

Automatic Radiance Level Data Quality and Signal-to-Noise Assessment for Hyperspectral Imaging

Simon A. Trim^{1,2}, Daniel Schläpfer¹

¹ ReSe Applications LLC, Wil, Switzerland - simon@rese-apps.com, daniel@rese-apps.com

² Remote Sensing Laboratories, Department of Geography, University of Zurich, Zurich, Switzerland

Keywords: Hyperspectral imaging, data quality, signal-to-noise ratio, radiance, L1 products.

Abstract

Raw hyperspectral imaging (HSI) data require multiple processing steps before physically meaningful reflectance products can be derived. Quality issues at the radiance level – e.g., detector artifacts, low-signal regions, or excessive noise – can propagate through the processing chain and compromise subsequent analyses. HSI workflows typically focus on calibration and correction but often lack standardized integrity checks prior to atmospheric compensation. To address this gap, we developed an automated Data Quality Assessment (DQA) toolbox within the EU Horizon m4mining project. The toolbox is integrated into INTACOR[®], an automated processing pipeline for airborne hyperspectral imagery developed by ReSe Applications LLC, and provides quality assessment of calibrated radiance data prior to subsequent processing. It identifies bad pixels and low-signal regions, derives vegetation information, and estimates spatially resolved apparent signal-to-noise ratio (SNR) using spectral differencing and local residual statistics, from which robust band-wise SNR statistics are derived. The toolbox produces spatially explicit quality layers and an automated report to support rapid assessment of acquired data. Application to airborne hyperspectral mining data demonstrates its ability to characterize spatial and spectral variations in data quality, providing a standardized radiance-level quality assessment stage within an operational HSI processing chain.

1. Introduction

Hyperspectral imaging (HSI) has become an essential tool for environmental monitoring, mineral exploration, and other Earth observation applications (Balz, 2022, Hong et al., 2026). Raw hyperspectral data must undergo several processing steps - including radiometric calibration, geocoding, and atmospheric correction - before physically meaningful reflectance retrieval is possible (Gao et al., 2009). However, quality issues at the radiance level can propagate through the entire workflow and compromise the validity of derived products.

Despite the critical importance of early-stage data quality checks, systematic and automated approaches remain underdeveloped. Most existing HSI processors provide calibration and correction but lack a standardized mechanism for assessing data integrity and signal quality prior to atmospheric compensation.

To address this gap, we have developed an automated Data Quality Assessment (DQA) toolbox as part of the m4mining project (m4mining, 2023-2026). The toolbox is integrated into INTACOR (ReSe Applications, 2026a), which was developed by ReSe Applications as the successor to the near real-time (NRT) processor described in (Trim and Schläpfer, 2024). With INTACOR providing an automated processing chain for airborne hyperspectral imagery, including geometric and atmospheric correction, the DQA toolbox fits in as a standardized quality diagnostics at the calibrated-radiance stage. The DQA generates a multi-layer ENVI file summarizing key quality metrics and a comprehensive PDF report with visualizations and recommendations for the end-user. The system is modular and designed to support field operations in mining environments, where rapid data quality evaluation is essential.

This paper presents the DQA workflow, describes its key functionality, and demonstrates its potential as a standardized quality control stage in operational HSI pipelines.

2. Data Quality Assessment Workflow

The DQA toolbox is implemented as a modular Python workflow consisting of preprocessing, quality metric extraction, visualization, and reporting components. Figure 1 illustrates the main processing chain. The workflow operates directly on calibrated radiance cubes and produces two main outputs:

- a multi-layer ENVI file containing pixel-wise quality metrics, and
- an automated PDF report with quality summaries, plots, and recommendations.

The workflow follows these main steps: (1) identification of bad pixels and low-signal regions, (2) apparent signal-to-noise ratio (SNR) estimation, (3) export of representative quality layers, and (4) automated report generation. The modular structure facilitates future extensions with additional quality metrics.

