

Abundance Estimation Methods in Spectral Unmixing for Real Data

Daniele Cerra, Miguel Pato, Emiliano Carmona

DLR, Earth Observation Center, 82234 Wessling, Germany - (daniele.cerra, miguel.figueiredovazpato, emiliano.carmona)@dlr.de

Keywords: Spectral Unmixing, Abundance Estimation, Sparse Reconstruction, Benchmark Datasets.

Abstract

Spectral unmixing estimates the fractional abundances of materials, having associated spectra called endmembers, in pixels acquired by imaging spectrometers. Validation of abundance estimation methods typically relies on synthetic data or comparisons to results obtained by other algorithms. This study considers results of typical abundance estimation algorithms on the DLR HySU (Hyperspectral Unmixing) benchmark dataset, which contains actual imaging spectrometer data acquired over several arrangements of known-size material patches for physically traceable validation. Abundance estimates are compared against measured target areas in pixels with different degrees of mixtures. We evaluate least squares and sparse unmixing methods across different noise scenarios on real data, and by contaminating the library through addition of non-relevant endmembers. Additionally, as a way to approximate hard sparsity constraints, we enforce cardinality constraints on endmember subsets, identifying those minimizing abundance errors relative to the full library. Results suggest that fully constrained least squares yields usually the best results, but struggles in cases of highly mixed pixels. Finally, we test quantization of abundance values as a way to enforce sparsity in non-negative least squares with limited but encouraging results. Overall, the increase in accuracy of results enforcing sparse solutions supports the use of computationally efficient sparse unmixing methods in practical scenarios, part of which may become feasible if quantum computing capabilities improve in the future.

1. Introduction

Hyperspectral imaging captures hundreds of contiguous spectral bands, enabling detailed material characterization (Bioucas-Dias et al., 2012). However, due to limited spatial resolution, pixels often contain mixtures of materials, requiring spectral unmixing to retrieve fractional abundances. In recent years, spaceborne hyperspectral platforms (such as DESIS, PRISMA, HISUI and EnMAP) operate at a typical ground sampling distance of 30m, with most land pixels being spectrally mixed at this scale. This makes spectral unmixing and its physical validation a relevant topic of research.

The validation of spectral unmixing remains a challenge, as real images having abundance maps to be used as ground truth are essentially non-existent. Therefore, most studies use either synthetic data (Iordache et al., 2011, Iordache et al., 2014, Eches et al., 2010, Li et al., 2017) or results obtained by other baseline algorithms (Zhu, 2017), both lacking physical reference. Recently, the lack of such dataset was addressed by (Haijen et al., 2025), where the authors propose a new benchmark dataset containing different mixtures of a total of 4 colors printed on paper, with known abundance for each printed mixture. Nevertheless, this dataset is still derived from laboratory measurements and limited to a single material (printed paper).

To overcome these limitations, Cerra *et al.* introduced the DLR HySU dataset in (Cerra et al., 2021), featuring an airborne image acquired by the HySpex sensor (both VNIR and SWIR) of five materials deployed as square patches of known sizes (with sides in the range 3 m–0.25 m) on a homogeneous background. The dataset allows physically measurable abundance values, although not directly: while the exact composition of each image element in terms of abundance values remains unknown, the areas of the objects in specific subsets of the image are known. This allows estimating the area from the total abundances of a given material in a given extent of the image, and comparing it to the available ground truth area.

The aim of this paper is to use the DLR HySU dataset to investigate the behaviour of typical abundance estimation algorithms in non-ideal conditions. Beyond classical least squares and sparsity promoting techniques, recent abundance estimation methods rely on probabilistic priors or other assumptions, such as spatially regularized unmixing, which exploits spatial correlations in real scenes (Iordache et al., 2012). Here, we opted to assess how representative least squares and sparse unmixing methods perform on real imaging conditions in a framework minimally conditioned by additional underlying assumptions.

Specifically, we quantify the impact of noise corruption and a more complex spectral library with respect to one containing the known spectra in the image, as provided in the DLR HySU benchmark. We additionally test cardinality constraints as a way to simulate ℓ_0 norm minimization, and assess the impact of quantizing abundance values to enforce sparsity in the results, with considerations on future capabilities of quantum computers to tackle this problem.

Results presented in this work are inherent to abundance estimation only, one of the three main steps in a spectral unmixing workflow. The other two, dimensionality estimation and endmember extraction, are not discussed here. The paper is structured as follows. Section 2 briefly introduces the DLR HySU dataset and the abundance estimation algorithms under test. Section 3 reports results on the different unmixing scenarios. We conclude in Section 4.

