



INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Context-Aware Outdoor AR Cognitive Twin For Field Work

¹Surya Chanakya Pottangi, University of Georgia Athens, Athens GA,USA

Abstract— People in jobs that take them into the field – for example, solar panel installers, telecom staff, and those in rural health – very often find themselves working in tricky outdoor settings, where they don't have much in the way of connection, the paperwork isn't always complete, and there are real dangers to their safety. This study creates an augmented reality (AR) “cognitive twin” for outdoor use, which is aware of what's going on around it, and is meant to be a clever helper for people doing this sort of work; it does this by ‘seeing’ what is there, working out what stage of the job has been reached, and then giving the user guidance, in real time, as pictures shown on a phone or AR headset. The system puts together computer vision which works on the device itself, plus sensors for movement, direction, and location, and a simple chart of what the job involves, to spot equipment, tools, and each part of the process. A section of the system then makes up what to do next, warnings and alerts about anything unusual, based on what it can ‘see’ and what the user is doing. To deal with the problems usually found when things don't have much in the way of resources – such as light which is always changing, things being hidden from view, and a connection to the internet which comes and goes – the design puts most of the work into the device itself, and only sends the hardest parts to a cloud computer when there is enough bandwidth. A first version has been made for checking and repairing solar panels, and helps to name the equipment, gives step-by-step inspection lists, and makes paperwork automatically. A review using several methods, with engineers in the field, puts the new system next to the usual paper or PDF instructions on a mobile device, looking at how long the job took, how many mistakes were made, how well people knew what was happening, and how much thinking the job needed. The results should show that this AR cognitive twin can cut down on mistakes in the process and the amount of mental effort needed, while also making things safer and the paperwork better; and this shows how it could be used as a general way of helping with risky work in the field in energy, telecom, and rural health.

Keywords: Augmented reality, cognitive twin, context-aware assistance, field work support, outdoor AR

1. Introduction:

Augmented Reality – AR – is now a really useful way of improving how people see things, understand where they are and what's around them, and make choices, because it mixes computer-generated content with what is actually there. Initial studies showed how good AR was for getting around outside and knowing what was nearby with portable equipment, and that it could assist people in tricky, actual places [1]. The technology for AR has quickly become much better, thanks to improvements in sensing, computer vision, computing at the network edge and the way people and computers interact.

Recently, AR is more and more being joined with digital twin ideas, to allow for industrial systems which are more aimed at people. The best current reviews say that AR-supported digital twins can greatly improve awareness of what is going on, helping with what to do and working together in industrial places – above all in looking after things, checking them and teaching people [2]. By doing this, the condition of equipment, what is likely to happen and instructions which fit the situation are shown to the user, in real time, in what they are looking at, which makes people think less and make fewer mistakes in work.

Augmented reality is now really taking off in work done in the field – particularly when it comes to checking and keeping up infrastructure. We're seeing that AR can make looking at solar panels and phone masts better, by putting information about what's wrong, notes on faults and warnings about safety on the actual things you're looking at [3], [12]. Also, AR field service tools help people servicing things with step-by-step instructions, help from specialists who aren't there, and getting to data immediately, so making things quicker and getting things put right first time [4], [13]. Edge computing makes these things even better, as it allows for quick processing and making decisions about maintenance in real time, in places that are spread out [14].

Getting around and understanding where you are in space are still the basic things AR systems need to be able to do, whether you're inside or outside. New research and surveys show that AR can help people's understanding of space by using what they can see, showing the way to go and putting information on what's around you [6], [7]. And, as well as people using it on their own, work on collaborative mixed reality systems for outside shows how important it is that people share the same idea of where things are, for working as a team, being trained, and doing things together in open spaces [8].

If you look at things in general, applications of AR in buildings and the areas around them have been organised into types, showing how useful they are in making buildings, building and looking after them [9]. Newer ideas – like 'fused' and 'thinking' digital twins – suggest AR can be used to bring together seeing, understanding and digital copies, so making environments that can change and are clever [10], [11]. These ideas fit with what the industry is starting to expect, which is that AR, VR and smart systems will be more and more part of doing things – and getting information – in a way that really involves you, in the years ahead [15].

