

# Spectra: 3D Multispectral Fusion and Visualization Toolkit

Peyman Moghadam<sup>1\*</sup>, Stephen Vidas<sup>2</sup>, and Obadiah Lam<sup>1</sup>

<sup>1</sup>Autonomous Systems, CSIRO  
Pullenvale, QLD Australia

Peyman.Moghadam@csiro.au, Obadiah.Lam@csiro.au

<sup>2</sup>School of Electrical and Electronic Engineering  
Nanyang Technological University, Singapore  
svidas@ntu.edu.sg

## Abstract

There is an increasing demand for accurate and intuitive representation of both thermographic and multispectral data in 3-Dimensional models. Existing fusion methods can be used to compensate for the typically low spatial resolution of multispectral images with high spatial resolution visible-spectrum images, overlaying the combined result on a 3D model. However, because of a lack of appropriate methods, thermal information has in the past been displayed on a separate model in a separate screen, which increases the cognitive workload of the user. To address this, we present a suite of novel mesh-based fusion methods that can represent both thermal and visible-spectrum data simultaneously on a single 3D model, highlighting regions of extreme temperature or informative temperature variation. These methods are targeted towards different applications and as a result have unique advantages. In addition, a 3D multispectral visualization toolkit “Spectra” has been developed. This visualizer allows the user to customize settings for each fusion method, including scaling and thresholding factors and many method-specific parameters. The “Spectra” toolkit is used to efficiently assess the effectiveness of the fusion methods for a variety of tasks.

## 1 Introduction

Remotely operated search and rescue robots may augment the capabilities of a conventional visible-spectrum camera by utilizing thermal sensors to detect the thermal signature of trapped victims [Jacoff *et al.*, 2000]. This thermal and visible data must be viewed intuitively by the remote user to control the robot. Thus we require a more photorealistic view of the scene [Se and Jasiobedzki, 2006]. To address the high cognitive effort required to view two separate modalities simultaneously, we propose methods of 3D multispectral fusion that




Figure 1: A multispectral 3D model of a compressor unit. The colour bar on the left shows the temperature range of the scene. The model was generated using the proposed visualization toolbox and its implementation of the novel “Intensity-Hue Mapping” method.

enable both thermal and colour data to be fused and visualized simultaneously in 3D.

The usefulness of thermal-visible fusion has already been identified in [Rasmussen *et al.*, 2009], where a fusion method using colour visible and thermal 2D images was developed. While existing work deals mostly with greyscale/thermal or greyscale/multispectral fusion [Schnelle and Chan, 2011] [Malpica *et al.*, 2013], we propose in this paper thermal/visible fusion methods that use colour-mapped thermal data instead of greyscale intensities (Figure 1). By increasing the amount of information that can be conveyed to the user, this also has the effect of allowing the representation of visible features (*i.e.*, high frequency features) such as writing or labels that are invisible in thermal-infrared domain.

3D models have several distinct advantages over 2D images, such as allowing the user to view the scene from different angles, including novel viewpoints not experienced during data capture. However, previous research has focused on pixel fusion. Even in the case of 3D data the fusion is conducted in 2D and then overlaid into the 3D model [Waxman *et*

\*Corresponding author

*et al.*, 2001]. This requires 2D colour pixel fusion followed by 3D registration [Ross *et al.*, 2000]. To overcome this limitation, this paper proposes 3D mesh fusion to reduce the uncertainty and error in the fusion as values are assigned individually to each vertex. This solves the ambiguity in pixel fusion which arises from positional error in assigning values from low resolution thermal data that has a larger pixel size than the higher resolution visible data [Xie *et al.*, 2004]. Spatial relationships between physical features within the scene can also be determined accurately, which in the case of 3D thermography enables quantitative analysis of spatially-dependent behaviour such as heat flow.

