https://doi.org/10.2352/ISSN.2470-1173.2021.18.3DIA-061

Two-Channel 3D Range Geometry Compression with Virtual Plane Encoding

Matthew G. Finley and Tyler Bell*

Department of Electrical and Computer Engineering, University of Iowa; Iowa City, Iowa 52242, USA
* ***tyler-bell@uiowa.edu***

Abstract

Modern computing and imaging technologies have allowed for many recent advances to be made in the field of 3D range imaging: range data can now be acquired at speeds much faster than real-time, with sub-millimeter precision. However, these benefits come at the cost of an increased quantity of data being generated by 3D range imaging systems, potentially limiting the number of applications that can take advantage of this technology. One common approach to the compression of 3D range data is to encode it within the three color channels of a traditional 24-bit RGB image. This paper presents a novel method for the modification and compression of 3D range data such that the original depth information can be stored within, and recovered from, only two channels of a traditional 2D RGB image. Storage within a traditional image format allows for further compression to be realized via lossless or lossy image compression techniques. For example, when JPEG 80 was used to store the encoded output image, this method was able to achieve an 18.2% reduction in file size when compared to a similar three-channel, image-base compression method, with only a corresponding 0.17% reduction in global reconstruction accuracy.

Introduction

Modern three-dimensional (3D) scanning devices and techniques allow us to capture high-precision 3D range data at speeds much faster than real-time. These advances in 3D scanning make this technology highly desirable in fields such as telemedicine, autonomous metrology, communications, and security. As resolution and capture speed increase, however, the quantity of data that must be stored or transmitted also increases.

This 3D data is traditionally stored in a mesh format such as OBJ or STL. These formats conveniently store 3D coordinate locations as well as information about how the coordinates must connect. Additionally, these formats often store auxiliary information about surface normals or texture coordinates, which is potentially convenient but comes at the cost of increased file sizes. As such, alternative methods of 3D range geometry compression have been explored.

One such approach to 3D range geometry compression is to store the 3D data within the three color channels of a traditional lossless (e.g., PNG) 2D RGB image, enabling the use of modern, well-defined image processing and compression techniques [1, 2, 3, 4]. If greater compression ratios are required, lossy compression standards such as JPEG may be used when storing the encoded output images, although care needs to be taken to ensure the encoding signals being utilized are robust to the effects of lossy compression.

In order to further reduce file sizes, several methods have been explored that aim to reduce the number of encoding signals, and therefore the number of color channels, required to represent a 3D range scan within a 2D color image [5, 6]. Although these methods were successful in using only two color channels to represent the 3D range data, and subsequently reducing file sizes when compared to similar three-channel compression methods, the nature of their encoding signals make them suitable only for lossless image storage.

To enable the use of *lossy* image compression standards for two-channel 3D range data compression, a method was proposed that utilized encodings robust to the effects of lossy compression [7]. This allows for high compression ratios to be achieved when compared to similar lossy three-channel compression approaches. However, the nature of this method's encoding signals forces ambiguity into the phase unwrapping process necessary to the decoding stage of image-based 3D range geometry compression, meaning that errors are present at regular intervals throughout the recovered geometry. Additional steps are needed to remove these errors. One approach was to first segment regions of assumed error, and then to fill these segmented regions with heavily filtered data. Regardless of the approach taken to correct the error, this method has a relatively high minimum RMS reconstruction error.

This paper proposes a novel method for the compression of 3D range geometry within only two of the three color channels available in a traditional 2D RGB image. The proposed encoding signals are robust to the effects of lossy compression, enabling high compression ratios to be achieved through the use of image standards such as JPEG. Further, this method's phase unwrapping process is robust, meaning that no error correction framework is necessary and a relatively low minimum reconstruction error is achievable. For example, the proposed method is able to achieve an 18.2% reduction in file size (compared to modern, robust three-channel image-based compression methods) for only a 0.17% reduction in global reconstruction accuracy when using the JPEG 80 compression standard to store the encoded output image.

