REALISTIC SENSOR NOISE FOR SIMULATED HYPERSPECTRAL IMAGES: A SENSOR-AGNOSTIC UNCERTAINTY MODEL TESTED ON ENMAP AND HYPERSCOUT-2

[Jamie Longino, Jaime Zabalza & Massimiliano Vasile]

University of Strathclyde, Glasgow, UK

ABSTRACT

Hyperspectral algorithms often train and test on simulated data, especially in cases where ground truth cannot be measured, such as for space objects. The noise in the data must behave like real sensor noise, growing with signal and represent fixed patterns, such as striping, rather than only one signal-to-noise ratio for each band. We present a sensor-agnostic noise model that derives noise from instrument parameters, considering each simulated pixel, split into fixed pattern and random parts. We test the model by adding noise to real HyperScout-2 and EnMAP images, then compare it with measured noise from EnMAP's less-processed Level-1B product. With only 12 parameters, the simulated noise matches measured to around 6% on test data and 8% on unseen data, against 11-32% for a single fixed noise level per band. Our testing also shows that smile and quantisation change between random and fixed pattern, and dual-gain readout must be modeled.

***Index Terms*—** Hyperspectral simulation, measurement uncertainty, sensor-agnostic modelling, noise estimation, EnMAP, HyperScout-2

1. INTRODUCTION

Simulated hyperspectral observations are used as measurements when large data volume is required or wherever ground truth cannot be established. Characterising resident space objects is an extreme case: the materials and attitude state of debris are unknown, so methods that identify characteristics from reflected light, spectral matching, colour indices or machine learning [1, 2, 3], have very little labelled flight data and commonly rely on laboratory spectra and simulated observations [1, 2]. For such work the simulated noise must behave like real sensor noise, realistic uncertainty-labelled errors make the training data, and therefore results obtained from them, more robust. Through this we are able to narrow the gulf between simulation and the real world that methods developed on simulated data must eventually cross.

Earth observation simulators often model sensor chains in detail, most of them for only one mission, such as end-to-end simulators like EeteS for EnMAP [4] and the CHIME observation performance simualtor, which covers its instrument and Level-0/1 processing [5]; generic instrument simulators also exist [6]. Earth observation metrology on the other hand, attaches uncertainty to each measurement, tracking how the errors correlate. This is seen in independent, structured and common effects in FIDUCEO [7], error-correlation information in CoMet toolkit [8], and also per-pixel uncertainty built from identified contributors in Sentinel-2's radiometric uncertainty tool [9]. The per-term model presented here combines both methods and is designed as sensor-agnostic. From a brief list of sensor parameters we compute, for each pixel and each observation band, the standard deviations for each noise source, we label them as *random* (changing from frame to frame, i.e. shot noise) or *fixed-pattern* (repeating in every frame, such as pixel-to-pixel gain differences).

A budget (cumulative noise) like this is only useful if it is realistic. To our knowledge, no simulator's uncertainty budget has been tested term by term against flight data, this is the primary goal of this paper. Its contributions are (i) a closure test which works by injecting (adding noise) every term physically into Level-1C (L1C) real products of two imagers of differing designs, before verifying the analytic budget, class by class, against the realised errors; (ii) a test against measured noise requiring no existing raw data: for EnMAP, the result is compared with the less-processed data of the Level-1B (L1B) product for the same looks. This allows for noise kept by resampling, on data not utilised in fitting the model, which is then checked against a signal-independent baseline; and (iii) the corrections these tests highlight for sensor-agnostic budgets: two terms whose classification depends on scan architecture or signal level, and additionally a readout-mode descriptor which the model initially lacked.

