

LA-UR- 00-4579

Approved for public release;
distribution is unlimited.

Title: GROUND COVER IDENTIFICATION AND MAPPING BY
CO2 LIDAR IMAGING

Author(s): Bernard R. Foy, C-6
Brian D. Mcvey, X-8
Roger R. Petrin, C-1
Joe J. Tiee, C-1
Carl Wilson, C-6

Submitted to: SPIE Proceedings
Boston November 5-8, 2000

Los Alamos

NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express 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 any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Ground cover identification and mapping by CO₂ lidar imaging

Bernard R. Foy, Brian D. McVey, Roger R. Petrin, Joe J. Tice, and Carl Wilson

Los Alamos National Laboratory
MS J567, Los Alamos, NM 87545

ABSTRACT

We report examples of the use of a scanning tunable CO₂ laser lidar system in the 9-11 μm region to construct images of vegetation and rocks at ranges of up to 5 km from the instrument. Range information is combined with horizontal and vertical distances to yield an image with three spatial dimensions simultaneous with the classification of target type. Reflectance spectra in this region are sufficiently distinct to discriminate between several tree species, between trees and scrub vegetation, and between natural and artificial targets. Limitations imposed by laser speckle noise are discussed.

Keywords: lidar, multispectral imaging, matched-filter analysis, vegetation mapping

1. INTRODUCTION

Infrared remote sensing instruments have been developed and applied to many different applications involving imagery of the earth's surface.¹ In principle, both passive instruments (detecting reflected or emitted light) and active instruments (using reflected laser light) can provide such spatial and spectral information, but only passive instruments are generally used because of their greater simplicity. In this paper, we demonstrate an application of a multispectral lidar platform to mapping of vegetation and rock. Active lidar imagery provides three spatial dimensions rather than two, simultaneously with the identification of the physical target sampled by the laser. At longer wavelengths accessible by a CO₂ laser, terrestrial features can be identified by their distinct spectral signatures.

Lidar sensing of vegetation offers some complementary characteristics to passive remote sensing. In the thermal infrared (8-12 μm), lidar interrogation of natural targets is essentially a pure reflectance measurement, unaffected by topographical shading and spatial variations in temperature. Differential measurements using CO₂ lasers have been known for some time to be useful in discriminating between vegetation types,²⁻⁴ tree species, and rock types.⁵ We have investigated the utility of a scanning CO₂ DIAL system in constructing vegetation maps for broad areas. Although the scanning of a narrow lidar beam is a time-consuming process compared to imaging sensors, the detailed spectral characteristics in the reflectance information may be useful in many applications.

2. EXPERIMENTAL APPROACH

The lidar system for these measurements utilizes a CO₂ laser developed at Los Alamos that combines advantages of high repetition rate, rapid and broad tunability, moderate pulse energy, and good quality spatial profile. The acousto-optically (AO) tuned laser has been described in detail,⁶ so only major features will be repeated here. The oscillator consists of an rf-excited waveguide laser operating at 5 kHz. A pair of AO modulators selects the laser wavelength and Q-switches the laser to provide pulses with temporal FWHM of ~ 200 ns. The wavelength is changed on every shot to a different CO₂ laser line, sweeping through both the 9.6 and 10.6 μm bands of CO₂ and providing 44 wavelengths. The 44-wavelength pattern is repeated to give a net frequency of 113 spectra per second. The pulse energy is amplified with two waveguide modules identical to the oscillator, producing an average pulse energy (over weak and strong lines) of ~ 0.8 mJ. The laser energy is measured on every shot with a thermoelectrically cooled HgCdTe detector at -40°C , and is termed here the "reference signal." Spatial profile and spectral content are monitored with on-line diagnostics. The beam is expanded with reflective optics to an initial beam waist of ~ 5 cm and a beam divergence of ~ 300 - 350 μrad ($2w_0$ parameter). Spatial profiles of the beam at different wavelengths were measured and are overlapped to within $\frac{1}{4}$ of a beam diameter. This is required so that the various wavelengths interrogate the same target or mixture of targets.

The receiver system of the lidar bench consists of a Cassegrain telescope with 30 cm aperture. A series of lenses focuses the return light with $f/1$ imaging onto a HgCdTe detector cooled to 35 K. An 8-12 μm bandpass filter cooled to 77 K is situated

in front of the detector to reduce background photons. The return signal is integrated with a 1 μ s boxcar gate, as is the reference signal for the pulse energy. Reference and return data for individual laser shots are stored on a data system. In all cases, the return signal was divided by the reference energy to account for minor variations of the pulse energy (typically 3-5% over a few minutes). For target ranging, time traces of the return signal are collected with a digitizer. Software is used to locate the peak of this averaged signal in time, monitor changes in the value, and feed this back to the data acquisition system in an automated fashion. The automated ranging system follows range changes as the lidar system is scanned across terrain with an accuracy of ~ 10 ns, equivalent to a range accuracy of ~ 1.5 m. The laser is vertically polarized when it exits the optical train, and no polarization filtering is used in the receiver.