Principle

Multiwavelength Depth Encoding

One image-based approach to the compression of 3D range geometries is the Multiwavelength Depth Encoding (MWD) method [4]. This method utilizes principles of phase-shifting in order to encode high-quality depth information within the three color channels of a traditional 2D RGB image. Mathematically,

the signals used to encode the depth information, Z, can be described as

$$I_1(i,j) = \frac{1}{2} + \frac{1}{2}\sin\left(2\pi \times \frac{Z(i,j)}{P}\right), \quad (1)$$

$$I_2(i,j) = \frac{1}{2} + \frac{1}{2}\cos\left(2\pi \times \frac{Z(i,j)}{P}\right), \quad (2)$$

$$I_3(i,j) = \frac{Z(i,j)}{\text{Range}(Z)}, \quad (3)$$

where P is a user-defined *fringe width* that determines the frequency of the sinusoidal encodings seen in Eqs. (1)-(2). Equation (3) is a low-frequency encoding representing the original depth information, and is used by the MWD decoding process in order to recover the encoded geometry.

Multiwavelength Depth Decoding

In order to recover the original depth information, the MWD decoding process first solves for the encoded phase information within Eqs. (1) and (2) using the inverse tangent function as

$$\phi_{HF}(i,j) = \tan^{-1}\left(\frac{I_1(i,j) - 0.5}{I_2(i,j) - 0.5}\right). \quad (4)$$

However, since the inverse tangent function is only defined in the range $[-\pi, \pi)$, $\phi_{HF}(i,j)$ will have sharp discontinuities at intervals of 2π. This is because the phase information encoded into the sinusoids has range greater than 2π, proportional to the number of encoding periods, $nstr$, defined as $nstr = \text{Range}(Z)/P$. As such, ϕ_{HF} is known as *wrapped phase* and will need to be unwrapped prior to the recovery of the original geometry.

In order to unwrap ϕ_{HF}, the MWD method first scales I_3 into the range of $-\pi$ to π via

$$\phi_{LF}(i,j) = I_3(i,j) \times 2\pi - \pi. \quad (5)$$

Next, a *stair image*, $K(i,j)$, can be calculated in order to determine how many 2π must be added or subtracted to the wrapped phase, ϕ_{HF}, in order to recover *absolute phase*. Mathematically, this stair image can be calculated as

$$K(i,j) = \text{Round}\left(\frac{\phi_{LF}(i,j) \times \text{Range}(Z)/P - \phi_{HF}(i,j)}{2\pi}\right). \quad (6)$$

It is important to note that the generation of this stair image, K, relies heavily on the representation of the original geometry stored in I_3. Significant degradation of this signal will result in an incorrect stair image, and a subsequently incorrectly recovered geometry. The absolute phase, $\Phi(i,j)$, can be recovered via

$$\Phi(i,j) = \phi_{HF}(i,j) + 2\pi \times K(i,j). \quad (7)$$

Finally, the absolute phase can simply be scaled into the original depth dimensions of the encoded geometry as

$$Z(i,j) = \frac{\Phi(i,j) \times P}{2\pi}. \quad (8)$$

Virtual Plane Encoding

This paper presents a novel method for the encoding of 3D range geometries within only two of the three available color channels present in a traditional 2D RGB image. The proposed method does not require the low-frequency signal commonly used to unwrap the phase information stored in the high-frequency sinusoidal encodings typical of phase-shifting compression methods. The method's proposed encodings can be described mathematically as

$$I_1(i,j) = \frac{1}{2} + \frac{1}{2}\sin\left(2\pi \times \frac{Z_{Modified}(i,j)}{P}\right), \quad (9)$$

$$I_2(i,j) = \frac{1}{2} + \frac{1}{2}\cos\left(2\pi \times \frac{Z_{Modified}(i,j)}{P}\right), \quad (10)$$

where P, as in the MWD method discussed previously, is a user-defined fringe width that determines the frequency of the sinusoidal encodings, and $Z_{Modified}(i,j)$ is a modified version of the original geometry.

