

Optimal band selection for target detection with a LWIR multispectral imager

Michael Zelinski,* Andrew Mastin, Vic Castillo, and Brian Yoxall

Lawrence Livermore National Laboratory, Livermore, California, United States

Abstract. Multispectral imaging can offer many benefits in cost, complexity, resolution, size, weight, and power, relative to hyperspectral imaging. When designing a multispectral system, spectral bandpasses can be selected using optimization algorithms configured to maximally separate target detection scores between target and background regions. A hyperspectral image (HSI) can serve as the source of data from which band groupings can be tested for optimality. The output of an adaptive cosine estimator target detection algorithm is used in an objective function. Three optimization algorithms are compared: particle swarm, dual annealing, and differential evolution. A global optimum is also found using a brute force approach on the Livermore Computing Syrah supercomputer. Three materials are investigated: calcite, gypsum, and limestone. This is done for 3-, 4-, and 5-band systems. The data originate from a longwave infrared HSI of a material display board. The optimization algorithms were run 30 times for every scenario. Performance statistics (maximum, minimum, mean, standard deviation, and median) based on the separation values are given. Additional characterization was performed using receiver operator characteristic (ROC) curves and the area under the ROC curve. While good performance was obtained for the three optimization algorithms, the dual annealing algorithm produced the highest and most consistent detection separation scores on average. © *The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.* [DOI: [10.1117/1.JRS.16.026505](https://doi.org/10.1117/1.JRS.16.026505)]

Keywords: multispectral; hyperspectral; longwave infrared; spectral filter design; spectral band selection; target detection; numerical optimization; supercomputing.

Paper 210551 received Aug. 30, 2021; accepted for publication Mar. 7, 2022; published online Apr. 14, 2022.

1 Introduction

Spectral imaging enables material classification based on spectral reflectance properties. Many materials, natural or manmade, have unique spectral properties that act as a fingerprint providing a link to the material's atomic or morphological state. When designing an imaging system, a very cost-effective way of introducing additional utility is to include a spectral imaging capability; in many cases this can be done without changing the size of the primary optic—often a driver of an image system's size, weight, and cost. Hyperspectral imaging systems can be designed to include several hundred contiguous spectral bands. However, acquiring this number of bands usually requires sophisticated optical components that are relatively fragile, complicated, and expensive. In addition, the optics might require additional cooling which increases the overall size, weight, and power (SWAP). These systems are often built with a low $F/\#$ to gather more photons; however, this is done as a tradeoff resulting in reduced spatial resolution.^{1,2}

Multispectral systems are often built with the intention of balancing system utility with constraints in SWAP. During the design process, decisions must be made about which spectral bandpasses to use. In the past, this was done by a subject matter expert, who used the known spectral signatures of intended targets and backgrounds, while also considering the impact of spectral atmospheric absorption and detector sensitivity.³

Over time, more sensors have been fielded, and researchers began using the newly acquired data and new experimental results to dictate the selection of bands in the next generation of designs. The advanced baseline imager aboard the GOES-R satellites has 16 bands that were

*Address all correspondence to Michael Zelinski, zelinski1@llnl.gov

derived from MODIS bands and other spaceborne systems.⁴ More recently, there has been an effort to use hyperspectral images (HSI) as scene truth and select corresponding bands from the image to simulate the performance of planned sensors.⁵ This idea was further advanced using an optimization algorithm to automatically select bands from the HSI. This is done by establishing a cost function around a given task, such as class separability,⁶ clustering,⁷ target detection and compressive sensing error minimization,⁸ and classification.⁹ This technique is used both by hyperspectral researchers to find out which bands carry high information content^{6,7} and by multispectral sensor designers to determine which bandpasses should be included in the system.^{8,9} The optimization algorithms previously employed for this task are: particle swarm optimization,⁶ dynamic programming,⁷ and differential evolution.⁸ Additional literature review is provided in Refs. 7 and 8. Currently, there has not been an attempt to apply these techniques to the long-wave infrared (LWIR); there has been no comparison between optimizing for a single material or task versus multiple materials or tasks; there are no references showing comparisons between different optimization algorithms; and there has been no work comparing the optimized results to known global optimum.

The method employed in this paper makes use of an LWIR HSI of a material display board. Groups of bands corresponding to multispectral bandpass filters are selected from the HSI image. These band groupings are averaged to form the bands of a multispectral image that are used with the adaptive cosine estimator (ACE) target detection algorithm.² This process is repeated, and performance is ranked against other multispectral band groupings synthesized from the same HSI image. Performance is monitored by comparing the target detection scores for pixels found over a target area to those over the background area, this what is referred to as “separation score.” Band selection is done using three methods: (1) a brute force approach exploring every band combination to find a global optimum set of bands, (2) a numerical optimization approach, and (3) there is one example of a user-defined set of bandpasses in Sec. 5 that is provided for demonstration purposes. The brute force approach is performed on LLNL’s Syrah supercomputer. Three numerical optimization approaches are examined: integer particle swarm, differential evolution, and dual annealing. An overview of this approach is provided in the flow-chart shown in Fig. 1.

The materials of interest are calcite, alabaster, and travertine. The materials are optimized for simultaneously, meaning that the separation score for the three materials is used in the cost function. The materials are also optimized for individually—each one is used alone in the cost function.

Section 2 focuses on the data processing. It includes a description of the data, band averaging, the normalization step used, the ACE algorithm, the separation score, and the parameterization used for optimization. Section 3 describes the brute force processing used to obtain a global optimum separation score. Section 4 describes the optimization algorithms. Section 5 provides

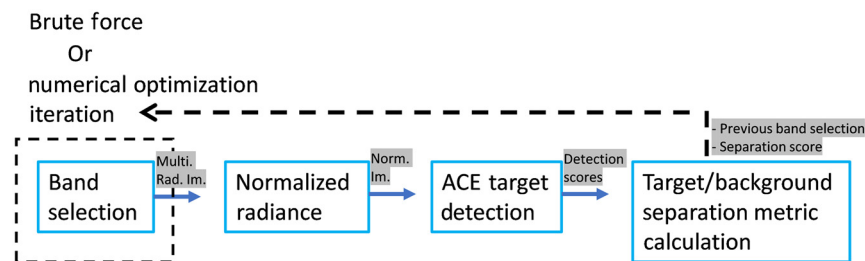


Fig. 1 Groups of bands are selected from an HSI and summed to form multispectral bands. The multispectral bands are then normalized in a step that uses the pixel’s brightness temperature. The normalized bands are then used with the ACE target detection algorithm. A target/background separation score is used to capture the utility of the selected group of bands. There are two primary options for band selection: the first is a brute force approach where every band combination is tested; the second is a numerical optimization approach where the optimization algorithm selects groups of bands with the objective of increasing the separation between target and background ACE scores. The gray boxes indicate which arrays are passed on to subsequent steps. Note, “Multi. Rad. Im.” stands for multispectral radiance image. And “Norm. Im.” stands for normalized image.

a description of the results. Section 6 contains a conclusion. Section 7 is an Appendix providing all the results for each of the optimization algorithms, materials, and band numbers.

