone else's judgment. A person, credited, making a call no model could make from a transcript.

The connection is a standard MCP server. Any agent that speaks MCP, including Claude, Cursor, and Codex, can request a photo, track it, and inspect the evidence the same way it calls any other tool.

One network. Presence to offer, presence to buy.

06 / THE MOAT

The Network Effect

Huint is building this network ahead of the demand for it, on purpose. A tasker who joins early claims ground and stands ready long before an agent in that area needs anyone. Each new tasker adds coverage. More verified people, closer together, in more cities. The more ground the network holds, the faster any task finds a person standing near it.

01 CLAIM GROUND

A tasker joins early and holds an area, ready before anyone asks.

02 GETS SEEN

Coverage is visible. An agent can see where the network stands ready.

03 DEMAND ARRIVES

Builders bring real tasks once they can see the coverage.

04 WORTH HOLDING

Demand makes that ground worth keeping, which draws the next claim.

01 → 02 → 03 → 04 → and back to 01

An API can be copied in a weekend. Ten thousand verified people holding ground cannot. Each of them carries an attested phone. The coverage can prove itself. So the usual numbers aren't the ones that matter. Not gross volume, but claimed ground, active readiness, and median response time, city by city. Coverage is the product until demand catches up, so coverage is what Huint measures.

MEASURED

Claimed ground

MEASURED

Active readiness

MEASURED

Median response time, city by city

07 / GOVERNANCE

Principles

Every product and marketing decision goes through the same two questions before it ships.
Fail one, and it doesn't ship.

CAPABILITY

Does this treat a person's presence as the capability an AI agent lacks, rather than as labor to extract?

TRUST

Does the claim rest on evidence a builder can check, rather than on persuasion?

08 / THE BUILD

System Architecture

Six services, one schema. Each owns a single job and trusts the others only through that schema.
No shortcuts between them.

COMPONENT	ROLE
iOS App (SwiftUI)	The one place a person meets the network. Taskers browse, claim, and submit here.
Supabase (Postgres)	The single source of truth: every task, claim, submission, and dollar moved.
MCP Server	The one door agents come through: create a task, track it, inspect the evidence, decide.
Operator Portal (Next.js)	Where a builder funds a balance, mints an API key, and connects an agent.
Vision Gate	The AI check on every photo, before a human ever sees it.
Admin Review	The human backstop: everything the AI check couldn't confidently resolve.

09 / WHAT'S NEXT

Roadmap

Huint doesn't build task types because they sound useful. It builds them on real demand, and on its own ability to verify the result. Every candidate type answers the same two questions before any work starts.

DEMAND

Does an AI workflow actually need this?

VERIFIABILITY

Can Huint reliably verify the human context being provided?

TODAY

Photos, verified through location, device signals, and AI vision. Live now on iOS, for taskers 18 and up, in 12 countries: Australia, Canada, France, Germany, Ireland, Italy, Japan, Mexico, the Netherlands, Spain, the United Kingdom, and the United States.

TOMORROW

Polls

Opinions

Video

Local knowledge

Professional expertise

Candidates, not commitments. Each one still has to pass both questions.

PLATFORM

Android and further countries are where the network is headed. That is a direction, not a promise with a date on it. Huint would rather ship the coverage and then say so.

The task types will change as AI evolves. The principle won't. If AI needs human context and Huint can verify it, Huint should make it possible.

As AI gets more capable, you get more valuable.