Generative Art and Communal Curation: A Proof of Concept

C.L. Fisher, BMath
Centennial College, Toronto, Canada; TUAS, Turku, Finland clfisher.com
email: clfisher@rogers.com

Abstract
One of the attractions of creating generative art is releasing control over the final outcome, allowing unexpected results. By involving a community in their process, the artist can separate themselves still further from the result, while maintaining a human element.
Inputs to this proof of concept included a process that creates new images through modified sortation of pixels from selected colour and greyscale photographs, with random factors informing the sort, and a selection of 25 photographs.
After the initial generation the images were presented pair-wise to a small set of people recruited from the artist's social media contacts and from the Arts Academy at TUAS. Using their aggregated preferences as additional input, a basic genetic algo- rithm was used to determine parameters for a new, smaller generation of images. Multiple iterations of voting and image generation produced a final piece.

Background
The project was an independent work for credit at Centennial College (Toronto), completed during an exchange term at TUAS (Turku). The image processing algo- rithms were developed from July 2018 - August 2019 as part of my ongoing explo- ration of randomness. Randomness attracts me because it divorces the final result from my own efforts - whether the art is appealing or not is not in my control. I bring the selection of the initial images, the decisions of whether to pre-set certain parame- ters, and the naming of the pieces as personal inputs to the process.
The infrastructure for voting, and for producing new images dependent on the voting results was developed from August to September 2019, and the creation of the final piece occurred from September 23 to October 22, 2019.

Overview of process
Figure 1: High level workflow

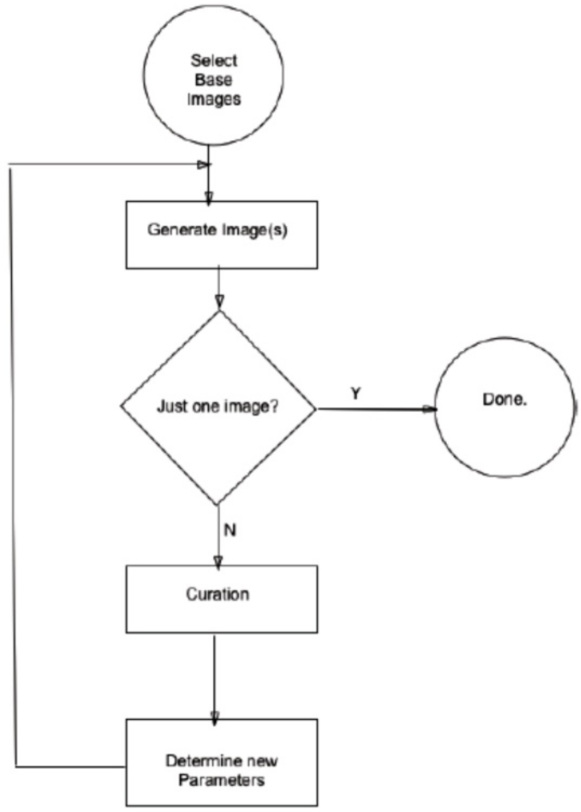

Initial expectations for final result
The outcome is the result of repeated voting rounds, resulting in a kind of consensus image. The element of chance means that sometimes the less popular elements go forward, but over many generations this should correct.

Selection of base images

For this project, I took base images from internet-available portraits of actresses from the first 50 years of American cinema - specifically the AFI Top 25 Legendsi. These images represent one kind of image manipulation — that of the studio system, and the process of creating the new images riffs off of that. The subject resonates with my long-standing interest in Golden Age Hollywood, dating from childhood.

When I found a pleasing colour image of the actress, I used it, even if the bulk of the actress's work was in black and white. This left 3 black and white images.

The process introduces enough change that copyright of the original images is not a concern.

The images are all of size 8.5" x 11".

Creation of images

Basic mechanism: Modified bubble sort

I create the images using a modified bubble sort of the pixels in the base image.

On each pass of the sort, I determine a random x and y offset for the whole pass. Each pixel is compared to the pixel indicated by this bounded offset, rather than be- ing compared to an adjacent pixel.

The offsets are randomly determined using an 8-input formula which was a "happy accident" from playing with the code. The resultant images are highly sensitive to the values of these inputs.

For some initial images, I pre-determined the inputs with values that are known to settle down and produce relatively smooth results. For others, the inputs were ran- domly determined within a wider range. I experimented with this range throughout the generation of the initial images.