2. SENSOR-AGNOSTIC UNCERTAINTY MODEL

Nothing in the uncertainty model is specific to a single instrument, a sensor generates error data only through descriptors of its optics, detector, readout and scan, taken from specifications or fitted to data. For each pixel and band, we convert the radiance into photoelectrons, calculate the standard deviation of each noise source from the instrument parameters, and add the random terms in quadrature into σ_T and the fixed-pattern terms into σ_S. The total uncertainty is $\sigma_C = (\sigma_T^2 + \sigma_S^2)^{1/2}$, expressed in radiance units. A radiance L in band b produces $N_e = R_b L$ electrons, with the responsivity $R_b = \frac{\pi\tau}{1+4F^2} A_{\mathrm{pix}}\, t\, \mathrm{QE}(\lambda_b)\, \Delta\lambda_b\, \frac{\lambda_b}{hc}$ set by the optics transmission τ, f-number F, pixel area A_{pix}, integration time t, quantum efficiency and band width $\Delta\lambda_b$. The detector terms follow the work of Konnik and Welsh [10]:

$$\sigma_T^2 = N_e + N_{\mathrm{dark}} + \sigma_{\mathrm{read}}^2 + \sigma_{\mathrm{kTC}}^2 + g^2/12, \tag{1}$$

$$\sigma_S^2 = (pN_e)^2 + (dN_{\mathrm{dark}})^2 + \sigma_{\mathrm{off}}^2, \tag{2}$$

in electrons, with N_{dark} the dark-signal electrons: photon and dark shot noise, read noise, reset noise (taken as zero under correlated double sampling, which causes it to be removed in CCDs but not always in CMOS sensors [10]; any residual is a component of the remaining effective read noise) and quantisation with gain g (e^-/DN) are temporal terms; photo-response non-uniformity (PRNU) p, dark-signal non-uniformity (DSNU) d and offset fixed-pattern noise are systematic. Aberration and atmospheric terms are tagged in the same way (Table 1). Spectral smile as $\sigma = |\partial L/\partial\lambda|\, s\, \delta\lambda$, with s the smile in units of channel spacing $\delta\lambda$, here it is tagged temporal, and for a target that moves across the field between frames; keystone as

Table 1. Terms of the uncertainty budget: the class given (T temporal, S systematic), how the verification method injects it, and the class of the realised error (Section 4). Bold: the tag is wrong, [a] when other temporal noise is below one ADC step, [b] in both push-broom and push-frame imagers.

Term	Injected as	Tagged	Realised
Photon, dark shot	Poisson draw	T	T
Read noise	Gaussian draw	T	T
Quantisation	ADC rounding	T	T or **S**[a]
PRNU, DSNU, offset	pattern per column	S	S
Smile	λ shift $\propto u^2$	T	**S**[b]
Keystone	spatial shift $\propto u$	S	S
Scintillation	log-normal draw	T	T
Extinction error	one draw per scene	S	S

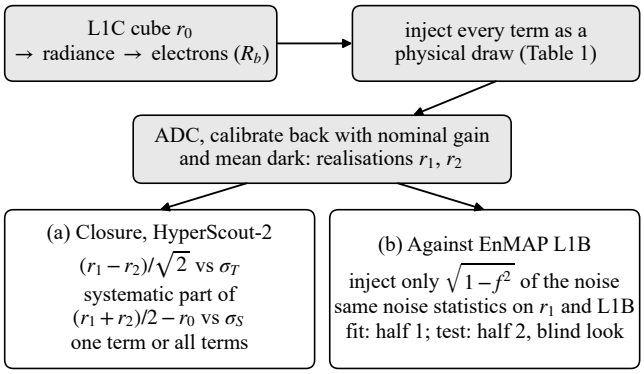


Fig. 1. Verification method. The analytic budget is tested against the realised errors (a), for EnMAP the synthetic noise against the noise measured on Level-1B (b).

$\sigma = L\,k\,f_g$, with k being the keystone in pixels and f_g as the fractional scene gradient, and is tagged as systematic. For through atmosphere observations, scintillation is temporal and the knowledge errors contributed by Rayleigh, aerosol and ozone extinction are systematic terms. For the Earth views tested here they are evaluated as being the two-way airmass of the solar and viewing zenith angles. Where a given instrument's values are unknown, the budget uses the following defaults: $s = k = 0.15$, $f_g = 0.1$ and, when used for a ground observer, a scintillation of 0.01 mag. Every simulated sample (a spectral response wavelength, or a pixel and band of a cube) is delivered with σ_T, σ_S and $\sigma_C = (\sigma_T^2 + \sigma_S^2)^{1/2}$.

