# Composite Digital Mosaics using Duotone Tiles

Gary R. Greenfield
Department of Mathematics & Computer Science
University of Richmond
Richmond, VA 23173, USA
ggreenfi@richmond.edu

## Abstract

We describe a method for combining a foreground image and a background image to make digital mosaics using duotone tiles. Unlike the tiles of thumbnail images used in photomosaics, our duotone tiles each consist of precisely two colors. Our method uses diffusion limited aggregation (DLA) to approximate via a curve the foreground image which is then composited over a pixellated enlargement of the background image.

## 1 Introduction

The current popularity of digital mosaics is surely due to the success of Silvers' well publicized, and widely marketed, Photomosaic technique [11]. Computer generated mosaics such as those by computer artist Robert Bosch that rely on sophisticated optimization algorithms [1] continue to advance the state of the art.

An ongoing challenge in digital mosaics is to make the resulting imagery more visually exciting by, for example, embedding messages, motifs, or subliminal images. The original Photomosaic technique of Silvers accomplishes this to a limited extent by (1) constructing high resolution mosaics, (2) using a large database of tiles, and (3) using a database all of whose images have the same thematic content. However, if we restrict ourselves to working with small sets of tiles as in the case of the domino mosaics of Knowlton [5], Knuth [6] and Bosch, or the automatically generated sets of grayscale tiles used in the photomosaics of Sah et al. [10], this issue becomes more problematic. Figure 1 shows two examples using small sets of tiles: a Mona Lisa domino mosaic portrait by Robert Bosch, and Andrew Pike's 2007 "Sierpinski Carpet", a high resolution portrait mosaic whose tiles were induced from gray scale versions of the Sierpinski carpet. Our goal is composite a background image and a foreground image into a digital mosaic in such a way as to make the background image almost subliminal. Further we would like to reduce the number of colors, and hence tiles, required.

Our approach to bringing multiple images into play for digital mosaics while simultaneously restricting the set of tiles that can be used is new in two significant ways. First, rather than drawing from only a small set of tiles, we use *duotone* tiles meaning each tile contains a pair of distinct colors. Second, the foreground image overlaid on our background image is approximated thanks to the use of simulated diffusion limited aggregation (DLA) by a continuous curve. Thus our overlay technique bears some similarity to the TSP art of Kaplan and Bosch [4], or the space filling curve mosaics of Knowlton (see http://www.knowltonmosaics.com/). We remark at the outset that the inspiration for our technique stems from a series of prints by I. Orosz exhibited at Bridges 2008 and reproduced in that exhibition pamphlet [9]. Online examples [12] of Orosz' prints of this type include the etchings "Dürer in the Forest" and "Dalí and the Holy Family." The most closely related work to ours that we are aware of appears in the thesis of Long [8]. Long explored the use of diffusion limited aggregation for special effects in commercial and representational imagery.

Stated formally, our task is to digitally composite two low resolution identically sized $m \times n$ images — a background image and a foreground image — in such a way as to produce a duotone tile digital mosaic. Since our results are reproduced here in grayscale we will focus on a limited number of examples. Figure 2 shows the Tanzanian sunset image, East African giraffe image, and author's portrait image we will use in

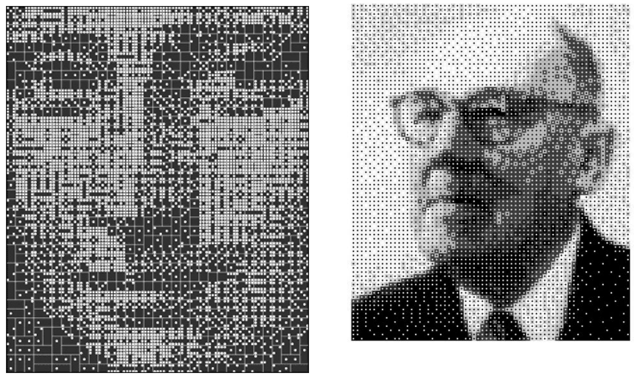

**Figure 1**: Left: Domino mosaic portrait of the Mona Lisa by Robert Bosch. Right: "Sierpinski Carpet" by Andrew Pike. Reprinted with permission.

this paper. Figure 3 shows the two duotone composites we obtained with the sunset and giraffe alternating as both foreground and background. To create them we cropped and downsampled the sunset and giraffe images to yield low resolution $70 \times 52$ pixel images. Downsampling reduces the number of colors we have to work with. Each duotone tile in these two composites is $21 \times 21$ pixels.

