ATLAS: Hyperspectral Image Compression via Adaptive Transfer of Low-rank Abundance Splatting

Jacob Rempel and Faisal Z. Qureshi
Faculty of Science
Ontario Tech University
2000 Simcoe St. N., Oshawa ON L1G 0C5, Canada
jacob.rempel@ontariotechu.net and faisal.qureshi@ontariotechu.ca

Abstract

LoR-SGS showed that 2D Gaussian splatting can compress hyperspectral images (HSI) by factorizing the cube into shared endmember spectra and low-rank abundance maps, so each Gaussian stores K abundance coefficients instead of λ bands. However, it does not employ densification, learned quantization, or training strategies that RGB splatting codecs use. We conduct a transfer study of nine such techniques in the HSI domain. Five transfer: LSQ+ quantization of abundances and covariances, distortion-driven densification, two-stage quantization-aware training, learnable endmembers, and gradient-guided initialization. Four fail: content-aware filters, position quantization, progressive quantization, and spatial delta coding; the failures are consistent with HSI abundance maps being spatially smooth, low-rank fields in which quantization error is amplified multiplicatively through spectral band coupling. The optimized pipeline, ATLAS, beats published LoR-SGS on 6 of 8 rate-distortion targets ($+1.20$ dB average PSNR), and a scaling study on scenes up to $2372 \times 2196 \times 128$ reaches 40.1 dB at 0.034 bits per pixel per band on a single GPU without tiling. Decoded cubes match or exceed original-cube classification accuracy at 106:1 to 280:1, and classifying the bitstream's abundance maps directly matches full-spectrum accuracy without reconstructing the cube.

1 Introduction

Hyperspectral sensors capture hundreds of contiguous spectral bands per pixel, but a single data cube holds hundreds of millions of values, a storage and transmission bottleneck for Earth observation missions [Rezasoltani and Qureshi, 2023]; downlink capacity is scarce enough that dedicated onboard HSI compression standards exist [Hernández-Cabronero et al., 2021]. Classical codecs such as JPEG2000 [Skodras et al., 2001] and PCA-DCT [Nian et al., 2016] underuse nonlinear inter-band correlation, tensor factorizations [Karami et al., 2012, Du et al., 2017] are restricted to linear mixtures, and implicit neural representations [Rezasoltani and Qureshi, 2023, Shi et al., 2024] trade their rate-distortion gains with slow per-image optimization and hard-to-compress MLP weights.

In the RGB domain, 2D Gaussian splatting [Zhang et al., 2024], derived from 3D Gaussian scene representations [Kerbl et al., 2023], represents an image as explicit optimizable Gaussians and decodes at over 1000 FPS. GaussianImage++ [Li et al., 2026] added distortion-driven densification, content-aware filters, and learned step size quantization (LSQ+) [Bhalgat et al., 2020], setting the state of the art for RGB splatting codecs. Applied naively to HSI, the approach collapses: with $\lambda \approx 200$ bands each Gaussian would carry 200 color coefficients. LoR-SGS [Li et al., 2025] resolved this by exploiting the low-rank structure of HSI: nonnegative matrix factorization (NMF) [Lee and Seung, 1999] splits the cube into K endmember spectra and abundance maps ($K \ll \lambda$), so each

40th Conference on Neural Information Processing Systems (NeurIPS 2026). Workshop: 2nd Workshop on Advances in Representation Learning for Earth Observation.

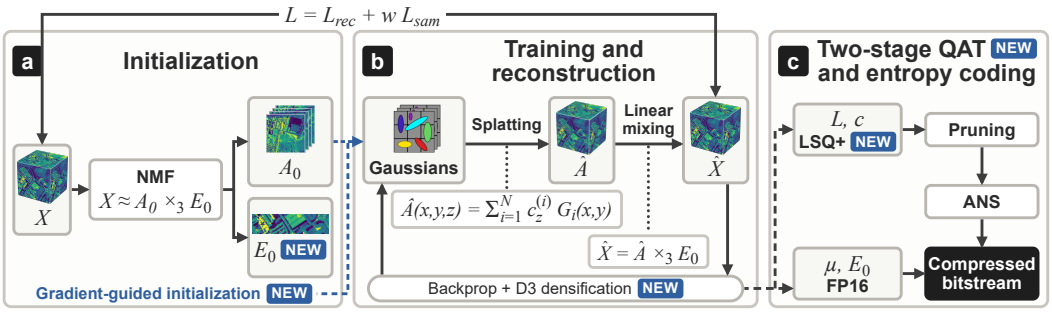


Figure 1: ATLAS pipeline, NEW marks transferred techniques. **(a)** NMF factors the cube $\mathbf{X} \approx \mathbf{A}_0 \times_3 \mathbf{E}_0$ into abundance maps $\mathbf{A}_0$ and (now learnable) endmembers $\mathbf{E}_0$. Gradient-guided initialization places the Gaussians from $\mathbf{A}_0$. **(b)** Each Gaussian has a position $\boldsymbol{\mu}$, Cholesky factors $\mathbf{L}$, and abundances $\mathbf{c}$. Splatting sums the Gaussian kernels G_i, weighted by their abundance coefficients $c_z^{(i)}$, into the abundance field $\hat{\mathbf{A}}$ (Eq. 1). Linear mixing with the endmembers gives $\hat{\mathbf{X}} = \hat{\mathbf{A}} \times_3 \mathbf{E}_0$. The loss $L = L_{rec} + w\,L_{sam}$ compares $\hat{\mathbf{X}}$ with $\mathbf{X}$ (reconstruction plus SAM, weight w); backpropagation updates all parameters, and D3 densification grows $N/2 \to N$ at high-error pixels. **(c)** Two-stage QAT: after FP32 warmup, LSQ+ quantizes $\mathbf{L}$ (8-bit signed) and $\mathbf{c}$ (6-bit unsigned), degenerate Gaussians are pruned, and both streams are ANS-coded; $\boldsymbol{\mu}$ and $\mathbf{E}_0$ stay FP16. The four streams form the `.atlas` bitstream.

Gaussian stores $K+5$ parameters regardless of λ. LoR-SGS still relies on residual vector quantization (RVQ) [Juang and Gray, 1982] and none of the RGB-specific training advances.

This paper asks *which* RGB-domain techniques transfer to the structurally different HSI setting, and which break down. We evaluate nine; five transfer and four fail, and the failures expose structural properties of HSI splatting that RGB experiments cannot reveal. The optimized pipeline, which we call ATLAS, beats published LoR-SGS on 6 of 8 rate-distortion targets (average $+1.20$ dB PSNR), and a scaling study on three larger scenes, including the 5.2-megapixel Chikusei cube, shows the Gaussian count grows sublinearly with pixel count (Section 4). Extended related work is in Appendix A.

