

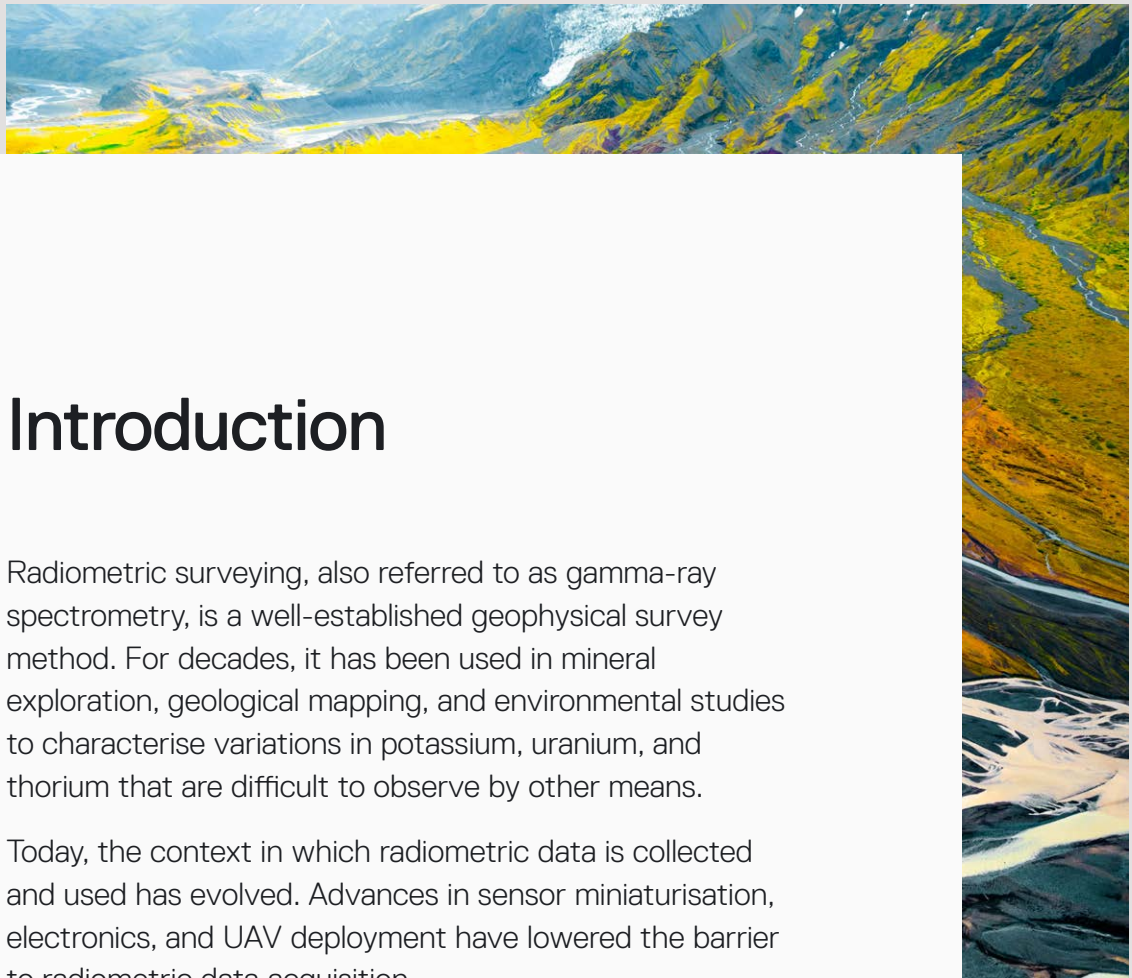
WHITE PAPER

RADIOMETRICS PROCESSING EXPLAINED

A practical guide to gamma-ray spectrometry
data preparation and interpretation



Introduction	XX
<hr/>	
The challenges and opportunities of modern radiometric surveys	XX
<hr/>	
Spectral smoothing with NASVD	XX
<hr/>	
Making eigenvector selection decisions in practice	XX
<hr/>	
A structured radiometrics data processing workflow in Oasis montaj	XX
<hr/>	
The future of radiometric processing	XX
<hr/>	



