Supported by European Union Horizon 2020, Project No. 810961

cudan.tlu.ee

Capturing aesthetic complexity in art using compression ensembles

Andres Karjus, Tillmann Ohm, Mar Canet Solà , Maximilian Schich

(ERA Chair for Cultural Data Analytics, Tallinn University)

Sebastian Ahnert

(University of Cambridge, The Alan Turing Institute)

The Conference on Complex Systems CCS 2021, Lyon

Questions → twitter.com/AndresKarjus

Slides and contact: andreskarjus.github.io

Quantification of visual aesthetics

- Has a long history (cf. Birkhoff 1933, Rigau etal 2007, Forsythe etal 2011, Tran etal 2018, Sigaki 2018, Lee etal 2020)
- Including visual complexity, using image compression (Rigau etal 2007, Müller etal 2018, Palumbo etal 2014; Bagrov etal 2020; also in adjacent domains, cf. Tamariz etal 2015, Miton etal 2021, Han etal 2021)
- The creative process of an artist as an algorithm
- We aim to capture these "algorithmic fingerprints", to quantify polymorphic family resemblance and the evolution of art in the complexity space
- Various methods, results against human judgements diverge
- Building on these ideas, we propose a novel general approach: "compression ensembles"
- Dataset: 70k paintings&drawings from Wikiart/art500k (Mao et al 2017)

How does it work?

- In the visual domain: instead of compressing just the image: produce an array of transformations & compress all of them
- Yields a vector of compression ratios; add various stats like colourfulness and fractal dimension estimators (n=109 total)
- Fit into a latent space using PCA to visualize and avoid collinearity where needed

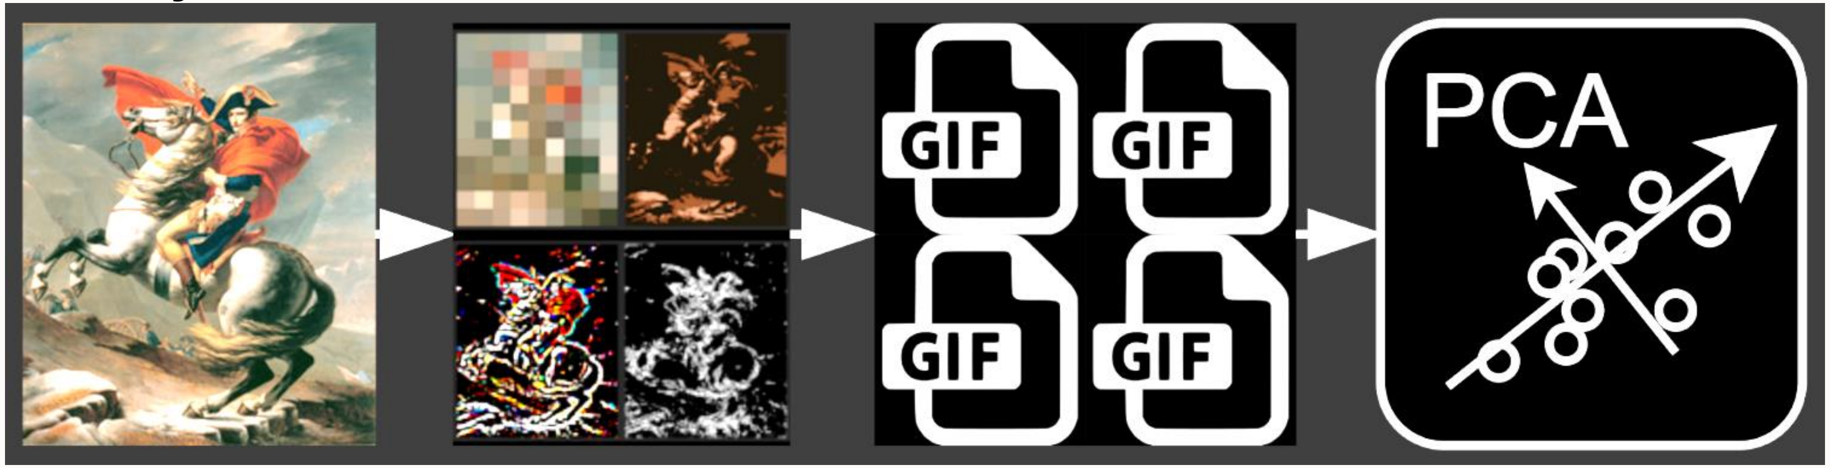

How does it work?