2 Preliminaries: LoR-SGS

LoR-SGS [Li et al., 2025] compresses a cube $\mathbf{X} \in \mathbb{R}^{h\times w\times \lambda}$ in three stages. NMF first decomposes it, $\min_{\mathbf{A}_0,\mathbf{E}_0\geq 0} \|\mathbf{X} - \mathbf{A}_0\mathbf{E}_0\|_F^2$, into abundance maps $\mathbf{A}_0 \in \mathbb{R}^{hw\times K}$ and an endmember matrix $\mathbf{E}_0 \in \mathbb{R}^{K\times\lambda}$ with $K \ll \lambda$ (typically 10–12); nonnegativity keeps both physically interpretable as mixing proportions and reflectance spectra [Bioucas-Dias et al., 2012]. The abundances are then represented by N 2D Gaussians, each parameterized by position $\boldsymbol{\mu} \in \mathbb{R}^2$, Cholesky factors $\{l_1, l_2, l_3\}$ of the covariance $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^\top$, and an abundance vector $\mathbf{c} \in \mathbb{R}^K$:

$$\hat{A}(x,y,z) = \sum_{i=1}^{N} c_z^{(i)}\, \mathcal{G}_i(x,y), \qquad \mathcal{G}_i(x,y) = \exp\!\left(-\tfrac{1}{2}\mathbf{d}^\top \boldsymbol{\Sigma}_i^{-1}\mathbf{d}\right), \tag{1}$$

with $\mathbf{d} = (x,y)^\top - \boldsymbol{\mu}_i$. The cube is reconstructed as $\hat{X}(x,y,b) = \sum_z \hat{A}(x,y,z)\,E_0(z,b)$, written $\hat{\mathbf{X}} = \hat{\mathbf{A}} \times_3 \mathbf{E}_0$, and optimized with the Frobenius reconstruction loss, so each Gaussian carries $K+5$ parameters. For compression, LoR-SGS stores positions and $\mathbf{E}_0$ at FP16, Cholesky parameters with asymmetric uniform quantization, and abundances with RVQ; the covariance and RVQ index streams are entropy-coded with asymmetric numeral systems. We use a unified kernel rasterizing all K channels in one pass instead of the original grouped multipass scheme; outputs are bit-identical and training is $1.35\times$ faster on average, with no novelty claimed (Appendix E).

3 Transfer Study

We evaluate nine techniques, five from GaussianImage++ [Li et al., 2026], two HSI-specific extensions, and two general compression methods, for transfer to HSI.

3.1 What Transfers

LSQ+ for abundances and covariances. We replace RVQ for abundances and uniform covariance quantization with attribute-separated LSQ+ [Bhalgat et al., 2020]: 6-bit unsigned with per-channel scale and offset for abundances, 8-bit *signed* for Cholesky factors. The sign is critical: raw l_1, l_2, l_3 may be negative before the positive semi-definite bound is added; unsigned quantization clips them.

Distortion-driven densification (D3). We start with $N/2$ Gaussians and grow to N by placing Gaussians at the highest-error pixels (spectral L_1) [Li et al., 2026], initializing abundances from the rendered map. Fewer, larger steps win: 2 steps by iteration 2000 beats 5 steps to iteration 5000 by $\sim$0.2 dB on PaviaU, likely because late Gaussians have fewer iterations to converge.

Two-stage quantization-aware training (QAT). Training with LSQ+ enabled from the start degrades quality badly. We instead warm up at FP32 with D3, then load the best checkpoint, enable the quantizers with step size and offset initialized from the current parameter distributions, and rebuild the optimizer. This loses only 0.88 dB on average versus the FP32 ceiling at 3–4$\times$ lower bitrate, and gains 2.26 dB on PaviaU over quantization without warmup.

Learnable endmembers. LoR-SGS freezes $\mathbf{E}_0$ after NMF. We make it learnable with a separate Adam optimizer at learning rate 10^{-4} ($50\times$ below the main rate), keeping nonnegativity via softplus reparameterization. Since $\mathbf{E}_0$ is stored at FP16 either way, this adds *zero* bitrate and $+0.34$ to $+1.68$ dB on PaviaU, Salinas, and Urban. The technique has no RGB equivalent.

Gradient-guided initialization. Blending Sobel gradient magnitudes of the NMF abundances with NMF reconstruction error biases initial placement toward spectral transitions. Gains over error-driven placement alone are modest ($+0.08$ dB on PaviaU, $+0.19$ dB on JasperRidge) but come at no cost.

3.2 What Fails, and Why

Table 4 (Appendix C) quantifies four failures, which we attribute to structural differences between RGB images and low-rank abundance fields. **Content-aware filters (CAF)** set per-Gaussian low-pass bounds $s_i = hw/(\alpha N_t)$ [Li et al., 2026], with $\alpha = 9\pi$ in the reference code; late-added D3 Gaussians receive $s_i \approx 0.28$, well below the fixed 0.5 of LoR-SGS, and the narrow Gaussians represent the spatially smooth rank-K mixing weights poorly (-0.68 to -1.72 dB across all four datasets, see Appendix C). **Position quantization** at 12 bits, effective for RGB, costs -3.29 dB versus FP16: a position error propagates through K abundance channels and then λ bands via $\mathbf{E}_0$, a multiplicative amplification absent when an error touches 3 color channels. **Progressive quantization**, ramping the quantization step size from $2.0\times$ to $1.0\times$, destabilizes training (PaviaU crashes to 23 dB): LSQ+ learns step size and offset through straight-through gradients [Esser et al., 2020], and overriding the step size mid-learning corrupts both. **Spatial delta coding** of Morton-ordered quantized indices *increases* bitrate by 14–47%: per-channel delta streams each pay fixed header and histogram overhead exceeding the entropy saved.

4 Experiments

Setup. We evaluate on four standard HSI compression benchmarks: PaviaU ($340\times340\times103$), Salinas ($180\times180\times204$), JasperRidge ($100\times100\times198$), and Urban ($307\times307\times162$). Metrics are PSNR, MS-SSIM [Wang et al., 2003], spectral angle mapper (SAM) [Yuhas et al., 1992], and bits per pixel per band (bpppb). The final configuration combines all five transferred techniques with SAM loss weight 0.05, pruning of degenerate Gaussians during QAT, and CAF disabled (details in Appendix B); the baseline is published LoR-SGS (Table II of Li et al. [2025]) at matched rates.

Main results. Table 1 shows the optimized pipeline beats published LoR-SGS on 6 of 8 targets, averaging $+1.20$ dB PSNR over 5 independent runs (std ≤ 0.17 dB, Appendix B); PaviaU@0.10 trails by 0.09 dB and Urban@0.05 by 0.04 dB. The largest gains are on Salinas ($+1.70$, $+2.20$ dB) and JasperRidge ($+0.91$, $+3.77$ dB); Appendix D gives a cumulative FP32 ablation. A rate-matched rerun of the released LoR-SGS code trails its published numbers by 0.6–3.4 dB (Appendix B).

Table 1: Comparison on four HSI datasets at two bitrate targets each. LoR-SGS rows reproduce Table II of Li et al. [2025]; bold marks the best PSNR per target; ours wins 6 of 8. HINER did not report Jasper Ridge quality metrics. Size in KB. Full table with all baselines in Appendix F.