The transmitter and receiver optical paths are made collinear before reflecting off a rotating gimbal mirror and exiting the lidar bench. The gimbal mirror is positioned with stepper motors to an accuracy of ~ 30 μ rad, or about 1/10 of the beam size. Angle steps of the mirror are controlled by the operating system in a manner coordinated with data collection. The step size was generally 350 μ rad, or about one increment of the beam size. Lidar images were obtained by sweeping the lidar line of sight across scenes in a raster manner. Two visible cameras were coaligned with the lidar line of sight, viewing visible light reflected from ZnSe beamsplitters. The visible scenes were videotaped to provide an exact record of the object interrogated at each mirror position, so that lidar data could later be overlaid on the visible image. Due to limitations of the gimbal controller, it is necessary to wait ~ 0.5 s at each location for the motion to cease before collecting data. The typical data accumulation time employed was one second at each mirror location, giving a total of 256 seconds for a 16X16 grid of points over a scene.

2.1 Signal Averaging Considerations

Noise in the measurement of the hard-target lidar return energy can arise from a variety of factors, including detector "dark current," thermal radiation from the scene, atmospheric turbulence, and laser instability. Laser speckle noise produced by the beam on the target can be a limiting factor for lidar measurements in the LWIR region. Goodman⁷ derived the following expression for the speckle signal-to-noise ratio in lidar measurements of this type.

$$S/N = \sqrt{M} \approx \left(\frac{S_m}{S_c} \right)^{1/2} = \sqrt{S_m} \frac{L}{z\lambda} \quad (1)$$

where M is approximately the number of detected speckles in the image of the beam spot (for $M \gg 1$), S_m is the telescope area, S_c is the speckle area on the target, λ is the wavelength, z is the range to target, and L is the width of the beam on target. L/z is the beam divergence. For our optical setup, the speckle S/N can be varied from 3.3 to 12.5 by changing the beam divergence, and for most of the work reported here it was set at ~ 8 , equivalent to speckle fluctuations of $\sim 13\%$.

Ideally, speckle noise and the other noise sources are reduced by signal averaging as the square root of the number of shots.⁸ Consequently, high-repetition rate is an important characteristic in producing good-quality lidar spectra. In a one-second measurement for a given target, the speckle noise would be reduced from 13% to 1.2% (assuming 113 shots/s at each wavelength). We observed, however, that ideal averaging was not obtained in all cases. Experiments on various natural targets such as trees, bushes, and grassy vegetation at short range (< 2 km) showed that shot-to-shot fluctuations could be reduced by averaging to 1-5%, but then reached a plateau. Similar behavior has been noted in early CO₂ lidar work by Menyuk and Killinger,⁹ who proposed that non-independent speckle patterns on consecutive laser shots impart a temporal correlation to the return signal which impedes signal averaging. While our measurements were performed at much higher repetition rate (5000 Hz vs. 10 Hz), similar considerations probably apply. We find that targets such as grass and trees, with reflecting elements that are small and mobile, exhibit better averaging ($\sim 1\%$) than targets such as dirt and insulation board ($\sim 3-5\%$), which are more fixed.

3. CO₂ LIDAR SPECTRA OF OBJECTS

A few example spectra obtained with our lidar instrument are shown in Figure 1. The spectra were corrected for instrument response by dividing them by the observed spectrum of a roughened (flame-sprayed) aluminum target board positioned at about the same range and oriented at near-normal incidence to the laser. Roughened aluminum has a reflectance of $\sim 80\%$ and a nearly flat spectral response in the 9-11 μ m range.^{10,11} There is only a slight upward slope of $< 1\%$ versus increasing wavelength. By taking a ratio to the spectrum of aluminum, atmospheric features are also approximately removed.