- Different transforms are informative of different aspects
- Pixelating a detailed image reduces its compression size
- Grayscaling an already grayscale image won't change size

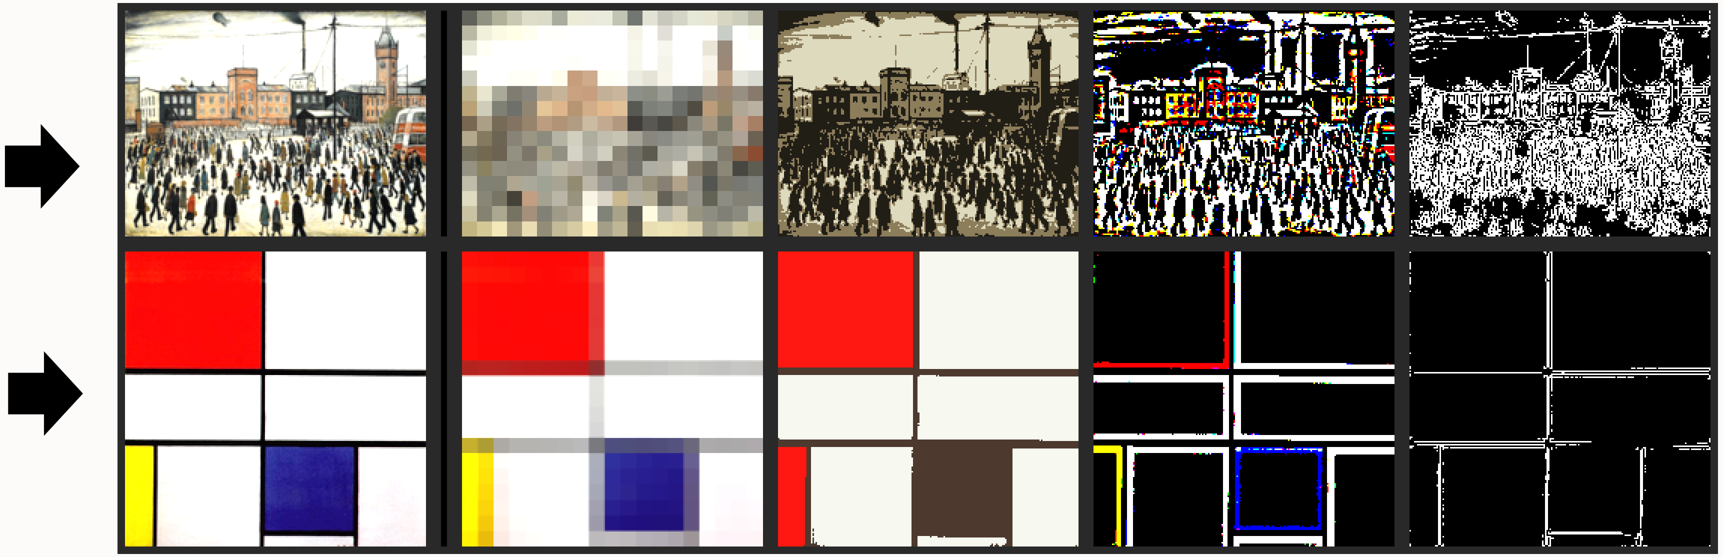


PC2: pattern and color complexity (pixelation and scramble transforms)
4
0
-4
-10
-5
0
5
10
PC1: overall compressibility (jpeg, gif, png), edge detection (comparison to blur, emboss filter, local adaptive thresholding)

the top right high-detail corner →

(high overall colour and pattern complexity)

the left-middle "portrait area" ↓

(medium colour and pattern complexity, below average overall and edge complexity)

How well does it work?

- These are pretty graphs but does this method actually capture what human beings perceive as visual complexity?
- Or, how human(or art-historian)-like is our mindless machine?
- We evaluate on two tasks: estimation of human visual complexity ratings, and automatic authorship and style retrieval

How well does it work?

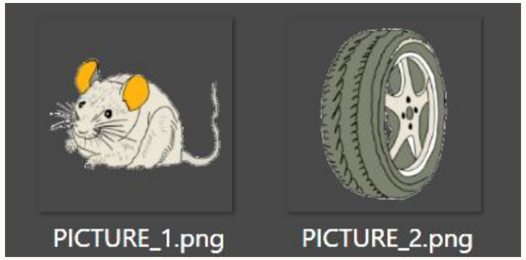

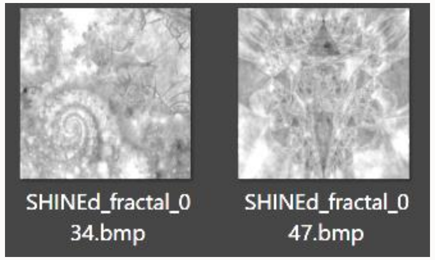

- Human complexity norms datasets, Multipic (n=750 per language; Duñabeitia et al 2018) and Fractals (n=400; Ovalle-Fresa et al 2020)
- Our method outperforms previous single-compression to human judgement correlations (and median prediction error 0.21 on Multipic is smaller than the average difference between the languages)