The verification work created two additions for the model. For imagers with a fixed detector (push-broom or push-frame), smile and keystone are applied *deterministically* as a wavelength shift $\propto u^2$ and a band-dependent cross-track shift $\propto u(\lambda_b - \bar{\lambda})$ in the field coordinate u (± 1 at field edges), scaled to the same full-field RMS as the variance terms. Additionally, prompted by the sensor gain characteristics of the EnMAP sensor, a readout-mode descriptor is added, detectors that switch between high- or low-gain readout have a σ_{read} and g that depend on whether the value of N_e is below a switch level.

3. VERIFICATION METHOD

3.1. Physical injection into Level-1C data

Fig. 1 outlines our method. Each L1C cube is converted into radiance and then into electrons with R_b. Each term is then drawn as a physical realisation (Table 1): Poisson shot noise, Gaussian read noise, a log-normal scintillation factor, rounding and clipping in the analogue-to-digital converter (ADC), fixed patterns per detector column (along-track stripes on the ground), the deterministic aberrations, and one extinction error per scene. The result is calibrated back along with the nominal gain and mean dark only. Two realisations share fixed patterns but not the temporal draws, so $(r_1 - r_2)/\sqrt{2}$ is a purely temporal product and $(r_1 + r_2)/2 - r_0$ is systematic plus half the temporal variance, which is subtracted. Both are compared with σ_T and σ_S evaluated via the clean signal r_0, as z-scores and coverage.

3.2. Comparison with a less-processed product

The EnMAP data we have used distributes each look as L1B (calibrated radiance in sensor geometry) and L1C (L1B orthorectified; and ours includes bilinear resampling). Both levels already contain EnMAP's real smile and keystone, so these are not injected. Resampling keeps some of the noise: L1C retains 72–73% of the L1B noise in the scenes we used, which is close to the $2/3$ expected for white noise, so only $\sqrt{1-(2/3)^2}$ of our modelled noise is injected. The generated synthetic cube and L1B data are then measured with identical statistics. Noise is estimated by spectral–spatial decorrelation [11, 12], regressing each band against its two neighbouring bands and on the pixel to its left (first and last bands are not estimated). We use 16×16-pixel blocks and then take the 25th percentile of the block residual scatter as signal bins, which gives signal-dependent noise curves [13, 14, 15]. Additionally we measure the adjacent-band correlation of the residuals as well as the along-track stripe amplitude.

L1B is not raw. Its processor already applies non-linearity, dark and gain-matching corrections, flat-field, straylight correction (this mixes neighbouring pixels and bands), defective-pixel interpolation and image-based destriping [16]. However it appears that its optional smile correction was not applied to our products. The comparison hereafter therefore tests the temporal noise model. PRNU and stripes are residuals left after the processing. The correlation between adjacent bands, which the straylight correction most likely introduces, is added phenomenologically by use of a three-tap spectral kernel.

The parameters of the model (read-noise configuration, residual PRNU per look, adjacent-band correlation and responsivity) are fit against the noise curves of the first along-track half of each look and then validated on the second half. A blind test fits against two looks and validates on the completely held half of the third.

3.3. Data

The two imagers differ in spectral separation, scan architecture and optics. HyperScout-2 on Φ-Sat-1 [17], is a push-frame imager using a linear variable filter: 15 L1C scenes (50 bands, 450–940 nm, 75 m) of varied land cover; its detector and optics parameters come from available public specifications, and its raw data are not public. EnMAP [18], is a prism push-broom spectrometer whose visible and near-infrared (VNIR) detector has a dual-gain readout: 91 bands (418–993 nm) across our three looks with L1B and L1C, Tyrol

(Alps, 2022-06-12, datatake 1049, tile 28), Nevada (desert, 2023-07-07, datatake 25905, tile 1) and Groningen (farmland with 46% water, 2022-08-10, datatake 2446, tile 2), all from processor version 01.05.06; we do not test the short-wave infrared detector.