RECEIVED

OCT 26 2000

DSTI

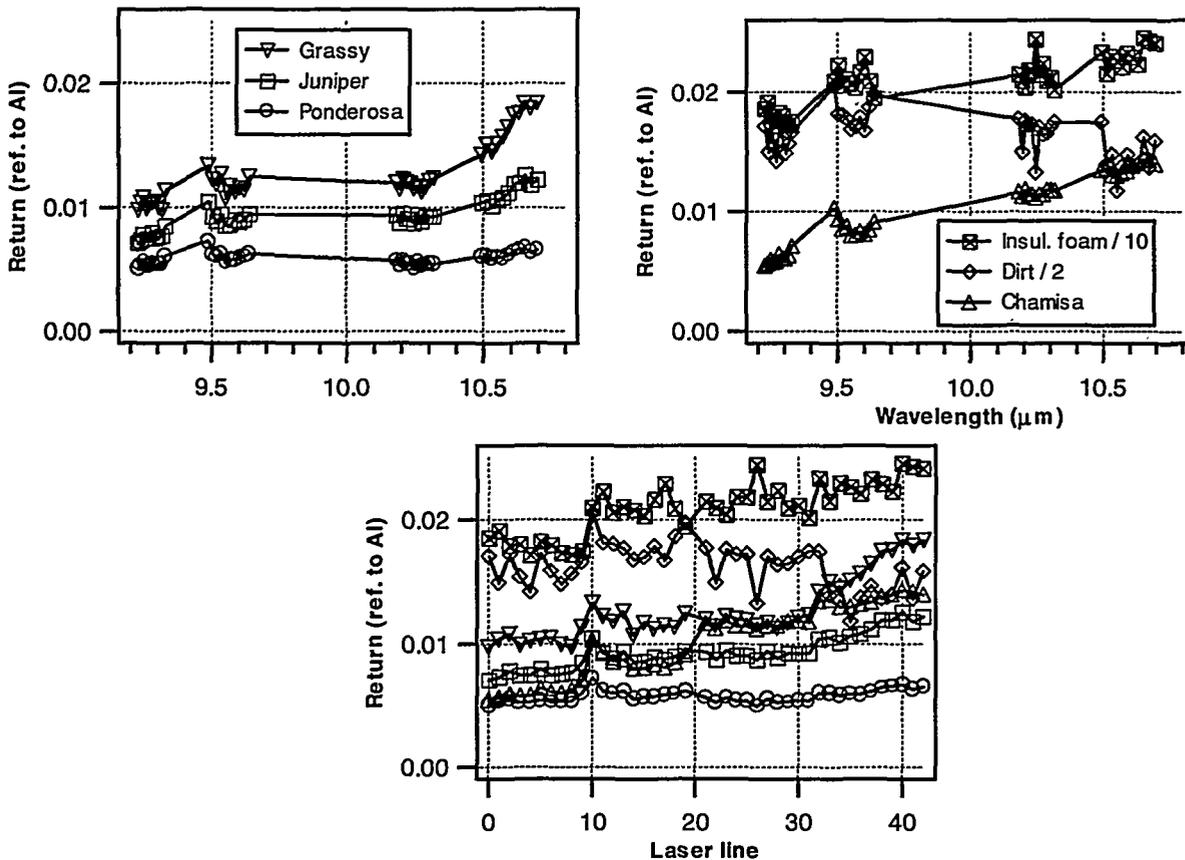


Figure 1. Lidar spectra of natural and man-made objects. Each spectrum is referenced to aluminum, which is assumed to be spectrally flat. The two spectra indicated in (b) were reduced to fit.

The materials of interest in Fig. 1 are grassy vegetation (a mixture of species dominated by Blue Grama Grass, *Bouteloua gracilis*), One-seed Juniper (*Juniperus monosperma*), Ponderosa Pine (*Pinus ponderosa*), a synthetic foam insulation board, a dirt/gravel mixture, and a shrub called Chamisa (*Chrysothamnus nauseosus*). Pinyon Pine (*Pinus edulis*) is not shown but appears spectrally similar to Ponderosa Pine. There are distinct differences in the overall reflectance of these objects, and slight differences in the spectral shape. The trees, shrubs, and grass have a low reflectance of ~1-3%, while the synthetic insulation foam (~30%) and the dirt/gravel mixture (~5%) are considerably higher. We find that these differences are somewhat reproducible among individuals of the same species, and can be used as a criterion for identification, but further work is necessary to clarify the statistical properties of large numbers of objects. Some variation in reflectance is expected from the dependence on angle of incidence.¹² The dirt and insulation foam spectra exhibit larger wavelength-to-wavelength and shot-to-shot variations than the other objects, which is a result of speckle noise described above. Even apart from amplitude differences, many spectra can be distinguished by shape. Grass, for example, has a distinctly rising spectral feature at longer wavelengths. The data were taken in early spring when the grasses were mostly dried stalks. The bottom panel displays the return signal versus laser line, which will be used in subsequent figures.

