

Automated pebble mosaic stylization of images

Lars Doyle¹ (✉), Forest Anderson¹, Ehren Choy¹, and David Mould¹

© The Author(s) 2019.

Abstract Digital mosaics have usually used regular tiles, simulating historical *tessellated* mosaics. In this paper, we present a method for synthesizing *pebble* mosaics, a historical mosaic style in which the tiles are rounded pebbles. We address both the tiling problem, of distributing pebbles over the image plane so as to approximate the input image content, and the problem of geometry, creating a smooth rounded shape for each pebble. We adopt simple linear iterative clustering (SLIC) to obtain elongated tiles conforming to image content, and smooth the resulting irregular shapes into shapes resembling pebble cross-sections. Then, we create an interior and exterior contour for each pebble and solve a Laplace equation over the region between them to obtain height-field geometry. The resulting pebble set approximates the input image while representing full geometry that can be rendered and textured for a highly detailed representation of a pebble mosaic.

Keywords non-photorealistic rendering; digital mosaics; image stylization; segmentation; image processing

1 Introduction

Mosaics are an art form that date back thousands of years. The earliest historical mosaics were *pebble mosaics* [1, 2], whose component pebbles were heterogeneous in size and shape. Pebble mosaics paved floors with pebbles, arranged so as to form an image or design. The craft of pebble mosaics continues into the 21st century [3] with new pebble mosaics being built by hobbyists and city planners.

Pebble mosaics, as well as the contemporaneous *chip mosaics* made of fragments of quarried stone [1],

use entirely irregular tiles. The archetypal mosaic is the *tessellated* mosaic, made of regular cubes of stone (tesserae). Such tessellated mosaics are the most familiar kind of mosaics and have been the most thoroughly studied in computer graphics. Tessellated mosaics have been dated to the third century BCE. However, pebble mosaics appeared in Greece hundreds of years earlier [1] and have not received much attention in computer graphics. In this paper, we propose a novel algorithm for constructing irregular pebble mosaics, using a variant of simple linear iterative clustering (SLIC) [4] to obtain an initial segmentation, smoothing the resulting boundaries, and using a Poisson solver to interpolate a smooth height field for each pebble which we can then render using conventional lighting and texturing.

For a mosaic to successfully convey an image, it is important to align tile edges with image edges. The use of square tiles imposes severe restrictions on the detail level that can be captured; our irregular tiles can convey considerable detail, including interior edges of figures, something often neglected in previous

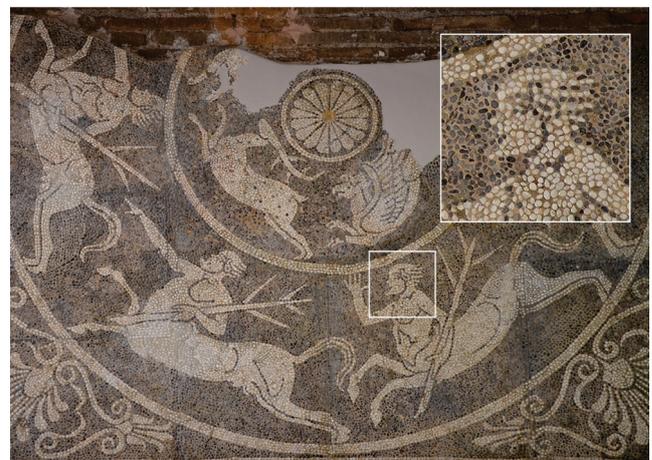


Fig. 1 Fragment of a pebble mosaic floor dating from the 4th century BCE.

¹ Carleton University, Ottawa, Canada. E-mail: lars.doyle@carleton.ca (✉).

Manuscript received: 2018-12-21; accepted: 2019-01-16



Fig. 2 An image progressing through our system. Left to right: input, segmentation, boundary smoothing, pebble geometry, lighting.

techniques. Our algorithm is entirely automatic; users can optionally guide the output by annotating the input image with an importance map or by manually adding decorative features in a preprocessing phase.

This paper makes two main contributions. Firstly, we adapt SLIC so that it is suitable for creating irregular, elongated pebble shapes. We estimate the local direction of the image and then bias the SLIC clustering distance according to a local coordinate system, producing natural-looking size and aspect ratio variations. Secondly, we compute smooth pebble geometry for the resulting tiles. We use the Laplace equation, setting up constraints and then solving to meet them, to produce smooth shapes resembling river pebbles. By creating and rendering this geometry, we bridge photorealism and non-photorealism.

This paper is organized as follows. Section 2 reviews previous work on computer-generated mosaics. Section 3 describes our algorithm in detail. Section 4 shows images created using our method and discusses its benefits and drawbacks. Finally, Section 5 summarizes the work and suggests future directions.

2 Related work

Battiato et al. [5] proposed a taxonomy of digital mosaic research in which the two initial branches divide *tile* mosaics from *multi-picture* mosaics. This distinction stems from the nature of the basic picture elements. In tile mosaics, the image plane is divided into small regions, each individually colored to represent the underlying input image. In contrast, multi-picture mosaics employ a dataset of images that are used to assemble an approximation to the input image based on local color and structure similarity; the typical result is a photomosaic [6]. We situate our current work within the tile mosaic branch.

In the seminal *Paint by Numbers* [7], Haeberli introduced many of the concepts that have since been used for mosaic emulation. His idea of using Voronoi diagrams for mosaics has been used in commercial

products and in subsequent research; centroidal Voronoi diagrams (CVDs) are particularly common. CVDs are often produced by Lloyd's algorithm, a relaxation process that repeatedly moves the Voronoi centres to the centroids of their regions. The CVD process has formed the basis for much work in mosaic and stipple creation [8–10], since it is a good way to distribute points on the plane.

