Turning blade data into smarter wind decisions

Words: Armando Costa, Founder & CEO, ArthWind

Inspection has become standard practice across wind fleets, but it is no longer the differentiator. The next phase of blade management will be defined by decision quality data, consistent analysis and models capable of translating condition insights into clear, risk-based actions.

We underestimated the blade. If you look back at how the wind industry approached asset reliability in its early years, blade operations & maintenance (O&M) barely registered as a priority problem. That was a mistake and an expensive one.

The pace at which turbine technology advanced left maintenance protocols behind. Blades grew faster than our understanding of how to manage them at scale.

Part of that is structural. Blades absorbed enormous cost pressure during manufacturing while simultaneously scaling in size and complexity. When you push a composite structure to its limits on both dimensions at once, you get consequences that don't show up immediately; they show up three or four years into operation, as progressive internal damage that nobody mapped.

Today blades are exceeding 70, 80, even 100 metres. A failure at that scale is not a maintenance event. It is a logistics crisis, a generation loss and in some cases, a safety incident. And yet, across a significant portion of the industry, inspection is still treated as a discretionary cost item rather than a core component of risk management.

The question I hear too often is whether to inspect annually or every two years. That is the wrong conversation entirely. The right question is how many inspection cycles per year give you the risk visibility you actually need to make intelligent decisions with your maintenance budget. Those are very different problems.

Armando Costa

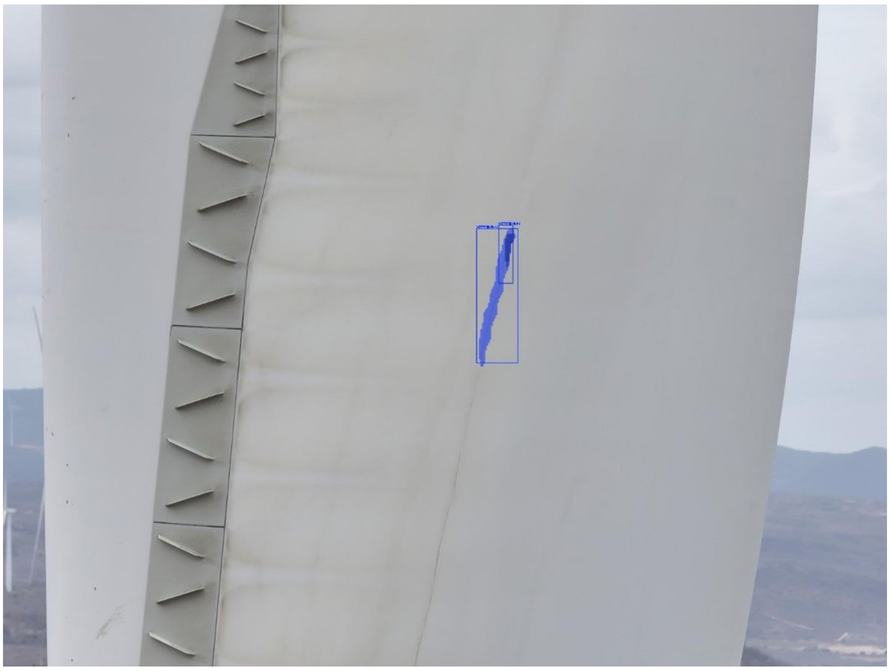

More data is coming. That is not automatically good news

The last several years have brought a real shift in how we collect blade condition data. Drones and internal inspection rovers are now established tools. Image-based inspection has become the industry standard for structural health assessment. Volumes have grown significantly and that is genuinely useful progress.

But the next wave is going to be much larger. Drone and rover technology is commoditising fast, both through service providers and through operators who are bringing inspection in-house. The data volumes coming in the next few years will be unlike anything the industry has processed before.

The problem is not storage or processing speed. Most inspection infrastructure was never designed to make decisions at scale and handling thousands of inspections with complex, multi-variable datasets requires exactly that.

More data does not automatically produce better decisions. It produces more noise unless the analytical infrastructure scales with it.

The competitive question for inspection companies is not who can collect the fastest. It is who can manage this incoming volume with enough precision to produce decisions that actually hold up under engineering scrutiny.

From where I sit, three specific patterns are working against that goal right now:

1. **Speed is killing image quality**

   The push to reduce inspection time has real consequences for capture quality: exposure, lighting and collection angle. These are not minor technical details. They are the parameters that determine whether a defect can be accurately characterised or even detected at all.