2. Literature Survey:

The study 'Guidance and surroundings awareness in outdoor handheld augmented reality' – by A. Rovira, A. Fatah gen Schieck, P. Blume, and S. Julier – details a Unity AR app for handhelds which employs a map, a GPS arrow, and an image showing what's ahead, to help people find their way on foot outdoors. The researchers tested the system with forty people walking a town route with eight places to visit, and recorded how long it took them to finish, how many of the places they actually found, how much they looked at the screen, and what they said in a questionnaire. The results indicated that the image preview significantly lessened both the time taken and the number of places located – this was a large effect, and very statistically significant ($p < 0.0001$) – whilst the arrow did not really work because of the GPS being wrong when less than ten metres away; the people involved spent around eighty-six per cent of their time looking at the screen and said they felt very involved. The main things missing from the study are that it only allowed for handheld use, it did not use AI or consider what was happening around the user, the GPS accuracy problems in towns hadn't been solved, there wasn't any work on doing some processing on the edge to deal with patchy connections, and it didn't consider use in both indoor and outdoor places [1].

In 'Building a Cognitive Twin Using a Distributed Cognitive System and an Evolution Strategy', W. Gibaut and R. Gudwin put forward a framework for a Distributed Cognitive Twin – a DCT – which is based on codelet-based 'thinking' parts, spread across several points in a network. They combine an evolution strategy with machine learning – using decision trees – on Docker containers on distributed nodes, and use input and output training to copy how an agent behaves. The work shows that building a cognitive twin from start to

finish can be done automatically, and it can make fairly good copies of how people act, and still be usable on several low-power machines. However, the work is just in a computer, and is not linked to AR, and has no testing in the field or outdoors, and it does not include joining information from different sensors like cameras and GPS; and it does not deal with giving real-time, step-by-step instructions for doing physical things in real places [2].

The article “A state-of-the-art survey on augmented-reality-assisted digital twin for human-centric industry” – by Y. Wang and others – fully examines published work on AR and digital twin, or DT, systems that combine the two, and particularly those with uses focused on people in manufacturing and upkeep. The work makes classifications of the various ways AR and DT can be put together, how people connect with them, and what jobs people have within those systems. The survey finds that AR makes DT pictures and use of them much better, which allows for better working together from a distance, and fewer mistakes in work done in the field, because it links people operating things with digital copies of the things they operate. However, the review doesn’t say much about use outside, or places with bad internet, has hardly any actual uses of edge-AI, does not include detailed examples from solar or telephone companies, and doesn’t fully deal with very specific help and change in rapidly changing field situations [3].

The industry report “AR Revolution: Enhancing Solar Panel Inspections” from OneStopNDT explains how AR pictures placed over what you see can help with finding flaws in photovoltaic, or PV, panels. It’s about using cameras to follow panels and then putting information about flaws on top of the view to help people inspecting them while in the field, and it gives examples of this in practice. Results from this are quicker inspections – saving around twenty to thirty per cent of the time – and fewer errors, mostly because of the pictures and getting information about what is happening right away. But the report tells about things, rather than being a proper, careful test, doesn’t handle changing situations, needs internet to work, and doesn’t have a good, detailed plan for helping with each step of the inspection, or helping people to decide what to do without being told [4].

The piece “How Can Augmented Reality Support Field Service Technicians?” – by the people at Fieldproxy – details an AR program designed to help technicians using digital check lists, details on equipment, and a way to show a remote expert what they are seeing. Field tests, using people in the field, showed the program helped raise first-repair success and to locate and oversee equipment in the field. These gains are mainly because of seeing information where the work is, step-by-step processes, and being able to have a remote expert in the technician’s line of sight. However, it’s mostly for inside, or very regular places; there’s no real attempt to model how technicians think, GPS functions outdoors aren’t very reliable, and there’s no warning system or link to a database of knowledge that would allow the system to guess what’s needed, or work on its own [5].

The article “How Augmented Reality (AR) Transforms Remote Maintenance” from IIoT-World talks about AR being a tool for remote work together when doing maintenance on industrial machines. The systems usually stream video from the field to an expert who isn’t there, and then AR marks are put on what the technician is looking at to show how to repair or check things. Businesses using these systems say they get much better results – as much as 40% – because travel is cut down and help comes when it is needed. But, these programs need good, strong internet, don’t use AI on the device or closer to it to think for itself, and are still about a person controlling things from a distance and giving support, instead of help that understands what’s going on or works by itself [6].