Introduction

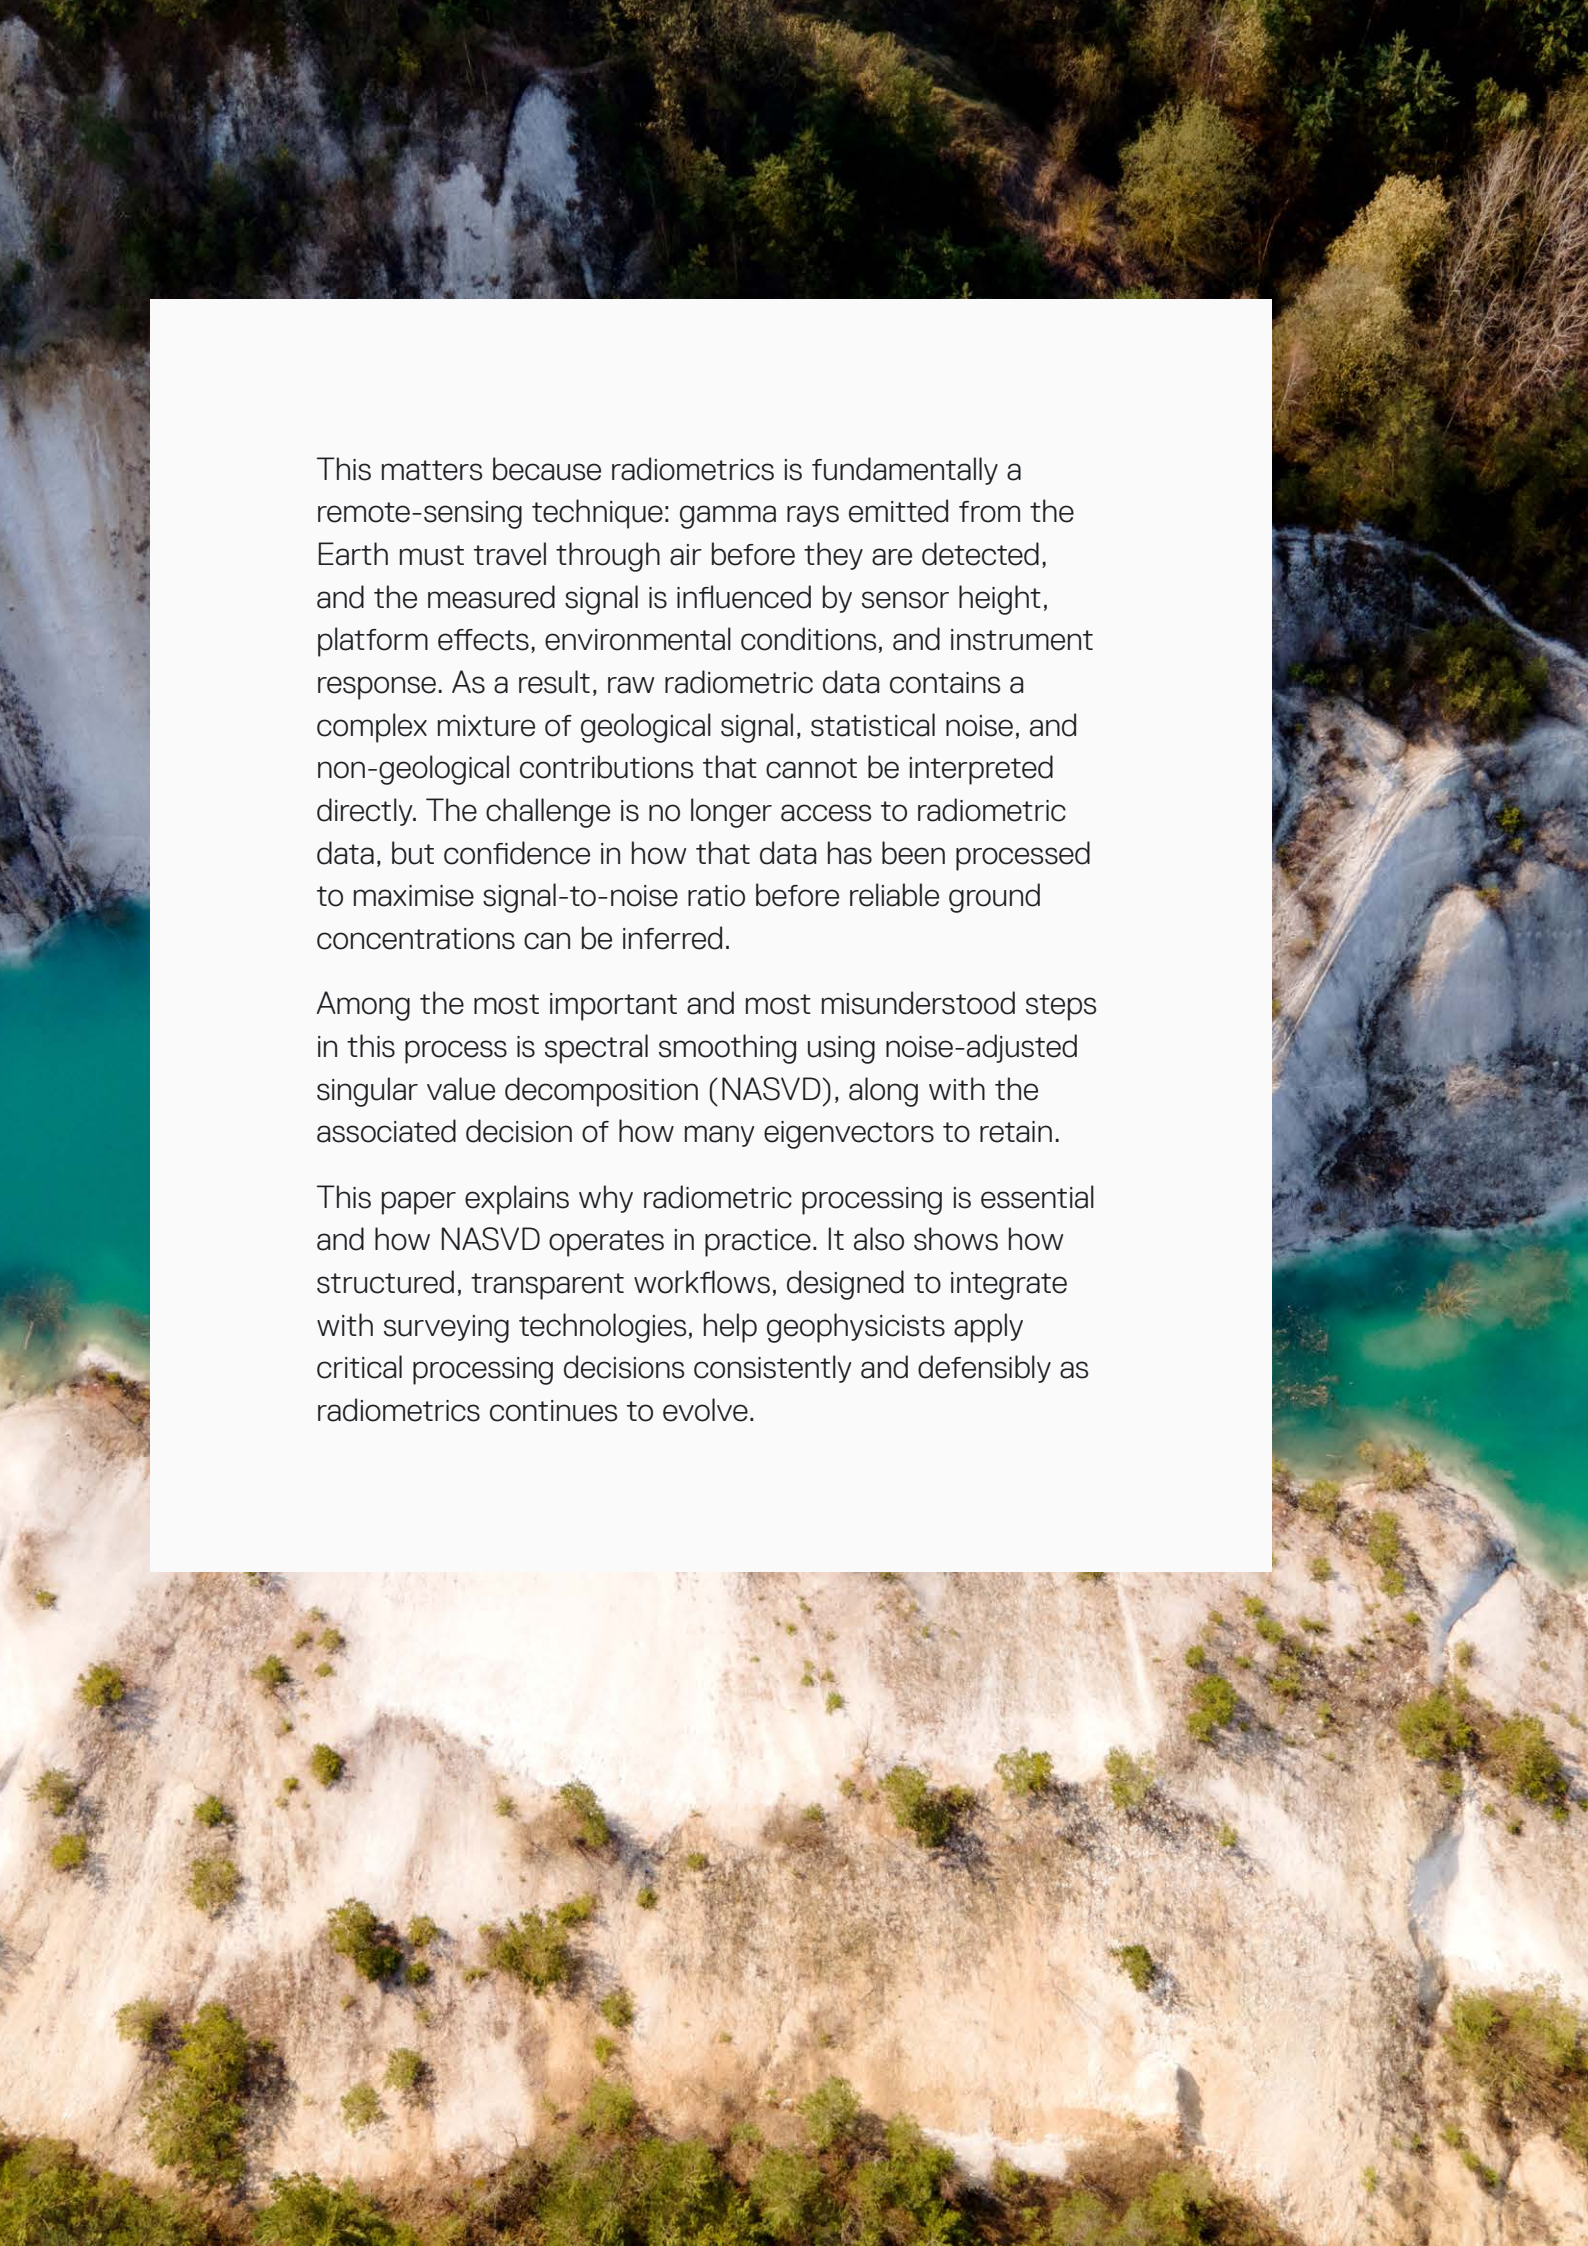
Radiometric surveying, also referred to as gamma-ray spectrometry, is a well-established geophysical survey method. For decades, it has been used in mineral exploration, geological mapping, and environmental studies to characterise variations in potassium, uranium, and thorium that are difficult to observe by other means.

