Generating passive NIR images from active LIDAR

Shea Hagstrom and Joshua Broadwater

Johns Hopkins University Applied Physics Lab, Laurel, MD

ABSTRACT

Many modern LIDAR platforms contain an integrated RGB camera for capturing contextual imagery. However, these RGB cameras do not collect a near-infrared (NIR) color channel, omitting information useful for many analytical purposes. This raises the question of whether LIDAR data, collected in the NIR, can be used as a substitute for an actual NIR image in this situation. Generating a LIDAR-based NIR image is potentially useful in situations where another source of NIR, such as satellite imagery, is not available. LIDAR is an active sensing system that operates very differently from a passive system, and thus requires additional processing and calibration to approximate the output of a passive instrument. We examine methods of approximating passive NIR images from LIDAR for real-world datasets, and assess differences with true NIR images.

Keywords: LIDAR, active, NIR, passive, approximation

1. INTRODUCTION

The development of airborne LIDAR systems has been important to the advancement of remote sensing. LIDAR is typically used in circumstances where mapping of 3D structure is important, and modern instruments are able to resolve sub-meter details.[1] However, most LIDAR systems provide not only this position information, but also intensity data related to the amplitude of the backscattered laser energy. Intensity information is often used only for context, and for this reason is often poorly calibrated in comparison to passive imaging systems.

Many commercially available LIDAR systems include or are paired with an RGB camera on aircraft. This pairing is useful for mapping applications, but not so much for spectral analysis because of the lack of channels outside the visible region. Most LIDAR systems operate with lasers in the near-infrared (NIR) region of the spectrum,[1,2] and an open question is whether the LIDAR intensity information can be used to produce useful NIR imagery. Spectral analysis requires image bands to be in a consistent, if not also calibrated, state, meaning that in order to be useful the LIDAR data must be processed to "look like" a passively collected image.

A pseudo-NIR (PNIR) image derived from LIDAR would be useful in circumstances where no other source of passively-collected data are available. This includes the previously mentioned LIDAR+RGB platforms, where satellite imagery is either not available or of insufficient resolution. Because LIDAR is an active system, PNIR imagery could be generated from data collected at night or under cloud cover, circumstances where collection of passive imagery is problematic or unavailable. Our goal is to examine what processing would be required to generate passive-like imagery from LIDAR, test these techniques on real-world data, and compare the results with actual passive images.

2. DATA

To demonstrate our image generation techniques we use two datasets provided by the National Ecological Observatory Network (NEON).[3] The first was collected over the Central Plains Experimental Range (CPER) in northern Colorado in June 2013. Airborne data from this collect includes a hyperspectral imager with 1m Ground Sample Distance (GSD), as well as linear-mode LIDAR from an Optech Gemini system at an average of 2.7 pulses/m^2 and a divergence of 0.8 mrad. The laser wavelength is 1064nm, and the LIDAR produces 4-return linear or 1ns waveform data. This particular dataset was chosen because of its large area and flat terrain, which

Author correspondence:
Shea Hagstrom: Shea.Hagstrom@jhuapl.edu
Joshua Broadwater: Joshua.Broadwater@jhuapl.edu

Laser Radar Technology and Applications XXI, edited by Monte D. Turner, Gary W. Kamerman, Proc. of SPIE Vol. 9832, 98320G · © 2016 SPIE
CCC code: 0277-786X/16/$18 · doi: 10.1117/12.2223850

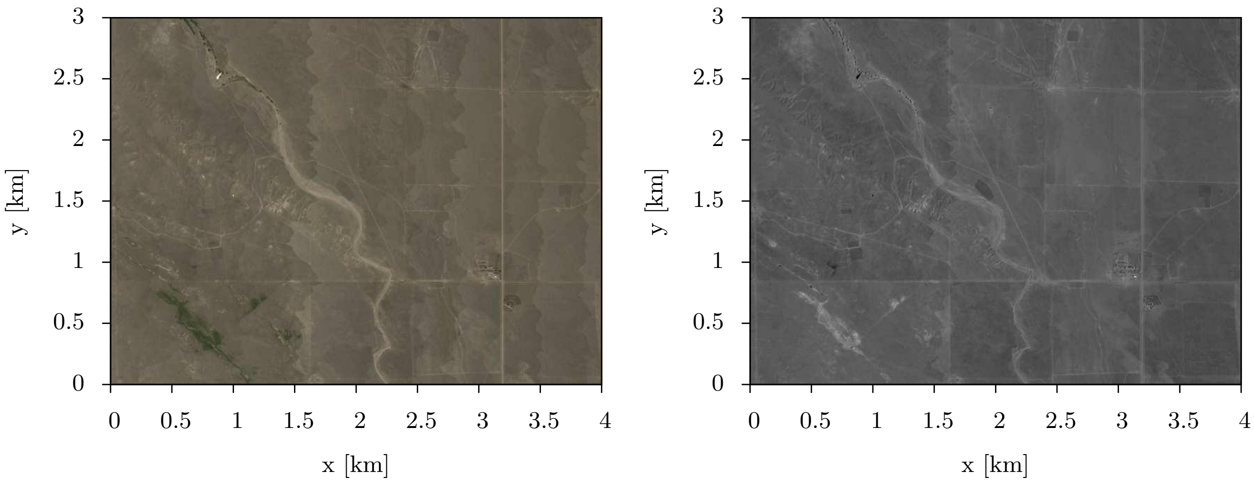

(a) RGB image

(b) 1064nm image

Figure 1. Passive images of the CPER study area. Some minor gaps in coverage are indicated by white in the RGB image.