2. Data and Methods

This section introduces the benchmark dataset used and the abundance estimation methods assessed in the experimental section.




Figure 1. DLR HySU: High resolution RGB image acquired over the targets of five different materials (bitumen, red metal, red fabric, blue fabric, green fabric) deployed over grass. The square targets have a side ranging from 0.25 m to 3 m. The insert shows spectra from airborne imaging spectrometer data (HySpex) available for the same area and spectrometer (SVC) measurements. For details, please refer to (Cerra et al., 2021).

2.1 DLR HySU Dataset

The dataset was acquired over DLR’s Oberpfaffenhofen site using the airborne HySpex instrument and covers the spectral range 416–992 nm (VNIR only) at 0.7 m spatial resolution (Cerra et al., 2021). Five uniform materials—bitumen, red painted metal, blue fabric, red fabric, and green fabric—were deployed in square patches on grass (Fig. 1). In this paper we solely use the “All Targets Masked” subset, including all deployments of the five targets on a relatively homogeneous grass background. The latter is included in the provided materials spectral library, which has therefore a total of six endmembers, directly selected from the image. Any portion of the image containing other materials, as verified by the creators of the dataset, is here masked out.

Each patch p of each size and material has a known physical area, $A_{\text{true},p}$, reported in (Cerra et al., 2021) and used here as reference for validation. Estimated abundances of the material m yield the total area estimate of the corresponding patch as:

$$A_{\text{est},p} = A_{\text{pix}} \sum_i a_{i,m}, \quad (1)$$

where A_{pix} is the area of a pixel on ground, $a_{i,m}$ is the abundance estimate for material m on pixel i and the sum includes the pixels in the relevant extent of the image. The relative area error (RAE), for each patch, is derived as:

$$\text{RAE}_p = 100 \times \frac{|A_{\text{est},p} - A_{\text{true},p}|}{A_{\text{true},p}}. \quad (2)$$

The resulting assessment is not pixel-based, and therefore prone to assessment errors derived by non-zero abundance values assigned to wrong materials in the areas of interest. Nevertheless, Eq. (2) allows estimating areas derived from abundance values, which to the best of our knowledge is unfeasible on other existing benchmark datasets. Furthermore, smaller targets occupy only a few pixels (namely 4 for 25 cm targets), which mitigates the mentioned problem. Finally, visual inspection of abundance

values for all large targets (2m and 3m) did not reveal significant abundance values for any target outside of its correct area.

2.2 Abundance Estimation Methods

Abundance estimation methods considered in this work include three representative and widely used spectral unmixing algorithms, namely non-negative least squares (NNLS), fully constrained least squares (FCLS), and sparse regression via the least absolute shrinkage and selection operator (LASSO), which estimate material fractions by minimizing spectral reconstruction error under different combinations of non-negativity, sum-to-one, and sparsity constraints:

- NNLS, enforcing non-negativity:

$$\min_a \|y - Ea\|_2^2 \quad \text{s.t.} \quad a_i \geq 0. \quad (3)$$

- FCLS, enforcing non-negativity and sum-to-one constraint:

$$\min_a \|y - Ea\|_2^2 \quad \text{s.t.} \quad a_i \geq 0, \|a\|_1 = 1. \quad (4)$$

- LASSO, enforcing the ℓ_1 -norm constraint and implemented using the SPAMS toolbox (Mairal et al., 2010, Mairal et al., 2009):

$$\min_a \|y - Ea\|_2^2 \quad \text{s.t.} \quad a_i \geq 0, \|a\|_1 \leq \lambda. \quad (5)$$

Non-linear unmixing algorithms are not considered here, as: topographical effects are negligible in DLR HySU; materials are not typically characterized by multiple scattering effects; and higher objects (e.g. building, trees), which could indirectly contribute to illumination received by targets through double scattering, are largely absent in the surrounding area. Furthermore, errors achieved for area estimations estimated with linear models in (Cerra et al., 2021) are considerably low to start with (around 2% for the case of the largest targets, increasing to around 20% for the case of the smallest targets).

3. Mixing Scenarios and Results

This section presents results obtained by the tested abundance estimation algorithms for different signal-to-noise ratios (SNRs) and considering a spectral library “contaminated” by the addition of non-relevant endmembers. Furthermore, we evaluate the effects of enforcing cardinality constraints as a way to approximate the ℓ_0 -norm problem, and of quantizing abundance values to remove non-relevant contributions.