4. RESULTS AND DISCUSSION

Figure 2 shows typical results for an image constructed from the lidar return signal strength. Two disparate wavelengths are illustrated, and the sum over all wavelengths is also shown. Included in this scene are a rocky cliff, trees, and grassy vegetation, at a range of 4.1 km. A photograph of the scene is shown in Figure 3. The lidar data were obtained over a 16X16 grid of points, with the separation between points being close to the beam size (~350 μ rad). At each pixel location, 128 shots were collected at each wavelength, corresponding to a dwell time of 1.1 s. Because of additional mirror settling time, the full scan required 8.7 minutes to complete. The outline of the trees is apparent in the sum image, but less distinct at the laser line 10P20 ($\lambda = 10.59 \mu$ m). The contrast between trees and grass and between grass and rock varies depending on wavelength. This would indicate that a multispectral analysis will yield more information than single-wavelength images, as expected.

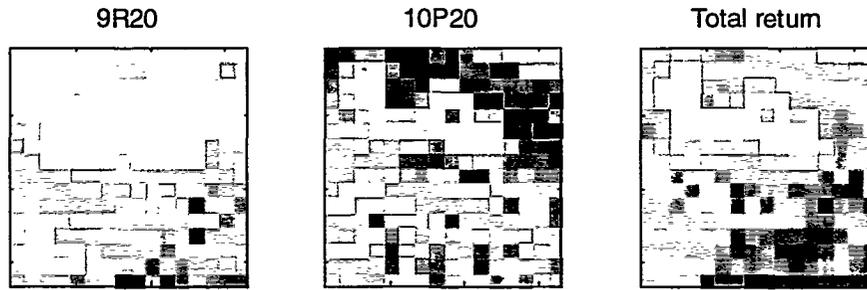


Figure 2. Images of lidar return energy for two selected wavelengths (9.27 μm and 10.59 μm), and summed over all wavelengths. White refers to lower return strength, and black to higher return strength. The scene is the same as shown in Fig. 3. The range is 4.1 km, and the dimensions of the image are 21 m in both horizontal and vertical directions.

A useful type of analysis that incorporates the spectral information is “k-means clustering,” which has frequently been employed in the analysis of passive images.¹ This approach essentially groups together pixels that have similar spectra, organizing them into k different groups or “clusters.” Mathematically, the spectra obtained at each pixel (by averaging for one second) are denoted as pixel vectors \mathbf{y}_i , with i denoting the pixel number; each vector has $p=40$ components corresponding to laser wavelengths. (In practice, a few noisy laser lines are discarded depending on laser performance.) The vectors \mathbf{y}_i have both direction and magnitude in p -dimensional space. They are grouped in the space in k groups and the mean vectors μ_k are found. The points for each pixel are moved between groups in an iterative fashion until the overall squared distance

$$d_{tot}^2 = \sum_{j=1}^k \sum_{i=1}^{n_j} (\mathbf{y}_{ij} - \mu_j)^t (\mathbf{y}_{ij} - \mu_j) \quad (2)$$

is minimized. Here, \mathbf{y}_{ij} is the i 'th pixel belonging to the cluster j , and n_j is the number of pixels in cluster j . The superscript t indicates the vector transpose. The Euclidean distance d is used here for simplicity. The application of this method yields the result shown in Figure 3. Four clusters are used in the analysis ($k=4$) because it corresponds to an intuitive breakdown of the scene elements into trees, grassy areas, and two types of rock, each indicated by a different symbol. In the rocky portion of the scene in the bottom of the image, detailed features are brought out by the analysis. An area of vegetation appears between the rock bands that is not obvious in the photograph. Field inspection confirmed that some sparse vegetation indeed occurs here. In the central upper portion of the photograph, an area of dead branches appears within the grouping of Juniper trees. That area is also revealed in the lidar data (circle symbol), the cluster most closely matching the spectrum of that pixel. In the region of rocks in the bottom half of the scene, the lidar data fall into two clusters. The crosses correspond to rock that is either uncovered by lichen or rock that is different in texture because it is partially broken into gravel.

The resultant spectra that correspond to the clusters in the image are shown in Figure 4. These are the mean spectra μ_j obtained for all of the pixels belonging to each cluster. Figure 4 displays spectra in two ways: as absolute return signal energies, and as differences from the mean spectrum of the entire image. When viewed in the first way, it is clear that all spectra share some common features, particularly evident in the 9 μm band. The attenuation in this region is well understood as absorption from atmospheric ozone (present at ~ 50 ppb).¹³ Atmospheric CO_2 and water vapor also contribute to the overall spectral shape. We have been able to synthesize these features, based on postulated concentrations, for similar spectra taken at ranges up to 7 km. In this way, the lidar system provides an accurate measure of atmospheric concentrations, averaged over the propagation length.

