Structure-aware bottle cap art

Leonardo Sacht

UFSC

SIBGRAPI 2022 / C&G paper

Motivation: plastic bottle caps

Image source: diveagainstdebris.org

- Long time to degenerate.

- Small to be mistaken for food by sea animals.

Motivation: plastic bottle caps

Image source: diveagainstdebris.org

- Long time to degenerate.
- Small to be mistaken for food by sea animals.
- Top 5 ocean trash items that are deadly for sea life (North Sea Foundation, 2017).

Motivation: low resolution images

Source: pngitem.com

Source: la-boite-a-trouvailles.com

Source: flickr.com

Contributions

- 1191
- 580
- 414
- 360
- 337
- 234
- 195
- 110
- 90
- 73

Caps

Input image

Source: whitehouse.gov

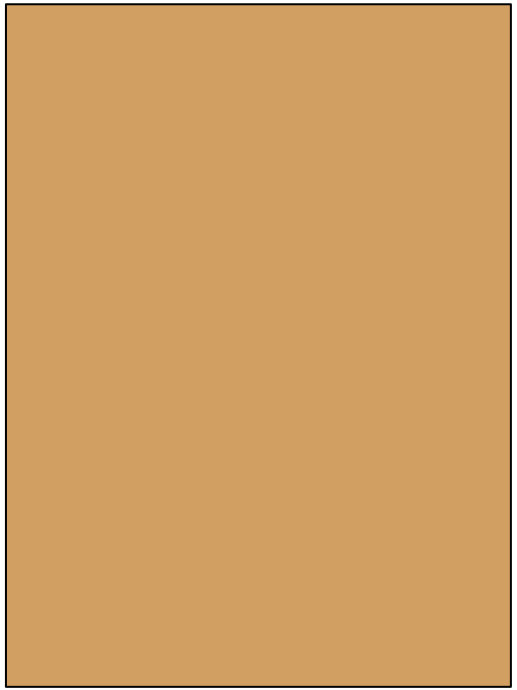

Canvas

New image processing problem

Contributions

- 1191
- 580
- 414
- 360
- 337
- 234
- 195
- 110
- 90
- 73

Caps

Input image

Source: whitehouse.gov

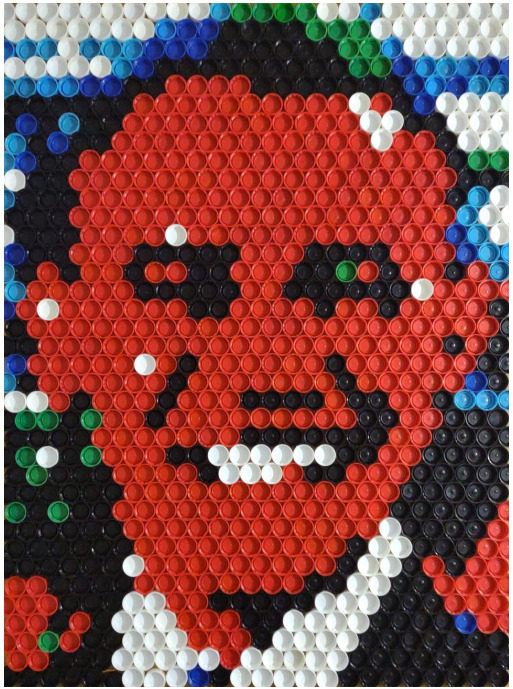

Output bottle cap art

New image processing problem

Optimization: simulated annealing strategy

Related work

Color quantization
[Xiang 1997]

Halftoning
[Velho and Gomes 1991]

Related work

Color quantization
[Xiang 1997]

Halftoning
[Velho and Gomes 1991]

Pixel art
[Gerstner et al. 2013]