$Z_{Modified}$ is generated by summing the original geometry with a virtual plane, and can be described mathematically as

$$Z_{Modified}(i,j) = Z(i,j) + Z_{Plane}(i,j). \quad (11)$$

If the magnitude of Z_{Plane} is sufficiently large when compared to the original geometry's depth range, then Z_{Plane} can be used to accurately describe the encoded geometry, $Z_{Modified}$. Further, if Z_{Plane} is a good approximation of $Z_{Modified}$, then the virtual plane can be used in the phase unwrapping process, eventually allowing for the recovery of the original depth information. This means that the need for an additional, low-frequency encoded signal such as the one seen in Eq. (3) is eliminated, and only the two high-frequency encoding sinusoids (Eqs. (9) and (10)) must be stored in the encoded output image. It is important to note that in order to regenerate the virtual plane, the four floating-point parameters required to represent a plane must be transmitted or stored alongside the encoded output image.

Figure 1 illustrates the proposed method applied to an ideal hemisphere with a depth range of 256 mm. Figure 1a is a 3D rendering of the original geometry to be compressed. Figure 1b is a 512×512 2D depth map, Z, representing the original 3D geometry. Figure 1c is the virtual plane, Z_{Plane}. In this case Z_{Plane} has a depth range of 819.2 mm, which is large enough when compared to the original, distinctly non-planar geometry to enable correct phase unwrapping. Figure 1d is a plot showing the cross-section of the original data (in red), the virtual plane (in dashed black) and the new, modified geometry to be encoded by Eqs. (9) and (10) (in blue). This plot further illustrates that the virtual plane (Z_{Plane}) is very large in magnitude when compared to the original geometry (Z), and that the modified geometry ($Z_{Modified}$) is closely approximated by only the virtual plane. Figure 1e is the modified geometry, $Z_{Modified}$, with a depth range of 887.3 mm. Figure 1f and 1g are the sinusoidal equations seen in Eq. (9) and Eq. (10), respectively. Figure 1h is the encoded output image stored in the PNG format, with Figs. 1f and 1g stored in the red and green color channels, respectively. The blue color channel is intentionally left empty in order to best illustrate the proposed method. This encoded output image has a file size of 93.1 KB.

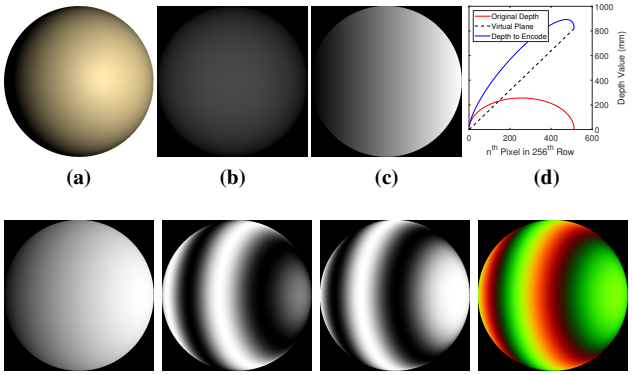

Figure 1. Encoding an ideal hemisphere with a depth range of 256 mm with the proposed method. (a) 3D rendering of the geometry to be encoded; (b) a 2D depth map representing the original geometry; (c) the virtual plane with a depth range of 819.2 mm; (d) a cross-section plot showing the original geometry in red, the virtual plane in dashed black, and the sum of the two in blue; (e) the modified geometry, generated by summing the original geometry with the virtual plane as shown in (d); (f) sinusoidal encoding of (e) corresponding to Eq. (9); (g) sinusoidal encoding of (e) corresponding to Eq. (10); (h) the encoded output image, in this case in the PNG format, generated by storing (f) and (g) in the red and green color channels, respectively.

Virtual Plane Decoding

The proposed method begins to decode the two-channel 2D image by first solving for the phase information contained within the two sinusoidal encodings, per the previously described Eq. (4). Next, the virtual plane must be regenerated using the four plane parameters that are transmitted or stored alongside the encoded output image.