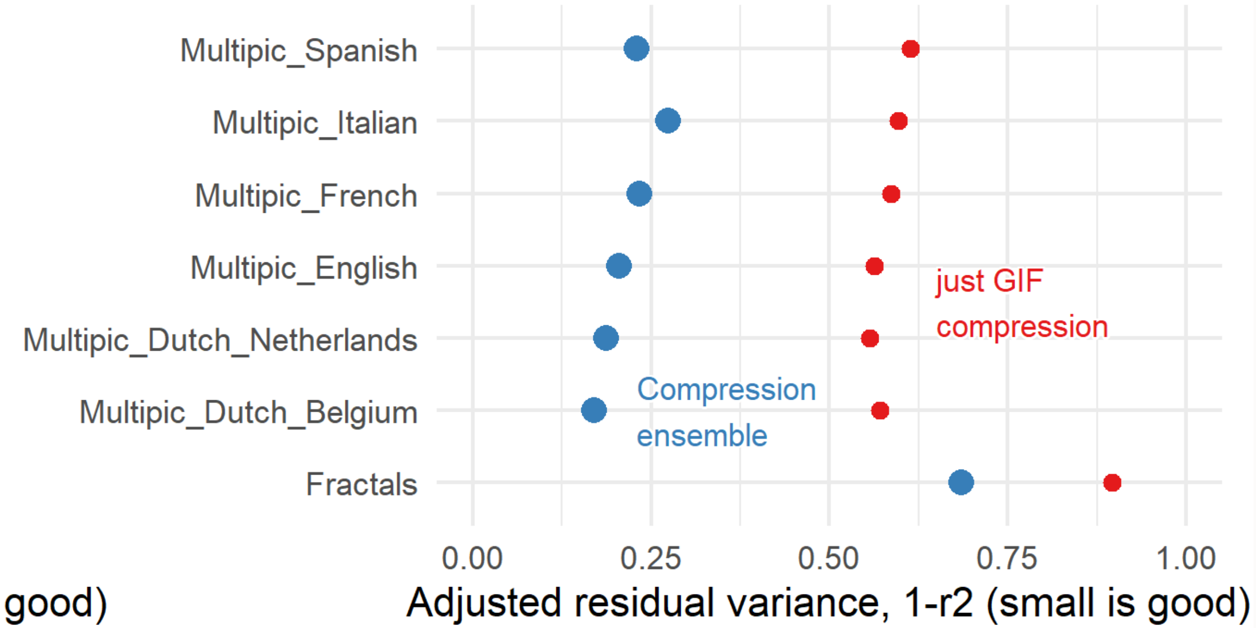

Type, media, author & style detection

- Simple model: linear discriminant analysis. Will report % of correct classifications as kappa scores, random chance adjusted accuracy (accuracy-baseline)/(100-baseline)
- In the 70k set, with enough training examples, easily distinguishes between landscape vs portrait, and oil vs drawing, both at ~85% accuracy or 70% kappa after adjusting for the 50% baseline
- Detects artist (in a subset of 91) about 40% of the time
- Correctly identifies style period and century ~30% of the time
- Shows that the compression ensembles capture similarities and differences between artwork types, authors and style periods

- Even with a handful of examples and a couple of transforms, classifies above the random baseline

Author (n=91)

N training examples:	2	20	100
+all 82 transforms	5	31	38%
+another 20 transforms	9	28	33
+next 20 transforms	9	23	27
+compress_jpeg0_x0.4	7	13	15
+fft1	6	12	13
+flood_hole_gif	6	12	13
+morph_pixelate20_gif	6	10	12
+fx_deskew_zoom_gif_x0.4	5	10	11
+stats_colorfulness_lab	5	9	10
+lines_division_gray_png_x0.4	5	7	8
+fft1_blur10	3	5	6
+blur10_png_x0.4	3	4	5
lines_division_gray_gif	1	2	2

Style period (n=13)