2.1 Bad pixel detection

Bad pixels, including saturated, zero-valued, NaN, and other no-data pixels, are automatically detected and combined into a single validity mask. Saturated pixels are identified relative to the valid radiometric range of the input data, while zero-valued, non-finite, negative, and metadata-defined no-data pixels are detected directly from the radiance cube. The resulting mask provides the end-user with an immediate spatial overview of detector artifacts and data dropouts and is propagated through subsequent DQA calculations so that invalid samples do not contribute to derived quality statistics.

Identifying these pixels at the radiance level is important because invalid measurements may otherwise propagate through




Figure 1. Overview of the automated DQA workflow, from input radiance cube to ENVI layers and PDF report. Solid arrows indicate the forward flow of data and module calls, while dotted arrows represent the feedback of metadata to the report generator for compilation of quality metrics and recommendations.

geocoding, atmospheric correction, and subsequent mineral mapping. The bad-pixel layer is therefore retained as an explicit DQA output rather than simply removing affected samples internally.

2.2 Low signal mask

Cast shadows and other low signal regions such as water pixels are automatically identified using percentile thresholding applied to either a near-infrared band (targeting the 850-870 nm wavelengths) for VNIR data (as shadows and low-signal areas exhibit dramatically reduced reflectance in the near-infrared region while maintaining good contrast with illuminated surfaces), or a band around 2100 nm wavelength for SWIR data. The 20th percentile of valid (non-bad) pixel values in the selected band serves as the threshold. A binary low signal mask is exported as a separate layer and used internally to restrict SNR analysis to areas with viable signal.

2.3 SNR layers

SNR is a critical metric for assessing the utility of hyperspectral data. While SNR is typically characterized under laboratory conditions (Woodward et al., 2009), image-based methods provide an indication of data quality under real acquisition conditions. On single images, SNR is often retrieved by comparing the average band signal to spatial noise, retrieved from filtered

homogeneous image areas (Gao, 1993). Noise characterization in hyperspectral imagery is complicated by the presence of multiple noise sources and their spectral and spatial characteristics (Rasti et al., 2018).

In HSI imagery, the high band-to-band correlation of the underlying spectral signal can instead be exploited to characterize noise in the spectral domain.

Thus, the DQA toolbox estimates *apparent* SNR using *spectral differencing*, exploiting the strong correlation of the underlying signal between adjacent spectral bands.

For each spatial pixel p and spectral band i , a spectral residual is calculated as:

$$d_i(p) = L_i(p) - \frac{L_{i-1}(p) + L_{i+1}(p)}{2},$$

where $L_i(p)$ is the measured radiance.

Assuming that the underlying spectrum varies smoothly across the three bands, this operation suppresses the spectrally correlated signal while retaining high-frequency components, including sensor noise. Applying the operation to all valid pixels produces a *difference image* - or equivalently a *spectral residual image* - for each band.




Figure 2. Visualizations produced by the DQA toolbox, based on vertical/nadir view HSI VNIR data acquired by m4mining in Greece. These images are included in the generated report and provide an overview of the regions of interest.

Considering only the noise contribution,

$$d_i = n_i - \frac{n_{i-1} + n_{i+1}}{2},$$

where n_i denotes the noise in band (i). Squaring gives

$$d_i^2 = n_i^2 + \frac{1}{4}n_{i-1}^2 + \frac{1}{4}n_{i+1}^2 + \frac{1}{2}n_{i-1}n_{i+1} - n_i n_{i-1} - n_i n_{i+1}.$$

Taking the expected value $E[\cdot]$, and assuming independent, zero-mean noise of approximately equal variance σ^2 in the three bands,

$$E[n_i^2] = E[n_{i-1}^2] = E[n_{i+1}^2] = \sigma^2,$$

while the cross terms vanish,

$$E[n_i n_{i-1}] = E[n_i n_{i+1}] = E[n_{i-1} n_{i+1}] = 0$$

The residual variance is therefore

$$\text{Var}(d_i) = \sigma^2 + \frac{1}{4}\sigma^2 + \frac{1}{4}\sigma^2 = \frac{3}{2}\sigma^2$$