Method	**Pavia University**					**Salinas**				
	bpppb	Size	PSNR	MS-SSIM	SAM	bpppb	Size	PSNR	MS-SSIM	SAM
JPEG2000	0.10	145	27.11	0.8790	0.1398	0.10	81	28.36	0.8804	0.0541
INR	0.10	145	33.15	0.9760	0.0802	0.10	81	37.38	0.9822	0.0257
HINER	0.10	145	32.20	0.9532	0.2582	0.10	81	31.35	0.9398	0.0609
	0.06	87	15.55	0.2040	0.6934	0.05	40	28.98	0.8975	0.0597
LoR-SGS	0.10	145	**38.95**	0.9926	0.0564	0.10	81	41.41	0.9916	0.0135
	0.06	87	35.76	0.9881	0.0697	0.05	40	37.83	0.9869	0.0170
Ours	0.096	140	38.86	0.9922	0.0493	0.098	79	**43.11**	0.9946	0.0118
	0.058	84	**36.48**	0.9866	0.0586	0.049	40	**40.03**	0.9884	0.0141
Method	**Jasper Ridge**					**Urban**				
	bpppb	Size	PSNR	MS-SSIM	SAM	bpppb	Size	PSNR	MS-SSIM	SAM
JPEG2000	0.10	24	18.71	0.5817	0.2800	0.10	186	23.70	0.8314	0.1616
INR	0.10	24	31.34	0.9818	0.0706	0.10	186	31.20	0.9733	0.0862
HINER	0.10	24	—	—	—	—	—	—	—	—
	0.06	15	—	—	—	0.05	93	24.66	0.8751	1.5708
LoR-SGS	0.10	24	37.19	0.9936	0.0380	0.10	186	38.22	0.9940	0.0468
	0.06	15	30.79	0.9804	0.0598	0.05	93	**34.24**	0.9883	0.0652
Ours	0.103	25	**38.10**	0.9949	0.0341	0.096	179	**38.61**	0.9945	0.0369
	0.064	15	**34.56**	0.9894	0.0418	0.047	88	34.20	0.9862	0.0532

Table 2: Scaling study, single run per scene, PaviaU repeated as reference. Raw (MB) is the cube size at 16 bits per value and Comp. (MB) the bitstream size, both in decimal megabytes. N/kpx is Gaussians per kilopixel.

Scene	$\mathbf{h \times w \times \lambda}$	**N**	**N/kpx**	**PSNR**↑	**MS-SSIM**↑	**SAM**↓	**Raw** (MB)	**Comp.** (MB)	**bpppb**↓	**Time** (min)
PaviaU (ref.)	$340 \times 340 \times 103$	14,481	125	40.68	0.9946	0.0426	23.8	0.23	0.151	2.0
Pavia Centre	$1096 \times 715 \times 102$	71,506	91	38.99	0.9939	0.0697	159.9	1.10	0.110	12.7
WHU-Hi-HongHu	$940 \times 475 \times 270$	55,934	125	35.32	0.9753	0.0783	241.1	0.88	0.058	22.0
Chikusei	$2372 \times 2196 \times 128$	157,749	30	40.06	0.9866	0.0419	1,333.5	2.80	0.034	34.3

Scaling to larger scenes. The benchmarks above are small crops (0.01–0.12 megapixels); we apply the same pipeline to three larger scenes: Pavia Centre ($1096\times715\times102$), WHU-Hi-HongHu [Zhong et al., 2020] ($940\times475\times270$), and Chikusei [Yokoya and Iwasaki, 2016] ($2372\times2196\times128$, among the largest publicly available airborne HSI scenes), each a single run at one operating point (protocol and caveats in Appendix B). Quality holds at scale (Table 2): 35.3–40.1 dB at 0.034–0.110 bpppb. The Gaussian budget grows *sublinearly* with pixel count: Chikusei has $45\times$ the pixels of PaviaU but needs only $11\times$ the Gaussians (30 versus 125 per kilopixel), since larger scenes are proportionally smoother. Peak GPU memory is instead dominated by the full $h \times w \times \lambda$ reconstruction, nearly independent of N: the 1.3 GB Chikusei cube compresses to 2.8 MB in 34 minutes on one 40 GB GPU without tiling, and the other two scenes train on a 10 GB GPU. WHU-Hi-HongHu (270 bands) gives the lowest PSNR at rank 12, suggesting rank should grow with the spectral dimension.

Downstream evaluation. Rate-distortion metrics do not guarantee analysis value [Penna et al., 2007, Kuester et al., 2022], so Appendix G scores per-pixel classification on original versus decoded data for three scenes under random and spatially blocked splits. Reconstructions at 106:1 to 280:1 match or exceed original-cube accuracy in every setting; JPEG2000 at matched rates is 6–12 dB worse and loses accuracy on the Pavia scenes under blocked splits. Classifying the bitstream's rank-K abundance maps without reconstructing the cube matches or exceeds full-spectrum accuracy with 8.5–$16.9\times$ smaller features; to our knowledge these are the first such results for a splatting HSI codec.

5 Conclusion

This paper proposes no new compression algorithm, it is a transfer study. Five RGB-side techniques transfer to HSI Gaussian splatting and beat published LoR-SGS on 6 of 8 targets; four fail, pointing to properties invisible in RGB experiments (smooth low-rank abundance fields, multiplicative error amplification through spectral coupling, learned-quantizer fragility) that should inform HSI-specific designs; open questions are in Appendix H. The scaling study suggests a concrete deployment: encoding needs one GPU and tens of minutes, so a gigabyte-scale cube shrinks to a few megabytes.

Acknowledgments and Disclosure of Funding

This research was enabled in part by support provided by the Digital Research Alliance of Canada `alliancecan.ca` through the Narval cluster. The work was also in part supported by NSERC DDG-2026-00031 Grant.

References

Yash Bhalgat, Jinwon Lee, Markus Nagel, Tijmen Blankevoort, and Nojun Kwak. LSQ+: Improving low-bit quantization through learnable offsets and better initialization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020.

José M. Bioucas-Dias, Antonio Plaza, Nicolas Dobigeon, Mario Parente, Qian Du, Paul Gader, and Jocelyn Chanussot. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):354–379, 2012.

Akshara Preethy Byju, Gencer Sumbul, Begüm Demir, and Lorenzo Bruzzone. Remote-sensing image scene classification with deep neural networks in JPEG 2000 compressed domain. *IEEE Transactions on Geoscience and Remote Sensing*, 59(4):3458–3472, 2021.

Inmaculada Dópido, Alberto Villa, Antonio Plaza, and Paolo Gamba. A quantitative and comparative assessment of unmixing-based feature extraction techniques for hyperspectral image classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):421–435, 2012.

Bo Du, Mengfei Zhang, Lefei Zhang, Ruimin Hu, and Dacheng Tao. PLTD: Patch-based low-rank tensor decomposition for hyperspectral images. *IEEE Transactions on Multimedia*, 19(1):67–79, 2017.

Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. Learned step size quantization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

European Space Agency. AI for image compression (Phi-sat-2). `https://www.esa.int/Applications/Observing_the_Earth/Phsat-2/AI_for_image_compression`, 2024. Accessed August 2026.

Fernando García-Vílchez, Jordi Muñoz-Marí, Maciel Zortea, Ian Blanes, Vicente González-Ruiz, Gustavo Camps-Valls, Antonio Plaza, and Joan Serra-Sagristà. On the impact of lossy compression on hyperspectral image classification and unmixing. *IEEE Geoscience and Remote Sensing Letters*, 8(2):253–257, 2011.

Christian Geiß, Patrick Aravena Pelizari, Henrik Schrade, Alexander Brenning, and Hannes Taubenböck. On the effect of spatially non-disjoint training and test samples on estimated model generalization capabilities in supervised classification with spatial features. *IEEE Geoscience and Remote Sensing Letters*, 14(11):2008–2012, 2017.

Miguel Hernández-Cabronero, Aaron B. Kiely, Matthew Klimesh, Ian Blanes, Jonathan Ligo, Enrico Magli, and Joan Serra-Sagristà. The CCSDS 123.0-B-2 "low-complexity lossless and near-lossless multispectral and hyperspectral image compression" standard: A comprehensive review. *IEEE Geoscience and Remote Sensing Magazine*, 9(4):102–119, 2021.