**Figure 2**: Left: Tanzanian sunset image. Center: East African giraffe. Right: Author's portrait.

## 2 Overview of the method

The "background mosaic" is simply a pixellated version of the background image. It is obtained by expanding each pixel of the low resolution background image to a block of pixels of the appropriate size. This yields a digital mosaic of the background image where each tile has a single color. To accomplish our overlay process, we rely on a modification of a simulated diffusion limited aggregation (DLA) technique for image *magnification* that was previously described by Greenfield [3]. We want to ensure that the DLA image that represents the foreground image will be perceived as a continuous curve so that much of the blockiness that normally arises when trying to expand our foreground image is suppressed. As a consequence both the low resolution background mosaic and the effectively higher resolution DLA foreground image can reinforce one another. To manage this we glue together several DLA structures deposited on each background tile in such a way that all tiles remain duotone. To this end, we invoke a three pass algorithm involving the foreground image. The first two passes deposit DLA structures on each background tile of the same color as the unique corresponding pixel of the foreground image. This second color makes every tile a duotone tile. The third and final pass glues these DLA structures together using additional DLA structures in such a way as to preserve the duotone nature of the tiles.

## 3 An example using the method

We provide details by working through an example. We will choose the sunset image and the author's portrait. To help reveal composite image detail, these images are first cropped and downsampled to yield $30 \times 22$ pixel images and the tiles in the resulting $30 \times 22$ duotone composite mosaics will each be $41 \times 41$ pixels. Figure 4 shows the results following the initial step where each image in turn is treated as the background image. It clearly reveals how each pixel in the background image is expanded into a monochrome tile that is a $41 \times 41$ block of pixels.

The first two of the three passes that use the foreground image aggregate a DLA structure within each tile by placing a seed particle at the center of each tile and then generating a DLA structure for the tile whose color corresponds to that of the associated pixel of the foreground image. The two structures are generated independently. The DLA aggregation method of Kobayashi et al. [7] as further refined by Greenfield [2] is invoked. To briefly review the aggregation process, for these two passes the center of the tile serves as the local origin, and one at a time additional particles are released and subsequently undergo a random walk until they either encounter the existing structure and adhere, or their time limit expires. When the structure grows to the appropriate size (here extending out to within one pixel of the tile's border), this process stops.

The third and final pass involving the foreground image again considers each pixel of the foreground image in turn. But now after it identifies the tile in the mosaic corresponding to the pixel it performs a search near the boundaries of the neighboring tiles looking for an aggregated particle. Finding such a particle means we have located the extremity of one of the DLA structures in some neighboring tile. If such a particle is found, then it serves as the local origin for a third DLA structure to be aggregated. This time we must be more careful. The purpose of this third structure is to glue existing aggregated structures together. Thus, in order to preserve the duotone property of the tiles, we must check before aggregating any subsequent particle to see which color must be assigned to it, i.e., to which tile it actually belongs. The other subtle point is that the search along the borders of neighboring tiles must be done pseudorandomly to ensure that no visual artifacts or biases will result. Figure 5 shows the results of these three aggregation passes by focusing on six foreground pixels. Notice in Figure 5 how in the last pass the three tiles on the left are joined to neighboring tiles that are not shown, but the three tiles on the right are joined to each other. Figure 6 shows the results using the sunset as the background and the giraffe as the foreground after each of the first two foreground passes. The results following the last pass are shown in Figure 7. For the sake of completeness, and for side by side comparison, in Figure 7 we show both duotone composites, i.e., the composites with foreground and background roles reversed.

## 4 Discussion

There are three key points that should be raised. First, because of the blockiness of the background image in the composite mosaic, care must be exercised when choosing background versus foreground to achieve an effective result. Second, we trust the reader realizes that we have used duotone tiles in order to eventually facilitate the construction of actual *physical* mosaics. At worst, if all the pixels in the downsampled $m \times n$ images had distinct colors, we would need sufficient quantities of $2mn$ distinct smaller tiles on hand to construct our composite mosaics. This is because the $p \times p$ enlarged duotone tiles that make up the composite can then each be formed using $p^2$ tiles selected from our repository of smaller ones. Of course, in practice it would be rare for all $2mn$ pixels to be be distinct. Third, there is no guarantee our search for a neighboring particle to serve as the seed during the final pass will be successful for every foreground image pixel, nor that all the structures aggregated during the first two DLA passes will be joined during this last pass. Empirically we have observed that the search for a suitable neighboring seed is almost always successful and that there is usually a large connected component that forms in the composite during this pass [3]. We have not taken the extra trouble to correct for these two possible shortcomings because it does not appear that the human eye is that discerning.