Hausner [8] presented an iterative algorithm for placing mosaic centres, using hardware-accelerated CVDs to distribute tiles. Hausner also identified a crucial issue in mosaics: tile edges should be aligned with image edges. Hausner resolved this in his work by having tiles move away from user-specified edges. An alternative method for achieving edge alignment was given by Elber and Wolberg [11], who arranged rows of tiles along streamlines parallel to initial user-specified curves. Yet another way of addressing edge alignment was given by Di Blasi and Gallo [12], who proposed to cut the rectangular tiles where they cross image edges. Liu et al. [13] used graph cuts rather than explicit edge detection to prevent tiles from crossing image edges.

Within the multi-picture mosaic branch a thread of research involves populating a set of container shapes with tiles, generally without any intention of providing interior image detail. Kim and Pellacini's *Jigsaw Image Mosaic* [14] is an example, where the method produces an irregular tiling of the image plane with predefined tiles, minimizing a set of error criteria including tile overlap and color mismatch. More recent work by Saputra et al. [15, 16] arranges figures within the container shape while seeking an aesthetic distribution rather than a full packing. Kwan et al. [17] accelerated partial-shape matching, through their *pyramid of arc-length descriptor*, for packing irregular shapes.

Other methods for distributing primitives and tiling the plane have been devised, and we briefly mention a few others. Smith et al. [18] focussed on coherent movement of tiles to create animated mosaics; later, Dalal et al. [19] used Fourier transforms to find good

packings of input primitives. Kaplan and Salesin [20, 21] worked on automatically controlling tile shapes to produce Escher-like tilings in which the tiles were close to an input goal shape. Similarly, Goferman et al. [22] extracted irregular regions of interest from a series of photographs and packed them in a puzzle-like manner within a chosen aspect ratio. Photo collage is a related area, but removes the constraint that an underlying image or containing shape must be represented. Using convolutional neural networks, Liu et al. [23] produced photo collages by grouping together images with similar content over the image plane.

3 Constructing pebble mosaics

In our approach, we tile the image plane using heterogeneous, 3D pebble-shaped objects. As in previous methods [8, 11, 12], our tiles avoid crossing image boundaries and are oriented to align with a direction map. However, we take a different approach towards this goal. Section 3.1 describes how we modify the simple linear iterative clustering (SLIC) algorithm [4] to produce oriented pebble shapes. We take advantage of the inherent boundary-avoiding quality of SLIC and thus have no need for explicit edge detection nor associated parameters or thresholds. We describe how we simplify the boundaries of the initial segmentation in Section 3.2 to produce smooth, “river-worn” pebbles. Finally, in Section 3.3 we construct a height field from the 2D boundaries to extend pebbles into 3D, before applying lighting to the resulting geometry. A schematic representation of our algorithm pipeline in Fig. 3 shows how an input image I is transformed into a pebble mosaic.

3.1 Segmentation

SLIC produces compact super-pixels by clustering pixels into groups, based on colour and spatial distance. Their tendency to adhere to image boundaries is beneficial for describing image content and forms the basis of our pebble shapes. In its original formulation, the spatial distance of a pixel p from a cluster center c can be described by an offset vector $\vec{v} = p - c$, allowing us to compute the l_2 distance as

$$D_s = \sqrt{v_x^2 + v_y^2} \quad (1)$$

where v_x and v_y are the components of \vec{v} parallel

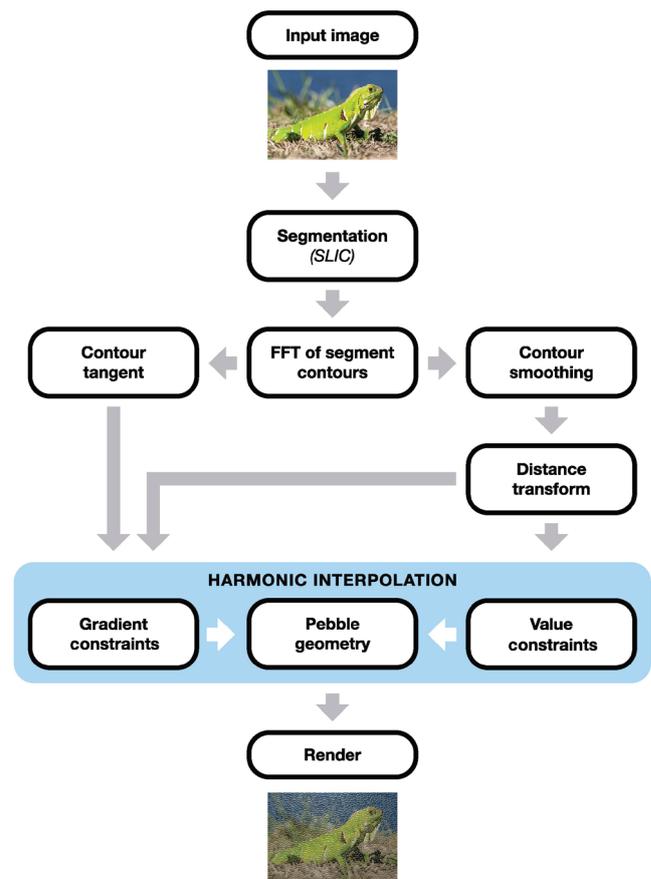


Fig. 3 Pipeline of our proposed method.

to the x - and y -axes. It is often the case in nature that pebbles are longer in one dimension than the other, forming approximate, oval-like boundaries as opposed to circular ones. As a first modification to Eq. (1) we can apply different scaling factors to the x and y components of \vec{v} . This results in the elongated super-pixels that are shown in the top right image of Fig. 4.