ISO/IEC JTC 1/SC 29. Information technology — JPEG AI learning-based image coding system — Part 1: Core coding system. ISO/IEC 6048-1, 2025.

Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

B.-H. Juang and A. H. Gray, Jr. Multiple stage vector quantization for speech coding. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 597–600, 1982.

Azam Karami, Mehran Yazdi, and Grégoire Mercier. Compression of hyperspectral images using discerete wavelet transform and Tucker decomposition. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):444–450, 2012.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (SIGGRAPH)*, 42 (4), 2023.

Jannick Kuester, Johannes Anastasiadis, Wolfgang Middelmann, and Michael Heizmann. Investigating the influence of hyperspectral data compression on spectral unmixing. In *Image and Signal Processing for Remote Sensing XXVIII*, volume 12267, page 122670H. SPIE, 2022.

Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3D Gaussian representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

Tiantian Li, Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Jun Zhang, and Yan Wang. Gaussian-Image++: Boosted image representation and compression with 2D Gaussian splatting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(8):6442–6449, 2026.

Tianyu Li, Ting Wang, Xile Zhao, and Chao Wang. LoR-SGS: Hyperspectral image compression via low-rank spectral Gaussian splatting. *IEEE Transactions on Geoscience and Remote Sensing*, 63: 1–13, 2025.

Jie Liang, Jun Zhou, Yuntao Qian, Lian Wen, Xiao Bai, and Yongsheng Gao. On the sampling strategy for evaluation of spectral-spatial methods in hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(2):862–880, 2017.

Jiahui Liu, Lili Zhang, Yang Gao, and Lele Qu. GS-HSI: cross-band prior-guided 2D Gaussian splatting for hyperspectral image compression. *Chinese Optics Letters*, 23(11):111106, 2025.

Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021.

Jakub Nalepa, Michal Myller, and Michal Kawulok. Validating hyperspectral image segmentation. *IEEE Geoscience and Remote Sensing Letters*, 16(8):1264–1268, 2019.

Yongjian Nian, Yu Liu, and Zhen Ye. Pairwise KLT-based compression for multispectral images. *Sensing and Imaging*, 17(1):3, 2016.

Barbara Penna, Tammam Tillo, Enrico Magli, and Gabriella Olmo. Transform coding techniques for lossy hyperspectral data compression. *IEEE Transactions on Geoscience and Remote Sensing*, 45 (5):1408–1421, 2007.

Shima Rezasoltani and Faisal Z. Qureshi. Hyperspectral image compression using implicit neural representations. In *Proceedings of the Conference on Robots and Vision (CRV)*, pages 248–255, 2023.

Junqi Shi, Mingyi Jiang, Ming Lu, Tong Chen, Xun Cao, and Zhan Ma. HINER: Neural representation for hyperspectral image. In *Proceedings of the 32nd ACM International Conference on Multimedia (ACM MM)*, 2024.

Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The JPEG 2000 still image compression standard. *IEEE Signal Processing Magazine*, 18(5):36–58, 2001.

Christopher Thirgood, Oscar Mendez, Erin Ling, Jon Storey, and Simon Hadfield. HyperGS: Hyperspectral 3D Gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

Róbert Torfason, Fabian Mentzer, Eiríkur Ágústsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Towards image understanding from deep compression without decoding. In *International Conference on Learning Representations (ICLR)*, 2018.

Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, volume 2, pages 1398–1402, 2003.

Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for Gaussian splatting. *Journal of Machine Learning Research*, 26(34):1–17, 2025.

Naoto Yokoya and Akira Iwasaki. Airborne hyperspectral data over Chikusei. Technical Report SAL-2016-05-27, Space Application Laboratory, University of Tokyo, 2016.

Roberta H. Yuhas, Alexander F. H. Goetz, and Joe W. Boardman. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (SAM) algorithm. In *Summaries of the Third Annual JPL Airborne Geoscience Workshop*, pages 147–149, 1992.

Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. GaussianImage: 1000 FPS image representation and compression by 2D Gaussian splatting. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2024.

Yunxiang Zhang, Bingxuan Li, Alexandr Kuznetsov, Akshay Jindal, Stavros Diolatzis, Kenneth Chen, Anton Sochenov, Anton Kaplanyan, and Qi Sun. Image-GS: Content-adaptive image representation via 2D Gaussians. In *ACM SIGGRAPH 2025 Conference Papers*, 2025.

Yanfei Zhong, Xin Hu, Chang Luo, Xinyu Wang, Ji Zhao, and Liangpei Zhang. WHU-Hi: UAV-borne hyperspectral with high spatial resolution (H2) benchmark datasets and classifier for precise crop identification based on deep convolutional neural network with CRF. *Remote Sensing of Environment*, 250:112012, 2020.

Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suya You, Zhangyang Wang, and Achuta Kadambi. Feature 3DGS: Supercharging 3D Gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

A Extended Related Work

Hyperspectral image compression. Conventional codecs such as JPEG2000 [Skodras et al., 2001] and PCA-DCT [Nian et al., 2016] ignore rich inter-band correlations and distort narrow-band spectral features. Decomposition-based methods like Tucker decomposition [Karami et al., 2012] and patch-based low-rank tensor decomposition [Du et al., 2017] factorize the cube into compact tensor representations but struggle with nonlinear spectral mixtures at material boundaries. Implicit neural representations [Rezasoltani and Qureshi, 2023, Shi et al., 2024] improve rate-distortion but need thousands of iterations per image and high GPU memory, and the resulting MLP weights are difficult to compress further.

Gaussian splatting for image compression. 3D Gaussian splatting [Kerbl et al., 2023] represents scenes as explicit Gaussians for real-time rendering; Compact3DGS [Lee et al., 2024] reduces its storage via learnable masking and codebook quantization. GaussianImage [Zhang et al., 2024] first adapted 2D splatting to single-image compression; GaussianImage++ [Li et al., 2026] added distortion-driven densification, content-aware filters, and attribute-separated LSQ+. In the hyperspectral domain, HyperGS [Thirgood et al., 2025] extends 3DGS to hyperspectral novel view synthesis but not single-image compression; LoR-SGS [Li et al., 2025] is the first to apply Gaussian splatting to HSI

compression via NMF decomposition, though still with RVQ [Juang and Gray, 1982]. Concurrently, GS-HSI [Liu et al., 2025] applies 2D splatting to HSI compression through cross-band prior reuse and adaptive resampling, evaluated against learned-codec baselines.

Quantization for neural and Gaussian-splatting compression. Splatting parameters are FP32 during training, so quantization is critical for file size [Nagel et al., 2021]. GaussianImage [Zhang et al., 2024] uses RVQ [Juang and Gray, 1982] with cascaded codebooks, which loses quality relative to scalar approaches [Li et al., 2026]. Quantization-aware training [Jacob et al., 2018] simulates quantization during training; LSQ [Esser et al., 2020] learns the step size via straight-through gradients, and LSQ+ [Bhalgat et al., 2020] adds a learnable offset. These techniques had been studied for RGB Gaussian attributes but not for HSI abundance features, whose statistics (nonnegative mixing weights) differ fundamentally.

B Configuration, Hardware, and Protocol Details

Datasets and operating points.