Existing software for visualization of 3D meshes is limited to the visible-spectrum, which has been sufficient for the overwhelming majority of 3D data. When multiple modalities are used, they are presented side-by-side in two separate screens, which slows the user as they have to process two separate streams of information at the same time, switching their gaze from one screen to the other. These visualizers often do not allow the remote user to easily change the visualization scheme and settings. This has been a significant motivating factor for the work presented in this paper, and is addressed by our second contribution.

This paper has two main contributions. First, we propose new 3D thermal/visible mesh fusion methods. A common theme of these methods is that they seek to both boost the spatial resolution of the 3D thermal representation by incorporating visible-spectrum data, while also highlighting regions of interest (e.g. thermal extremities) to the end user. Eight proposed 3D mesh fusion methods, including the Intensity-Hue Mapping method shown in Figure 1, can be found in Section 3.

The second major contribution is the “Spectra” 3D multi-spectral visualization toolkit, which is presented to visualize the fused meshes. This visualizer application is the first 3D multispectral visualizer, as well as the first application to offer fusion implementation and visualization in real-time. This was developed in order to improve the intuitiveness of 3D fused imagery to suit the needs of the user and application. Users can customise the visualization to their personal preferences using a Graphical User Interface (GUI) and see the effects of their changes immediately. This visualization toolkit is demonstrated in Section 4.

The effectiveness of both the proposed fusion methods and the visualization toolkit are discussed in Section 5. The demonstration videos of the proposed 3D multispectral visualization toolkit can be found online.<sup>1</sup>

## 2 HeatWave: 3D Thermography System

Recently, we developed a new mobile hand-held 3D thermography device called HeatWave [Vidas and Moghadam, 2013b] [Moghadam and Vidas, 2014]. HeatWave is a small hand-



Figure 2: HeatWave 3D thermography device.

held 3D thermography and ranging device that is able to generate a precise 3D model of an object or scene, which is overlaid with view-independent temperature and appearance information in real-time. HeatWave consists of a thermal camera, a range sensor and visible camera, all rigidly attached together in an ergonomic form factor (Figure 2).

As an operator moves the device around target objects, the range sensor captures the geometric information and generates a precise 3D model of the environment by incorporating information from each new frame and simultaneously estimating the device position and orientation (i.e. pose) with respect to the model in real-time. Given the estimate of the pose of the device, appearance information and temperature distributions from visible and thermal cameras respectively are overlaid on the 3D model.

The thermal-infrared camera in our current HeatWave system is an Optris PI450 with a resolution of 384 by 288 pixels, captured at 80 frames per second. This camera contains an uncooled microbolometer to detect radiation emitted from the surfaces of objects relative to their temperature with wavelengths within the long-wave infrared (LWIR) band of the electromagnetic spectrum (7–14 $\mu$ m). A low cost ASUS Xtion Pro Live RGB-D (colour and range) sensor operating at 30 frames per second and a resolution of 640 by 480 pixels has been used in the current HeatWave system. The Xtion has an effective range of up to four meters which is ideal for close-range 3D thermography applications.

The current hand-held HeatWave device is designed for first responders and rescue crews where they hold the device while walking and mapping the environment in 3D. It can also be mounted on a robotic arm attached to a mobile rescue robot for fine motion control. The base mobile robot and robotic arm can be tele-operated or autonomously moved to various poses around the environments to complete the 3D model. The output of HeatWave is a 3D multispectral model with high accuracy, view independent augmented temperature and visible information. Visualization options for these rich and complex 3D multispectral point clouds are currently very limited. Typically, the 3D point clouds are displayed to an expert user or a remote operator in their separate modalities (visible 3D model or thermal 3D model) independently with two side-by-side windows, as seen in the ThermalMapper Project

<sup>1</sup><https://wiki.csiro.au/display/ASL/HeatWave>

[Borrmann *et al.*, 2012] and [Nagatani *et al.*, 2014]. However, such visualization increases the cognitive workload of the operator by presenting the information in two separate models. Users will need to switch their gaze from one screen (with one modality) to the other to gain knowledge of the robots surroundings and status. Insufficient knowledge in a search and rescue task may result in the loss of human lives [Keyes *et al.*, 2010].