Control Points

Instead of sorting across the entire image, uninterrupted, or being constrained by predetermined lines, the sort is governed by a number of control pointsii.

On each comparison, the process checks which of the two points being compared is nearer to the control point closest to the point being evaluated. It uses the nearer point as the base for the comparison. In other words, if the offset point is closer to the control point, the result for the comparison will be inverted. This is what pro- duces the starburst effects seen in some of the images in Appendix III.

Tint

One input into the process is whether to tint. If tinted, the output is treated using a process where each pixel receives a hue corresponding to its greyscale value. Only

2 of the original set of images (#7, #15) were tinted.

Communal Curation

The set of voters for communal curation of this piece was drawn from two groups: my social media contacts (people with technical backgrounds ranging in age from mid-20s to early 60s, plus some recent art school graduates), and art students and instructors at TUAS.

The voters visit a website where they are presented images, pair-wise. Their pref- erences are stored with no identifying information. The voters may keep voting until they have seen all the images for the active generation.iii

After a set period, the voting for a round closesiv, and I feed the results into the process for generating the next, smaller set of images. The vote/generate cycle is repeated until there is only one image left.

I don't know how many of the voters saw all the pairs, or how many saw every pair.

That information was not kept.v

Some voters reported a strong preference for certain pieces, but there is no direct data on preference strength.

Titling the images

I considered labelling the images with their database ids, but wanted something a voter could use to discuss an image more comfortably. I felt using the actress's names could strongly influence the results by association. So, the titles of the first set of images were single words from the title of a movie the actress was in ("Rebec- ca"), from a line spoken by the actress ("Seatbelts") or from a particular scene asso- ciated with the actress ("Budgie"). Names in subsequent rounds still relate to the ac- tresses, but in equally obscure ways. I felt confident that these titles would not be directly associated with the actresses by most of the votersvi.

Anecdotally, we know the names of the pieces influenced at least one voter, though not through association with the subject matter.

Producing the next generation

Once voting for a round was completed, a program tallies the votes, ranks the im- ages and determines breeding pairs for the next generation. The highest-ranked im- age is paired with the last, 2nd with second-last, etc.. If there are an odd number of images, the first-place image is bred with the images in the last two places instead. vii

A number of parameters are bred for input into the next generation: base image; base-colour or tintviii; number of control points; control points (x and y separately); parameters for generating the bubble-sort offsets; maximum number of sort passes; stop threshold by % of pixels moved in the pass; and maximum time for the sort.ix

The process selects the values for the next generation image according to a weight- ed chance on each parameter — if Image A has m votes, and Image B has n votes, the chance that Image A will provide that parameter is m/(m+n).

Two parameters have additional adjustments — a chance at mutation in the breeding process.

Base Image mutation

Once chosen, we process the base image, randomly flipping individual bits in some pixels on a tiny chance. This produces a small variation in the base image, which accumulates and becomes a discernible component of the output image over several generations.x

Control point variation

The chosen control point's x and y values have a 10% chance of being nudged up to 10% towards an edge. This introduces more variation in the new generations than would be present with a strict copying of the original control points.

Observations:

Feedback into voting process

After round 1, the TUAS students recommended adding a progress bar, moving the titles above the images, and changing from radio buttons and a single vote button to individual voting buttons for each image. These suggestions were implemented for round 2.

Attrition

Although we cannot know, except anecdotally, who voted in each round, we can es- timate the rate of attrition by tracking the ratio of total votes to number of available pairs of images to vote on. This is an imperfect measure because the number of people completing the entire set of images pairs should increase as the number of images decreases.

Table 1: Votes by Round

Round	Number of Images	Number of Image Pairs	Number of Votes	Votes / Image Pairs
1 (Sep 23 - 25, 2019)	25	300	4431	14.77
2 (Oct 1 - 3)	13	78	785	10.06
3 (Oct 6 - 8)	7	21	299	14.24
4(Oct 9 - 11)	4	6	78	13
5(Oct 19-22)	2	1	15	15

Anecdotally, causes for the dip in round 2 included travel and technical difficulties. These might be reduced with a longer voting period, but a short voting period en- courages prompt action. The voting for round 1 was a scheduled event for the TUAS students, but not for subsequent rounds.

Overall, the rate of attrition was within expected bounds - we had 22 reported voters in the first round, and 15 in the final. The voting was purely a volunteer activity, and it took place over a relatively short period of time — one month for the 5 rounds. The burden of fully participating decreased with time, as the number of pairs available decreased rapidly - from 300 pairs in the first round to 1 in the last.