Artists often take advantage of pebble shapes and will emphasize image edges by aligning the long side of a pebble parallel to an edge. We can approximate this effect with one further modification to our distance metric. First, we construct a structure tensor [24] at each super-pixel center by integrating the matrix field, $\nabla I \nabla I^T$, weighted by a Gaussian function. The tensor’s unit eigenvectors e_1 and e_2 , associated with eigenvalues $\lambda_1 \geq \lambda_2$, are parallel and perpendicular to the smoothed image gradient. We can now use these vectors as a new basis in our distance calculation. Furthermore, applying a larger weight to the component of \vec{v} parallel to e_1 than e_2 allows super-pixels to spread tangent to image

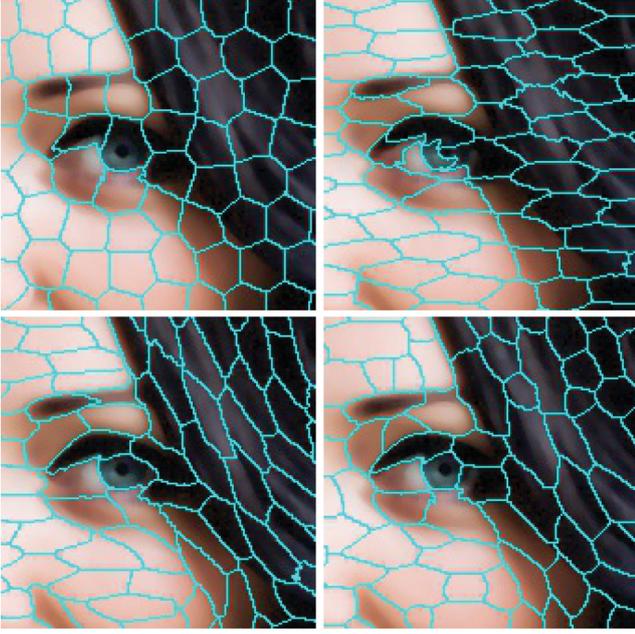


Fig. 4 Top left: original SLIC. Top right: scaling v_y in Eq. (1) by $\alpha = 3$. Bottom left: scaling $\vec{v} \cdot \vec{b}_1$ in Eq. (3) by $\alpha = 3$. Bottom right: scaling using random scaling in Eq. (3).

edges. This effect can be seen in Fig. 4(bottom left). In flat or corner regions, where there is inadequate orientation information, we simply assign a default direction. This assessment is made by thresholding an orientation coherence estimate:

$$C = \sqrt{\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + K}} \quad (2)$$

where K is a constant chosen to avoid division by zero and to de-emphasize weak tensors.

The final distance metric is

$$D_s = \sqrt{\alpha_1 \vec{v} \cdot \vec{b}_1 + \alpha_2 \vec{v} \cdot \vec{b}_2} \quad (3)$$

where α_1 and α_2 are scaling factors, controlling both the aspect ratio and the overall size of each cluster. The vectors \vec{b}_1 and \vec{b}_2 correspond either to the local image orientation, if a strong local orientation exists, or a default direction. The decision is made by comparing C to a threshold T_{coh} as follows:

$$\vec{b}_1 = \begin{cases} \vec{e}_1, & C > T_{\text{coh}} \\ \vec{d}_1, & C \leq T_{\text{coh}} \end{cases}, \quad \vec{b}_2 = \begin{cases} \vec{e}_2, & C > T_{\text{coh}} \\ \vec{d}_2, & C \leq T_{\text{coh}} \end{cases} \quad (4)$$

The vectors d_1 and d_2 comprise a default orthonormal basis. In our examples we set T_{coh} to 0.5 and d_1 to the y -axis.

The scaling factors α_1 and α_2 are selected individually for each super-pixel guided by a random

process, such that

$$\alpha_1 = \phi_{a1} \phi_s, \quad \alpha_2 = \phi_{a2} \phi_s \quad (5)$$

Through experimentation, we chose to compress the aspect ratio perpendicular to edges by $\phi_{a1} = 3$. The other terms are determined by two uniform random numbers $r_1, r_2 \in [0, 1]$. We then set $\phi_{a2} = (\phi_{a1} - 1)r_1^2 + 1$ and set the scale term $\phi_s = r_2^2 + 1$.

The local distance metric D_s is used in the SLIC process to oversegment the image. We refer to the resulting oversegmented image as P , and each segment, $P_i \in P$, is a *pebble*.

3.2 Boundary smoothing

The pebbles constructed in Section 3.1 contain many irregularities that depart from the smooth pebble shapes that we wish to create. Hence, we apply a low-pass filter in the frequency domain [25] to each pebble's outer contour $c_o(k)$, for $k = 0, 1, \dots, K - 1$. This process effectively reconstructs a contour from L Fourier coefficients, where $L < K$. In Fig. 5, we illustrate a contour reconstructed with various values of L . Note that at this stage we can obtain a resolution-independent, tiled image P by applying a scale factor to the Fourier coefficients, obtaining a larger (or smaller) c_o as needed. As seen in Fig. 6, one advantage of rendering at a higher resolution is the increased surface area that can be used for adding texture and lighting. Finally, each pebble $P_i \in P$ is updated by flood filling its corresponding c_o .