4. RESULTS

4.1. Closure and tagging

Injected one by one into HyperScout-2 scene 02D1, the six terms used reproduce their analytic σ to within 1.5% (photon shot, read noise, PRNU and scintillation 1.000, offset pattern 1.015, quantisation 0.987). Quantisation matches in magnitude but not in class, with no other noise, both realisations round identically, so the error has no temporal part at all (Section 4.3). Stripe amplitudes match their prediction to 1.01–1.13 and the atmospheric terms draw to within 0.97–1.03 in all of the 15 scenes. With every term on, the budget slightly over-predicts, the combined z-scores have a standard deviation of 0.75–0.86 and 76–82% of errors fall within 1σ, against 68% for a Gaussian.

The aberration terms are responsible for this(Fig. 2). With smile and keystone switched off, the other terms together close in scene 02D1 (z-score standard deviation 1.00, 68.2% within 1σ). This is due to smile being tagged temporal, when in a push-frame or push-broom imager it is instead fixed per detector column, causing it to repeat in every frame. With the existing budgets tags the temporal closure drops to 0.87–0.94 (and to 0.54 at 450 nm in 02D1) wherever the spectrum is very steep. Moving smile across to the systematic budget brings it to 1.000 in every scene. The over-prediction itself appears due to the variance forms. With the field dependence assumed (smile $\propto u^2$ and zero at field centre, keystone $\propto u$) and the 120-km tiles focused on the detector, covering the central 40%. The deterministic errors are 15–19% (smile) and 9–48% (keystone) of the budget's σ; the measured fractional scene gradient, 0.014–0.091, is also underneath the default of 0.1. These ratios are however, conditional, at the field edge the smile ratio would reach 1.1 and ratio of keystone would grow by 2.7-fold. For a smile defined relative to the field-mean wavelength, like the measured smile in EnMAP's data, a centred tile sees 0.4–0.95 of the budget's σ. The default smile itself (0.15 channel spacings) is instead a specification-level value (EnMAP specifies $<$ 0.2 pixel [19]), 7–11 times EnMAP's measured VNIR smile. A single σ per term cannot describe an error that depends on field position, shape and instrument; it has to instead be computed at each field position.

4.2. Validation against EnMAP L1B

Table 2 and Fig. 3 compare three noise models on data not used in fitting. All inject noise into L1C in the same way and differ only in the noise model:

- *Per-term model, 96 parameters*: the budget of Section 2 with dual-gain readout, a free responsivity scale per band (the per-band freedom of a given empirical noise fit), residual PRNU per look and one adjacent-band correlation;
- *Per-term model, 12 parameters*: the same as the 96 parameter model, with the logarithm of the responsivity scale constrained to a degree-4 polynomial in wavelength;
- *SNR knob, 89 parameters*: one signal-independent σ per band, as a signal-to-noise ratio (SNR) specified at only one reference signal would give, fitted the same way.