This virtual plane is then scaled into the range of the wrapped phase by

$$\phi_{Plane}(i,j) = \frac{Z_{Plane}(i,j)}{\text{Range}(Z_{Plane})} \times 2\pi - \pi. \tag{12}$$

Then, ϕ_{Plane} and ϕ_{HF} can be used to calculate the stair image K, which will be used to unwrap ϕ_{HF}. Mathematically, the stair image can be described for each pixel (i,j) as

$$K_{Plane} = \text{Round}\left(\frac{\phi_{Plane} \times \text{Range}(Z_{Modified})/P - \phi_{HF}}{2\pi}\right). \tag{13}$$

Next, the stair image is used to unwrap the highly discontinuous phase information in ϕ_{HF} via

$$\Phi_{Modified}(i,j) = \phi_{HF}(i,j) + 2\pi \times K_{Plane}(i,j). \tag{14}$$

This unwrapped phase information, $\Phi_{Modified}$, is called absolute phase and can be scaled into the dimensions of the modified geometry via

$$Z_{Modified}(i,j) = \frac{\Phi_{Modified}(i,j) \times P}{2\pi}. \tag{15}$$

Finally, the original depth information can be recovered by subtracting the virtual plane from this recovered modified geometry

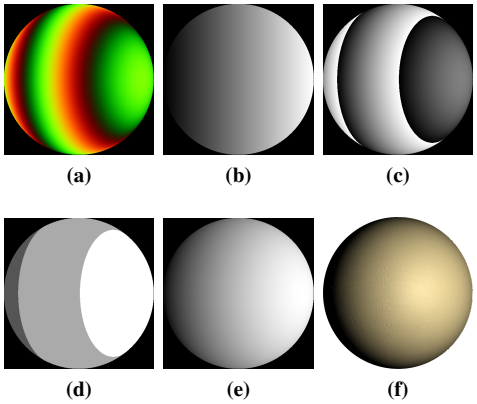

Figure 2. Decoding an ideal hemisphere with the proposed method. (a) The encoded PNG image output by the encoding process described previously; (b) the virtual plane, regenerated from the four floating-point parameters stored alongside the encoded image; (c) the high-frequency wrapped phase calculated from the two sinusoidal signals stored in (a); (d) the stair image generated using (b) and (c); (e) the recovered modified geometry; (f) a 3D rendering of the recovered original geometry, generated by removing (b) from (e).

using the following equation:

$$Z(i,j) = Z_{Modified}(i,j) - Z_{Plane}(i,j). \tag{16}$$

Figure 2 illustrates the proposed method's decoding process on the same ideal hemisphere presented previously. Figure 2a is the encoded output image generated via the proposed encoding process. Figure 2b is the virtual plane, Z_{Plane}, regenerated using the four plane parameters that are stored or transmitted alongside the encoded output image. Figure 2c is the high frequency phase information, ϕ_{HF}, calculated from the two sinusoidal equations stored in the red and green color channels of the encoded output image via Eq. (4). Figure 2d is the stair image, K_{Plane}, generated through the use of ϕ_{HF} and Z_{Plane} via Eq. (13). Figure 2e is $Z_{Modified}$, recovered by using the stair image to unwrap the high frequency, discontinuous phase information and then scaling into the correct dimensions. Figure 2f is a 3D rendering of the original geometry, Z, recovered by subtracting the virtual plane (Z_{Plane}) from the recovered $Z_{Modified}$ via Eq. (16). The recovered geometry has an RMS error of 0.162 mm, which corresponds to a 99.94% global RMS reconstruction accuracy when compared to the total depth range of 256 mm.