N training examples:	2	20	100	1000
+all 86 transforms	5	13	24	28%
+another 20 transforms	5	15	23	26
+next 20 transforms	4	16	22	24
+colors_grayscale_gif	5	14	17	18
+lines_bw_canny_gif	5	14	17	18
+stats_angleentropy	5	14	16	17
+lines_cartoon_gif	5	11	14	15
+lines_edge2_color_gif	5	11	13	14
+lines_edge5_gray_gif_x0.4	4	10	12	12
+color_chroma_divide_gif	5	9	11	11
+stats_colorfulness_rgb	5	8	10	9
+color_luminance_divide_gif	4	6	6	7
lines_edge_lat_gif	2	4	4	5

Drawing or oil painting

	2	20	100	1000
+all 81 transforms	16	39	54	69%
+another 20 transforms	11	28	58	66
+next 20 transforms	13	26	57	62
+colors_p10_gif	12	43	52	55
+color_chroma_divide_gif	9	41	52	56
+colors_round_gif	10	42	52	54
+emboss_conv_grayp_gif	12	44	52	54
+blur10_gif	13	45	52	52
+lines_bw_canny_gif	9	26	35	38
+colors_add2_gif	7	24	32	35
+colors_grayscale_gif_x0.4	3	12	16	17
+compress_gif	4	13	17	18
colors_quantize5_gif	3	10	15	15
N training examples:	2	20	100	1000

Landscape or portrait

	2	20	100	1000
+all 89 transforms	15	43	53	71%
+another 20 transforms	10	31	62	71
+next 20 transforms	9	27	61	66
+emboss_conv_grayd_gif	14	44	54	57
+flood_hole_gif_x0.4	16	44	54	56
+stats_colorfulness_rgb	13	44	53	56
+fx_deskew_zoom_gif	15	45	54	54
+blur30_gif	17	47	53	54
+lines_edge5_gray_png	18	47	53	52
+lines_color_conv_gif	19	46	50	50
+lines_hough40_gif	25	47	49	50
+color_luminance_divide_gif	23	46	49	49
fractaldim2	32	48	49	49
N training examples:	2	20	100	1000

*the total number of transforms varies, as constant and collinear variables for a given subset are removed from the classifier

*transforms ordered by a rough estimate of variable importance

Author (n=91)

	2	20	100
+all 82 transforms	5	31	38%
+another 20 transforms	9	28	33
+next 20 transforms	9	23	27
+compress_jpeg0_x0.4	7	13	15
+fft1	6	12	13
+flood_hole_gif	6	12	13
+morph_pixelate20_gif	6	10	12
+fx_deskew_zoom_gif_x0.4	5	10	11
+stats_colorfulness_lab	5	9	10
+lines_division_gray_png_x0.4	5	7	8
+fft1_blur10	3	5	6
+blur10_png_x0.4	3	4	5
lines_division_gray_gif	1	2	2
N training examples:	2	20	100

Style period (n=13)

	2	20	100	1000
+all 86 transforms	5	13	24	28%
+another 20 transforms	5	15	23	26
+next 20 transforms	4	16	22	24
+colors_grayscale_gif	5	14	17	18
+lines_bw_canny_gif	5	14	17	18
+stats_angleentropy	5	14	16	17
+lines_cartoon_gif	5	11	14	15
+lines_edge2_color_gif	5	11	13	14
+lines_edge5_gray_gif_x0.4	4	10	12	12
+color_chroma_divide_gif	5	9	11	11
+stats_colorfulness_rgb	5	8	10	9
+color_luminance_divide_gif	4	6	6	7
lines_edge_lat_gif	2	4	4	5
N training examples:	2	20	100	1000

Century (n=7)

	2	20	100	1000
+all 87 transforms	5	13	28	35%
+another 20 transforms	5	16	27	32
+next 20 transforms	5	18	26	29
+fft2_blur10	4	14	19	20
+colors_saturate_gif	4	14	18	20
+fft1_blur10	4	15	18	20
+stats_colorfulness_lab	3	15	18	20
+stats_contrastsd	3	11	14	15
+lines_hough50_gif	4	11	14	15
+color_chroma_divide_gif	3	8	11	13
+lines_color_conv_gif	4	8	11	13
+emboss_gray4_gif_x0.4	3	8	10	10
color_luminance_divide_gif	3	5	6	7
N training examples:	2	20	100	1000

- Looking inside the style classifier (all transforms, 1000 examples)
- Expressionism is hardest to classify
- Impressionism is the easiest
- Complexity profiles of Baroque, Rococo, Romanticism are often confusable

