GEONADIR

Rehabilitation monitoring beyond snapshots

Dr. Karen Joyce
Co-founder & Head of Product and Science, GeoNadir

Contents

Introduction

Rehabilitation monitoring is often treated as an annual reporting obligation - something to be completed, documented, and filed away each year. But ecosystems do not recover in annual steps.

They respond to rainfall, disturbance, temperature, and management actions in ways that are gradual, uneven, and highly seasonal. When monitoring programs rely too heavily on snapshots or single reporting milestones, important context can be lost. We also know that timing can influence interpretation, meaning that our decisions may be shaped by anniversary appearance rather than trajectory.

Rehabilitation monitoring is often treated as an annual reporting obligation – something to be completed, documented, and filed away each year. But ecosystems do not recover in annual steps.

This ebook sets out GeoNadir's perspective on what effective rehabilitation monitoring looks like in practice - and why the dominant approaches often fall short. Across the following sections, we examine:

- why focusing on "seeing weeds" can distract from more meaningful questions
- what good rehabilitation looks like when assessed through time
- why analogue sites provide essential context
- how sampling design shapes the story we tell
- how different data sources can work together in a monitoring hierarchy
- and how reporting can move beyond snapshots toward evidence built over time

The goal is not to add complexity or workload. It is to strengthen the foundations for judgement, improve defensibility, and design monitoring systems that reflect how ecosystems actually behave.

At the end of this ebook, you'll find a short Rehabilitation Monitoring Health Check. It is intended as a practical tool to help you reflect on your own monitoring approach and identify small changes that could deliver greater confidence in outcomes.

Rehabilitation monitoring works best when it is intentional, contextual, and built on evidence that accumulates through time. The following pages outline what that can look like in practice.

Banner image by Leiada Krözjhen on Unsplash

01

Rehabilitation monitoring and the problem with “seeing weeds”

In almost every conversation about remote sensing for mine rehabilitation, **“seeing weeds”** quickly becomes the focal point.

“Can I see the weeds?”
“Can I count the weeds?”
“Which species are there?”

These sound like reasonable questions. But rather than take them at face value, I see them as shorthand for a bunch of concerns about risk, compliance, and confidence: *Will weed pressure undermine rehabilitation success? Will problems be detected early enough to intervene? Will there be defensible evidence if decisions are questioned?*

Those are legitimate concerns. But when “seeing weeds” becomes the dominant criterion for evaluating monitoring approaches, it often nudges rehabilitation programs in the wrong direction.

Why “seeing weeds” *feels* like the right focus

Weeds are tangible. You can point to them in the field. They feature prominently in management plans and risk registers. And for anyone who has watched a rehabilitation area fail because of weed dominance, they are an obvious warning sign.

It also doesn’t help that much of the current technology narrative reinforces this focus. High spatial resolution imagery, AI-driven classifiers, and automated species detection maps promise the ability to *see* more, in more detail, than ever before. It’s easy to assume that if we can identify weeds from drones, for exaple, we must be reducing risk.

But visibility and confidence are not the same thing.

What rehabilitation monitoring needs to achieve

Mine site rehabilitation is rarely about the complete absence of weeds. Regulators are not looking for pristine, weed-free landscapes. What they are looking for is evidence that a site is:

- on a credible trajectory toward the approved final land use;
- structurally and functionally stable;
- increasingly self-sustaining; and
- comparable, over time, to appropriate reference (analogue) sites.

These are outcomes assessed through...

 **patterns**

 **trends**

 **trajectories**

...not through single observations.

A rehabilitated area can contain weeds and still be progressing well. In early stages of succession, weeds can even be helpful in stabilising soils and preventing erosion. Conversely, a site can appear uniformly green in a snapshot and be fundamentally unstable. Perhaps it’s dominated by opportunistic species that respond quickly to rainfall but do little to build longterm ecosystem function.

When monitoring programs prioritise what is easiest to see at one moment in time, they can miss

The limits of snapshot thinking

Much of the emphasis on “seeing weeds” comes from a snapshot mindset: if a problem is present, it should be visible in an image – or through ground survey.

The trouble is that rehabilitation systems are dynamic. Seasonal growth and senescence, rainfall variability, establishment phases, and disturbance legacies all shape how vegetation looks at any given moment. A single image or site check, no matter how detailed, captures only a sliver of that story.

In practice, this can lead to two common pitfalls:

False reassurance
shortlived greenness
interpreted as success

False alarm
temporary weed presence
interpreted as failure

Neither is particularly helpful for decisionmaking.

A more useful focus? Behaviour over time

If the real concern behind “seeing weeds” is risk and defensibility, then a more useful focus is not visibility, but *behaviour*.

Over time, functioning rehabilitation areas tend to show:

- seasonal responses that align with climate and rainfall;
- recovery after stress events such as drought;
- decreasing volatility as systems mature; and
- convergence toward the behaviour of analogue sites.

Areas at risk often show the opposite: exaggerated responses, persistent divergence from reference conditions, or increasing instability through time.

These patterns are often detectable well before problems are obvious on the ground. And long before a once-a-year survey or flight would capture them.

Photo by Judy Beth Morris on Unsplash

An important caveat…

None of this is an argument against weed mapping in general.