Generation of Virtual Plane

As discussed in the previous sections, the magnitude of the virtual plane that is summed with the original geometry must be large when compared to the original geometry's depth range, such that $Z_{Modified}$ can be correctly approximated using only Z_{Plane}. In order to determine whether this relationship is satisfied for some Z_{Plane}, it is necessary to generate K_{Ideal}, the stair image generated using $Z_{Modified}$. This can be done for each pixel (i,j) via

$$K_{Ideal} = \text{Round}\left(\frac{\phi_{Ideal} \times \text{Range}(Z_{Modified})/P - \phi_{HF}}{2\pi}\right), \tag{17}$$

where ϕ_{Ideal} is the low-frequency phase data calculated using the modified geometry being encoded. Mathematically, this term can

be defined as

$$\phi_{Ideal}(i,j) = \frac{Z_{Modified}(i,j)}{\text{Range}(Z_{Modified})} \times 2\pi - \pi. \quad (18)$$

In order for the phase unwrapping process to function correctly, the stair map generated referencing the virtual plane (via Eq. (13)) must be equal to the stair map generated referencing the depth map modified by the virtual plane (via Eq. (17)). This equality, defined mathematically as

$$K_{Plane}(i,j) = K_{Ideal}(i,j), \quad (19)$$

must hold for every pixel, (i,j). In practice this simply means that the virtual plane is large enough that it can serve as a good approximation of the modified geometry. If this equality is *not* true, i.e., $K_{Plane}(i,j) \neq K_{Ideal}(i,j)$, then the magnitude of Z_{Plane} must be increased.

In general, any virtual plane can be represented with only four parameters. Practically, for the experiments in this manuscript, $Z_{Plane}(i,j)$ was generated to be horizontally increasing by setting the value of each column of pixels in Z_{Plane} equal to the product of the column position, the image width, and some scaling factor. This scaling factor was modified as necessary in order to produce correctly unwrapped phase information via the equality given in Eq. (19). Mathematically, this process can be described using

$$Z_{Plane}(i,j) = j \times \text{ImageWidth} \times \text{SF}, \quad (20)$$

where j is the column index, ImageWidth is the depth map's width, in pixels, and SF is the user-defined scaling factor.

Experimental Results

The proposed method of two-channel 3D range geometry compression was evaluated with several experiments. In the first set of experiments a complex piece of data, in this case a 3D scan of a bust of George Washington [8] with a depth range of 337.6 mm, was compressed into a 713×632 2D image format, as shown in Fig 3.

Figure 3a shows Z, the floating-point depth map representing the original 3D range geometry. Figure 3b is Z_{Plane}, the virtual plane, with a depth range of 607.7 mm. Figure 3c is $Z_{Modified}$, the sum of Z and Z_{Plane}. In this case, $Z_{Modified}$ has a depth range of 792.2 mm. Figure 3d is the 2D image generated by sinusoidally encoding $Z_{Modified}$ according to Eqs. (9) and (10) and storing these encodings into the red and green color channels of a PNG image, respectively. These sinusoidal encodings were generated with two encoding periods, where the number of encoding periods (nstr) can be determined according to the formula $nstr = \text{Range}(Z_{Modified})/P$. For the purposes of this experiment, the blue color channel was left empty in order to best illustrate the potential for file size savings. Figure 3e is the wrapped phase recovered from the two sinusoidal encodings stored in Fig. 3d. Figure 3f is the stair image, K_{Plane}, which is calculated using the regenerated virtual plane and the wrapped phase as described in Eqs. (12) and (13). Figure 3g is the recovered depth information, generated by first calculating the absolute phase, rescaling it, and then removing the regenerated virtual plane as described by Eqs. (14)-(16), respectively. Figure 3h is a 3D rendering of the recovered depth information. This recovered geometry has an RMS

