

Pixelated Image Abstraction with Integrated User Constraints

Timothy Gerstner^a, Doug DeCarlo^a, Marc Alexa^c, Adam Finkelstein^b, Yotam Gingold^{a,d}, Andrew Nealen^a

^a*Rutgers University*
^b*Princeton University*
^c*TU Berlin*
^d*Columbia University*

Abstract

We present an automatic method that can be used to abstract high resolution images into very low resolution outputs with reduced color palettes in the style of pixel art. Our method simultaneously solves for a mapping of features and a reduced palette needed to construct the output image. The results are an approximation to the results generated by pixel artists. We compare our method against the results of two naive methods common to image manipulation programs, as well as the hand-crafted work of pixel artists. Through a formal user study and interviews with expert pixel artists we show that our results offer an improvement over the naive methods. By integrating a set of manual controls into our algorithm, we give users the ability to add constraints and incorporate their own choices into the iterative process.

Keywords: pixel art, image abstraction, non-photorealistic rendering, image segmentation, color quantization

1. Introduction

We see pixel art every day. Modern day handheld devices such as the iPhone, Android devices and the Nintendo DS regularly utilize pixel art to convey information on compact screens. Companies like Coca-Cola, Honda, Adobe, and Sony use pixel art in their advertisements [1]. It is used to make icons for desktops and avatars for social networks. While pixel art stems from the need to optimize imagery for low resolution displays, it has emerged as a contemporary art form in its own right. For example, it has been featured by Museum of Modern Art, and there are a number of passionate online communities devoted to it. The “Digital Orca” by Douglas Coupland is a popular sight at the Vancouver Convention Center. France recently was struck by a “Post-it War”¹, where people use Post-It notes to create pixel art on their windows, competing with their neighbors across workplaces, small businesses, and homes.




Figure 1: Examples of pixel art. “Alice Blue” and “Kyle Red” by Alice Bartlett. Notice how faces are easily distinguishable even with this limited resolution and palette. The facial features are no longer proportionally accurate, similar to deformation in a caricature.

What makes pixel art both compelling and difficult is the limitations imposed on the medium. With a significantly limited palette and resolution to work with, the task of creating pixel art becomes one of carefully choosing the set of colors and placing each pixel such that the final image best depicts the original subject. This task is particularly difficult as pixel art is typically viewed at a distance

Email address: timgerst@cs.rutgers.edu (Timothy Gerstner)

¹<http://www.postitwar.com/>




Figure 2: Pixel art images simultaneously use very few pixels and a tiny color palette. Attempts to represent image (a) using only 22×32 pixels and 8 colors using (b) nearest-neighbor or (c) cubic downsampling (both followed by median cut color quantization), result in detail loss and blurriness. We optimize over a set of superpixels (d) and an associated color palette to produce output (e) in the style of pixel art.

where the pixel grid is clearly visible, which has been shown to contribute to the perception of the image [2]. As seen in Figure 1, creating pixel art is not a simple mapping process. Features such as the eyes and mouth need to be abstracted and resized in order to be represented in the final image. The end product, which is no longer physically accurate, still gives the impression of an identifiable person.

However, few, if any methods exist to automatically or semi-automatically create effective pixel art. Existing downsampling methods, two of which are shown in Figure 2, do not accurately capture the original subject. Artists often turn to making pieces by hand, pixel-by-pixel, which can take a significant amount of time and requires a certain degree of skill not easily acquired by novices of the art. Automated and semi-automated methods have been proposed for other popular art forms, such as line drawing [3, 4] and painting [5]. Methods such as [6] and [7] not only abstract images, but do so while retaining salient features.

We introduce an entirely automated process that transforms high resolution images into low resolution, small palette outputs in a pixel art style. At the core of our algorithm is a multi-step iterative process that simultaneously solves for a mapping of features and a reduced palette to convert an input image into a pixelated output image. In the first part of each iteration we use a modified version of an image segmentation proposed by Achanta et al. [8] to map regions of the input image to output pixels. In the second step, we utilize an adaptation of mass-constrained deterministic annealing [9] to find an optimal palette and its association to out-

put pixels. These steps are interdependent, and the final solution is an optimization of both the spatial and palette sizes specified by the user. Throughout this process we utilize the perceptually uniform CIELAB color space [10]. The end result serves as an approximation to the process performed by pixel artists (Figure 2, right).

This paper presents an extended edition of Pixelated Image Abstraction [11]. In addition to an expanded results section, we have added a set of user controls to bridge the gap between the manual process of an artist and the automated process of our algorithm. These controls allow the user to provide as much or as little input into the process as desired, to produce a result that leverages both the strengths of our automated algorithm and the knowledge and personal touch of the user.

Aside from assisting a class of artists in this medium, applications for this work include automatic and semi-automatic design of low-resolution imagery in handheld, desktop, and online contexts like Facebook and Flickr, wherever iconic representations of high-resolution imagery are used.

2. Related Work

One aspect of our problem is to reproduce an image as faithfully as possible while constrained to just a few output colors. Color quantization is a classic problem wherein a limited color palette is chosen based on an input image for indexed color displays. A variety of methods were developed in the 1980's and early 1990's prior to the

94 advent of inexpensive 24-bit displays, for exam- 146
 95 ple [12, 13, 14, 15]. A similar problem is that of
 96 selecting a small set of custom inks to be used in
 97 printing an image [16]. These methods rely only 147
 98 on the color histogram of the input image, and are 148
 99 typically coupled to an independent dithering (or 149
 100 halftoning) method for output in a relatively high 150
 101 resolution image. In our problem where the spatial 151
 102 resolution of the output is also highly constrained, 152
 103 we optimize simultaneously the selection and place- 153
 104 ment of colors in the final image. 154

105 The problem of image segmentation has been 155
 106 extensively studied. Proposed solutions include 156
 107 graph-cut techniques, such as the method proposed 157
 108 by Shi and Malik [17], and superpixel-based meth- 158
 109 ods QuickShift [18], Turbopixels [19], and SLIC [8].
 110 In particular, SLIC (Simple Linear Iterative Clus-
 111 tering) produces regular sized and spaced regions 159
 112 with low computational overhead given very few in-
 113 put parameters. These characteristics make SLIC 160
 114 an appropriate starting point for parts of our 161
 115 method. 162

116 Mass-constrained deterministic annealing 163
 117 (MCDA) [9] is a method that uses a probabilistic 164
 118 assignment while clustering. Similar to k-means, it 165
 119 uses a fixed number of clusters, but unlike k-means 166
 120 it is independent of initialization. Also, unlike 167
 121 simulated annealing [20], it does not randomly 168
 122 search the solution space and will converge to the 169
 123 same result every time. We use an adapted version 170
 124 of MCDA for color palette optimization. 171

125 Puzicha et al. [21] proposed a method that 172
 126 reduces the palette of an image and applies half- 173
 127 toning using a model of human visual perception. 174
 128 While their method uses deterministic annealing 175
 129 and the CIELAB space to find a solution that 176
 130 optimizes both color reduction and dithering, our 177
 131 method instead emphasizes palette reduction in 178
 132 parallel with the reduction of the output resolution. 179

133 Kopf and Lischinski [22] proposed a method that 180
 134 extracts vector art representations from pixel art. 181
 135 This problem is almost the inverse of the one pre- 182
 136 sented in this paper. However, while their solution 183
 137 focuses on interpolating unknown information, con- 184
 138 verting an image to pixel art requires compressing 185
 139 known information. 186

140 Finally, we show that with minor modification 187
 141 our algorithm can produce “posterized” images, 188
 142 wherein large regions of constant color are sepa- 189
 143 rated by vectorized boundaries. To our knowledge, 190
 144 little research has addressed this problem, though 191
 145 it shares some aesthetic concerns with the *artistic* 192

thresholding approach of Xu and Kaplan [23].