## 5 Conclusions and future work

We have shown that our method for generating duotone composite mosaics is viable. Clearly, some skill is required in selecting, cropping, and downsampling source imagery in order to achieve superior results. Future work should be directed towards mitigating the effects of the blockiness of the background image in the composite and developing methods for handling imagery with different height and width dimensions.

## 6 Acknowledgment

This work was supported in part by a University of Richmond 2008 Arts & Sciences Summer Fellowship.

## References

[1] Bosch R., Constructing domino portraits, *Tribute to a Mathemagician*, Cipra B. (Ed.), 2004, 251–256.

[2] Greenfield G., Composite diffusion limited aggregation paintings, *Bridges Donostia: Bridges 2007: Mathematical Connections in Art, Music and Science*, Sarhangi R. and Barrallo J. (Eds.), 2007, 15–20.

[3] Greenfield G., Connectivity and a diffusion limited aggregation digital image magnification technique, *Proceedings of the 13-th International Conference on Geometry and Graphics*, Weiss G. (Ed.), CD-ROM (ISBN 978-3-86780-042-6), 2008.

[4] Kaplan C. S. and Bosch R., TSP art, *Renaissance Banff: Bridges 2005: Mathematical Connections in Art, Music and Science*, Sarhangi R. and Moody R. V. (Eds.), 2005, 301-308.

[5] Knowlton, K., Knowlton Mosaics: Computer Assisted Portrait Art, http://www.knowltonmosaics.com/ (accessed 1 April 2009).

[6] Knuth, D., The Stanford GraphBase: A Platform for Combinatorial Computing, *ACM Press*, New York, 1993.

[7] Kobayashi Y., Niitsu T., Takahashi K., and Shimoida S., Mathematical modeling of metal leaves, *Mathematics Magazine*, Volume 76, Number 4, 2003, 295–298.

[8] Long J., Modeling dendritic structures for artistic effects, *M.Sc. Thesis*, University of Saskatchewan, 2007.

[9] Orosz I., Istvan Orosz Bridges 2008, *Exhibition Publication*, Utisz Bt., 2008.

[10] Sah S. B., Ciesielski V., Sousa D., and Berry M., Comparison between genetic algorithm and genetic programming performance for photomosaic generation, *Simulated Evolution and Learning 2008 Proceedings*, Li X. et al. (Eds.), Springwer Verlag LNCS 5361, 2008, 259–268.

[11] Silvers R., *Photomosaics*, Henry Holt and Company, New York, NY, 1997.

[12] UTISZ, The homepage of Istvan Orosz, http://web.axelero.hu/utisz/page.htm (accessed March 2009).