Thus, the apparent single-band noise standard deviation is estimated from the residual standard deviation as:

$$\hat{\sigma}_i = \frac{\sigma_{d,i}}{\sqrt{3/2}}$$

To obtain spatially resolved apparent SNR, the residual standard deviation is estimated within a local spatial neighborhood $\Omega(p)$. The local residual variance is

$$\sigma_{d,i,\text{local}}^2(p) = E_{\Omega(p)}[d_i^2] - (E_{\Omega(p)}[d_i])^2,$$

where $E_{\Omega(p)}[\cdot]$ denotes the local mean within the spatial neighborhood centered on pixel p . The local residual standard deviation is therefore

$$\sigma_{d,i,\text{local}}(p) = \sqrt{E_{\Omega(p)}[d_i^2] - (E_{\Omega(p)}[d_i])^2}.$$

Applying the $\sqrt{3/2}$ correction gives the spatially resolved apparent SNR:

$$\text{SNR}_i(p) = \sqrt{\frac{3}{2}} \left(\frac{L_i(p)}{\sigma_{d,i,\text{local}}(p)} \right).$$

Local moments are calculated using valid pixels only, and a 9×9 pixel neighborhood is used as a compromise between statistical stability and spatial localization. Pixels identified as low-signal are excluded from the SNR statistics used for DQA assessment.

A representative band-wise apparent SNR is subsequently derived from the spatial SNR estimates as the median across valid pixels outside low-signal regions:

$$\text{SNR}_{i,\text{band}} = \text{median}_{p \in V_i} [\text{SNR}_i(p)],$$

where V_i denotes the set of valid, non-low-signal pixels in band i . The 25th and 75th percentiles (P_{25} and P_{75}) of the same spatial SNR distribution are additionally calculated to characterize within-scene variability. This provides a robust spectral summary derived from the same local estimator as the SNR maps, without estimating a single residual variance across the entire heterogeneous scene.




Figure 3. Examples of band-wise apparent SNR as a function of wavelength for VNIR and SWIR data, as generated by the DQA toolbox. The plotted SNR represents the median of the spatially resolved apparent-SNR estimates within each band, using valid pixels outside low-signal regions. The 25th and 75th percentiles characterize the spatial variability of SNR within each wavelength.

In this way, apparent SNR is characterized at complementary spatial and spectral scales: local SNR maps reveal spatial variations in data quality within individual bands (see Figure 2), while their robust band-wise summaries characterize apparent SNR as a function of wavelength. Representative SNR maps are additionally exported for 10 selected wavelengths. These bands are chosen according to sensor type (VNIR or SWIR) and the needs of mining applications. For example, wavelengths between 400–460 nm are relevant for jarosite detection, 600–1100 nm for iron oxides, and 2100–2450 nm for other key mineral features.

Figure 3 shows examples of band-wise apparent SNR as a function of wavelength for VNIR and SWIR data. The median of the spatially resolved apparent-SNR estimates is displayed for each wavelength, with the 25th and 75th percentiles indicating spatial variability within the assessed image regions.

As an additional consistency assessment, apparent SNR was independently evaluated for two spatial subsets representing predominantly bright and dark terrain within the same acquisition (see Figure 4). The subsets were selected as fixed image regions rather than by thresholding individual pixel radiance, allowing the response of the estimator to different scene signal levels to be examined while retaining the spatial and spectral variability present within each region.

3. NDVI layer

To assist geologists in avoiding vegetation-contaminated pixels, the workflow computes and exports an NDVI layer. This enables masking or down-weighting of vegetated areas when generating geological products, ensuring that quality assessment is tailored to the needs of mining surveys.