Problem setup

Input image

580
414
195
360
1191
234
337
110
90
73

Caps

Canvas

Problem setup

Input image

580
414
195
360
1191
234
337
110
90
73

Caps

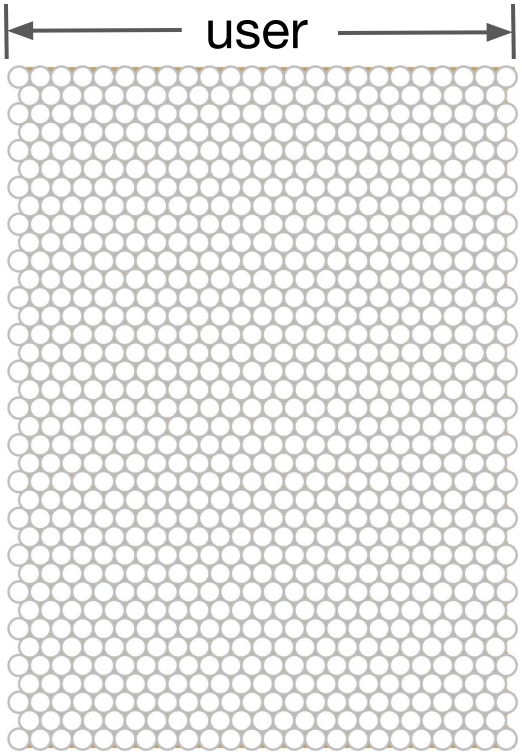

Bottle cap grid

Problem setup

Input image

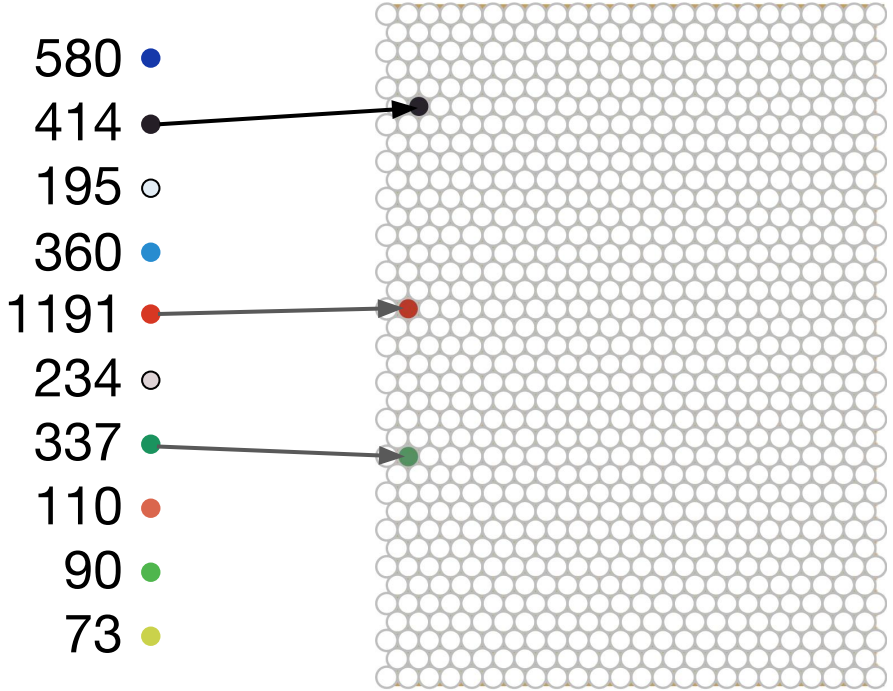

Caps

Bottle cap grid

Problem setup

Input image

580 ●
414 ●
195 ○
360 ●
1191 ●
234 ○
337 ●
110 ●
90 ●
73 ●

Caps

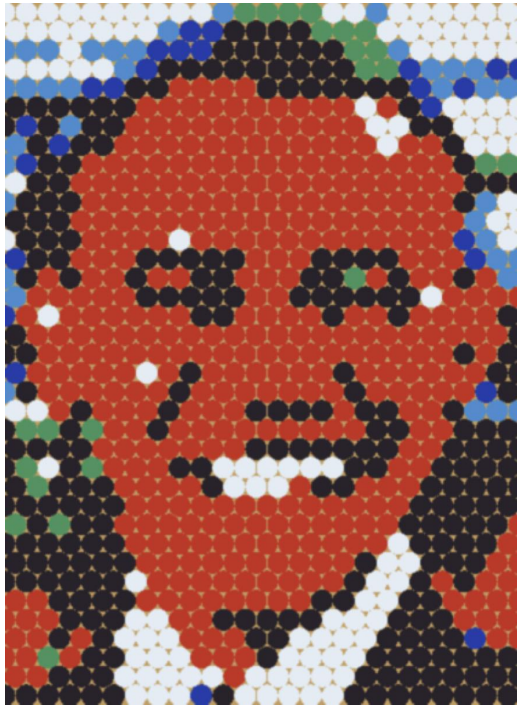

Computer-generated bottle cap art

Problem setup

Input image

580
414
195
360
1191
234
337
110
90
73

Caps

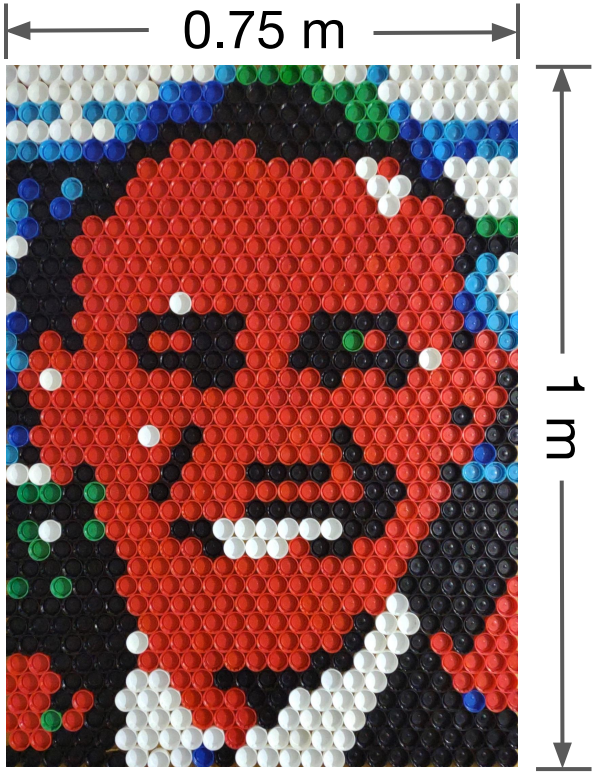

Physically assembled bottle cap art

Naive solution

Input image

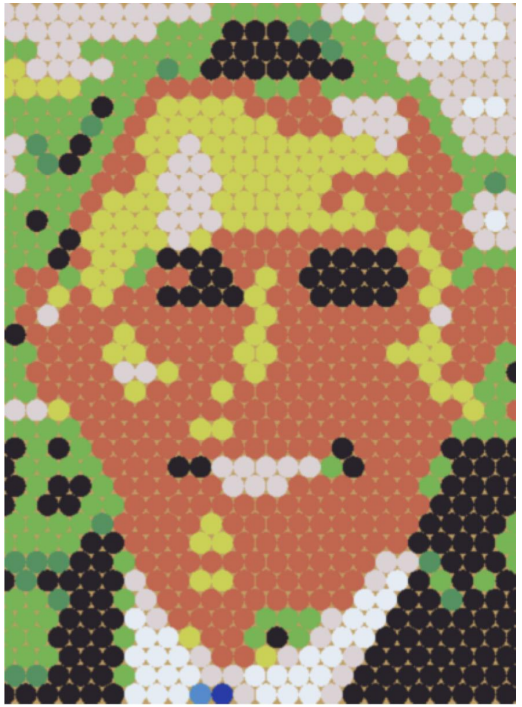

Naive solution

Spatial discretization

Full resolution input

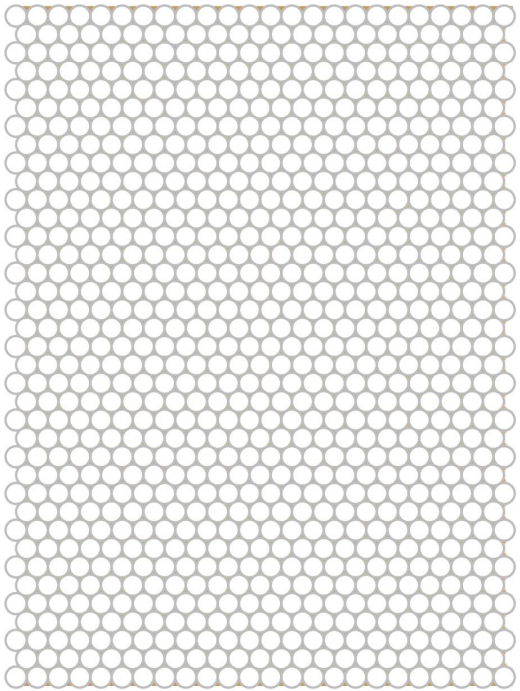

Bottle cap grid

Spatial discretization

Full resolution input

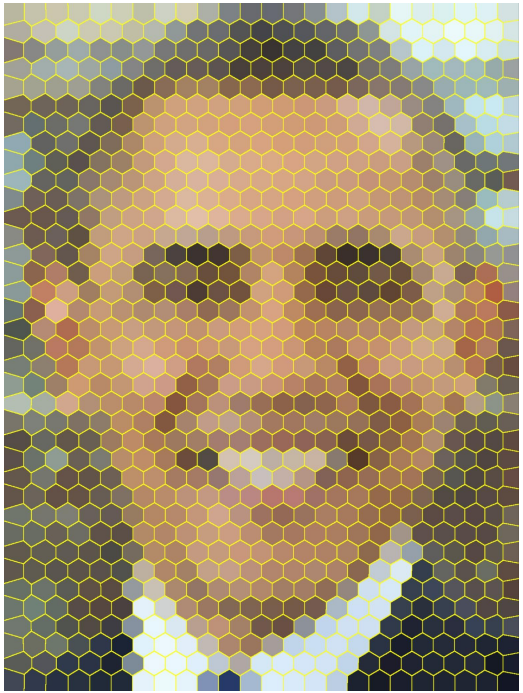

Voronoi diagram

Discretization

Spatial discretization

Full resolution input

SLIC with alpha=50
[Achanta et al. 2012]

Discretization

Spatial discretization

Full resolution input

SLIC with alpha=5
[Achanta et al. 2012]

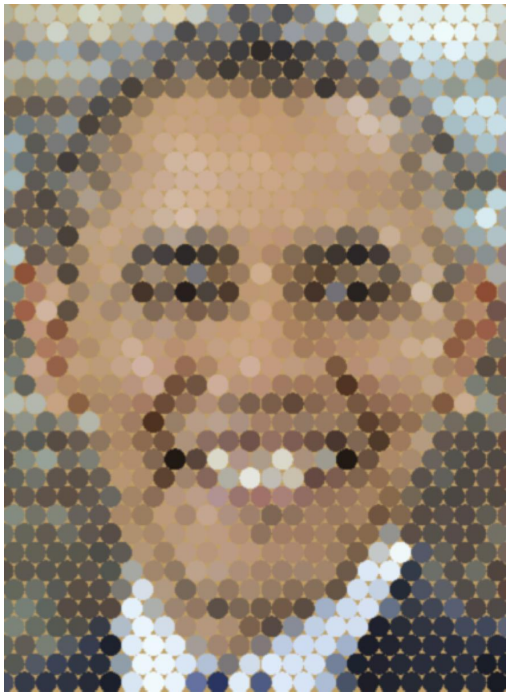

Discretization

Spatial discretization

Full resolution input

SLIC with alpha=50
[Achanta et al. 2012]