In precision agriculture, biosecurity, or targeted eradication programs, identifying specific weed species is often exactly the right focus. In those contexts, spatial precision is critical. And where weeds are genuinely affecting mine site rehabilitation, knowing where they are and what they are is important to target interventions.

But mine rehabilitation operates at different spatial and temporal scales, with different regulatory objectives. Here, confidence comes less from knowing exactly what is present in one place at one time, and more from understanding how systems are behaving through time.

Confidence comes less from knowing exactly what is present in one place at one time, and more from understanding how systems are behaving through time.

Can we ask better questions?

“Seeing weeds” persists as a focal point because it feels concrete. But it is ultimately a proxy for deeper concerns about risk, accountability, and outcomes.

When monitoring programs are designed around better questions, the resulting evidence is often more useful, more defensible, and more aligned with how rehabilitation success is actually assessed. How about:

Is this area stabilising? Is it responding to seasons as expected?
Is it trending towards becoming more like its analogue, year by year?

02

What does good ecosystem rehabilitation look like over time?

As established in the previous section, focusing on *“seeing weeds”* can miss key outcomes of rehabilitation monitoring. Not because weeds don’t matter, but because hyperfocus on a single component of the ecosystem at a single moment tells us surprisingly little about whether a system is proceeding towards its desired state and if our rehabilitation efforts are succeeding.

That raises a more important question:

If not the tangibility of weeds, then what does good rehabilitation actually look like?

And more specifically, *what does it look like over time? And how do we better understand and respond to the ecosystem trends?*

Rehabilitation is a process, not a state

One of the quiet traps in rehabilitation monitoring (and not just for ecosystems either...) is the idea that there is a point at which a site *is* rehabilitated. It becomes a compliance milestone rather than an ongoing process.

In reality, rehabilitation is a process of transition. Systems move from disturbance toward stability, passing through phases of establishment, consolidation, and maturation. Along the way, they fluctuate. They respond to seasons. They respond to stress. They sometimes move backwards before they move forwards again.

Good rehabilitation is therefore not defined by how a site looks on a particular day, but by how it behaves across many days, seasons, and years.

And yes, weeds may be one component of that. But it doesn't mean to say that counting and identifying species is necessarily the answer either.

In reality, rehabilitation is a process of transition. Systems move from disturbance toward stability, passing through phases of establishment, consolidation, and maturation.

The importance of seasonal behaviour

One of the strongest indicators of system function is how vegetation responds to normal seasonal cycles.

In functioning systems, growth and senescence tend to follow predictable patterns driven by rainfall, temperature, and day length. There is variation from year to year (climate change will affect this), but it sits within an expected envelope.

Photo by Chris Lawton on Unsplash

Early-stage rehabilitation areas often show exaggerated responses – rapid green-up after rainfall, followed by equally rapid decline. Over time, as soils develop and plant communities establish, these responses usually become more moderated and more consistent.

This moderation is a sign of increasing resilience.

A snapshot image can capture 'greenness' or even species distribution, but it cannot tell you whether that greenness is fleeting or sustained.

Further, an ill-timed snapshot may completely miss the presence of certain species of interest.

Recovery after stress matters more than peak condition

Another defining feature of good rehabilitation is how systems respond to stress.

Droughts, heatwaves, and extreme rainfall events are not exceptions. They are part of the operating environment. A system that only looks good in favourable conditions is not a robust one.

Healthy or well-functioning systems tend to:

- decline during stress events;
- recover afterward at a rate consistent with their surrounding environment (or compared to analogue sites); and
- return toward their previous condition rather than resetting entirely.

By contrast, systems at risk often show delayed recovery, incomplete recovery, or increasing sensitivity to each new stress event.

These patterns are rarely obvious in single surveys. They emerge only when observations are stitched together through time, and with suitable observation frequency.

Stability does not mean sameness

A common misconception is that stability implies uniformity. In ecological systems, stability usually means the opposite: variability within bounds.

As rehabilitation progresses, good systems tend to show:


reduced extreme swings


more consistent responses year to year


fewer abrupt internal contrasts

Importantly, they do not stop changing.
They simply change and respond in more predictable ways.

Monitoring approaches that focus on isolated observations can mistake normal variability for failure, or miss early warning signs of instability.

Context and comparison matters

It is difficult to judge rehabilitation in isolation.

Climate variability alone can make one year look wildly different from the next. Without context, it becomes hard to tell whether a change reflects progress, setback, or simply local weather.

Without context, it becomes hard to tell whether a change reflects progress, setback, or simply local weather.

This is why comparison is so important.

Good rehabilitation does not need to look identical to surrounding undisturbed land. But over time, it should begin to *behave* in similar ways, responding to seasons and stress in broadly comparable patterns.

When rehabilitated areas consistently diverge from nearby reference systems under the same conditions, that divergence is meaningful.

From snapshots to trends and trajectories

Taken together, these characteristics point to a simple shift in perspective.

Good rehabilitation is not defined by peak condition, species presence, or momentary appearance. It is defined by **behaviour** and **direction**.

Is the system becoming more stable? Is its response to seasons becoming more consistent? Is its recovery after stress improving?