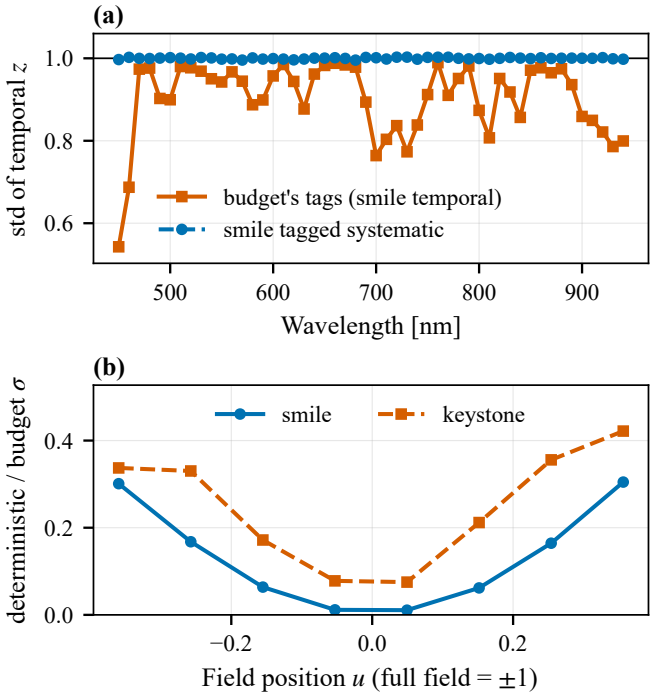


Fig. 2. HyperScout-2 scene 02D1. (a) Standard deviation of temporal z-score per band: with the budget's tags (smile temporal) and with smile tagged as systematic. (b) RMS of the deterministic smile and keystone errors relative to the variance terms of the budgets, against across-track field position.

With just 12 parameters, our per-term model reproduces the noise in L1B on the held-out data to around 6% RMS, across all bands/ signal levels, and to 8% on a look it wasn't fitted against. Replacing 89 free responsivity scales by a smooth curve causes little change, so the result is not due to per-band freedom. The SNR knob is fitted to each band's noise level, but cannot follow the dependence on signal (slope -0.19 to -0.38). It over-predicts dark pixels by up to 49% and under-predicts bright ones by about 18%, its error reaches as high as 32% on the blind look. This is also true for the part of the noise that L1C processing removed and the model must supply, the injected excess variance is 1.03–1.08 of the removed variance for our 12-parameter uncertainty model and 0.60–1.50 for the SNR knob. Adjacent-band correlation does not separate the results, due to the synthetic cubes largely inheriting it from the L1C input, and neither does along-track stripes, their synthetic amplitude is 0.91–1.15 of L1B's for all three models. Inputting the measured L1C retention of 0.72 instead of $2/3$ changes the error by at most 0.3 on the held-out halves and 0.7 when applied to the blind look.

With publicly available EnMAP parameters and the budget's defaults the synthetic noise was 1.9 times that of L1B (Tyrol. A single read noise struggled to fit bright scenes, as the noise barely changes with signal, together with the shot-limited Alps scene. However, adding dual-gain readout reduced the residual by 14%. Refitting on a finer grid (steps of $10\,\mathrm{e}^-$ in high gain and $25\,\mathrm{e}^-$ in low gain), with a block bootstrap over bands, the three-look fit with free responsivities puts the switch point at $10^5\,\mathrm{e}^-$, the high-gain read noise at $50\,\mathrm{e}^-$ (68% interval 40–70 e^-) and the low-gain read noise at $325\,\mathrm{e}^-$ (300–350 e^-). The specified system read noise, including correlated double sampling, is 70 and $200\,\mathrm{e}^-$ [20]; ours are effective values