Reference image

Structural similarity

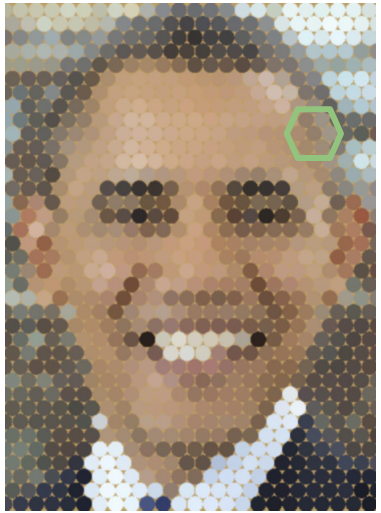

Reference **X**

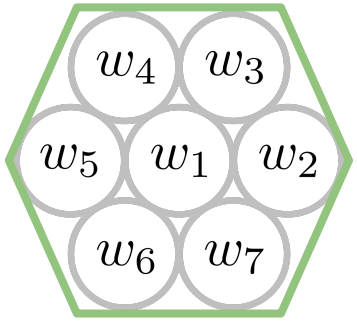

Window and weights

Local color mean

$$\mu_x = \sum_{i=1}^{7} w_i \mathbf{x}_i$$

Local color deviation

$$\sigma_x = \left(\sum_{i=1}^{7} w_i (\mathbf{x}_i - \mu_x) \right)^{\frac{1}{2}}$$

Structural similarity

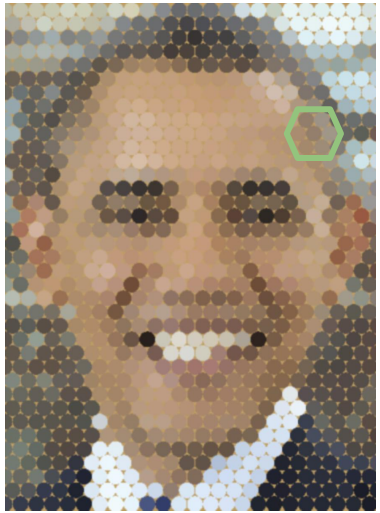

Reference **X**

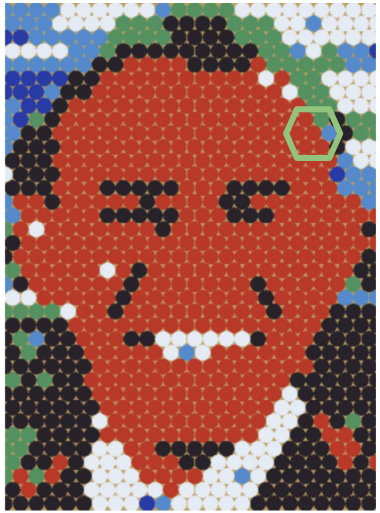

State **Y**

$$\mu_x = \sum_{i=1}^{7} w_i \mathbf{x}_i \qquad \sigma_x = \left(\sum_{i=1}^{7} w_i (\mathbf{x}_i - \mu_x) \right)^{\frac{1}{2}}$$

$$\mu_y = \sum_{i=1}^{7} w_i \mathbf{y}_i \qquad \sigma_y = \left(\sum_{i=1}^{7} w_i (\mathbf{y}_i - \mu_y) \right)^{\frac{1}{2}}$$

$$\sigma_{xy} = \sum_{i=1}^{7} w_i (\mathbf{x}_i - \mu_x)(\mathbf{y}_i - \mu_y)$$

$\text{SSIM} \approx 0$ 👎

$\text{SSIM} \approx 1$ 👍

$$\text{SSIM} = \frac{(2\mu_x\mu_y + \gamma_1)}{(\mu_x^2 + \mu_y^2 + \gamma_1)} \cdot \frac{(2\sigma_{xy} + \gamma_2)}{(\sigma_x^2 + \sigma_y^2 + \gamma_2)}$$

Structural similarity

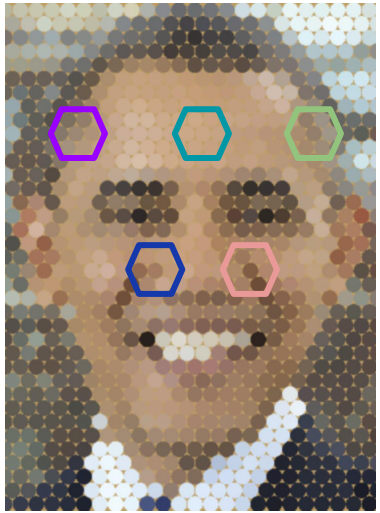

Reference **X**

State **Y**

$\mathrm{SSIM} \approx 0$

$\mathrm{SSIM} \approx 1$

Loop over all windows

$$\mathrm{MSSIM}(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{j=1}^{M} \mathrm{SSIM}(\mathbf{X}_j, \mathbf{Y}_j)$$

Structural similarity

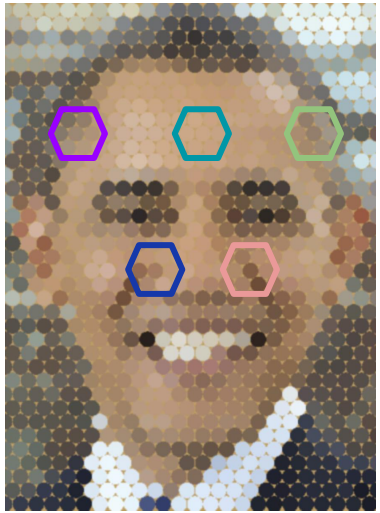

Reference **X**

State **Y**

Our final energy

$$E(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{j=1}^{M} \mathrm{SSIM}(\mathbf{X}_j, \mathbf{Y}_j) \cdot \mathrm{Pen}(j)$$

has a penalty term to encourage constant regions in the reference to be constant in the result.

Optimization

Reference **X**

- 1191
- 580
- 414
- 360
- 337
- 234
- 195
- 110
- 90
- 73

Caps

- Caps given by user

$$E(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{j=1}^{M} \mathrm{SSIM}(\mathbf{X}_j, \mathbf{Y}_j) \cdot \mathrm{Pen}(j)$$

Optimization

Reference **X**

- **1191**
- **580**
- **414**
- **360**
- **337**
- 234
- **195**
- 110
- 90
- 73

Caps

- Caps given by user
- Some cap colors are discarded as a preprocess

$$E(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{j=1}^{M} \mathrm{SSIM}(\mathbf{X}_j, \mathbf{Y}_j) \cdot \mathrm{Pen}(j)$$

Optimization

Reference **X**

- 1191
- 580
- 414
- 360
- 337
- 234
- 195
- 110
- 90
- 73

Caps

Initial solution $\mathbf{Y}_0$

- Caps given by user
- Some cap colors are discarded as a preprocess
- Naive initial solution

$$E(\mathbf{X}, \mathbf{Y}) = \frac{1}{M} \sum_{j=1}^{M} \mathrm{SSIM}(\mathbf{X}_j, \mathbf{Y}_j) \cdot \mathrm{Pen}(j)$$