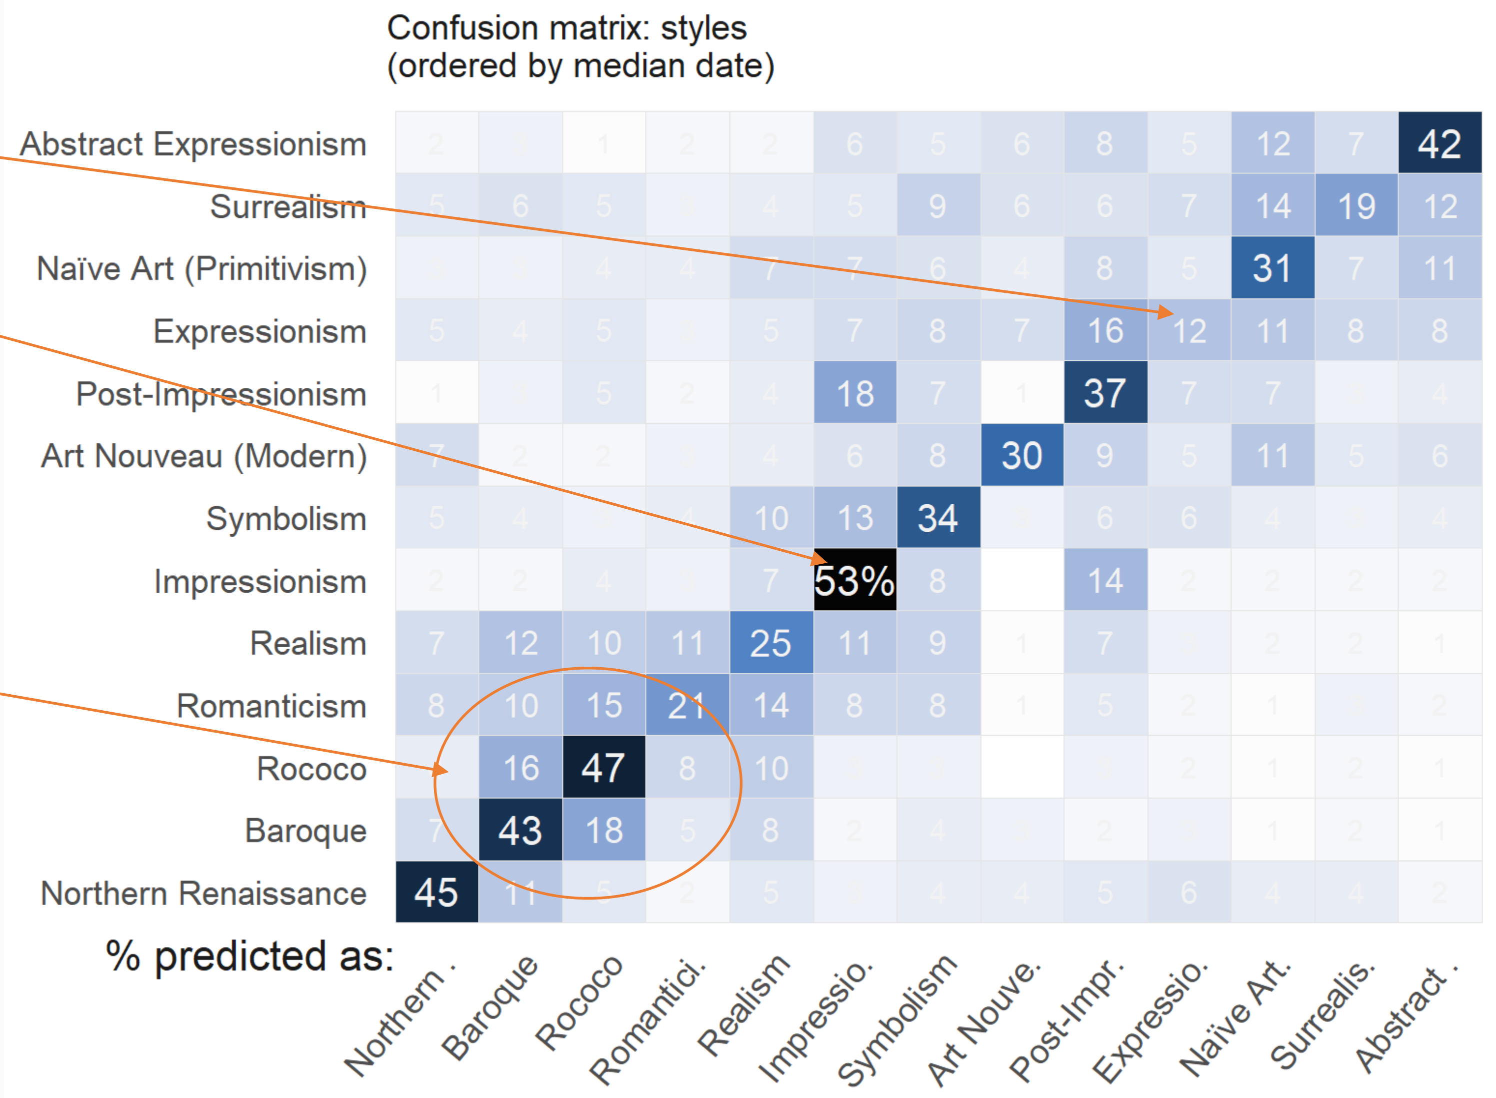

- Also, not surprisingly, distinguishing post-impressionism from impressionism is not always easy either
  - (note that all these labels come from the art500k/wikiart metadata, which is not itself an absolute authority on art history and style classification)