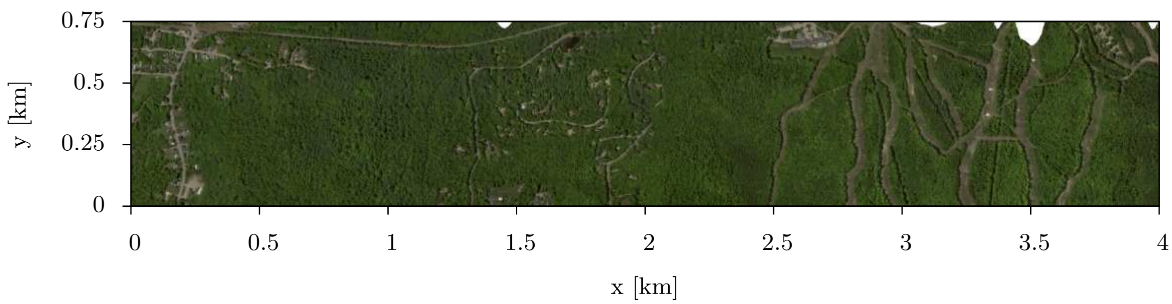

(a) RGB image

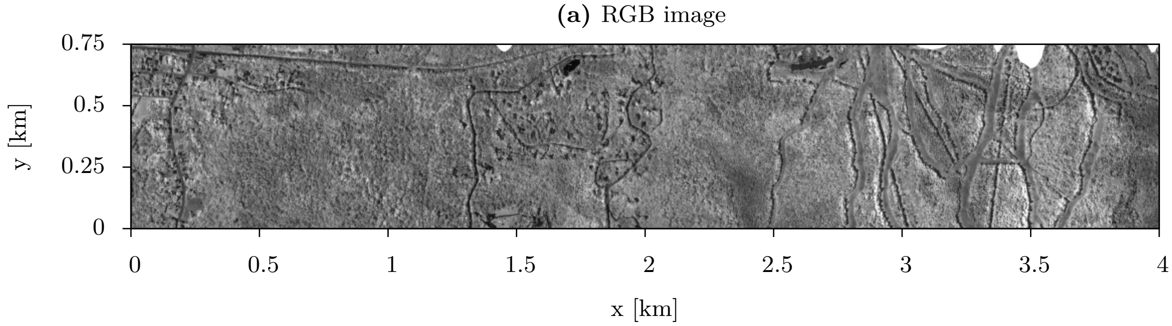

(b) 1064nm image

Figure 2. Passive images of the Bartlett study area. Some minor gaps in coverage are indicated by white in the RGB image.

has very little geometric complexity and should generally produce a single return per LIDAR pulse. Our study area corresponds to a 4 x 3 kilometer subset of the full area, and is shown in Figure 1.

The second dataset was collected from the Bartlett Experimental Forest in New Hampshire in June 2014. 1m hyperspectral data was again collected, along with both linear-mode and waveform LIDAR from a Gemini system at 1.9 pulses/m^2. An overview of the study area is shown in Figure 2, and consists of a 4km x 750m subset of the full area. This dataset was chosen because of the high geometric complexity due to foliage, which provides a challenging case for image generation, and also because there are several flat grassy areas which can be used for comparison. Note that for both datasets half of the LIDAR points were removed to arrive at the above numbers, due to an optical alignment issue that resulted in decreased intensities along one scan direction.

3. METHODS

In order to estimate a passive image from LIDAR intensity information we need to consider the differences between the two modalities. Figure 3 shows our workflow for estimating passive images from LIDAR, and approximating the effects not present in the LIDAR modality. Prior work has been done in combining LIDAR with passive imagery by Bandyopadhyay,[4] but by only considering the raw LIDAR intensities.

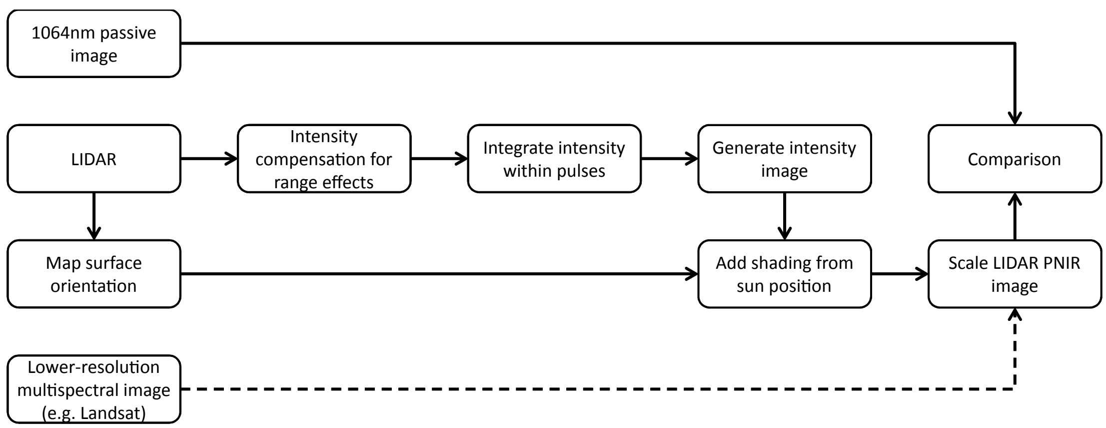

Figure 3. Workflow for our process of estimating a passive image from LIDAR. Lower-resolution multispectral data could potentially be used in the calibration process, but is indicated as optional by the dotted line here.