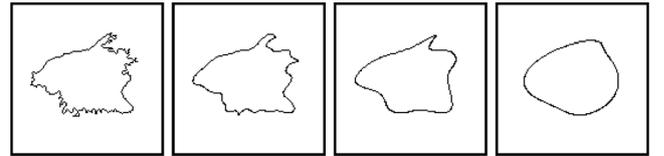


Fig. 5 Left to right: original contour, and reconstructed contours using $L = 37, 17,$ and 7 Fourier coefficients.



Fig. 6 High-resolution pebbles rendered at 5 times the input resolution. Left: pebble shapes. Right: 3D rendered pebbles.

3.3 Pebble geometry

We construct a height field for each pebble by means of harmonic interpolation over the domain, Ω , that resides between two contours (Fig. 7(left)). The outer contour, c_o , is described above. We obtain the inner contour, c_i , by thresholding the normalized distance transform of P_i by $T_{\text{dist}} \in (0, 1)$. We set a zero gradient at the inner contour, thus creating a small flat face to each pebble which then curves downwards to the image plane. In all examples, we set $T_{\text{dist}} = 0.85$.

Our height field is the solution to the Laplace equation [26]:

$$\Delta P_i = 0 \text{ over } \Omega \tag{6}$$

with boundary value constraints $P_i|_{c_o} = 0$ and $P_i|_{c_i} = 1$. Additionally, we set gradient constraints at the boundaries such that $|\nabla P_i| = 0$ on c_i . The gradient on c_o is constructed as follows. Returning to the Fourier transform of Section 3.2, we note that the derivative $c_o'(k)$ of the sampled function $c_o(k)$ can be computed in the Fourier domain. This process provides us with a sequence of vectors that are tangent to the curve, one for each sample point. Rotating each vector 90° inwards gives us a gradient orientation that is orthogonal to the boundary. The gradient magnitude is chosen as follows:

$$|\nabla P_i| = \frac{\beta}{T_{\text{dist}} D_{\text{max}}} \tag{7}$$

where D_{max} is the maximum value of the distance transform. The parameter β determines the shape of the resulting pebble; various settings are illustrated in Fig. 8. We choose $\beta = 2$ to construct the pebble profile curving downward into the surrounding area in our examples. Notice that setting β too high results in the gradient overshooting its target at the inner contour resulting in a depression at the center as seen in the bottom row of Fig. 8.

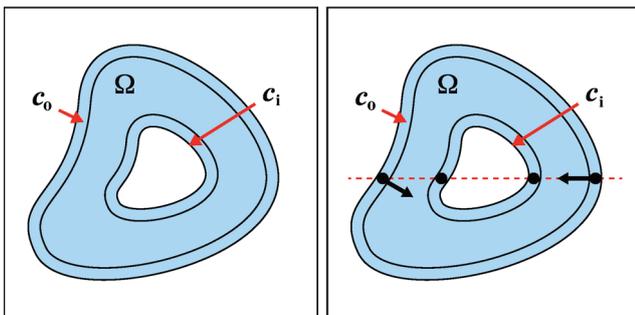


Fig. 7 Left: the domain, Ω , and boundaries (c_o and c_i) of P_i . Right: the gradient orientation on c_o (arrows) and zero-gradient on c_i (dots).

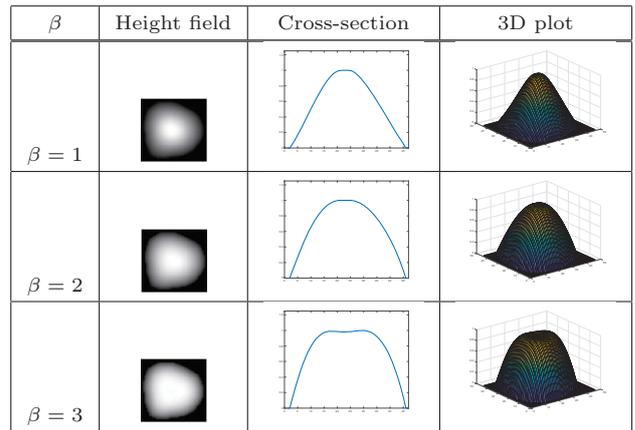


Fig. 8 Constructing a height field at varying scales of gradient magnitude on c_o .

3.4 Rendering

We apply Phong shading to the resulting height field. We use the average colour in I under P_i as the pebble's surface color. Optionally, we can apply a rock texture to the pebble as well. The texture image is randomly sampled for each pebble and combined with the luminosity channel using a multiply blend. Example mosaics produced using this scheme, with and without texture, are shown in Figs. 9(above) and 9(below) respectively.



Fig. 9 Above: result without texture. Below: textured result.

4 Results and discussion

We demonstrate our method on photographs containing various subject matter in Fig. 10, using 2000 pebbles in each example. The original source images are shown in Fig. 20. Rendering time for a 1.5 megapixel image is 28 s using our unoptimized CPU implementation. The majority of this time (25 s) is spent solving 2000 N_i^2 sparse linear systems in order to construct the geometry of the pebbles. Increasing the pebble count leads to solving smaller matrices and thus faster execution time; for example, using 3000 pebbles reduces the solving time to 16 s.