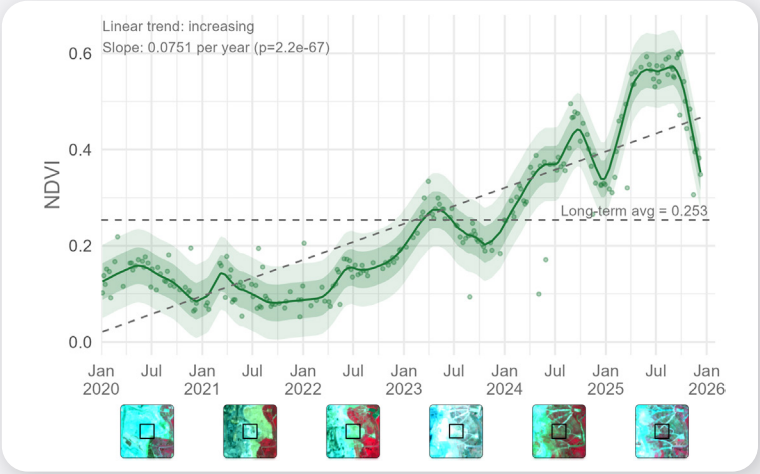

These are trend and trajectory questions, not visibility questions.

The figure above illustrates this point based on a rehabilitation site in Australia, and using NDVI values derived from Sentinel-2 data. What I want you to notice though is that it's not about absolute NDVI values, or about identifying species. It's about direction, stability, and response to seasons. And how observing in this way allows us to join the dots between single snapshots.

Asking the right follow-up questions

Once monitoring focuses on behaviour through time, different and often more useful questions emerge:

- Where are systems improving steadily, even if slowly?
- Where are they plateauing?
- Where are they becoming more volatile or diverging from reference behaviour?

These questions support earlier intervention, more targeted fieldwork, and more defensible decision-making.

They also align closely with how rehabilitation success is actually assessed – through evidence of progress, resilience, and sustainability rather than perfection.

The following section examines one of the most powerful tools for answering these questions: comparing rehabilitated areas to their analogue sites.

Banner image by kazuend on Unsplash

03

Why analogue sites matter in rehabilitation monitoring

As established in the previous section, good rehabilitation is not defined by how a site looks at its peak, or at a certain snapshot, but by how it behaves over time. That naturally leads to the next question:

Behaves compared to what?

This is where analogue (or reference) sites come in, and where rehabilitation monitoring often becomes more powerful and defensible.

In most rehabilitation plans, analogue or reference sites are already defined. They are the nearby areas that represent the *target condition* the rehabilitated land is expected or desired to move towards. They may also have a set of criteria related to the percentage weed cover or thresholds for slope to reduce erosion risk for example.

This does not mean that they are ideals. They are not pristine. And they are not static.

Analogue sites are simply functioning systems operating under the same climate, soils, and disturbance regime as the rehabilitation area. And they provide a crucial source of context.

Why comparison matters

Rehabilitation doesn't occur in a vacuum. Rainfall varies from year to year. Droughts and heatwaves occur. Seasons with optimal or poor conditions may alternate. So looking at a rehabilitated area in isolation makes it very difficult to know whether a change reflects progress, failure, or simply weather. Climate change adds a further layer of complexity: as baseline conditions shift over decadal timescales, the behaviour of analogue sites will also change, which is an important consideration when interpreting long-term rehabilitation trajectories.

By comparing rehabilitation areas to analogue sites exposed to the same conditions, much of this ambiguity disappears. This is because both areas experience the same weather, and should also share similar exposure and topographical characteristics. Therefore, if they behave differently over time, that difference is meaningful.

But asking whether a rehabilitated area is *behaving like* its analogue does not mean asking whether it looks the same.

It means asking questions such as:

- Do both areas green up and senesce at roughly the same times?
- Do they respond similarly to rainfall events?
- Do they recover after drought at comparable rates?
- Is variability increasing, decreasing, or remaining stable over time?

These are behavioural questions. They are about patterns and responses, not about species lists or pixel-level detail.

A case study

To make this more concrete, consider a rehabilitation area at an Australian site, compared with a nearby area selected as an analogue based on similar landform and soils.

Over the past five years, both sites have been exposed to the same climate variability and seasonal cycles, which is clearly reflected in their vegetation index time series below. The analogue site shows a relatively stable pattern, with consistent seasonal oscillations around a higher long-term average (0.568). The rehabilitation site, by contrast, sits at a much lower baseline (long term average 0.230) and does not have the same seasonality. The rehabilitation site is largely exposed or bare ground over this time period, with some grasses and weeds showing green-up at times.

Importantly, there is only a weak positive trend in the rehabilitation signal over this period. This suggests that the system is still in a very early stage, where meaningful change is limited. What matters at this point is not convergence in magnitude, but the fact that both sites are responding to the same environmental drivers, and that the rehabilitation trajectory can now be interpreted in that shared context.

At this stage, the value of the comparison lies less in demonstrating rehabilitation success, and more in establishing a baseline for how the rehabilitation area behaves relative to its analogue, against which future change can be meaningfully assessed.

