

AI in the Trades: Why Judgment Is Your Career Moat

TYGART MEDIA *STRATEGY*

VOLUME I • ISSUE 04
QUARTERLY POSITION

BY WILL TYGART • LONG-FORM POSITION • PRACTITIONER-GRADE

There is a version of the AI transition story that gets told constantly, and it goes like this: AI will automate jobs, workers will be displaced, and the people who adapt will be the ones who learn to use AI tools. This version is not wrong exactly. It's just missing the part that matters most for the people who actually work in the trades.

The people who build things, fix things, assess damage, run field operations, and carry years of hard-won judgment in their bodies and their hands — these are not knowledge workers whose jobs can be uploaded to a language model. Their work requires physical presence, sensory intelligence, and the kind of contextual judgment that comes from doing something 500 times in conditions that were never twice the same.

But the transition is real, and it's happening around them whether they're paying attention or not. The question isn't whether AI changes the trades. It's which trades workers end up on the right side of that change — and why.

The answer is not "the ones who learn to code." It's not "the ones who get an AI certification." It's the ones who understand what AI can't do without them, and position themselves as the irreplaceable layer between the intelligence and the outcome.

That's the Wire and Fire Guy. And the window to become one is shorter than most people realize.

01 What the Wire and Fire Guy Actually Is

In electrical work, the wire and fire guys are the experienced field technicians who come in after the rough work is done. They're not project managers. They're not estimators. They're the people who look at what the system is supposed to do, look at what's actually been installed, and bridge the gap between the plan and the physical reality. They troubleshoot. They adapt. They make judgment calls that no blueprint anticipated.

The name is an archetype, not a job title. It describes a class of worker who exists in every trades field: the senior technician in water damage who knows from the smell and the color of the staining that the timeline is longer than the moisture readings suggest. The fire restoration veteran who can read a smoke pattern and tell you which rooms were occupied and which weren't before the alarm triggered. The field supervisor who looks at an estimate and spots the three line items that will blow up into supplements before the job starts.

These people carry knowledge that cannot be extracted from documentation because it was never documented. It lives in their sensory memory, their accumulated pattern recognition, their feel for how this specific type of situation typically develops. AI systems trained on the documentation don't have it. AI systems that have processed thousands of job files come closer but still don't have the physical dimension — the reading of a space that happens in the first ten minutes of being in it.

That knowledge — embodied, sensory, judgment-based — is the moat. And right now, most of the people who have it don't know it's a moat.



02 The 18-Month Window

Here is what is true right now, in April 2026: AI systems can write estimates. They can process moisture readings. They can identify scope items from photos. They can draft communications to adjusters. They can route jobs. They can flag outliers in a dataset of completed claims. They can do all of this faster and cheaper than a human doing the same work.

Here is what is also true: every one of those AI outputs needs a human to verify it against physical reality before it becomes an action. The estimate needs someone on-site who can see what the AI couldn't. The moisture readings need someone who can read the environment around the reading — the substrate, the airflow, the odor, the age of the damage. The scope items need someone who can look at the photo and then look at the actual wall and tell you what the photo didn't capture.

That verification layer — the human in the loop between the AI's output and the physical world — is not going away. What is going away, over the next 18 to 36 months, is everything on the other side of that line. The data entry. The scheduling calls. The status updates. The form-filling. The paperwork that currently consumes a significant portion of every field technician's non-field time.

The technician who understands this transition has a clear path: move toward the verification layer, away from the data layer. Develop the judgment that makes the AI's output trustworthy or correctable. Become the person the AI reports to, not the person doing the work the AI can do.

The technician who doesn't understand it will find their job slowly hollowed out — not eliminated suddenly, but compressed, devalued, and increasingly focused on the tasks that AI hasn't gotten to yet, which is a shrinking list.

03 Why Judgment Is the Moat

Judgment is not the same as experience. Experience is a prerequisite for judgment but not a guarantee of it. Judgment is what happens when experience meets a situation that doesn't match any template and produces a correct decision anyway.

AI systems are template-matching engines at their core. They are extraordinarily good at situations that resemble situations in their training data. They fail — sometimes silently, which is worse — when the situation deviates from the distribution they've seen. A water damage job in a 1920s Craftsman with non-standard framing, original plaster walls, and an HVAC system that was retrofitted

twice is a deviation. An AI trained on modern residential restoration data will produce an estimate and a timeline. A Wire and Fire Guy with 15 years of experience will look at the same job and know the estimate is wrong and the timeline is optimistic, because they've been inside enough 1920s Craftsmans to know what those walls hold.

This is the moat. Not the ability to use an AI tool — that's table stakes within 18 months. The ability to know when the AI tool is wrong, and why, and what to do about it instead. That requires the tacit knowledge that only physical experience builds. It cannot be trained into a model. It cannot be acquired from a certification. It grows from doing the work in conditions the documentation never anticipated, enough times to develop the pattern recognition that operates below conscious awareness.

The trades worker who wants to be on the right side of the AI transition doesn't need to compete with the AI on the AI's terms. They need to become the irreplaceable layer between the AI's output and the physical world. That layer is called judgment, and building it is a career strategy.