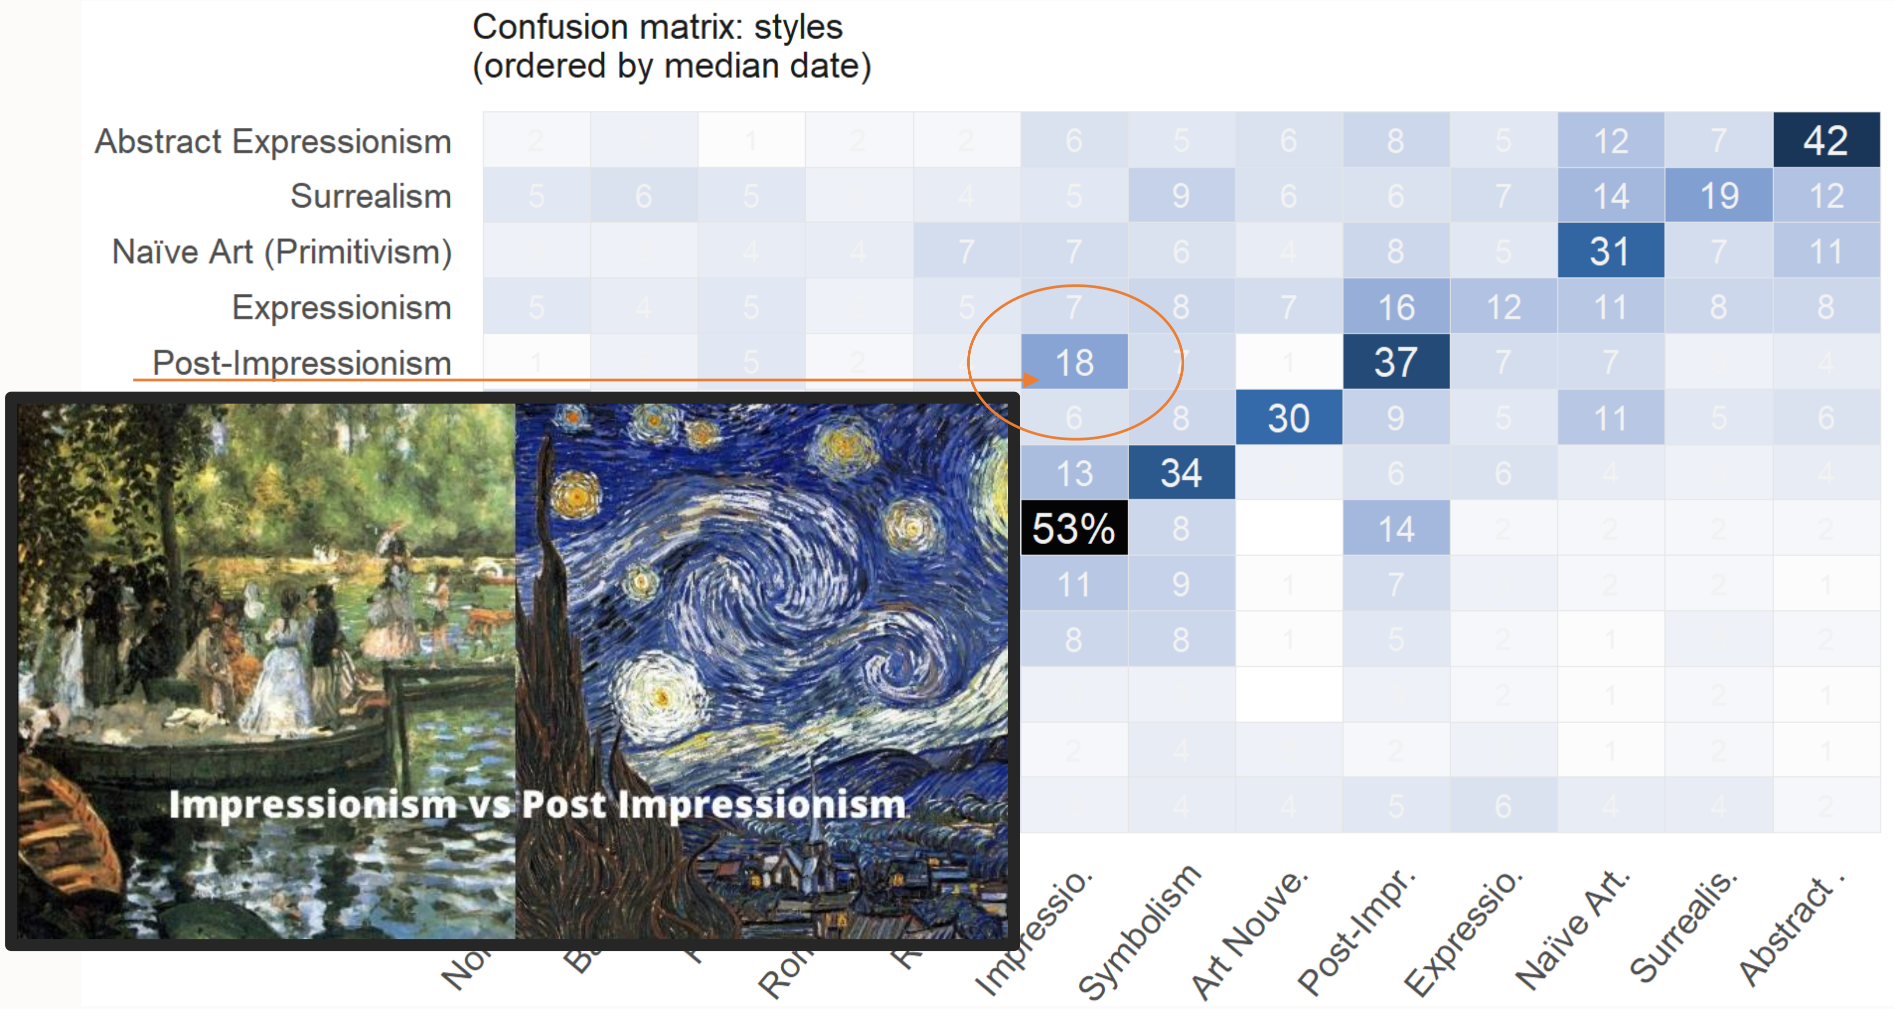

- Some artists are harder to distinguish than others
- E.g. Thomas Lawrence vs George Romney:

- Or William Hogarth vs Nicolas Poussin:

Partial confusion matrix: most often misclassified authors (ordered by median date)

% predicted as:	Nicolas P.	William H.	George Ro.	Thomas La.	Thomas Co.	Albert Bi.	Paul Klee	Paul Jaco.
Paul Jacoulet							20	56
Paul Klee		1					20	2
Albert Bierstadt	1				22	24		
Thomas Cole	1				67%	11		
Thomas Lawrence			22	52				
George Romney			49	30				
William Hogarth	17	32	1				1	
Nicolas Poussin	36	1				2	1	

- *depicted: top similar artworks for these artist pairs – note that this is just based on the compression vectors, no object detection or other machine learning here.

Explorations: art evolution

- Our 70k art sample has slowly shifted towards including more complex art over time (PC1), there's more variation in the pattern and colour complexity dimensions (PC2, PC3), and later centuries contain more concrete/contrast-wise less complex images (PC3)