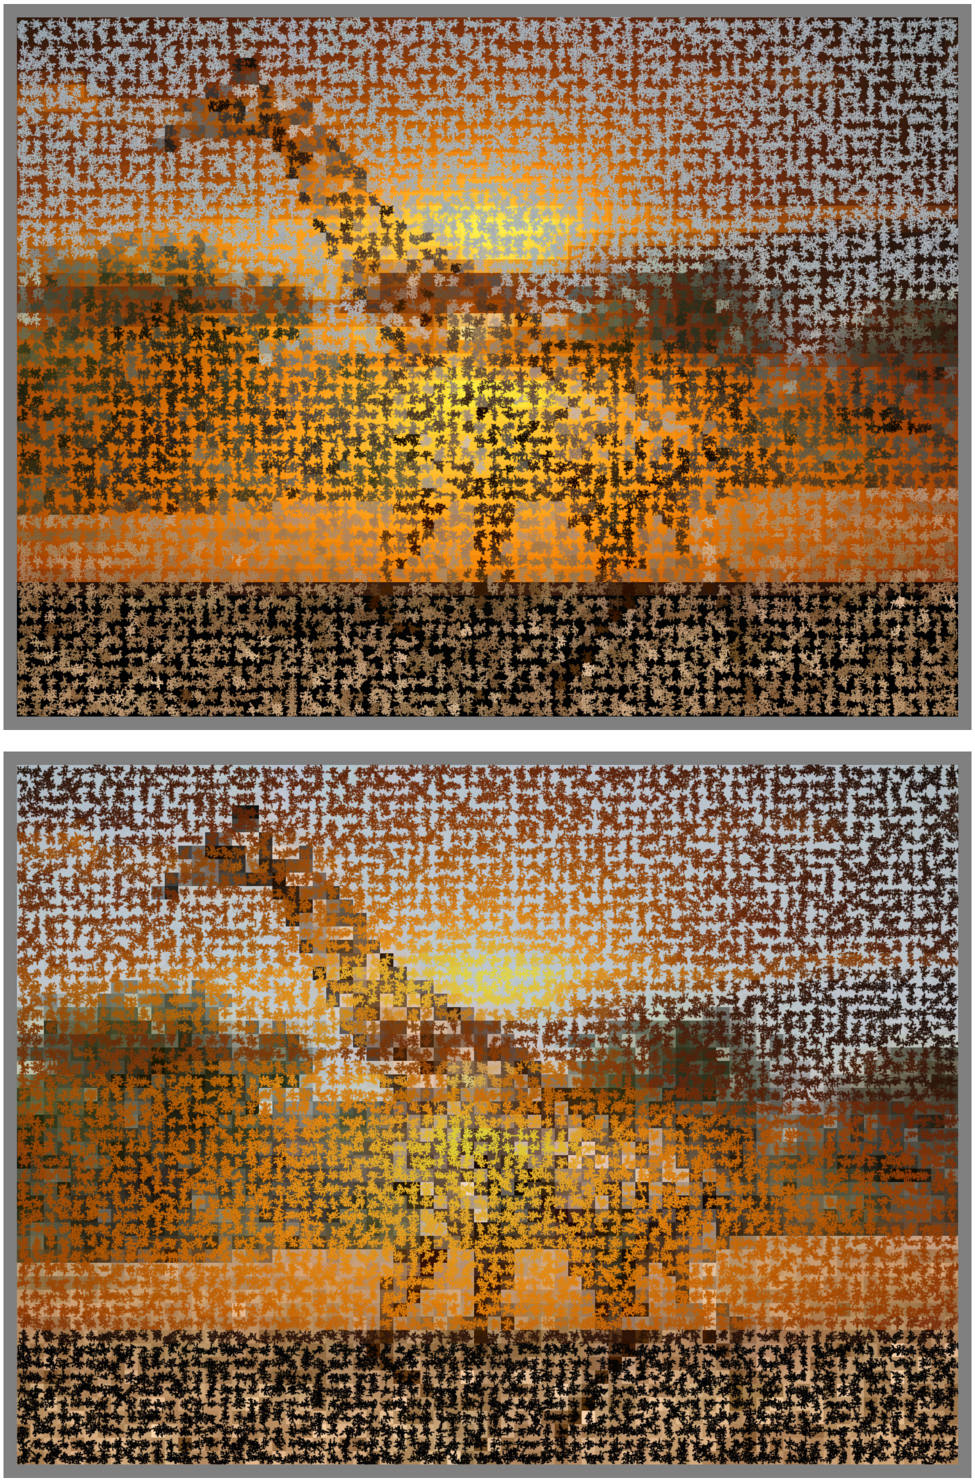

**Figure 3**: Our $72 \times 52$ duotone composite mosaics. Top: Sunset as background, giraffe as foreground. Bottom: Sunset as foreground, giraffe as background.