04 The Context Layer as Job Security

There is a more technical version of this argument, and it's worth understanding even if you never write a line of code.

AI systems are dramatically more useful when they have context — specific knowledge about the situation, the history, the people involved, and the standards that apply. A generic AI asked to write an estimate for a water damage job produces a generic estimate. An AI given the job address, the

property age, the adjuster's history with this contractor, the specific moisture readings, and the known quirks of the local building code produces something much better.

The person who provides that context — who knows enough about the job to load the AI with the information that makes its output accurate — is not replaceable. They are, in fact, more valuable as AI systems get better, because better AI systems reward better context. The technician who can brief an AI the way a good editor briefs a writer — specific, accurate, anticipating the failure modes — gets dramatically better results than the technician who types a query and accepts whatever comes back.

This is what "human in the loop" actually means in practice. It's not a compliance checkbox. It's the functional requirement that the AI's output is verified, corrected, and contextualized by someone who has the embodied knowledge to know when it's right and when it isn't. That someone, in the trades, is the Wire and Fire Guy.

05 From Field Tech to AI Supervisor: What the Career Path Looks Like

This is not a story about leaving the trades. It's a story about moving up the value stack within them.

The field technician who wants to make this transition has three things to develop, in order of how quickly they compound:

Domain depth first. The judgment moat requires genuine expertise. The technicians who end up in the verification layer are the ones who actually know the work at the level where deviation from documentation is visible and

meaningful. This is built by doing the work, paying attention, and developing the habit of asking "why does this job look different from what the estimate anticipated?"

AI literacy second. Not coding. Not machine learning theory. The practical ability to give an AI system a useful brief, evaluate its output for the specific failure modes common to your domain, and correct it with the context that changes the answer. This is learnable in weeks, not years, and it compounds quickly once the domain depth is in place to evaluate the output.

Communication between the two layers third. The ability to translate between the physical world — what you're seeing in the field — and the data layer that the AI operates on. This is partly documentation discipline (logging what you observe in terms that AI systems can use later) and partly the ability to communicate your corrections and their reasoning so the system improves over time rather than repeating the same errors.

The career path is not: field tech → project manager → estimator → office. That path still exists but it's compressing as AI handles more of what project managers and estimators do. The path that compounds in an AI-native industry is: field tech with deep domain knowledge → field tech who understands AI output → field supervisor who runs AI-assisted teams → operations role that owns the verification layer for a company's AI systems.

That last role doesn't have a standard job title yet. In three years it will. The people who get those roles will be the ones who understood the transition early enough to position themselves correctly — and who built the judgment depth that no model can replicate.



06 A Note on Pinto

This is the article I wanted to write since we published the original Wire and Fire Guys piece. That piece named the archetype. This one tries to give it a career map.

Pinto — who handles the infrastructure layer in this operation, the GCP deployments, the Cloud Run services, the database architecture — is the Wire and Fire Guy of AI infrastructure. He doesn't just run the code. He understands what it's supposed to do, sees when it deviates from that, and bridges the gap between the plan and the physical reality of production systems. The AI produces the output. Pinto verifies it against what the system is actually doing and knows why they differ.

That's the role. That's the moat. The window to build it is open. It won't be open forever.

07 Frequently Asked Questions

Does this apply outside the restoration industry?

Yes. The Wire and Fire Guy archetype exists in every trades field and every industry where physical reality diverges from documentation. Construction, manufacturing, healthcare, agriculture, logistics — any field where experienced human judgment is applied to physical conditions that AI systems observe indirectly through data. The timeline and the specific skills differ by domain. The structure of the argument is the same.

What's the minimum AI literacy a trades worker needs to develop?

Three things: the ability to give an AI system a specific, accurate brief for a task; the ability to evaluate the output for domain-specific failure modes (the things AI typically gets wrong in your industry); and the discipline to log corrections in a way that builds context over time rather than each correction being one-off. None of this requires programming knowledge. It requires domain expertise applied to a new kind of tool.

How urgent is the 18-month window?

The 18–36 month range is where most of the data entry, scheduling, and communication tasks that currently consume field technician time will be substantially automated in adoption-leading companies. The companies that adopt early set the new baseline for what's competitive. Workers in those companies develop the verification-layer skills first and build the largest knowledge lead. The window is not a cliff — it's a slope — but the slope is steeper now than it will be in three years when the transition is mostly complete in leading companies and everyone is catching up.

What about union rules and job protections?

Job protections can slow the transition but don't reverse the value dynamics. The worker who has built genuine verification-layer expertise is more valuable whether or not the AI transition is delayed by contract. And the worker who hasn't built it is less valuable on the same timeline. The protection is in the skill, not the rule.



Wire and Fire: The AI Transition Career Cluster

- [Why Judgment Is the Moat: What AI Can't Replace in the Trades](#)
- [The Context Layer as Job Security](#)
- [From Field Tech to AI Supervisor: The Career Path That Doesn't Have a Name Yet](#)

Related: [The Human Distillery](#) — the methodology for capturing the tacit knowledge this cluster describes.