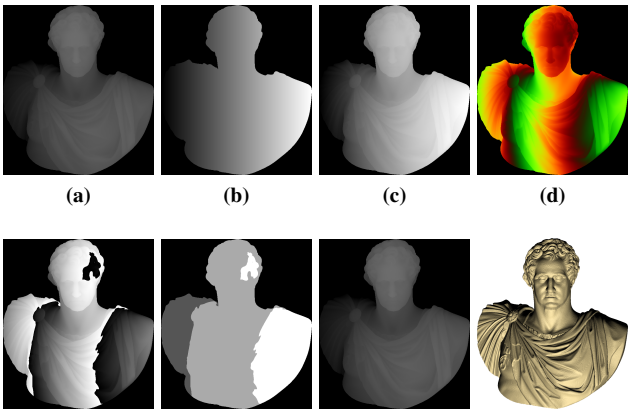

Figure 3. The proposed method applied to a 3D scan of a bust of George Washington. (a) 713×632 2D depth map representing the original 3D geometry; (b) a virtual plane with a depth range of 607.7 mm; (c) the modified geometry, generated by summing (a) with (b); (d) the encoded output image, stored in the PNG format; (e) the high-frequency wrapped phase generated by solving the sinusoidal signals stored in (d); (f) the stair image generated using (e) and the regenerated virtual plane; (g) the original geometry, recovered by first unwrapping (e) using (f), and then subtracting the virtual plane; (h) 3D rendering of the recovered geometry.

error of 0.144 mm, which gives a global RMS reconstruction accuracy of 99.96% when compared to the original depth range of 337.6 mm. The file size associated with the two-channel encoding of the geometry is 163.1 KB.

In order to illustrate the proposed method's potential for file size savings, these numerical results can be compared to the results obtained when encoding the same original geometry (seen in Fig. 3a) with the MWD method described in the Principle section of this paper. When the same number of encoding periods ($nstr = 2$) was utilized and the PNG image format was used to store the encoded output, the MWD method achieved a global RMS reconstruction accuracy of 99.98% with an associated file size of 247.7 KB. It can be seen that the proposed method of two-channel 3D range geometry compression was able to achieve a file size reduction of 34.2%, when compared to an equivalent three-channel compression algorithm, for only a 0.02% reduction in global RMS reconstruction accuracy.

The next experiment performs a similar comparison between the proposed method and the MWD method while varying the image storage format in order to demonstrate the proposed method's ability to utilize both lossless and lossy image compression standards in storing its encoded output. Table 1a reports the results of this comparison when two encoding periods were utilized and the PNG format was used to store the encoded outputs, the result of which has already been discussed. Table 1b shows the results of this comparison when two encoding periods were utilized and the JPG 100 format was used to store the encoded output images. As can be seen, the proposed method is able to achieve a 14.0% reduction in file size when compared to the MWD method, for only a 0.08% reduction in global RMS reconstruction accuracy. Table 1c gives the results when two encoding periods were utilized and the JPG 80 format was used to store the encoded output

Table 1. Performance of the proposed method compared to the MWD method when encoding a 337.6 mm scan of a bust of George Washington. (a) the comparison when the PNG image format was used to store the encoded output images; (b) the comparison when the JPG 100 image format was used to store the encoded output images; (c) the comparison when the JPG 80 image format was used to store the encoded output images.

(a)

PNG	Global Reconstruction Accuracy	File Size (KB)	Reduction in Accuracy	Reduction in File Size
MWD	99.98%	247.7	—	—
VPE	99.96%	163.1	0.02%	34.2%

(b)

JPG 100	Global Reconstruction Accuracy	File Size (KB)	Reduction in Accuracy	Reduction in File Size
MWD	99.72%	182.4	—	—
VPE	99.64%	156.9	0.08%	14.0%

(c)

JPG 80	Global Reconstruction Accuracy	File Size (KB)	Reduction in Accuracy	Reduction in File Size
MWD	99.60%	42.8	—	—
VPE	99.43%	35.0	0.17%	18.2%

images. It can be seen that a similar trend holds here, as well: the proposed method is able to achieve a 18.2% reduction in file size when compared to the MWD method, for only an associated 0.17% reduction in global RMS reconstruction accuracy. This experiment highlights the ability of the proposed two-channel range geometry compression method to achieve significant file size reduction when compared to equivalent three-channel compression algorithms, while maintaining very high (above 99%) global RMS reconstruction accuracy.