Notice that even at this coarse scale, most of the important image features are still recognizable. The elongated pebble shapes add an impression of motion to the results. This is most noticeable in the cat image at the top left where the pebbles follow the fur orientation. In the portrait image (second row, left) we see how random pebble scaling can add visual interest to otherwise flat image regions. This brings to mind the activity of a mosaicist using tiny pebbles to fill the empty spaces left between larger stones. In the bottom row, adding texture supports the transition from the synthetic 3D shapes in the top rows to a more natural-looking material.



Fig. 10 Results. Top two rows: without texture. Bottom row: using marble texture.

Inspired by historical mosaics, such as the one depicted in Fig. 1, we demonstrate our method on the ornamental designs shown in Fig. 11. Due to the high contrast in these images, the pebbles adhere well to the image content, creating a striking re-representation of the input.

4.1 Degrees of freedom

Our system has five notable degrees of freedom that can influence the outcome of the final rendered mosaic: color, shape, texture, orientation, and size. We briefly discuss each here.

Color. Following the tradition in tile mosaics [7, 8, 11, 12] we render each pebble with the average color under the corresponding image region. Alternatively, we could allow color to vary over the pebble region, guided by the input image.

Shape. Pebble shape can be influenced by the low-pass filter used in the smoothing process discussed in Section 3.2 and illustrated in Fig. 5. We chose to retain seven Fourier coefficients, resulting in smooth oval-like pebble shapes. However, less smoothing would provide more shape variety.

Texture. We currently limit pebble texture to a single sample but there is potential for more development along this dimension. For example, a database of texture swatches could be employed to match pebble texture with the underlying image. This addition would provide further connection with



Fig. 11 Results for ornamental motifs. Left: input images. Right: results.

the input image and increase recognizability.

Orientation. Pebbles are oriented parallel to image edges, as is common in both traditional and digital mosaics [8, 11, 12]. As explained in Section 3.1, we determine orientation through a structure tensor field, defaulting to a fixed orientation where inadequate information is present. We could also ask the user to provide a vector field in place of a single default direction.

Size. We discuss pebble size in the following subsections, first talking about local variation in pebble dimensions and then discussing size more generally, including the option of varying pebble size based on an importance map.

4.1.1 Pebble dimensions

In Section 3.1 and Eq. (5) we describe a random process that determines the aspect ratio and relative size of individual pebbles. We now show how varying these parameters can influence the resulting mosaic; the images in the top row of Fig. 12 provide a visual example. In Fig. 12(top left) we fix $\phi_s = 1$ to maintain a constant scale and vary the aspect ratio using a random number. Here we increase ϕ_{a1} to 5 and calculate ϕ_{a2} as before. The long thin pebbles work well in this situation where we connect them with the cat's fur. Compare this result to the cat in Fig. 10. There, setting ϕ_{a1} to 3 shows less movement in the cat's fur, but randomly changing ϕ_s brings out more variation and liveliness. In Fig. 12(top right), we fix the aspect ratio to $\phi_{a1} = \phi_{a2} = 1$ and allow the scale parameter to vary. We set $\phi_s = 5r^2 + 1$, where r is a uniform random number in $[0, 1]$. Without orientation information it is more difficult to identify the image. Also, such

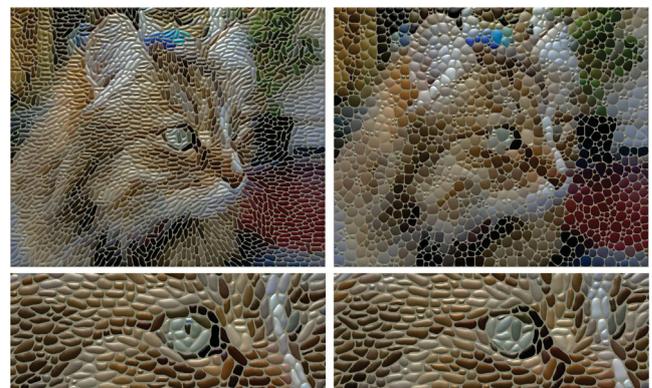


Fig. 12 Top left: randomly varying pebble aspect ratio, fixed scale. Top right: randomly varying pebble scale, fixed aspect ratio. Bottom: rendering is nondeterministic due to random scale parameter.

extreme variability in pebble size is distracting since the sizes are chosen randomly rather than based on image content. In Fig. 12(bottom, left and right) we demonstrate the impact of the random factors in the scaling parameters: note the different outcomes for two runs, using identical parameters.

4.1.2 Pebble size

In Fig. 13 we vary the number of pebbles that make up a mosaic image. On the left we see a detailed result using 3000 pebbles. Many traditional mosaics, such as the one depicted in Fig. 1, were constructed with this high level of detail. Next, we see a result using 1000 pebbles. Even at this larger size, much of the image remains clear owing to SLIC's tendency to adhere to image boundaries. Finally, the pebble size on the right has probably been pushed too far, making it difficult to recognize the main figure in the result. See Fig. 15 for a rendering of this image using 2000 pebbles.

We can also vary the pebble size by use of an importance map. The mask in the inset of Fig. 14 indicates regions to be rendered with smaller, more numerous pebbles. This technique is useful for drawing attention to important regions and provides



Fig. 13 Varying pebble size. Left to right: 3000, 1000, 500 pebbles.

