

Work Intelligence at the Point of Work

Why industrial knowledge has to move from the classroom to the equipment - and what it takes to put it there



Table of Contents

Summary.....	2
1. The gap between training and work.....	2
2. Why the obvious approaches have not scaled.....	2
3. The approach: recognise, guide, assist, assess.....	3
Recognise.....	3
Guide: symbols, not labels.....	3
Assist.....	3
Assess.....	3
4. The flywheel.....	4
5. What decides whether this can be deployed.....	4
Hardware, by environment.....	4
Offline operation.....	5
Certification.....	5
Data residency and clearance.....	5
6. What is demonstrated today.....	5
7. Conclusion.....	5

Summary

Industrial organisations invest heavily in training people before they work, and very little in helping them while they work. The gap is not a content problem. Procedures, drawings and permits already exist. The problem is that none of it is available at the moment and place where the decision is made.

This paper describes an approach that closes that gap: a worker raises a device toward equipment, the system recognises what it is looking at, marks the components with the engineering symbols the worker already knows, opens the relevant procedure and permit, assists during the task, and captures how the work was actually performed.

It also describes the engineering constraints that decide whether such a system can be deployed at scale — recognition without per-asset modelling, offline operation, hazardous-area hardware, and the data rights that govern learning from captured work.

1. The gap between training and work

A procedure is written once, taught in a classroom, and then expected to survive the journey to a plant, a wellhead or a compressor deck. In practice it does not travel well.

The experienced operator who knows the asset is rarely standing beside the person who needs to know. The permit is a document somewhere else. The drawing is a PDF on a laptop in an office. The worker, in front of the actual equipment, has the least access to the knowledge at the exact moment it is worth the most.

The cost of this gap is not primarily training cost. It is rework, standby time, near misses, and the slow loss of expertise as experienced workers retire without any mechanism that captures what they knew.

2. Why the obvious approaches have not scaled

Documentation systems put the content online but not at the point of work. They require the worker to know what to look for and to leave the task to find it.

Digital twins are accurate and valuable, but each asset must be modelled before anything can be shown. A refinery has tens of thousands of tagged items. Authoring a bespoke twin for each is not a programme that finishes.

Augmented reality pilots have repeatedly demonstrated well and deployed poorly. The demonstrations rely on a single pre-modelled asset, in good light, with connectivity, on hardware that is not certified for the area where the work actually happens.

The common failure is the same in all three: they require the world to be prepared in advance.

3. The approach: recognise, guide, assist, assess

The alternative is to prepare nothing and recognise everything.

Recognise

The device identifies equipment and its components from the live camera view. This uses a general-purpose vision model, constrained to an industrial component vocabulary, rather than a model trained per asset. Nothing about the specific unit needs to exist in a database beforehand.

Recognition also applies to engineering drawings. A piping and instrumentation diagram is itself a component list with tags, and reading it gives the system the vocabulary of a specific plant without any modelling work.

Guide: symbols, not labels

Recognised components are marked not with text labels but with the idealised symbols defined in ISA-5.1 and ISO 10628 — the notation an industrial worker already reads fluently.

This choice carries more weight than it first appears. A label has to be read; a symbol is recognised at a glance, at distance, in poor light, through a visor. It also means the camera view and the engineering drawing speak one language, so a worker moving between the drawing and the equipment is not translating between two representations.

Critically, symbols give the visual quality of a digital twin without the authoring cost of one. The system is not rendering the asset. It is annotating reality.

Assist

Each recognised component becomes an anchor for the knowledge attached to it: the procedure step, the permit for today's job, the preflight simulation, and the captured expertise of operators who have done this work before.

The most valuable form of assistance is the least technological. An experienced worker's judgement, captured once and surfaced at the right component at the right moment, is worth more than any rendering.

Assess

The final step is what distinguishes work intelligence from work instruction. The system captures how the task was actually performed — passively or actively, as the situation allows.

Assessment operates at two horizons. **At the point of work**, immediate feedback can correct a sequence error before it becomes an incident. **After the fact**, a structured debrief shows what went well and what did not, across a team and across a site.

Both the positive and the negative delta matter. The surprisingly good performance is as instructive as the error, and organisations almost never capture it.

4. The flywheel

Captured work is the only proprietary asset in this chain. Models are increasingly commoditised; industrial work data is not.

The loop runs: work is captured at the point of performance, the relevant components are extracted, the extraction is distilled into training material, and a domain model is trained on it. That model then improves recognition and assistance for the next worker.

This is why the acquisition layer matters more than the model. A general frontier model is available to everyone. A model trained on how this industry actually performs this work is not, and cannot be replicated by anyone without the same access to the field.

Two consequences follow. First, capture should begin as early as possible, because work not recorded cannot be recovered later. Second, the right to learn from customer work data is a contractual matter that must be settled before deployment, not after. It is a clause, not a feature, and it determines whether the flywheel is legally able to turn.

5. What decides whether this can be deployed

A system of this kind lives or dies on four practical constraints that have nothing to do with the quality of its recognition.

Hardware, by environment

There is no single device that covers every industrial setting. Deployment is a fleet decision, not a product decision.

Environment	Practical hardware today	Role
Training centre, safe indoor	Immersive headsets	Highest visual quality; best for learning and rehearsal
Field, non-hazardous	Phones and ruggedised tablets	The deployable path for recognition and capture
Hazardous area, Zone 1	Intrinsically safe phones and tablets; assisted-reality headsets	Hands-free procedures where certification governs
See-through AR glasses	Pilot only	Promising form factor; certification and spatial quality not yet there

Content should therefore be built once against an open standard and deployed across device classes, rather than rebuilt per headset.

Offline operation

Most of the field has no reliable signal. A system that requires connectivity to identify a component is a system that works in the office and fails at the wellhead. Recognition models, procedures and assets must be able to run on the device, with captured work synchronised when a connection returns.

Certification

A hazardous-area site will not authorise thousands of uncertified devices. Intrinsic safety certification is not a refinement to be added later; it is the gate through which any large field deployment must pass, and it operates on procurement timescales rather than software ones.

Data residency and clearance

Large industrial operators increasingly require in-country hosting and security clearance for systems that touch operational data. These are calendar items, not engineering ones, and they should be started in parallel with development rather than after it.

6. What is demonstrated today

Honesty about maturity is more useful to a serious buyer than enthusiasm.

Working and verified on device: recognition of industrial equipment from a live camera view, with components identified and returned to the device; a full multi-step procedure with per-component references loaded from an authoring system; a three-dimensional asset rendered and anchored in the real environment, operating without connectivity once loaded.

In development: symbol overlay in place of text labels; knowledge portals behind each recognised component; integration of simulation and captured expertise; and assessment capture into the learning loop.

Not yet addressed: hazardous-area certification of a specific device fleet, and the contractual framework for learning from customer work data.

7. Conclusion

The industrial sector does not need more training content. It needs the content it already has to arrive at the right equipment, at the right moment, in a form a working person can act on without stopping.

Recognition makes that possible without modelling the world first. Symbols make it legible without teaching anyone a new language. Assessment makes it improve rather than merely repeat.

The organisation that captures how work is actually performed, and is permitted to learn from it, will hold something that no general-purpose model provider can obtain: an accurate account of how its industry really does the job.