Multispectral Thermal Imaging with a LWIR+VIS Camera

Jacob Hines
Department of Applied Physics
Stanford University
hines@stanford.edu

Evan Wang
Department of Electrical Engineering
Stanford University
wangevan@stanford.edu

Abstract—Thermal imaging, operating at long-wave infrared (LWIR) wavelengths, is important for a wide range of room temperature applications. However, cameras at this wavelength are expensive compared typical visible (VIS) cameras. We built a multispectral camera that captures a high-resolution visible image to enhance a low-resolution thermal image at a fraction of the cost of higher-end thermal cameras.

Index Terms—thermal, multispectral, registration, fusion

I. INTRODUCTION AND MOTIVATION

According to Wien's displacement law, objects at room or body temperature emit black body radiation with a peak radiance in the range of 9-10 μm [1]. This range of so-called "thermal radiation" – or long-wave infrared (LWIR) – serves a critical role for defense and security detection, product inspection, and other room temperature imaging purposes [2]. Despite the abundance of infrared light to sense, LWIR cameras are very expensive compared to their visible counterparts, with entry level units ranging from $1000 - $3000 [3].

FLIR Systems produces the Lepton, a compact and significantly more affordable LWIR camera module that is designed to fit inside a smartphone [4]. However, due to its low cost ($259) and small size (10.50 x 12.7 x 7.14 mm), the camera has a limited resolution of only 160 x 120 pixels. Such a low-resolution is not ideal for imaging applications.

The objective of this project is to combine the Lepton LWIR camera with a high-resolution visible light camera and compare multispectral fusion methods that combine the spectral information from the thermal image with the gradient information from the thermal image. Such a device would serve as an inexpensive thermal camera with reasonable image quality for standard imaging applications.

II. RELATED WORK

Multispectral imaging over visible and infrared (IR) bands is already an area with great interest. A significant focus of work in this area is related to hyperspectral object identification [5, 6], pedestrian detection [7], and product inspection [8]. Techniques in these applications range from simply overlapping information from both visible and thermal cameras to create images with more channels of information, to using thermal images to computationally remove backgrounds from

TABLE I
COMPARISON OF THERMAL CAMERAS

	FLIR Boson 640	FLIR Lepton 3 + R Pi
Cost	$3000	$260 + $190
Resolution	640 x 512	160 x 120 IR 3200 x 2400 VIS
Frame rate	60 Hz	9 Hz
Sensitivity	50 mK	50 mK
Horizontal FOV	Up to 95°	57°
Spectral range	7.5 - 13.5 μm	8 - 14 μm

visible images in order to isolate objects of interest [6]. However, these applications rely on expensive high-resolution thermal cameras.

Once a LWIR+VIS image pair has been captured, the two images must be registered. Due to large spectral differences in image content, standard descriptor-based registration techniques do not work well [9]. To solve the registration problem, many application-specific techniques have been proposed [10–12]. These use methods include contour and segemnt detection [13, 14] and hybrid feature extraction from urban environments [15].

Finally, there are a variety of works that explore pansharpening, the use of high-resolution panspectral images to sharpen low-resolution multispectral images [16, 17]. Notable methods include highpass filtering, wavelet decomposition [18–21], and a variety of multichannel fusion techniques [22–26].

III. METHOD OVERVIEW

Our compact system for multispectral fusion is designed around the FLIR Lepton thermal camera module, which we interfaced with a Raspberry Pi Zero W board. For the visible camera we used a Raspberry Pi Camera Module V2, which has a resolution of 2464 x 3280 pixels. In total, the completed camera cost just under $500, about half of which comprises the Lepton. The remaining half comprises the Raspberry Pi, visible camera, and assorted accessories. Table I shows a comparison between of various features of the FLIR Lepton 3 and the FLIR Boson 640, a representative high quality thermal camera module [27]. Our camera is explained in detail in Section IV.

Image capture is synchronized via the Raspberry Pi. When the capture button is pressed, the Raspberry Pi captures images from the thermal and visible cameras, then saves the two

images to memory. The images are later offloaded to an external computer for post-processing. The post-processing can be divided into two steps: image registration and image fusion.

Since the two cameras are offset by approximately 1 inch in the camera housing, the captured visible and thermal images experience parallax. Additionally, the field-of-view (FOV) of the visible camera is slightly larger than that of the thermal camera. Therefore it necessary to register the images prior to image fusion. To accomplish this, we first perform edge detection on both images and compute the cross-correlation between the edges. Maximizing the edge-edge correlation yields a shift that best aligns the two images. The registration algorithm is described in detail in Section V.

There are many methods to fuse the registerd VIS+LWIR images. A simple single channel method involves injecting the high frequency components of the visible image into the thermal image. A more complex method preserves the intensity information of the thermal image while transferring gradient information from the visible image. We present these methods, along with a luminance-chrominance method, in Section VI. We then discuss their results in Section VII.

IV. CAMERA

A. Hardware Design

The VIS+LWIR camera is built around the a Raspberry Pi Zero W board which serves as the main controller of the camera. The Raspberry Pi is connected to a PiCamera V2 visible camera and the FLIR Lepton 3 thermal camera. The two cameras are mounted adjacent to each other on the inside of a aluminum box which houses the Raspberry Pi controller and cameras. Figure 1 shows a front view of the completed camera. The Lepton and PiCamera can be seen mounted in the housing.

On the rear of the device, a 5-inch touchscreen panel is mounted to the exterior of the housing. The screen allows the user to access the Raspberry Pi and run the camera software as well as see previews of the captured images. The Raspberry Pi is controlled via a wireless keyboard and mouse which connect via a USB receiver. An external USB flash drive allows images to be transferred off the camera to a desktop computer for processing. Figure 2 shows a block diagram of the camera module.

In total, the camera cost just under $500 including the needed mounts, cables, and peripherals. Appendix gives the itemized breakdown of parts used in the completed camera.

B. Image Capture

To capture images we use the Python package `picamera` [28] to interface with the PiCamera and the C module `Lepton3Module` [29, 30] to interface with the FLIR Lepton 3. The `Lepton3Module` image capture function is compiled into an executable file which can be run from within a Python script.

In the camera's idle state the live preview feed from the PiCamera is displayed on the touchscreen display. When the
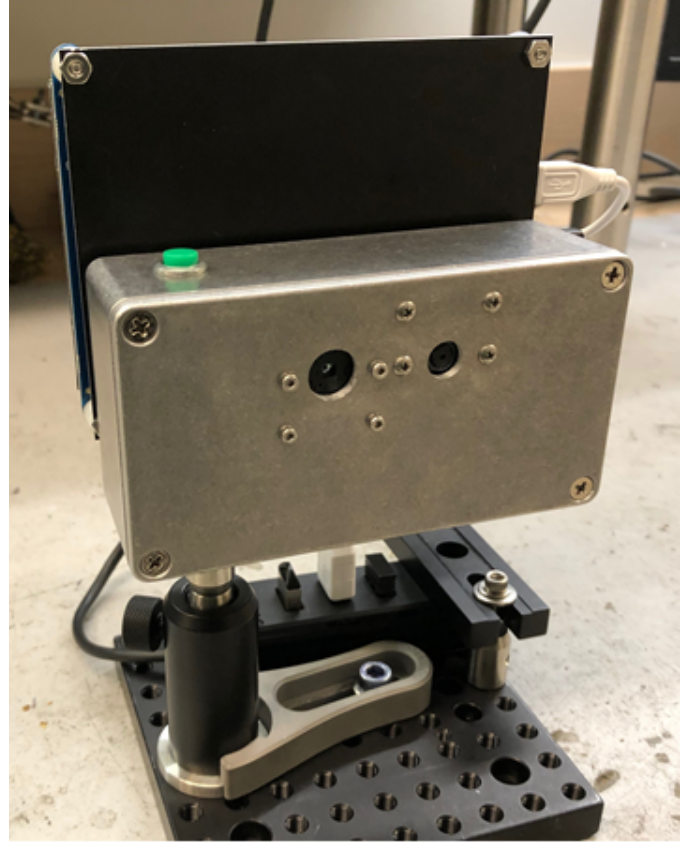

Fig. 1. Front view of finished camera module. The FLIR Lepton (left) and Raspberry Pi Camera V2 (right) can be seen mounted in the camera housing.
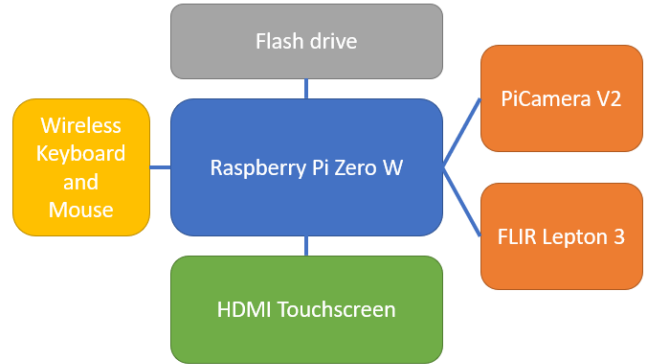

Fig. 2. Block diagram of the LWIR+VIS camera.

capture button is pressed, the live preview is stopped and capture from the FLIR Lepton 3 begins. Image data is from the Lepton is transmitted in a series of 4 packets. However, due to the poor quality of connections, packets are often dropped. The thermal image capture function waits until at least one of each packet has arrived and displays the image constructed from the most recent of each packet.