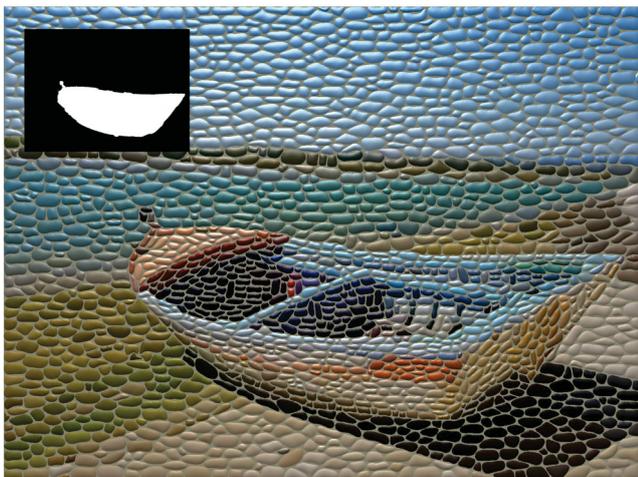


Fig. 14 Pebbles within the important area (inset) are rendered at a higher frequency.

a more detailed representation of the content.

4.2 Comparison with related work

Figure 15 shows a comparison between our method and Hausner's [8] using 2000 pebbles. Here, we turn off the lighting effects and make the comparison based on tile shape alone. (The color shift between the two examples is due to using different source photographs of the painting). By using heterogeneous shapes, image content can be more accurately portrayed than when using an equal number of 2D homogeneous primitives. In our result, the pebble shapes cleanly outline the contours of the figure and its drapery. Where smaller pebbles are needed to fill an image region, our method is not restricted to a uniform pebble size. Both these properties stem from our use of SLIC as the initial segmentation method. Of course, both our method and Hausner's can use smaller primitives in regions specified by users.

Similarly, we compare our method with three previous tile mosaic algorithms on a common image in Fig. 16. Our result is on the bottom right using 3000 pebbles. At the top left, Di Blasi and Gallo [12] obtain clean lines and uniform spacing by cutting tiles that overlap perceptual guidelines and neighbouring tiles. The edges in our rendering are obtained through SLIC which adhere well to step edges but fail when perceptual boundaries are not matched with a strong color discontinuity. An example can be seen in the thin strand of feathers above the brim of Lena's hat where pebbles are not constrained to this narrow region. This is a case in which perceptual edge detection would benefit our segmentation. Schlechtweg et al.'s [27] *RenderBots* show fine detail by using 9000 primitives but the placement is uneven and rendering took one hour to complete.



Fig. 15 Comparison. Left: Hausner's method. Right: ours. Both results use 2000 tiles

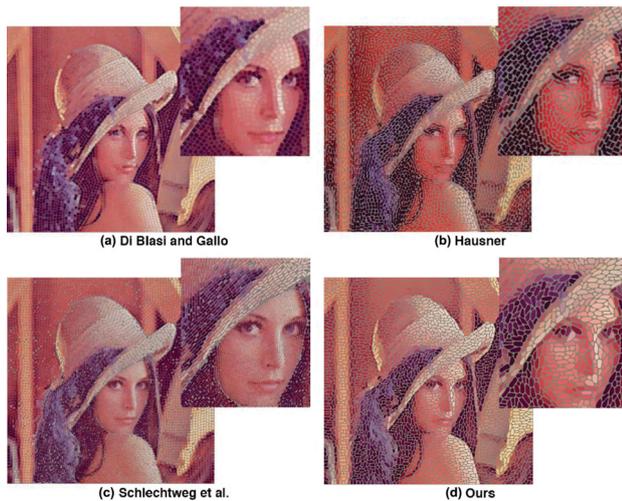


Fig. 16 Comparison with previous tile mosaic algorithms.

Recently, there has been much attention given to using convolutional neural networks for image stylization [28, 29]. In Fig. 17(center) we show a result obtained from *deepart.io*, a popular online implementation of Gatys et al.’s method [28]. The high-level semantic features used in neural style transfer preserve image features better than the low-level color features that we use; compare the detail images in Figs. 17(bottom left) and 17(bottom right). The iguana’s eye clearly highlights the advantage of using semantic features: style transfer reproduced the eye using a single pebble, improving recognizability. Our method, in contrast, uses a number of pebbles dependent on the SLIC super-pixel size; it artificially breaks the eye into three pebbles. The advantage of our method lies in explicitly modeling pebble shapes. The texture produced by neural style transfer in Fig. 17(center) only roughly approximates that found in the style example in Fig. 17(top left). For example, the definition of individual pebbles is completely lost in parts of the background and the side of the iguana’s head. In contrast, our method explicitly models individual pebble geometry and can output well-defined shapes at any resolution.

4.3 Limitations

Our method performs best on images with high contrast and clear distinctions between regions of differing semantic content. Due to the relatively large scale of the pebbles, some subtle image features or tiny details can be lost. Figure 18(top) shows an image dominated by high-frequency content. In our rendering in Fig. 18(right), only a large-scale