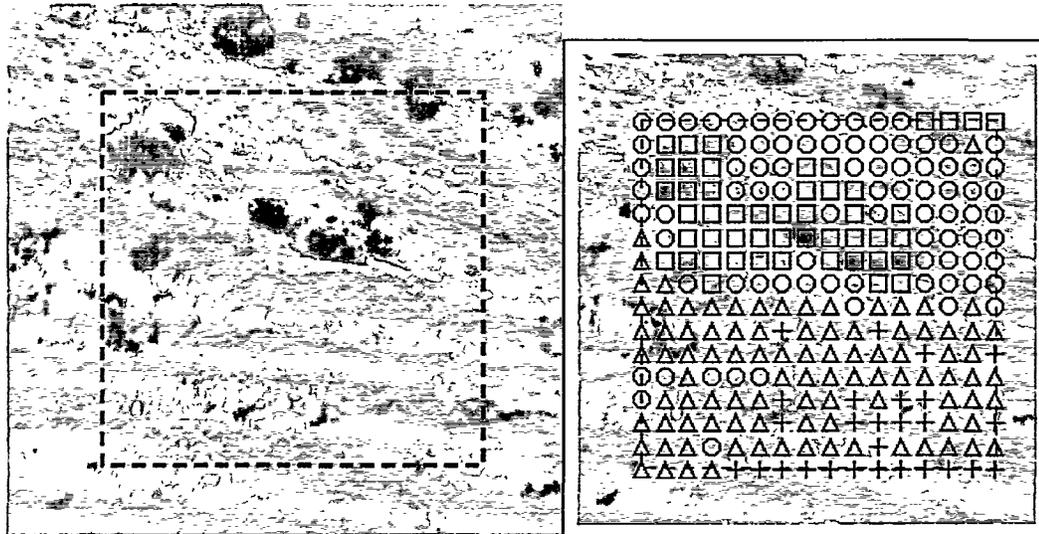


Figure 3. Photograph of scan area and cluster image representation, generated by the k-means cluster analysis of the lidar data. The dashed line in the left photograph indicates the scan area. On the right, symbols corresponding to the four clusters are overlaid on the same picture.

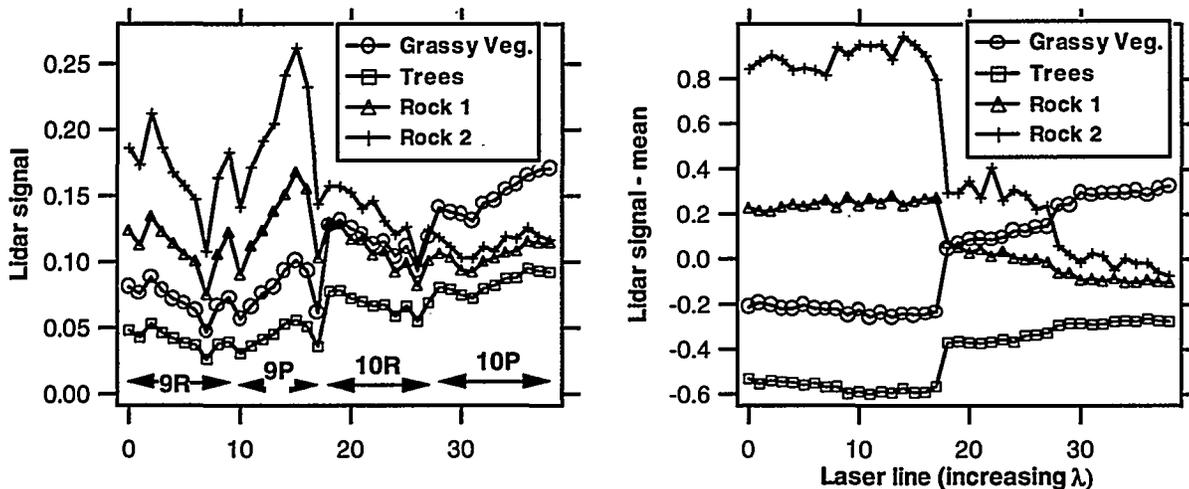


Figure 4. Spectra of the clusters indicated in Fig. 3. Each spectrum is the mean spectrum for all of the pixels belonging to that cluster. In the top panel, the absolute spectra are plotted (referenced to aluminum). In the right panel, the spectra are plotted as the difference from the mean spectrum of the entire image. Note that there is a wavelength gap near the center, as shown in Fig. 1.