Today, the context in which radiometric data is collected and used has evolved. Advances in sensor miniaturisation, electronics, and UAV deployment have lowered the barrier to radiometric data acquisition.

As a result, radiometrics is being applied across a growing range of use cases, from site-scale investigations and environmental studies to tailings assessment and soil characterisation, often by teams working under tighter timelines and with smaller detectors.

This shift has important consequences.

Smaller sensors and lower-altitude surveys tend to produce noisier spectra, even as spatial resolution and data volumes increase. At the same time, radiometric datasets are being generated and processed by a broader range of users, many of whom rely on default processing settings without fully understanding their implications.



This matters because radiometrics is fundamentally a remote-sensing technique: gamma rays emitted from the Earth must travel through air before they are detected, and the measured signal is influenced by sensor height, platform effects, environmental conditions, and instrument response. As a result, raw radiometric data contains a complex mixture of geological signal, statistical noise, and non-geological contributions that cannot be interpreted directly. The challenge is no longer access to radiometric data, but confidence in how that data has been processed to maximise signal-to-noise ratio before reliable ground concentrations can be inferred.

Among the most important and most misunderstood steps in this process is spectral smoothing using noise-adjusted singular value decomposition (NASVD), along with the associated decision of how many eigenvectors to retain.

This paper explains why radiometric processing is essential and how NASVD operates in practice. It also shows how structured, transparent workflows, designed to integrate with surveying technologies, help geophysicists apply critical processing decisions consistently and defensibly as radiometrics continues to evolve.



CHAPTER 1

The challenges and opportunities of modern radiometric surveys

Advances in detector miniaturisation by innovators such as [Medusa Radiometrics](#), combined with the growing use of unmanned aerial vehicles (UAVs), are enabling radiometric surveys across a wider range of applications. These developments make it practical to survey smaller, constrained, or hard-to-reach areas that would previously have been inaccessible or cost-prohibitive using traditional methods, often with higher spatial resolution but more challenging counting statistics.

At the same time, radiometric data is being used by a broader range of practitioners and applied to more diverse problems. Environmental investigations, tailings assessments, soil characterisation, and site-scale studies often require fast turnaround and consistent outputs, leaving limited opportunity for iterative processing or specialist review.

In these settings, established workflows and default processing parameters are commonly relied upon to manage complexity and maintain efficiency.

The challenge is that radiometric processing is not a purely mechanical step. Decisions made during noise reduction and correction directly affect the reliability of the final data, and when applied without a clear understanding of their implications, subtle distortions can be introduced that may not be obvious during routine quality checks.

As practitioners have long recognised, confidence in radiometric interpretation depends as much on how noise is managed as on how data is acquired. With modern surveys operating closer to the limits imposed by detector size and survey geometry, the margin for error during

processing has narrowed. Choices that were once conservative or implicit now have a greater impact on final results.

These changes have practical implications for spectral smoothing. In legacy airborne surveys with large detectors and high-count rates, eigenvector selection for NASVD could often be treated as relatively stable across datasets. In modern surveys, particularly those using smaller detectors, shorter integration times, or low-altitude platforms, the separation between signal and noise is less distinct.

As a result, eigenvector selection becomes more sensitive to detector size, count statistics, and survey conditions, and less transferable between projects. NASVD, therefore, becomes not just a standard processing step, but a point at which informed professional judgement is increasingly required.

About Medusa Radiometrics

Medusa Radiometrics develops compact gamma-ray spectrometers designed for flexible deployment on foot, by vehicle, or from low-altitude airborne platforms. These systems prioritise portability and high spatial resolution, often operating with smaller detector volumes and lower count statistics than conventional airborne surveys.

As a result, Medusa Radiometrics prioritises spectral stability and processing methods, such as NASVD, to manage noise and preserve meaningful spectral structure in low-count survey environments.





CHAPTER 2

Spectral smoothing with NASVD

Because all subsequent corrections and calculations depend on the stability of the gamma-ray spectrum, noise reduction must be addressed at the very start of radiometric processing.

Spectral smoothing is a foundational step that determines the reliability of every downstream result. Once the signal has been distorted or suppressed at this stage, it cannot be recovered through later filtering or correction.

Among the methods used for spectral smoothing, NASVD has become a standard approach. Rather than smoothing individual channels independently, NASVD evaluates

the full spectral shape and separates a coherent signal from statistically incoherent noise. This allows noise to be reduced while preserving the relationships between spectral channels that are critical for accurate stripping and concentration calculations.

In the sections that follow, the focus is not only on how NASVD works, but on how to make defensible, survey-specific choices about eigenvector selection in practice, moving beyond default settings to decisions grounded in detector characteristics, count statistics, and acquisition conditions.

What is an eigenvector in radiometric processing?

In NASVD-based smoothing, an eigenvector corresponds to a distinct pattern of variation in gamma-ray spectra. The strongest eigenvectors capture coherent spectral structures associated with the geological signal, while higher-order eigenvectors increasingly represent random noise.

This is similar to listening to an orchestra in a concert hall, where the strings and bass carry the music, while audience movement and ambient sounds add background interference. NASVD separates these dominant spectral components from incidental noise.

If too few eigenvectors are retained, however, legitimate but lower-intensity parts of the “music” may be discarded along with the noise, leaving only the loudest elements and suppressing subtle geological variation. Eigenvector selection, therefore, requires care to reduce noise without stripping away meaningful signal.