Selection of images and "rule of three"
I thought the composition of the images might play a strong role in what people vot- ed for. The control points correspond to the centres of the starburst elements, which are the strong centres of interest in the 'well-behaved' cases.
However, basic analysis shows that the voters did not tend strongly to vote for the images that had control points closer to the thirds. (See Appendix VI). If there was a preference for images with elements according to the rule of three, it was over- whelmed by other factors.
To what extent did each generation resemble the winners from the previous one?

Base Image

The base image used naturally has an outsized influence on the image. The process uses the whole (albeit mutated) base image from one of the parents, In the first round, only 6 of the winning images had their base images passed on to the
next generation, despite an overall advantage in votes of 2800 to 1872.xi
Although
the winning image from the first round was bred twice, with over a 70% weight each time, its base image did not survive. See Appendix IV.
Tint
Only one of the tinted images was a winner in the first round. Tint was eliminated as a factor in succeeding rounds.

Number of control points
20/27 pairings produced the same number of control points as the winning element.

Control points
10-20% of the points were "nudged" on each round, so we do not expect a clear cor- respondence across generations. In 16 out of 27 transitions, the distance between the control points in the image with more votes and the chosen points was less than the distance between the control points with fewer votes and the chosen points. This is the amount we would expect with the ratio of votes. See Appendix V.
Offset parameters
The winners of the first round, with one exception, were all created with the well-be- haved offsets. An inspection of the round 2 images shows that there was sufficient input from the wilder images for visibly looser behaviour in 3 cases. 2 of the wilder images won in that round, but the images became tamer as the rounds progressed. (See Appendix III for image thumbnails)
Compositional elements
At the start of the project, I wondered whether the final composition would cleave more closely to traditional rules such as "rule of three" than the initial images did.

Table 2: Convergence to thirds check

Generation	Mean distance to thirds - X	Mean distance to thirds - Y
1	0.122	0.143
2	0.118	0.159
3	0.116	0.138
4	0.145	0.187
5	0.134	0.179
6	0.111	0.297

There appears to be no convergence to rule of three. The "nudging" might be play- ing too big a role in the positioning for any meaningful convergence to take place.

Going forward:
Base Image
The base image weighs heavily in the appearance of the images. With weighted randomness, favoured images can, and did, end up essentially being lost to the process.

To mitigate this, we could create new base images from the parent images by ran- dom combination of the two images, on a weighted basis. This introduces a risk of introducing aesthetic difficulties — some colour palettes just don't work well together. It might be instructive to do a dry run based on the votes in the first round.

Tint Factor

The tint factor was eliminated after the first round, although both the images it was used on made it through to the second round. Including more tinted images (with a variety of base hues), would increase the likelihood of the tint being a factor in later rounds.

Convergence?

On the next iteration, the nudging factor should be reduced. There is a balance be- tween introducing enough motion in the points, and overpowering the will of the vot- ers. It is not certain that the will of the voters would be toward convergence.

Community Composition

This proof of concept was produced with the help of an artificial, mixed community. The next iteration could be produced with a tighter community, working with initial images that are meaningful to that community.

Breeding mechanism

I was dissatisfied with the loss of voter-favoured elements after the first round. This could be mitigated by promoting the top few images from early rounds into subse- quent rounds — at a cost of larger voting pools, and possibly higher attrition

To investigate in future iterations

1/ Effect of language: The voters were asked to choose the image they preferred. What if they were asked which was best? The most striking? Or which made them feel more?

2/ Mechanisms for measuring individual strength of preference. There is a difference between preferring one image over another, and strongly liking the preferred image. Could we measure that in a user-friendly way? What would it mean?

3/ Being able to look at an individual voters choices (still anonymized). It would be interesting to see whether voters are consistent in their choices - whether someone who prefers A to B and B to C might prefer C to A, for example.