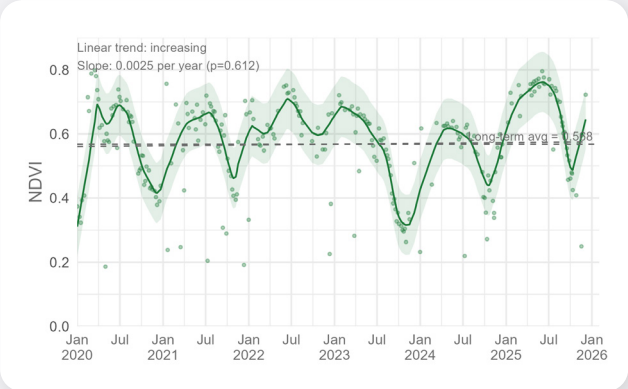

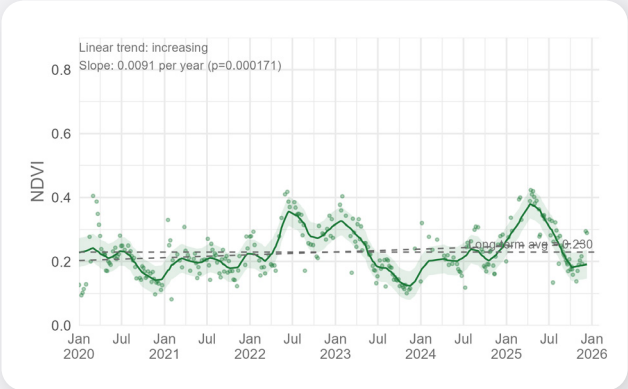

What divergence tells us

Just as importantly, divergence from analogue behaviour is often an early warning. We need to look out for when a rehabilitated area:

- remains greener for longer than the analogue,
- responds too strongly to small rainfall events, or
- becomes increasingly volatile through time.

These conditions suggest that different processes are dominating. This can reflect weed-driven dynamics, shallow rooting, poor soil development, or other underlying issues. Crucially, these signals often appear well before problems are obvious on the ground. And they still don't require species level identification to raise an alert.

Instead, we want to move towards demonstrating that a system is:

- trending toward a reference condition,
- becoming more stable through time, and
- responding appropriately to environmental variability.

This gives us confidence to demonstrate that rehabilitation is on a credible path to closure. Comparing with an analogue site provides a clear, climate-normalised way to do this.

From comparison to confidence

When monitoring focuses on analogue behaviour, confidence comes from context rather than detail.

Instead of asking whether a particular feature (e.g. weed species) can be seen in an image, the question becomes whether the system as a whole is moving in the right direction under real-world conditions. I have previously argued that focusing on counting weeds and species detection can detract from the bigger picture.

When monitoring focuses on analogue behaviour, confidence comes from context rather than detail.

This shift does not eliminate the need for field surveys or targeted high-resolution data. It does, however, make those efforts more strategic. We can focus attention where divergence is emerging rather than spreading effort evenly across the landscape or only conducting repeat quadrats / transects.

Ultimately, *"Is this area behaving like its analogue?"* asks whether rehabilitation is becoming more resilient, more predictable, and more self-sustaining over time. And this is exactly what rehabilitation monitoring is meant to establish.

The following section examines why temporal resolution is an undervalued dimension in rehabilitation monitoring design.

Banner image by Aron Visuals on Unsplash

04

Why time is the most undervalued dimension in rehabilitation monitoring

As established in the previous section, rehabilitation success is better understood through behaviour over time than through what can be seen in a single snapshot, and that comparison to analogue sites provides essential context for interpreting change.

There is one dimension that quietly underpins all of this, yet is still routinely undervalued in monitoring design. Time. Not just how long rehabilitation has been underway, but *how often* and *when* we observe it.

There's no argument our ability to track progress is key to 'ticking the box' on rehabilitation success. On the way to meeting this, I've had many conversations with people about the optimal spatial dimensions (ground sample distance and spatial extent) for a particular metric. And people will debate the value of different drone altitudes for image capture. But I rarely see the same consideration given to time.

This section examines what happens when the temporal dimension is not optimised, and what monitoring programs can do differently.

The same site, many possible stories

To explore this, I looked at vegetation change at an Australian mine rehabilitation site using NDVI as a simple, familiar proxy for vegetation condition. NDVI isn't the point here though. The same logic would apply to any index, metric, or measurement. The actual location is also not important for this analysis.

Using all valid satellite observations over the site from 2020 to 2025 (>250), a clear picture emerges. The site shows:

- a strong seasonal cycle
- increasing peak and baseline values over time
- recovery after periods of decline
- an overall upward trajectory

Taken together, these signals are consistent with a system that is progressing, but still dynamic. This is the most complete story we can tell from the data.

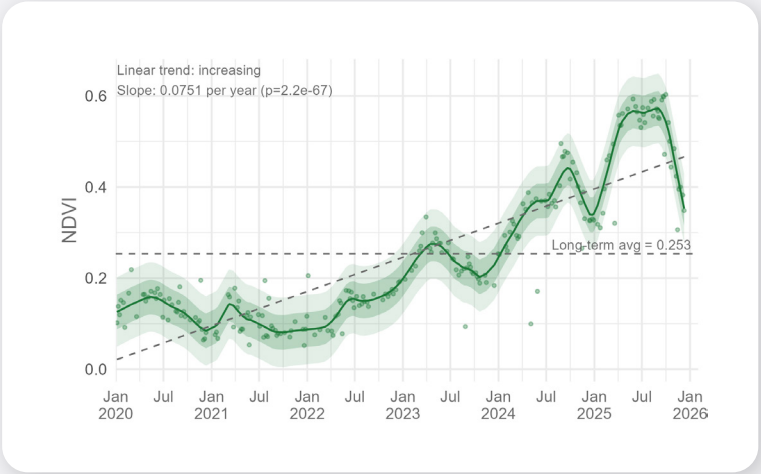

What happens when we sample less often?