3.1 Performance on Noisy Data

Results of NNLS, FCLS, LASSO and other abundance estimation algorithms on the dataset are reported in (Cerra et al., 2021). Therein, FCLS yields in general the most accurate results, but leads to large errors in the estimation of the area of smaller targets, which are highly mixed. In this section, we explore the behaviour of least squares and sparse unmixing methods in presence of strong noise corruption. Additive Gaussian noise was applied to simulate SNR levels in the range 1–25 dB. Each algorithm was applied to the noisy data to estimate abundance maps. At low SNR (≤ 10 dB), all algorithms show significant errors, especially for small (0.25–0.5 m) patches where pixels are highly mixed. With increasing SNR, the area errors

rapidly decrease. For $\text{SNR} \geq 25$ dB, NNLS, FCLS, and LASSO converge to similar accuracies, confirming robustness consistent with (Cerra et al., 2021, Bioucas-Dias et al., 2012).

Figure 2 summarizes the mean RAE across target sizes and algorithms. Larger patches yield sub-5% errors at high SNR, while small ones are noise-dominated. In general, FCLS yields the best results out of the four algorithms, although it struggles (as observed in (Cerra et al., 2021)) in scenarios with highly mixed pixels, and for high SNR values yields slightly inferior results with respect to sparse methods (LASSO) and NNLS. Furthermore, we can observe that setting the λ parameter in LASSO to 1.0 with respect to 1.2 results in lesser distortion for very noisy images on small targets, but slightly larger ones on large targets.

In order to assess the impact of any specific material on the results, errors per materials averaged over all target sizes (from 0.25 to 3 m) are reported in Fig. 3. By observing the different ranges of RAE, it is obvious that some materials result more difficult to correctly estimate. For example bitumen, which is characterized by a rather flat and featureless spectrum, yields very high errors, especially for low SNR values. Also green fabric exhibits in general higher reconstruction errors, probably as its absorption features are in general shallower, and may be partially correlated to other endmembers (the confusion of this endmember with grass is discussed in the next paragraph). Red fabric and red metal sheets, being characterized by stronger absorption features, yield the smaller errors. Another aspect of interest is the different performance of the algorithms on the different materials: while NNLS yields the worst results in general for bitumen and green fabric at low SNR values, for red metal sheets all algorithms behave similarly, while sparse algorithms yield the worst results on blue fabric.

To give an idea about the quality of the reconstruction results, we report in Fig. 4 abundance maps for the 6 materials as estimated by NNLS for values of SNR ranging from 5 to 25, in steps of 5, and with no noise added. The six abundance maps are here combined in a single RGB representation, which is made possible by the spatial disjunction of pairs of materials in the scene. The abundances (stretched from 0 to 1) are color-coded as follows: red metal and red fabric in red, green fabric and grass in green, bitumen and blue fabric in blue. The grass abundance was rescaled from 0 to 0.3 in order not to hide the boundary with green fabric. Even in presence of strong noise corruption, being the endmembers in the library not noisy themselves, results look visually acceptable already starting from $\text{SNR} = 15$. At $\text{SNR} = 25$ the difference with the results on the dataset with no noise added appear small. At low SNR values, the confusion between green fabric and grass is evident. As shown in Fig. 1, apart from a peak at green wavelengths, green fabric has also an increase in the near infrared, which could resemble stressed vegetation. At very low SNR values, it could be assumed that the featureless bitumen endmember is used by the algorithm to match the overall brightness of the spectra and specific spectral features, compensating noise.

3.2 Expanded Spectral Library

In order to assess the behaviour of the considered algorithms when a higher number of spectra is available in the input library for the abundance estimation step, we expanded the library from 6 to 18 materials, where the additional 12 are unrelated to the targets of interest. The number is comparable to the dimensionality estimation output of the HySIME algorithm (P Nascimento and Bioucas-Dias, 2007) reported in (Cerra et al., 2021),

and spectra were manually selected from the full HySpex subset of DLR HySU. These include asphalt, trees, stressed grass, mixed pixels including the white reference panel, metal, and other artificial objects, selected to represent realistic entries that may occur in airborne scenes. The aim of this experiment is to verify how robust the abundance estimation algorithms are against the extensiveness of the spectral library, which could be exploited to better fit a given spectrum under the observed constraints by the abundance estimation algorithms under test.