Appendix I : Acknowledgements

Thanks to the people who made this project a success:

Votersxii

Sofia Aarmio, Diane Bayley, Brian Peter Dickson, Jesper Dolgov, Katharine Draper Quinn, Pene Gerber, Saara Hast, Kati Immonen, Heli Janhiainen, Venla Kaasinen, Aapo Kotilainen, Sakari Kyyrönen, Matti Lankireen, Jussi Lipasti, Eero Merimaa, Mehtap Memmi Mertdoğan, Natalie Plociennik, Bob Shaland, Michelle Wehrle, Maria West, Anita Woodard, Shiyu Zhang

Testers

Seonaid Lee, Robert Quinn

Advisors

Eero Merimaa (TUAS)

Lisa Binnie (Centennial College)

David McClyment (Centennial College)

Appendix II: Distribution of votes

Table 3: Distribution of votes, Generation 1

Image number	For	Against	% for	Image number	For	Against	% for
1	134	233	36.5	**13**	164	201	44.9
2	104	247	29.6	**14**	194	171	53.2
3	228	124	64.8	**15**	212	142	59.9
4	91	265	25.6	**16**	104	247	29.6
5	219	133	62.2	**17**	117	231	33.6
6	195	155	55.7	**18**	209	150	58.2
7	122	228	34.9	**19**	177	173	50.6
8	187	177	51.4	**20**	240	121	66.5
9	127	217	36.9	**21**	244	114	68.2
10	184	167	52.4	**22**	188	162	53.7
11	206	145	58.7	**23**	201	151	57.1
12	196	161	54.9	**24**	212	147	59.1
				25	173	182	48.7

Table 4: Distribution of votes, Generation 2

Image number	For	Against	% for	Image number	For	Against	% for
26	61	60	50.4	**33**	54	67	44.6
27	62	65	48.8	**34**	74	42	63.8
28	50	71	41.3	**35**	75	47	61.5
29	59	59	50	**36**	47	73	39.2
30	51	70	42.1	**37**	77	43	64.2
31	54	68	44.3	**38**	73	47	60.8
32	76	45	62.8				

Table 5: Distribution of votes, Generation 3

Image number	For	Against	% for	Image number	For	Against	% for
39	41	44	48.2	**43**	58	27	68.2
40	37	48	43.5	**44**	44	41	51.8
41	30	56	34.9	**45**	45	41	52.3
42	44	42	51.2				

Table 6: Distribution of votes, Generation 4

Image number	For	Against	% for	Image number	For	Against	% for
46	20	19	51.3	**48**	19	20	48.7
47	19	20	48.7	**49**	20	19	51.3

Table 7: Distribution of votes, Generation 5

Image number	For	Against	% for	Image number	For	Against	% for
50	5	10	33.3	**51**	10	5	66.7