Now imagine a more typical monitoring scenario: one observation per year.

Using the same underlying dataset, I extracted a single value per year, repeating this exercise for each calendar month. In other words, I asked: *What trend would we infer if we always sampled in January? In February? In August?*

Each of these sample months produces a different linear trend. Some suggest rapid improvement. Others imply much slower progress. Some have larger variation than others.

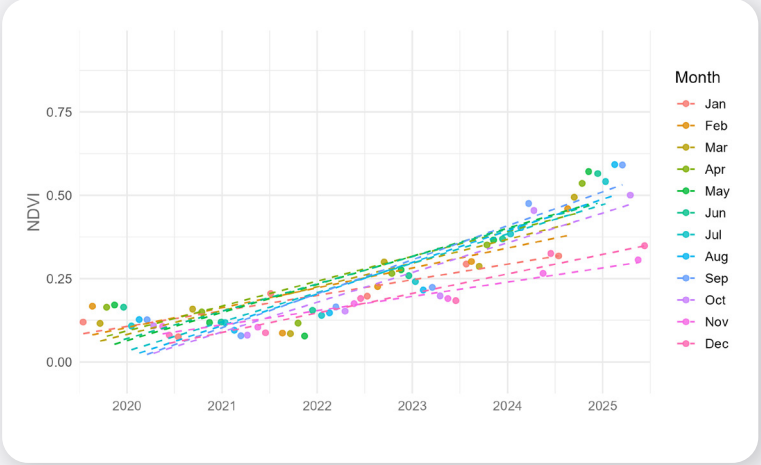

All are technically correct. And all are incomplete.

The key difference between them is *when* the sample was taken.

This isn't a failure of statistics or remote sensing. It's a reminder that when systems are seasonal, sparse sampling collapses a rich time series into a single, highly conditional narrative.

Why this matters for decision-making

In rehabilitation, trends are often used to support decisions and milestones. To illustrate this, I took the linear trend from each monthly sampling scenario and asked a deliberately simple question:

Based on this trend, how long would it take for the site to reach an NDVI value of 0.6?

This threshold is arbitrary, but it serves as a stand-in for the kind of numeric trigger that often appears in reporting frameworks.

The results vary dramatically. Depending on which month is used for annual sampling, the estimated time to reach this threshold differs by years.

In some cases, the site appears close to achieving the target. In others, it looks much further away.

Yet nothing about the site itself has changed, only the sampling strategy.

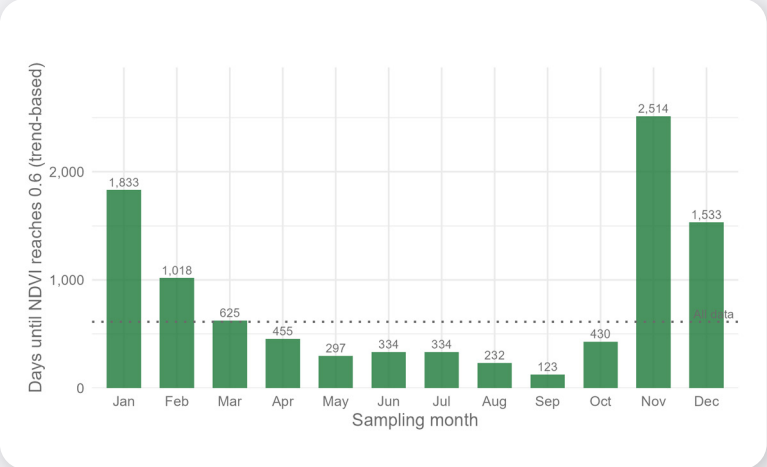

A necessary caveat and disclaimer

These projections are based on simple linear models. Real systems are not linear, and thresholds can be crossed earlier or later depending on future conditions.

That caveat is important. But it does not weaken the point. If anything, it strengthens it. Because if even a simple model produces wildly different conclusions depending on sampling month, then confidence based on sparse observations alone should be treated with caution.

It's also important to be clear about what this analysis is *not* saying. It is not an argument against annual sampling. Annual sampling is often necessary. Field surveys, drone campaigns, and detailed assessments are expensive, logistically complex, and sometimes risky. They provide information that high-frequency data cannot. Things like species composition, structure, and on-the-ground context.

The issue is not that annual data are wrong. It's that annual data are conditional. When we rely on a small number of observations, the timing of those observations matters enormously. Without additional context, it's easy to mistake seasonal effects for trends, or to over- or under-estimate progress.

The issue is not that annual data are wrong. It's that annual data are conditional.

High-frequency data provides context, not replacement

High-temporal-resolution data whether from satellites, in-situ sensors, or other sources are not a replacement for detailed fieldwork. But the value in frequent observations lies in providing context, joining the dots, and letting us know *how* we get from point A to point B. Frequent observations allow us to:

- see seasonal patterns rather than single values;
- distinguish short-term variability from long-term change;
- understand how systems respond to stress and recovery; and
- interpret sparse measurements within a broader trajectory.