NASVD in practice: Balancing noise reduction and signal preservation

The key decision in NASVD lies in how many eigenvectors to retain. This choice determines the balance between noise reduction and signal preservation. Retaining too few eigenvectors can oversmooth the data, suppressing legitimate variability and reducing sensitivity to subtle geological features. Retaining too many can introduce residual noise that propagates into subsequent processing steps and compromise the stability of derived radioelement concentrations.

This balance is particularly critical for uranium, which presents specific challenges in radiometric processing. In conventional

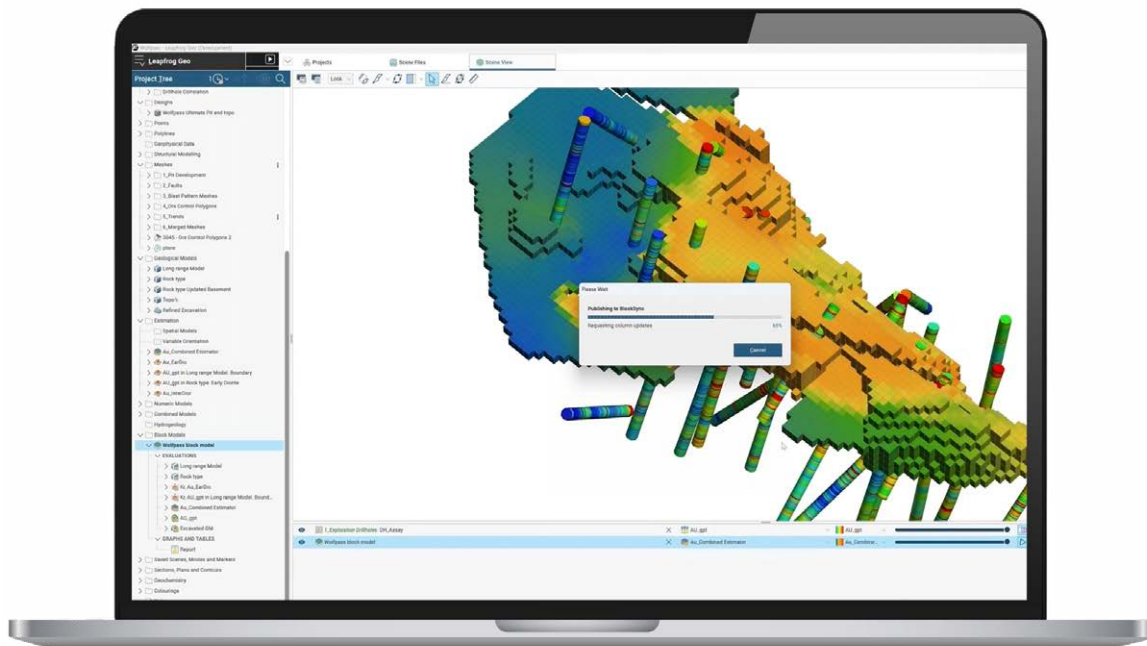
window-based approaches, the uranium signal is poorly separated from the thorium background and is further complicated by atmospheric radon, both of which can introduce instability if not properly managed.

In modern surveys, especially those using smaller detectors or lower count statistics, these effects can be amplified, thereby increasing uncertainty in the derived uranium channel. NASVD-based spectral smoothing plays an important role in stabilising the full spectrum prior to stripping and correction, thereby reducing noise while preserving genuine geological variability in uranium responses.

There is no single correct number of eigenvectors that applies to all surveys. Optimal selection depends on detector size, count rates, survey geometry, and acquisition conditions. For this reason, eigenvector selection is a professional judgement call.

Importantly, NASVD is distinct from conventional filtering. Rather than simply

suppressing high-frequency variation, it preserves the underlying spectral structure while removing statistically incoherent components. Practitioners assess eigenvector selection by examining eigenvalue spectra, comparing reconstructed outputs, and considering how smoothing affects downstream corrections and derived products.





CHAPTER 3

Making eigenvector selection decisions in practice

There is currently no consensus among experts on how many eigenvectors should be retained in NASVD processing. In practice, selection has often relied on careful and sometimes time-consuming evaluation of multiple reconstructions, comparison of downstream products, and analyst experience.

Such iterative approaches can provide insight, but they are not always practical in operational environments where consistency and efficiency are required.

Oasis montaj addresses this challenge by providing a clear statistical framework for eigenvector selection. Eigenvalues plotted against eigenvector number typically exhibit a transition from steep decay to a plateau. This inflection point provides an explicit, data-driven indication of where a statistically coherent signal gives way to noise-dominated components.

While professional judgement remains important, this statistical approach reduces reliance on arbitrary defaults and improves consistency between datasets. Rather than selecting eigenvectors based on visual smoothness alone, practitioners are supported by a transparent and defensible method aligned with the statistical structure of the data itself.



CHAPTER 4

A structured radiometrics data processing workflow in Oasis montaj

A structured workflow is particularly important in radiometrics because many corrections are interdependent. Errors introduced early, such as unstable spectra or inconsistent windowing, can trickle through the processing chain and only become apparent during interpretation.

Making each step explicit allows these issues to be identified and addressed early, when they are easier to diagnose and correct.



Benefits of a radiometrics data processing workflow

A structured radiometric processing workflow provides operational and technical advantages that become increasingly important in modern survey contexts:



Reduced processing risk

Each correction step produces its own output channel, making the impact of background, height, and stripping corrections visible rather than implicit.