The primary contributions of this paper are: (1) a demonstration of a multispectral band optimization applied to an LWIR hyperspectral dataset. (2) A comparison between individual optimization and combined optimization for calcite, gypsum, and limestone. (3) A comparison of three integer optimization algorithms applied to multispectral band selection. (4) A comparison of optimized bands to the global optimal solution for three- and four-band cases.

2 Hyperspectral Data Processing

2.1 Data

The baseline image is a 49-band LWIR hyperspectral radiance image spanning 7.93 to 12.56 μm window. The material display board was imaged under a cold sky. The exact geometry of the setup is not known; the board was angled at 45 deg relative to the surface normal. The viewing geometry was slightly downward looking at an approximate 10-deg angle from a high rise building at an unknown distance. The sensor had a slight focusing issue during the data collection. Readers can see a spectral radiance plot of calcite taken from this dataset in Fig. 2 along with a three-channel image of the material display board.

The measured radiance is affected by the material spectral emissivity, bidirectional reflectance distribution, geometry, and temperature. Atmospheric radiance and transmission also affect the measured radiance; these factors are influenced by the atmospheric gas' transmission, pressure, number density, concentration, and temperature.

In this work, an assumption is made that all materials are diffuse reflectors. The Lambertian (diffuse) LWIR radiance model² is defined as

$$L = B(T)\epsilon\tau + (1 - \epsilon)L_d\tau + L_u, \quad (1)$$

where $B(T)$ is a Planckian blackbody at temperature T , ϵ is the material emissivity, τ is the atmospheric transmission, L_d is the incident radiance from the surrounding environment (downwelling radiance), L_u is the self-emitted radiance between the target and the detector. Each of these terms includes a wavelength dependence which has been omitted for simplicity.




Fig. 2 (a) A calcite emissivity spectrum is plotted with data from a scene pixel (in radiance) that contains calcite (see red box in Fig. (b)). (b) An RGB image was created from three hyperspectral bands to give the reader context of the material display board. The squares are blocks of various materials. The warm and emissive materials show up as bright white squares; metallic materials are reflecting cold sky and appear as black squares; and the bottom row shows several purple squares, these materials have spectral features located over the three hyperspectral bands used to make this image.




Fig. 3 The target material spectra used in this study have a diverse set of spectral features of varying widths and depths.

The target spectra came from the Arizona State University Spectral Library and the Jet Propulsion Lab EcoStress Spectral Library (the original measurement was made by USGS). A plot of the target spectra is shown in Fig. 3:

- Calcite: CALCITE C12¹⁰
- Alabaster: GYPSUM S5¹⁰
- Travertine: rock.sedimentary.limestone.solid.all.ro337.usgs.perknic.spectrum¹¹

2.2 Band Selection and Multispectral Band Creation

Groups of contiguous bands are selected from the hyperspectral radiance cube. The groups are selected such that the edge of the adjacent group does not overlap. Each group is then averaged spectrally. The resulting set is intended to simulate data collected by a multispectral system. Figure 4 provides an example of a four-band case.

It should be noted that detector noise characteristics for the multispectral system were not simulated.

2.3 Data Normalization by Brightness Temperature

To use the multispectral data in a target detection algorithm, a normalization technique was used as an initial processing step to help reduce variability caused by the material temperature. This is

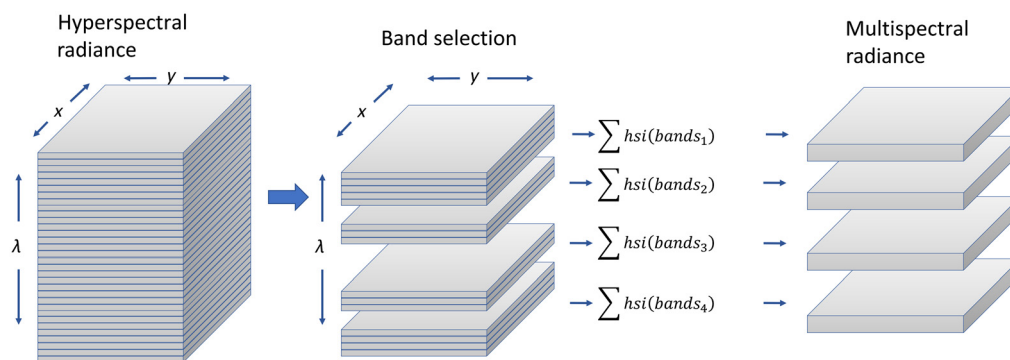


Fig. 4 The original dataset is a full hyperspectral cube, groups of contiguous bands are then selected from this cube, and these groups are averaged to create the individual multispectral bands. “ N_x ” defines the number of HSI bands in the group.

done by first computing the brightness temperature (T_{\max} [K]) of every multispectral pixel radiance vector $im(p)$:

$$T(p) = \frac{hc}{\lambda k \log((2 hc^2 \lambda^{-5}) / (im(p) * 1E4) + 1)}, \quad T(p)_{\max} = \max(T(p)), \quad (2)$$

where $h = 6.62606896e-34$ (J * s), $c = 299,792,458$ (m/s), $k = 1.3806504e-23$ (J * K⁻¹), λ is the wavelength band center in microns, and $im(p)$ has units of microflicks [$\mu\text{W}/(\text{cm}^2 \text{sr } \mu\text{m})$].

The brightness temperature is then used to normalize the data by dividing each pixel by a Planck curve defined by the brightness temperature:

$$x = \frac{im(p)}{10^{-4} * (2hc^2 \lambda^{-5}) * (e^{hc/\lambda k T_{\max}} - 1)^{-1}}. \quad (3)$$

This approach works well for objects that are emitting more radiance than they are reflecting—as a sidenote, proportionally larger amounts of reflected atmospheric water and ozone prevent the selection of an accurate maximum brightness temperature. x will be used in the ACE target detection step. The reader may notice that this equation bears some resemblance to the emissivity equation.²

2.4 ACE

The ACE algorithm is a popular hyperspectral target detection algorithm.^{2,6,12} The algorithm uses the scene data multivariate mean and covariance in a process referred to as whitening to suppress the data's direction of maximum variability. The ACE algorithm whitens and normalizes the pixel vectors (x) and target vector(s) (t) and then computes cosine angle between the two. The target spectra are defined in Sec. 2.1. ACE is defined as

$$\begin{aligned} & \frac{(x - \hat{x})^T \Sigma^{-1} (t - \hat{x})}{\sqrt{(x - \hat{x})^T \Sigma^{-1} (x - \hat{x})} \sqrt{(t - \hat{x})^T \Sigma^{-1} (t - \hat{x})}} \\ &= \left(\frac{\Sigma^{-\frac{1}{2}} (x - \hat{x})}{\sqrt{(x - \hat{x})^T \Sigma^{-1} (x - \hat{x})}} \right)^T \left(\frac{\Sigma^{-\frac{1}{2}} (t - \hat{x})}{\sqrt{(t - \hat{x})^T \Sigma^{-1} (t - \hat{x})}} \right), \end{aligned} \quad (4)$$

where \hat{x} is a spectral mean, Σ is the spectral covariance matrix, and ^T denotes a matrix transpose. The authors refer to ACE score as the value computed by Eq. (4). This value is computed for every pixel in the scene.