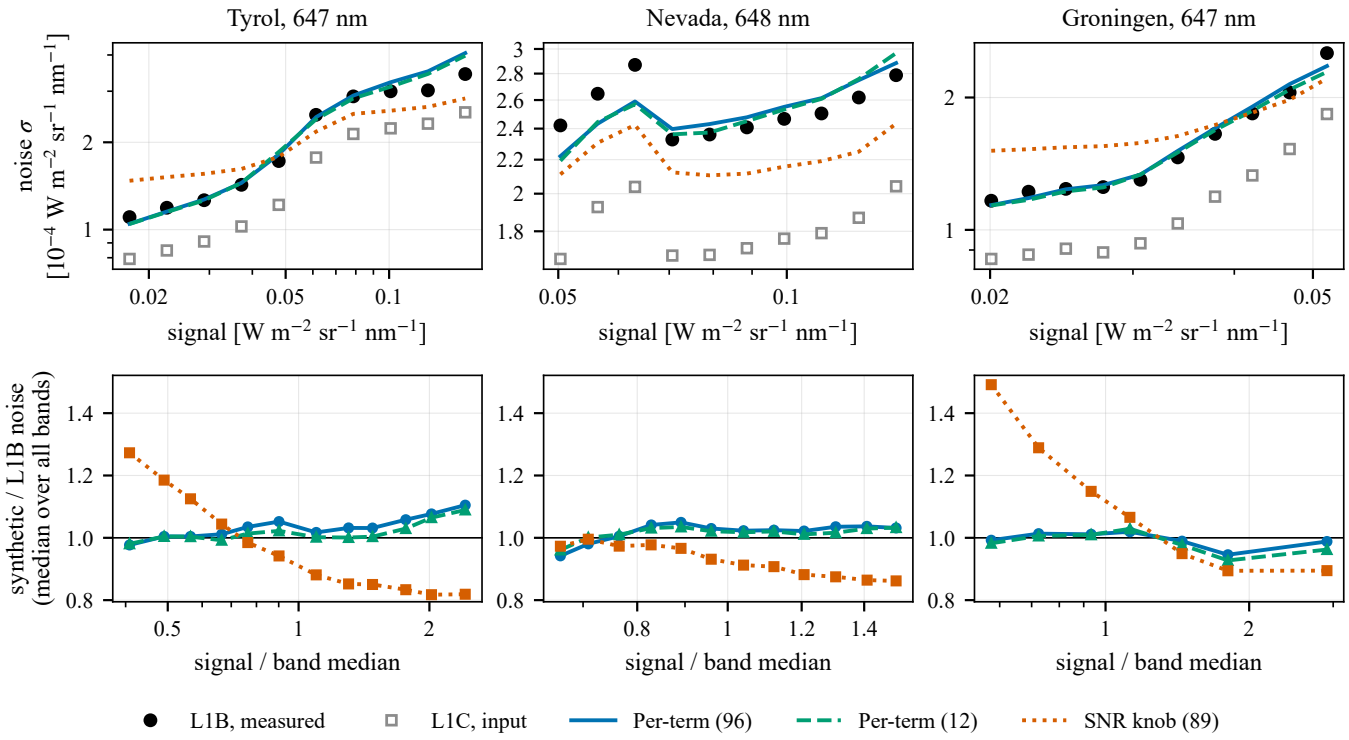


Fig. 3. EnMAP, held-out halves. Top: noise against signal in the band nearest 647 nm, measured on L1B, on the L1C input, and on the synthetic cubes of the three noise models (parameter counts in brackets). Bottom: synthetic/L1B noise against signal relative to each band's median signal, as the median over all bands.

Table 2. EnMAP validation. Fit: mean squared log residual on the first halves. Error: $100\times$ RMS of $\ln(\sigma_{\text{synthetic}}/\sigma_{\text{L1B}})$ over all bands and signal bins on the held halves (Tyrol / Nevada / Groningen) and on Groningen's second half when Groningen is left out of the fit (blind). Slope: median over bands of $d\ln(\sigma_{\text{synthetic}}/\sigma_{\text{L1B}})/d\ln L$ on the held halves. Errors and slopes are means over three temporal random seeds; the seed-to-seed standard deviation of the error is at most 0.12.

Model (param.)	Fit	Error, held out	Blind	Slope
Per-term (96)	0.023	6.9 / 6.4 / 5.8	7.1	$+0.02$ to $+0.05$
Per-term (12)	0.029	6.4 / 5.5 / 6.5	7.7	$+0.01$ to $+0.04$
SNR knob (89)	0.155	14.7 / 10.7 / 23.6	32.1	-0.38 to -0.19

after gain matching and the straylight correction. Without Groningen's dark water the high-gain value barely gets constrained, being over 0–140 e^- the two-look cost varies by less than 4%. The same terms thus serve both imagers, aside from the dual-gain readout and the adjacent-band correlation that EnMAP's processing most likely adds, only specific values differ.