The objective of this paper is to address these limitations and to propose an intermediary step by keeping the “human-in-the-loop.” “Spectra”, a novel 3D fusion and visualization toolkit is presented, which enables the dense, high-resolution 3D multispectral model to be visualized in either modality independently or both modalities simultaneously.

### 3 Thermal-Visible Fusion Methods

In order to decrease the response time and reduce the cognitive workload of the user or remote user, a number of fusion methods are proposed in this paper and implemented in “Spectra”. Each of these fusion methods is customisable according to the user’s preferences, and each method has advantages in specific scenarios and applications. For example, the thermal-only mesh shown in Figure 3a is sufficient to convey thermal information and highlights hot and cold regions well, but lacks the visible information seen in Figure 3b which includes useful details such as the text on the screen and the label on the phone.

The 3D thermal-visible fusion methods presented here have been developed with the end goal of highlighting irregularities and conveying information unambiguously. In order to compensate for varying environments of operation, scale factors are included in the methods to adapt the fusion methods to each unique situation. These fusion methods not only highlight irregular regions of interest in the datasets, they also allow the user to read and examine important features such as writing or labels. The presented fusion methods can be used to combine a visible colour mesh with a single corresponding thermal-infrared mesh.

Throughout this paper, the following notations have been used.  $F(x, y, z)$  denotes the fused output mesh at the vertex  $(x, y, z)$ . Similarly,  $V(x, y, z)$  refers to the visible mesh, and  $T_{map}(x, y, z)$  refers to the thermal colour-mapped mesh.  $T_{raw}(x, y, z)$  refers to the temperature of the vertex in the raw thermal mesh (a single constant value), with the temperature measured in degrees Celsius.

#### 3.1 Thermal-Visible Blend

A simple way to fuse and display both modalities is the Thermal-Visible Blend [Bican *et al.*, 2002]. We have extended this method from 2D pixel fusion and applied it to 3D mesh fusion. It is low on computation time but offers limited customization and can reduce contrast. The implementation is shown in Equation 1 and an example output image is presented in Figure 4a.

$$F(x, y, z) = T_{map}(x, y, z)\alpha + V(x, y, z)(1 - \alpha) \quad (1)$$

The  $\alpha$  value is the only configurable variable for this method.

#### 3.2 Monochrome Threshold

In the 2D fusion case, Monochrome Thresholding has been shown by Rasmussen to reduce the user’s cognitive workload [Rasmussen *et al.*, 2009]. This advantage is expected to carry over to our 3D adaptation. The monochrome thresholding blends the visible RGB mesh with a monochrome thermal mesh, mapped to varying intensities of a chosen colour based on the temperature information. However, the blended result is only displayed above a certain threshold temperature. For any vertices under the threshold, the visible information is displayed. This is useful to quickly highlight high temperature regions of the scene. Equations 2, 3 and 4 detail the implementation, and results are shown in Figure 5b. The monochrome thermal mapping does not correspond to any thermal palette, so the palette bar to the left of the adjacent figure does not apply.

$$I(x, y, z) = \frac{T_{raw}(x, y, z) - T_{min}}{T_{max} - T_{min}} \quad (2)$$

where  $T_{min}$  and  $T_{max}$  are the minimum and maximum temperatures in the scene.

$$B(x, y, z) = HSV(H, I(x, y, z), I(x, y, z)) \quad (3)$$

$$F(x, y, z) = \begin{cases} B(x, y, z)\alpha + V(x, y, z)(1 - \alpha) & \text{if } T_{raw}(x, y, z) > T_{threshold} \\ V(x, y, z) & \text{otherwise} \end{cases} \quad (4)$$

The user can set the  $\alpha$  value as in Method 3.1, but can also optimize the threshold temperature  $T_{threshold}$ . By offering a choice of the colour hue used, we allow the user to compensate for varying environments and backgrounds with different dominant colours in the visible domain. The user is also able to invert the method and highlight low temperatures instead of high temperatures.