3.1 Compensating for range effects

The LIDAR instrument records the backscattered laser pulse energy as a function of time. Based on the work of Wagner et al.[5] we can express the laser power reflected back to the system as the simplified function

$$P(t) = \left(\frac{ra_s cos(\theta)}{\pi}\right) * (P_l(t)) \left(\frac{4a_r}{\theta^2 R^4 \pi}\right) \tag{1}$$

where r is the surface reflectance, a_s is the area of the surface, θ is the angle between the laser beam and the surface normal, $*$ is the convolution operation, P_l is the outgoing laser pulse, a_r is the receiver optics area, and R is the range to the surface. We have again assumed Lambertian surfaces for simplicity. Combining all system terms to a constant gives the relative signal intensity for a surface as

$$I \propto \left(\frac{ra_s cos(\theta)}{R^4}\right) \tag{2}$$

where I is the signal intensity, and multiple intensity values may be produced from surfaces intersected by a single laser pulse.

Note that because of the beam divergence, the laser footprint size is proportional to R^2. Assuming a surface is large enough to fill the footprint, the signal falls off as only $1/R^2$. Signals from smaller surfaces will decrease more quickly, which is an argument for preferring smaller beam sizes.

The best we can do in this case is to apply a compensation factor to each return intensity of

$$I_c = I_0 \left(\frac{R}{R_{ref}} \right)^2 \tag{3}$$

where I_c is the corrected intensity, I_0 is the raw intensity, and R_{ref} a reference range that prevents values from growing too large. Slightly improved compensation can be achieved by compensating intensity based on surface orientation as well,[6,7] but we elect to ignore this effect because of our high incidence angles and low surface variability.

3.2 Pulse signal integration

The measured radiance from a surface within the passive instrument GSD, as defined by Schott,[8] assuming only solar illumination and that transmission and bias effects have been removed, is

$$L_{total} = [rKE_{sun}cos(\theta)]\,\frac{1}{\pi} \tag{4}$$

where L_{total} is the full radiance reaching the sensor, r is the surface reflectance, K is the fraction of sunlight hitting the surface, E_{sun} is the solar irradiance at the surface, and θ is the angle between the surface normal and the sun. Note that a Lambertian surface is assumed, adding the $1/\pi$ term, and that wavelength dependencies have been omitted for conciseness. When multiple surfaces fall within the GSD, we can express this as

$$L_{total} = \sum_{i=1}^{N} [r_i a_i K_i E_{sun} cos(\theta_i)]\,\frac{1}{\pi} \tag{5}$$

where a_i is the fractional area of surface i within the GSD, and we use summation as a convenient means of expressing the integral over a finite number of surfaces.

By calibrating the radiance, we can produce the total reflectance r_{total} using

$$r_{total} = \sum_{i=1}^{N} r_i a_i \tag{6}$$

where the cosine term is typically dropped during the process due to lack of knowledge. This is the form often used in analysis such as spectral unmixing, where steady-state illumination integrated over the sensor exposure time is implicitly assumed.

Note the similarity to the earlier LIDAR intensity equation (3), with the primary difference being the LIDAR's dependence on range and time. Assuming range effects have already been compensated for, integrating the LIDAR signal over time is equivalent to summing the signals of all surfaces within the footprint of a single laser pulse, which we can express as

$$I_{sum} = \sum_{i=1}^{N} \left(\frac{r_i a_i cos(\theta_i)}{R_i^4} \right). \tag{7}$$

Performing such a summation is a method of approximating a passive system signal from an active LIDAR. Note that in-scattering from adjacent objects and other light sources is not part of the recorded LIDAR signal, and thus will be a source of error when compared to passive images.

An important consideration is how to handle the positioning of integrated LIDAR pulses in 3D space, because the original pulse may consist of multiple points along the laser path for both linear-mode and waveform LIDAR. We replace each pulse with a single point containing the summed intensity of that pulse, at the position of highest intensity.

We have thus far assumed that the entire backscattered LIDAR signal is recorded so that it can be used during integration. However, this is only true for waveform LIDAR. Linear-mode LIDAR produces a small number of discrete intensity values that are derived from the signal level at an instantaneous signal time, using a triggering mechanism. This triggering typically occurs at the leading edge of Gaussians, but in the case of overlapping or broadened signals may fail. Through this failure to trigger, or due to a limited number of returns, the discretization process may "lose" some of the signal energy. This will tend to occur most in geometrically complex areas, such as foliage, and when the beam diameter is large and tends to encompass more surfaces.

The intensity values from discrete LIDAR are assumed to have been based on signals without a bias level, and any bias has been removed from our waveform data. Attenuation of the return signal occurs due to atmospheric effects,[9] but as long as the range of nearest to farthest detections remains small this effect should remain roughly constant. In practice we will ignore it since the intensity will be scaled later anyway. The remaining variation in the result is ideally due to primarily reflectance, meaning that LIDAR intensity is expected to be most closely correlated with passive reflectance, rather than radiance images with potential upwelling bias effects.