The reflectance spectrum of the lidar target is superimposed on the atmospheric absorption spectrum. Thus, it is somewhat easier to recognize spectral differences in target types if the atmospheric contribution is first removed. While that can be done in principle by using known absorption cross sections for atmospheric constituents, it is somewhat tedious since it must be recalculated for each new set of atmospheric concentrations and pathlengths (corresponding to target range). Instead, we used a simple method that is approximately equivalent to atmospheric removal: subtraction of the mean spectrum of the entire image, with results shown in Fig. 4b. (The difference spectrum is subsequently divided by the mean, yielding a spectrum with fractional difference from the mean.) Here it can be seen that different objects in the scene have slowly varying spectra with distinct slopes and amplitudes, in agreement with the data of Fig. 1.

The extent of spectral differences between target types can be visualized with a plot of the signal levels for two selected wavelengths, as shown in Figure 5a. The pixels are seen to fall into distinct groups separated by the hand-drawn boundaries (dashed curves). Within one cluster, the spread is quite large, often larger than the distance between cluster centers. Note that the classification analysis is performed over the full space of p wavelengths, not just the two shown in the plot (in this example, $p=39$). Despite the large spread, we find that the clustering results are quite reproducible. When the same scene is scanned days and even weeks apart, only a few percent of the observed pixels change identity. The amplitude differences

between clusters are also reproducible. This suggests that overall return signal amplitude is an important criterion in distinguishing between target types.

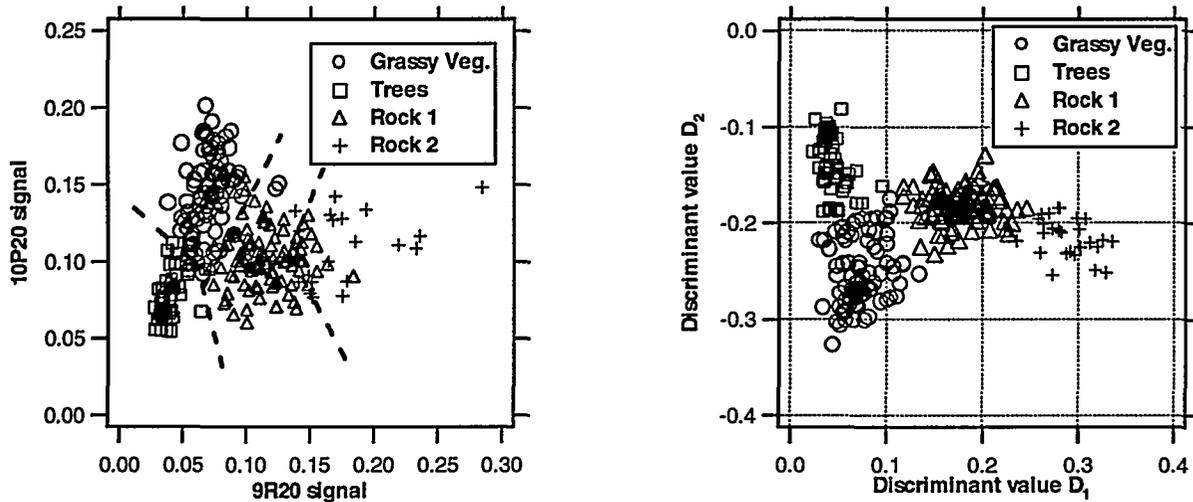


Figure 5. (a) "Scatter plot" of lidar signal strengths at two wavelengths, 10P20 ($\lambda=10.59 \mu\text{m}$) and 9R20 ($\lambda=9.27 \mu\text{m}$). (b) Plot of first two discriminant values, D_2 vs. D_1 for each pixel, labeled by cluster assignment.

More separation exists between pixels than indicated in Fig. 5a if one considers all wavelengths. One way to illustrate this in a two-dimensional fashion is to use discriminant analysis, as it is termed in multivariate statistics. As described by Rencher,¹⁴ we construct discriminant functions which when applied to the individual pixel vectors \mathbf{y}_i yield a scalar result that indicates the degree of separation between the points. The discriminant functions are labeled $D_i = \mathbf{a}_i^T \mathbf{y}_i$. The value of D_i is a scalar obtained as a dot product between the discriminant vector \mathbf{a}_i and the data vector (or pixel spectrum) \mathbf{y}_i . The optimal discriminants are obtained as eigenvectors of the matrix $\mathbf{E}^{-1}\mathbf{H}$, where \mathbf{E} and \mathbf{H} are similar to covariance matrices for the data within and between clusters. Specifically, the matrices are defined by the following.