While the thermal image is being captured, a snapshot of the visible image is displayed on the screen. Once the thermal
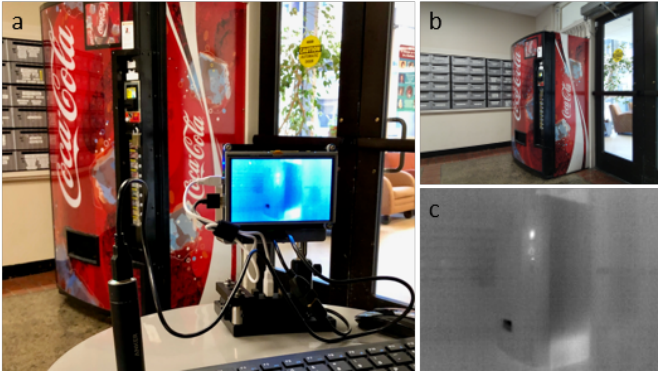
Fig. 3. Example of image capture. a) Rear view of camera module, briefly displaying the captured thermal image. b) Visible and c) thermal images of the captured scene.

Fig. 4. Examples of captured visible and thermal image pairs.

image is assembled, a second visible image is captured in order to minimize the time difference between thermal and visible images. Once the synchronized images are captured, the thermal image is briefly displayed on the display before the camera returns to its idle state. Both images are timestamped and saved on the Raspberry Pi. Figure 3 shows an example of the camera module capturing a scene. The touchscreen displays the captured thermal image of the Coca Cola vending machine. The corresponding visible and thermal images are shown as well. Figure 4 shows three more representative visible and thermal image pairs that were captured using the camera module.

Fig. 5. Staggered packet acquisition results in choppy images of moving objects.
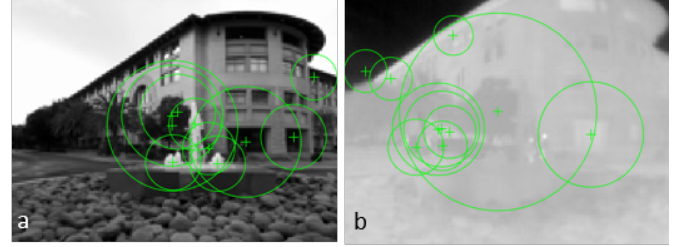

Fig. 6. SURF features extracted from visible and thermal images.

Due to the limitation of dropped packets in the current setup, the camera module can only reliably capture images of still scenes and it is incapable of displaying a live thermal feed. The result of trying to capture to capture a thermal image of moving objects is shown in Figure 5. Packet loss can potentially be minimized by shortening the connections between the Lepton and the Raspberry Pi.

V. IMAGE REGISTRATION

Visible and thermal images store salient features in different ways [9]. For a typical visible image, much of the important information can be extracted from local gradients; an image of a parrot still resembles a parrot after extreme adaptive histogram equalization. The most important features of a thermal image, on the other hand, are typically intensity maxima and minima[1]. Thus while similar keypoints might be extracted from visible and thermal images (see Figure 6), their associated descriptors will in general be completely different. This makes image registration difficult.

Our first attempt at image registration applied the RANSAC algorithm to SURF keypoints extracted from the two images. While refining the code, however, we realized that we could use the fixed camera gemoetry to our advantage. While RANSAC would be necessary for images taken at an unknown relative position, our images were constrained by the fact that the camera modules are mounted next to each other in the camera housing. This allowed us to limit the scope of our registration transformtaion to a rescaling by s followed by a translation by $\vec{r}$.

Our registration algorithm begins by resizing the visible image to match the size of the thermal image scaled by an

---

[1]This mismatch in information between the two spectral bands is related to the scattered vs emissive origins of the visible and thermal image, respectively.
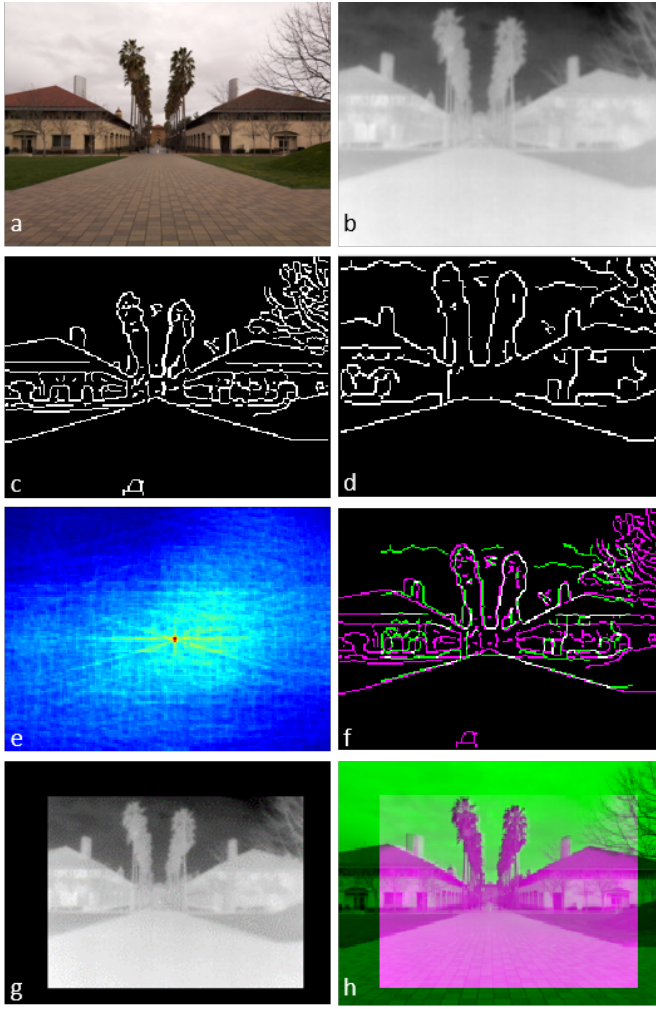
Fig. 7. Edge-based image registration algorithm. a) Visible image. b) Thermal image. c) Canny-detected visible edge mask. d) Thermal edge mask. e) Cross-correlation between edge maps after the thermal image is rescaled by an empirical scale factor s. f) Rescaling and shift associated with the strongest edge-edge correlation. g) Registered thermal image. h) Registered image pair.

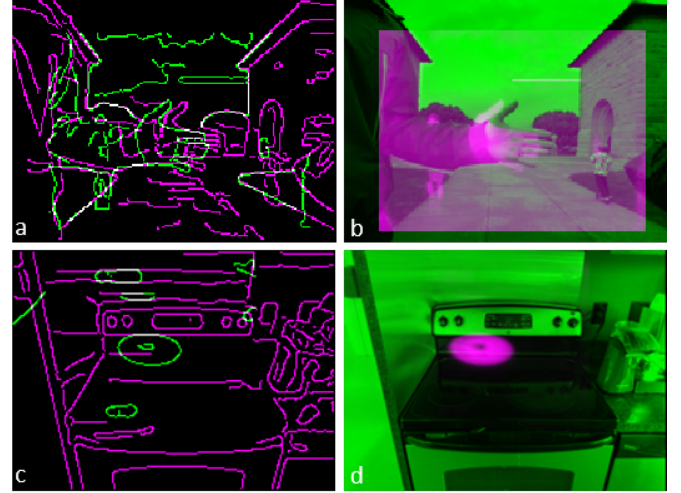| Edge Method	Success Rate
Canny	**0.92**
Log	0.82
zerocross	0.82
Prewitt	0.73
Sobel	0.73
Roberts	0.66

Fig. 8. Failure modes of the edge-based registration. a) and b) Strong edges in the background result in offset foreground objects. c) and d) No common edges means that maximizing the edge-edge correlation does not correctly regsiter the images.

empirical[2] factor s. The Canny edges are then extracted from each, as shown in Figures 7c and 7d. The cross correlation between the two edge masks is calculated from their rotated convolution, as shown in Figure 7e. The maximum correlation value is recorded. This process is repeated as s is scanned between 1.1 and 1.3. The scaling s that maximizes the correlation is chosen and the associated shift $\vec{r}$ is determined (see Figure 7f). With the optimum s and $\vec{r}$ obtained, the thermal image is registered to the visible image and the images are displayed together for verification (Figures 7g and 7h).