Faster onboarding and knowledge transfer

A fixed correction sequence reflects accepted radiometric practice, while still exposing key decision points such as NASVD eigenvector selection.



Improved quality control efficiency

Before-and-after channels allow instability in uranium or ratio products to be traced back to earlier processing choices.



Repeatability and auditability

Explicit parameters and channel naming allow processing decisions to be reviewed, justified, and reproduced across surveys.



Scalability through automation

Once validated, the same radiometric workflow can be scripted and applied consistently without hiding professional judgement.

In Oasis montaj, radiometric processing is implemented through the Radiometric Processing System (RPS), a guided, step-by-step workflow that aligns with established best practices for airborne gamma-ray spectrometry.¹

Rather than treating processing as a collection of isolated tools, the RPS organises each stage in a logical sequence that reflects how corrections and transformations interact.

While individual surveys may differ in detector configuration or objectives, the workflow's underlying structure remains consistent. This consistency is critical for ensuring that key decisions, particularly those related to noise reduction and correction, are applied deliberately rather than implicitly.

Data import and preparation

Radiometric processing begins with the import of raw spectrometer data, including the full gamma-ray spectrum and associated navigation information such as position and altitude. Preserving the full spectrum at this stage is essential, as subsequent steps rely on spectral shape rather than pre-windowed counts.

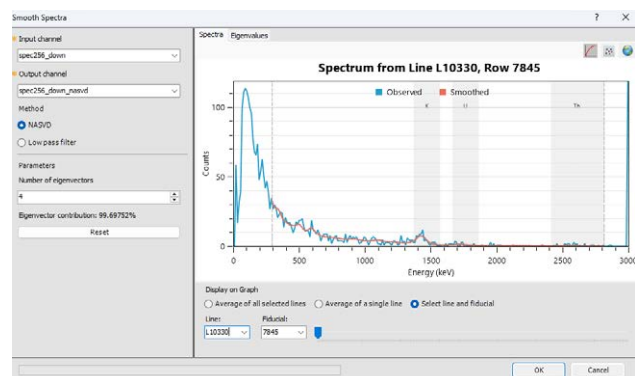
Early inspection of the imported data allows practitioners to identify obvious acquisition issues, verify metadata, and confirm that the dataset is suitable for further processing before any irreversible transformations are applied.

Spectral smoothing using NASVD

NASVD-based spectral smoothing is typically applied early in the workflow, as it directly influences the stability of all downstream

corrections. At this stage, the practitioner selects how many eigenvectors to retain when reconstructing the spectra, based on an assessment of noise structure, count rates, and survey conditions.

Oasis montaj makes this step explicit by allowing direct comparison between raw (obtained) and smoothed spectra. In a well-chosen reconstruction, high-frequency statistical noise is reduced while primary photopeaks and overall spectral shape are preserved. Over-smoothing, by contrast, may dampen peak structure or flatten subtle features, while under-smoothing leaves residual noise visible in the reconstructed curve.



The Smooth Spectra dialog from the RPS showing raw spectra (blue) and NASVD-smoothed spectra (orange).

Selecting the number of eigenvectors

Visual comparison alone, however, is not sufficient to justify the number of eigenvectors retained.

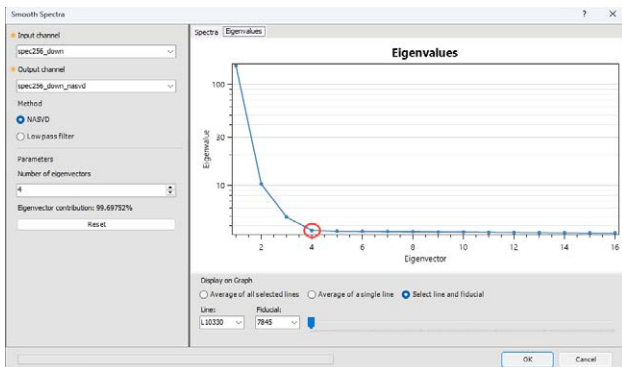
After calculating eigenvectors, the Eigenvalues tab in the Smooth Spectra dialog provides an explicit statistical basis for selecting how many components to retain.

1. International Atomic Energy Agency (IAEA). Guidelines for radioelement mapping using gamma ray spectrometry data.

The plot displays the eigenvalue number against the eigenvalue, where each eigenvalue represents the amount of variance captured by that component. In most datasets, the curve shows a steep initial decay followed by a transition to a flatter plateau. This inflection marks the point at which a statistically coherent signal gives way to noise-dominated components.

Oasis montaj also reports the Eigenvector Contribution, indicating the percentage of the original spectral variance preserved in the reconstructed output.

Selecting the number of eigenvectors near the transition from steep decay to plateau allows users to retain statistically meaningful variance while discarding incoherent noise. This provides a transparent, empirical alternative to traditional rules of thumb, grounding the decision directly in the behaviour of the dataset.