147 3. Background

148 Our method for making pixel art builds upon two 149
 150 existing techniques, which we briefly describe in this 151
 152 section. 153

154 **SLIC.** Achanta et al. [8] proposed an iterative 155
 156 method to segment an image into regions termed 157
 158 “superpixels.” The algorithm is analogous to k-
 means clustering [24] in a five dimensional space
 (three color and two positional), discussed for ex-
 ample in Forsyth and Ponce [25]. Pixels in the input
 image p_i are assigned to superpixels p_s by minimiz-
 ing

$$159 \quad d(p_i, p_s) = d_c(p_i, p_s) + m\sqrt{\frac{N}{M}}d_p(p_i, p_s) \quad (1)$$

160 where d_c is the color difference, d_p is the posi- 161
 162 tional difference, M is the number of pixels in the 163
 164 input image, N is the number of superpixels, and 165
 166 m is some value in the range $[0, 20]$ that controls 167
 168 the relative weight that color similarity and pixel 169
 170 adjacency have on the solution. The color and po- 171
 172 sitional differences are measured using Euclidean 173
 174 distance (as are all distances in our paper, unless 175
 176 otherwise noted), and the colors are represented in 177
 178 LAB color space. Upon each iteration, superpixels 179
 180 are reassigned to the average color and position of 181
 182 the associated input pixels. 183

184 **Mass Constrained Deterministic Anneal- 185
 186 ing.** MCDA [9] is a global optimization method 187
 188 for clustering that draws upon an analogy with the 189
 190 process of annealing a physical material. We use 191
 192 this method both for determining the colors in our 193
 194 palette, and for assigning one of these palette colors 195
 196 to each pixel—each cluster corresponds to a palette 197
 198 color. 199

200 MCDA is a fuzzy clustering algorithm that prob- 201
 202 abilistically assigns objects to clusters based on 203
 204 their distance from each cluster. It relies on a 205
 206 temperature value T , which can be viewed as pro- 207
 208 portional to the expected variance of the clusters. 209
 210 Initially, T is set to a high value T_0 , which makes 211
 212 each object equally likely to belong to any cluster. 213
 214 Each time the system locally converges T is low- 215
 216 ered (and the variance of each cluster decreases). 217
 218 As this happens, objects begin to prefer favor par- 219
 220 ticular clusters, and as T approaches zero each ob- 221
 222 ject becomes effectively assigned to a single cluster, 223
 224 at which point the final set of clusters is produced. 225

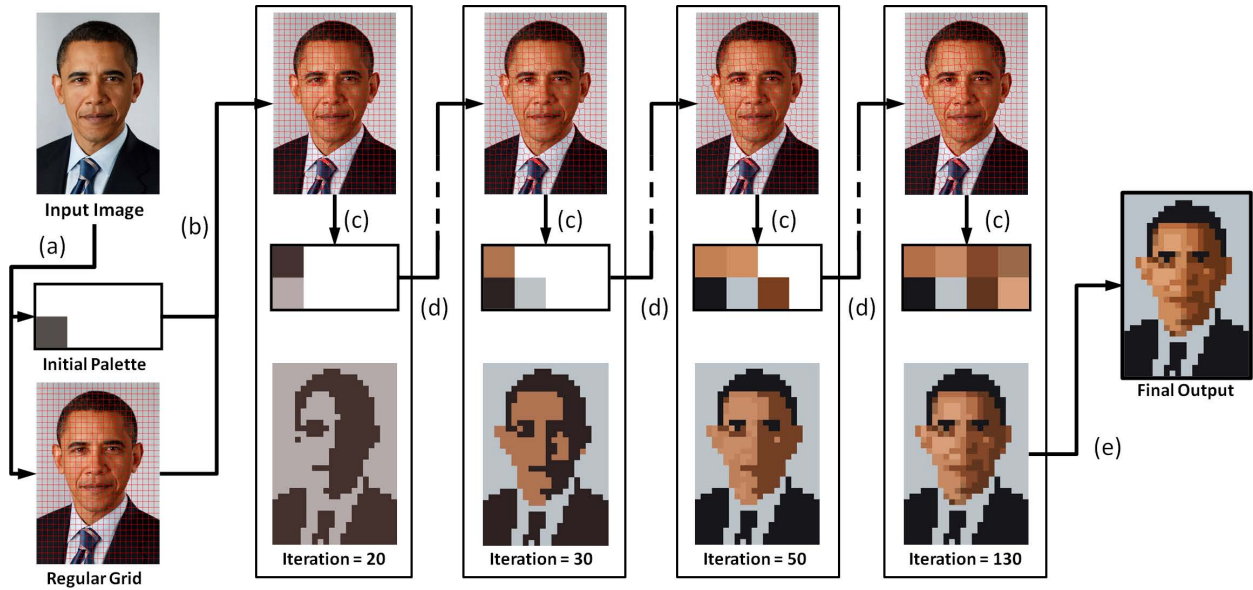


Figure 3: The pipeline of the algorithm. The superpixels (a) are initialized in a regular grid across the input image, and the palette is set to the average color of the M input pixels. The algorithm then begins iterating (b). Each iteration has two main steps: (c) the assignment of input pixels to superpixels, and (d) the assignment of superpixels to colors in the palette and updating the palette. This not only updates each color, but may also add new colors to the palette. After convergence, the palette is saturated (e) producing the final output.

In Section 4.3 we provide a formal definition of the conditional probability we use to assign superpixels to colors in the palette.

Since at high T having multiple clusters is redundant, MCDA begins with a single cluster, represented internally by two sub-clusters. At the beginning of each iteration these sub-clusters are set to slight permutations of their mean. At a high T these clusters converge to the same value after several iterations, but as the temperature is lowered they begin to naturally separate. When this occurs, the cluster is split into two separate clusters (each represented by their own sub-clusters). This continues recursively until the (user specified) maximum number of clusters is reached.