Results reported in Fig. 5 for an 18 endmember library expand on the previous findings in Fig. 2 for a 6 endmember library. While FCLS yields in general the best reconstructions, it is here more obvious how highly mixed pixels become more difficult to handle for this algorithm. The additional challenge of having available several endmembers to employ in order to enforce the sum-to-one constraint results in higher distortions for the 50 cm and 25 cm targets, also at higher SNR values.

3.3 Approximating ℓ_0 -norm with Cardinality Constraints

Imposing a constraint on the ℓ_0 -norm in spectral unmixing has been shown to improve the solution for the abundance estimation step (Ben Mhenni et al., 2018). Nevertheless, being the problem NP-hard, this is usually relaxed by imposing constraints on the ℓ_1 - or ℓ_2 -norms, or by finding different approximations, such as the one based on the introduction of an arctan penalty proposed in (Tang et al., 2016). Cardinality constraints are discussed in (Ben Mhenni et al., 2018). Very recent works at the time of writing either use Branch-and-Bound methods (Latif et al., 2025), or present to the user a set of best solutions with fixed cardinality k for the sparse solution (Foix-Colonier et al., 2025). The mentioned methods show promising results on synthetic data, but to the best of the authors' knowledge none has been applied to real data.

In this section we consider the enlarged library of 18 spectra, of which 12 are non-relevant, and perform FCLS unmixing using all subsets of fixed cardinality k , with k ranging from 1 to 6. The subset yielding the smallest mean residual vector was chosen as the optimal reduced library. The sparse reconstruction problem can then be expressed as:

$$\min_a \|y - Ea\|_2^2 \quad \text{s.t.} \quad a_i \in [0, 1], \|a\|_1 = 1, \|a\|_0 \leq k. \quad (6)$$

Results are reported in Tab. 1. As a help to interpret the results, Fig. 6 adapted from (Cerra et al., 2021) is reported, which shows the expected number of endmembers contributing to the different targets for every image element. Therein, black corresponds to zero contributions from endmembers related to the five targets, while colors ranging from dark red to white represent one to five targets present in a pixel, respectively.

The table shows that enforcing extremely low cardinality ($k \in \{1, 2\}$) has low impact on the most dominant endmembers, as the majority of the pixels are pure and distortions should be limited to missing non-dominant materials in a mixed pixel. As expected, for smaller targets with highly mixed pixels the RAE is larger (up to 100%), as usually the "grass" endmember is dominating most of the image elements and is likely chosen by the algorithm. Furthermore, the total number of non-zero elements (NZE) in the area of interest is, as expected, at its minimum. As the fixed cardinality k increases, results in terms of average RAE improve, roughly matching FCLS results for $k = 5$.




Figure 2. Average RAE across target sizes and algorithms as a function of SNR, when using the original 6 endmember spectral library.




Figure 3. Average RAE across target materials and algorithms as a function of SNR, when using the original 6 endmember spectral library.




Figure 4. Abundance maps estimated by NNLS for different values of SNR, ranging from 5 to 25 in steps of 5 (left to right), plus the case with no noise added (right). The six abundance maps are color-coded as follows and stretched from 0 to 1: red metal sheets and red fabric in red, bitumen and blue fabric in blue, green fabric and grass in green. Grass abundances are multiplied by 0.3 not to hide their edge with green fabric.




Figure 5. Average RAE across target sizes and algorithms as a function of SNR, when using as input a library containing 18 spectra.

Table 1. Average RAE across target sizes for different constraints on cardinality of endmembers using FCLS and an input spectral library of 18 spectra. Average RAE across all targets and number of nonzero-elements (NZE), inversely correlated to sparsity, are also reported. The last row shows the benchmark case (FCLS with input library of 18 spectra).

Cardinality k	RAE [%]					Average	NZE
	3 m	2 m	1 m	0.5 m	0.25 m		
1	4.98	11.29	40.99	98.04	100	51.06	309
2	4.96	6.84	14.25	28.77	100	31.03	609
3	3.75	4.04	12.37	31.03	75.03	25.24	885
4	3.19	6.27	11.18	10.32	81.63	22.52	1123
5	3.81	5.76	15.02	12.86	67.47	20.98	1324
6	4.30	6.80	14.61	16.35	37.17	15.85	1479
18 (benchmark)	3.92	6.68	13.65	17.59	59.41	20.25	1594




Figure 6. Expected degree of mixture (as derived from NNLS unmixing) across all targets for the DLR HySU dataset.