2.5 Separation Score

Ideally, the target detection algorithm will have a good rate of true positive detection and a low rate of false positive detection. While it is not exactly analogous, if the algorithm is performing well, the output scores over target areas will be higher than those over background areas. The assumption used here is that a greater separation in ACE scores between target and background will result in an improvement in detection performance (either a decrease in false positive rate or an increase in true positive rate).

As mentioned in Sec. 1.1, the sensor experienced a spatial focusing issue which resulted in some pixels being mixed with both target and background. In this work, the number of on-target pixels was much less than the number of background pixels. Figure 5(a) shows a detection plane for the calcite target. Plotting a histogram of the ACE scores for on-target and background regions, we see these two groups overlap where the pixels are mixed with either target or background [see Fig. 5(b)].

There is a relatively small number of target pixels and many of the ACE scores in the target-region are statistical outliers. For this reason, a median is used to characterize the target-region ACE scores. The background region, which contains a much greater number of pixels, was characterized by taking a mean value for this region. Capturing the amount of separability between

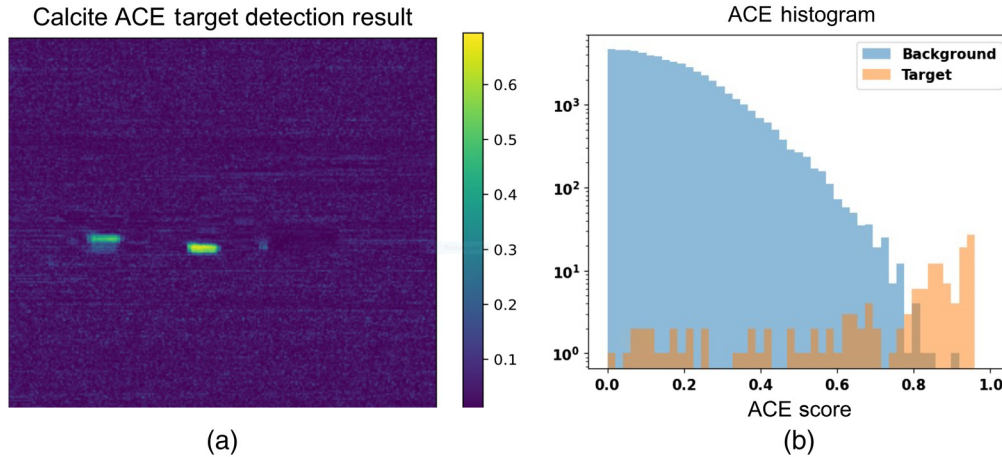


Fig. 5 (a) Observe the ACE detection plane for the calcite target. The high detection values (bright yellow square) span the target region. The blue/green region is entirely background. (b) The histogram of ACE scores for the target area and background area shows perfect separation between the pure target and background areas. Pixels that are mixed with target and background overlap either class. The y-axis uses a log scale to account for the significant imbalance between the pixel counts of the two classes.

the target and background is done by differencing the median of the target area ACE scores and the mean of the background ACE scores:

$$\text{Score} = \text{median}(\text{ACE}_{\text{target}}) - \text{mean}(\text{ACE}_{\text{background}}). \quad (5)$$

2.6 Parameterization

Parameterizing this problem is done by defining the lower and upper bounds of each multispectral bandpass. An illustration of this is provided for the four-band case in Fig. 6. The multispectral system’s bandpasses can be defined by a vector $Z = (z_1, z_2, \dots, z_N)$, which is a vector of hyperspectral band indices. N.B. the N defined here is $2 \times$ the number of multispectral bands.

There are three constraints that are used to define this problem. **Constraint 1:** Vector $Z \in \mathbb{Z}^+$, that is, it is an integer vector since it contains hyperspectral band indices. **Constraint 2:** Vector $Z_n \in [1, 49]$, the HSI used here has 49 bands and therefore the band indices (n) must be between 1 and 49. **Constraint 3:** $z_1 < z_2, z_2 \leq z_3, \dots, z_{N-1} < z_N$, the multispectral bandpasses must have a width of at least one band and they are allowed to be contiguous. Parameterizing Eq. (5), we can define the separation Score for a given scenario, Z , as:

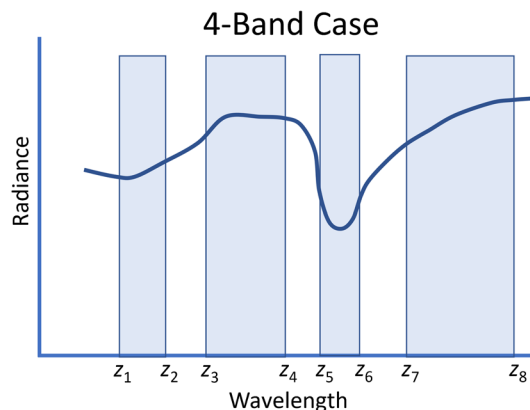


Fig. 6 An illustration of the four-band case. The blue rectangles represent the bandpass filters. The bandpass filters are constrained to have some width, potentially contiguous but not overlap.

$$\text{Score}(Z) = \text{median}(\text{ACE}_{\text{target}}(Z)) - \text{mean}(\text{ACE}_{\text{background}}(Z)). \quad (6)$$

Equation (6) applies directly to the single material case. When there is more than one material, the median of the target ACE scores and the mean of the background ACE scores is computed for each material. Then, the mean is taken for the targets and backgrounds, this value is then differenced. For a scenario with M materials:

$$\text{Score}(Z) = \sum_m^M (\text{median}(\text{ACE}_{\text{target},m}(Z)))/M - \sum_m^M (\text{mean}(\text{ACE}_{\text{background},m}(Z)))/M. \quad (7)$$

3 Brute Force Data Processing

The number of multispectral configurations for the three-, four-, and five-band cases is given as three-band 1.8×10^7 , four-band 7.5×10^8 , and five-band 1.9×10^{10} . Computing a brute force solution beyond four bands is not currently possible in a reasonable amount of time using the resources at LLNL. Computing all configurations for the four-band case took ~ 1 week using 128 nodes on LLNL's Syrah supercomputer.¹³ The Syrah computer has Intel Xeon E5 processors with 16 cores. Python and Dask were used to run the jobs and log results; 16 jobs were run on each node, with a total of 2048 jobs running simultaneously.