4. Method

Our automated algorithm is an iterative procedure—an example execution is shown in Figure 3. The process begins with an input image of width w_{in} and height h_{in} and produces an output image of width w_{out} and height h_{out} which contains at most K different colors—the palette size. Given the target output dimensions and palette size, each iteration of the algorithm segments the pixels in the input into regions corresponding to

pixels in the output and solves for an optimal palette. Upon convergence, the palette is saturated to produce the final output. In this section, we describe our algorithm in terms of the following:

Input Pixels The set of pixels in the input image, denoted as p_i where $i \in [1, M]$, and $M = w_{\text{in}} \times h_{\text{in}}$.

Output Pixels The set of pixels in the output image, denoted as p_o where $o \in [1, N]$, and $N = w_{\text{out}} \times h_{\text{out}}$.

Superpixel A region of the input image, denoted as p_s where $s \in [1, N]$. The superpixels are a partition of the input image.

Palette A set of K colors c_k , $k \in [1, K]$ in LAB space.

Our algorithm constructs a mapping for each superpixel that relates a region of input pixels with a single pixel in the output, as in Figure 4. The algorithm proceeds similarly to MCDA, with a superpixel refinement and palette association step performed upon each iteration, as summarized in Algorithm 1. Section 5.1 describes how the algorithm can be expanded to allow a user to indicate important regions in the input image.

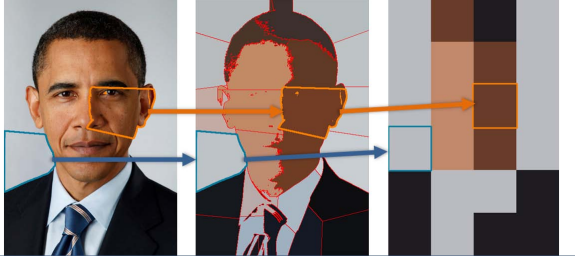


Figure 4: Pixels in the input image (left) are associated with superpixel regions (middle). Each superpixel region corresponds to a single pixel in the output image (right).

Algorithm 1

- ▷ **initialize** superpixels, palette and temperature T (Section 4.1)
 - ▷ **while** ($T > T_f$)
 - ▷ **refine** superpixels with 1 step of modified SLIC (Section 4.2)
 - ▷ **associate** superpixels to colors in the palette (Section 4.3)
 - ▷ **refine** colors in the palette (Section 4.3)
 - ▷ **if** (palette converged)
 - ▷ **reduce** temperature $T = \alpha T$
 - ▷ **expand** palette (Section 4.3)
 - ▷ post-process (Section 4.4)
-

4.1. Initialization

The N superpixel centers are initialized in a regular grid across the input image, and each input pixel is assigned to the nearest superpixel (in (x, y) space, measured to the superpixel center). The palette is initialized to a single color, which is set to the mean value of the M input pixels. All superpixels are assigned this mean color. See Figure 3, step (a).

The temperature T is set to $1.1T_c$, where T_c is the critical temperature of the set of M input pixels, defined as twice the variance along the major principal component axis of the set in LAB space [9]. The T_c of a set of objects assigned to a cluster is the temperature at which a cluster will naturally split. Therefore, this policy ensures that the initial temperature is easily above the temperature at which more than one color in the palette would exist.

4.2. Superpixel refinement

This stage of the algorithm assigns pixels in the input image to superpixels, which correspond to



Figure 5: Our method uses palette colors when finding superpixels. Using the mean color of a superpixel works when the palette is unconstrained (left), but fails when using a constrained palette (middle). This is because the input pixels cluster into superpixels based on colors that do not exist in the final image, which creates a discrepancy. Using the palette colors to represent the superpixels (right) removes this discrepancy.

pixels in the output image—see steps (b) and (d) in Figure 3.

To accomplish this task, we use a single iteration of our modified version of SLIC. In the original SLIC algorithm, upon each iteration, every input pixel is assigned to the superpixel that minimizes $d(p_i, p_s)$, and the color of each superpixel is set to the mean color value of its associated input pixels, m_s . However, in our implementation, the color of each superpixel is set to the palette color that is associated with the superpixel (the construction of this mapping is explained in Section 4.3). This interdependency with the palette forces the superpixels to be optimized with respect to the colors in the palette rather than the colors in the input image. Figure 5 shows the results of using the mean color value instead of our optimized palette used in Figure 2.

However, this also means the color error will be generally higher. As a result, we’ve found that minimizing $d(p_i, p_s)$ using a value of $m = 45$ is more appropriate in this case (Achanta et al. [8] suggest $m = 10$). This increases the weight of the positional distance and results in a segmentation that contains superpixels with relatively uniform size.

Next, we perform two steps, one modifies each superpixel’s (x, y) position for the next iteration, and one changes each superpixel’s representative color. Each step is an additional modification to the original SLIC method and significantly improves the final result.

As seen in Figure 6 (left), SLIC results in superpixel regions which tend to be organized in 6-connected neighborhoods (i.e. a hexagonal grid). This is caused by how the (x, y) position of each



Figure 6: Without the Laplacian smoothing step, the superpixels (left) tend to have 6-connected neighborhoods. This causes small distortions in the output (center), which are particularly noticeable on the ear, eye and mouth, when compared to original output that uses the superpixels that included the smoothing step (right).

superpixel is defined as the average position of the input pixels associated with it. This hexagonal grid does not match the neighborhoods of the output pixels, which are 8-connected (i.e. a rectangular grid) and will give rise to undesirable distortions of image features and structures in the output, as seen in Figure 6(center).

We address this problem with Laplacian smoothing. Each superpixel center is moved a percentage of the distance from its current position to the average position of its 4-connected neighbors (using the neighborhoods at the time of initialization). We use 40%. As seen in Figure 2 (d), this improves the correspondence between the superpixel and output pixel neighborhoods. Specifically, it helps ensure that superpixel regions that are adjacent in the input map are also adjacent pixels in the output. To be clear, it is only in the next iteration when the superpixels will be reassigned based on this new center, due to the interleaved nature of our algorithm.

In our second additional step, the color representatives of the superpixels are smoothed. In the original SLIC algorithm, the representative color for each superpixel is the average color m_s of the input pixels associated with it. However, simply using the mean color can become problematic for continuous regions in the image that contain a color gradient (such as a smooth shadowed surface). While this gradient appears natural in the input image, the region will not appear continuous in the pixelated output.

To remedy this, our algorithm adjusts the values of m_s using a bilateral filter. We construct an image of size $w_{\text{out}} \times h_{\text{out}}$ where each superpixel is assigned the same position as its corresponding

output pixel, with value m_s . The colors that results from bilaterally filtering this image, m_s' are used while iterating the palette.

4.3. Palette refinement

Palette iteration is performed using MCDA [9]. Each iteration of the palette, as seen in step (c) in Figure 3, can be broken down into three basic steps: **associating** superpixels to colors in the palette, **refining** the palette, and **expanding** the palette. The associate and refine steps occur every iteration of our algorithm. When the palette has converged for the current temperature T , the expand step is performed.