Optimization

Reference **X**

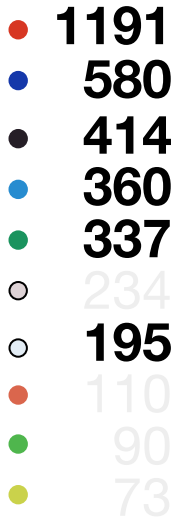

Caps

Initial solution $\mathbf{Y}_0$

Intermediate **Y**

- Simulated annealing strategy
- Update neighborhood with a single color

Optimization

Reference **X**

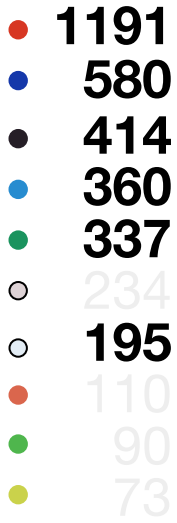

Caps

Initial solution $\mathbf{Y}_0$

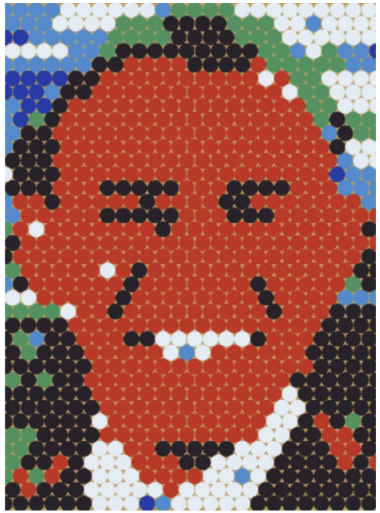

Intermediate **Y**

- Simulated annealing strategy
- Update neighborhood with a single color

Optimization

Reference **X**

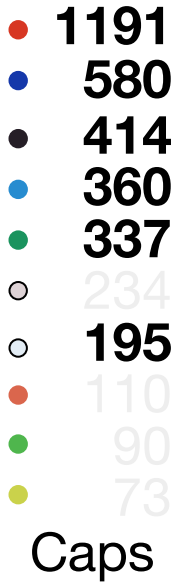

Caps

Initial solution $\mathbf{Y}_0$

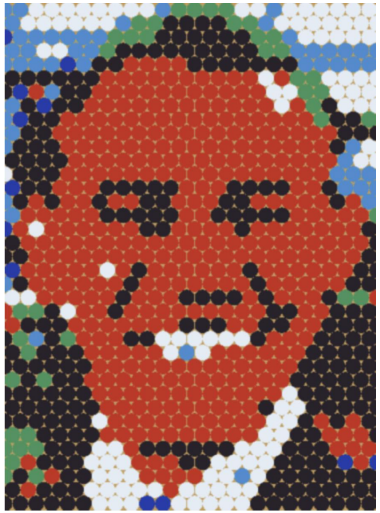

Intermediate **Y**

Optimization

Reference **X**

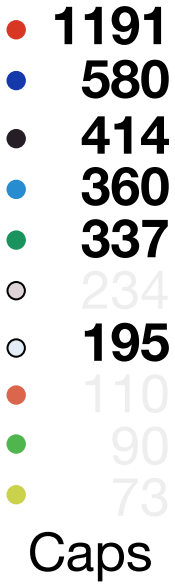

Caps

Initial solution $\mathbf{Y}_0$

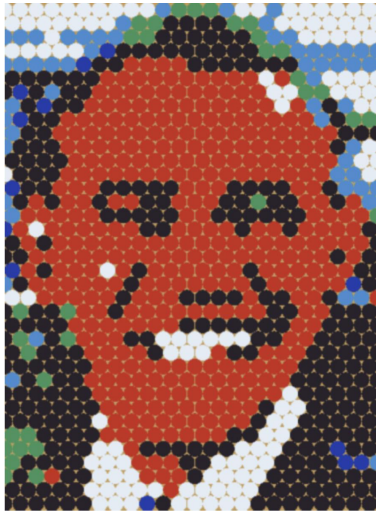

Final result

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

Girl with a Pearl Earring by Johannes Vermeer

Results

1191 580 414 360 337 234 195 110 90 73

Input image

Girl with a Pearl Earring by Johannes Vermeer

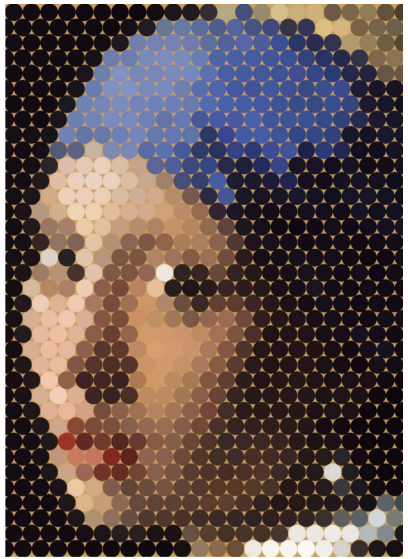

Spatial discretization

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

Girl with a Pearl Earring by Johannes Vermeer

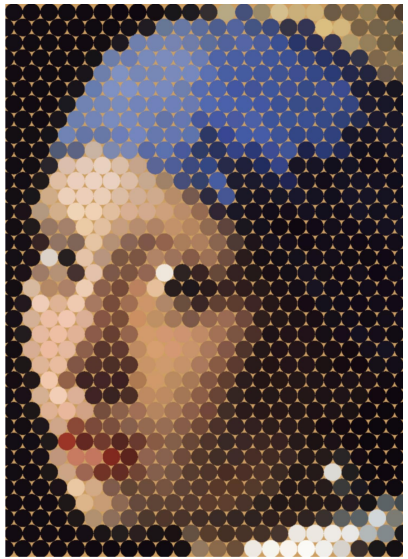

Spatial discretization

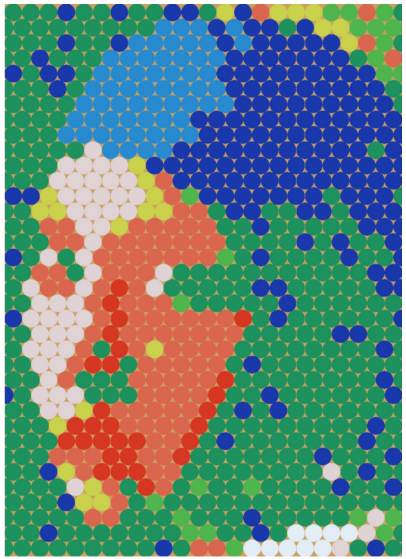

Initial solution

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

Girl with a Pearl Earring by Johannes Vermeer

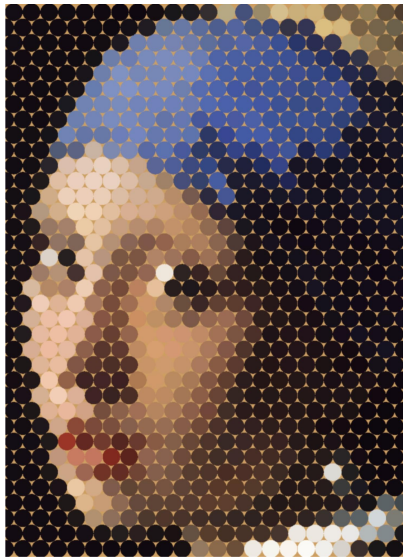