   A fast inspection with poor imagery is not cheaper than a slower one. It is just worse data, which produces worse decisions downstream.

2. **Outsourced labelling with no unified standard**

   Cost pressure pushed a lot of high-volume operations to cascade the labelling function across multiple outsourced layers. The result is that the training data feeding your AI models was classified by different teams using different criteria. You cannot build a reliable decision model on top of that foundation. The inconsistency is baked in from the start.

3. **Scale without product knowledge**

   Chasing volume without deep blade knowledge produces recommendations that are generically conservative and barely usable.

   A standard defect catalogue does not apply to every blade type and model. The operators I work with know this. They have received reports that flag the same finding as critical on one platform and low-priority on another. That is not an AI problem. That is a training data and product knowledge problem.

What nine years of keeping analysis in-house actually produces

At ArthWind, we made a deliberate decision early on to keep 100% of the data analysis process internal. Not because it was cheaper, quite the opposite. But because we understood that the value of inspection data is directly tied to the interpreter's knowledge.

Knowing what you are looking at inside a blade requires understanding how it was

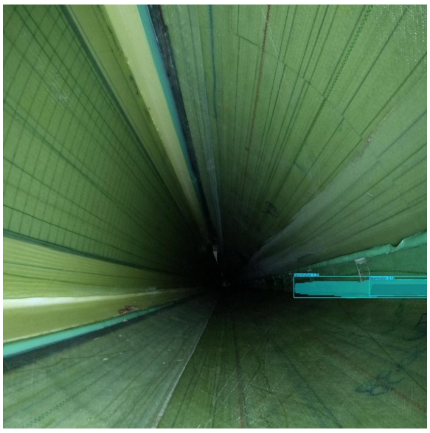

built. Markets with deep roots in blade manufacturing, where generations of engineers developed expertise in composite materials, structural behaviour, and large-scale production, produce a fundamentally different quality of analysis than teams trained primarily on pattern recognition of images.

That manufacturing knowledge is what tells you whether an anomaly is a defect or a residue from the production process. It is not something you can train in quickly.

Over nine years, that approach has resulted in over 130,000 blade inspections and a database of millions of catalogued points of interest, each one linked to blade model, campaign, and progression history. In 2026 alone, we crossed 12,000 inspections in a market that is not seeing significant new capacity additions. That number is worth pausing on.

The growth is not coming from new turbines going up. It is coming from operators who have stopped treating inspection as a response to problems and started treating it as a standard monitoring practice. That is a real shift in how the industry thinks about risk management.

The hard problem is not detection. It is the recommendation

Building an algorithm that identifies defects in blade images is not particularly difficult anymore. The technology is accessible. What remains genuinely hard and what I rarely hear discussed at the industry level, is the recommendation engine. The step between identifying what you found and telling an operator what to do about it.

Defect categorisation standards vary widely across companies and campaigns. Without a consistent, traceable classification logic, engineering recommendations default to conservative generalisations that disconnect from the asset's actual condition and history. Operators end up with reports that don't help them prioritise, and expensive maintenance decisions are made on incomplete information.

The methodology we have developed and used since 2019, KIN (Keep, Increased and New), is our answer to this. Every defect found in an inspection gets classified into one of three states relative to the previous campaign: it stayed the same (Keep), it got worse (Increased), or it is a new finding (New). That longitudinal tracking is what transforms inspection data into something usable for risk-based decision-making.

Severity becomes a function of progression, not just current appearance. A defect that looks serious but has been stable for three years carries a different risk profile than one that has doubled in size in twelve months. A static defect catalogue cannot capture that. A model trained on longitudinal, progression-linked data can.

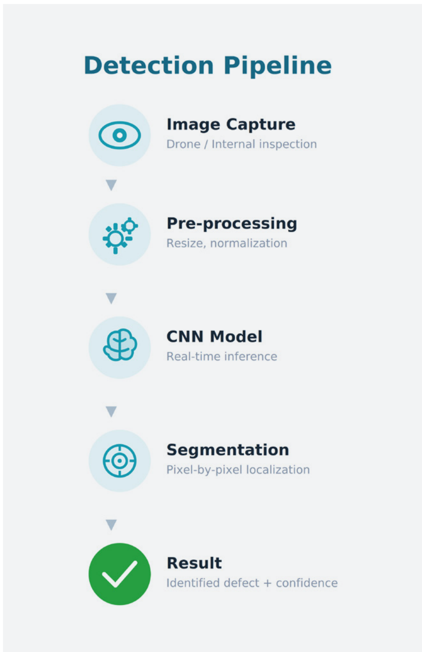

If we are serious about cost reduction and operational reliability in wind energy, the conversation needs to move from inspection frequency to model quality.