It is important to note how we handle the sub-clusters mentioned in Section 3: we treat each sub-cluster as a separate color in the palette, and keep track of the pairs. The color of each c_k is the average color of its two sub-clusters. When the maximum size of the palette is reached (in terms of the number of distinct colors c_k), we eliminate the sub-clusters and represent each color in the palette as a single cluster.

Associate. The MCDA algorithm requires a probability model that states how likely a particular superpixel will be associated with each color in the palette. See Figure 7. The conditional probability $P(c_k|p_s)$ of a superpixel p_s being assigned color c_k depends on the color distance in LAB space and the current temperature, and is given by (after suitable normalization):

$$P(c_k|p_s) \propto P(c_k) e^{-\frac{\|m_s' - c_k\|}{T}} \quad (2)$$

$P(c_k)$ is the probability that color c_k is assigned to any superpixel, given the existing assignment.




Figure 7: Each superpixel (left) is associated by some conditional probability $P(c_k|p_s)$ to each color in the palette (middle). The color with the highest probability is assigned to the superpixel and its associated output pixel in the final image (right).

368 Upon initialization, there is only one color, and 413
 369 thus this value is initialized to 1. As more colors 414
 370 are introduced into the palette, the value of this 415
 371 probability is computed by marginalizing over p_s : 416

$$372 \quad P(c_k) = \sum_{s=1}^N P(c_k|p_s)P(p_s) \quad (3) \quad 417$$

373 For the moment, $P(p_s)$ simply has a uniform distri- 421
 374 bution. This will be revisited in Section 5.1 when 422
 375 incorporating user-specified importance. The val- 423
 376 ues of $P(c_k)$ are updated after the values of $P(c_k|p_s)$ 424
 377 are computed using Equation 2. Each superpixel is 425
 378 assigned to the color in the palette that maximizes 426
 379 $P(c_k|p_s)$. Intermediate results of this assignment 427
 380 can be seen in Figure 3 (bottom row). The expo- 428
 381 nential distribution in Equation 2 tends towards a 429
 382 uniform distribution for large values of T , in which 430
 383 case each superpixel will be evenly associated with 431
 384 every palette color. As T decreases, superpixels fa- 432
 385 vor colors in the palette that are less distant. At 433
 386 the final temperature, the generic situation after 434
 387 convergence has $P(c_k|p_s) = 1$ for a single color 435
 388 in the palette and $P(c_k|p_s) = 0$ for the rest. In 436
 389 this case, deterministic annealing is equivalent to 437
 390 k-means clustering. 438

391 **Refine.** The next step is to refine the palette
 392 by reassigning each color c_k to a weighted average 439
 393 of all superpixel colors, using the probability of 440
 394 association with that color:

$$395 \quad c_k = \frac{\sum_{s=1}^N m_s' P(c_k|p_s) P(p_s)}{P(c_k)} \quad (4) \quad 441$$

396 This adapts the colors in the existing palette given
 397 the revised superpixels. Such changes in the palette
 398 can be seen in Figure 3, as the computation pro- 447
 399 gresses.

400 **Expand.** Expansion only occurs during an 448
 401 iteration if the palette has converged for the current 449
 402 temperature T (convergence is measured by the 450
 403 total change in the palette since last iteration being 451
 404 less than some small value $\epsilon_{\text{palette}}$). First, the 452
 405 temperature is lowered by some factor α (we use 453
 406 0.7). Next, the palette is expanded if the number 454
 407 of colors is less than the number specified by the 455
 408 user. For each c_k we check to see if the color 456
 409 needs to be split into two separate colors in the 457
 410 palette. As per MCDA, each color in the palette 458
 411 is represented by two cluster points c_{k_1} and c_{k_2} . 459
 412 We use $\|c_{k_1} - c_{k_2}\| > \epsilon_{\text{cluster}}$ (where $\epsilon_{\text{cluster}}$ is 460

a sufficiently small number), to check for palette
 separation. If so, the two cluster points are added
 to the palette as separate colors, each with its own
 pair of cluster points. As seen in Figure 3, over the
 course of many iterations, the palette grows from a
 single color to a set of eight (which is the maximum
 number specified by the user in this example).

After resolving any splits, each color is repre-
 sented by two sub-clusters with the same value
 (unless the maximum number of colors have been
 reached). In order for any color’s sub-clusters to
 separate in the following iterations, c_{k_1} and c_{k_2}
 must be made distinctly different. To do so, we
 perturb the sub-clusters of each color by a small
 amount along the principal component axis of the
 cluster in LAB space. Rose [9] has shown this to be
 the direction a cluster will split. This perturbation
 allows the sub-clusters of each color to merge when
 $T > T_c$ and separate when $T < T_c$.

Algorithm 1 is defined so that the superpixel and
 palette refinement steps are iterated until conver-
 gence. The system converges when the tempera-
 ture has reached the final temperature T_f and the
 palette converges. We use $T_f = 1$ to avoid trun-
 cation errors as the exponential component of the
 Equation 2 becomes small.

4.4. Palette Saturation

441 As a post-processing step, we provide the option
 442 to saturate the palette, which is a typical pixel
 443 artist technique, by simply multiplying the a and b
 444 channels of each color by a parameter $\beta > 1$. This
 445 value used in all our results is $\beta = 1.1$. Lastly,
 446 by converting to from LAB to RGB space, our
 algorithm outputs the final image.

5. User Controls

The algorithm described in Section 4 completely
 automates the selection of the color palette. This
 stands in marked contrast to the traditional, man-
 ual process of creating pixel art, where the artist
 carefully selects each color in the palette and its
 placement in the image. Therefore, we propose a
 set of user controls that leverage the results of our
 algorithm and bridges the gap between these two
 extremes. These controls allow the user to have in
 as much or as little control over the process as they
 want. This combines the power and speed of our au-
 tomated method with the knowledge and creativity
 of the user.

461 The first user control, originally proposed in Pix-
462 elated Image Abstraction [11], is an “importance
463 map” that acts as an additional input to our algo-
464 rithm and lets the user emphasize areas of the im-
465 age they believe to be important. The second and
466 third controls we propose, pixel and palette con-
467 straints, are used after the automated algorithm
468 initially converges. Using these two controls, the
469 user can directly edit the palette colors and their
470 assignment in the output image, giving them full
471 control over the result. After each set of edits, the
472 user can choose to have our automated algorithm
473 continue to iterate using the current result as its
474 starting point with the user’s edits as constraints
475 (see Section 5.4). To demonstrate the effectiveness
476 of these user controls, we developed a user interface
477 that was used to generate the results in Figure 16.