Dataset	Spatial	Bands	Rank	N (high / low)
PaviaU	340×340	103	12	9,148 / 5,427
Salinas	180×180	204	12	4,858 / 2,276
JasperRidge	100×100	198	10	1,434 / 751
Urban	307×307	162	12	11,685 / 5,721

Final configuration. 6-bit unsigned LSQ+ for abundances, 8-bit signed LSQ+ for Cholesky covariance, FP16 positions, D3 densification from $N/2$ to N by iteration 2000, two-stage QAT with 50% warmup, gradient-guided initialization (Sobel gradient blended with NMF reconstruction error), SAM loss weight 0.05, learnable endmembers at learning rate 10^{-4}, pruning interval 100, content-aware filters disabled. Total iterations are $1.5\times$ the baseline: 10,500 for PaviaU, Salinas, and Urban; 12,000 for JasperRidge.

Baseline. Published LoR-SGS results from Table II of Li et al. [2025] (RVQ quantization). All quality metrics are compared at matched bitrate targets.

Controlled rerun of LoR-SGS. Table 1 compares with the numbers published in Table II of Li et al. [2025]. To control for training length and bitrate matching, we reran the authors' released code, averaged over 5 runs. From their Table I we used the Gaussian counts and codebook sizes and the per-dataset commitment weight, and ran for 7000 or 8000 iterations. We re-enabled the RVQ step, which the released code leaves commented out, and kept the Adan optimizer of the released code rather than the Adam the paper names, since Adam is 0.86 dB lower on average. Table 3 shows that the rerun at the same N reproduces the published PSNR, but at 1.2 to 1.4 times the published bpppb; the paper's own Tables V and VI report 0.11–0.13 bpppb for configurations its Table II labels 0.10, so we read those labels as rounded targets. Training $1.5\times$ longer adds 0.02–0.24 dB. With N instead reduced until the measured bpppb matches the published one, PSNR falls 0.6–3.4 dB below the published figure in seven of eight cells; ATLAS exceeds this rate-matched rerun on all eight by $+2.56$ dB on average.

Run variance. Table 1 reports means over 5 independent runs per target; the per-target PSNR standard deviation over the runs is 0.06–0.17 dB. Every win in Table 1 exceeds twice this spread (the smallest, Urban at 0.096 bpppb, is $+0.39$ dB against a 0.10 dB standard deviation), while the two deficits are within it: Urban@0.05 (-0.04 dB) is under one standard deviation and PaviaU@0.10 (-0.09 dB) about 1.5.

Hardware. Benchmark runs use one NVIDIA RTX 3080 (10 GB). The Chikusei scaling run uses one NVIDIA A100 (40 GB) on the Narval cluster of the Digital Research Alliance of Canada.

Scaling protocol. Each large scene is a single run at one operating point with fixed seeds, using the final configuration with per-scene schedules: 12,000 iterations for Pavia Centre and WHU-Hi-

Table 3: Two reruns of the authors' released LoR-SGS code, averaged over 5 runs, against their published Table II. *Same* N: at the Gaussian counts their Table I lists; ΔPSNR is rerun minus published, and $\Delta_{1.5\times}$ is the change in PSNR when the same run is trained $1.5\times$ longer. *Same bitrate*: with N reduced until the measured bpppb matches the published one. The PaviaU 0.06 same-N cell collapsed on four seeds of five (its fifth seed gave 0.061 bpppb and 35.38 dB); the Urban 0.05 same-bitrate cell collapsed on two.

	Published		**Rerun, same N**					**Rerun, same bitrate**				
Dataset	bpppb	PSNR	N	bpppb	PSNR	ΔPSNR	$\Delta_{1.5\times}$	N	bpppb	PSNR	ΔPSNR	$\Delta_{1.5\times}$
PaviaU	0.10	38.95	23,000	0.130	38.94 ± 0.10	-0.01	$+0.13$	17,398	0.100	37.81 ± 0.11	-1.14	$+0.03$
PaviaU	0.06	35.76	10,000	0.053	33.25 ± 1.19	-2.51	$+0.03$	10,853	0.066	35.97 ± 0.09	$+0.21$	$+0.13$
Salinas	0.10	41.41	13,000	0.138	41.44 ± 0.13	$+0.03$	$+0.24$	8,479	0.099	40.56 ± 0.16	-0.85	$+0.17$
Salinas	0.05	37.83	5,000	0.057	37.77 ± 0.06	-0.06	$+0.03$	4,162	0.051	37.20 ± 0.06	-0.63	$+0.02$
JasperRidge	0.10	37.19	2,750	0.135	36.99 ± 0.26	-0.20	$+0.05$	1,868	0.098	33.82 ± 0.17	-3.37	$+0.02$
JasperRidge	0.06	30.79	1,000	0.075	31.46 ± 0.08	$+0.67$	$+0.02$	476	0.060	28.77 ± 0.09	-2.02	$+0.05$
Urban	0.10	38.22	29,000	0.122	38.49 ± 0.40	$+0.27$	$+0.13$	23,148	0.098	36.94 ± 0.04	-1.28	$+0.06$
Urban	0.05	34.24	14,500	0.067	35.25 ± 0.25	$+1.01$	$+0.14$	10,412	0.046	32.37 ± 1.38	-1.87	$+0.03$

HongHu and 10,000 for Chikusei, QAT starting at 50%, densification complete by iteration 2500 (3000 for Chikusei), and rank 16 for Chikusei, matching its 128 bands (rank 12 otherwise). Chikusei is cropped from 2517×2335 to its largest rectangle free of empty flight-strip borders. Repeats of the PaviaU reference varied by up to 0.4 dB from nondeterministic gradient accumulation in the rasterizer's backward pass, so single-run PSNRs carry a few tenths of a dB of run-to-run variation. No published compression results exist for these scenes at comparable rates, so Table 2 reports absolute metrics with PaviaU as a reference point rather than a baseline comparison.

Code. Code is available at `https://github.com/vclab/atlas-hsi`

C Failed Transfer Measurements

Table 4: Techniques that fail to transfer from RGB to HSI (Section 3.2), each evaluated at its RGB-default setting. Baselines differ by row because each technique was tested against the pipeline configuration current at that stage of the study (FP32 D3 training for CAF, the single-stage quantized pipeline for position quantization, the final two-stage pipeline for progressive quantization and delta coding); each row is a controlled comparison within one configuration.

Technique	**Dataset**	**Metric**	**Baseline**	**With**	**Change**
CAF	PaviaU	PSNR (dB)	40.34	39.63	-0.71
	Urban	PSNR (dB)	39.61	37.89	-1.72
Position quant. (12-bit LSQ+)	PaviaU	PSNR (dB)	37.33	34.04	-3.29
Progressive quant. (500-iter ramp)	PaviaU	PSNR (dB)	38.86	23.27	-15.59
Delta coding	JasperRidge	bpppb	0.1031	0.1520	$+47.4\%$
	PaviaU	bpppb	0.0959	0.1097	$+14.4\%$

Why CAF fails. We ran two FP32 tests on PaviaU at fixed N=14,500 with CAF off, setting the minimum Gaussian width to 0.25, 0.28, 0.35, 0.5, 0.75, and 1.0 yields 40.07, 40.36, 41.04, 41.24, 41.29, and 41.34 dB (averaged over 2 runs, repeats differ by at most 0.1 dB): PSNR falls monotonically as the floor drops below 0.5, and 0.28 is the floor CAF assigns to late-added Gaussians. With D3 (7,250 to 14,500), CAF loses -0.83 dB (39.94 versus 40.77 dB), and clamping its floor at 0.5 produces 40.78 dB.