Finally, setting $k = 6$ outperforms FCLS applied without constraints. This suggests that this step implicitly removes endmembers that, while rarely dominating, introduce distortions in mixed pixels. Considering NZE values, results indicate that an increase in sparsity of around 10% in the solution (as changing from benchmark solution to $k = 6$) is decreasing the average error by almost 20%, with the smallest 0.25 m targets (highly mixed) benefiting the most from this approach.

Although results are not as clear as in the synthetic experiments reported in (Latif et al., 2025, Ben Mhenni et al., 2018, Foix-Colonier et al., 2025), cardinality constraints seem to deliver promising results also when applied to real data.

3.4 Abundance quantization

An alternative way to induce sparsity without ℓ_1 or ℓ_2 constraints is to quantize the possible abundance values in NNLS. In practice, this amounts to write each abundance value between 0 and a given maximum value $a_{\max}$ (typically slightly above 1 to allow for noise) as

$$a = a_{\max} \sum_{q=0}^{N_b-1} 2^q b_q, \quad (7)$$

where N_b is the number of bits and $b_q \in \{0, 1\}$. With quantized abundances, NNLS becomes a quadratic unconstrained binary optimization (QUBO) that can be solved either classically or approximately using quantum solvers. To illustrate the potential of this approach, we solved the quantized NNLS problem for the DLR HySU dataset without added noise and the original six endmember library running the QUBO quantum sampler implemented in the quark package (Lobe and Windgätter, 2025) in a conventional computer. Although this sampler does not run on a real quantum processor, it reproduces the search behavior based on quantum annealing, which relies on quantized formulations. Therefore, the reported experiment considers a possible future direction rather than current operational applicability of quantum approaches to unmixing problems.

The abundance estimation errors using $a_{\max} = 1.2$, 25000 samples and 6, 7 or 8 bits are shown in Tab. 2 together with the plain NNLS results. As expected, quantization gradually induces sparsity as the number of bits is reduced. The average errors for 7 and 8 bits are around 10% better than the stand-

Table 2. Average RAE across target sizes for different number of bits used in the quantized NNLS problem with no added noise and an input spectral library with 6 endmembers. Average RAE across all targets and number of nonzero-elements (NZE), inversely correlated to sparsity, are also reported. The last row shows the benchmark case (plain NNLS).

Number of bits N_b	RAE [%]						NZE
	3 m	2 m	1 m	0.5 m	0.25 m	Average	
6	3.66	3.01	14.57	21.71	24.98	13.59	825
7	3.65	2.42	11.24	21.96	16.05	11.07	874
8	3.50	2.50	11.89	22.45	16.43	11.35	918
NNLS	3.38	3.08	10.63	21.26	25.24	12.72	940

ard NNLS results, driven by a slight improvement in the 2 m targets and a significant improvement in the highly mixed case of 0.25 m targets. These interesting findings should be further investigated in future work, but it is fair to point out that the modest improvement over NNLS comes at the price of a much longer computational time. Solving quantized NNLS is still manageable in our case, given the small size of the problem (6 endmembers, 6–8 bits, hundreds of pixels). The computational resources needed for significantly larger problems (e.g., tens or hundreds of endmembers as in standard libraries) grow exponentially and quickly become prohibitive in normal circumstances. However, this approach could become feasible in the future, as capabilities of quantum computing are expected to dramatically improve.

4. Conclusions

We presented a physically traceable validation of hyperspectral unmixing using the DLR HySU benchmark based on real imaging spectrometer data. Area-based comparison of estimated and true target sizes enables objective accuracy assessment under varying noise conditions and library configurations. Classical algorithms (NNLS, FCLS, LASSO) converge for $\text{SNR} \geq 25$ dB, confirming the reliability of the DLR HySU benchmark. Errors quantified after adding to the input spectral library non-relevant endmembers generally favor FCLS. Nevertheless, NNLS provides a robust baseline across scenarios, while FCLS yields the most accurate results in moderately mixed conditions but degrades for highly mixed pixels. This should be attributable to the sum-to-one constraint, which becomes increasingly restrictive when several endmembers contribute to a single pixel, which is particularly clear when using the extended spectral library. Sparse methods represented by LASSO show competitive performance when library size increases or noise levels are high. Subsequently, we consider hard constraints on cardinality as an approximation of the ℓ_0 -norm, which yields a slight improvement in the results, as observed (with superior results) by other works relying on synthetic data (Foix-Colonier et al., 2025). Finally, we consider quantization of abundances as a different approach to enforce sparsity, evaluating the consequent improvement in the results: this leads to considerations on future capabilities of quantum computers to tackle the unmixing problem, where potential alternative methods related to sparsity enforcement could be efficiently implemented. Although this study relies on airborne hyperspectral data, the investigated aspects are relevant for spaceborne hyperspectral missions such as EnMAP, PRISMA, and DESIS. Here, noise corruption, expanded or overcomplete spectral libraries, and especially highly mixed pixels need to be considered in most of practical applications employing these data.