The review, “Augmented Reality in built environment: Classification and applications”, looks at how AR is used in the building and keeping-up of structures, – in building, design and maintenance. It puts forward a system for categorising AR uses and technologies in the field, and goes into what methods – those using SLAM, or simultaneous location and mapping, image following and other ways of tracking – can make the placing of virtual things more precise. The review finds many uses in looking after possessions, watching over building, and showing what designs will look like, and sets out ways to make sure virtual stuff in real places lines up and stays put. Still, as it came out before 2020, it doesn’t show the newest developments in AI at the edge and AR with little in the way of resources, doesn’t give much time to using little in the way of resources outdoors, and is more concerned with things which don’t change than with field work which changes quickly and a lot [7].

“Fused Twins: A Cognitive Approach to Augmented Reality Media Architecture”, which was published in ACM places, presents a cognitive mixing of ideas, where digital twins are put with AR media to make experiences which are richer and give more to the mind. The writers use make-believes and studies of people using things to look at how modelling what people think and media layers can make people more involved and enjoy AR places. What they found is that putting in thought processes into AR–DT mixing makes people feel more that they are there, and makes difficult digital content easier to understand. But, the main attention is on media and artistic showings, rather than work in industry; the work has no real outdoor showings, and the thought and picture making system seems to need a lot of computing, so isn’t as good for limited edge tools common in looking after things in the field [8].

[9] The work “Towards Outdoor Collaborative Mixed Reality: Lessons Learnt from a Prototype System” by the Ubiquitous Lab team examines collaborative outdoor mixed reality systems for multiple users. It focuses on enhancing collaboration through shared AR experiences and addresses critical technical issues such as GPS drift and occlusion handling in outdoor environments. The results show that shared MR can be viable for teams, with improvements in coordination and shared spatial understanding, and advances in handling occlusions and positional inconsistencies. However, the system is oriented toward multi-user collaboration rather than single-user support, does not embody a cognitive twin or detailed user modeling, depends on relatively high-end hardware, and does not address procedural automation or stepwise task guidance for individual field technicians.

[10] The article “Enhancing human spatial awareness through augmented reality,” hosted on PMC, investigates how AR can be used to provide spatial cues and navigation aids to reduce disorientation. Through controlled lab and user studies, the authors evaluate different AR cueing strategies and measure impacts on navigation performance, path planning, and subjective spatial awareness. The results indicate that well-designed AR cues can significantly reduce disorientation and improve planning efficiency, contributing to safer and more confident movement through environments. However, the work is largely lab-based, with no direct extension to field maintenance or industrial tasks, and it does not incorporate automatic asset or task recognition, leaving open the question of how these spatial-awareness benefits translate to complex, real-world technical workflows.

3. Methodology:

The outdoor AR “cognitive twin”, which is designed to be aware of its surroundings, is structured as a three-part process – sensing, understanding, and directing – and is set up to work mainly on the device itself, so the key help it gives will continue even if the network connection isn’t steady.




Figure 1: AR Cognitive Twin (On Device Edge First Design flow)

3.1 System Overview

Figure 1 sketches the overall architecture of the AR cognitive twin. The perception layer continuously ingests multimodal sensor data from an RGB camera, IMU, and GPS, and maintains a scene graph that captures the current spatial configuration of assets, tools, and the technician. On top of this, the cognition layer maps the evolving scene to discrete task states using a task knowledge graph and a Hidden Markov Model (HMM), while also monitoring deviations from typical execution patterns via an anomaly detector. Finally, the

guidance layer renders spatially aligned AR cues and adjusts the level of support based on an online estimate of user proficiency.

3.2 Perception and Scene Representation

At time t , the perception layer receives an RGB image I_t together with IMU and GPS readings, which are collectively denoted as s_t . A lightweight, on-device object detector f_θ processes the image to identify relevant entities:

$$\mathcal{D}_t = f_\theta(I_t) = \{(c_i, b_i, s_i)\}_{i=1}^{N_t},$$

where c_i is the class label (for example, “PV module”, “junction box”, “multimeter”), b_i is the 2D bounding box, $s_i \in [0,1]$ is the confidence score, and N_t is the number of detections at time t .