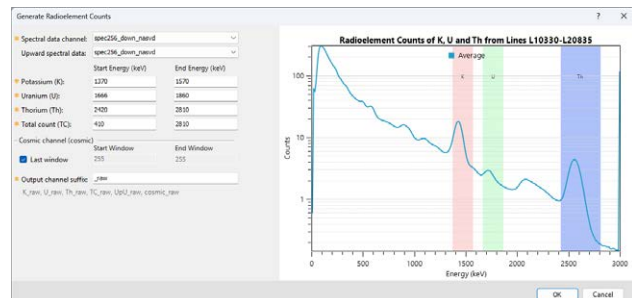


The Eigenvalues tab in the Smooth Spectra dialog, showing Eigenvalues plotted against the eigenvector number. The transition from steep decay to plateau provides a statistical indication of where a signal gives way to noise. The Eigenvector Contribution field reports the percentage of spectral variance retained in the reconstructed output.

Generation of radioelement window counts

Once spectra have been smoothed, counts are generated for the standard potassium, uranium, thorium, and total count energy windows. Window positioning can be adjusted to reflect the characteristics of the input spectra, ensuring accurate capture of photopeaks.

Because these windowed counts form the basis for all subsequent corrections, their quality is directly dependent on the choices made during spectral smoothing.

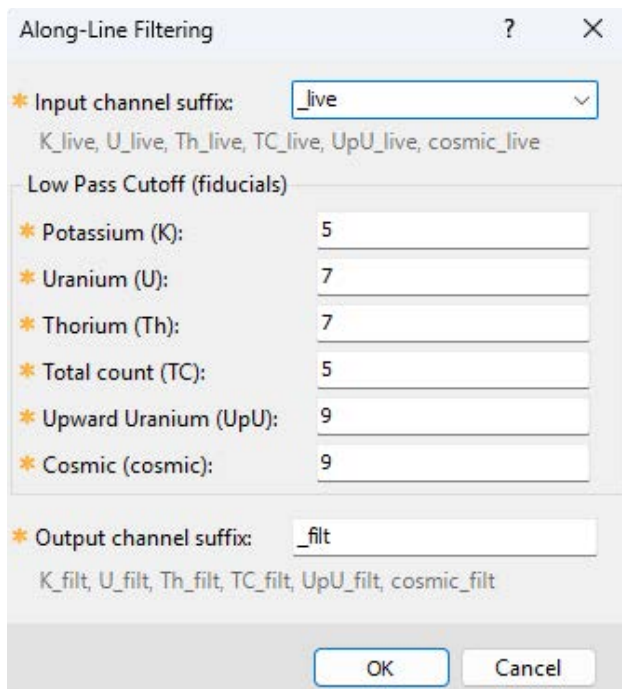


The Generate Radioelement Counts dialog, showing gamma-ray spectra and the energy windows used to calculate potassium, uranium, and thorium counts.

Filtering and noise stabilisation

Low-pass filtering is applied to improve the statistical reliability of selected channels, including cosmic background and atmospheric radon estimates. Filtering may also be used to stabilise corrected radioelement channels prior to calculating ratios.

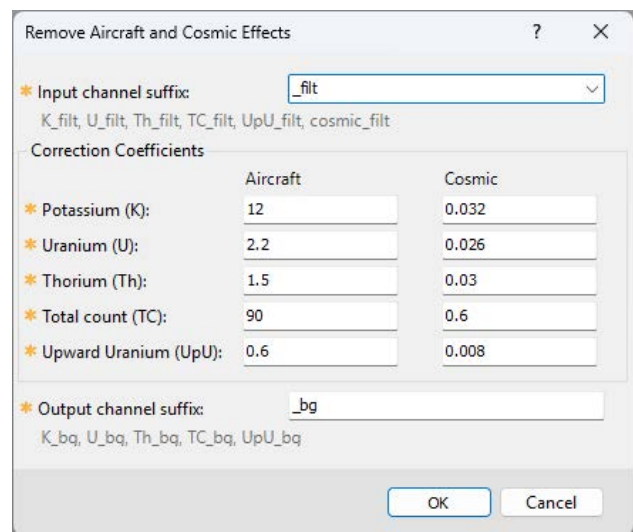
Importantly, filtering is designed as a controlled enhancement step rather than a substitute for proper smoothing or correction. Applying it within a structured workflow helps prevent over-filtering and ensures that its effects are understood and traceable.



The Along-Line Filtering dialog, showing low-pass filtering applied to radioelement and cosmic channels to improve statistical stability.

which these corrections are applied matters, as each step assumes that previous influences have already been addressed.

The RPS enforces this sequencing, reducing the risk that corrections are skipped, reordered, or applied inconsistently across datasets.



The Remove Aircraft and Cosmic Effect dialog, showing correction of aircraft background and cosmic radiation prior to downstream processing. Parameters shown are illustrative. Correction coefficients are survey-specific and derived using established airborne gamma-ray spectrometry procedures (e.g., IAEA).

Applying corrections

A series of essential corrections is then applied to the windowed counts. These include correcting data for variations in flight height and atmospheric conditions and the removal of various effects, including aircraft and cosmic background radiation, atmospheric radon, and Compton scattering. The order in

Calculation of radioelement concentrations and ratios

After correction, windowed counts are converted into apparent ground concentrations of potassium, uranium, and thorium. Because uranium and thorium are derived from daughter isotopes, these values are expressed as equivalent concentrations.