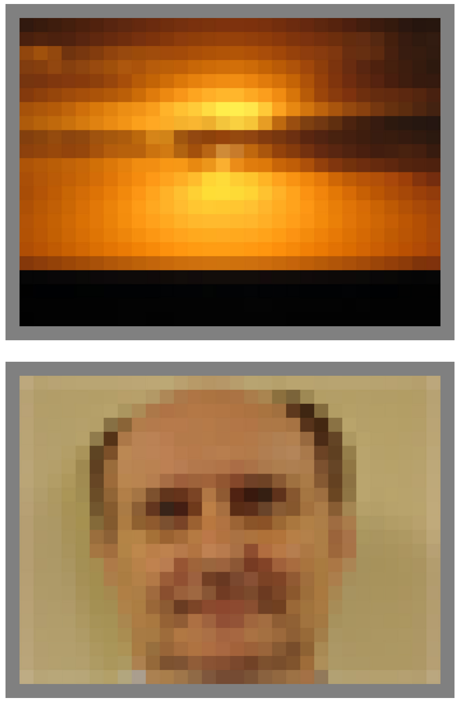

**Figure 4**: Initial pass for the background images yielding $30 \times 22$ mosaics with $41 \times 41$ tiles.

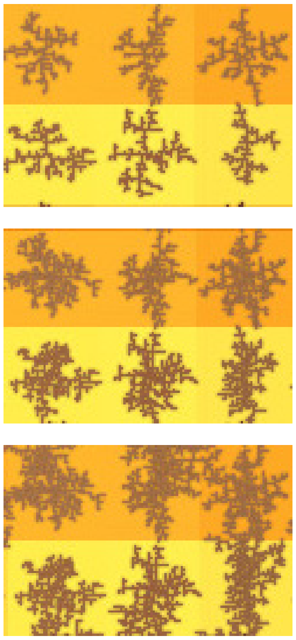

**Figure 5**: An example of the cumulative aggregation of the DLA structures used to construct the foreground image after each of the passes.

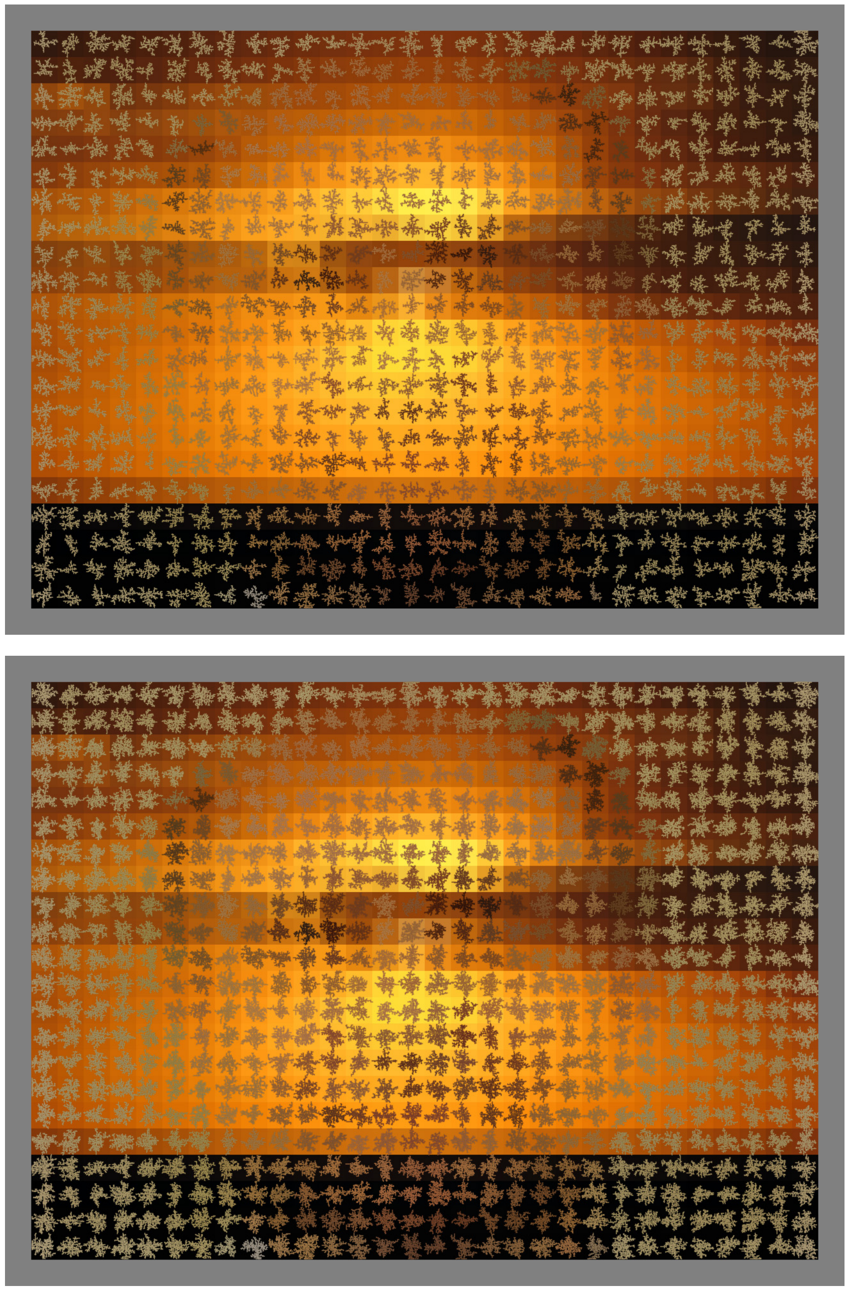

**Figure 6**: The first two passes used to build up the DLA structure for the foreground portrait image over the background sunset image.