Spatial discretization

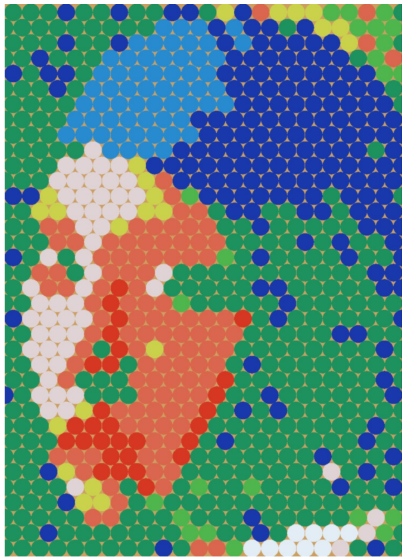

Initial solution

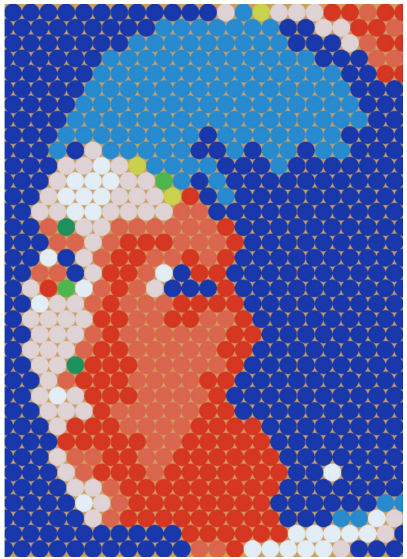

Computer-generated bottle cap art

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

Girl with a Pearl Earring by Johannes Vermeer

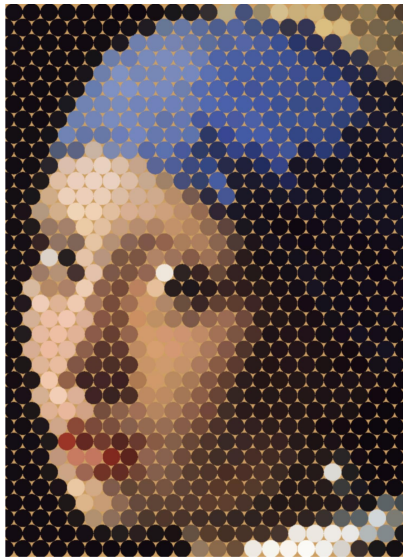

Spatial discretization

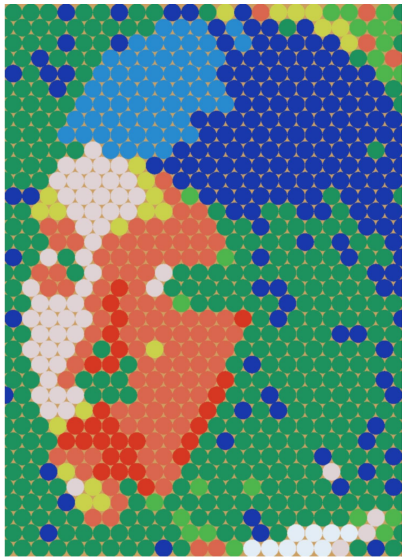

Initial solution

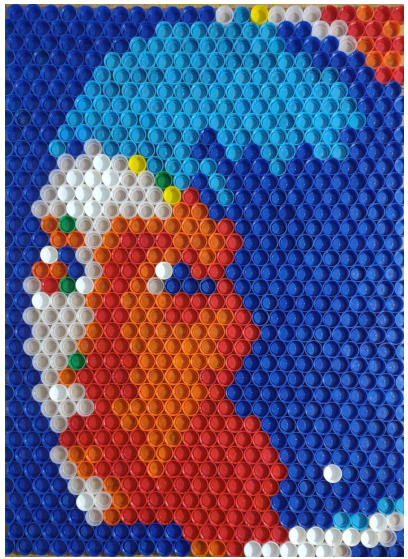

Physically assembled bottle cap art

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

By Flickr user zaphad1

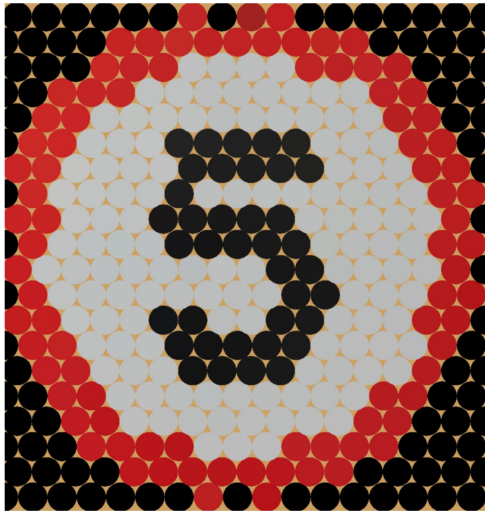

Spatial discretization

(Assembled) initial/final solution

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

By Flickr user Mike Sayre

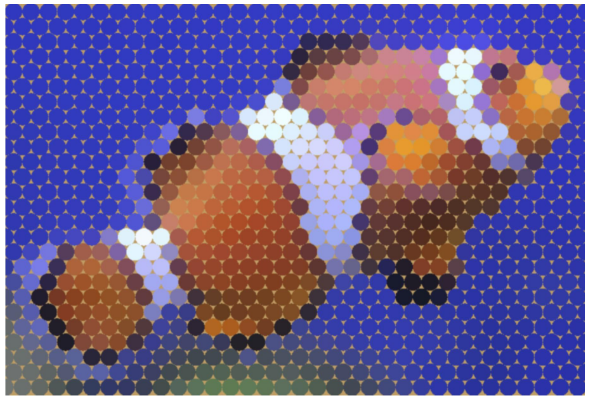

Spatial discretization

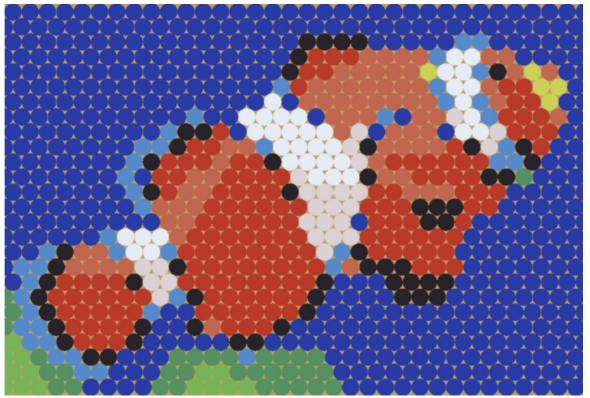

Computer-generated bottle cap art

Results

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

By Flickr user Christopher Gabbard

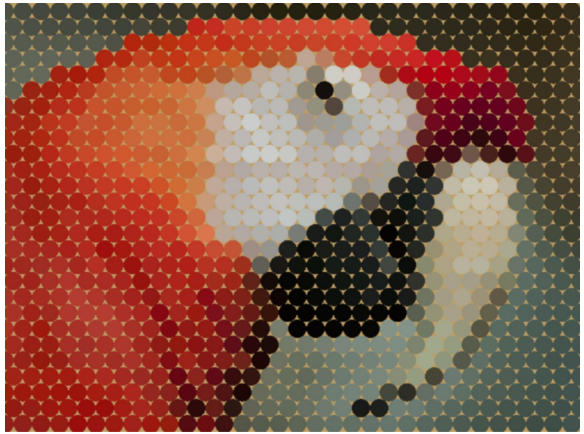

Spatial discretization

Physically assembled bottle cap art

Pixel art methods

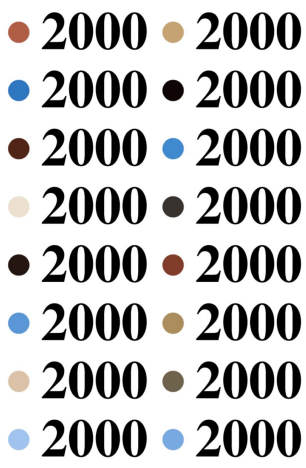

Caps

Input image

By Flickr user
William Warby

Our result

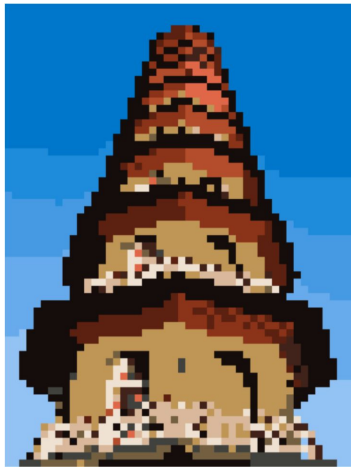

[Gerstner et al. 2013]

[Kopf et al. 2013]

Implementation: Matlab program

- Modification of Matlab's `slic` built-in function
- Vectorization
- Re-use of local SSIM values