4 Numerical Optimization Algorithms

Three integer programming options are available to guide the selection of optimal band configuration—integer particle swarm optimization, dual annealing, and differential evolution. As addressed in Sec. 5.2, each algorithm is run 30 times for every scenario investigated. These algorithms were run on PCs, the computational requirements are miniscule compared to what was required for the brute force optimization. A single optimization run would generally complete in minutes.

4.1 Integer Particle Swarm Optimization

Particle swarm optimization (PSO) interrogates a problem with a population (or swarm) of individual potential solutions (particles). As the algorithm makes iterative steps toward an optimal solution, it is influenced by the previous global best solution as well the individual particle's best previous solution. Each particle has a velocity that influences its momentum as it traverses the solution space. Having momentum as well as influence from the global best solution helps prevent particles from getting stuck in local minima.

The implementation used in this work closely follows¹⁴ which is similar to continuous PSO; however, the dependent variables are rounded to their nearest integer and the constraints mentioned in Sec. 2.6 are enforced. The optimization was setup to use: 50 particles, 200 iterations ($M = 200$), a cognitive weight (c_1) of 0.6, a social weight (c_2) of 0.6, a velocity constriction factor (χ) of 0.729, inertia weight (ω) of 0.9, and an inertial weight decay function of:

$$\Omega(i) = \omega * \left(e^{-\frac{i}{M/2}} * \sin\left(\frac{3\pi}{2} * \frac{(M-i)}{M}\right) \right) + .001, \quad (8)$$

where i is an iteration index.

4.2 Dual Annealing

Simulated annealing is a well-known heuristic for unconstrained optimization problems. It is an improvement over basic local search methods because it allows the incumbent solution to move in a direction that is not locally improving with a certain probability, thus allowing the solution to escape local optima. Dual annealing is a type of simulated annealing algorithm that combines

classical simulated annealing and fast simulated annealing. Dual annealing is a standard routine in the SciPy Optimize package.^{15,16}

Simulated annealing routines are generally applied to unconstrained optimization problems. Since the problem has some constraints—all bands have lower and upper bounds, bands cannot overlap—the objective function was modified to return a poor score if one of the constraints was violated. In the experiments, the total number of iterations was limited to 300, and the visit parameter was set to 2.9 (instead of the default value of 2.62). This allows the algorithm to jump to greater distances in the search space. The remaining parameters were set to their default values in SciPy Optimize.

4.3 Differential Evolution

Differential evolution is an evolutionary algorithm. Starting with a randomly selected set of candidate solutions, the algorithm iterates through a sequence of mutation, recombination, and selection steps until a termination criterion is met. The algorithm is also included in the SciPy Optimize package.¹⁷

The SciPy Optimize package implementation of differential evolution allows constraints to be specified. We added constraints to limit the upper and lower bounds of the bands and constraints to ensure that the bands do not overlap. We set the maximum number of iterations to 300 and used default values in SciPy Optimize for all other parameters.

5 Results

5.1 Brute Force Processing

The brute force global optimum separation scores and associated bandpasses for the three- and four-band scenarios are provided in Table 1 of the Appendix. The bandpasses are also displayed graphically in Figs. 7 and 8. While the separation score is a useful metric for the optimization algorithm, detection performance is best captured by examining the probability of detection (PD) and the probability of false alarm (PFA). This done using a receiver operator characteristic (ROC) curve. Overall performance can be quantified by taking the area under the ROC (AUROC) curve. Figures 7 and 8 show ROC curves for the three- and four-band cases. Figure 7 shows expected behavior for calcite and gypsum, where the optimal spectral band selection is aligned with dominant spectral features. However, this does not occur with the individual case of limestone or for the three-material case used in the optimization. It is unclear why this happened. In the individual limestone case, the other subtle spectral features (8.5, 9.2, and 9.8 μm) may have offered more information relative to the background spectra than the dominate feature (11 μm). In the three-material case, one material may have gained substantially from the selected band configuration and mask performance gains (or losses) for the other materials. Figure 8 shows more consistent behavior for the four-band case, where bands fall on dominant spectral features. Examining the ROC curve, in both cases, the contiguous multispectral bands perform worse than the hyperspectral system. An exciting result is that a multispectral system with band optimization for three materials can produce better performance than the original hyperspectral system. Optimizing for each individual material results in further improvement for the material under inspection; however, detection performance for the other materials declines substantially. It should be noted that this observation applies to only this dataset. Another dataset with the same target materials, but different background, could result in a different optimal band combination. Future work could examine the extensibility of these bandpasses, comparing these results to results collected over different backgrounds or different times of day (different thermal conditions).

5.2 Numerical Optimization Results

The detailed optimization results are provided in Table 1 of the Appendix. For each scenario, 30 optimization runs were performed. The mean, standard deviation, median, minimum, and maximum values are given for these 30 optimizations. The bandpasses for the maximum separation score are also given.