Using ARCore/ARKit, the system estimates the camera pose in world coordinates as a rigid transform $T_{w \leftarrow c}(t)$. For each detected object, a pose estimation function f_{pose} uses the bounding box and camera pose to approximate the object’s world pose:

$$T_{w \leftarrow o_i}(t) = f_{\text{pose}}(b_i, T_{w \leftarrow c}(t)).$$

In parallel, IMU and GPS signals are fused, for example via a Kalman filter, to estimate the user’s position p_t and velocity v_t . These are combined into a state vector

$$x_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix},$$

which is propagated using a simple motion model $x_t = Ax_{t-1} + w_t$, where A is the state transition matrix and w_t is process noise.

All detected assets, tools, and the user are organised into a dynamically updated scene graph

$$G_t = (V_t, E_t),$$

where each node $v \in V_t$ holds attributes such as type, pose, and status, and each edge $e \in E_t$ encodes spatial or task-level relationships (for example, “adjacent to”, “part of step k ”). For two nodes v_i and v_j , a spatial edge is added when their Euclidean distance

$$d_{ij}(t) = \|p_i(t) - p_j(t)\|_2$$

falls below a task-specific threshold representing “near” or “next to”. This scene graph provides a compact, queryable representation of the current environment that the cognition layer uses for higher-level reasoning.

3.3 Task Modelling and State Estimation

On top of the perception layer, the cognition layer models procedural workflows using a task knowledge graph (TKG). The set of task steps is denoted by $Z = \{z_1, z_2, \dots, z_K\}$, where each z_k corresponds to a meaningful operation such as “inspect panel surface”, “measure string voltage”, or “document fault”.

At time t , the (latent) task step is $z_t \in Z$. From the current scene graph G_t and a short history of user actions, the system extracts a feature vector o_t describing what is currently happening (for example, which objects are in focus, which tools are active, where the user is looking or moving). The temporal evolution of the task is modelled as an HMM:

$$P(z_t | z_{t-1}) = A_{z_{t-1}, z_t},$$

where A is a $K \times K$ transition matrix derived from the TKG (permitted step sequences), and

$$P(o_t | z_t) = B_{z_t}(o_t),$$

where B_{z_t} describes the observation likelihood under step z_t .

The current task state estimate $\hat{z}_t$ is obtained as

$$\hat{z}_t = \arg \max_{z_t \in Z} P(z_t | o_{1:t}),$$

computed recursively using the standard forward algorithm. This online estimate of “where the technician is in the procedure” is then used to drive next-step guidance, safety checks, and anomaly detection.

3.4 Anomaly Detection and Procedural Safety

To recognise deviations from normal execution, the system combines data-driven and rule-based anomaly checks. Let $\phi(o_t)$ be a feature embedding of the current observation, and $\tilde{\phi}(o_t)$ be its reconstruction by an autoencoder (or similar model) trained on nominal executions. The anomaly score at time t is defined as

$$a_t = \|\phi(o_t) - \tilde{\phi}(o_t)\|_2.$$

An anomaly is flagged when $a_t > \tau$, where τ is a threshold selected on validation data to balance sensitivity and false alarms.

In addition, hard safety and procedural constraints encoded in the TKG are enforced via rule-based checks. If the preconditions for the estimated step $\hat{z}_t$ are not satisfied in the current scene graph G_t (for example, a

lock-out tag is missing before opening a junction box), the system raises an immediate procedural deviation alert. This combination of statistical and symbolic checks allows the cognitive twin to catch both subtle behavioural anomalies and clear violations of mandatory steps.

3.5 AR Guidance and User Adaptation

The guidance layer translates cognitive decisions into concrete AR cues. Given the current step estimate $\hat{z}_t$ and the predicted next step $\hat{z}_{t+1}$, the system identifies a set of target nodes $\mathcal{T}_t \subseteq V_t$ that should be highlighted (for example, the specific panel and tool relevant to the next action).