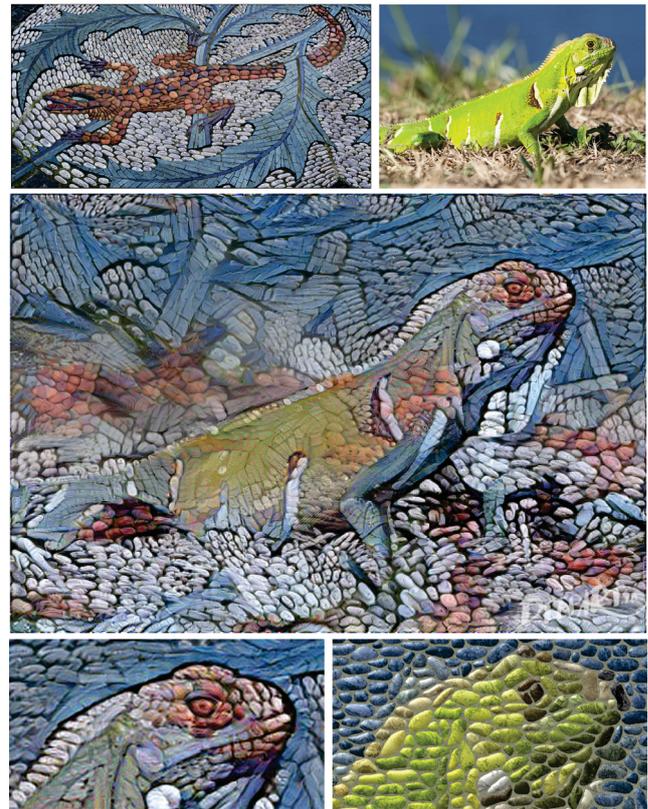


Fig. 17 Comparison with neural style. Top left: style example. Top right: input image. Center: pebble mosaic rendered with neural style [28] as implemented at *deepart.io*. Bottom left: detail. Bottom right: detail of Fig. 9(bottom).

impression of the scene is captured. Reducing pebble size is only a limited option since, past a certain scale, the cement between the pebbles will feature as prominently as the pebbles themselves. In Fig. 18(bottom) the facial features are poorly represented. SLIC does not effectively cope with lighting changes in the area of the man’s nose, for example. Either more sophisticated low-level processing or learning-based semantic segmentations could improve on our results, and both are promising directions for future work.

Continuing our discussion on color, we also note that our resulting images would be difficult to recognize based on pebble layout only. See Fig. 19 for an example of a black and white pebble layout. Without colorization, the orientation and pebble boundaries only hint at the underlying image. More work could be done to emphasize the structural content of the image by varying pebble shape and size, linking size and shape variation to image content instead of varying pebble dimensions with random factors. At the same time, it might be possible to



Fig. 18 Limitations of our method. Top: high-frequency features. Bottom: semantic content.



Fig. 19 Pebble layout without colorization.

improve our pebble colors. Because we add lighting effects to a base color derived from the image, the final pebble color distribution is not necessarily very close to the desired color. We could improve the mosaic by better integration of the lighting process and the selection of base color.

Processing time is also an issue. Our main bottleneck is determining the numerous matrices that construct the height field. Taking advantage of parallelization would help. Also, solving at a lower resolution and smoothing the results could improve timing.



Fig. 20 Input images used in Figs. 9, 10, 13, 14, and 15.

Although we think that smooth river-worn pebbles are the most commonly type used for pebble mosaics, more varied rock types in principle could be used, and this paper did not attempt to treat these.

5 Conclusions

In this paper we present a method to render 3D pebble mosaics. Digital mosaics have been presented in the NPR literature previously, but only in the context of tiling a 2D surface; here, we not only create a tiling representing pebbles, but also generate a height field for the pebbles so that they can be rendered.

Our method starts by segmenting the image plane with SLIC, equipped with a modified distance metric. The resulting super-pixels adhere to image boundaries and hence no further edge detection is required. By varying the size, orientation, and aspect ratio of the super-pixels, we obtain pebble shapes that are highly expressive in their depiction of image content.

We construct the geometry of each pebble by solving a Laplace equation on the domain between two contours. The resulting height field can then be rendered using a variety of lighting techniques beyond the simple Phong shading model we use in this paper. In addition, since we have synthesized 3D geometry, our pebble mosaics can be used in novel applications, from 3D virtual environments to physical 3D printed objects.

In the future we would like to use semantic segmentation to improve the initial super-pixel clustering. Important image regions, especially on the human face, could benefit by constraining clustering to regions of similar content. Better use of low-level image features could improve on the SLIC segmentation. Pebble texture could also be customized to suggest image details at a scale below the size of individual pebbles. This addition would bridge the gap between tile and multi-picture mosaics, as defined by Battiato et al. [5], and strengthen the connection between the original image and its mosaic representation.

Acknowledgements

We would like to thank the anonymous reviewers for many insightful comments. We also thank members of the Graphics, Imaging and Games Lab for productive comments and discussions. Funding for this work was

provided by NSERC, OGS, and Carleton University.

We used many images from Flickr under a Creative Commons license. Thanks to the numerous photographers who provided material: Douglas Scortegagna (landscape), bDom (b&w portrait), Julio Romero (iguana), Peat Bakke (t-rex), Gábor Lengyel (portrait), Tommie Hansen (canal), Theen Moy (cat), JB Banks (dark woods), Richard Messenger (Yemeni), Greg Myers (tomatoes), sicknotepix (toque).