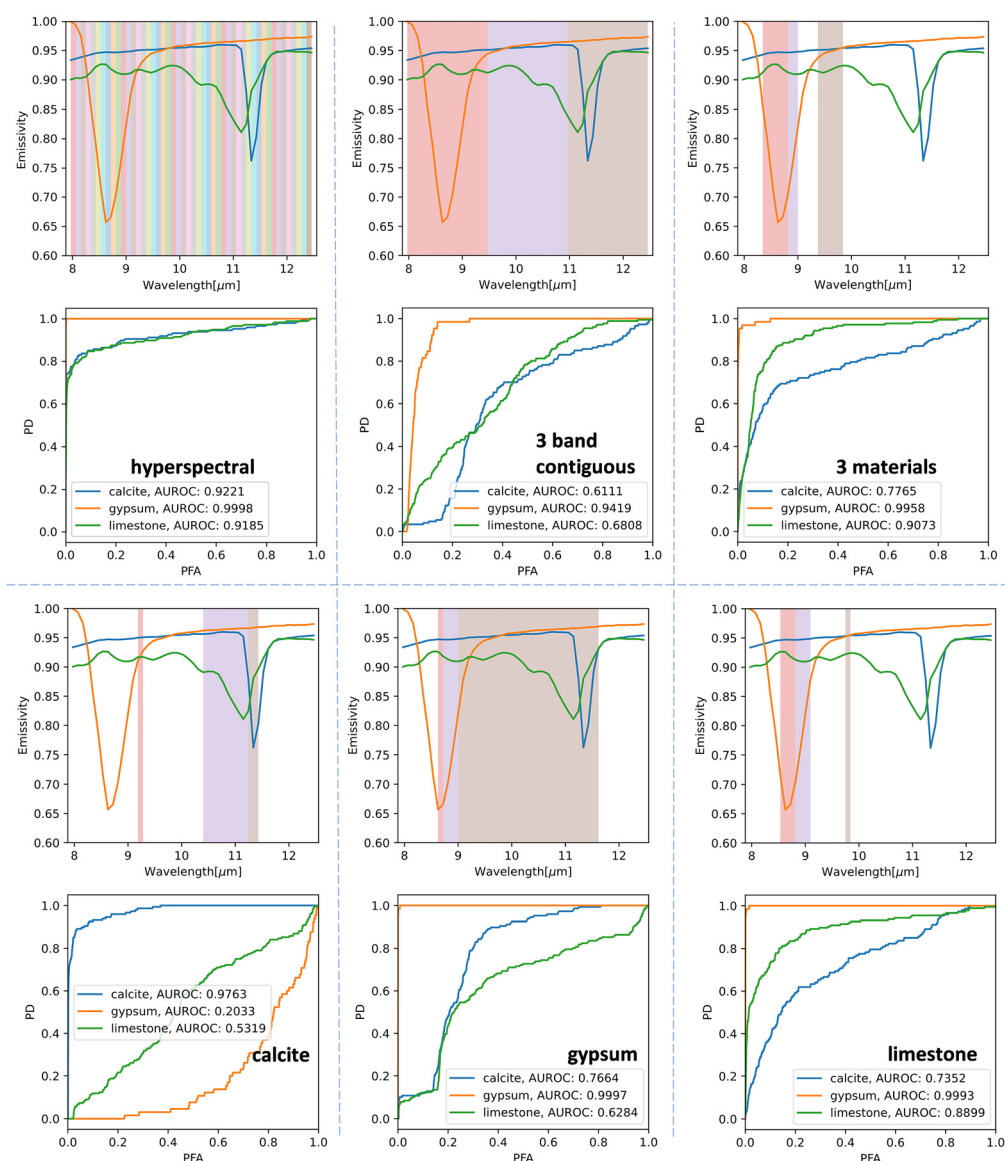


Fig. 7 The brute force result for the three-band case is shown here. Each vertical pair of plots corresponds to an optimization scenario. For each plot pair, the upper plot shows the three material spectra of interest and selected spectral bandpasses. The lower plot shows the ROC curves and AUROC curve values for each of the three materials of interest. This shows a comparison between the original hyperspectral data, three-band contiguous data, and three-band optimized data.

Figure 9 provides a summary of algorithm performance. The mean is calculated for the four material types of a single band number and optimization type. For example, the mean of the three-band particle swarm algorithm 30-run means is $\text{mean}(0.518, 0.576, 0.886, 0.504) = 0.621$. And the mean of the 30-run standard deviations is $\text{mean}(0.031, 0.105, 0.019, 0.028) = 0.046$.

For the three-band case, all three algorithms performed well and found the global maximum for several materials. The particle swarm algorithm found the global maximum for all materials; however, the separation score standard deviation was highest for all materials using this algorithm making it the most inconsistent performer. The differential evolution algorithm produced the best results on average with a higher mean separation score and a lower standard deviation.

In the four-band case, each algorithm found the global maximum in two out of the four material scenarios. The differential evolution algorithm had the highest mean separation score and the lowest mean standard deviation for the 30-run means.

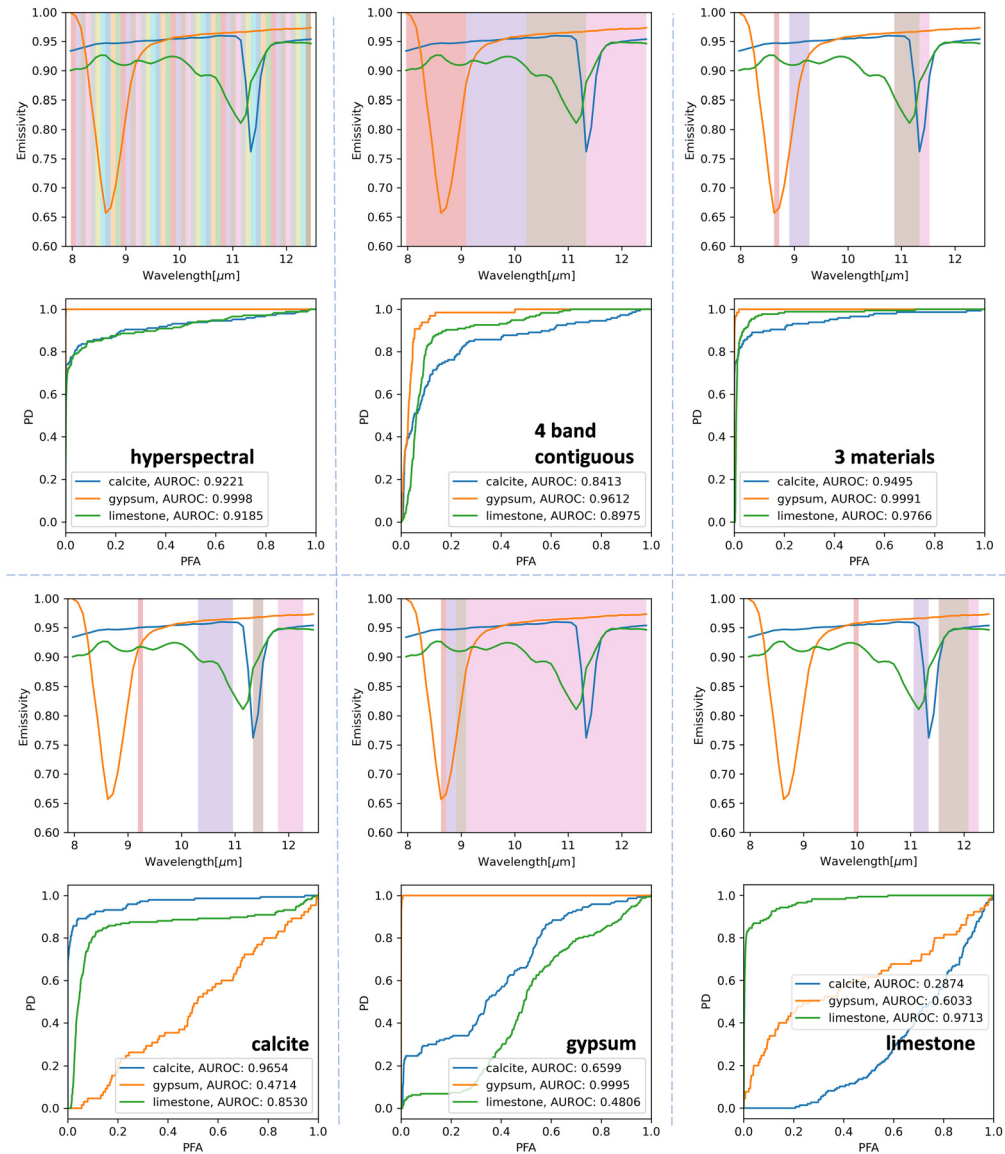


Fig. 8 The brute force results from the four-band case.