#### 3.3 Colour Mapped Threshold

A different implementation of Method 3.2 replaces the use of a monochrome threshold with a colour mapped threshold. The visible and thermal colour mapped meshes are blended, and the temperature threshold is then applied as before. This method highlights irregularities well with temperature detail within the highlighted regions. Details of the implementation are in Equation 5, and an example result is shown in Figure 5a.

$$F(x, y, z) = \begin{cases} T_{map}(x, y, z)\alpha + V(x, y, z)(1 - \alpha) & \text{if } T_{raw}(x, y, z) > T_{threshold} \\ V(x, y, z) & \text{otherwise} \end{cases} \quad (5)$$





Figure 3: (a) Output 3D temperature mesh model using the Rainbow colour palette. (b) Output 3D colour mesh model.




Figure 4: (a) Output 3D Thermal-Visible Blending fusion. (b) Output 3D Intensity-Hue Mapping fusion.




Figure 5: (a) Output 3D Colour Mapped Threshold fusion. The false colour mapping scheme used for this figure was designed to maximize both contrast and intuitiveness [Vidas and Moghadam, 2013a]. (b) Output 3D Monochrome Threshold fusion.




Figure 6: Comparison of Thermal Highlighting and Histogram-based Thresholding fusion methods.

### 3.4 Thermal Highlighting

The Thermal Highlighting fusion method is a novel contribution. Often, highlighting just hot or cold regions is inadequate. In many cases it is beneficial to view all the temperature extremities by highlighting the hot *and* the cold regions simultaneously. This fusion method is most useful when the detection of extreme hot and cold temperatures is more important than the colour information in the visible mesh. The implementation is explained in Equations 6 and 7, and results are shown in Figures 6a and 7b.

$$W(x, y, z) = \begin{cases} k \cdot \left| \frac{T_{raw}(x, y, z) - T_{high}}{T_{max} - T_{high}} \right|^2 & \text{if } T_{raw}(x, y, z) > T_{high} \\ k \cdot \left| \frac{T_{raw}(x, y, z) - T_{low}}{T_{max} - T_{low}} \right|^2 & \text{if } T_{raw}(x, y, z) < T_{low} \end{cases} \quad (6)$$

where  $T_{low}$  and  $T_{high}$  are the low and high temperature thresholds.

$$F(x, y, z) = \begin{cases} T_{map}(x, y, z) \cdot W(x, y, z) + V(x, y, z)(1 - W(x, y, z)) & \text{if } (T_{raw}(x, y, z) > T_{high}): \\ T_{map}(x, y, z) \cdot W(x, y, z) + V(x, y, z)(1 - W(x, y, z)) & \text{if } (T_{raw}(x, y, z) < T_{low}): \\ V(x, y, z) & \text{otherwise} \end{cases} \quad (7)$$

Note that instead of using a flat  $\alpha$  blend, a weighting factor  $W(x, y, z)$  is calculated based on the temperature of the vertex. This provides a smoothing effect in the transition band between the regions on either side of the threshold temperature. To simulate the scalable  $\alpha$  value in the previous fusion methods, a user-selected  $k$ -value acts as a scaling factor.

Because the thresholds in the highlighting method are also customizable, it can substitute for the threshold methods as long as the user is aware of the smoothing effect in the transition band.

### 3.5 Intensity-Hue Mapping

Intensity-Hue Mapping is another novel method proposed in this paper to unambiguously convey temperature information for all mesh faces, whilst still incorporating visual spatial information. The Intensity-Hue Mapping method makes use of the Hue-Saturation-Lightness (HSL) or Hue-Saturation-Value (HSV) colour spaces to best represent all the data (red, green, blue, temperature) in the RGB output mesh. The application of this fusion method is to allow the user to quickly interpret all of the temperatures present in the scene while allowing geometric information, texture and other critical information such as writing and labels to be easily identified and read [Vidas *et al.*, 2013]. Essentially, it increases the spatial resolution while retaining the spectral information [Al-Wassai *et al.*, 2011].