478 5.1. Importance Map

479 As stated in Section 4 our automated method
480 does not favor any image content. For instance,
481 nothing is in place that can distinguish between
482 foreground and background objects, or treat them
483 separately in the output. However, user input (or
484 the output of a computer vision system) can easily
485 be incorporated into our algorithm to prioritize the
486 foreground in the output. Thus, our system allows
487 additional input at the beginning of our method.
488 Users can supply a $w_{\text{in}} \times h_{\text{in}}$ grayscale image of
489 weights $W_i \in [0, 1]$, $i \in [1, M]$, used to indicate
490 the importance of each input pixel p_i . In our
491 interface, this is done by using a simple brush to
492 mark areas with the desired weight. We incorporate
493 this map when iterating the palette (Section 4.3) by
494 adjusting the prior $P(p_s)$.

495 Given the importance map, the value $P(p_s)$ for
496 each superpixel is given by the average importance
497 of all input pixels contained in superpixel p_s (and
498 suitable normalization across all superpixels):

$$499 P(p_s) \propto \frac{1}{|p_s|} \sum_{\text{pixel } i \in p_s} W_i \quad (5)$$

500 $P(p_s)$ thus determines how much each superpixel
501 affects the resulting palette, through Equations 3
502 and 4. This results in a palette that can better
503 represent colors in the regions of the input image
504 marked as important.

505 5.2. Pixel Constraints

506 In traditional pixel art, the artist needs to manu-
507 ally choose the color of each pixel in the output.

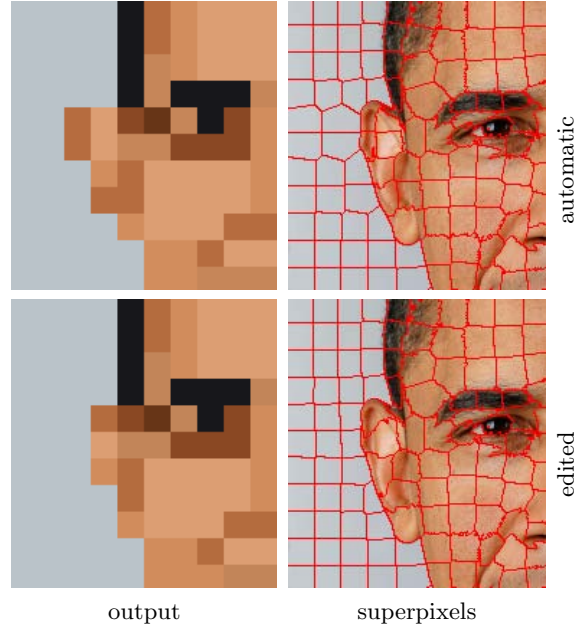


Figure 8: When the user provides constraints to the image, future iterations of the algorithm will update the superpixels in a way that seeks to decrease error under the new constraints. In this example, the original image (top left), is modified by constraining several pixels of the ear to the background color (bottom left). As a result, the superpixels (top right) are redistributed to match the constraints (bottom right). The superpixels that used to be part of the ear now form segments of the background, and neighboring pixels in the output have changed to accommodate the new superpixel distribution.

508 In contrast, our automated algorithm makes the
509 choice entirely for the user. By adding a simple
510 constraint in to our program, we can allow the user
511 to work in the area between these two extremes. For
512 each pixel in the output, we allow the user to select
513 a subset of colors in the palette. For each color not
514 in this subset, we set the conditional probability of
515 that color for this pixel, $P(c_k|p_s)$, to zero.
516 This restricts the color assigned to the output pixel
517 to the color with the highest conditional probability
518 within the subset. Note this has the convenient
519 property of being equivalent to the manual process
520 when the subset is a single color, and to the auto-
521 matic process when the subset is the entire palette.

522 As explained in Section 4.2, superpixels are rep-
523 resented using the color in the palette with the
524 highest conditional probability, $P(c_k|p_s)$. There-
525 fore, adding these constraints will affect the assign-
526 ment of input pixels to superpixels in future iter-
527 ations. As a result, when constraints are added
528 by the user, neighboring superpixels will naturally

529 compensate as the algorithm attempts to decrease 578
530 error under these constraints, as seen in Figure 8. 579

531 In our interface, we implement this tool as a paint 580
532 brush, and allow the user to select one or more 581
533 colors from the palette to form the subset as they 582
534 paint onto the output image. Using the brush, they 583
535 are able to choose the precise color of specific pixels, 584
536 restrict the range of colors for others, and leave the 585
537 rest entirely to our algorithm. 586

538 5.3. Palette Constraints

539 Similarly, in traditional pixel art the artist needs 587
540 to manually choose each color of the palette. We 588
541 again provide a set of constraints to give the user 589
542 control over this process while using our algorithm. 590
543 After the palette has initially converged, the user 591
544 has the option to edit and fix colors in the palette. 592
545 This is done in one of two ways. The first is a trivial 593
546 method; the user directly modifies a specific color 594
547 in the palette. The second utilizes the information 595
548 already gathered by our algorithm. By choosing a 596
549 color in the palette c_k , and then a superpixel p_s 597
550 formed by our algorithm, we set c_k to the mean 598
551 color of that region, m_s , as found in Section 4.2. 599
552 While the first method allows the user to have direct 600
553 control, the second provides them with a way of 601
554 selecting a relatively uniform area of the original 602
555 image from which to sample a color, and without 603
556 having to specify specific values. 604

557 In addition to changing the color, the user has the 605
558 option to keep these colors fixed or free during any 606
559 future iterations of the algorithm. If they are fixed, 607
560 they will remain the same color for the rest of the 608
561 process. If they are not fixed, they will be free to 609
562 converge to a new value as our algorithm iterates, 610
563 starting with the initial color provided by the user’s 611
564 edit. This gives the users another dimension of 612
565 palette control in addition to the ability to manually 613
566 choose the colors. 614

567 Note that when a color is changed in the palette, 615
568 areas of the original image may no longer be well 616
569 represented in the palette. Fortunately, during any 617
570 future iterations, our algorithm will naturally seek 618
571 to reduce this discrepancy by updating the unfixed 619
572 colors in the palette as it attempts to minimize error 620
573 and converge to a new local minimum. 621

574 5.4. Reiterating

575 After using any of the tools described in this sec- 623
576 tion, the user has the option of rerunning our algo- 624
577 rithm. However, rather than starting from scratch, 625

the algorithm begins with the results of the pre-
vious iteration, subject to the constraints specified
by the user. When rerunning the algorithm, the
temperature remains at the final temperature T_f
it reached at convergence, and continues until the
convergence condition described in Section 4.3 is
met again. Note that while iterating, the algo-
rithm maintains the user’s constraints. Therefore
the user can decide what the algorithm can and
cannot update. Also note that since the algorithm
is not starting from scratch, it is generally close to
the next solution, and convergence occurs rapidly
(usually less than a second). After the algorithm
has converged, the user can continue making edits
and rerunning the algorithm until satisfied. In this
way the user becomes a part of the iterative loop,
and both user and algorithm work to create a final
solution.