The five-band case has many more potential combinations, and it was therefore not possible to compute a brute force result. The particle swarm algorithm produced the best results on average; however, there was large variability in the scores. The dual annealing algorithm produced the most consistent results but only by a small margin compared with the differential evolution. The performance margins here are closer than in the three-band case.

Figure 10 shows the best five-band optimization results for each material. The bandpasses for the individual materials follow the dominant spectral features. However, the three-material optimization lacks a bandpass over the dominant limestone feature. If the bandpasses are modified to include a band over the dominant limestone feature an interesting observation is made (see “User Defined” plot in Fig. 10). The AUROC values for limestone increase while the values of calcite and gypsum both decrease. On average, the AUROC values for the optimized bands are higher than the user defined bands. Three-material AUROC average: $(0.9879 + 0.9907 + 0.9484)/3 = 0.9757$, user defined AUROC average: $(0.9569 + 0.9854 + 0.96168)/3 = 0.9680$. The optimization algorithm would have likely favored improved performance for calcite and gypsum at the expense of the limestone performance.




Fig. 9 Bar chart shows a summary of the separation scores for each algorithm. The bar values are produced by taking the mean value of the mean separation scores for the four material types for each algorithm. The error bars are produced by taking the mean value of the standard deviation of the separation scores. The dual annealing and differential evolution algorithms produced the highest and most consistent results for the three-band case. For the four- and five-band cases, all algorithms are similar. It is worth noting that the particle swarm algorithm has higher variance than the other two algorithms, which is a good thing if the optimization can be run many times—as a higher maximum score could be more attainable with this algorithm.

6 Conclusion

6.1 Brute Force Computing

A brute force search of all possible spectral band combinations was performed for a three- and four-band system using a 49-band HSI as a data source. Executing this task required using 128 nodes on LLNL's Syrah supercomputer and 1 week to complete. Given the successful performance of the optimization algorithms, there is little to be gained by brute force searches of this kind in future work; the resources required are significantly greater than what is used for the numerical optimization approach.

6.2 Numerical Optimization

Integer programming, which is a type of numerical optimization that utilizes integer valued variables, was used to select band combinations from the 49-band HSI. Three algorithms were used in the optimization: particle swarm, dual annealing, and differential evolution. Overall, the performance for all three algorithms is similar. There are many instances where these algorithms found the global optimum solution. The differential evolution algorithm produced high separation scores more consistently than the other algorithms but only by a small margin. The particle swarm algorithm produced a higher standard deviation of separation scores, making it the most inconsistent performer. That said, it did achieve some of the best scores; most notably for the three-band case, where the global maximum was found for all materials.

Plotting the optimized bandpasses against the target vectors revealed that the results were not always intuitive; there are some instances of dominant spectral features not being selected for bandpasses. The most unusual example of this being the three-band limestone optimization where smaller features were selected over the dominant feature. There are several factors that could have caused this, particularly the scene whitening step that is used in ACE which can be heavily influenced by the background. This behavior was also observed in the three-material optimizations for three bands and five bands. When three materials are used in the optimization, it is possible that substantial performance gains of one material might overshadow the decreased performance of the other two.

6.3 User Defined Optimization

In Fig. 10, the authors modified the three-material optimization by moving one band to the dominant (and unused) limestone spectral feature. The band that was moved existed beyond $12 \mu\text{m}$

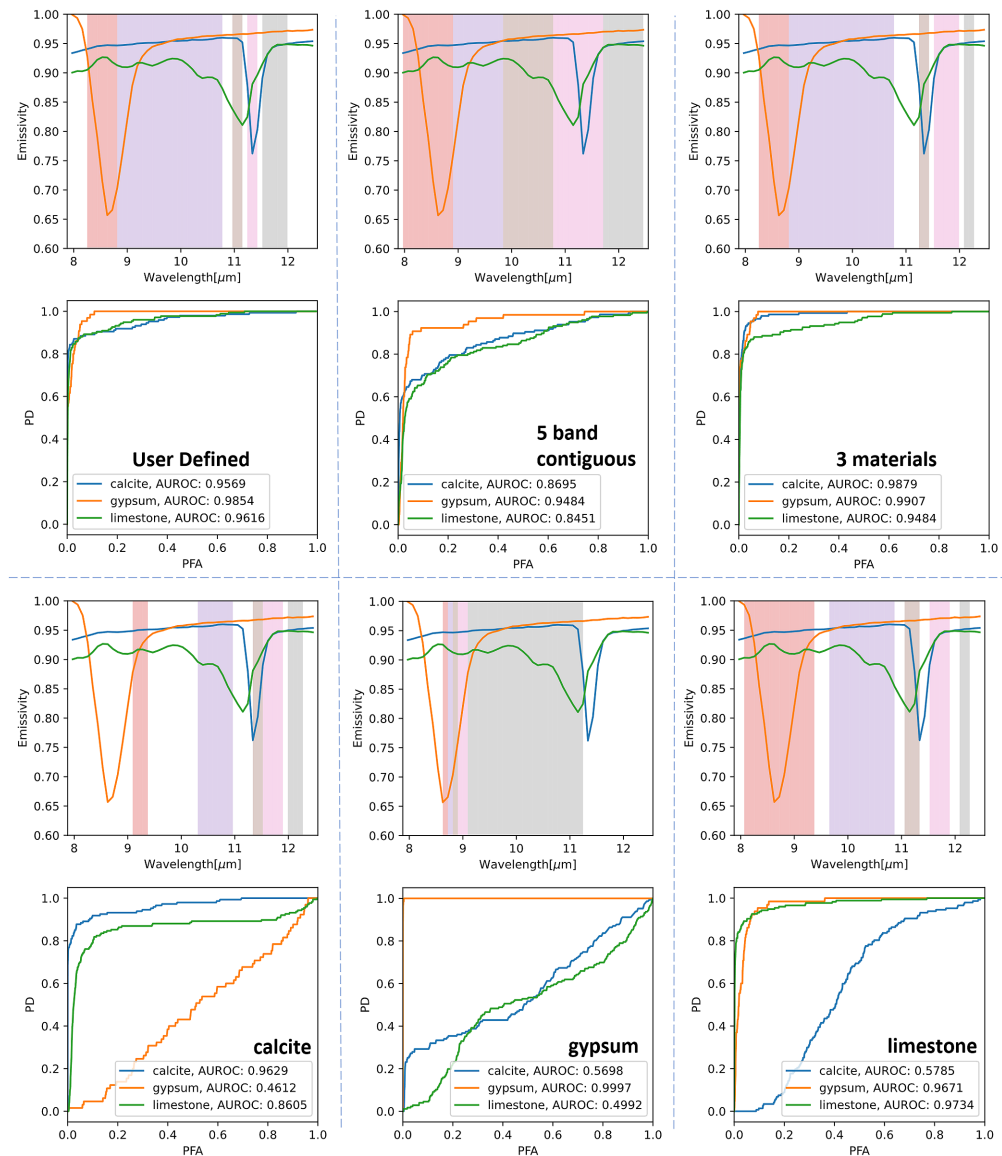


Fig. 10 Five-band results obtained from the best numerical optimizations, along with a “User Defined” set of bandpasses in the upper left.