First, the colours in the colour mesh are rescaled according to Equation 8.

$$V_{scaled}(x, y, z) = k \cdot V(x, y, z) + 255 \left( \frac{1 - k}{2} \right) \quad (8)$$

Next, the scaled colour mesh  $V_{scaled}(x, y, z)$  is converted to greyscale  $V_{grey}(x, y, z)$ . These greyscale values are used to find the changes in luminance applied to the thermal colour mapped mesh from the visible information, as shown in Equation 9.

$$L(x, y, z) = 2 \left( \frac{V_{grey}(x, y, z)}{255} - 0.5 \right) \quad (9)$$

The application of the luminance change is shown in Equation 10, and an example result is shown in Figures 4b and 7a.

$$F(x, y, z) = \begin{cases} T_{map}(x, y, z) + L(x, y, z) \cdot (255 - T_{map}(x, y, z)) & \text{if } (L(x, y, z) > 0): \\ T_{map}(x, y, z) + L(x, y, z) \cdot T_{map}(x, y, z) & \text{if } (L(x, y, z) \leq 0): \end{cases} \quad (10)$$




Figure 7: Detailed views of thermal irregularities showing readability of writing and labels.



Figure 8: Inverse Intensity-Hue Mapping fusion.

In the same way as in the Thermal Highlighting method, an RGB scale factor  $k$  has been included to give the user control over the visualization.

### 3.6 Inverse Intensity-Hue Mapping

The application of this Inverse Intensity-Hue Mapping fusion is to highlight areas of high temperatures by scaling the visible intensities in those regions. While the Intensity-Hue Mapping method seeks to retain as much information as possible by scaling the colour mapped data with the visible (greyscaled) information, this method does the inverse.

The change in intensities does affect the perceived colour, so this method is not meant to preserve all the visible information. Instead, it allows the user to quickly identify high temperature regions and read/examine relevant visible details. This is unhindered by the rest of the model which has lower temperatures, as these tend to fade to grey.

$H(x, y, z)$  and  $S(x, y, z)$  are transformed from  $V(x, y, z)$  and retained for the fused model. The thermal intensity  $I(x, y, z)$  is substituted for the third channel, as shown in Equations 11 and 12. An example image generated using this method is

shown in Figure 8.

$$I(x, y, z) = 255 \left( \frac{T_{raw}(x, y, z) - T_{min}}{T_{max} - T_{min}} \right)^k \quad (11)$$

$$F(x, y, z) = HSV(H, S, I(x, y, z)) \quad (12)$$

The exponential weighting value  $k$  is chosen by the user to scale the proportion of the model that is in focus. During the design of the visualizer this exponential weighting was found to be more intuitive to use as opposed to a linear scale factor.

### 3.7 Histogram Bin Thresholding

The final proposed fusion method creates a histogram that can be displayed for the user to examine. The thresholding is applied to the histogram: the  $n$  bins with the lowest count are displayed as colour mapped thermal data and all the others are displayed as regular visible data, which has the effect of highlighting regions with temperature discrepancies. In general, this has a similar effect to Thermal Highlighting: the most uncommon temperature regions tend to be the hottest and coldest areas. However, in select cases, such as medical applications, the histogram can show that the less common temperatures are those that occur in the middle of the temperature range. In these cases, this fusion method offers the only way to quickly retrieve this information and highlight the relevant areas. An example image is shown in Figure 6b.

The user can set the resolution of the histogram, that is, the number of bins calculated, as well as specifying the threshold number  $n$ . Note that the method can be inverted to highlight the regions with the most common temperatures as well, if so desired.