$$\mathbf{H} = \sum_{i=1}^k n_i (\boldsymbol{\mu}_i - \boldsymbol{\mu}_{tot})(\boldsymbol{\mu}_i - \boldsymbol{\mu}_{tot})^T \quad (3)$$

$$\mathbf{E} = \sum_{i=1}^k \sum_{j=1}^{n_i} (\mathbf{y}_{ij} - \boldsymbol{\mu}_i)(\mathbf{y}_{ij} - \boldsymbol{\mu}_i)^T$$

Here, $\boldsymbol{\mu}_{tot}$ is the mean spectrum of the whole image. Thus, \mathbf{H} is a weighted sum of matrices related to the differences between clusters, while \mathbf{E} is related to the sums of squares between pixel vectors belonging to a given cluster. The largest eigenvalue of $\mathbf{E}^{-1}\mathbf{H}$ corresponds to the first eigenvector \mathbf{a}_1 , and the second largest to \mathbf{a}_2 , and so on. Application of these to each pixel in the data produces discriminant function values D_1 and D_2 . Duda and Hart¹⁵ also describe this type of discriminant analysis. Values of the first two discriminant functions for the set of pixels is shown in Fig. 5b. It can be seen that there is slightly better separation between the four clusters than is apparent in Fig. 5a, where simply the signal values for two wavelengths were plotted. The discriminant plot also shows that the various pixels seem to be correctly assigned to clusters, with the possible exception of a few near the boundaries between clusters.

Examination of Fig. 5 shows that there is clearly significant variation within the four clusters of data. This is due to physical variation in the scene. For example, the Juniper trees vary in foliage density. The grassy area in the scene is inhomogeneous due to a natural distribution of plants on the ground. One might ask whether this extent of variation can be quantified by the lidar measurement. One way to do this is to examine the range of values of the discriminant functions within a given cluster. An alternative way is to construct matched filters corresponding to each individual cluster and apply the filters to the whole scene, looking for variations in the matched filter output. To accomplish this for a given cluster, we segregate the data into target (the cluster of interest: trees for example) and background pixels (all other pixels in the scene: rock, grass). If the target data has \mathbf{y}_{targ} and sample covariance matrix \mathbf{S}_{targ} , the background data has mean \mathbf{y}_{bkg} and sample covariance matrix \mathbf{S}_{bkg} , then an optimal matched filter for the target is

$$\mathbf{a}_i = \frac{\mathbf{S}_{ib}^{-1}(\mathbf{y}_{i\text{arg}} - \mathbf{y}_{bkg})}{(\mathbf{y}_{i\text{arg}} - \mathbf{y}_{bkg}) \mathbf{S}_{ib}^{-1}(\mathbf{y}_{i\text{arg}} - \mathbf{y}_{bkg})} \quad (4)$$

where the appropriate covariance matrix is the pooled, or weighted average covariance matrix between target and background:

$$\mathbf{S}_{ib} = \frac{1}{n_i + n_b - 2} [(n_i - 1)\mathbf{S}_{i\text{arg}} + (n_b - 1)\mathbf{S}_{bkg}] \quad (5)$$

and n_i and n_b are the number of pixels belonging to target and background, respectively. For each of the four clusters in the current analysis, we choose target and background pixels and generate the four matched filter vectors $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, \mathbf{a}_4$. Performing a dot-product application of each of these to all of the image pixels yields a set of four images as shown in Fig. 6. Now we can see that within each cluster, there is a range of values of filter output, i.e. varying degrees of membership to that particular cluster.

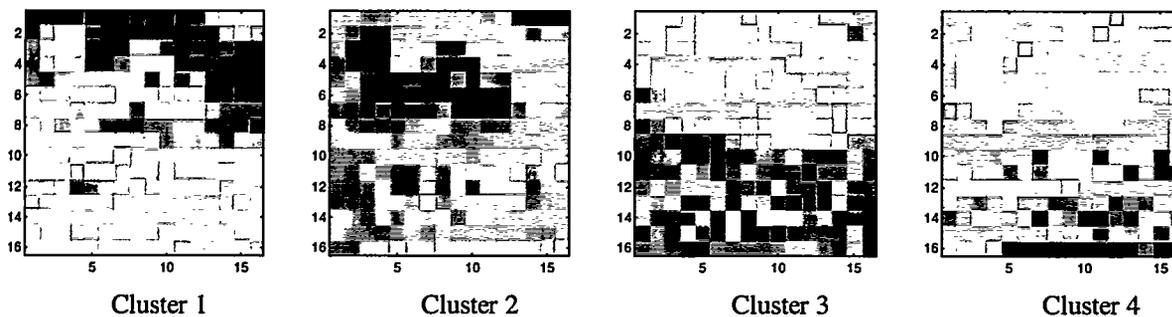


Figure 6. Matched-filter images of data from Fig. 3. Black represents high values of filter output, white is low. Figures from left to right are the results for a filter corresponding to cluster 1, 2, 3, 4, respectively.