in a location that was spectrally flat for the targets used here. Moving this band improved performance for limestone detection but resulted in decreased performance for calcite and gypsum. This serves as a reminder that flat target spectra can still be informative relative to the background spectral shape. In this case, the best detection performance overall was obtained by ignoring the primary limestone feature.

6.4 Future Work

Given the good optimization performance, these tests should be expanded to datasets collected under different thermal conditions, different backgrounds, and single pixel detection. Future work might also include modifying the objective function to penalize optimization steps that result in large performance gains for one material at the expense of other materials.

7 Appendix

The optimal separation scores and the resulting spectral band passes for each optimization algorithm and the brute force runs are presented in Table 1.

Table 1 Optimal separation scores and associated band passes.

	Score mean	Score stdev	Score median	Score min.	Score max	Band1	Band2	Band3	Band4	Band5
BF 3 materials	—	—	—	—	0.55515	8.300 to 8.767	8.767 to 8.954	9.328 to 9.795	—	—
BF calcite	—	—	—	—	0.674105	9.141 to 9.234	10.356 to 11.197	11.197 to 11.384	—	—
BF gypsum	—	—	—	—	0.895479	8.580 to 8.673	8.673 to 8.954	8.954 to 11.571	—	—
BF limestone	—	—	—	—	0.560547	8.487 to 8.767	8.767 to 9.047	9.701 to 9.795	—	—
PS 3 materials	0.518364	0.031276	0.542802	0.453742	0.55515	8.300 to 8.767	8.767 to 8.954	9.328 to 9.795	—	—
PS calcite	0.575677	0.105161	0.668617	0.452116	0.674105	9.141 to 9.234	10.356 to 11.197	11.197 to 11.384	—	—
PS gypsum	0.886491	0.019354	0.895479	0.842728	0.895479	8.580 to 8.673	8.673 to 8.954	8.954 to 11.571	—	—
PS limestone	0.504094	0.028198	0.492389	0.467229	0.560547	8.487 to 8.767	8.767 to 9.047	9.701 to 9.795	—	—
DA 3 materials	0.487246	0.030931	0.475659	0.45688	0.551799	8.487 to 8.673	8.767 to 9.328	9.608 to 9.701	—	—
DA calcite	0.670912	0.003703	0.671809	0.663004	0.674105	9.141 to 9.234	10.356 to 11.197	11.197 to 11.384	—	—
DA gypsum	0.880705	0.016615	0.886127	0.849579	0.895479	8.580 to 8.673	8.673 to 8.954	8.954 to 11.571	—	—
DA limestone	0.551234	0.004991	0.004991	0.536583	0.560056	9.888 to 9.982	11.010 to 11.290	11.758 to 12.038	—	—
DE 3 materials	0.489944	0.030557	0.475659	0.470109	0.55515	8.300 to 8.767	8.767 to 8.954	9.328 to 9.795	—	—
DE calcite	0.671463	0.001447	0.671307	0.667396	0.673753	9.141 to 9.234	10.449 to 11.197	11.197 to 11.384	—	—
DE gypsum	0.895471	0.000028	0.895479	0.895335	0.895479	8.580 to 8.673	8.673 to 8.954	8.954 to 11.571	—	—
DE limestone	0.550541	0.004138	0.54967	0.542579	0.560056	9.888 to 9.982	11.010 to 11.290	11.758 to 12.038	—	—
BF 3 materials	—	—	—	—	0.66004	8.580 to 8.673	9.047 to 9.234	10.823 to 11.290	11.290 to 11.477	—

Table 1 (Continued).

	Score mean	Score stdev	Score median	Score min.	Score max	Band1	Band2	Band3	Band4	Band5
BF calcite	—	—	—	—	0.749689	9.141 to 9.234	10.262 to 10.916	11.290 to 11.477	11.758 to 12.225	
BF gypsum	—	—	—	—	0.880633	8.580 to 8.673	8.673 to 8.860	8.860 to 9.047	9.047 to 12.412	
BF limestone	—	—	—	—	0.649635	9.888 to 9.982	11.010 to 11.290	11.477 to 12.038	12.038 to 12.225	
PS 3 materials	0.638347	0.033596	0.655121	0.50705	0.66004	8.580 to 8.673	9.047 to 9.234	10.823 to 11.290	11.290 to 11.477	
PS calcite	0.73189	0.010351	0.730067	0.712649	0.748904	9.141 to 9.234	10.262 to 10.916	11.290 to 11.477	11.571 to 12.225	
PS gypsum	0.871592	0.024929	0.878666	0.741122	0.880633	8.580 to 8.673	8.673 to 8.860	8.860 to 9.047	9.047 to 12.412	
PS limestone	0.610239	0.032027	0.626191	0.501694	0.631205	8.393 to 8.673	8.767 to 9.234	10.823 to 11.290	11.290 to 11.571	
DA 3 materials	0.651425	0.010202	0.65415	0.617299	0.66004	8.580 to 8.673	9.047 to 9.234	10.823 to 11.290	11.290 to 11.477	
DA calcite	0.745436	0.006351	0.747494	0.727042	0.749689	9.141 to 9.234	10.262 to 10.916	11.290 to 11.477	11.758 to 12.225	
DA gypsum	0.833707	0.036388	0.834192	0.775509	0.877118	8.580 to 8.673	8.673 to 8.767	8.767 to 8.954	8.954 to 11.664	
DA limestone	0.628085	0.003743	0.629031	0.615428	0.633646	9.047 to 9.141	9.888 to 10.169	11.010 to 11.290	11.571 to 12.038	
DE 3 materials	0.652239	0.006196	0.654658	0.634456	0.659007	8.580 to 8.673	8.860 to 9.234	10.823 to 11.290	11.290 to 11.477	
DE calcite	0.736941	0.006073	0.736679	0.724731	0.749689	9.141 to 9.234	10.262 to 10.916	11.290 to 11.477	11.758 to 12.225	
DE gypsum	0.877151	0.003713	0.878898	0.86791	0.880633	8.580 to 8.673	8.673 to 8.860	8.860 to 9.047	9.047 to 12.412	
DE limestone	0.623406	0.006076	0.623769	0.611997	0.63817	9.047 to 10.075	11.103 to 11.290	11.477 to 11.851	12.131 to 12.225	
PS 3 materials	0.677283	0.008323	0.678158	0.648876	0.688449	8.206 to 8.393	8.487 to 8.767	8.767 to 9.047	10.823 to 11.197	11.290 to 11.571
PS calcite	0.769051	0.0086	0.770449	0.747497	0.784381	9.047 to 9.328	10.262 to 10.916	11.290 to 11.477	11.477 to 11.851	11.944 to 12.225

Table 1 (Continued).