D Cumulative FP32 Ablation

Table 5 isolates each transferred technique's contribution at FP32 (no quantization). D3 contributes most on large datasets (Salinas $+1.05$ dB). Error-driven initialization and learnable endmembers add

consistent gains across all datasets. SAM loss has minimal PSNR impact (-0.02 to $+0.15$ dB) but substantially improves the SAM metric.

Table 5: Cumulative FP32 ablation: incremental PSNR (dB) of each transferred technique. Baseline: our reproduction with the unified kernel only. Deltas are relative to the previous row.

Technique	JasperRidge	PaviaU	Salinas	Urban
Baseline (unified kernel)	38.59	39.81	43.28	38.82
+ D3 densification	-0.53^*	$+0.50$	$+1.05$	$+0.77$
+ Error-driven init	$+0.57$	$+0.35$	$+0.03$	$+0.34$
+ Learnable endmembers	$+0.11$	$+0.27$	$+0.11$	$+0.28$
+ SAM loss ($w = 0.05$)	-0.02	$+0.08$	$+0.15$	$+0.12$
Total (FP32 ceiling)	**38.72**	**41.01**	**44.60**	**40.32**

*D3 on JasperRidge starts at $N/2 = 750$ Gaussians on a 100×100 image, which is very constrained; error-driven initialization (next row) more than recovers the deficit. The init row measures error-driven placement alone; the Sobel blend of Section 3 adds a further $+0.08$ to $+0.19$ dB on top. Per-technique deltas are rounded independently and may not sum exactly to the measured totals.

Quantization gap. Two-stage QAT loses on average 0.88 dB versus the FP32 ceiling while reducing bitrate 3–$4\times$. Single-stage QAT, with the quantizers active from the first iteration and no FP32 warmup, reaches 37.33 dB on PaviaU versus 39.59 dB for the two-stage approach, a 2.26 dB gap, showing that the warmup stage is critical.

E Unified Rasterization Kernel

The original LoR-SGS implementation splits the K abundance channels into three groups of at most four and invokes the rasterizer once per group. We rasterize all K channels in a single pass. We claim no novelty: Feature 3DGS introduced a parallel N-dimensional rasterizer for 3D scenes [Zhou et al., 2024], gsplat supports N-dimensional feature rendering [Ye et al., 2025], and Image-GS renders 2D Gaussians with variable channel counts [Zhang et al., 2025].

Table 6 isolates the kernel from every other change: both implementations run their production training loop at FP32 with fixed N (no densification, quantization stages, or pruning) on identical data, initialization, loss, and optimizer, averaged over five runs. The unified kernel cuts wall-clock training time by 11–39% (mean $1.35\times$), with the largest gains on small scenes, where the fixed cost of each launch is a larger share of the total time. Rendered outputs are bit-identical to the multipass implementation at both ranks used in this paper ($K = 10$ and $K = 12$), so every quality result would be unchanged under the original kernel.

Table 6: Kernel benchmark: multipass (three rasterizer launches per forward pass) vs unified single-pass. Identical FP32 fixed-N training; N matches the high-bitrate operating points of Table 1; mean $\pm$ std over 5 runs on the RTX 3080.

Dataset	N	Iterations	Multipass (s) ↓	Unified (s) ↓	Speedup ↑
PaviaU	9,148	7000	45.7 ± 1.0	39.7 ± 0.8	$1.15\times$
Salinas	4,858	7000	26.2 ± 0.5	17.8 ± 0.5	$1.47\times$
Jasper Ridge	1,434	8000	26.3 ± 0.4	15.9 ± 0.7	$1.65\times$
Urban	11,685	7000	56.6 ± 0.5	50.5 ± 1.2	$1.12\times$
Mean					$\mathbf{1.35}\times$

F Full Comparison Table

Table 7 extends Table 1 with the remaining baselines reported by Li et al. [2025]: JPEG, PCA-DCT, and HP-INR. HINER did not report Jasper Ridge quality metrics.

Table 7: Full compression comparison on four HSI datasets, including all baselines of [Li et al., 2025]. Bold PSNR marks the winner at each bitrate target. Size in KB. LoR-SGS rerun: the authors' released code, mean of 5 runs, the Gaussian count was adjusted to the published bpppb (Appendix B).

Method	Pavia University					Salinas				
	bpppb	Size	PSNR	MS-SSIM	SAM	bpppb	Size	PSNR	MS-SSIM	SAM
JPEG	0.10	145	23.76	0.7019	0.2454	0.10	81	25.46	0.7858	0.1112
JPEG2000	0.10	145	27.11	0.8790	0.1398	0.10	81	28.36	0.8804	0.0541
PCA-DCT	0.10	145	24.70	0.7999	0.2270	0.10	81	21.94	0.6955	0.0981
INR	0.10	145	33.15	0.9760	0.0802	0.10	81	37.38	0.9822	0.0257
HP-INR	0.06	87	32.83	0.9747	0.0812	0.05	40	37.30	0.9821	0.0258
HINER	0.10	145	32.20	0.9532	0.2582	0.10	81	31.35	0.9398	0.0609
	0.06	87	15.55	0.2040	0.6934	0.05	40	28.98	0.8975	0.0597
LoR-SGS	0.10	145	**38.95**	0.9926	0.0564	0.10	81	41.41	0.9916	0.0135
	0.06	87	35.76	0.9881	0.0697	0.05	40	37.83	0.9869	0.0170
LoR-SGS rerun	0.100	145	37.81	0.9916	0.0604	0.099	80	40.56	0.9903	0.0145
	0.066	96	35.97	0.9887	0.0684	0.051	41	37.20	0.9856	0.0178
Ours	0.096	140	38.86	0.9922	0.0493	0.098	79	**43.11**	0.9946	0.0118
	0.058	84	**36.48**	0.9866	0.0586	0.049	40	**40.03**	0.9884	0.0141
Method	**Jasper Ridge**					**Urban**				
	bpppb	Size	PSNR	MS-SSIM	SAM	bpppb	Size	PSNR	MS-SSIM	SAM
JPEG	0.10	24	21.15	0.8267	0.2177	0.10	186	22.15	0.7358	0.2535
JPEG2000	0.10	24	18.71	0.5817	0.2800	0.10	186	23.70	0.8314	0.1616
PCA-DCT	0.10	24	28.91	0.9605	0.1018	0.10	186	23.47	0.8141	0.2758
INR	0.10	24	31.34	0.9818	0.0706	0.10	186	31.20	0.9733	0.0862
HP-INR	0.06	15	31.25	0.9816	0.0709	0.05	93	30.79	0.9721	0.0874
HINER	0.10	24	—	—	—	—	—	—	—	—
	0.06	15	—	—	—	0.05	93	24.66	0.8751	1.5708
LoR-SGS	0.10	24	37.19	0.9936	0.0380	0.10	186	38.22	0.9940	0.0468
	0.06	15	30.79	0.9804	0.0598	0.05	93	**34.24**	0.9883	0.0652
LoR-SGS rerun	0.098	24	33.82	0.9894	0.0470	0.098	183	36.94	0.9925	0.0526
	0.060	14	28.77	0.9672	0.0719	0.046	86	32.37	0.9831	0.0728
Ours	0.103	25	**38.10**	0.9949	0.0341	0.096	179	**38.61**	0.9945	0.0369
	0.064	15	**34.56**	0.9894	0.0418	0.047	88	34.20	0.9862	0.0532

G Downstream Task Evaluation

Motivation. Rate-distortion metrics do not reliably predict the effect of lossy compression on analysis tasks [Penna et al., 2007, Kuester et al., 2022], and the Earth-observation community has historically distrusted lossy coding for exactly this reason. In practice, the European Space Agency (ESA) deployed Phi-Sat-2 with learned lossy compression onboard, they validated this choice by a downstream task: building-segmentation masks predicted from the decompressed images matched those from the originals [European Space Agency, 2024]. We therefore evaluate ATLAS on the standard protocol of that literature: score the same analysis pipeline on the original cube and on the reconstruction, and report the difference [García-Vílchez et al., 2011]. To our knowledge no Gaussian-splatting codec has reported downstream numbers for HSI; the only learned-codec precedent is HINER, which loses 2.1 points of overall accuracy (OA) at 109:1 on PaviaU [Shi et al., 2024].