6. Results

We tested our algorithm on a variety of input
images at various output resolutions and color
palette sizes (Figures 9–16). For each example, we
compare *our method* to two naive approaches:

- *nearest method*: a bilateral filter followed by
median cut color quantization, followed by
nearest neighbor downsampling,
- *cubic method*: cubic downsampling followed
by median cut color quantization. Unless
otherwise stated, the results are generated
using only our automated algorithm, and no
user input was integrated into the result.

All of our results use the parameter settings from
Section 4. Each result was produced in generally
less than a minute on an Intel 2.67Ghz i7 processor
with 4GB memory. Each naive result is saturated
using the same method described in Section 4.4.
Please note it is best to view the results up-close or
zoomed-in, with each pixel being distinctly visible.

In Figure 9, we show the effects of varying
the number of colors in the output palette. Our
automatic method introduces fewer isolated colors
than the nearest method, while looking less washed
out than the cubic method. As the palette size
shrinks, our method is better able to preserve
salient colors, such as the green in the turban.
Our method’s palette assignment also improves the
visibility of the eyes and does not color any of the
face pink.



Figure 9: Varying the palette size (output images are 64×58).




Figure 10: Varying the output resolution (palette has 16 colors).




Figure 11: Examples of very low resolution and small palette sizes.

626 Similar results are seen in Figure 10 when we
 627 vary the output resolution. Again we see that the
 628 cubic method produces washed-out images and the
 629 nearest method has a speckled appearance. At all
 630 resolutions, our method preserves features such as
 631 the goggles more faithfully, and consistently chooses
 632 more accurate skin tones for the faces, whereas both
 633 naive methods choose gray.

634 Using our automated algorithm, the image of Barack Obama is recognizable even at extremely
 635 small output resolutions and palette sizes (Figure 11). At 22×32 and 11×16 , our method more
 636 clearly depicts features such as the eyes while coloring regions such as the hair and tie more consistently.
 637 At 11×16 , the nearest method produces a result that appears to distort facial features, while
 638 the cubic method produces a result that “loses” the eyes. At 6×4 , results are very abstract, but our
 639 method’s output could still be identified as a person or as having originated from the input.
 640

641 In Figure 12, we compare our automated output to manual results created by expert pixel artists.
 642 While our results exhibit the same advantages seen in the previous figures over the naive methods, they
 643 do not match the results made by artists. Expert artists are able to heavily leverage their human
 644 understanding of the scene to emphasize and de-emphasize features and make use of techniques such
 645 as dithering and edge highlighting. While there are many existing methods to automatically dither an
 646 image, at these resolutions the decision on when to apply dithering is nontrivial, and uniform dithering
 647 can introduce undesired textures to surfaces (such as skin).
 648 Figure 13 contains additional results computed using various input images. Overall, our automated
 649 approach is able to produce less noisy, sharper images with a better selection of colors than the
 650 naive techniques we compared against.
 651 To verify our analysis, we conducted a formal user study with 100 subjects using Amazon Mechanical Turk.
 652 Subjects were shown the original image and the results of our automated method and the two naive methods.
 653 The results were scaled to approximately 256 pixels along their longest dimension using nearest neighbor
 654 upsampling, so that users could clearly see the pixel grid. We asked subjects the question, “Which of the
 655 following best represents the image above?” Subjects responded by choosing a result image. The stimulus
 656 sets and answer choices were randomized to remove bias. The study consisted of the example images and
 657 parameters shown in our paper, excluding the results generated by our method.

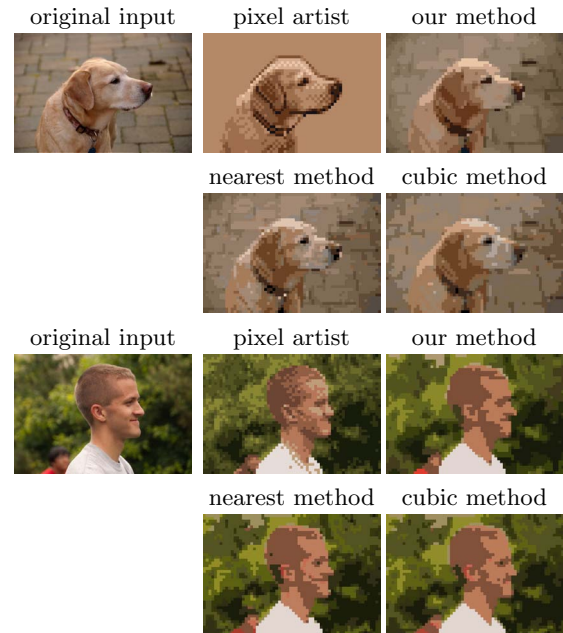


Figure 12: Comparing to the work of expert pixel artists (64×43). The results generated from our method and the naive methods use 16 colors in the first example, 12 in the second. The pixel artists use 8 colors in the first example, 11 in the second.

654 as dithering and edge highlighting. While there are many existing methods to automatically dither an
 655 image, at these resolutions the decision on when to apply dithering is nontrivial, and uniform dithering
 656 can introduce undesired textures to surfaces (such as skin).
 657

658 Figure 13 contains additional results computed using various input images. Overall, our automated
 659 approach is able to produce less noisy, sharper images with a better selection of colors than the
 660 naive techniques we compared against.

661 To verify our analysis, we conducted a formal user study with 100 subjects using Amazon Mechanical Turk.
 662 Subjects were shown the original image and the results of our automated method and the two naive methods.
 663 The results were scaled to approximately 256 pixels along their longest dimension using nearest neighbor
 664 upsampling, so that users could clearly see the pixel grid. We asked subjects the question, “Which of the
 665 following best represents the image above?” Subjects responded by choosing a result image. The stimulus
 666 sets and answer choices were randomized to remove bias. The study consisted of the example images and
 667 parameters shown in our paper, excluding the results generated by our method.




Figure 13: Additional results at various resolution and palette sizes. Columns (left to right): input image, output of our algorithm, output of the nearest method, output of the cubic method.

erated using user input, and each stimulus was duplicated four times (sixty total).

We accounted for users answering randomly by eliminating the results of any subject who gave inconsistent responses (choosing the same answer for less than three of the four duplicates) on more than a third of the stimuli. This reduced the number of valid responses to forty. The final results show that users choose our results 41.49% of the time, the nearest method 34.52% of the time, and the cubic method 23.99% of the time. Using a one-way analysis of variance (ANOVA) on the results, we found a p value of 2.12×10^{-6} , which leads us to reject the null hypothesis that subjects all chose randomly. Using Tukey’s range test we found that our automated method is significantly different from the nearest method with a 91% confidence interval, and from the cubic method with a 99% confidence interval. While we acknowledge that the question asked is difficult one given that it is an aesthetic judgment, we believe the results of this study still show subjects prefer the results of our method over the results of either naive method.