We tried different edge detection methods for our algorithm and the results are summarized in Table II. The Canny edge method was able to correctly register 57 out of 62 captured image pairs.

While it was tempting to circumvent the problem of image registration by implementing a fixed transformation, the finite

separation of the camera modules meant that no single transform could compensate for parallax at all depth planes. For instance, a transformation that overlapped objects at infinity would leave closer objects misregistered. By maximizing the edge-edge correlation, our algoritm corrects for parallax at the depth plane with the strongest edges visible in both images. This can result in images that are correclty registered at one depth but offset at another, as seen in Figures 8a and 8b.

Our algorithm also fails when the thermal and the visible images have no common edges. This is the case in the picture of the electric range, shown in Figures 8c and 8d. Here the visible image could not see the outline of the heating element while the thermal image could only see the ouline of the heating element.

VI. IMAGE FUSION METHODS

Once the images are registered we can proceed to image fusion. We begin by considering the important aspects of each image. Steady state solutions to the heat equation are typically diffuse, so in general our thermal image will not have sharp gradients[3]. Thus the salient features of a thermal image are typically intensity maxima and minima that highlight warm

---

[2]This rescaling comes from the camera modules' different fields of view.

[3]One common exception is at object edges, where an object at one temperature ends and a background object at a different temperature appears.

and cool regions, respectively. In contrast, visible images carry much more information in their gradients. A solid object, for example, could have strong visible gradients in a small region. The same object would likely have much slower thermal gradients.

Under these assumptions, our image fusion methods should combine the intensity information from the thermal image with the gradient information from the visible image. We outline three such methods below, before presenting results and discussion in Section VII.

A. YCbCr Merge

In the RGB color space intensity and gradient information is spread across all color channels. Because we are interested in fusing information from two sources, it would be helpful to work in a color space where the salient features of each image are confined to a single channel. This motivates our first image fusion method based on the YCbCr color space.

We begin by converting the grayscale thermal image into an RGB image with a desired colormap. This encodes the spectral intensity information of the grayscale thermal image into chrominance information of the RGB thermal image. Next we convert the thermal and visible images from RGB to YCbCr. We replace the luminance (Y) component of the thermal image with the corresponding component from the visible image. Finally, we convert the fused image from YCbCr to RGB. The final product has the gradient information from the visible image encoded in its luminance and the intensity information from the thermal image encoded in its chrominance.

B. High Pass Filtering

The YCbCr method relies on having multiple channels in the fused image. If this was not an option, we must consider of fusing the images within a single channel.

Given that we are interested in the gradient information from the thermal image and the intensity information from the thermal image, one simple fusion method is to extract the high frequency information from the thermal image and inject it into the inherently low frequency thermal image. This is called the High Pass Filtering (HPF) method.

We begin by applying a Gaussian low pass filter to the visible image $\mathbf{v}$:

$$\mathbf{v}_{LP} = \texttt{imgaussfilt}(\mathbf{v}, \sigma) \qquad (1)$$

where σ is the standard deviation of the Gaussian kernel. Next we isolate the high frequency components of the visible image by subtracting its low passed components:

$$\mathbf{v}_{HP} = \mathbf{v} - \mathbf{v}_{LP} \qquad (2)$$

Finally we form the fused image $\mathbf{x}$ by injecting the high frequency visible components into the thermal image $\mathbf{t}$:

$$\mathbf{x} = \mathbf{t} + \mathbf{v}_{HP}. \qquad (3)$$

Whereas the YCbCr method is fully constrained, the HPF method has one free parameter, σ. In order to make our algorithm independent of input image sizes, we introduce the variable α to parameterize σ:

$$\sigma(\alpha) = \alpha \times \frac{1}{2} \times \frac{\texttt{size}(\mathbf{v})}{\texttt{size}(\mathbf{t})} \qquad (4)$$

When $\alpha = 1$, this parameterization returns the optimum cutoff frequency σ considering the maximum frequency component present in the thermal image.

C. Gradient Transfer Function

In order to avoid the spectral distortion introduced by HPF at strong α, we consider one final method. Recalling that we want to preserve the intensity information of the thermal image while injecting the gradient information from the visible image, we construct the following objective function [9]:

$$\varepsilon(\mathbf{x}) = \frac{1}{2}|\mathbf{x} - \mathbf{t}|_2^2 + \lambda|\boldsymbol{\nabla}\mathbf{x} - \boldsymbol{\nabla}\mathbf{v}|_1 \qquad (5)$$

The first term in this objective function seeks to match the intensity of the fused image to the intensity of the thermal image. The second term wants to inject the gradient information from the visible image into the fused image. The relative weight of these terms' contributions is set by λ. For a given value of λ, the fused image $\mathbf{x}$ is found by minimizing the objective function $\varepsilon(\mathbf{x})$:

$$\mathbf{x} = \arg\min_{\mathbf{x}} \varepsilon(\mathbf{x}) \qquad (6)$$

Our implementation of this method uses an ADMM+TV solver in which we have modified the regulation term to minimize the total variation between the visible and the thermal image.

D. Other Methods

We also experimented with wavelet, Brovey, bilateral, and high pass modulation fusion methods. Because these methods are discussed extensively in existing review papers [16, 17] and our results were not substantially different from HPF and GTF, we have omitted them from this report.

VII. FUSION EVALUATION

A. Metrics

Before analyzing the results of our single channel methods we must first identify a set of useful evaluation metrics. In addition to the standard PSNR and SSIM, we are also interested in looking at metrics that quantify the amount and type of information in each image. One popular metric is the quality with no reference (QNR) index [16]. In general this index quantifies both spectral and spatial distortion of a multispectral image. Because our thermal image is the only band in our "multispectral" image (the visible image is playing the role of the high resolution panspectral image), we are only concernd with the spatial distortion term:

$$QNR = 1 - |Q(\mathbf{f}, \mathbf{v}) - Q(\mathbf{t}, \mathbf{v}_{LP})| \qquad (7)$$

Here the Q index [31], defined for two signals x and y, is given by

$$Q(x, y) = \frac{4\sigma_{xy}\bar{x}\bar{y}}{\left(\sigma_x^2 + \sigma_y^2\right)\left[(\sigma_x)^2 + (\sigma_y)^2\right]} \qquad (8)$$

Finally, because we are ultimately interested in the amount of information contained in the final image, we will include its entropy in the list of metrics. Larger numbers are better for all metrics.

B. Results

Representative results of the YCbCr, HPF, and GTF fusion methods are shown in Figure 11.

We begin by considering the YCbCr results qualitatively. Because the contributions from the thermal and visible images are contained in separate channels in the YCbCr color space, the information is displayed without thermal ambiguity (i.e. warm regions have a warm chrominance, regardless of their visible texture). More images are shown in Appendix C, Figure 12.

Qualitatively the HPF and GTF methods look very similar for $\alpha = 1$ and $\lambda = 0.1$. In Figure 11, one of the only clear differences between the two methods is a loss of thermal gradient detail in the image of the electric range for the GTF method. It is worth noting, however, that the stove had been turned on shortly before the image was taken; the gradients likely would have gotten weaker had we given the range time to thermalize.

Figure 9 shows the metrics for the HPF and GTF methods over a range of α and λ for the heat vent scene shown in Figures 10a and 10b. Representative images for a range of parameter values can be found in Appendix C, Figures 13 and 14. All metrics are nomalized with respect to the original thermal image, i.e. their values when $\alpha, \lambda = 0$. Here the visible image has been blurred and downsampled to match the size of the original thermal image. Prior to fusion, the thermal image was blurred and downsampled by a factor of 2. The original thermal image was used as ground truth.

As α and λ increase, the growing entropy indicates that the visible image is transferring information into the fused result. While the PSNR and SSIM decrease for both methods, it is worth noting that their saturation values are lower for the HPF method than the GTF method. The QNR index shows the largest difference between the two methods, with the HPF QNR becoming negative while the GTF QNR saturates at a finite positive value.

These limit results can be seen qualitatively in Figure 10. Of note is the intensity difference between the heat vent in the foreground and the buildings in the background. In Figure 10c we see that the HPF method has matched the intensities of the cold buildings with that of the warm vent, distorting the thermal information. This is consistent with the strong decrease in the HPF QNR. In contrast, Figure 10d shows a suppression of the building intensity, enforced by the $|\mathbf{x} - \mathbf{t}|_2^2$ term in the objective function. We do, however, also see a decrease in the intensity of the heat vent, likely due to its strong visible gradients dominating in the $|\nabla\mathbf{x} - \nabla\mathbf{v}|_1$ term of the objective function.