Future work will explore additional abundance estimation algorithms in the same validation framework, consider additional

recently defined benchmark dataset with true endmember abundances, and possibly further establish formulations of this problem to be run in a quantum simulator.

References

- Ben Mhenni, R., Bourguignon, S., Ninin, J., Schmidt, F., 2018. Spectral Unmixing with Sparsity and Structuring Constraints. *IEEE Whispers (2018) : 9th Workshop on Hyperspectral Image and Signal Processing*, Amsterdam, Netherlands.
- Bioucas-Dias, J. M., Plaza, A., Dobigeon, N. et al., 2012. 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.
- Cerra, D., Pato, M., Alonso, K., Köhler, C., Schneider, M., de los Reyes, R., Carmona, E., Richter, R., Kurz, F., Reinartz, P., Müller, R., 2021. DLR HySU—A Benchmark Dataset for Spectral Unmixing. *Remote Sensing*, 13(13). <https://www.mdpi.com/2072-4292/13/13/2559>.
- Eches, O., Dobigeon, N., Altmann, Y., Tournier, J.-Y., 2010. Bayesian Estimation of Linear Mixtures Using the Normal Compositional Model. *IEEE Transactions on Geoscience and Remote Sensing*, 48(11), 3966–3978.
- Foix-Colonier, N., Bourguignon, S., Schmidt, F., Latif, M., 2025. Guaranteed best sparse solutions for spectral unmixing. *Proceedings of the 13th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS 2025)*, IEEE, Barcelona, Spain. 12–14 November 2025.
- Haijen, X., Koirala, B., Tao, X., Scheunders, P., 2025. A benchmark linear unmixing dataset with spectral variability and ground truth. *Proceedings of the 13th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS 2025)*, IEEE, Barcelona, Spain. 12–14 November 2025.
- Iordache, M. D., Bioucas-Dias, J. M., Plaza, A., 2011. Sparse Unmixing of Hyperspectral Data. *IEEE Transactions on Geoscience and Remote Sensing*, 49(6), 2014–2039.
- Iordache, M.-D., Bioucas-Dias, J. M., Plaza, A., 2012. Total Variation Spatial Regularization for Sparse Hyperspectral Unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 50(11), 4484–4502.
- Iordache, M.-D., Bioucas-Dias, J. M., Plaza, A., 2014. Collaborative sparse regression for hyperspectral unmixing. *IEEE Trans. Geosci. Remote Sens.*, 52(1), 341–354.

Latif, M., Foix-Colonier, N., Bourguignon, S., Schmidt, F., 2025. Branch-and-bound algorithm for exact ℓ_0 -norm sparse spectral unmixing. *Proceedings of the 33rd European Signal Processing Conference (EUSIPCO 2025)*, EURASIP, Lyon, France. To appear.

Li, C., Liu, Y., Cheng, J., et al., 2017. Sparse Unmixing of Hyperspectral Data with Noise Level Estimate. *Remote Sens.*, 9(11), 1166.

Lobe, E., Windgätter, L., 2025. quark - quantum application reformulation kernel.

Mairal, J., Bach, F., Ponce, J., Sapiro, G., 2009. Online dictionary learning for sparse coding. International Conference on Machine Learning.

Mairal, J., Bach, F. R., Ponce, J., Sapiro, G., 2010. Online Learning for Matrix Factorization and Sparse Coding. *J. Mach. Learn. Res.*, 11, 19-60.

P Nascimento, J. M., Bioucas-Dias, J. M., 2007. Hyperspectral signal subspace estimation. *2007 IEEE International Geoscience and Remote Sensing Symposium*, 3225–3228.

Tang, W., Shi, Z., Borjigen, D., 2016. Sparse Hyperspectral Unmixing Using an Approximate ℓ_0 Norm. *Remote Sensing*, 8(3), 187.

Zhu, F., 2017. Hyperspectral unmixing: ground truth labeling, datasets, benchmark performances and survey. *arXiv preprint arXiv:1708.05125*.