A slight variation of this method allows the user to calculate thresholds based on dataset statistics. The thresholds are determined to be  $k$  standard deviations away from the mean  $\mu$ , as shown in Equations 13 and 14. The final effect is very similar to thermal highlighting as the mean tends to lie in the middle of the temperature range (assuming a Gaussian distribution of temperatures). A resulting image and corresponding






Figure 9: Histogram showing threshold limits for the distribution threshold method.

histogram are shown in Figures 6c and 9 respectively.

$$t_{min} = \mu - k\sigma \quad (13)$$

$$t_{max} = \mu + k\sigma \quad (14)$$

The user can select the value of  $k$  interactively.

## 4 Visualization Toolkit

We have developed a visualization toolkit that can be used to perform fusion using the methods listed in Section 3 in real-time. This allows users to customize the visualization to their personal preferences and see the effects of their changes immediately. Even for larger meshes the rendering of a new fusion scheme takes only seconds.

“Spectra” is the first 3D multispectral visualization toolkit that allows real-time visualization of a variety of multispectral fusion methods. It is also possible for a user to extend the toolkit by adding new fusion methods. All of these are presented to the user through the “Spectra” Graphical User Interface which is intuitive and easy to use.

Some other key features of “Spectra” include:

- Standard thermography options such as automatic temperature ranging, setting of min/max temperatures, displaying a temperature scale bar and changing colour palettes
- Allowing a user to view thermal-only, visible-only or fused outputs

A screenshot of “Spectra” is shown in Figure 10. This displays the general preferences dialog (Label 1) as well as the background colour dialog for adjusting the colour of the background (Label 2). The bar at the bottom of the screenshot (Label 3) displays the statistics relevant to each dataset: number of vertices, minimum and maximum temperature, and camera information. The colour bar on the side of the screen displays the temperature scale and is updated dynamically if the thermal colour palette or temperature range is changed (Label 4).




Figure 11: Flowchart of software dependencies in the “Spectra” multispectral visualization toolkit.

The toolkit has been constructed in C++ for a Linux (Ubuntu 12.04) system using open source libraries. Figure 11 shows the libraries used by “Spectra”:

- Point Cloud Library (PCL 1.6.0) - 3D point cloud and polygon mesh handling
- Visualization Toolkit (VTK 5.10.1) - used by PCL for 3D rendering
- Qt 4.8.4 - Graphical User Interface development framework
- QVTK Widget - Bridge between Qt and VTK, used for rendering and interactor

The application can read standard polygon format files (.ply). 3D output models from “Spectra” can be saved in standard 3D polygon file formats and opened with 3rd-party 3D viewers. Thermal infrared data is represented using standard thermography colour palettes. Users can customize the visualization to their personal preferences and see the effects of their changes immediately.

## 5 Results

All of the images throughout this paper were generated and saved using the proposed “Spectra” visualization toolkit. (Figure 12)<sup>2</sup>. The first 3D mesh shows a workstation with two computer monitors at higher temperatures, and a drink bottle on the left side at a lower temperature. Once again, it must be emphasized that the aim of these fusion methods is to highlight the regions of interest to the operator, which will lower their overall cognitive workload [Rasmussen *et al.*, 2009].

These results qualitatively support the statements made about each fusion method’s strengths and weaknesses. In Figure 4a, the Thermal-Visible Blend method shows that the temperatures and visible spectrum information can be observed on the computer monitors. However, the scene has reduced

<sup>2</sup><https://www.youtube.com/watch?v=QB2xpjDcyKw>




Figure 10: Screenshot of the “Spectra” 3D multispectral visualization toolkit. The preferences menu allows users to optimize the appearance of the model by varying method-specific parameters. 3D thermal model is from [Vidas *et al.*, 2013].



Figure 12: Screenshot from the “Spectra” visualization toolkit, showing the Intensity-Hue Mapping fusion on an industrial HVAC system. More detail is in the video attached to the paper.



contrast compared to Intensity-Hue Mapping, which achieves the same outcomes without the loss of contrast, as shown in Figure 4b.