Finally, we note that the metrics in Figure 9 for the heat vent image are represenatative of the entire dataset, as shown in Appendix C, Figure 15.
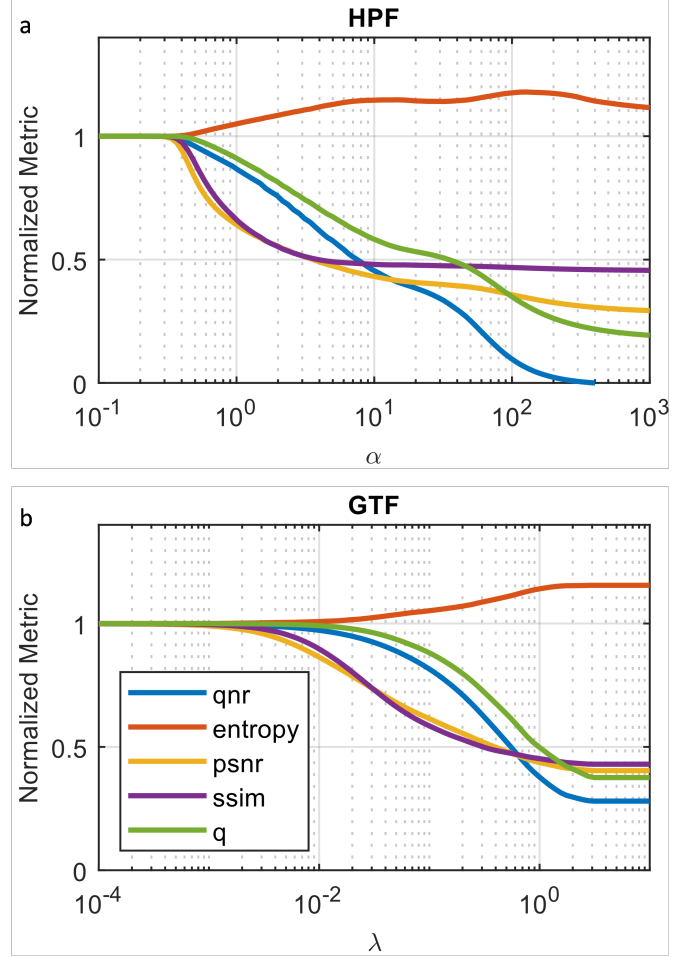

Fig. 9. Evaluation metrics for a) HPF fusion and b) GTF fusion. For each method, the metrics are normalized to their values at $\alpha, \lambda = 0$: QNR = 0.9899, entropy = 6.1625, PSNR = 41.6, SSIM = 0.9781, Q = 0.9968.
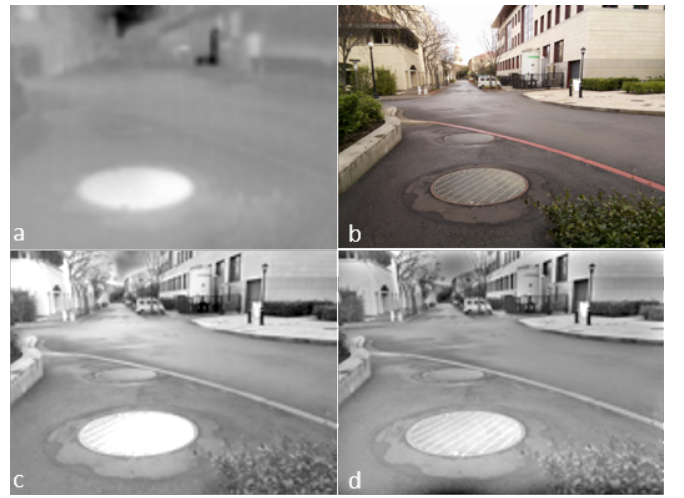

Fig. 10. Limiting cases for HPF and GTF. a) Original thermal image. b) Original visible image. c) HPF fusion with $\alpha = 100$. d) GTF fusion with $\lambda = 100$.
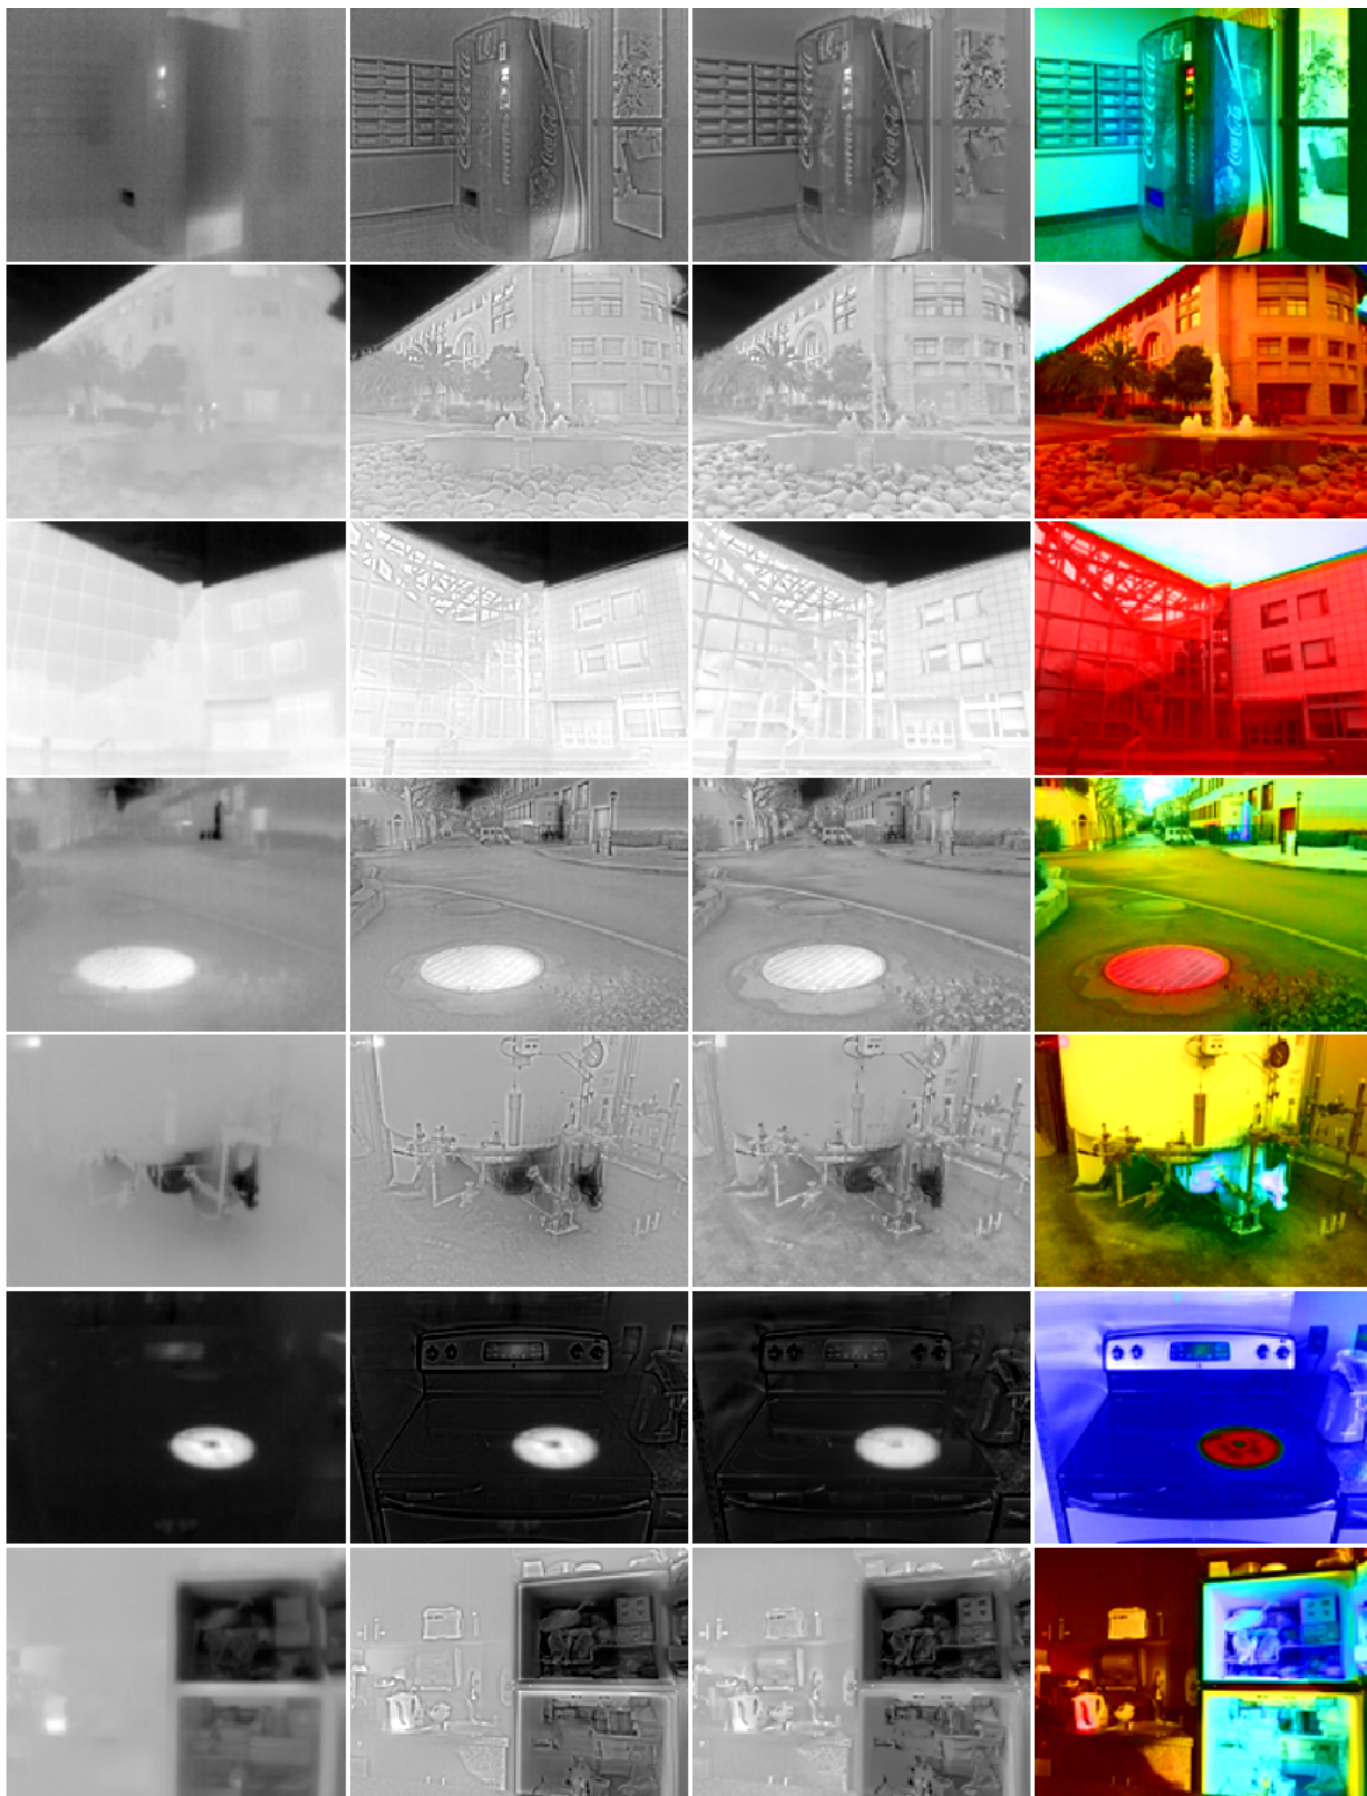
Fig. 11. Fusion montage. From left to right: original thermal image, HPF fusion ($\alpha = 1$), GTF fusion ($\alpha = 0.1$), YCbCr fusion.