	Score mean	Score stdev	Score median	Score min.	Score max	Band1	Band2	Band3	Band4	Band5
PS gypsum	0.839	0.028362	0.838225	0.802116	0.869461	8.580 to 8.673	8.673 to 8.767	8.767 to 8.860	8.860 to 9.047	9.047 to 11.290
PS limestone	0.669682	0.025642	0.675329	0.544562	0.694275	8.206 to 8.954	9.795 to 10.823	11.010 to 11.290	11.477 to 11.851	12.038 to 12.225
DA 3 materials	0.68012	0.007025	0.682831	0.65738	0.689815	8.487 to 8.767	8.767 to 10.729	11.197 to 11.384	11.477 to 11.944	12.038 to 12.225
DA calcite	0.77466	0.003191	0.774262	0.767682	0.78433	9.047 to 9.234	10.262 to 10.916	11.290 to 11.477	11.477 to 11.851	11.851 to 12.225
DA gypsum	0.799155	0.00483	0.801723	0.787525	0.8075	8.487 to 8.580	8.580 to 8.673	8.954 to 9.421	9.421 to 10.075	10.075 to 12.505
DA limestone	0.675165	0.01315	0.66831	0.658336	0.701529	8.019 to 9.328	9.608 to 10.823	11.010 to 11.290	11.477 to 11.851	12.038 to 12.225
DE 3 materials	0.673066	0.005558	0.673045	0.662165	0.686213	8.580 to 8.767	8.860 to 9.328	11.197 to 11.384	11.477 to 11.944	12.038 to 12.225
DE calcite	0.766065	0.004502	0.766098	0.755009	0.776484	9.047 to 9.141	10.356 to 11.103	11.290 to 11.477	11.571 to 11.664	11.944 to 12.038
DE gypsum	0.804689	0.017791	0.798503	0.783957	0.851933	8.580 to 8.673	8.673 to 8.860	8.860 to 9.047	9.047 to 9.141	9.234 to 9.328
DE limestone	0.676839	0.006006	0.676276	0.665775	0.694299	8.113 to 9.141	9.795 to 10.823	11.010 to 11.290	11.477 to 11.851	11.944 to 12.318

Acknowledgments

Thank you to Janice Lawson and the team at Lawrence Livermore National Laboratory who collected this data.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes. Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344. Release number: LLNL-JRNL-825597.

References

1. R. D. Fiete, "Image quality and lambda FN/p for remote sensing," *Opt. Eng.* **38**(7), 1229–1240 (1999).
2. J. Schott, *Remote Sensing: The Image Chain Approach*, 2nd ed., Oxford University Press (2007).
3. G. Begni, "Selection of the optimum spectral bands for the spot satellite," *Photogramm. Eng. Remote Sens.* **48**(10), 1613–1620 (1982).
4. T. J. Schmit et al., "Introducing the next-generation Advanced Baseline Imager on GOES-R," *Bull. Am. Meteorol. Soc.* **86**(8), 1079–1096 (2005).
5. N. Longbotham et al., "Prelaunch assessment of worldview-3 information content," in *IEEE, WHISPERS* (2014).
6. H. Su et al., "Optimized hyperspectral band selection using particle swarm optimization," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **7**(6), 2659–2670 (2014).
7. Q. Wang, F. Zhang, and X. Li, "Optimal clustering framework for hyperspectral band selection," *IEEE Trans. Geosci. Remote Sens.* **56**(10), 5910–5922 (2018).
8. T. Doster et al., "Designing manufacturable filters for a 16-band plenoptic camera using differential evolution," *Proc. SPIE* **10198**, 1019803 (2017).
9. B. Meyer et al. "Selecting optimal spectral bands for improved detection of autofluorescent biomarkers in multiphoton microscopy," *J. Biomed. Opt.* **25**(7), 071206 (2020).
10. <https://speclib.asu.edu/>
11. <https://speclib.jpl.nasa.gov/ecospeclibdata/rock.sedimentary.limestone.solid.all.ro337.usgs.perknic.spectrum.txt>
12. D. Manolakis et al., "Is there a best hyperspectral detection algorithm?" *Proc. SPIE* **7334**, 733402 (2009).
13. <https://hpc.llnl.gov/hardware/platforms/syrah>
14. E. C. Laskari, K. E. Parsopoulos, and M. N. Vrahatis, "Particle swarm optimization for integer programming," in *Proc. IEEE Congr. Evol. Comput.*, Vol. 2, pp. 1582–1587 (2002).
15. Y. Xiang et al., "Generalized simulated annealing algorithm and its application to the Thomson model," *Phys. Lett. A*, **233**, 216–220 (1997).
16. https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.dual_annealing.html
17. R. Storn and K. Price, "Differential evolution: a simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optim.* **11**, 341–359 (1997).

Michael Zelinski is a research scientist at Lawrence Livermore National Laboratory whose interests include image system modeling, Fourier optics, hyperspectral imaging, high-performance

computing, and machine learning. He received his master's degree in remote sensing from Rochester Institute of Technology in 2009.

Andrew Mastin is an operations research scientist at Lawrence Livermore National Laboratory. He received his PhD in electrical engineering and computer science from Massachusetts Institute of Technology.

Vic Castillo is a scientist and former group leader in the Computational Engineering Division at Lawrence Livermore National Laboratory with a background in computational physics, machine learning, and system dynamics with over 30 years of experience in industry and government research. He received his PhD from UC Davis in applied engineering in 1999.

Brian Yoxall is a mechanical engineer with deep expertise in high performance precision optical systems, microscale assemblies, and challenging operational environments. He received his PhD in mechanical and aeronautical engineering from UC Davis before joining LLNL in 2011. LLNL experience includes designing and building target assemblies for the NIF, design and build of LWIR hyperspectral sensor systems, and as a technical advisor to DOE's Defense Nuclear Nonproliferation R&D office.