3.1 Geometric quality assessment

A complementary quality metric - concerning pixel geometry after geocoding - can be assessed based on the image geometry map (IGM) output from PARGE (ReSe Applications, 2026b). It quantifies the relation between along- and across-track pixel sampling distance (see Figure 5).

Further quantities for geometric quality assessment are about to be included as standard outputs. These are the relative number of non-geocoded pixels (e.g. due to lack of DSM information), the relative accuracy of overlap regions, and the pixel smear of non-stabilized imagery.

4. Outputs: ENVI Layers and PDF Report

The DQA toolbox produces a consolidated ENVI file with 13 layers:

- a bad pixel mask
- a shadow/low signal mask
- 10 SNR maps at representative wavelengths
- an NDVI layer

In addition, the toolbox generates a comprehensive PDF report that summarizes the analysis. The report includes band-by-band SNR curves, various visualizations such as for the shadow mask and SNR and NDVI maps, and recommendations for further processing or data re-acquisition if thresholds are not met.

Figure 6 shows example excerpts from the automated report.

5. Results and Discussion

The DQA toolbox provides a standardized and automated means of assessing HSI radiance data prior to atmospheric correction. By operating at the radiance level, it enables detector artifacts, noise characteristics, and low-signal effects to be identified before they propagate into reflectance products. This capability is particularly relevant for mining surveys, where re-acquisition costs are high and rapid decision-making is required.

The implemented SNR approach uses a standard-deviation-based estimate of residual variability to derive apparent SNR at local spatial scales. Unlike a single scene-wide residual variance, local estimation limits the influence of scene-wide radiometric heterogeneity. The resulting estimates nevertheless remain image-based measures rather than laboratory-characterized sensor




Figure 4. Consistency assessment of the apparent SNR estimator using spatial subsets representing predominantly bright and dark terrain. Left: RGB representation of the assessed image showing the locations of the two spatial subsets. Right: VNIR and SWIR band-wise apparent SNR for the dark subset, full assessed region, and bright subset. Band-wise values represent the median of the locally estimated apparent SNR. Shaded intervals indicate the 25th–75th percentile range. The bright subset exhibits higher apparent SNR than the dark subset across all evaluated bands.




Figure 5. Geometric quality check, with (from left to right): SWIR image of quarry, across-track pixel size [0-0.2m], along-track pixel size [0-0.2m], and relation across- to along-track [0-3].

SNR and rely on locally smooth spectral behaviour and approximately independent noise between adjacent bands. Genuine spectral curvature and inter-band noise correlation may therefore influence the estimates. Band-wise apparent SNR is consequently reported as the median of the local SNR distribution, with its interquartile range characterizing spatial variability.

Higher apparent SNR was generally observed for the SWIR data than for the VNIR data. This is qualitatively consistent with the different detector and sampling characteristics of the Mjolnir VS-620 channels, for which the SWIR sensor is specified with substantially higher peak SNR (Norsk Elektro Optikk AS, 2026). However, the image-derived values should not be directly compared with laboratory sensor specifications, as they additionally depend on scene radiance, spectral structure, and local spatial variability.

Regarding the consistency assessment, apparent SNR was consistently higher in the bright spatial subset than in the dark subset, with this relationship observed in all spectral bands for both VNIR and SWIR data (Figure 4). The median bright-to-dark SNR ratio across wavelengths was 1.48 for VNIR and 1.86 for SWIR. Despite the different absolute SNR levels, the sub-

Report summary for VNIR data cube

SNR Statistics

Metric	Value
Mean SNR	128.47
Median SNR	138.83
Min SNR	11.34
Max SNR	208.84
% Excellent SNR (≥ 100)	65.00%
% Good SNR (30-99)	32.50%
% Poor SNR (< 30)	2.50%
Overall Quality Score	100.00/100
Bad Pixels (total)	0.00%
% Shadow Pixels	18.29%
% Non-Shadow Pixels	81.71%