VIII. Future Work

From a hardware perspective, the first goal in future work would be to establish a better connection protocol with the Lepton. Currently, the poor data transmission severely limits the frame rate and prevents the camera from capturing any scenes with motion. Utilizing the full nine frames per second capability of the camera opens the door for a variety of other goals. Namely, it would be possible to perform basic image registration and fusion on the Raspberry Pi itself. Image registration is already a very fast technique, but we could use a simple method for image fusion such as HPF or YCbCr that is computationally inexpensive to generate a live feed of the enhanced thermal images. On-board, real-time image fusion would allow the dual visible-thermal camera module to more closely replicate the functionality of more expensive thermal cameras.

From a physical perspective, it would be a simple matter to calibrate the readouts from the thermal camera to real temperatures. Then, the thermal images would no longer be displayed on an arbitrary grey scale, but as real object temperatures – a significantly more useful metric.

Finally, there are many potential avenues for improving the algorithms used to register and merge the images. One improvement to our registration algorithm is the segmentation of the thermal image prior to registration. Each segment would then be independently registered according to the depth of its dominant edges. This local parallax correction would allow more objects at different depths to be correctly registered. Because the thermal image is typically sparse, the registered image segments could then easily be interpolated back together to form the final registered image.

If we had access to the full nine frames per second of the camera, it would then make sense to include runtime with the metrics used to compare the fusion methods. It is likely that the YCbCr method would win based on calculational simplicity and usefulness in real-time imaging applications.

IX. Conclusion

Overall, the combination of a low-cost thermal camera with a visible camera produced promising results. Using image fusion algorithms, we are able to successfully inject information from the visible image into the thermal image. Qualitatively the images look much better to the human eye, but the HPF and GTF methods can introduce false temperature information into the fused image. Thus, significant work remains if this setup is desired to replace actual higher resolution thermal cameras, especially for applications in which high spectral accuracy is required.

Appendix A
Author Contributions

Jacob Hines (EE 367, EE 368):

- Machined and assembled the camera housing
- Implemented image registration and fusion
- Wrote the second half of the final report

Evan Wang (EE 367):

- Wrote code to capture and store images from both camera modules on the Raspberry Pi
- Wrote the first half of the final report

Any task not listed above was shared equally between the authors.

Appendix B
Bill of Materials

TABLE III
Camera module part list

Item	Cost
Lepton 3.5	$259.00
Lepton Breakout Board	$39.99
Raspberry Pi Zero W Kit	$59.95
R Pi Camera Module V2	$24.68
R Pi Zero v1.3 Camera Cable	$5.95
Adafruit 5" HDMI Display	$74.95
Wireless Keyboard and Mouse	$19.95
Total	**$484.47**

References

[1] SpectralCalc.com. Spectral calculator, 2019. http://www.spectralcalc.com/blackbody_calculator/blackbody.php.

[2] JEAN BRUNELLE DALSA, TELEDYNE. Thermal Imaging Spurs New Applications.

[3] FLIR Boson Compact LWIR Thermal Camera OEM-Cameras.com | The Highest Quality OEM and Thermal Cameras.

[4] Lepton LWIR Micro Thermal Camera Module | FLIR Systems.

[5] L. St-Laurent, X. Maldague, and D. Prevost. Combination of colour and thermal sensors for enhanced object detection. In *2007 10th International Conference on Information Fusion*, pages 1–8, July 2007.

[6] James W. Davis and Vinay Sharma. Background-subtraction using contour-based fusion of thermal and visible imagery. *Computer Vision and Image Understanding*, 106(2):162–182, May 2007.

[7] S. Hwang, J. Park, N. Kim, Y. Choi, and I. S. Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1037–1045, June 2015.

[8] L. Bienkowski, C. Homma, K. Eisler, and C. Boller. Hybrid Camera and Real-View Thermography for Non-destructive Evaluation. In *Proceedings of the 2012 International Conference on Quantitative InfraRed Thermography*. QIRT Council, 2012.

[9] Jiayi Ma, Chen Chen, Chang Li, and Jun Huang. Infrared and visible image fusion via gradient transfer and total variation minimization. *Information Fusion*, 31:100–109, September 2016.

[10] Johan Johansson, Martin Solli, and Atsuto Maki. An Evaluation of Local Feature Detectors and Descriptors for Infrared Images. In Gang Hua and Hervé Jégou, editors, *Computer Vision – ECCV 2016 Workshops*, Lecture

Notes in Computer Science, pages 711–723. Springer International Publishing, 2016.

[11] Barbara Zitová and Jan Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, October 2003.

[12] Leila M G Fonseca and B S Manjunath. Registration Techniques for Multisensor Remotely Sensed Imagery. page 8, 1996.

[13] and B. S. Manjunath and S. K. Mitra. A contour-based approach to multisensor image registration. *IEEE Transactions on Image Processing*, 4(3):320–334, March 1995.

[14] Enrique Coiras, Javier Santamaria, and Carlos Miravet. Segment-based registration technique for visual-infrared images. *Optical Engineering*, 39(1):282–290, January 2000.

[15] Jungong Han, Eric J. Pauwels, and Paul de Zeeuw. Visible and infrared image registration in man-made environments employing hybrid visual features. *Pattern Recognition Letters*, 34(1):42–51, January 2013.

[16] G. Vivone, L. Alparone, J. Chanussot, M. Dalla Mura, A. Garzelli, G. A. Licciardi, R. Restaino, and L. Wald. A Critical Comparison Among Pansharpening Algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, 53(5):2565–2586, May 2015.

[17] Xiangchao Meng, Huanfeng Shen, Huifang Li, Liangpei Zhang, and Randi Fu. Review of the pansharpening methods for remote sensing images based on the idea of meta-analysis: Practical discussion and challenges. *Information Fusion*, 46:102–113, March 2019.