The final experiment uses the same 3D scan of a bust of George Washington, augmented with an additional, disjoint, 3D scan of a human face in order to illustrate the ability of the proposed method to determine absolute phase. This augmented 3D range scan is 713×1112 pixels with a depth range of 362.5 mm. Figure 4 demonstrates the proposed method of two-channel 3D range geometry compression applied to this disjoint 3D scan. Figure 4a is Z, the original floating-point depth information to be compressed. Figure 4b is the virtual plane, Z_{Plane}, in this case generated with a range of 719.8 mm. Figure 4c is $Z_{Modified}$, the sum of Z_{Plane} and Z. This modified geometry has a depth range of 945.8 mm. Figure 4d is the 2D image (in this case, stored in the PNG format) output by the proposed method after sinusoidally encoding the modified deph information, $Z_{Modified}$, using Eqs. (9) and (10). Figure 4e is the wrapped phase calculated from the two sinusoidal signals stored in the encoded output image via Eq. (4). Figure 4f is the stair image, K_{Plane}, generated via Eq. (13). Figure 4g is the depth map, Z, recovered by calculating the absolute phase, rescaling it, and removing the virtual plane as described by Eqs. (14)-(16), respectively. Finally, Figure 4h is a 3D rendering of the recovered geometry, Z. This experiment highlights the ability of the proposed method to accurately encode and decode true depth information related to spatially isolated surfaces within a scene.

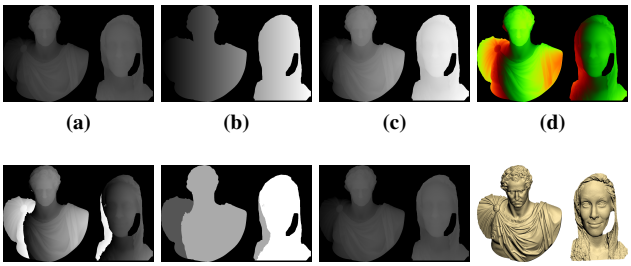

(a) **(b)** **(c)** **(d)** **(e)** **(f)** **(g)** **(h)**

Figure 4. The proposed method applied to a disjoint 3D scan of a bust of George Washington augmented with a 3D scan of a human face. This augmented scan is 713×1112 pixels with a depth range of 362.5 mm. (a) The original range geometry to be compressed; (b) the virtual plane, with range 719.8 mm; (c) the modified geometry generated by summing (a) with (b); (d) the encoded output image stored in the PNG format; (e) the recovered high-frequency wrapped phase; (f) the stair image, K_{Plane}; (g) the recovered original geometry, calculated by unwrapping (e) with (f) and subtracting the regenerated virtual plane; (h) 3D rendering of the recovered original geometry.

Summary

This paper has presented a novel method for the compression of 3D range geometry into only two of the three available color channels of a traditional 2D RGB image. The proposed method is able to achieve significant file size reduction when compared to similar image-based compression algorithms that require all three color channels to store the 3D data, and is able to achieve nearly the same global RMS reconstruction accuracy. Further, the proposed method's encoding signals are robust to the effects of storage in a lossy image format such as JPEG, allowing further file size reduction to be achieved, making this method suitable for applications where very high compression ratios are desired.

References

[1] Nikolaus Karpinsky and Song Zhang. Composite phase-shifting algorithm for three-dimensional shape compression. *Optical Engineering*, 49(6):1 – 6, 2010.

[2] Song Zhang. Three-dimensional range data compression using computer graphics rendering pipeline. *Appl. Opt.*, 51(18):4058–4064, 2012.

[3] Pan Ou and Song Zhang. Natural method for three-dimensional range data compression. *Appl. Opt.*, 52(9):1857–1863, 2013.

[4] Tyler Bell and Song Zhang. Multiwavelength depth encoding method for 3d range geometry compression. *Appl. Opt.*, 54(36):10684–10691, 2015.