Appendix III: Image Thumbnails

Figure 2: Generation 1 Image Thumbnails

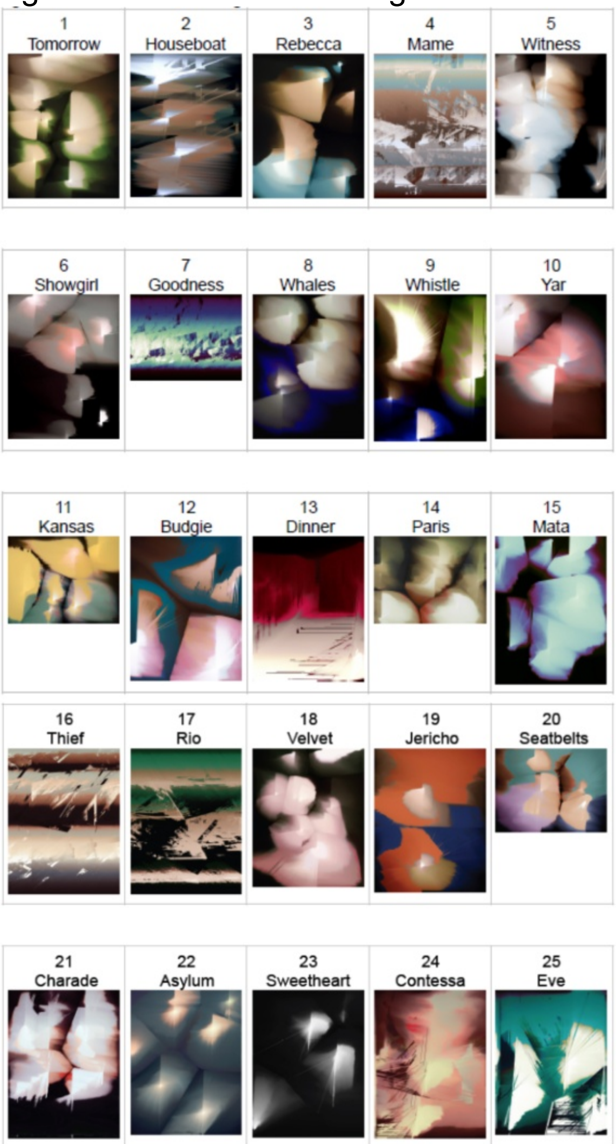

Figure 3: Generation 2 Image Thumbnails

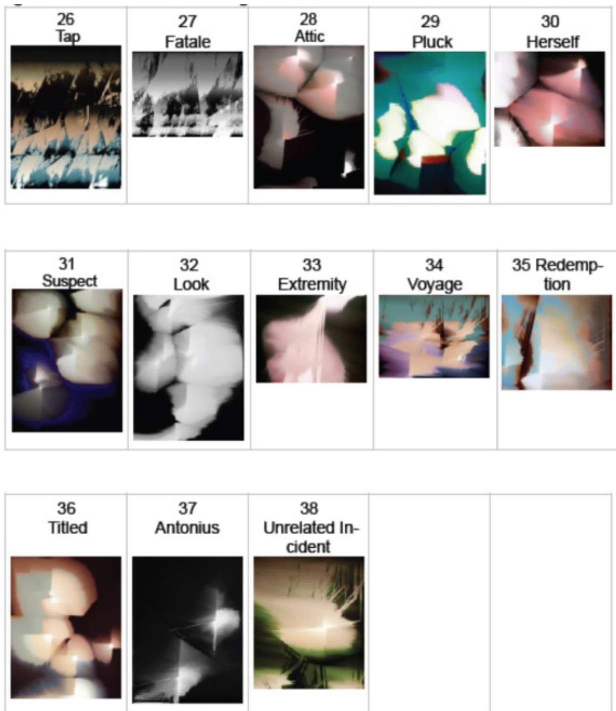

Figure 4: Generation 3 Image Thumbnails

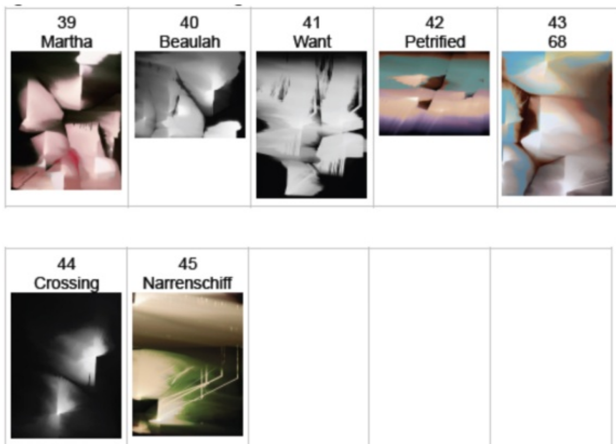

Figure 5: Generation 4 Image Thumbnail

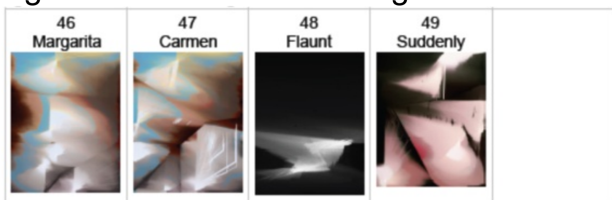

Figure 6: Generation 6 Image Thumbnails

50 University Av-enue	51 Castilleja de la Cuesta			
				

Figure 7: Generation 7 Image Thumbnail

52 Finale				
				

Appendix IV: Base Image Inheritance

Table 8: Inherited base images Generation 1 to Generation 2

Image 1	Image 2	% chance for image 1 to prevail	Inherited base image	Round 2 im-age
21	16	70.1	16	36
20	2	69.8	20	34
3	17	66.1	3	26
5	7	64.2	7	27
24	1	62.5	1	38
15	9	61.3	15	32
18	13	56	18	33
11	25	54.4	25	29
23	19	53.2	23	37
12	10	51.6	10	30
6	22	50.8	6	28
14	8	51	8	31
21	4	72.8	4	35

Table 9: Inherited base images Generation 2 to Generation 3

Image 1	Image 2	% chance for image 1 to prevail	Inherited base image	Round 3 im-age
37	31	58.7	37	44
35	28	60	35	43
34	36	61.2	34	42
38	30	58.9	38	45
27	29	51.2	27	40
26	33	53	33	39
34	32	62.2	32	41

Table 10: Inherited base images Generation 3 to Generation 4

Image 1	Image 2	% chance for image 1 to prevail	Inherited base image	Round 4 im-age
43	40	61.1	43	46
45	39	52.3	39	49
44	42	50	44	48
43	41	65.9	43	47