4.3. Tagging depends on the regime

Quantisation is tagged temporal ($g/\sqrt{12}$), which assumes our other temporal noise dithers the signal across the ADC steps. This is not always true. Fitted to Groningen alone, the model settles on zero read noise and a single-gain readout with a step of 122-e^-. Over water at bands 900–1000 nm the shot noise becomes 0.09–0.29 of a step, and the rounding error repeats itself. The systematic error then exceeds its predicted value by over 4-fold at 0.29 of a step and by 20-fold at 0.09 (the predictions there are only about 4 and 1 e^-; where the ratios do not change on different random seeds), while the ratio stays between 0.9 and 1.2 in every band until 890 nm. This is because of the model rather than EnMAP, whose read noise dithers its quantisation, but it applies to any low-noise detector with coarse quantisation viewing dark targets. Like with smile, quantisation has to be tagged by regime and not just term.

4.4. Scope and limitations

Two imagers with silicon detectors do not establish complete sensor independence: whisk-broom and snapshot architectures, as well as infrared detectors remain untested. The EnMAP comparison tests the temporal noise of the VNIR detector. Where L1B has already been flat-fielded and destriped, so the systematic terms are checked only by the closure test, the fitted read noise, PRNU and adjacent-band correlation are effective values of the processed product. Of the three looks, only Groningen's dark water constrained the high-gain read noise. The HyperScout-2 closure depends upon publicly available detector and optics parameters, and the size of its aberration ratios on the unknown tile position. Space-object observation adds regimes not tested here, unresolved targets viewed through turbulent atmosphere and above all, it shares the low-signal regime of Section 4.3.

5. CONCLUSION

We tested a sensor-agnostic, hyperspectral uncertainty model against flight data from two hyperspectral imagers of differing designs. Each noise source tested alone reproduces its predicted analytic size. On EnMAP the model reproduces the signal-dependent Level-1B noise to around 6% RMS on test data, and 8% on unseen scenes with 12 parameters, where a model with one fixed noise level per-band errs by as much as 32%. The per-term errors thus narrow the gap between simulated and real noise by two- to four-fold in this case. Which should make algorithms trained on simualted data more robust. The tests also located three corrections: smile is fixed-pattern when used on a push-frame or push-broom imager, ADC quantisation becomes systematic when all other temporal noise is under a single step, and dual-gain readout must be modelled. For our sensor-agnostic budget, the tags must therefore be derived directly from the instrument's details and the regime of the signal rather than fixed per term; faint space objects seen by a low-noise detector exist in the regime where quantisation changes class. The test against Level-1B needs no raw data and can be applied to any available mission data that distributes both a sensor-geometry and an orthorectified product. Next we will test the model on raw data, such as on HYPSO-1's or Level-1A, and recover injected smile and keystone by using a scene-based estimator, where the image-based counterpart of laboratory characterisation [21] is scored with co-registration metrics [22, 23].

6. REFERENCES

[1] M. Vasile, L. Walker, A. Campbell, S. Marto, P. Murray, S. Marshall, and V. Savitski, "Space object identification and classification from hyperspectral material analysis," *Scientific Reports*, vol. 14, 2024, Art. no. 1570.

[2] J. A. Reyes, H. M. Cowardin, and M. Velez-Reyes, "Analysis of spacecraft materials discrimination using color indices for remote sensing for space situational awareness," *The Journal of the Astronautical Sciences*, vol. 70, no. 5, 2023, Art. no. 33.

[3] K. De Alba, E. De Alba, Z. Gazak, and J. Fletcher, "M2S2: Multispectral dataset for material segmentation of satellites," in *Proc. 15th Workshop on Hyperspectral Image and Signal Processing (WHISPERS)*, Bellaterra, Spain, 2025, pp. 1–5.

[4] K. Segl, L. Guanter, C. Rogass, T. Kuester, S. Roessner, H. Kaufmann, B. Sang, V. Mogulsky, and S. Hofer, "EeteS—the EnMAP end-to-end simulation tool," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 5, no. 2, pp. 522–530, 2012.

[5] N. Lamquin et al., "The CHIME observation performance simulator (OPSI) software system: development and status at critical design review," in *Proc. 14th Workshop on Hyperspectral Image and Signal Processing (WHISPERS)*, Helsinki, Finland, 2024.