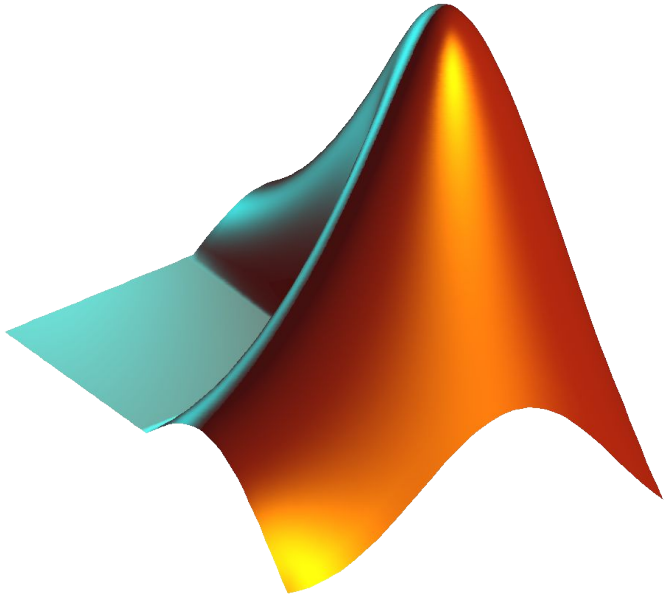

Matlab's logo (Mathworks)

Timings

	(Preprocessing)	(Optimization)	(Physical)
Result	**Secs. 4.1 to 4.4**	**Section 4.5**	**Assembling**
Figure 1	0.33 secs	192.97 secs	1h 45 min
Figure 2	0.31 secs	133.21 secs	1h 43 min
Figure 10	0.72 secs	129.51 secs	1h 43 min
Figure 11	0.72 secs	222.87 secs	1h 27 min
Figure 12	0.77 secs	18.58 secs	35 min
Figure 13	0.67 secs	83.77 secs	37 min

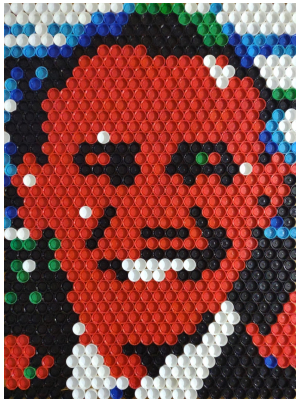

Figure 2

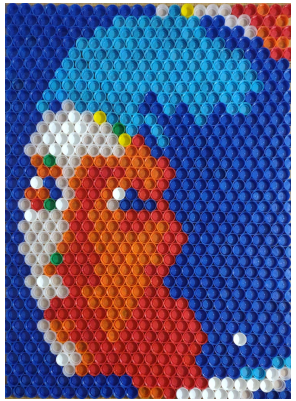

Figure 11

Figure 1

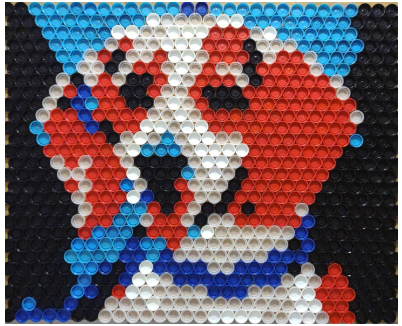

Figure 10

Figure 12

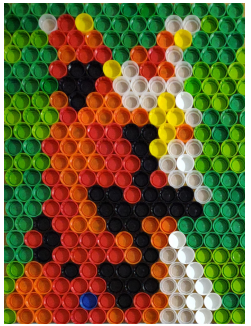

Figure 13

Limitation

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

Input image

By Flickr user Clint Budd

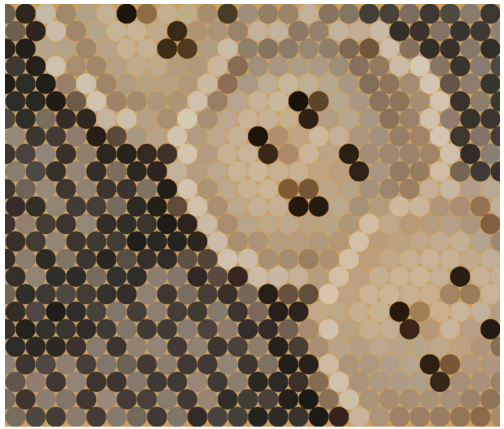

Spatial discretization

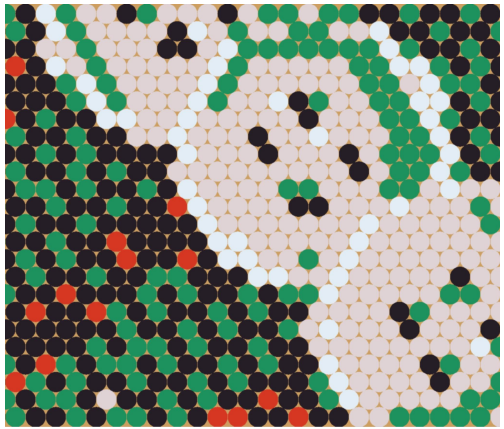

Computer-generated bottle cap art

Future works

- Faster optimization and run the method on image collections

(Optimization)
Section 4.5
192.97 secs
133.21 secs
129.51 secs
222.87 secs
18.58 secs
83.77 secs

Future works

- Faster optimization and run the method on image collections

- Other cap grids and cap sizes

Mandela,
by artist Denise Hugues

Future works

- Faster optimization and run the method on image collections
- Other cap grids and cap sizes
- Other graphics problems that use wasted material

Thank you!

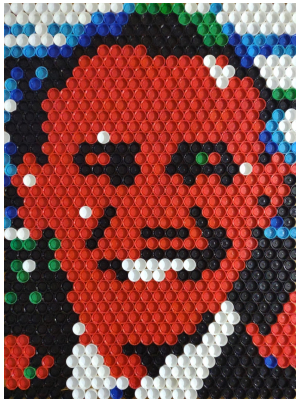

For more, visit the website:

mtm.ufsc.br/~leo/bottle_cap_art/

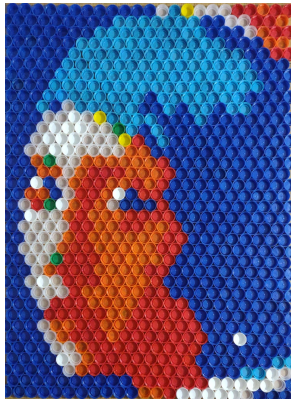

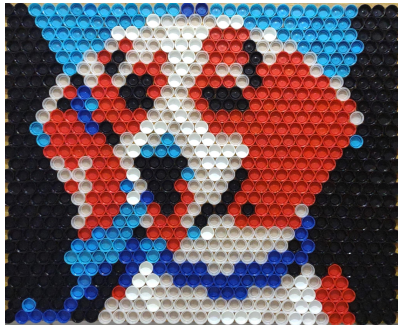

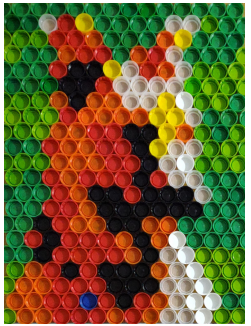

Backup slides

Varying sizes

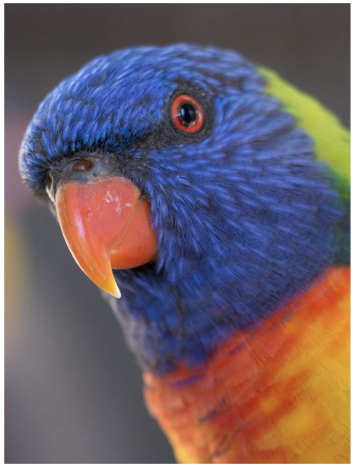

Input image

By Flickr user troposa1

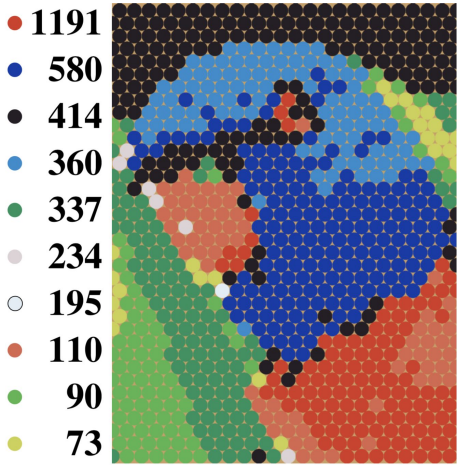

1 meter by 0.75 meters

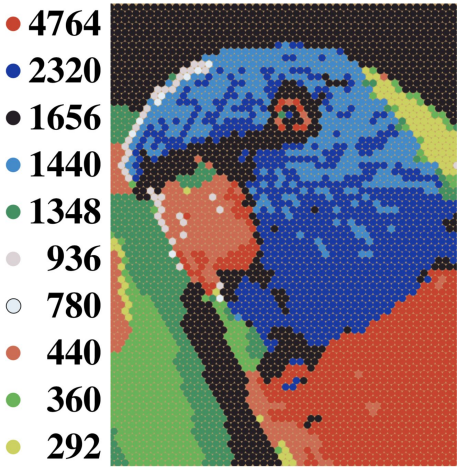

2 meters by 1.5 meters

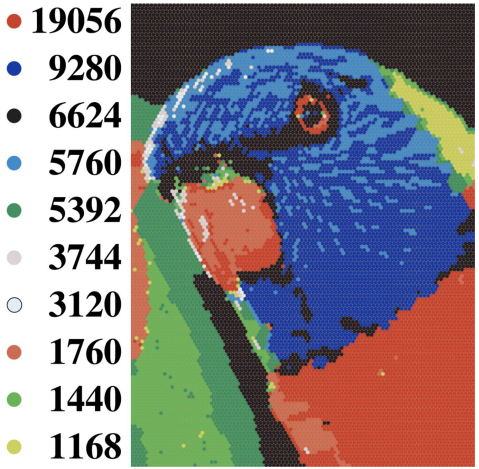

4 meters by 3 meters

Other caps

Input image

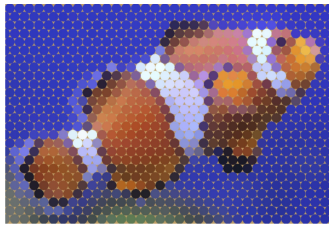

Discretization

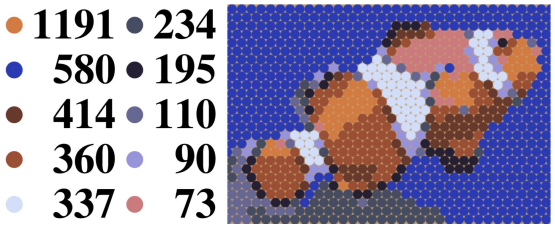

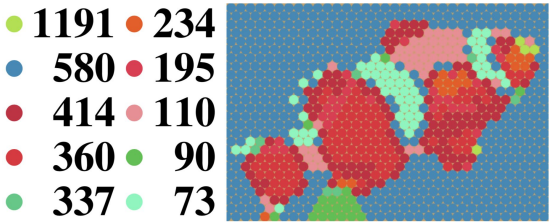

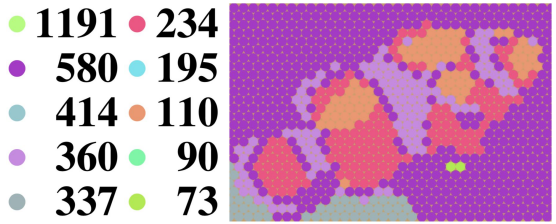

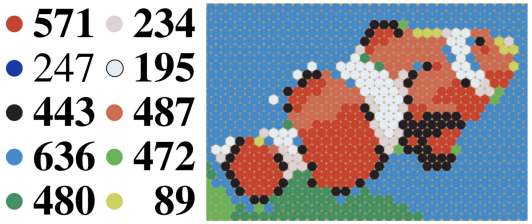

Caps preprocessing

1191 580 414 360 337 234 195 110 90 73

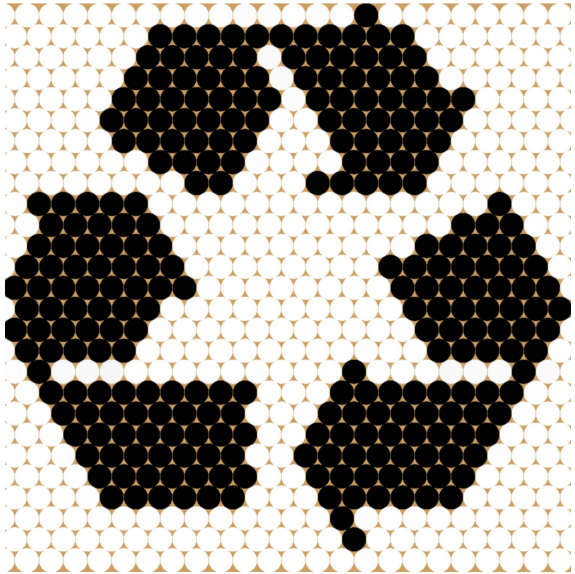

Spatial discretization

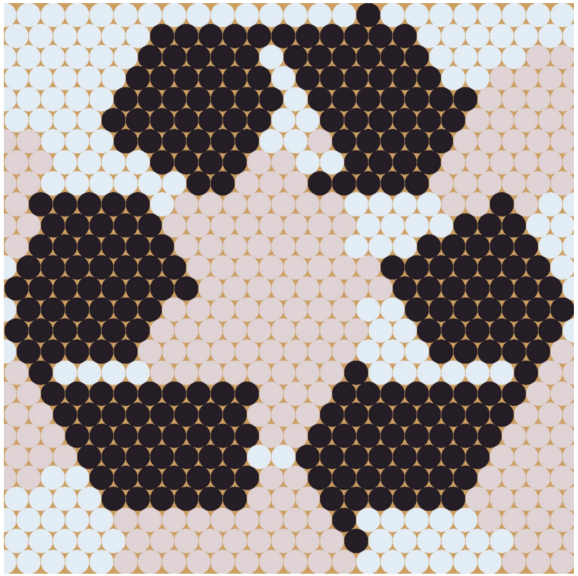

Result using all colors

Caps preprocessing

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