In all of these images, three dimensional spatial information is obtained concurrently with the spectral information in each scene component. The third dimension of "depth" or range from instrument to target is determined by necessity in the lidar measurement. The best way to display this information is in a color plot to designate cluster identity.

Finally, we note that chemical plumes can be imaged in the same fashion as landscape features. If the plume is situated between the instrument and the lidar target, it can be detected as an absorption feature. By using clustering and other techniques such as Principal Component Analysis,^{1,14} the chemical signature can be optimally retrieved from a background with a varying spectral signature. This has been discussed recently in connection with passive infrared imaging.¹⁶

ACKNOWLEDGMENTS

A large team of people at Los Alamos combined efforts to make this work possible. We thank Chuck Fite, Paul Johnson, John Jolin, Jerry Romero, Tom Shimada, Dave Thompson, Dennis Remelius, Cliff Hewitt, George Busch, Ed MacKerrow, and Mark Schmitt. This work was sponsored by the Department of Energy.

REFERENCES

1. J.A. Richards and X. Jia, *Remote Sensing Digital Image Analysis*, Springer, New York, 1999, p. 223.
2. M.S. Shumate, S. Lundqvist, U. Persson, and S.T. Eng, "Differential reflectance of natural and man-made materials at CO₂-laser wavelengths," *Appl. Opt.* **21**, 2386-2389 (1982).
3. P.V. Cvijin, D. Ignjatijevic, I. Mendas, M. Sreckovic, L. Pantani, and I. Pippi, "Reflectance spectra of terrestrial surface materials at CO₂-laser wavelengths: effects on DIAL and geological remote-sensing," *Appl. Opt.* **26**, 4323-4329 (1987).
4. R.M. Narayanan, and S.E. Green, "Field-measurements of natural and artificial targets using a midinfrared laser reflectance sensor," *IEEE Phot. Tech. Lett.* **6**, 1023-1026 (1994).
5. T.J. Cudahy, L.B. Whitbourn, P.M. Conner, P. Mason, and R.N. Phillips, "Mapping surface mineralogy and scattering behavior using backscattered reflectance from a hyperspectral midinfrared airborne CO₂ laser system (MIRACO(2)LAS)," *IEEE Trans. Geosci. Remote Sens.* **37**, 2019-2034 (1999).

6. D.C. Thompson, G.E. Busch, C.J. Hewitt, D.K. Remelius, T. Shimada, C.E.M. Strauss, C.W. Wilson, and T.J. Zaugg, "High-speed random access laser tuning," *Appl. Opt.* **38**, 2545-2553 (1999).
7. J.W. Goodman, "Statistical properties of laser speckle patterns," in *Laser Speckle and Related Phenomena*, J.D. Dainty, ed., pp. 9-54, Springer-Verlag, Berlin, 1984.
8. E.P. MacKerrow, M.J. Schmitt, and D.C. Thompson, "Effect of speckle on lidar pulse-pair ratio statistics," *Appl. Opt.* **36**, 8650-8669 (1997).
9. N. Menyuk, D.K. Killinger, and C.R. Menyuk, "Limitations of signal averaging due to temporal correlation in laser remote-sensing measurements," *Appl. Opt.* **21**, 3377-3383 (1982).
10. The ASTER spectral library is available at the web site of the Jet Propulsion Laboratory.
11. R.J. Nemzek, Los Alamos National Laboratory, unpublished results.
12. M.A. Jarzembski and V. Srivastava, "Comparison of continuous-wave CO₂ lidar calibration by use of earth-surface targets in laboratory and airborne measurements," *Appl. Opt.* **37**, 7120-7127 (1998).
13. *Laser Monitoring of the Atmosphere*, ed. E.D. Hinkley, Springer-Verlag, New York, 1976.
14. A.C. Rencher, *Methods of Multivariate Analysis*, Wiley, New York, 1995, p. 303.
15. R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*, Wiley, New York, 1973, p.118.
16. C.C. Funk, J. Theiler, D.A. Roberts, and C.C. Borel, "Clustering to improve matched filter detection of weak gas plumes in hyperspectral thermal imagery," to be submitted, Los Alamos National Laboratory.