[6] B. Borguet and V. Moreau, "A versatile hyperspectral instrument simulator, application to the CALHYS spectrometer," in *Proc. 14th Workshop on Hyperspectral Image and Signal Processing (WHISPERS)*, Helsinki, Finland, 2024.

[7] J. Mittaz, C. J. Merchant, and E. R. Woolliams, "Applying principles of metrology to historical Earth observations from satellites," *Metrologia*, vol. 56, no. 3, 2019, Art. no. 032002.

[8] National Physical Laboratory, "CoMet toolkit (punpy and obsarray)," `https://www.comet-toolkit.org`, 2026, Accessed 22 Sep. 2026.

[9] J. Gorroño, N. Fomferra, M. Peters, F. Gascon, C. I. Underwood, N. P. Fox, G. Kirches, and C. Brockmann, "A radiometric uncertainty tool for the Sentinel 2 mission," *Remote Sensing*, vol. 9, no. 2, 2017, Art. no. 178.

[10] M. Konnik and J. Welsh, "High-level numerical simulations of noise in CCD and CMOS photosensors: review and tutorial," arXiv:1412.4031, 2014.

[11] R. E. Roger and J. F. Arnold, "Reliably estimating the noise in AVIRIS hyperspectral images," *Int. J. Remote Sens.*, vol. 17, no. 10, pp. 1951–1962, 1996.

[12] L. Gao, Q. Du, B. Zhang, W. Yang, and Y. Wu, "A comparative study on linear regression-based noise estimation for hyperspectral imagery," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 6, no. 2, pp. 488–498, 2013.

[13] L. Alparone, M. Selva, B. Aiazzi, S. Baronti, F. Butera, and L. Chiarantini, "Signal-dependent noise modelling and estimation of new-generation imaging spectrometers," in *Proc. 1st Workshop on Hyperspectral Image and Signal Processing (WHISPERS)*, Grenoble, France, 2009, pp. 1–4.

[14] N. Acito, M. Diani, and G. Corsini, "Signal-dependent noise modeling and model parameter estimation in hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 8, pp. 2957–2971, 2011.

[15] J. Meola, M. T. Eismann, R. L. Moses, and J. N. Ash, "Modeling and estimation of signal-dependent noise in hyperspectral imagery," *Applied Optics*, vol. 50, no. 21, pp. 3829–3846, 2011.

[16] DLR EnMAP Ground Segment, "Level 1B processor (systematic and radiometric correction) ATBD," Tech. Rep. EN-PCV-TN-4006, issue 1.11, DLR, 2026.

[17] M. Esposito, S. S. Conticello, M. Pastena, and B. Carnicero Domínguez, "In-orbit demonstration of artificial intelligence applied to hyperspectral and thermal sensing from space," in *Proc. SPIE 11131, CubeSats and SmallSats for Remote Sensing III*, 2019, Paper 111310C.

[18] T. Storch et al., "The EnMAP imaging spectroscopy mission towards operations," *Remote Sens. Environ.*, vol. 294, 2023, Art. no. 113632.

[19] DLR, "EnMAP HSI instrument specification," `https://www.enmap.org/data/doc/EnMAP_Specs.pdf`, 2026, Accessed 23 Sep. 2026.

[20] eoPortal, "EnMAP (environmental mapping and analysis program)," `https://www.eoportal.org/satellite-missions/enmap`, 2026, Accessed 22 Sep. 2026.

[21] W. H. Bakker, H. M. A. van der Werff, and F. D. van der Meer, "Determining smile and keystone of lab hyperspectral line cameras," in *Proc. 10th Workshop on Hyperspectral Image and Signal Processing (WHISPERS)*, Amsterdam, The Netherlands, 2019, pp. 1–5.

[22] T. Skauli, "An upper-bound metric for characterizing spectral and spatial coregistration errors in spectral imaging," *Optics Express*, vol. 20, no. 2, pp. 918–933, 2012.

[23] IEEE, "IEEE standard for characterization of hyperspectral imaging devices—ultraviolet through shortwave infrared," IEEE Std 4001-2025, 2025.