References

- [1] Dunbabin, K. M. D. *Mosaics of the Greek and Roman World*. Cambridge University Press, 1999.
- [2] Ling, R. *Ancient Mosaics*. British Museum Press, 1998.
- [3] Howarth, M. *The Complete Pebble Mosaic Handbook*. Firey Books, 2003.
- [4] Achanta, R.; Shaji, A.; Smith, K.; Lucchi, A.; Fua, P.; Süsstrunk, S. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 34, No. 11, 2274–2282, 2012.
- [5] Battiato, S.; di Blasi, G.; Farinella, G. M.; Gallo, G. Digital mosaic frameworks—An overview. *Computer Graphics Forum* Vol. 26, No. 4, 794–812, 2007.
- [6] Silvers, R. *Photomosaics*. Henry Holt and Co., Inc., 1997.
- [7] Haeberli, P. Paint by numbers: Abstract image representations. In: Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques, 207–214, 1990.
- [8] Hausner, A. Simulating decorative mosaics. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, 573–580, 2001.
- [9] Secord, A. Weighted Voronoi stippling. In: Proceedings of the 2nd International Symposium on Non-Photorealistic Animation and Rendering, 37–43, 2002.
- [10] Hiller, S.; Hellwig, H.; Deussen, O. Beyond stippling—Methods for distributing objects on the plane. *Computer Graphics Forum* Vol. 22, No. 3, 515–522, 2003.
- [11] Elber, G.; Wolberg, G. Rendering traditional mosaics. *The Visual Computer* Vol. 19, No. 1, 67–78, 2003.
- [12] Di Blasi, G.; Gallo, G. Artificial mosaics. *The Visual Computer* Vol. 21, No. 6, 373–383, 2005.
- [13] Liu, Y.; Veksler, O.; Juan, O. Generating classic mosaics with graph cuts. *Computer Graphics Forum* Vol. 29, No. 8, 2387–2399, 2010.
- [14] Kim, J.; Pellacini, F. Jigsaw image mosaics. *ACM Transactions on Graphics* Vol. 21, No. 3, 657–664, 2002.
- [15] Saputra, R. A.; Kaplan, C. S.; Asente, P.; Měch, R. FLOWPAK: Flow-based ornamental element packing. In: Proceedings of Graphics Interface Conference, 8–15, 2017.
- [16] Saputra, R. A.; Kaplan, C. S.; Asente, P. RepulsionPak: Deformation-driven element packing with repulsion forces. In: Proceedings of Graphics Interface Conference, 10–17, 2018.
- [17] Kwan, K. C.; Sinn, L. T.; Han, C.; Wong, T. T.; Fu, C. W. Pyramid of arclength descriptor for generating collage of shapes. *ACM Transactions on Graphics* Vol. 35, No. 6, Article No. 229, 2016.
- [18] Smith, K.; Liu, Y.; Klein, A. Animosoics. In: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, 201–208, 2005.
- [19] Dalal, K.; Klein, A. W.; Liu, Y.; Smith, K. A spectral approach to NPR packing. In: Proceedings of the 4th International Symposium on Non-Photorealistic Animation and Rendering, 71–78, 2006.
- [20] Kaplan, C. S.; Salesin, D. H. Escherization. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, 499–510, 2000.
- [21] Kaplan, C. S.; Salesin, D. H. Dihedral escherization. In: Proceedings of the Conference on Graphics Interface, 255–262, 2004.
- [22] Goferman, S.; Tal, A.; Zelnik-Manor, L. Puzzle-like collage. *Computer Graphics Forum* Vol. 29, No. 2, 459–468, 2010.
- [23] Liu, L. J.; Zhang, H. J.; Jing, G. M.; Guo, Y. W.; Chen, Z. G.; Wang, W. P. Correlation-preserving photo collage. *IEEE Transactions on Visualization and Computer Graphics* Vol. 24, No. 6, 1956–1968, 2018.
- [24] Brox, T.; van den Boomgaard, R.; Lauze, F.; van de Weijer, J.; Weickert, J.; Mrázek, P.; Kornprobst, P. Adaptive structure tensors and their applications. In: *Visualization and Processing of Tensor Fields. Mathematics and Visualization*. Weickert, J.; Hagen, H. Eds. Springer Berlin Heidelberg, 17–47, 2006.
- [25] Gonzalez, R. C.; Woods, R. E. *Digital Image Processing*, 3rd edn. Prentice Hall, 2008.
- [26] Pérez, P.; Gangnet, M.; Blake, A. Poisson image editing. *ACM Transactions on Graphics* Vol. 22, No. 3, 313–318, 2003.
- [27] Schlechtweg, S.; Germer, T.; Strothotte, T. RenderBots—Multi-agent systems for direct image generation. *Computer Graphics Forum* Vol. 24, No. 2, 137–148, 2005.
- [28] Gatys, L. A.; Ecker, A. S.; Bethge, M. Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 2414–2423, 2016.

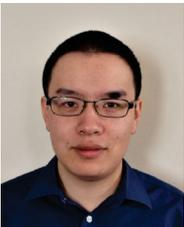
- [29] Johnson, J.; Alahi, A.; Fei-Fei, L. Perceptual losses for real-time style transfer and super-resolution. In: *Computer Vision—ECCV 2016. Lecture Notes in Computer Science, Vol. 9906*. Leibe, B.; Matas, J.; Sebe, N.; Welling, M. Eds. Springer Cham, 694–711, 2016.



Lars Doyle is a Ph.D. student in the School of Computer Science at Carleton University where he works in the Graphics, Imaging, and Games Lab. His research interests focus on image processing, image stylization, and super-resolution. He received his master and bachelor degrees in computer science from Carleton University. Previously, he worked as a graphic designer.



Forest Anderson is a fourth-year undergraduate student at Carleton University, majoring in computer science. He is interested in image processing, game development, and computer security.



Ehren Choy is a software developer at Ultra Electronics. He received his bachelor degree of software engineering from the University of Waterloo, and his master degree of computer science from Carleton University.



David Mould received his Ph.D. degree in computer graphics from the University of Toronto in 2002. Prior to this, he earned his B.Sc. degree in combined physics and computer science from the University of British Columbia. He was a faculty member at the University of Saskatchewan for six years, and is presently a professor at Carleton University, where he co-directs the Graphics, Imaging, and Games Lab in the School of Computer Science. Dr. Mould's research interests include game design, non-photorealistic rendering, and procedural natural phenomena.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.