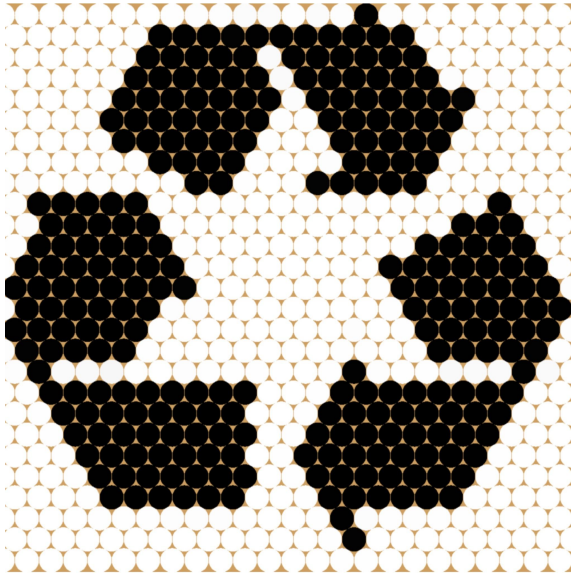

Spatial discretization

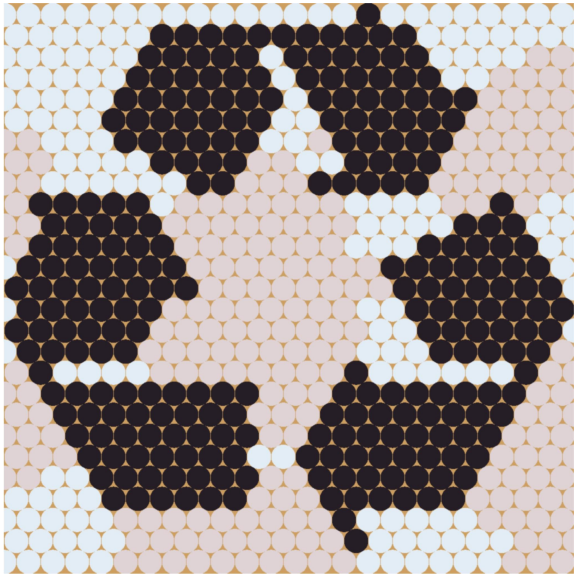

Result using all colors

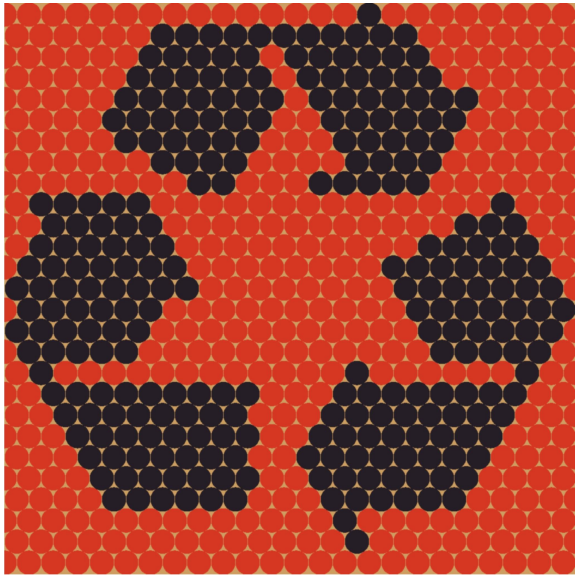

Result

Optimization with full and low resolution reference

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

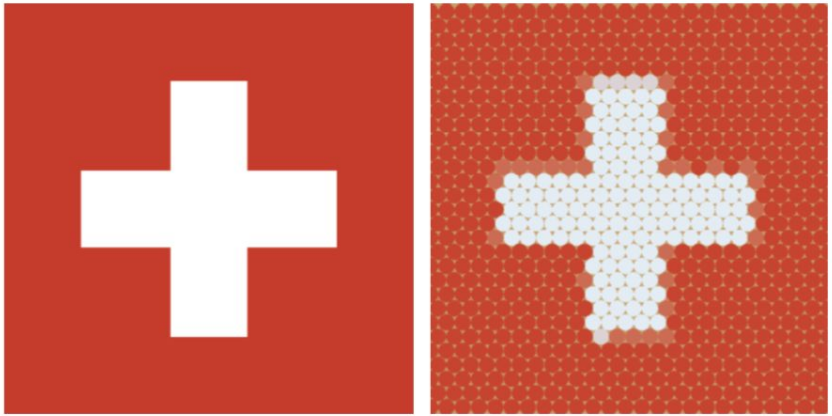

Optimization with
full resolution reference

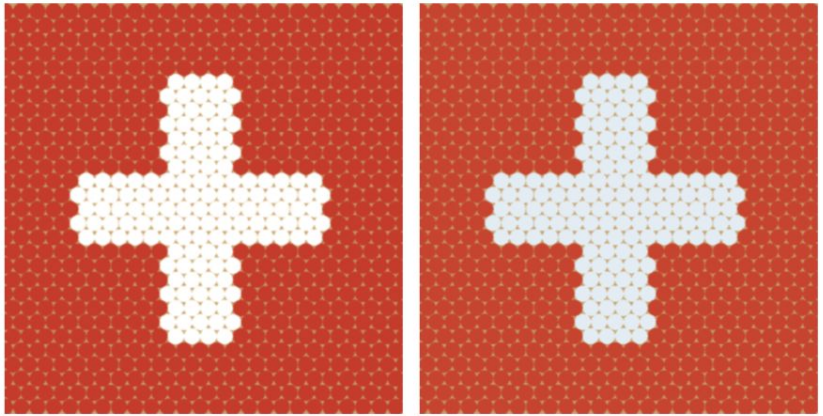

Optimization with
low resolution reference

Varying alpha

● 1191 ● 580 ● 414 ● 360 ● 337 ○ 234 ○ 195 ● 110 ● 90 ● 73

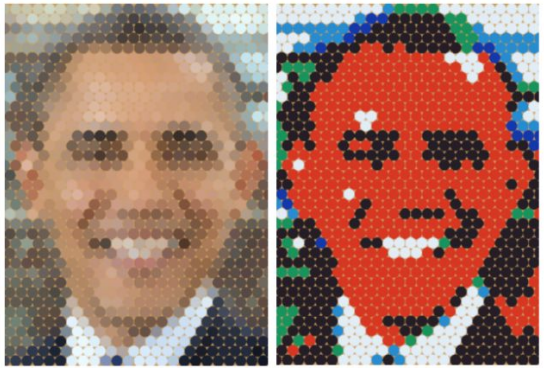

alpha = 500

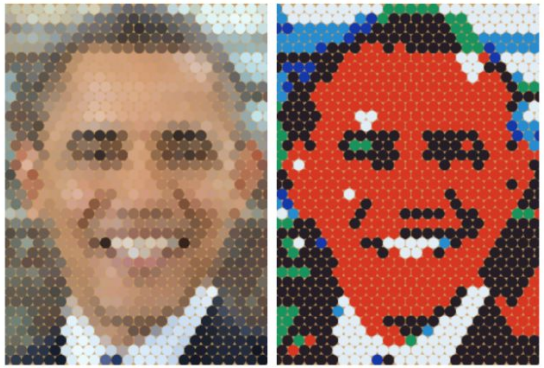

alpha = 100

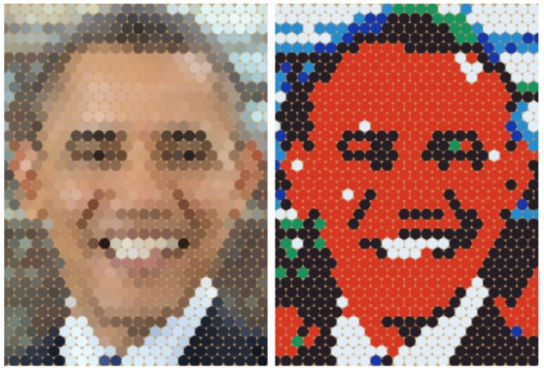

alpha = 50
(our choice)

Penalty term

1191 580 414 360 337 234 195 110 90 73

Input image

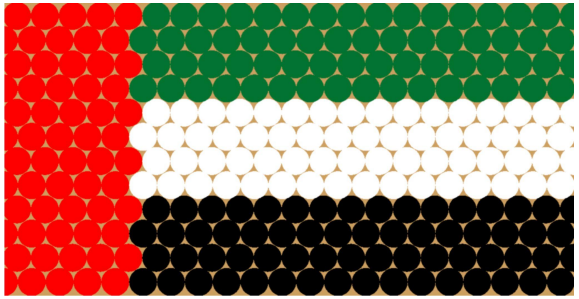

Spatial discretization

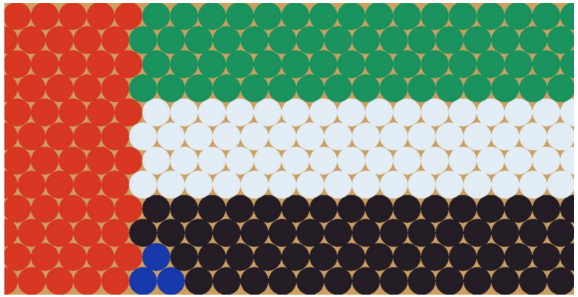

Result without penalty

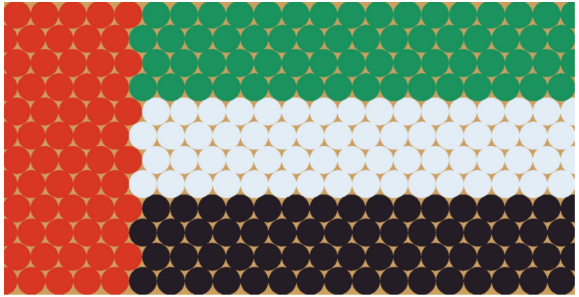

Result with penalty