3.3 Shading

LIDAR intensity images are effectively shadowless, because the transmitter-receiver angle is zero. In order to approximate a passive image, shading must be added as if from passive illumination. At NIR wavelengths typically used by LIDAR the contribution of light scattered by the atmosphere is negligible, so we will limit ourselves to handling only the primary solar source for simplicity.

Local surface normals are calculated from the LIDAR point clouds at a resolution of 2 meters using weighted data from within 4 meters. We chose to accomplish this using a voxelized approach,[10,11] but point-based methods are common as well.[6,12]

We can use Equation (4), substituting integrated intensity for reflectance, ray-tracing to estimate the occlusion factor K, and the sun-surface angle for θ, to arrive at an estimate of the illuminated dataset. The sun position is chosen to match the time and location of the passive image we are comparing our PNIR image to. All surfaces are assumed to be Lambertian, because we have insufficient information to do otherwise, but it is important to note that specular surfaces are likely to produce incorrect results.

3.4 Scaling the PNIR image

Note that the PNIR image with shading effects applied is still in uncalibrated units, and is expected in ideal circumstances to be equivalent to passive reflectance with an unknown constant factor applied. Obtaining the factor to scale to reflectance is a non-trivial task, which could be accomplished in several ways.

Likely the easiest method is to use in-scene calibration targets to derive the proper scale factor. However, this severely limits the general-purpose use.

If the target area is covered by lower-resolution multispectral or hyperspectral imagery with a band near the LIDAR wavelength, the PNIR image average can be scaled to match the average of the overlapping lower-resolution image. This method relies on the PNIR image being a good relative estimate of a passive image. Potentially non-coincident lower resolution images could be used for this process, assuming at a minimum that scene changes are minimal and illumination conditions are similar.

If the primary intention is to combine the LIDAR-derived image with simultaneously collected RGB images, lower-resolution multispectral images could also be used to derive the ratio between, for example, the average of the red and NIR bands. The PNIR image could then be scaled in relation to the red channel collected from the LIDAR platform to satisfy this ratio. Again this relies on the LIDAR-derived PNIR image being a good estimate of a passive image overall, and is potentially sensitive to different timing or illumination conditions of the lower-resolution reference images.

What these techniques share in common is an assumption that the aforementioned processing steps are successful in applying passive image properties to the LIDAR data. This is a necessary first step, and will be the focus of our results rather than the potential end-to-end comparison after incorporating external data, which we leave to future work.

4. RESULTS

4.1 Central Plains Experimental Range

The CPER dataset represents what should be the most straightforward case of generating a passive image from LIDAR, because of the flat terrain which produces primarily one-return pulses.

A passive-derived reflectance image at 1064nm is compared to our LIDAR-derived image in Figure 4. Note that the reflectance image shows some minor artifacts at edges between flightlines of the mosaic. These artifacts occur where reflectance retrieval was unable to fully compensate for the wide range of viewing angles and changing illumination conditions, but do not affect the final result by more than a few percent.

The LIDAR-derived pseudo-NIR image in Figure 4b shows the same overall shapes and features as the reflectance image, but is different in subtle ways. Surface shading appears more pronounced in the PNIR image, which we attribute to a combination of adjacency and scattering effects. Some brighter areas, particularly those with foliage of any kind, appear darker than expected as compared to the reflectance image. This is most