[5] Zhiling Hou, Xianyu Su, and Qican Zhang. Virtual structured-light coding for three-dimensional shape data compression. *Optics and Lasers in Engineering*, 50(6):844 – 849, 2012.

[6] Yajun Wang, Lianxin Zhang, Sheng Yang, and Fang Ji. Two-channel high-accuracy holoimage technique for three-dimensional data compression. *Optics and Lasers in Engineering*, 85:48 – 52, 2016.

[7] Matthew G. Finley and Tyler Bell. Two-channel depth encoding for 3d range geometry compression. *Appl. Opt.*, 58(25):6882–6890, 2019.

[8] Smithsonian Institution: 3D Digitization Collection. George Washington. `https://3d.si.edu/object/3d/george-washington:d8c63fde-4ebc-11ea-b77f-2e728ce88125`, 2021.

Author Biography

Matthew G. Finley is a Ph.D. student in the Holo Reality Lab at the University of Iowa, studying applied signal and image processing. Finley earned his B.S. and M.S. from the University of Iowa in 2018 and 2019, respectively. He has given several invited talks on his research, his work has been featured on the cover of Applied Optics, and he was awarded a Best Student Paper Award at Electronic Imaging 2020 (3DMP).

Prof. Tyler Bell is an Assistant Professor of Electrical and Computer Engineering at the University of Iowa. He leads the Holo Reality Lab and is a faculty member of the Public Digital Arts (PDA) cluster. Tyler received his Ph.D. from Purdue University in 2018. His current research interests include high-quality 3D video communications; high-speed, high-resolution 3D imaging; virtual reality, augmented reality; human computer interaction; and multimedia on mobile devices.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

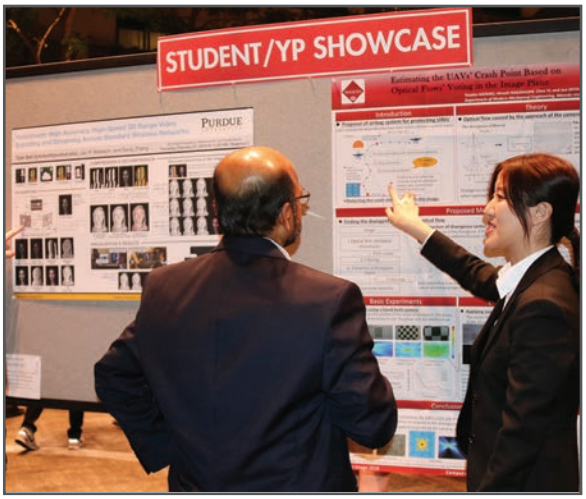

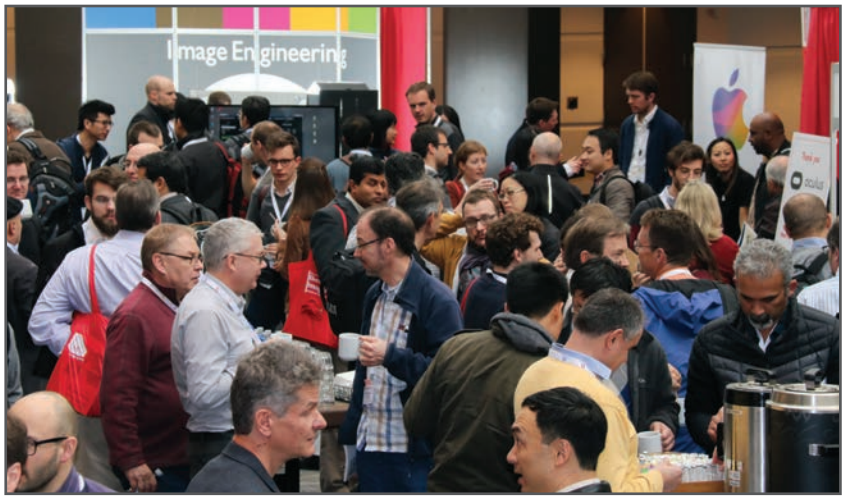

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org