Table 11: Inherited base images Generation 4 to Generation 5

Image 1	Image 2	% chance for image 1 to prevail	Inherited base image	Round 5 im-age
49	48	51.3	48	51
46	47	51.3	47	50

Table 12: Inherited base images Generation 5 to Generation 6

Image 1	Image 2	% chance for image 1 to prevail	Inherited base image	Final image
51	50	66.7		52

Appendix V: Point position distances between generations

Table 13: Distance between source and resultxiii

Images(most votes..least votes..result)	Distance measure of result from 'winner'	Distance measure of result from 'loser'	Result is closer to 'winner'
21..16..36	1.17	2.20	true
20..2..34	0.88	1.15	true
3..17..26	0.54	1.90	true
24..1..38	0.00	2.51	true
15..9..32	1.03	1.05	true
18..13..33	0.00	0.87	true
37..31..44	0.00	1.05	true
35..28..43	0.87	2.54	true
27..29..40	0.28	1.86	true
34..32..41	0.00	1.95	true
43..40..46	0.78	1.17	true
44..42..48	0.20	0.49	true
43..41..47	0.43	1.67	true
49..48..51	0.00	0.90	true
46..47..50	0.39	0.82	true
50..51..52	0.86	1.52	true
5..7..27	2.14	0.09	false
11..25..29	0.92	0.55	false
23..19..37	0.59	0.38	false
12..10..30	1.68	0.98	false
6..22..28	0.98	0.94	false
14..8..31	3.28	0.00	false
21..4..35	1.96	1.12	false
34..36..42	0.94	0.90	false
38..30..45	1.65	1.56	false
26..33..39	0.17	0.00	false
45..39..49	2.60	0.31	false

Appendix VI: Distance from rule of thirds

Table 14: Distance between images and rule of thirdsxiv

Pair	Distance of winner	Distance of loser	Winner is closer
21..16	0.091	0.129	true
20..2	0.119	0.102	false
3..17	0.135	0.131	false
5..7	0.149	0.107	false
24..1	0.2	0.104	false
15..9	0.128	0.147	true
18..13	0.091	0.2	true
11..25	0.119	0.151	true
23..19	0.181	0.101	false
12..10	0.182	0.178	false
6..22	0.142	0.137	false
14..8	0.15	0.087	false
21..4	0.091	0.172	true
37..31	0.11	0.087	false
35..28	0.18	0.159	false
34..36	0.131	0.096	false
38..30	0.2	0.187	false
27..29	0.13	0.122	false
26..33	0.135	0.177	true
34..32	0.131	0.132	true
43..40	0.174	0.119	false
45..39	0.215	0.134	false
44..42	0.11	0.09	false
43..41	0.174	0.131	false
49..48	0.169	0.108	false
46..47	0.117	0.195	true
50..51	0.147	0.169	true

Notes

i"AFI's 100 YEARS...100 STARS." American Film Institute, American Film Institute, https://www.afi.- com/afis-100-years-100-stars/.

ii These were randomly determined for the initial images.

iii They could have gone to another machine, or cleared their cookies, and repeated the process. For a proof of concept, safeguards against that behaviour were deemed overkill.

iv The voting was not actually disabled, which caused some data analysis headaches later.

v Four students volunteered the information that they completed all the pairs in the first round.

vi I am the only old-film buff among my social media contacts. My experience with the student de- mographic suggests they have little knowledge of pre-1980 popular culture.

vii This was thought to provide the best chance for the more popular image components to survive into the next generation.

viii There were two tinted images in the first generation. That was the last generation where tint was a factor.

ix Obviously, only one of maximum number of sort passes; stop threshold by % of pixels moved in the pass; and maximum time for the sort will actually be relevant. We don't know which one will be relevant before creating the image, though.

x The change is 1 in 10000 — far greater than my original intention of mimicking human gene muta- tion rates, but sufficient for noticeable effects over time. It is visible in the final image, if you already know it is there. It is just barely discernible in the "black and white" image the generation before.

xi That includes the votes for the first place image twice - because it was matched up to the two low- est -vote receivers.

xii

There was no requirement that the voters identify themselves, and the total pool of possible voters was significantly larger than this list.

xiii This is the sum of distances between corresponding points in the list of points, If there are more points in one image than the other, the surplus points are ignored.

xiv This is the sum over the points, of the distance from the closest horizontal third (x = 0.333, x =

0.667) and the distance from the closest vertical third (y = 0.333, y - 0.667)