[18] H. Li, B. S. Manjunath, and S. K. Mitra. Multisensor Image Fusion Using the Wavelet Transform. *Graphical Models and Image Processing*, 57(3):235–245, May 1995.

[19] J. Zhou, D. L. Civco, and J. A. Silander. A wavelet transform method to merge Landsat TM and SPOT panchromatic data. *International Journal of Remote Sensing*, 19(4):743–757, January 1998.

[20] J. Nunez, X. Otazu, O. Fors, A. Prades, V. Pala, and R. Arbiol. Multiresolution-based image fusion with additive wavelet decomposition. *IEEE Transactions on Geoscience and Remote Sensing*, 37(3):1204–1211, May 1999.

[21] Krista Amolins, Yun Zhang, and Peter Dare. Wavelet based image fusion techniques — An introduction, review and comparison. *ISPRS Journal of Photogrammetry and Remote Sensing*, 62(4):249–263, September 2007.

[22] Alan R Gillespie, Anne B Kahle, and Richard E Walker. Color enhancement of highly correlated images. II. Channel ratio and "chromaticity" transformation techniques. *Remote Sensing of Environment*, 22(3):343–365, August 1987.

[23] Pat S Chavez. Comparison of Three Different Methods to Merge Multiresolution and Multispectral Data:Landsat TM and SPOT Panchromatic. *PHOTOGRAMMETRIC ENGINEERING*, page 9, 1991.

[24] M. Gonzalez-Audicana, J. L. Saleta, R. G. Catalan, and R. Garcia. Fusion of multispectral and panchromatic images using improved IHS and PCA mergers based on wavelet decomposition. *IEEE Transactions on Geoscience and Remote Sensing*, 42(6):1291–1299, June 2004.

[25] Hamid Reza Shahdoosti. MS and PAN image fusion by combining Brovey and wavelet methods. *arXiv:1701.01996 [cs]*, January 2017. arXiv: 1701.01996.

[26] Mohamed Ghadjati, Abdelkrim Moussaoui, and Abdelhak Boukharouba. A novel iterative PCA–based pansharpening method. *Remote Sensing Letters*, 10(3):264–273, March 2019.

[27] FLIR. Boson, 2019. https://www.flir.com/products/boson/.

[28] Dave Jones. picamera. https://github.com/waveform80/picamera, Nov 2018.

[29] Novacoast. Lepton-3-module. https://github.com/novacoast/Lepton-3-Module, Aug 2018.

[30] GroupGets. Leptonmodule. https://github.com/groupgets/LeptonModule, Oct 2017.

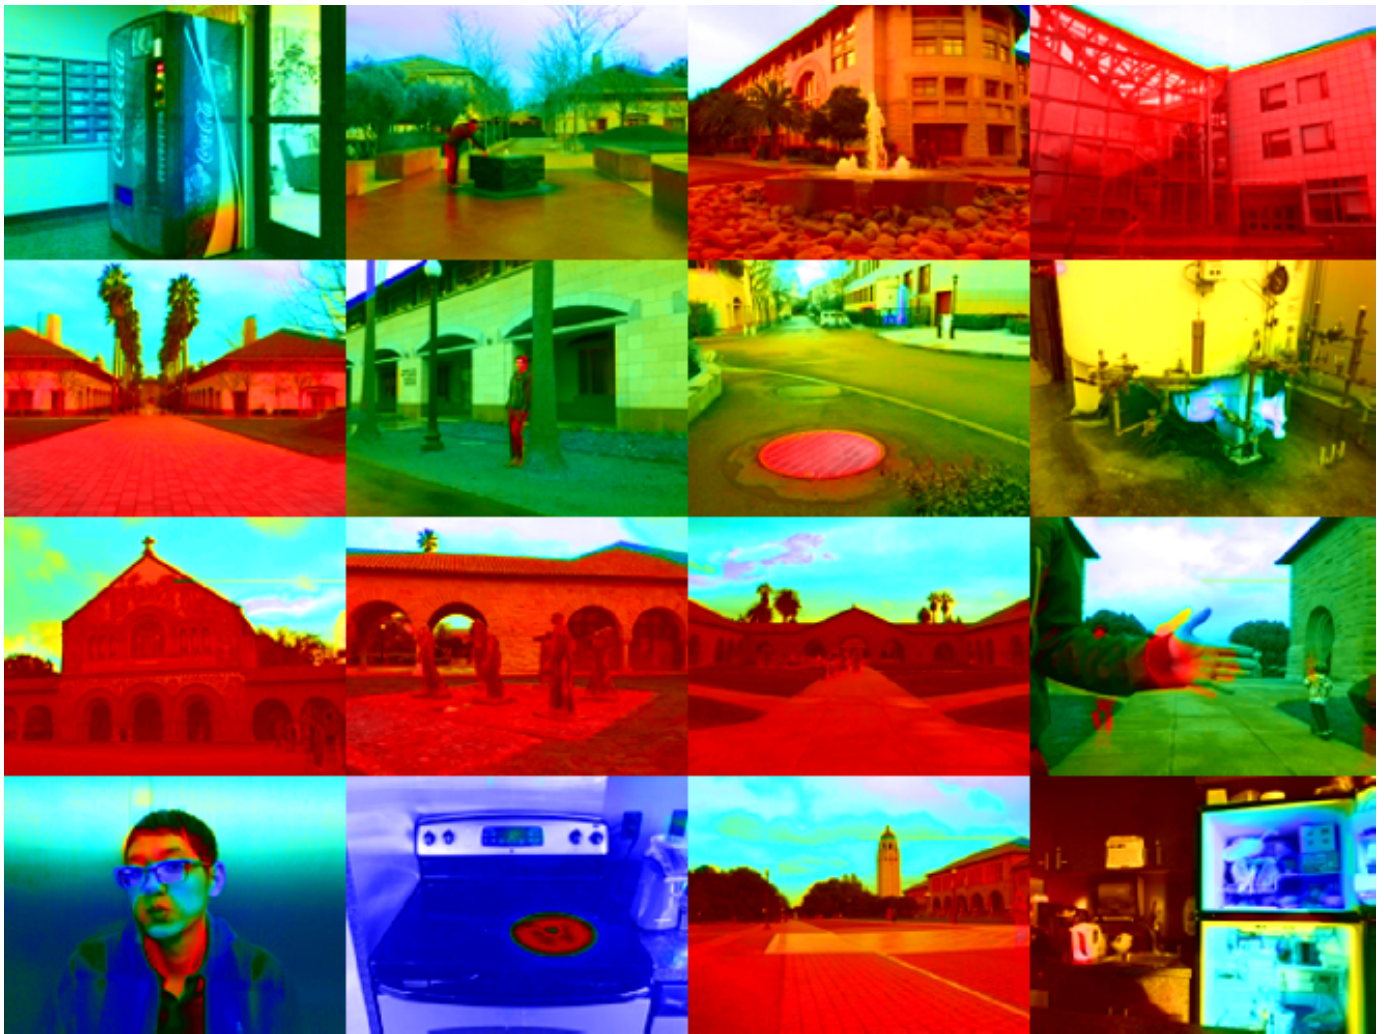[31] A universal image quality index - IEEE Journals & Magazine.