Scenes. PaviaU (our 340×340 crop; class 5 is absent from the crop, leaving 8 classes and 30,556 labeled pixels, so its numbers are self-relative and not comparable to published full-scene accuracies), Pavia Centre (full scene, 9 classes, 148,152 labeled pixels), and WHU-Hi-HongHu (full scene, 22 classes, 386,693 labeled pixels) [Zhong et al., 2020]. Ground-truth rasters are cropped exactly as the cubes are loaded; we verified all label counts against the crops.

Conditions. (1) The original cube, band-max normalized, which is both the training target and what the codec reconstructs. (2) The ATLAS reconstruction, decoded from the entropy-coded bitstream at the operating points reported in the main paper: 41.0 dB at 0.150 bpppb (106:1) for PaviaU, 38.7 dB at 0.108 (149:1) for Pavia Centre, 35.3 dB at 0.057 (280:1) for HongHu (these single-run figures differ from Table 2 by up to 0.4 dB, within the run-to-run variation noted in Appendix B). (3) JPEG2000 Part 1 at the same rate target per scene (single multi-component codestream, global

rate allocation over all bands, uint16 input; this is the standard JPEG2000 baseline configuration for HSI). JPEG2000 reaches 29.0, 28.1, and 29.1 dB at those rates, 6–12 dB below ATLAS. Its rate is measured on the whole file while ATLAS reports entropy payload; the difference is small relative to the PSNR gap. (4) The splatted rank-K abundance maps decoded from the ATLAS bitstream without the endmember mixing step, i.e. partial decoding: the splat still runs, but the $K \times \lambda$ mixing product is never formed and the $h \times w \times \lambda$ cube never reconstructed. The resulting per-pixel features are K abundances instead of λ bands: 12 instead of 102 or 103 on the Pavia scenes ($\sim 8.5\times$ fewer) and 12 instead of 270 on HongHu ($22.5\times$; $16.9\times$ at the rank-16 setting used in the main text). (5) A control of rank-K NMF abundances computed on the original cube ($\mathbf{A}_0$). The compressed-domain features differ from the spectra by two transformations, rank reduction and then Gaussian fitting with quantization. The accuracy drop from spectra to $\mathbf{A}_0$ is the cost of rank reduction, and the drop from $\mathbf{A}_0$ to the splatted abundances is the cost of Gaussian fitting and quantization.

Classifiers and splits. Per-pixel support vector machine (SVM) with a radial basis function (RBF) kernel (C and γ grid-searched by 3-fold cross-validation on the training pixels only; features standardized with statistics from the training pixels only) and a 200-tree random forest as a second classifier. Each condition randomly samples 100 training pixels per class and tests on the rest, repeated for 10 seeded samples; we report mean and standard deviation over the samples (OA, average accuracy (AA), and Cohen's kappa; kappa tracks OA in every condition and is omitted from the table for space). Sample k uses the same seed, and therefore the same training and test pixels, in every condition, so the per-sample differences between conditions are paired observations.

Random versus blocked splits. Across the field, random stratified splits are the default. A per-pixel classifier sees only the pixel it is classifying, so adjacent training pixels cannot influence a test prediction [Liang et al., 2017, Geiß et al., 2017, Nalepa et al., 2019]. A lossy reconstruction blends each pixel with its neighbours, so a decoded test pixel carries signal from adjacent training pixels, which usually share its label. We therefore also report a spatially blocked protocol: the scene is tiled into 32 px blocks, whole blocks are assigned to a train or test pool (resampled per seed), training pixels are sampled only from train blocks, and the test set is every labeled pixel in test blocks. Unused train-block pixels are discarded rather than tested. Blocks are not separated by a discarded margin, so mixing across borders is reduced rather than eliminated (see Leakage control).

Results. Table 8 reports the SVM results; the random forest reproduces every ordering and prefers the compressed-domain features even more strongly (e.g. HongHu original spectra 65.9% OA versus 12-d splatted abundances 73.8%). Full per-sample records for both classifiers are in the experiment repository.

Reconstruction versus original. The ATLAS reconstruction matches or exceeds original-cube accuracy on every scene under both protocols (blocked OA $+2.1$, $+0.8$, and $+0.3$ points respectively; on PaviaU the paired per-draw difference is positive on 10 of 10 random-split draws and 9 of 10 blocked draws). Seed-paired Wilcoxon signed-rank tests over the 10 draws give $p \leq 0.019$ on five of the six scene–protocol combinations; the sixth, HongHu under blocked splits ($+0.4$ OA points, $p = 0.25$), is a statistical tie: there ATLAS matches the original rather than exceeding it, which is why we claim "match or exceed" and not "exceed." The gain is consistent with the low-rank splat acting as a spectral denoiser: each reconstructed spectrum is a smooth rank-K mixture, which suppresses per-band noise that per-pixel classifiers are sensitive to. This contrasts with HINER's reported 2.1-point OA loss at a comparable ratio on PaviaU [Shi et al., 2024].

JPEG2000 at matched rates. JPEG2000 loses 6–12 dB of PSNR, and once spatial leakage is controlled it also loses accuracy on both Pavia scenes (blocked OA -5.1 and -1.6 points versus the original). JPEG2000's one win, HongHu, is explained by the blur control: a $\sigma = 1.5$ Gaussian blur of the original, with no compression at all, matches or exceeds JPEG2000 on every metric under both protocols. HongHu's labels are assigned per crop field, so every pixel in a field shares one class, and spatial smoothing pulls each pixel toward its field's typical spectrum, which helps any per-pixel classifier. JPEG2000's HongHu score shows the benefit that smoothing gives even as the run loses 6–12 dB of PSNR, so on scenes with large uniform fields accuracy alone cannot measure how much the codec preserved. Therefore, we report PSNR, both split protocols, and the blur control together: PSNR measures preserved information directly, blocked splits remove label leakage across the train/test split, and the blur control shows what smoothing achieves on its own.

Table 8: Downstream classification on three scenes at 106:1 to 280:1 compression. ATLAS reconstructions match or exceed original-cube accuracy under both split protocols. Per-pixel SVM (RBF kernel), 100 training pixels per class, mean $\pm$ std over 10 seed-paired samples; OA is the percentage of test pixels classified correctly, AA the unweighted mean of per-class accuracies. Random splits sample test pixels anywhere in the scene; blocked splits train and test on disjoint 32 px blocks, suppressing spatial leakage. *Original*: per-pixel spectra of the uncompressed cube. *ATLAS*: reconstruction decoded from the ATLAS bitstream. *JPEG2000*: standard multi-component codestream at the same bitrate (6–12 dB lower PSNR). *Splatted abundances*: rank-K abundance maps decoded from the bitstream without reconstructing the spectral cube. *NMF $\mathbf{A}_0$ control*: rank-K NMF abundances of the original cube, isolating rank reduction from the codec. *Gaussian blur* (HongHu only): original cube blurred at $\sigma = 1.5$ to match JPEG2000's smoothing, a no-compression control. Pavia University uses our 8-class crop, so its values are not comparable to published full-scene accuracies.