We also received feedback from three expert pixel artists on our automated method; each concluded that the automated results are, in general, an im-

provement over the naive approaches. Ted Martens, creator of the Pixel Fireplace, said that our algorithm “chooses better colors for the palette, groups them well, and finds shapes better.” Adam Saltzman, creator of Canabalt and Flixel, characterized our results as “more uniform, more reasonable palette, better forms, more readable.” Craig Adams, art director of Superbrothers: Sword & Sworcery EP, observed that “essential features seem to survive a bit better [and] shapes seem to come through a bit more coherently. I think the snowboarder’s goggles are the clearest example of an essential shape—the white rim of the goggle—being coherently preserved in your process, while it decays in the ‘naive’ process.”




Figure 14: The original pixel art image (© Nintendo Co., Ltd.) is converted to a vectorized version using Kopf and Lischinski’s method [22]. The vectorized version is then converted back to a pixelated version using our automated method and the two naive methods.

720 In Section 2, we mentioned that the method of
 721 Kopf and Lischinski [22] is essentially the inverse
 722 process of our method; it takes a pixel art piece and
 723 converts it to a smooth, vectorized output. To see
 724 how well our method actually serves as the inverse
 725 process, we took the vectorized output of their
 726 method as the input of our automated algorithm,
 727 setting the size of output image and palette to the
 728 same as their input. The results, compared to
 729 those of the naive methods, are shown in Figure 14.
 730 Visually our method appears to outperform either
 731 naive method, and obtains a result that is similar
 732 to their original input. To quantify the effectiveness
 733 of our method, we took the sum of the Euclidean
 734 distance in LAB space of every pixel between our
 735 output and their input. We did the same for
 736 the naive methods. We found the total error for
 737 our method, the nearest method and the cubic
 738 method to be 1.63×10^3 , 3.05×10^3 and 9.84×10^3 ,
 739 respectively. In other words, our method has 47%
 740 less error than the nearest method, and 83.5% less
 741 error than the cubic method.

742 In Figure 15, we present results from our method
 743 using an the importance map as an additional in-
 744 put to the algorithm. The results are closer to those
 745 created by expert pixel artist. Figure 15(left) al-
 746 locates more colors in the palette to the face of the
 747 person, similar to the manual result in Figure 12.
 748 Figure 15(right) also shows an improvement over
 749 the non-weighted results in Figure 9. For both ex-
 750 amples, the importance map emphasizes the face
 751 and de-emphasizes the background; consequently,
 752 more colors are allocated to the face in each exam-
 753 ple at the expense of the background.

754 In Figure 16, we demonstrate the advantage of
 755 allowing the user to also place pixel and palette
 756 constraints during our iterative process. In Fig-
 757 ure 16(top), the user provides minimal, but effective
 758 changes, such as improving the jawline, and remov-
 759 ing a skin color in favor of a blue in the palette for
 760 the tie. They also introduce a simple striped pat-
 761 tern into the tie, which still represents the original
 762 image, but no longer has a direct correspondence,
 763 and would not be achievable by our algorithm alone.
 764 These changes took less than a minute to make.

765 The improved result in Figure 16(middle row) is
 766 achieved by interleaving multiple steps of user con-
 767 straints and iterations of our algorithm. The user
 768 is also able to incorporate the advanced techniques
 769 observed in Figure 12 such as dithering and edge
 770 highlighting, which are not natively built into our
 771 algorithm.



772 Figure 15: Results using an importance map. (top) 64×58 ,
 773 16 colors (bottom) 64×43 , 12 colors

774 The image in Figure 16(bottom) is a failure case
 775 for our automated algorithm, due to the lighting
 776 and high variation in the background. However,
 777 even with this initially poor output, by interleaving
 778 the iterative process with user constraints (such as
 779 restricting the background to a single color) the
 780 results are significantly improved.

781 Finally, while not the direct goal of our work,
 782 we briefly mention a secondary application of our
 783 method, image *posterization*. This artistic tech-
 784 nique uses just a few colors (originally motivated
 785 by the use of custom inks in printing) and typi-
 786 cally seeks a vectorized output. Adobe Illustrator
 787 provides a feature called LiveTrace that can poster-
 788 ize the image in Figure 2(a), yielding Figure 17(a)
 789 with only 6 colors. To our knowledge, little research
 790 has addressed this problem, though it shares some
 791 aesthetic concerns with the *artistic thresholding* ap-
 792 proach of Xu and Kaplan [23]. A simple modifica-
 793 tion to our optimization that omits the smoothing
 794 step (Figure 6-left) and then colors the original im-
 795 age via associated superpixels gives us Figure 17(b),
 796 which makes a more effective starting point for vec-
 797 torization. The resulting Figure 17(c) offers im-
 798 proved spatial and color fidelity, based on a good
 799 faith effort to produce a similar style in Illustrator.

798 7. Conclusion, Limitations and Future work

799 We present a multi-step iterative process that si-
 800 multaneously solves for a mapping of features and a
 801 reduced palette to convert an input image to a pixe-
 802 lated output image. Our method demonstrates sev-



Figure 16: The results of the automatic method compared to the results obtained by integrating user input into the iterative process with our interface. Note that the user can choose to make only a few key edits (top), or they can leverage their understanding of the image to drastically improve images that are otherwise difficult for the automated algorithm (bottom).

803 eral advantages over the naive methods. Our results
 804 have a more vibrant palette, retain more features
 805 of the original image, and produce a cleaner output
 806 with fewer artifacts. While the naive methods pro-
 807 duce unidentifiable results at very low resolutions
 808 and palette sizes, our approach is still able to cre-
 809 ate iconic images that conjure the original image.
 810 Thus our method makes a significant step towards
 811 the quality of images produced by pixel artists.

812 Nevertheless, our method has several limitations
 813 which we view as potential avenues for future
 814 research. While pixel artists view the results of our
 815 automated algorithm as an improvement, they also
 816 express the desire to have a greater control over the
 817 final product.

818 To address these concerns, we implemented sev-
 819 eral controls that allow the user to give as much or
 820 little feedback into the automated process as they
 821 desire. By incorporating an importance map we
 822 give the user the ability to guide the palette selec-
 823 tion, and by giving the user the ability to provide
 824 pixel and palette constraints and interleave them




Figure 17: (a) vectorized photo posterized with Illustrator (6 colors). (b) Optimization without Laplacian smoothing, coloring associated input pixels (6 colors). (c) Vectorizing b in Illustrator yields similar style with better spatial and color fidelity than a .

825 with our algorithm, we remove the gap between the
 826 manual and automated methods of producing pixel
 827 art.

828 The results of combining these user constraints
 829 into our iterative algorithm are encouraging. For
 830 future work, we wish to expand on our proposed
 831 method and user controls to increase the interaction
 832 between the automated algorithm and the user.
 833 Our goal is to create a complete system that
 834 incorporates the speed and power of an automated
 835 method to assist artists in their entire process,
 836 without restricting the control of the artists over
 837 the final result.