When you train AI on data structured this way, the output changes fundamentally. The recommendation is no longer a label attached to an image. It is an engineering decision grounded in how that defect has actually behaved, traceable, auditable and connected to what happened in previous campaigns.

erode is the quality of the analytical foundation underneath the system, the dataset, the labelling consistency and the product knowledge embedded in the model.

The organisations that define the next decade of blade O&M will be those whose models can

conversation needs to move from inspection frequency to model quality. Inspect more often, yes, but only if the data coming out of those inspections is precise enough, consistently interpreted and feeding a recommendation engine that actually understands the product it is evaluating.

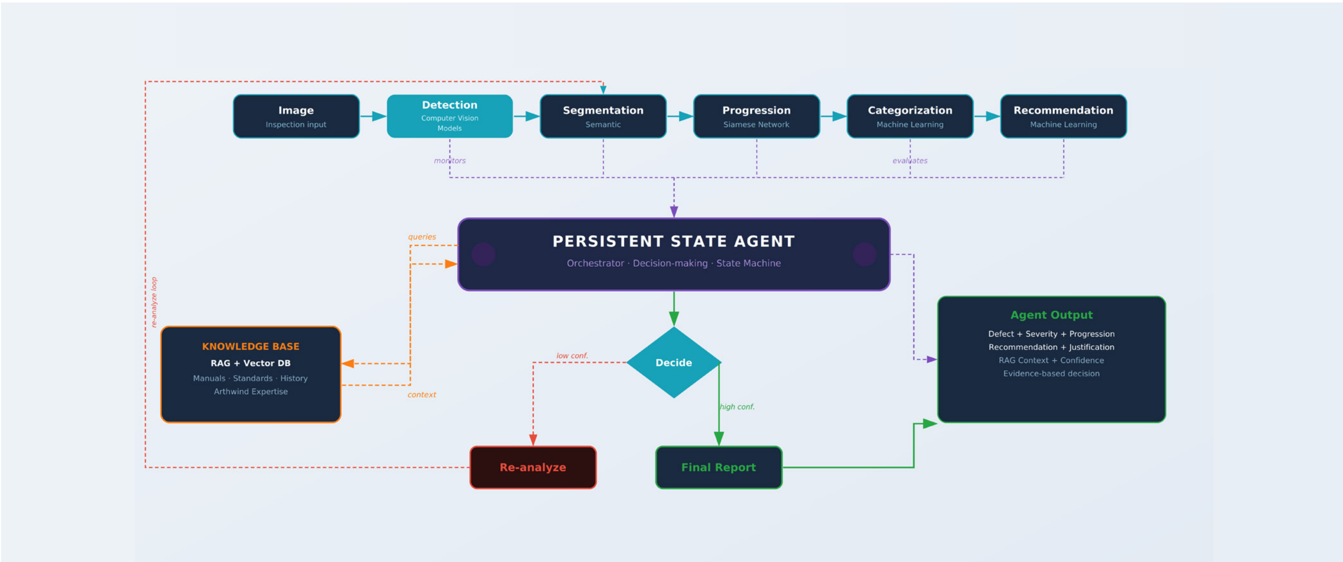

That is the standard the industry needs to hold AI recommendations to.

The future belongs to decision models, not data collectors

As inspection moves toward remote and automated delivery, the question of who wins is not really about the scale of collection. Hardware gets cheaper. Drones get faster. Those advantages erode. What does not

answer operational questions autonomously: which assets need continuous monitoring, which ones should go into repair campaigns this year and how should a constrained maintenance budget be allocated across a multi-site fleet. These are not reporting questions. They are decision questions and they need decision-grade inputs.

If we are serious about cost reduction and operational reliability in wind energy, the

The companies that will get there are not necessarily the biggest or the fastest. They are the ones who spent the last decade building data infrastructure with enough technical rigour to train models that can actually decide. That is the work we have been doing at ArthWind, and the direction we are continuing to push in 2026.

arthwind.com.br