	Random split		**Blocked split**	
Condition	OA (%)	AA (%)	OA (%)	AA (%)
Pavia University (8-class crop, ATLAS at 106:1)				
Original	89.3 ± 0.7	90.5 ± 0.5	84.1 ± 4.4	86.9 ± 2.1
ATLAS	91.3 ± 1.2	92.1 ± 0.6	86.2 ± 3.6	87.9 ± 1.5
JPEG2000	90.7 ± 0.7	88.9 ± 0.6	79.0 ± 3.7	75.5 ± 2.8
Splatted abundances (12-d)	92.4 ± 0.8	92.5 ± 0.7	87.7 ± 3.6	89.3 ± 1.8
NMF $\mathbf{A}_0$ control (12-d)	88.3 ± 1.1	90.0 ± 0.3	84.2 ± 3.3	87.2 ± 1.0
Pavia Centre (9 classes, ATLAS at 149:1)				
Original	97.4 ± 0.3	94.6 ± 0.3	96.5 ± 0.6	91.7 ± 1.5
ATLAS	98.0 ± 0.3	95.5 ± 0.4	97.3 ± 0.5	93.3 ± 1.4
JPEG2000	96.3 ± 0.4	93.9 ± 0.6	94.9 ± 0.7	89.5 ± 2.0
Splatted abundances (12-d)	98.2 ± 0.2	96.0 ± 0.4	97.3 ± 0.4	93.6 ± 1.2
NMF $\mathbf{A}_0$ control (12-d)	97.4 ± 0.2	94.3 ± 0.3	96.6 ± 0.5	91.8 ± 1.5
WHU-Hi-HongHu (22 classes, ATLAS at 280:1; rank 16 at 233:1)				
Original	75.1 ± 0.7	72.1 ± 0.2	72.6 ± 1.6	66.5 ± 1.1
ATLAS	75.9 ± 1.2	73.6 ± 0.5	73.0 ± 1.6	66.7 ± 0.6
JPEG2000	90.9 ± 0.3	90.3 ± 0.3	85.3 ± 1.0	75.0 ± 1.8
Gaussian blur (no codec)	91.1 ± 0.4	91.5 ± 0.2	87.4 ± 1.7	82.0 ± 2.0
Splatted abundances (12-d)	73.8 ± 1.0	72.0 ± 0.4	71.6 ± 2.0	65.5 ± 1.3
NMF $\mathbf{A}_0$ control (12-d)	65.8 ± 1.5	63.9 ± 0.5	64.5 ± 2.1	59.5 ± 1.1
Splatted abundances (16-d)	77.1 ± 0.8	76.2 ± 0.5	74.9 ± 1.1	68.4 ± 1.7
NMF $\mathbf{A}_0$ control (16-d)	68.4 ± 1.4	65.1 ± 0.3	66.9 ± 1.9	60.9 ± 1.1

Classification without reconstruction. Classifying the splatted rank-K abundance maps, without reconstructing the spectral cube, matches or exceeds full-spectrum accuracy on the Pavia scenes at 12 dimensions and on HongHu at 16 dimensions (blocked OA $+3.6$, $+0.8$, and $+2.3$ points respectively). The $\mathbf{A}_0$ control splits this into two causes: rank-12 reduction removes 1.0 OA point on the Pavia scenes but 9.3 points on 22-class HongHu under random splits, and the Gaussian fitting step recovers 8.0 of those points at equal dimensionality, leaving the 12-d features 1.3 points below the full spectrum and indicating that the splatted abundances carry spatial context that per-pixel NMF lacks. Raising the abundance rank to 16 removes the remaining 1.3-point deficit and puts HongHu above the full spectrum (+2.0 OA points under random splits, +2.3 under blocked), at a 20% higher bitrate (0.0686 versus 0.0572 bpppb, still 233:1) and with 0.5 dB higher PSNR. We take this as a tested design rule: rank 12 suffices up to 9 classes; 22 classes need rank 16. For deployment, the memory argument is direct: the analysis reads 12–16 values per pixel instead of 102–270 bands, the decode never allocates the $h \times w \times \lambda$ cube (1.3 GB for a Chikusei-scale scene), and the work the analysis does, no longer grows with the number of bands. Abundance features as classifier input follow established practice [Dópido et al., 2012], and task inference on codec latents follows Torfason et al. [2018] and ISO/IEC JTC 1/SC 29 [2025]; compressed-domain analysis for remote sensing is surveyed in Byju et al. [2021].

Leakage control. Every ATLAS gain under random splits is also a gain under blocked splits, so the gains are not an artifact of labels leaking across the train/test split. Enlarging the blocks from 32 to 96 px lowers HongHu's blocked OA to 67.4% for the original and 75.6% for JPEG2000, shrinking the gap from $+12.7$ to $+8.2$ points, so at most a third of the 32 px gap can be border leakage; because

96 px blocks also leave rare classes with too few pixels on one side of the split, 32 px is the reported protocol.

Scope. Operational missions today compress losslessly at 1.4:1 to 3.4:1; the 100:1 to 500:1 regime studied here is the target of aggressive learned compression, not current practice. JasperRidge and Urban carry unmixing rather than classification ground truth; an unmixing-fidelity evaluation in the style of García-Vílchez et al. [2011] is left to future work, as is classification from per-Gaussian parameters directly.

H Open Questions

Spectral-aware densification. D3 places new Gaussians at pixels with the highest spatial reconstruction error, a criterion inherited from RGB. An HSI-specific variant could weight errors by spectral importance or target bands where the endmember decomposition is poorest, potentially recovering quality on the two targets where we underperform.

Adaptive bit allocation. The pipeline uses a fixed bit budget (6-bit abundance, 8-bit covariance) for all Gaussians. Allocating more bits to Gaussians at material boundaries, where abundance gradients are steepest and quantization error propagates most severely, could improve quality at no bitrate cost.

Endmember quantization. $\mathbf{E}_0$ is stored at FP16, a small but non-negligible fraction of the bitstream. Whether it tolerates lower precision without degrading spectral fidelity across all bands is untested.

Scaling beyond single-GPU memory. The binding constraint at scale is the full-cube reconstruction held in GPU memory during training, which is nearly independent of the Gaussian count. Scenes whose $h \times w \times \lambda$ reconstruction exceeds single-GPU memory, such as full satellite swaths, would require tiled or streamed training, which we have not attempted.

Content-aware filters revisited. CAF fails because it lets late-added Gaussians become narrower than 0.5 pixels, too narrow to overlap into a smooth surface. Appendix C shows that leaving CAF on but not letting any Gaussian go narrower than 0.5 removes the loss entirely. A CAF with a learned minimum width rather than the fixed $hw/(9\pi N)$ formula might keep its spatial adaptivity, and since a minimum width of 1.0 scored a further $+0.10$ dB at fixed N, the best value for HSI may be above 0.5.