Both the Monochrome and Colour Mapped Threshold fusion methods are similar, with the key difference being the representation of the thermal data (see Figures 5b and 5a). Upon inspection, the visible information on the monitor is clearer with Monochrome Thresholding, even though the  $\alpha$  value is the same. However, it is also impossible to discern the temperatures using this method due to the lack of any scale. Therefore, if specific temperature information is critical, the Colour Mapped Threshold method must be used, but the Monochrome Threshold method is still useful to highlight irregularities.

The fusion methods in Figure 6 improve on the threshold methods by also highlighting the cold water bottle. The iron colour palette is demonstrated in this figure. All three output meshes in this figure achieve the same outcome: they highlight the high temperature monitor and the low temperature bottle. If the user is interested in identifying regions outside a specific temperature range, for instance in the scenario where a mechanic is targeting overheating components, the Thermal Highlighting method can be used, as in Figure 6a. The Histogram Fusion methods can also be used by the operator for the same effect and offer different means of specifying the threshold temperatures, with default settings that do not require the user to know the exact temperatures of the thermal extremities (Figures 6b and 6c).

Figure 8 shows results using the Inverse Intensity-Hue Mapping scheme. This quickly highlights high temperature regions. The computer monitor is more vivid and the user's attention is drawn to that region, but the colder parts of the scene are faded.

Another dataset is shown in Figure 7, with a more detailed view of the effects of three fusion methods. As before, the high temperature laptop is highlighted, as is the cold bottle. The Intensity-Hue Mapping method shows the writing on both the laptop and the monitor (see Figure 7a). Both the Thermal Highlighting and the Histogram Distribution Thresholding methods offer the same success in highlighting irregularities but in this case the Thermal Highlighting method blurs the words on the laptop monitor slightly, as observed in Figure 7c. This is due to the smoothing effect in the transition band from the weighting factor.

## 6 Conclusion

This paper has proposed several new 3D multispectral fusion and visualization methods. Dense, high-resolution 3D multispectral mesh models can be visualized in the thermal and visible modalities either independently or simultaneously using the proposed interactive visualization toolkit "Spectra". We used 3D mesh fusion techniques to represent both modalities (thermal and visible) in a single 3D model. This can reduce the cognitive load on the operator by allowing them

to focus on a single 3D multi-modal model. Users can easily and efficiently interact with the 3D model and select and change different parameters of fusion methods, with results presented to the operator on-the-fly. By utilizing 3D fusion, visible and thermal data collected from an autonomous mapping robot can be visualized without requiring two separate screens, reducing the remote operator's cognitive workload. Building inspectors investigating structural defects or heat loss can quickly identify areas of interest, with the thermal information augmented by the visible data. Remote operators of search and rescue robots can more quickly identify victims based on their thermal signature.

Future areas of study may result in extensions of the software application, including the implementation of other fusion methods, such as segmentation based on regions of common temperature. A method of autonomously selecting a fusion method and the corresponding settings to maximize the information displayed to the operator will also be developed. In the immediate future, a user study will also be conducted to identify which (if any) colour palettes, fusion methods and settings are more popular and intuitive. This will yield quantitative results to support the relevance of this work. Feedback from this user study will also be used to further improve the visualization toolkit.