In this role, high-frequency data don't compete with annual sampling, but they help to explain observations, or even to target when sampling should occur.

Designing monitoring with time in mind

We argue endlessly about what and where, and almost never about when and how often. We even talk about resolution as if there is only one type – and it's definitely spatial! Yet if rehabilitation monitoring is about understanding trajectories rather than snapshots, then temporal design deserves the same attention we give to spatial resolution.

Questions worth asking include:

- How often do we need to observe this system to separate trend from seasonality?
- What context is required to interpret our field measurements?
  This may include specific flowering or green-up stages of vegetation.
- Where does sparse sampling carry the greatest risk of misinterpretation?

These questions are more important to answer than ‘at what time is the ecologist the cheapest?’

The bigger picture

When monitoring programs underperform, it is rarely because the data are wrong. More often, it’s because the questions being asked are not well aligned with the dynamics of the system, or because the way data are collected does not match the questions we actually need answered.

Stepping back can be valuable. Are we asking questions that reflect how this system changes through time? Are our sampling strategies designed to capture those dynamics, or are they driven by convenience and habit? Which metrics genuinely matter for decision-making, and what is the most appropriate way to measure them?

In short, the temporal dimension of our data is critical for turning observations into understanding. But it also needs to be handled thoughtfully. Despite the case for higher-frequency observations, *more* is not always better. Data toxicity is real, and without clear questions, well-designed analytical workflows, and careful interpretation, dense time series can leave us in a tangled mess rather than sharpening insight.

In the next section, I’ll step back from individual data types and look at how satellites, drones, and field surveys can be combined into a monitoring hierarchy that balances cost, spatial resolution, and temporal insight.

05

Designing a monitoring hierarchy that actually works

Across the earlier sections in this eBook, I've argued for a shift in how we think about rehabilitation monitoring. To move away from snapshots and visual checklists, and towards behaviour, context, and trajectories through time.

The previous sections have established that good rehabilitation monitoring depends on understanding behaviour through time, using analogue sites for context, and treating temporal resolution with the same rigour as spatial resolution.

Once you accept the framing presented in this series, a practical question follows:

How do we design a monitoring program that captures what matters, without becoming prohibitively expensive or analytically overwhelming?

The answer is not to choose between field surveys, drones, and satellites. It's to use them together in a monitoring hierarchy.

Not all data need to do the same job

A common source of frustration in rehabilitation monitoring outcomes is derived from expecting a single data source to answer every question. I am frequently asked to compare field survey, drones, and satellites – and for my opinion on which is 'best'. And of course the answer can only be 'it depends!'

Most times, those asking for the comparison have already developed their own perspectives on the answer. They consider field surveys to be the 'source of truth' and a regulatory requirement. They want high spatial resolution drone imagery to provide wall-to-wall assessments akin to the field survey. And typically satellites are dismissed because they can't "see" individual plants.

However, these expectations and assumptions create false trade-offs. The reality is that different tools are good at different things, and a robust monitoring program recognises this rather than fights it.

The role of high-frequency, lower spatial resolution data

High-temporal resolution data such as satellite time series are most valuable for understanding context. This includes both spatial and temporal context. That means that these data are helpful for understanding what is happening in nearby locations (spatial context), as well as the times in-between sampling (temporal context).

Frequently acquired satellite data are well suited to:

- characterising seasonal behaviour,
- identifying longer-term trajectories,
- distinguishing short-term variability from persistent change,
- highlighting periods or areas of divergence from expected behaviour, and
- providing archival data when and where you may otherwise have data gaps.

Data at these scales may not tell us *why* something is happening. But it is very good at telling us *where* and *when* to look more closely.

The role of targeted, high-spatial resolution data

Drone imagery and other high-spatial resolution aerial photography excel at spatial detail. Drone technology has advanced rapidly and offers genuinely powerful capabilities for spatial data capture.

They are well suited to:

Mapping topographic structure and land / vegetation cover

Identifying specific features or management issues

Supporting measurements that require spatial precision

Used strategically, high-spatial resolution data can be deployed where broader monitoring indicates something interesting is happening. It is also by far the best option for mapping landforms. It is less well suited to being spread thinly across entire sites on fixed schedules and can be quite expensive to acquire and process for that. Targeted captures makes these data more informative and more cost-effective.

The role of field surveys

Field-based surveys remain essential. While drone mapping could also be considered field survey, I am not referring to it in this context. By field survey, I mean observing and taking measurements on the ground (or in-water) by a suitably qualified person.

Ecological field surveys provide information that no remote sensing product can fully replace: species composition, condition, soil properties, and ecological context. They also provide the grounding needed to interpret remotely sensed signals with confidence (AKA calibration and validation or cal/val).

In a hierarchical design, field surveys play a confirming and explanatory role. They help answer *why* patterns observed in broader datasets look the way they do, and whether those patterns reflect desirable or undesirable processes.