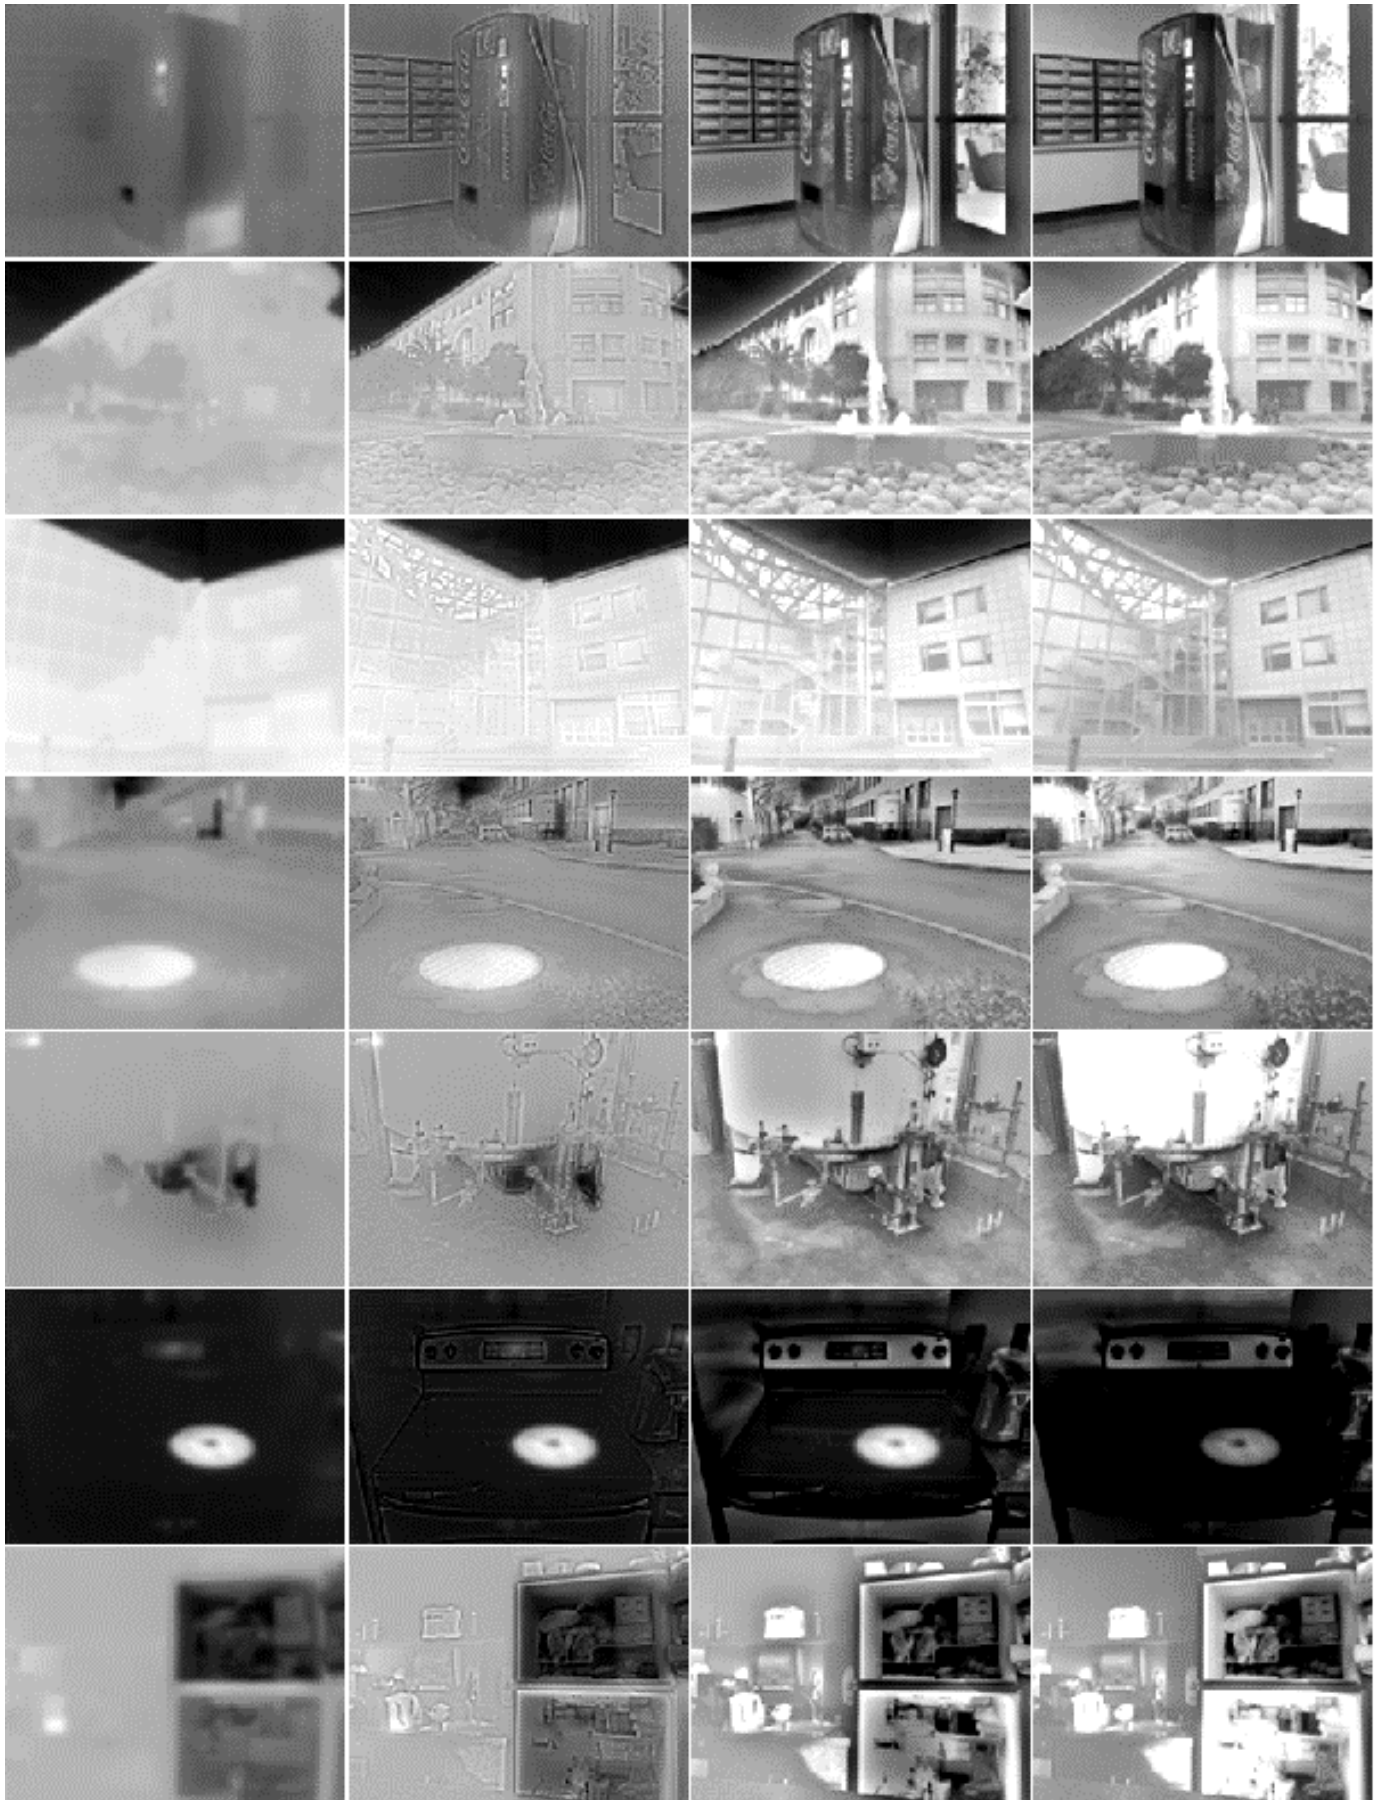Fig. 12. YCbCr fusion for a variety of image pairs.