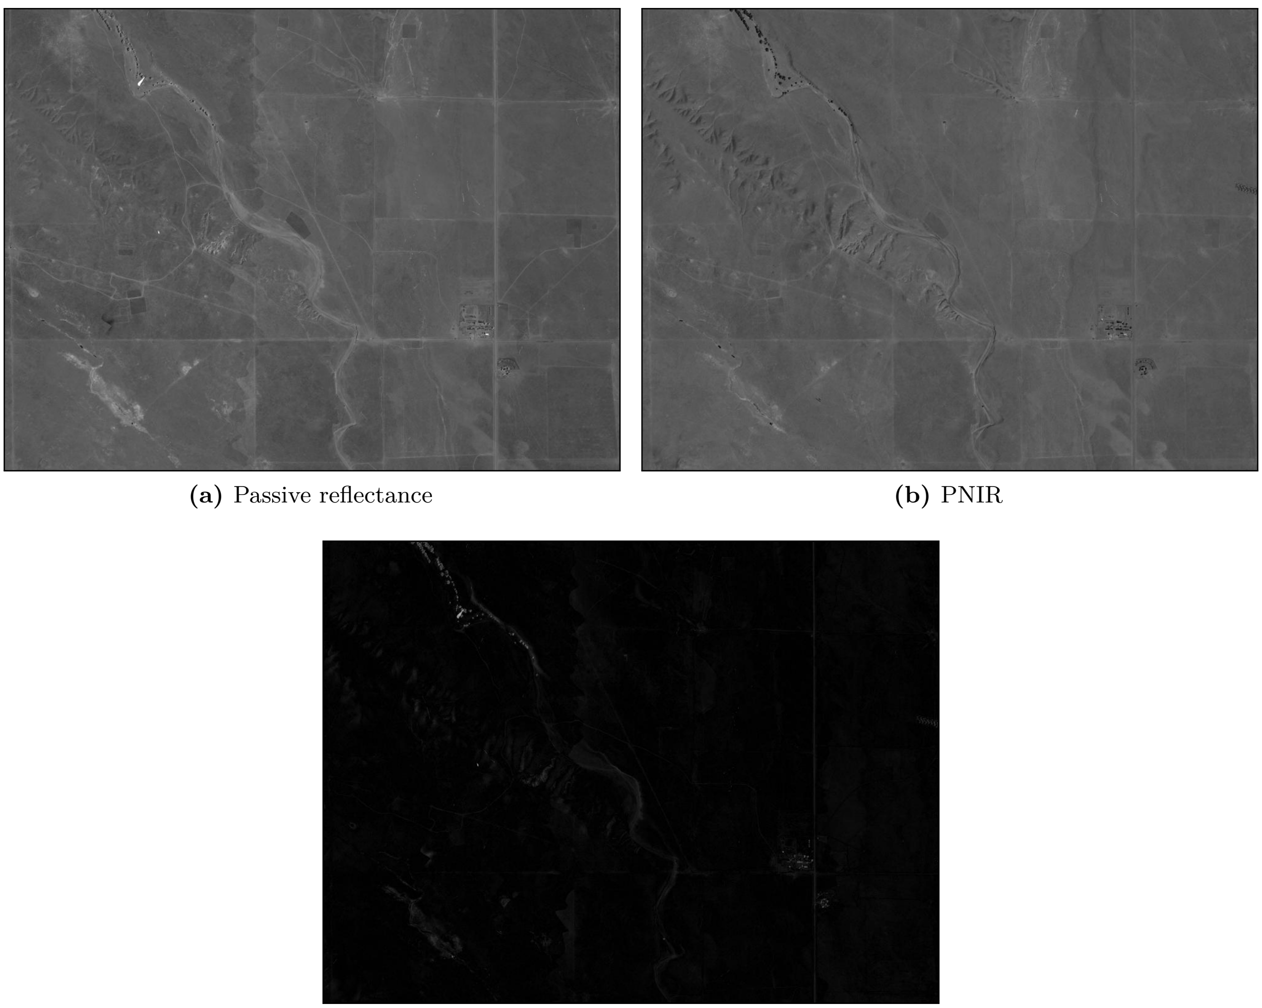

(a) Passive reflectance

(b) PNIR

(c) Difference

Figure 4. Reflectance and PNIR images. Small gaps in the reflectance image coverage appear as white, and do not contribute to processing or statistics. Reflectance and PNIR images are scaled such that 0-1 reflectance fills the image range, and the absolute difference image is unscaled.

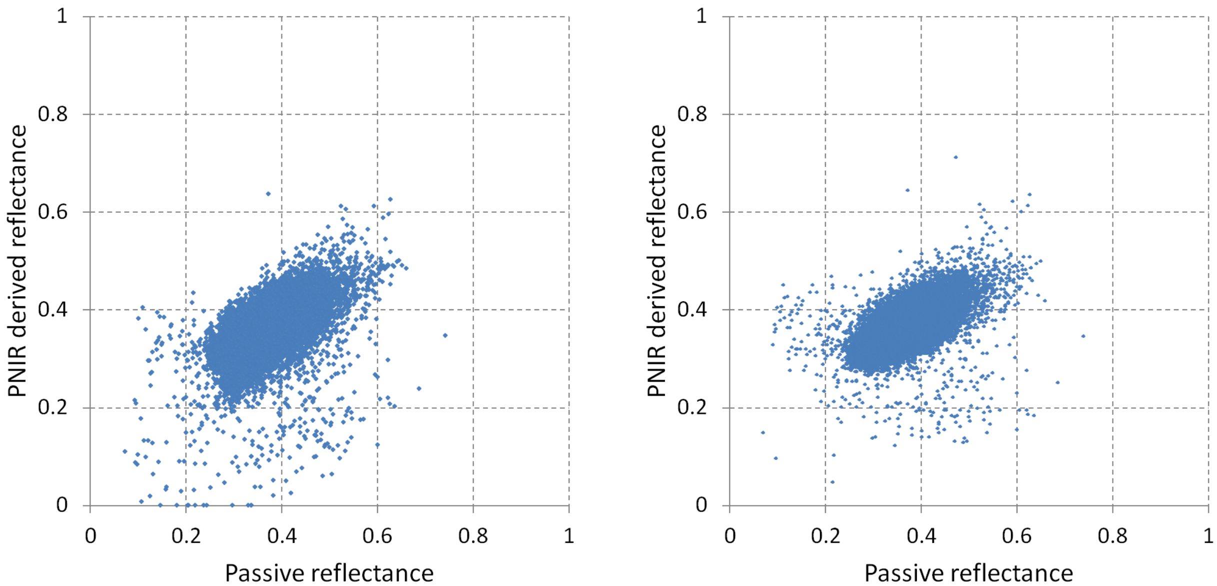

(a) PNIR with shading **(b)** PNIR without shading

Figure 5. Reflectance VS PNIR values for the CPER dataset.

pronounced with trees and the large lighter feature in the lower-left, and we attribute this darkening to lowered intensities as a result of LIDAR temporal spreading from tall and short foliage respectively. The remainder of the difference between the two images may also be caused by specular effects, such as along roads and areas where water runoff has left rocky surfaces cleaner than the surroundings.