Putting the pieces together

In practice, an effective monitoring hierarchy might look like this:

1. **Broad, frequent observations** establish context and trajectories, and may inform appropriate sampling timing
2. **Targeted, higher-resolution data** investigate specific signals or areas of concern
3. **Field surveys** validate interpretation and inform management responses

Each layer informs the next, but also doesn't need to occur in this order as feedback loops are common. Importantly, no single layer is expected to carry the entire burden of assessment, and therefore, each must help the others.

Hierarchical monitoring shifts effort from routine data collection towards informed decision-making. This means that instead of asking whether a site looks acceptable at a particular moment, the focus moves to whether it is behaving as expected under real-world conditions. The behaviour then allows us to consider when and where intervention or further investigation is actually warranted.

Hierarchical monitoring shifts effort from routine data collection towards informed decision-making.

This approach also aligns well with practical constraints. High-frequency satellite data provide continuity without excessive cost, while more expensive or labour-intensive methods are used where they add the most value.

Good monitoring design is less about the tools themselves than about the questions they are meant to answer. When time, context, and scale are considered explicitly, satellites, drones, and field surveys stop competing for attention and instead become complementary parts of a coherent system.

Rehabilitation monitoring works best when it moves away from checklists and toward understanding behaviour through time.

Designing monitoring as a hierarchy allows each data source to do what it does best. High-frequency data provide continuity and context, targeted high-spatial resolution data add spatial detail where it matters most, and field surveys ground interpretation in ecological reality. None of these layers is sufficient on its own, but together they allow patterns to be interpreted with far greater confidence.

Rehabilitation monitoring works best when it moves away from checklists and toward understanding behaviour through time. A hierarchical approach doesn't just reduce cost or complexity, but instead improves decision-making by ensuring that effort is directed where it adds the most value, and that change is interpreted in context rather than in isolation.

Banner photo by Yvette Goldberg on Unsplash

06

Rehabilitation monitoring beyond annual reports

Across this eBook, I've argued for a shift in how we think about rehabilitation monitoring. I've suggested moving away from snapshots and checklists, and toward behaviour, context, and trajectories through time. I've also argued that good monitoring depends less on individual tools and more on how information is combined to answer the right questions.

So if we change the way that we think about rehabilitation monitoring, how does that fit with reporting cycles? How should we report rehabilitation outcomes when the underlying evidence is inherently dynamic? Keeping in mind of course that we want better foundations for judgement, not more work.

Why annual reports became the norm

Annual reports (like annual field surveys) exist for good reasons. Regulatory cycles, budgeting processes, and audit requirements all operate on yearly rhythms. Annual reports provide a clear, accountable document that we can review, archive, and compare over time.

There is nothing inherently wrong with this structure. The challenge arises when we treat annual reports as *the evidence itself*, rather than as a summary of evidence collected through time. Annual reports have a convenient snapshot mindset.

The limitation of snapshot-based reporting

Natural ecosystems and of course rehabilitation sites are dynamic. The vegetation and landform respond to rainfall, temperature, disturbance, and management actions on timescales that rarely align neatly with reporting deadlines.

So when we base assessment on a single annual snapshot, several risks emerge:

- seasonal conditions can dominate interpretation,
- early warning signals can be missed,
- short-term variability can be mistaken for long-term trend, and
- confidence can be overstated or understated depending on timing.

These limitations do not necessarily mean that reporting has failed. Instead, they reflect a mismatch between system dynamics and how we accumulate evidence.

Evidence built over time

Alternatively, we can separate *evidence accumulation* from *reporting frequency*.

Building evidence over time means we collect and curate observations more frequently than we formally interpret or report them. This allows patterns, trajectories, and anomalies to be preserved in the record, even if interpretation occurs only periodically or in response to specific questions or triggers. Instead of reconstructing behaviour from a few snapshots, we can draw on a continuous history when it is needed (this is the satellite super power).

Importantly, this does not require constant interaction with data or complex visual interfaces. In many cases, the outcome remains a simple assessment (e.g. a status, threshold, or traffic-light indicator) but one that is supported by a much richer, contextualised, and more defensible evidence base.

Building evidence over time means we collect and curate observations more frequently than we formally interpret or report them.

When we accumulate evidence through time, changes in system behaviour become visible earlier, even if interpretation occurs only periodically. This allows emerging issues to be detected before they escalate into compliance problems, and supports management responses that are more targeted and proportionate rather than reactive.

At the same time, positive trajectories can be assessed with greater confidence. Instead of relying on a single observation to demonstrate progress, we gradually build confidence by observing persistent behaviour under a range of conditions.

From a risk perspective, this shifts rehabilitation monitoring away from a periodic test at fixed reporting milestones. Instead we move towards a more robust process of ongoing validation, where a growing body of evidence supports our conclusions - and interventions.

Aligning monitoring with value

When we design monitoring around evidence built over time, value accrues continuously rather than episodically. Each new observation adds context to what came before it, even if formal reporting occurs less frequently.

This approach prioritises defensible outcomes over continuous interaction, ensuring that insight is available when needed without requiring ongoing engagement with analytical tools. In practice, this means confidence builds quietly in the background, rather than demanding constant attention.

This model better reflects how rehabilitation actually unfolds. It is gradual, uneven, and responsive to external conditions. Monitoring and reporting can therefore stay in sync with system behaviour, supporting clearer and more defensible decisions throughout the life of a site.