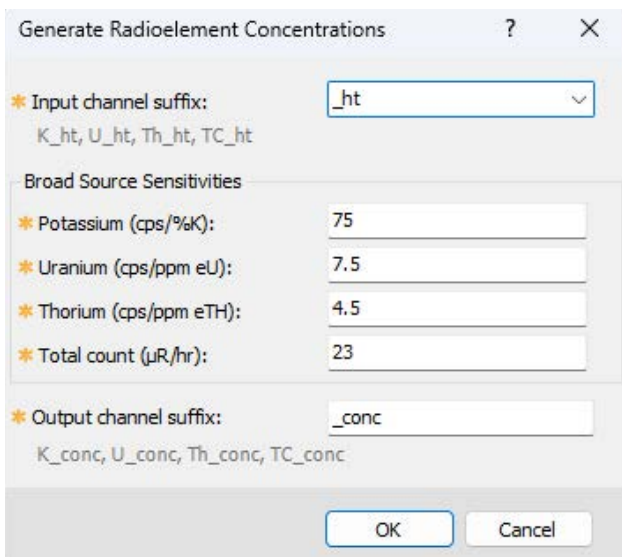
Radioelement ratios are then calculated to highlight subtle variations that may not be apparent in individual concentration channels. The reliability of these products depends directly on the stability of earlier processing steps, particularly smoothing and correction.

Final QC focuses on internal consistency: comparison between channels, inspection of ratio behaviour, and review of spatial coherence. Importantly, QC at this stage is diagnostic rather than corrective. Issues identified here typically point back to earlier decisions, particularly in smoothing and correction.

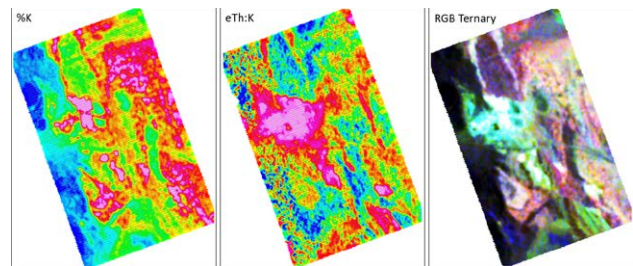
deliverables, such as concentration maps, ratio images, and ternary images that combine potassium, uranium, and thorium responses into a single visual representation.

At this point, the value of a structured workflow becomes clear: each product can be traced back through a documented sequence of processing decisions.

This transparency supports quality control, peer review, and repeatability, especially important when datasets are revisited, reprocessed, or compared across surveys.



The Generate Radioelement Concentrations dialog, showing survey-specific source sensitivities used to calculate potassium, uranium, and thorium concentrations.



Examples of final radiometric data products, from left to right: (a) potassium (%), (b) equivalent thorium-to-potassium ratio, and (c) an RGB ternary image.

Creation of final products and documentation

The final stage of the workflow involves gridding the processed data and generating

Why structure matters

By embedding radiometric processing within a guided, transparent workflow, Oasis montaj helps ensure that critical processing decisions are visible rather than hidden behind defaults. This is particularly important for NASVD-based smoothing, where the consequences of eigenvector selection extend throughout the entire processing chain.



CHAPTER 5

The future of radiometric processing

Confidence in radiometric interpretation comes from understanding what each processing step does, recognising where professional judgement is required, and applying that judgement consistently within a transparent workflow. This need becomes more pronounced as radiometrics is applied across a wider range of platforms, detector configurations, and acquisition environments.

Structured processing environments help make these decisions visible, repeatable, and defensible.

In Oasis montaj, radiometric workflows embed established processing sequences, NASVD-based smoothing, and quality checks in a way that reflects long-standing best practice while remaining adaptable to modern datasets. This provides a stable processing foundation as survey methods continue to evolve.

At the same time, technologies such as Medusa Radiometrics are expanding where and how radiometric data can be collected, enabling surveys on foot, by vehicle, or from low-altitude airborne platforms. While this flexibility increases access, it also introduces greater variability in acquisition conditions, reinforcing the importance of disciplined, transparent processing to maintain confidence in the final results.

Within the broader Seequent ecosystem, this alignment supports connected workflows that carry radiometric data from acquisition through processing and interpretation without obscuring critical decisions. Rather than automating judgment away, the emphasis is on making those decisions explicit, reviewable, and defensible as radiometric applications continue to expand.

Next steps and related resources

The following resources provide practical methods for exploring and applying the workflows discussed in this paper.

Work through a complete radiometric processing workflow

Access the [Radiometric Mastery Session](#) to work through real datasets and see how key processing decisions affect interpretation. Have a question about NASVD or eigenvector choice? Continue the discussion on the [Seequent Community](#).

Discuss modern radiometric workflows with Seequent

[Connect with our experts](#) to discuss current survey challenges, best-practice processing, and how radiometrics fits into integrated interpretation workflows.

See how these workflows apply to different data scenarios

[Request a demo](#) and explore how radiometric processing workflows in Oasis montaj support different survey types and data quality scenarios.

See how Medusa Radiometrics integrates with Oasis montaj

Read more about the [Seequent x Medusa Radiometrics](#) collaboration and how modern radiometric acquisition is supported by transparent, standards-aligned processing workflows in Oasis montaj.

Understand the underground to build a better world.

Seequent, The Bentley Subsurface Company, helps organisations to understand the underground, giving them the confidence to make better decisions faster.

Seequent builds world-leading technology that is at the forefront of Earth sciences, transforming the way our customers work.

Every day we help them develop critical mineral resources more sustainably, design and build better infrastructure, source renewable energy, and reduce their impact on the environment.

Seequent operates in 145+ countries while proudly maintaining headquarters in New Zealand.