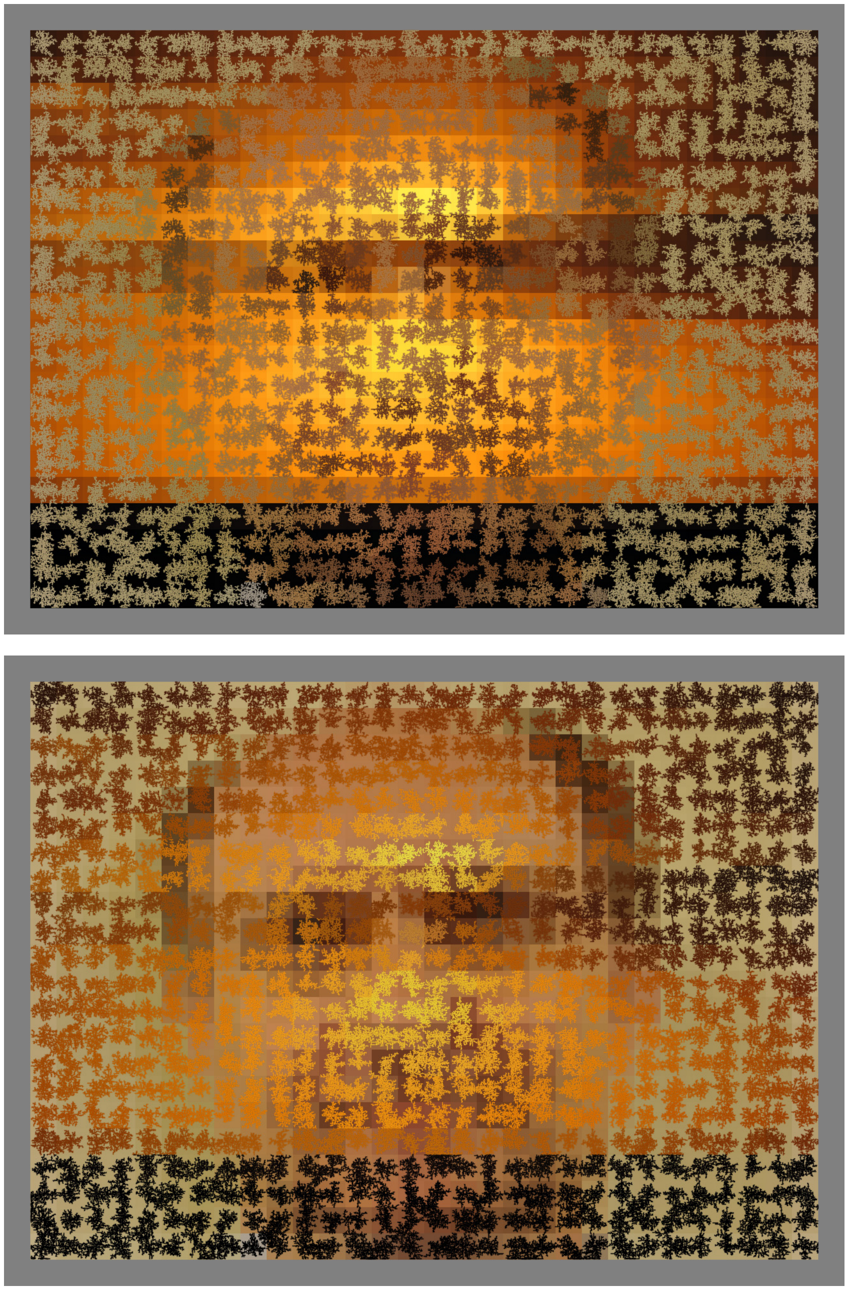

**Figure 7**: Top: Final pass gluing previously aggregated DLA structures of the foreground portrait together over the background sunset while preserving the duotone nature of the $41 \times 41$ enlarged tiles. Bottom: The result obtained by reversing the roles of foreground and background.