The PNIR image of Figure 4 was scaled such that its average was equal to the reflectance image average. Plotting each reflectance image pixel against each PNIR image pixel gives the plot shown in Figure 5a. Though the overall variability of reflectance in the scene is low, the data clearly show a nearly 1:1 relationship with a correlation coefficient of 0.60. Noise is relatively consistent at about $\pm 6\%$. The fact that scaling the PNIR average to the reflectance image average produces a good relationship reinforces our earlier prediction that reflectance is the primary contributor to intensity variation after range effects have been compensated for, at least in the case of flat terrain.

Visual comparison of the passive reflectance and PNIR images shows more pronounced shading in the PNIR image, so we also plotted the pixel values of the unshaded intensity image against passive reflectance in Figure 5b. The unshaded PNIR average was scaled to the passive image average. This decreased the overall variability slightly, but also increased the lower PNIR values and decreased the higher values such that the overall distribution no longer follows the 1:1 line as well. This means that our shading is helping match the LIDAR-derived image to the passive image, but at the expense of some small additional error due to imperfections in the process.

The average PNIR results are consistent enough that our reflectance-derivation technique from Section 3.4 using lower-resolution external images could likely be used. However, the differences from the reference passive image at the smaller scale suggest that it would be of limited use for achieving increased resolution.

4.2 Bartlett, NH

The Bartlett dataset contains additional complexity in the form of forests and elevation changes, making it much more challenging to estimate a passive image from LIDAR. A high fraction of laser pulses within the study area contain multiple returns, which from our earlier analysis we would expect to cause issues.

(a) Passive reflectance

(b) PNIR

(c) PNIR scaled using exposed ground

Figure 6. Reflectance and PNIR images. Small gaps in the reflectance or PNIR image coverage appear as white, and do not contribute to processing or statistics. Reflectance and PNIR images are scaled such that 0-1 reflectance fills the image range.

Figure 6 shows the passive reflectance band centered about 1064nm and the PNIR image derived from the integrated linear-mode LIDAR. The PNIR image is scaled so that its average is the same as the reflectance image's average. The match between images is very poor in this case, and it is visually apparent that the forested areas are too dark while simultaneously the exposed ground is much too light.

From both our theoretical analysis and the earlier CPER example we expect the best estimate of the reflectance from LIDAR to occur in the areas containing exposed ground, because these areas will produce only a single return per laser pulse. Figure 6c shows the result of scaling the PNIR image average to the reflectance average using only the exposed ground areas in both images. This results in a good match in those ground areas, but also foliage that is far too dark. The overall correlation coefficient is 0.37. Note that using small exposed ground areas is an unrealistic method of calibration if relying on lower-resolution passive images, but is included here for demonstration purposes.

One of our hypotheses from the earlier theoretical analysis is that linear-mode LIDAR would appear darker than expected in areas with multiple returns due to scattering effects and loss of return signal power. A waveform LIDAR system should show very little of the signal loss inherent to a linear-mode system, so we generated a PNIR image from waveform LIDAR as seen in Figure 7.

The waveform intensity image was derived by integrating over the entire return signal of each pulse as outlined in Section 3.2, and the result after applying illumination and scaling by averages shown in Figure 7b. Though still not a good match to the passive reflectance reference image, the waveform-derived image is at least more similar than the linear-mode image, particularly in foliated areas. The overall correlation coefficient is 0.53.

Like the prior linear-mode image, we also attempted scaling by using only the exposed ground areas as seen