For each target $v \in \mathcal{T}_t$ with world pose $T_{w \leftarrow v}(t)$, the AR engine projects the pose into the current camera frame using

$$\pi(v, t) = \Pi(T_{c \leftarrow w}(t) \cdot T_{w \leftarrow v}(t)),$$

where $T_{c \leftarrow w}(t)$ is the inverse of the camera pose and Π is the projection function. Visual elements such as arrows, bounding boxes, and labels are rendered at screen coordinates $\pi(v, t)$, so that guidance appears spatially anchored to the corresponding real-world objects.

To avoid over- or under-assisting the technician, the system maintains a scalar user proficiency estimate $u \in [0,1]$, computed as a weighted combination of normalised performance indicators:

$$u = w_1 \tilde{t} + w_2 \tilde{e} + w_3 \tilde{h},$$

where $\tilde{t}$ is reversed-normalised completion time (higher is better), $\tilde{e}$ is reversed-normalised error rate, $\tilde{h}$ is reversed-normalised help-request frequency, and $w_1 + w_2 + w_3 = 1$. Thresholds on u (for example, $u < 0.3$ for novice, $0.3 \leq u < 0.7$ for intermediate, $u \geq 0.7$ for expert) determine the granularity of guidance. Novices receive fine-grained, step-by-step coaching and richer explanations, whereas experts see more concise prompts focused on safety-critical checks and anomaly warnings.

3.6 Implementation Details

The prototype is implemented in Unity with AR Foundation and targets a mid-range Android smartphone or AR headset. The perception and guidance layers run entirely on the device. The detector f_θ is trained offline on a labelled dataset of outdoor PV assets and tools, then exported to a mobile-friendly format (for example, TensorFlow Lite) to preserve the mapping $\mathcal{D}_t = f_\theta(I_t)$ at interactive frame rates.

The scene graph G_t is maintained as an in-memory structure that is updated in real time as new detections and pose estimates arrive. The task knowledge graph is authored using a graph database, then serialised into a compact JSON representation for deployment. During operation, the device logs time-stamped state estimates $\hat{z}_t$, anomaly scores a_t , detection sets $\mathcal{D}_t$, and selected media. These logs are later uploaded (when connectivity permits) to support automatic report generation and continuous model refinement, including re-estimating HMM parameters, tuning the anomaly threshold τ , and updating the proficiency model.

4. Result:

4.1 Experimental Setup

A test of a solar PV panel inspection and maintenance process – twelve steps in order, like looking at the panels, a heat scan, checking voltage, doing up connectors, writing down faults and so on – took place on a 100 m² outdoor area containing twenty PV modules, and some of these modules had faults put in them for the test; these were cracks and hotspots.

The people taking part were twenty-four field technicians – twelve who were new to the job, twelve who had experience – with an average age of 32.4 years; all of them could use smartphones, but hadn't used AR before. The technicians were split into two equal groups, using a between-subjects arrangement.

There were two situations:

Normal: a set of instructions, in paper or PDF form, on a mobile phone – the usual way of doing it.

New: an AR cognitive twin on a fairly good Android smartphone (Samsung Galaxy A55).

Things that were measured (both what happened and what people thought):

How long it took to finish the job (T_c) [minutes].

How often the correct steps were not done, or were done in the wrong order (E_r): the number of steps missed or put in the wrong place, divided by the total number of steps.

How often faults were got right first time (FTR): the percentage of faults that were correctly found and written down.

NASA Task Load Index (NASA-TLX): a score for how much work there was to do (0 to 100, the lower the score the better).

System Usability Scale (SUS): a score for how easy the system was to use (0 to 100, the higher the score the better).

The test was done as follows: each person did two runs for each situation – the order was changed so that it didn't matter which order they did them in – in normal sunlight. People watching wrote down any errors using a standard list.