Rehabilitation does not progress in annual steps, and evidence of its success rarely arrives neatly aligned with reporting cycles. Monitoring systems that recognise this reality allow reporting to remain structured and familiar, while ensuring that conclusions are grounded in a genuine understanding of system behaviour through time.

In the end, better reporting does not come from more documents, but from better evidence beneath them.

In the end, better reporting does not come from more documents, but from better evidence beneath them.

Build defensible rehabilitation monitoring – beyond snapshots

Rehabilitation monitoring is often treated as an annual reporting obligation - something to be completed, documented, and filed away each year. But ecosystems do not recover in annual steps.

They respond to rainfall, disturbance, temperature, and management actions in ways that are gradual, uneven, and highly seasonal. When monitoring programs rely too heavily on snapshots or single reporting milestones, important context can be lost. We also know that timing can influence interpretation, meaning that our decisions may be shaped by anniversary appearance rather than trajectory.

GeoNadir will analyse your rehabilitation area and show you:

- How your site is tracking over time
- Whether it's behaving like its analogue site
- Where there are early signs of divergence or risk
- What you can confidently stand behind in reporting
- How to interpret a rehabilitation report for your site

We'll run an initial analysis for free, using your site data and provide a clear, defensible view of condition and trajectory.

REQUEST YOUR REHABILITATION REPORT

Health check

This short health check is designed to help you reflect on how well your current rehabilitation monitoring approach aligns with the realities of dynamic systems.

It is not a compliance audit, and it is not about whether your program is “right” or “wrong”. Instead, it is intended to surface strengths, gaps, and assumptions that may affect how defensible and useful your monitoring really is.

You can complete this individually or as a team. Most people take 10–15 minutes.

How to use this health check

For each statement below, indicate how strongly you agree using the following scale:

1	2	3	4	5
Strongly disagree	Disagree	Neutral/Unsure	Agree	Strongly agree

Be honest. The value comes from reflection, not scoring highly.

A Purpose and decision-making

STATEMENT	1	2	3	4	5
We are clear about the decisions our rehabilitation monitoring is meant to support.	○	○	○	○	○
Our monitoring questions focus on behaviour and trajectories, not just end states.	○	○	○	○	○
If an indicator changed, we know what management response it would trigger.	○	○	○	○	○
We can explain why each metric is being measured, not just how.	○	○	○	○	○

B Time and sampling design

STATEMENT	1	2	3	4	5
Our monitoring design explicitly considers seasonality.	○	○	○	○	○
We understand how the timing of data collection affects interpretation.	○	○	○	○	○
We can distinguish short-term variability from longer-term trends.	○	○	○	○	○
Our confidence in outcomes does not rely on sampling occurring in a “good” month.	○	○	○	○	○

C Context and analogue sites

STATEMENT	1	2	3	4	5
We use reference or analogue areas to help interpret change.	○	○	○	○	○
We compare trajectories and patterns, not just absolute values.	○	○	○	○	○
Our interpretation accounts for climate variability shared between sites.	○	○	○	○	○
We can explain why a site may look different but still be on a valid trajectory.	○	○	○	○	○

D Data hierarchy and roles

STATEMENT	1	2	3	4	5
We use different data sources for different purposes.	○	○	○	○	○
We don't expect any single dataset to answer every monitoring question.	○	○	○	○	○
We use high-frequency data to provide context, not just snapshots.	○	○	○	○	○
We deploy high spatial resolution data strategically rather than uniformly.	○	○	○	○	○

E Reporting and defensibility

STATEMENT	1	2	3	4	5
Our reports summarise evidence built over time rather than single observations.	○	○	○	○	○
We can trace how a reported status was reached if questioned later.	○	○	○	○	○
We are confident defending conclusions outside formal reporting periods.	○	○	○	○	○
Our reporting reflects uncertainty as well as progress.	○	○	○	○	○

Interpreting your responses

Rather than producing a numerical score, this health check is most useful for identifying patterns in where confidence is high and where uncertainty remains. There is no right or wrong answer here.

A simple way to start is to look at which sections contained the most disagreement or uncertainty as below:

Mostly A & B issues

Often points to a monitoring program that is strongly influenced by snapshots and sampling timing.

Mostly C issues

May indicate that interpretation lacks sufficient context spatially or temporally.

Mostly D issues

May point to inefficiencies or unrealistic expectations of individual tools or data sources.

Mostly E issues

Often highlights risk around defensibility and reporting rather than data availability.

Programs rarely sit neatly in one category. Mixed responses are normal.

What next?

This health check is intended as a starting point for conversation. Areas of uncertainty or disagreement are often the most productive places to focus discussion. Improving rehabilitation monitoring does not usually require *more* data, though *different* data is common. It more often requires clearer questions, better alignment between tools and purpose, and stronger links between evidence and reporting.

Use this exercise to identify where small changes in design could lead to much greater confidence in outcomes. Contact karen@geonadir.com for more information.

About the author

Dr. Karen Joyce is Co-founder and Head of Product and Science at GeoNadir. Her work sits at the intersection of remote sensing, ecology, and data platforms, with a focus on making spatial data more actionable for environmental monitoring and reporting. She has worked extensively across Australia and overseas to build remote sensing workflows and evidence-based reporting. GeoNadir's platform was built to support the kind of continuous, contextualised monitoring described in this ebook.