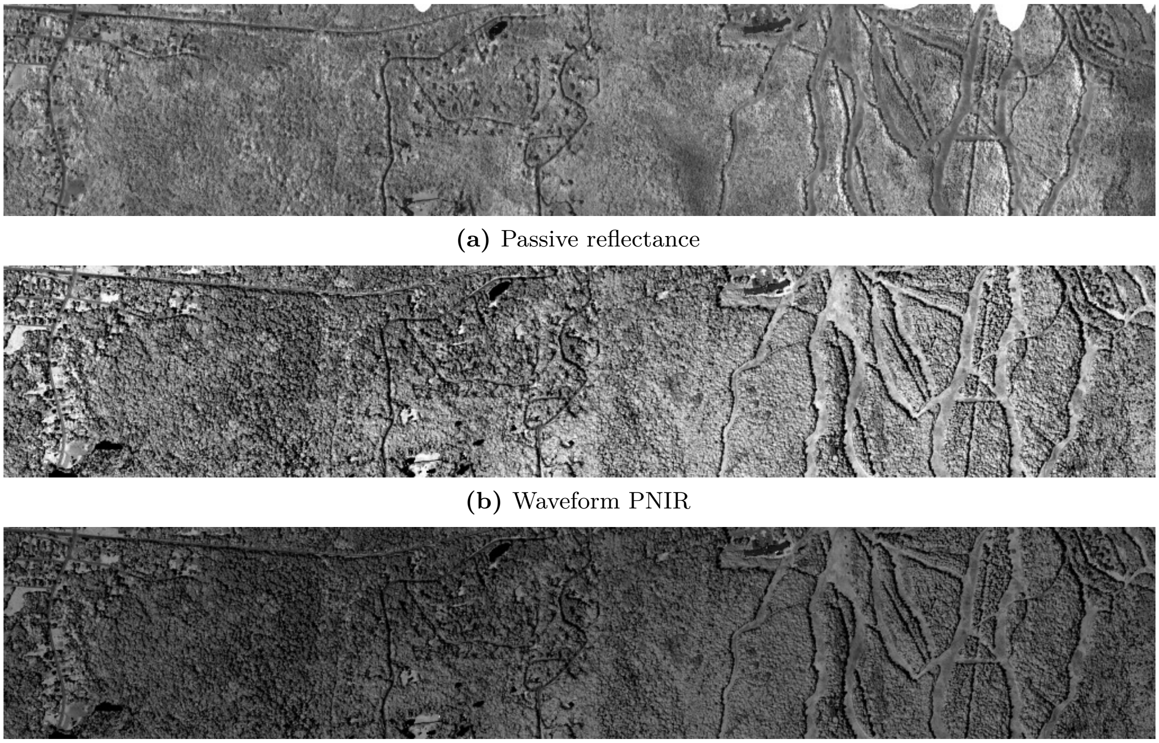

(a) Passive reflectance

(b) Waveform PNIR

(c) Waveform PNIR scaled using exposed ground

Figure 7. Reflectance and PNIR images. Small gaps in the reflectance or PNIR image coverage appear as white, and do not contribute to processing or statistics. Reflectance and PNIR images are scaled such that 0-1 reflectance fills the image range.

in Figure 7c. This results in a better match than the linear-mode result of Figure 6c, particularly in the foliated areas, but still is not a great match to the passive image.

Figure 8 compares the distribution of PNIR values derived from linear-mode and waveform LIDAR, after both have been scaled by their respective averages. The bimodal distribution of the discrete LIDAR result is due to the exposed ground and forested areas. Because the majority of the image is forest, and appears abnormally dark, scaling by the average value artificially increases the intensity of the exposed areas not subject to this darkening effect. In contrast, the waveform distribution is much more consistent, indicating a more complete recording of the backscattered signal was captured. Neither distribution is 1:1, but the waveform result at least shows a bounded and somewhat linear and correlated result for the very complex case of foliage.

Another way to visualize the difference between the linear-mode and waveform LIDAR images is by considering the energy loss due to the linear discretization process. We compute the ratio of waveform to discrete LIDAR integrated intensity after equalizing them using exposed grassy areas, the results of which are shown in Figure 9. This image indicates that flat areas show nearly equal intensity for both LIDARs, and that the waveform shows significantly higher intensity in any areas containing foliage. Several areas actually show the waveform data being about 10 times higher, though it is unclear exactly what causes this localized large discrepancy. The overall ratio is 1.90, meaning that each linear-mode pulse missed on average about half of the backscattered energy.

However, even the increased data provided by the waveform LIDAR was insufficient to produce an accurate passive image estimate overall. The remaining difference is likely a contribution of several effects, including inaccuracy in lighting estimates, the lack of in-canopy scattering, and specular effects.

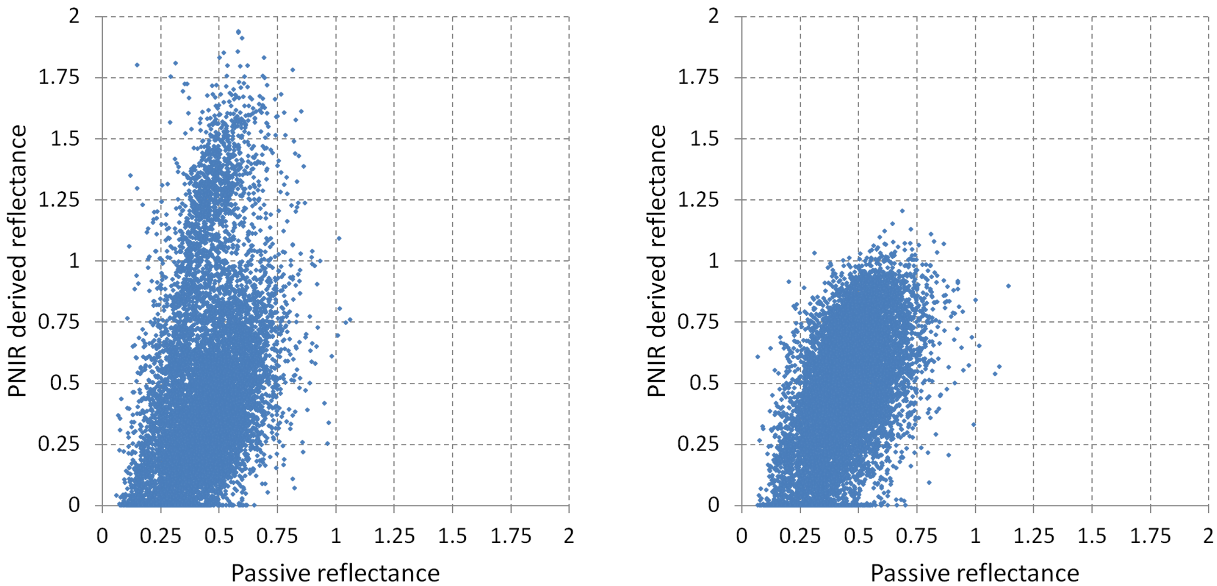

(a) Scaled PNIR from discrete LIDAR

(b) Scaled PNIR from waveform LIDAR

Figure 8. Reflectance VS PNIR values for the Bartlett dataset, using both linear-mode and waveform LIDAR. Correlation coefficients are 0.37 and 0.53 for (a) and (b) respectively. The scale is extended here to show the full extent of the data distribution.