838 As such, the next step is to explore how the
 839 user’s feedback can help inform more advanced
 840 pixel art techniques in our algorithm, such as those
 841 that would produce edge highlighting and dither-
 842 ing. We’d also like to look into ways of automati-
 843 cally performing palette transfers, which would al-
 844 low potential applications of this work to include,
 845 for example, reproduction of an image in repeat-
 846 ing tiles like Lego, or design for architectural fa-
 847 cades composed of particular building materials like
 848 known shades of brick. Currently, our algorithm
 849 is limited to working with colors that are similar
 850 to the original image due to the nature of how we
 851 minimize error, and such an application is not pos-
 852 sible without the user applying a large number of
 853 constraints.

854 Acknowledgments

855 We thank the anonymous reviewers for their help-
 856 ful feedback. We also wish to thank pixel artists
 857 Craig Adams, Ted Martens, and Adam Saltsman
 858 for their advice and comments. This research is

859 supported in part by the Sloan Foundation, the 918
 860 NSF (CAREER Award CCF-06-43268 and grants 919
 861 IIS-09-16129, IIS-10-48948, IIS-11-17257, CMMI- 920
 862 11-29917, IIS-09-16845, DGE-05-49115), and gen- 921
 863 erous gifts from Adobe, Autodesk, Intel, mental 922
 864 images, Microsoft, NVIDIA, Side Effects Software, 923
 865 and the Walt Disney Company. The following copy- 924
 866 righted images are used with permission: Figure 1 925
 867 by Alice Bartlett, Figure 9 by Louis Vest, Figure 13 926
 868 (giraffe) by Paul Adams, and Figure 13 (pagoda) 927
 869 by William Warby. The pixel art in Figure 12 is 928
 870 copyright Adam Saltsman (top) and Ted Martens 929
 871 (bottom). 930

872 References

- 873 [1] Vermehr K, Sauerteig S, Smital S. eboy. 938
 874 <http://hello.eboy.com>; 2012. 939
- 875 [2] Marr D, Hildreth E. Theory of edge detection. Inter- 940
 876 national Journal of Computer Vision 1980;. 941
- 877 [3] DeCarlo D, Finkelstein A, Rusinkiewicz S, Santella A. 942
 878 Suggestive contours for conveying shape. ACM Trans 943
 879 Graph 2003;22(3):848–55. 944
- 880 [4] Judd T, Durand F, Adelson EH. Apparent ridges for 945
 881 line drawing. ACM Trans Graph 2007;26(3):19–. 946
- 882 [5] Gooch B, Coombe G, Shirley P. Artistic vi- 947
 883 sion: painterly rendering using computer vision tech- 948
 884 niques. In: Non-Photorealistic Animation and Ren- 949
 885 dering (NPAR). ISBN 1-58113-494-0; 2002, p. 83–90. 950
 886 doi:<http://doi.acm.org/10.1145/508530.508545>. URL 951
 887 <http://doi.acm.org/10.1145/508530.508545>. 952
- 888 [6] DeCarlo D, Santella A. Stylization and abstraction 953
 889 of photographs. ACM Trans Graph 2002;21:769–76. 954
 890 doi:<http://doi.acm.org/10.1145/566654.566650>. URL 955
 891 <http://doi.acm.org/10.1145/566654.566650>. 956
- 892 [7] Winnemöller H, Olsen SC, Gooch B. Real-time video 957
 893 abstraction. ACM Trans Graph 2006;25:1221–6. 958
- 894 [8] Achanta R, Shaji A, Smith K, Lucchi A, Fua P, 959
 895 Süssstrunk S. SLIC Superpixels. Tech. Rep.; IVRG 960
 896 CVLAB; 2010. 961
- 897 [9] Rose K. Deterministic annealing for clustering, 962
 898 compression, classification, regression, and related 963
 899 optimization problems. Proceedings of the IEEE 964
 900 1998;86(11):2210–39. doi:10.1109/5.726788. 965
- 901 [10] Sharma G, Trussell HJ. Digital color imaging. IEEE 966
 902 Transactions on Image Processing 1997;6:901–32. 967
- 903 [11] Gerstner T, DeCarlo D, Alexa M, Finkelstein A, 968
 904 Gingold Y, Nealen A. Pixelated image abstraction. 969
 905 In: Proceedings of the International Symposium on 970
 906 Non-Photorealistic Animation and Rendering (NPAR). 971
 907 2012;. 972
- 908 [12] Gervautz M, Purgathofer W. Graphics gems. chap. A 973
 909 simple method for color quantization: octree quantiza- 974
 910 tion. ISBN 0-12-286169-5; 1990, p. 287–93. 975
- 911 [13] Heckbert P. Color image quantization for frame buffer 976
 912 display. SIGGRAPH Comput Graph 1982;16:297–307. 977
- 913 [14] Orchard M, Bouman C. Color quantization of images. 978
 914 IEEE Trans on Signal Processing 1991;39:2677–90. 979
- 915 [15] Wu X. Color quantization by dynamic programming 980
 916 and principal analysis. ACM Trans Graph 1992;11:348– 981
 917 72. 982

- [16] Stollnitz EJ, Ostromoukhov V, Salesin DH. Reproduc- 983
 ing color images using custom inks. In: Proceedings of 984
 SIGGRAPH. ISBN 0-89791-999-8; 1998, p. 267–74. 985
- [17] Shi J, Malik J. Normalized cuts and image segmenta- 986
 tion. IEEE Transactions on Pattern Analysis and Ma- 987
 chine Intelligence 1997;22:888–905. 988
- [18] Vedaldi A, Soatto S. Quick shift and kernel methods 989
 for mode seeking. In: In European Conference on 990
 Computer Vision, volume IV. 2008, p. 705–18. 991
- [19] Levinshtein A, Stere A, Kutulakos KN, Fleet DJ, 992
 Dickinson SJ, Siddiqi K. Turbopixels: Fast superpixels 993
 using geometric flows. 2009. 994
- [20] Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by 995
 simulated annealing. Science 1983;220:671–80. 996
- [21] Puzicha J, Held M, Ketterer J, Buhmann JM, Fellner 997
 DW. On spatial quantization of color images. IEEE 998
 Transactions on Image Processing 2000;9:666–82. 999
- [22] Kopf J, Lischinski D. Depixelizing pixel art. ACM Trans 1000
 Graph 2011;30(4):99–. 1001
- [23] Xu J, Kaplan CS, Mi X. Computer-generated paper- 1002
 cutting. In: Proceedings of Pacific Graphics. 2007, p. 1003
 343–50. 1004
- [24] MacQueen JB. Some methods for classification and 1005
 analysis of multivariate observations. In: Proceedings 1006
 of 5th Berkeley Symposium on Mathematical Statistics 1007
 and Probability. 1967, p. 281–97. 1008
- [25] Forsyth DA, Ponce J. Computer Vision: A Modern 1009
 Approach. Prentice Hall; 2002. 1010