Fig. 13. High Pass Filtering fusion. From left to right: $\alpha = 0, 1, 10, 100$.
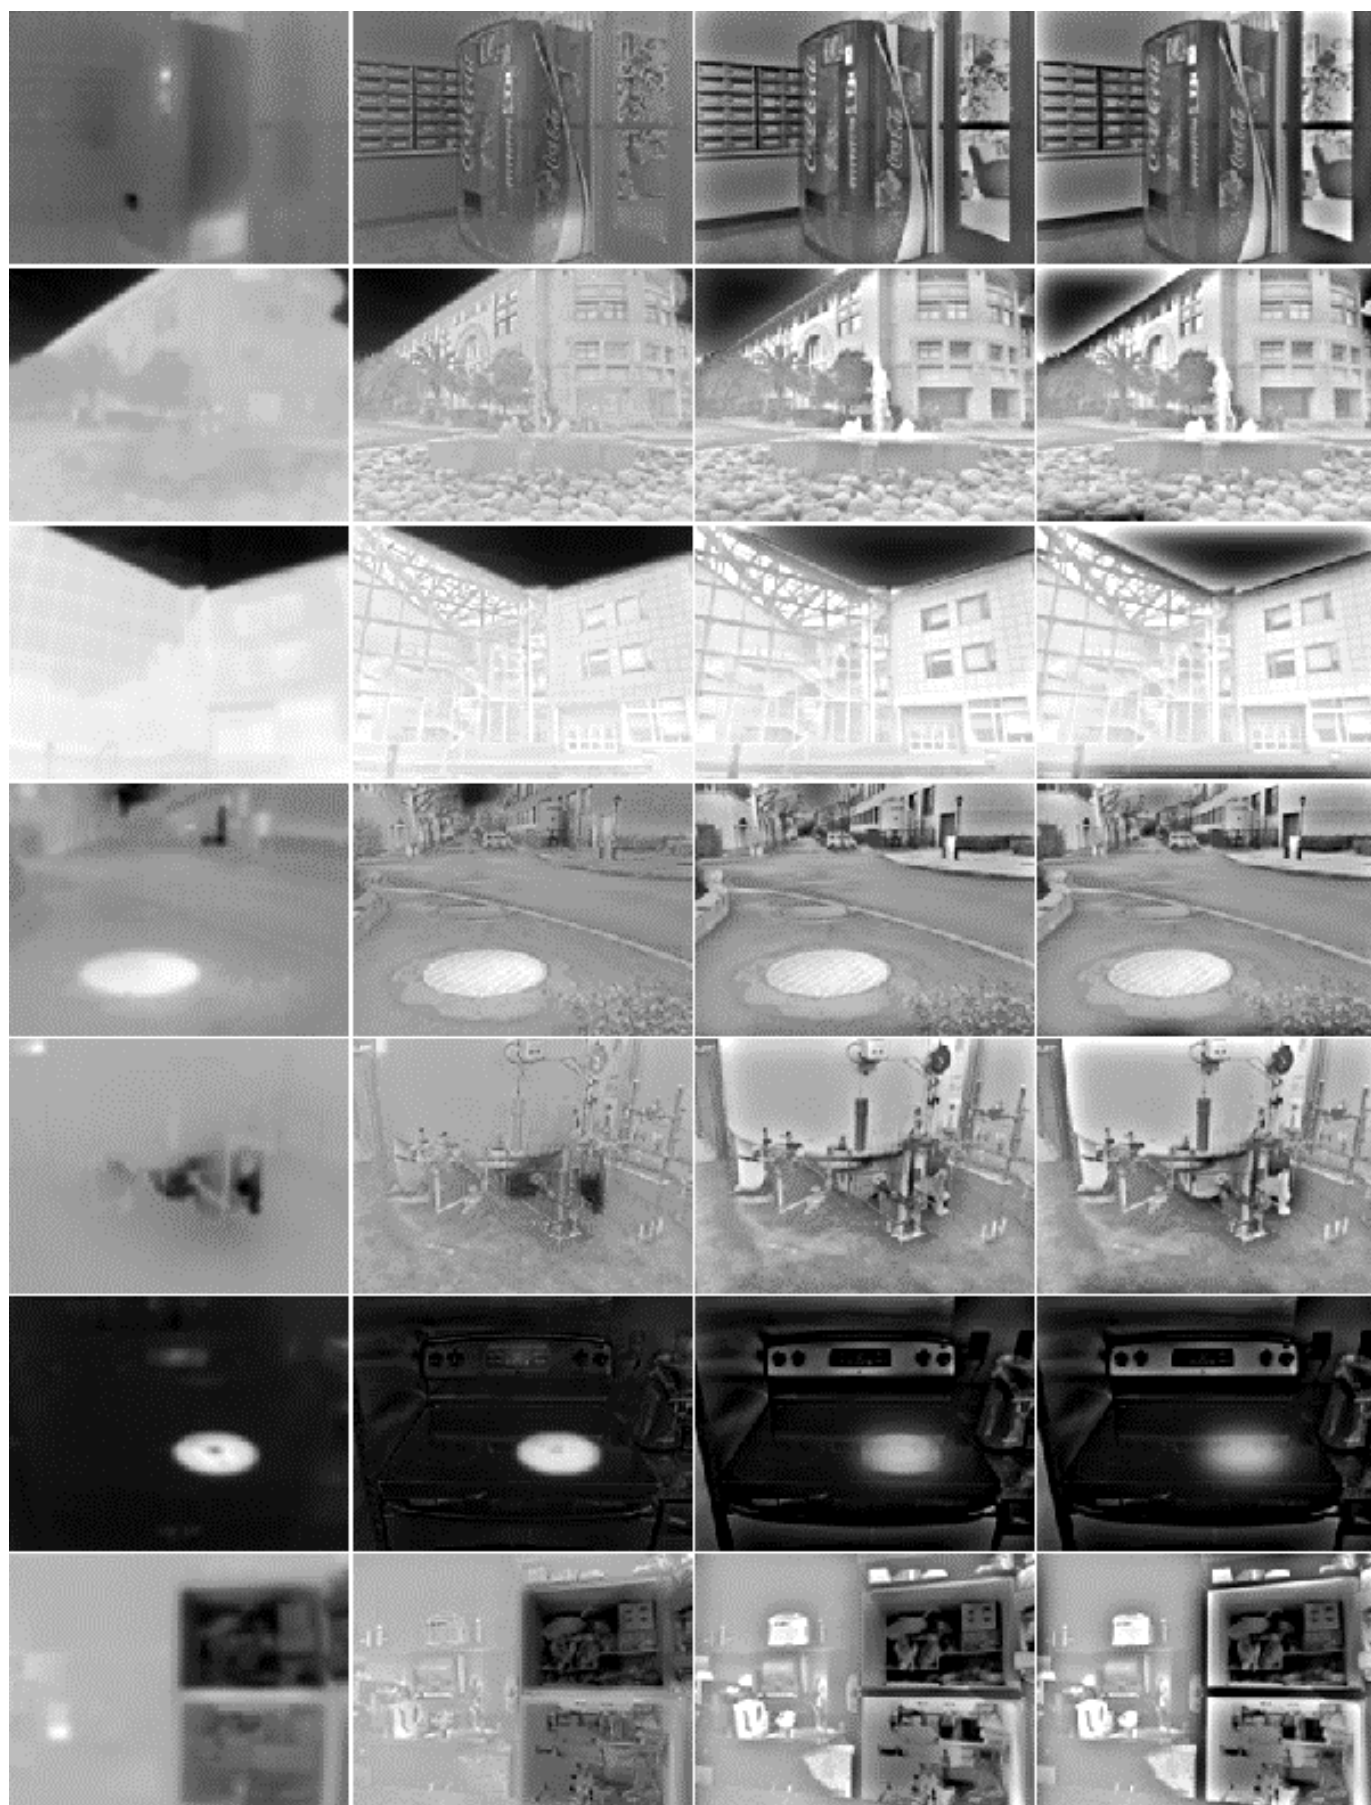
Fig. 14. Gradient Transfer Function fusion. From left to right: $\lambda = 0, 0.1, 1, 10$.
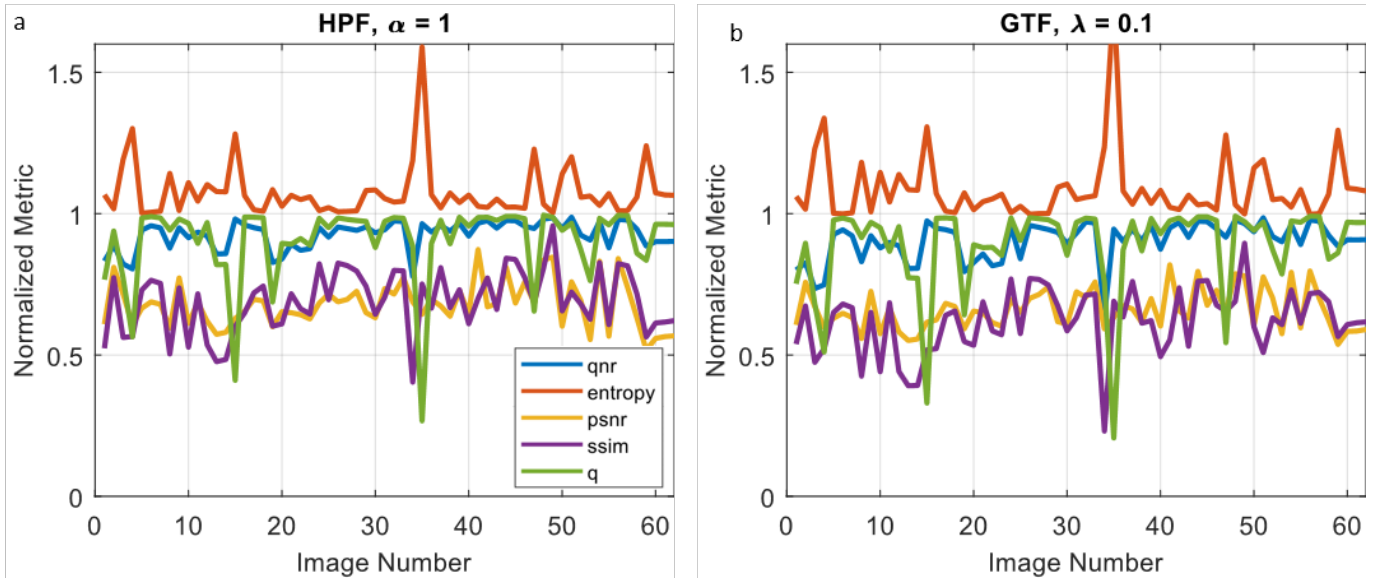
Fig. 15. Fixed fusion paramer across all images. a) HPF fusion with $\alpha = 1$. b) GTF fusion method with $\lambda = 0.1$. Images that report large spikes in metrics were incorrectly aligned by the registration algorithm.