## References

- [Al-Wassai *et al.*, 2011] Firouz Abdullah Al-Wassai, N. V. Kalyankar, and Ali A. Al-Zuky. The IHS transformations based image fusion. *Journal of Global Research in Computer Science*, 2(5):70–77, 2011.
- [Bican *et al.*, 2002] Jakub Bican, Daniel Janeba, Katerina Tbarsk, and Jaroslav Vesel. Image overlay using alpha-blending. *Nuclear Medicine Review*, 5(1):53, 2002.
- [Borrmann *et al.*, 2012] Dorit Borrmann, Andreas Nchter, Marija DJakulovi, Ivan Maurovi, Ivan Petrovi, Dinko Osmankovi, and Jasmin Velagi. The project ThermalMapper-Thermal 3D mapping of indoor environments for saving energy. In *Proceedings of the 10th International IFAC Symposium on Robot Control (SYROCO)*, volume 10, 2012.
- [Jacoff *et al.*, 2000] Adam Jacoff, Elena Messina, and John Evans. A standard test course for urban search and rescue robots. In *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, pages 499–503, 2000.
- [Keyes *et al.*, 2010] Brenden Keyes, Mark Micire, Jill L. Drury, and Holly A. Yanco. *Improving Human-Robot Interaction through Interface Evolution*. InTech, 2010.
- [Malpica *et al.*, 2013] Jos A. Malpica, Mara C. Alonso, Francisco Pap, Antonio Arozarena, and Alex Martinez De Agirre. Change detection of buildings from satellite imagery and lidar data. *International Journal of Remote Sensing*, 34(5):1652–1675, March 2013.

- [Moghadam and Vidas, 2014] Peyman Moghadam and Stephen Vidas. Heatwave: the next generation of thermography devices. In *Thermosense: Thermal Infrared Applications XXXVI*, volume 9105. Proc. SPIE, 2014.
- [Nagatani *et al.*, 2014] Keiji Nagatani, Kazuki Otake, and Kazuya Yoshida. Three-dimensional thermography mapping for mobile rescue robots. In *Field and Service Robotics*, pages 49–63. Springer, 2014.
- [Rasmussen *et al.*, 2009] Nathan D. Rasmussen, Bryan S. Morse, Michael A. Goodrich, and Dennis Eggett. Fused visible and infrared video for use in wilderness search and rescue. In *Workshop on Applications of Computer Vision (WACV)*, page 18, 2009.
- [Ross *et al.*, 2000] W. D. Ross, Allen M. Waxman, W. W. Streilein, M. Aguiar, J. Verly, F. Liu, M. I. Braun, P. Harmon, and S. Rak. Multi-sensor 3D image fusion and interactive search. In *Proceedings of the Third International Conference on Information Fusion*, volume 1, 2000.
- [Schnelle and Chan, 2011] Stephen R. Schnelle and Alex Lipchen Chan. Enhanced target tracking through infrared-visible image fusion. In *Proceedings of the 14th International Conference on Information Fusion*, 2011.
- [Se and Jasiobedzki, 2006] Stephen Se and Piotr Jasiobedzki. Photo-realistic 3D model reconstruction. In *The 2006 IEEE International Conference on Robotics and Automation*, pages 3076–3082, 2006.
- [Vidas and Moghadam, 2013a] Stephen Vidas and Peyman Moghadam. Ad Hoc Radiometric Calibration of a Thermal-Infrared Camera. *2013 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8, November 2013.
- [Vidas and Moghadam, 2013b] Stephen Vidas and Peyman Moghadam. Heatwave: a handheld 3d thermography system for energy auditing. *Energy and Buildings*, 66:445–460, 2013.
- [Vidas *et al.*, 2013] Stephen Vidas, Peyman Moghadam, and Michael Bosse. 3D thermal mapping of building interiors using an rgb-d and thermal camera. In *The 2013 IEEE International Conference on Robotics and Automation*, 2013.
- [Waxman *et al.*, 2001] Allen M. Waxman, J. G. Verly, David A. Fay, Fang Liu, Michael I. Braun, Benjamin Pugliese, W. D. Ross, and W. W. Streilein. A prototype system for 3D color fusion and mining of multisensor/spectral imagery. In *4th International Conference on Information Fusion, Montreal*, 2001.
- [Xie *et al.*, 2004] Huchen Xie, Guang Li, Holly Ning, Cynthia Menard, C. Norman Coleman, and Robert W. Miller. 3D voxel fusion of multi-modality medical images in a clinical treatment planning system. In *Computer-Based Medical Systems, 2004. CBMS 2004. Proceedings. 17th IEEE Symposium on*, page 4853, 2004.