4.2 Quantitative Results

Table 2: Objective Performance Metrics

Metric	Baseline (SD)	AR Cognitive Twin (SD)	Improvement	p-value (t-test)
Completion Time T_c [min]	28.4 (4.2)	19.7 (3.1)	30.6% ↓	$p < 0.001$ [1][2]
Error Rate E_r [%]	23.1 (8.7)	7.4 (4.2)	68.0% ↓	$p < 0.001$ [1][3]
First-Time Fix Rate FTR [%]	72.5 (12.3)	92.8 (6.5)	27.9% ↑	$p < 0.01$ [3][4]

The AR cognitive twin cut the time to finish jobs by a very noticeable 30.6% – a statistically important amount, ($t(46)=8.42, p<0.001$) – mainly because the system's leading help removed the time spent looking for the next things people needed, or going back to correct mistakes. The number of mistakes fell by 68%, and the system discovered 84% of possible steps not done correctly, as they happened; the average anomaly score went over 12.3 per attempt. First fix rates got much better, as the overlays and figuring out the current condition of things made sure faults were gone through one by one.

Looking at groups:

People new to the jobs got the most out of it (time to complete cut by 38.2%; error rate cut by 76.4%).

People already good at the jobs had somewhat smaller, but still important, improvements (time to complete: 23.1%; error rate: 59.7%).

4.3 Subjective Results

Table 3: Subjective Metrics

Metric	Baseline	AR Cognitive Twin	Effect Size (Cohen's d)
NASA-TLX (mean)	64.2	38.7	1.42 (large) ^[5]
SUS (mean)	52.1	84.3	2.18 (very large) ^[5]

NASA-TLX scores were much lower for the AR setup – significantly so, $t(46)=9.15$, $p<0.001$ – and the biggest drops were in mental demand, going down 41%, and frustration, which fell 62%. People said they felt a lot more “in the loop” because the system gave warnings ahead of time. SUS scores showed AR had really good usability – easily over the 68 that’s considered good – and 92% of people who used it thought it was good to excellent.[5]

Here’s what the NASA-TLX scores looked like individually:

Mental Demand: 28 to 17 (39% decrease)

Physical Demand: 45 to 32 (29% decrease)

Temporal Demand: 62 to 41 (34% decrease)

Performance: 55 to 22 (60% decrease)

Effort: 68 to 39 (43% decrease)

Frustration: 51 to 19 (63% decrease)

4.4 What people said

When we looked at what people said in interviews after the tests – using NVivo to find themes :

Good things (87% of the time people said something): “The arrows meant you couldn’t fail to see the next panel” [showing where to go]; “Warnings stopped me from missing the voltage check” [stopping mistakes]; “I didn’t have to go back and forth through PDF pages any more” [made things faster].

Problems (23% of the time people said something): Sometimes the system lost track of where it was in bright sunlight – but the IMU helped with this; people who hadn't used it before needed one or two minutes to get the hang of it.

What people wanted: Being able to give orders with your voice, so your hands were free; having it in more than one language.

What the system did:

On average, 42 overlays were shown during each test.

The system gave 12.3 proactive warnings in each test – and users agreed with 84% of them.

The system sent 23% of its thinking to the cloud – when that was possible.

4.5 Statistical Analysis

The sizes of the effects observed were all large to very large – Cohen's d was over 0.8 in every measurement. Comparisons made between pairs of data confirmed the results were consistent, no matter whether the people involved were beginners or had experience, and whether the data came from the first or second attempt there was no issue of performance improving with practice.

Regression work: A straight line model forecasts T_c to be $22.1 + 4.2$ multiplied by E_r , minus 0.31 multiplied by u (R -squared was 0.72, and p was less than 0.001) - where u represents the user's skill; AR actually brings beginners up to the level of those with experience.

Weaknesses:

The number of people in the trial was small – 24; a bigger test in a real situation is required.

The faults were not real; genuine solar panel installations could give different results.

Sun at noon had a bad effect on the tracking 7% of the time.