SNR by Spectral Region (VNIR)

Region	Mean SNR	Median SNR	Min SNR	Max SNR
400-500nm	116.13	118.45	40.52	173.39
500-600nm	180.16	188.39	86.93	208.64
600-700nm	180.07	193.65	81.82	208.84
700-800nm	139.19	154.57	11.34	184.49
800-900nm	100.92	95.68	57.89	141.21
900-1000nm	45.21	41.07	22.65	68.49

Region Analysis:

- Best SNR region: 500-600nm (Mean SNR: 180.16)
- Worst SNR region: 900-1000nm (Mean SNR: 45.21)

Quality Assessment

HIGH QUALITY: This dataset has excellent SNR characteristics

Recommendations

- No specific processing corrections needed - the dataset shows good SNR characteristics
- Proceed with standard processing workflow; data quality is suitable for most analysis purposes

Report summary for SWIR data cube

SNR Statistics

Metric	Value
Mean SNR	138.69
Median SNR	90.42
Min SNR	7.07
Max SNR	389.72
% Excellent SNR (≥ 100)	46.67%
% Good SNR (30-99)	33.67%
% Poor SNR (< 30)	19.67%
Overall Quality Score	85.78/100
Bad Pixels (total)	6.14%
% No-Data Pixels	6.14%
% Shadow Pixels	16.12%
% Non-Shadow Pixels	83.88%

SNR by Spectral Region (SWIR)

Region	Mean SNR	Median SNR	Min SNR	Max SNR
900-1200nm	257.98	279.46	44.35	356.20
1200-1500nm	166.38	115.30	8.25	389.72
1500-1800nm	230.72	239.28	67.31	311.09
1800-2100nm	37.78	26.55	7.07	111.67
2100-2400nm	67.89	69.63	30.52	112.52

Region Analysis:

- Best SNR region: 900-1200nm (Mean SNR: 257.98)
- Worst SNR region: 1800-2100nm (Mean SNR: 37.78)

Quality Assessment

HIGH QUALITY: This dataset has excellent SNR characteristics

Recommendations

- Flag or remove bands with SNR < 10 for analysis
- Consider spectral binning to improve SNR in low-quality bands
- Note: 6.14% of pixels were identified as bad pixels
- These bad pixels were excluded from SNR and shadow/non-shadow calculations

Figure 6. Example excerpts from the automated DQA PDF report. The report summarizes key quality metrics including SNR statistics, bad pixel and shadow masks, and provides visualizations and recommendations derived from the metadata feedback of the workflow.

sets retained coherent wavelength-dependent behaviour, supporting the consistency of the image-based estimator across regions with contrasting signal levels.

The DQA workflow is already integrated into INTACOR (ReSe Applications, 2026a), the automated HSI processing pipeline developed by ReSe Applications. This enables radiance-level quality assessment to form part of the operational processing chain rather than being performed as a separate pre-processing step. The current DQA implementation has reached Technology Readiness Level (TRL) 5 through validation on relevant mining data. Although currently validated using NEO's HySpex Mjolnir VS-620 (Norsk Elektro Optikk AS, 2026), the workflow is designed to be sensor-agnostic and adaptable to other VNIR/SWIR imaging systems.

The modular architecture allows the scope of the DQA to be expanded beyond its current emphasis on radiometric and signal-related quality. Planned extensions include characterization of spectral smile and spatial keystone, providing information on spatial variation in wavelength calibration and spectral-spatial misregistration, respectively (Yokoya et al., 2010). Further development will include additional raw-data checks, geometric and reflectance-level quality assessment, and support for additional sensors. Together, these developments will extend the DQA toward a more comprehensive characterization of hyperspectral data quality.