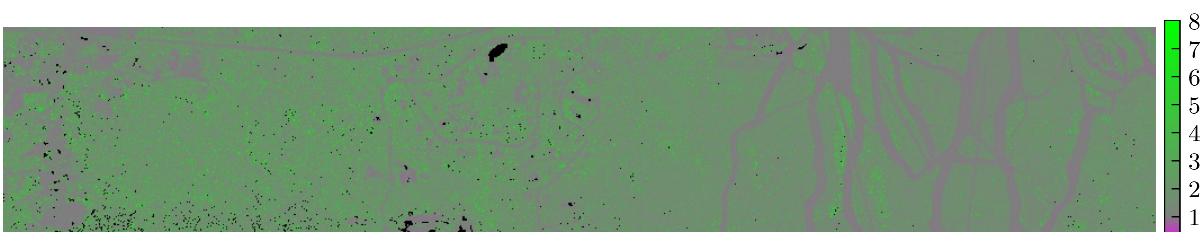

Figure 9. Ratio of waveform to linear-mode integrated intensity for the Bartlett study area. Areas missing either discrete or waveform data are shown as black. Note the non-linear scale used for color. The overall ratio is 1.90.

5. CONCLUSIONS

Our primary goal of producing images from LIDAR that are equivalent to passive images was only partially successful. The simpler geometry of the flat plains in the CPER dataset showed that there is a decent correlation of 0.60 between our derived image and the reference passive image, which also appears as nearly 1:1 after our processing despite the low overall variability of the scene. The primary sources of error in this case appear to be specular effects and overemphasizing the shadowing due to surface angle as compared to the reference image.

The discrete LIDAR-derived image for the much more complex Bartlett scene showed a much lower correlation of 0.37 with the reference passive image, and is also not a good 1:1 match. By examining waveform LIDAR collected over the same area, we found that approximately half of the integrated energy per pulse was on average lost due to the use of the linear LIDAR intensity, and in some areas was much higher. Using the waveform LIDAR data to generate an image improved both the correlation and visual appearance, but was obviously still not a good match to the reference image. We attribute the remaining error to being unable to account for the complex in-canopy scattering and shading effects.

The process of generating imagery from LIDAR remains interesting in that the effect of adding shadows and shading due to illumination can make them more readable, but unfortunately does not perform well in approximating a true passive image. Several potential options for using lower-resolution multispectral images to aid in calibration were outlined, but without good results in the direct PNIR-passive comparison of this paper they are not viable options.

The fact that the method performs most poorly in foliated areas, where it would arguably be most useful in estimating statistics such as NDVI, means that LIDAR is unlikely in general to be a useful source of spectral data complementary to passive RGB images. However, this does not preclude multi-wavelength LIDAR from being used without passive imagery, and the analysis of signal losses in discrete LIDAR systems may prove important in future work.

ACKNOWLEDGMENTS

The authors would like to thank the National Ecological Observatory Network for making these datasets available for our use.

REFERENCES

[1] Lemmens, M., "Airborne lidar sensors," *GIM International* **23**(2) (2009).

[2] Baltsavias, E. P., "Airborne laser scanning: existing systems and firms and other resources," *ISPRS Journal of Photogrammetry and Remote Sensing* **54**(2-3), 164 – 198 (1999).

[3] National Ecological Observatory Network, Available on-line http://data.neoninc.org/ from NEON, Boulder, CO, USA. Accessed: Dec. 2015 (Bartlett), Feb. 2016 (CPER).

[4] Bandyopadhyay, M., *Quantifying the urban forest environment using dense discrete return LiDAR and aerial color imagery for segmentation and object-level biomass assessment*, PhD thesis, Rochester Institute of Technology (April 2015).

[5] Wagner, W., Ullrich, A., Ducic, V., Melzer, T., and Studnicka, N., "Gaussian decomposition and calibration of a novel small-footprint full-waveform digitising airborne laser scanner," *ISPRS Journal of Photogrammetry and Remote Sensing* **60**(2), 100 – 112 (2006).

[6] Jutzi, B. and Gross, H., "Normalization of lidar intensity data based on range and surface incidence angle," in [*In Proceedings of Laserscanning 09*], 213–218 (2009).

[7] Shaker, A., Yan, W. Y., and El-Ashmawy, N., "The effects of laser reflection angle on radiometric correction of the airborne lidar intensity data," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XXXVIII-5/W12**, 213–217 (2011).

[8] Schott, J., [*Remote Sensing : The Image Chain Approach, 2nd Edition*], Oxford University Press (2007).

[9] Measures, R. M., [*Laser Remote Sensing: Fundamentals and Applications*], Wiley-Interscience (1984).

[10] Hagstrom, S., *Voxel-based LIDAR analysis and applications*, PhD thesis, Rochester Institute of Technology (July 2014).

[11] Hagstrom, S. and Broadwater, J., "Atmospheric and shadow compensation of hyperspectral imagery using voxelized lidar," in [*Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International*], 2959–2962 (July 2015).

[12] Maas, H.-G. and Vosselman, G., "Two algorithms for extracting building models from raw laser altimetry data," *ISPRS Journal of Photogrammetry and Remote Sensing* **54**(2-3), 153 – 163 (1999).