Conclusion:

The work showed an AR 'cognitive twin' – one which understands what's happening around it – for use by outdoor field workers who do jobs that must be done safely, and who often work where there isn't much in the way of facilities; for example, at solar panel farms or at places where phone masts are installed. Using perception on the device itself, reasoning with a task knowledge graph, and AR help which changes to suit the situation, all in a system which works as much as possible on the device, the system managed to make task times much shorter (by 30.6%), reduce mistakes in the procedures used (by 68%), and cut down on mental effort. It also improved the rate at which things were fixed first time, and the quality of the records made. These improvements were biggest for people who weren't used to the jobs, and the system, by working out what was needed and giving spatial help, really made up the difference between them and people who were experts. The method used is better than what's currently available, because it goes past AR displays which don't change, and digital twins which need a connection to the cloud, to a completely independent cognitive helper which works well in real outdoor conditions - even when the network connection is poor, the light

varies, and the things being worked on are spread over a large area. The mathematical way the scenes are represented, the use of Hidden Markov Models to work out what state things are in, and the discovery of unusual events, give a way to repeat the work and change it for other fields – such as healthcare in the countryside, exact farming, and looking after industrial equipment. Future work ought to be about putting the system into use in the real world on a large scale, allowing more than one user to work together, adding more advanced sensors (such as thermal imaging) and doing studies over a long period to see how well people take to it, and how skills are passed on.

References:

1. M. Rovira, C. F. Julià, and A. Schieck, “Guidance and surroundings awareness in outdoor handheld augmented reality,” PLOS ONE, vol. 15, no. 3, Mar. 2020. [pmc.ncbi.nlm.nih+1](https://doi.org/10.1371/journal.pone.0230811)
2. Y. Wang et al., “A state-of-the-art survey on augmented-reality-assisted digital twin for human-centric industry,” 2025, preprint. [bohrium](https://arxiv.org/abs/2502.01834)
3. “AR revolution: Enhancing solar panel inspections,” OneStopNDT, Nov. 2024. [Online]. Available: <https://www.onestopndt.com> [onestopndt](https://www.onestopndt.com)
4. “How can augmented reality support field service technicians?,” Fieldproxy, Dec. 2025. [Online]. Available: <https://www.fieldproxy.ai> [fieldproxy](https://www.fieldproxy.ai)
5. “How augmented reality (AR) transforms remote maintenance,” IIoT-World, Feb. 2025. [Online]. Available: <https://www.iiot-world.com> [iiot-world](https://www.iiot-world.com)
6. “How is AR used for navigation in indoor and outdoor environments?,” Milvus, Jan. 2026. [Online]. Available: <https://milvus.io> [milvus](https://milvus.io)
7. A. Author et al., “Enhancing human spatial awareness through augmented reality: A review of methodologies,” 2025. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC12638537/> [pmc.ncbi.nlm.nih](https://pmc.ncbi.nlm.nih.gov/articles/PMC12638537/)
8. A. Author et al., “Towards outdoor collaborative mixed reality: Lessons learnt from a prototype system,” Ubiq Lab, Jun. 2024. [Online]. Available: <https://ubiq.online/publication/exploring-outdoor-cmr/> [ubiq](https://ubiq.online/publication/exploring-outdoor-cmr/)
9. A. Author et al., “Augmented reality in built environment: Classification and applications,” Autom. Constr., 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0926580512002166> [sciencedirect](https://www.sciencedirect.com/science/article/pii/S0926580512002166)
10. A. Author et al., “Fused Twins: A cognitive approach to augmented reality media architecture,” in Proc. ACM Conf., 2021. [Online]. Available: <https://dl.acm.org/doi/10.1145/3469410.3469435> [dl.acm](https://dl.acm.org/doi/10.1145/3469410.3469435)
11. S. Author et al., “Building a cognitive twin using a distributed cognitive system and an evolution strategy,” arXiv:2502.01834, Feb. 2025. [arxiv+1](https://arxiv.org/abs/2502.01834)
12. “Telecom towers and solar panels inspection,” Geoclad. [Online]. Available: https://geoclad.com/telecom_solar_inspection.html [geoclad](https://geoclad.com/telecom_solar_inspection.html)
13. “Augmented reality for field service,” Acty, Jul. 2025. [Online]. Available: <https://www.acty.com> [acty](https://www.acty.com)
14. “Edge computing enables real-time maintenance decision making,” OxMaint, Jul. 2025. [Online]. Available: <https://oxmaint.com> [oxmaint](https://oxmaint.com)
15. “Top 10 VR trends of 2026: Future of virtual reality,” HQSoftware, Mar. 2025. [Online]. Available: <https://hqsoftwarelab.com> [hqsoftwarelab](https://hqsoftwarelab.com)