6. Conclusion and Outlook

We have presented an automated DQA toolbox for HSI, first developed within the m4mining project and now integrated into the INTACOR automated HSI processing pipeline. The toolbox provides bad-pixel detection, low-signal masking, NDVI calculation, spatially resolved apparent-SNR estimation, robust band-wise SNR statistics, and automated reporting. By operating directly on calibrated radiance data, its outputs allow users to identify problematic spatial regions and spectral ranges before proceeding with geocoding and atmospheric correction.

The image-based SNR approach combines spectral differencing with local residual statistics, enabling spatial SNR maps and robust spectral summaries to be derived from a common estimator. Although these estimates do not replace laboratory sensor characterization, they provide an indication of signal quality under actual acquisition and scene conditions.

Future work will focus on validation across additional sensors and operational field campaigns beyond mining and on expanding the range of quality metrics. Planned developments include characterization of spectral smile and spatial keystone, geometric quality assessment, and extension of the DQA framework to reflectance-level products. By standardizing early-stage quality assessment within an automated processing chain, the DQA toolbox provides a foundation for more reliable downstream hyperspectral processing and analysis.

7. Acknowledgments

This work has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

References

Balz, T., 2022. New Developments in Remote Sensing for the Environment [Special issue]. *Remote Sensing*, 14.

Gao, B.-C., 1993. An operational method for estimating signal to noise ratios from data acquired with imaging spectrometers. *Remote Sensing of Environment*, 43(1), 23 – 33. [https://doi.org/10.1016/0034-4257\(93\)90061-2](https://doi.org/10.1016/0034-4257(93)90061-2).

Gao, B.-C., Montes, M. J., Davis, C. O., Goetz, A. F., 2009. Atmospheric correction algorithms for hyperspectral remote sensing data of land and ocean. *Remote Sensing of Environment*, 113, S17-S24. Imaging Spectroscopy Special Issue. <https://doi.org/10.1016/j.rse.2007.12.015>.

Hong, D., Li, C., Yokoya, N., Zhang, B., Jia, X., Plaza, A., Gamba, P., Benediktsson, J. A., J., C., 2026. Hyperspectral imaging. *Nat Rev Methods Primers*, 6(19), <https://doi.org/10.1038/s43586-026-00470-x>.

m4mining, 2023-2026. Multi-scale, Multi-sensor Mapping and dynamic Monitoring for sustainable extraction and safe closure in Mining environments. <https://www.m4mining.eu/home/> (accessed September 6th, 2026).

Norsk Elektro Optikk AS, 2026. HySpex Mjolnir VS-620. <https://www.hyspex.com/hyspex-products/hyspex-mjolnir/mjolnir-vs-620> (accessed September 6th, 2026).

Rasti, B., Scheunders, P., Ghamisi, P., Licciardi, G., Chanussot, J., 2018. Noise Reduction in Hyperspectral Imagery: Overview and Application. *Remote Sensing*, 10(3), <https://doi.org/10.3390/rs10030482>.

ReSe Applications, 2026a. INTACOR. <https://intacor.com/> (accessed September 6th, 2026).

ReSe Applications, 2026b. PARGE. <https://parge.com/> (accessed September 6th, 2026).

Trim, S. A., Schlöpfer, D., 2024. Development of an Integrated Physical Hyperspectral Imaging (HSI) Point Cloud Data Processor. 2024 14th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS), 1–5. <https://doi.org/10.1109/WHISPERS65427.2024.10876528>.

Woodward, J., Brown, S., Smith, A., Lykke, K., 2009. Hyperspectral Imager Characterization and Calibration. 2009 IEEE International Geoscience and Remote Sensing Symposium, Cape Town, -1, <https://doi.org/10.1109/IGARSS.2009.5418005>.

Yokoya, N., Miyamura, N., Iwasaki, A., 2010. Preprocessing of hyperspectral imagery with consideration of smile and keystone properties. *SPIE Proceedings*, 7857, 78570B. <https://doi.org/10.1117/12.870437>.