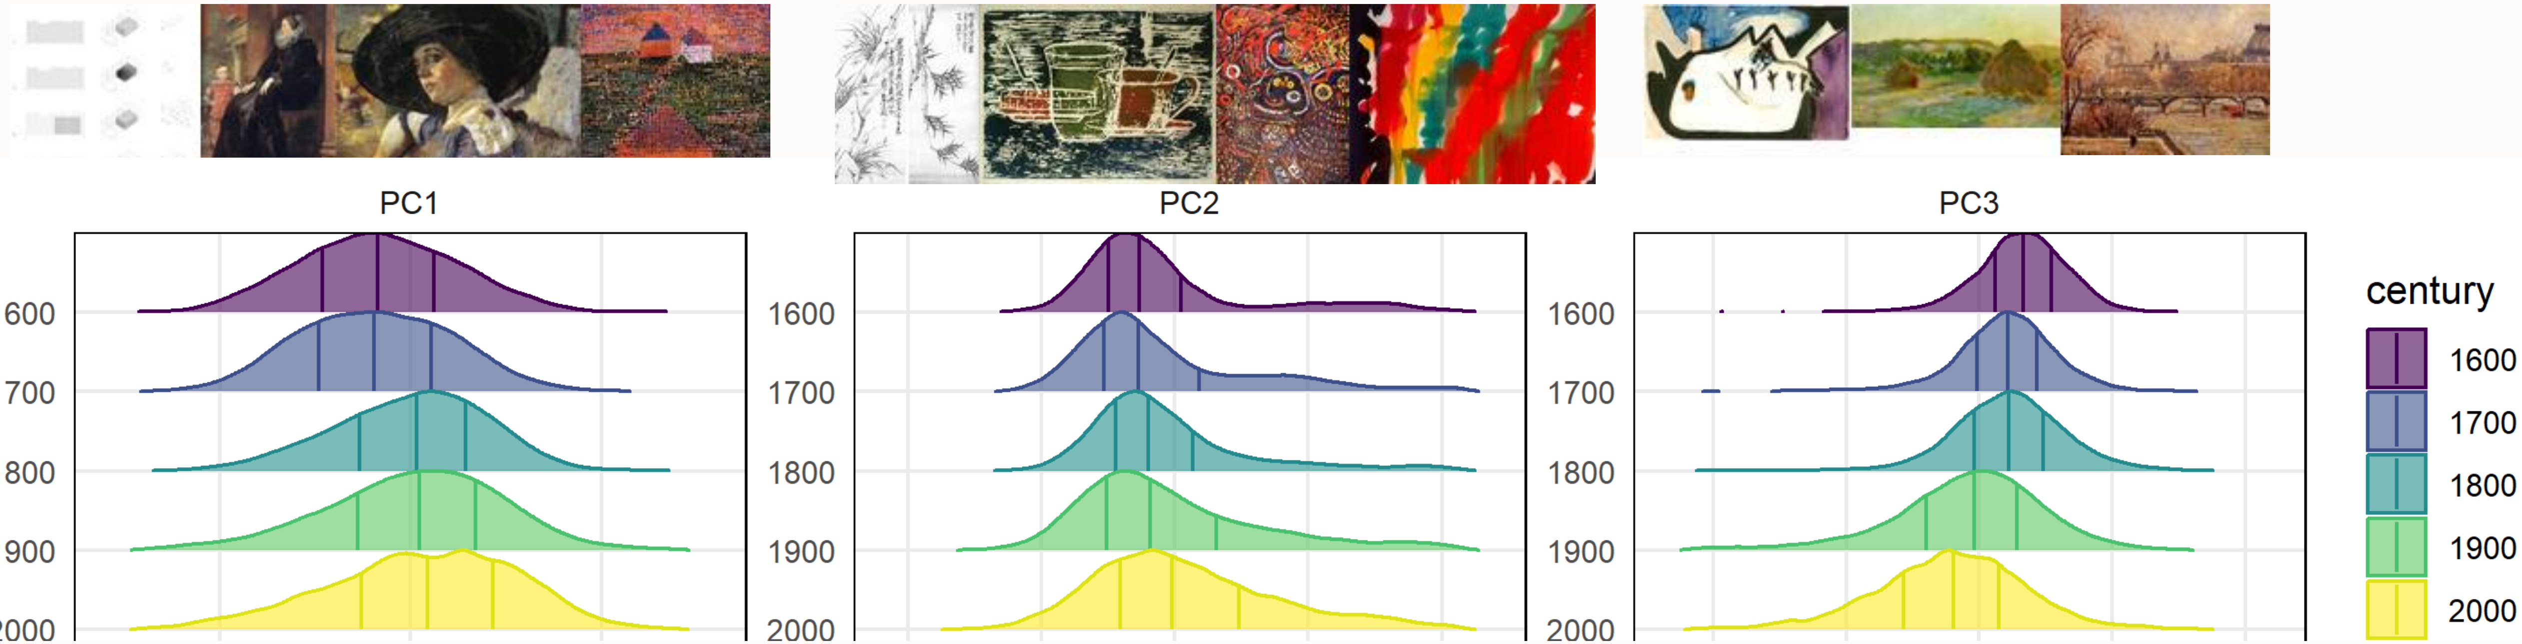

Explorations: artistic trajectories

- Towards quantifying artistic career trajectories and innovations
- We can observe how artists move through the complexity space over the course of their careers
- Allows asking questions like, what happens when you hit the boundary of what's considered art in your time period? Which artists play safe and who are ahead of their time?
- Example: innovating in the space is quantifiable as local density of a given artwork (Euclidean distance to n% of nearest artists' works), contrasting the density before its creation, and after (here, arbitrary 20 year window)

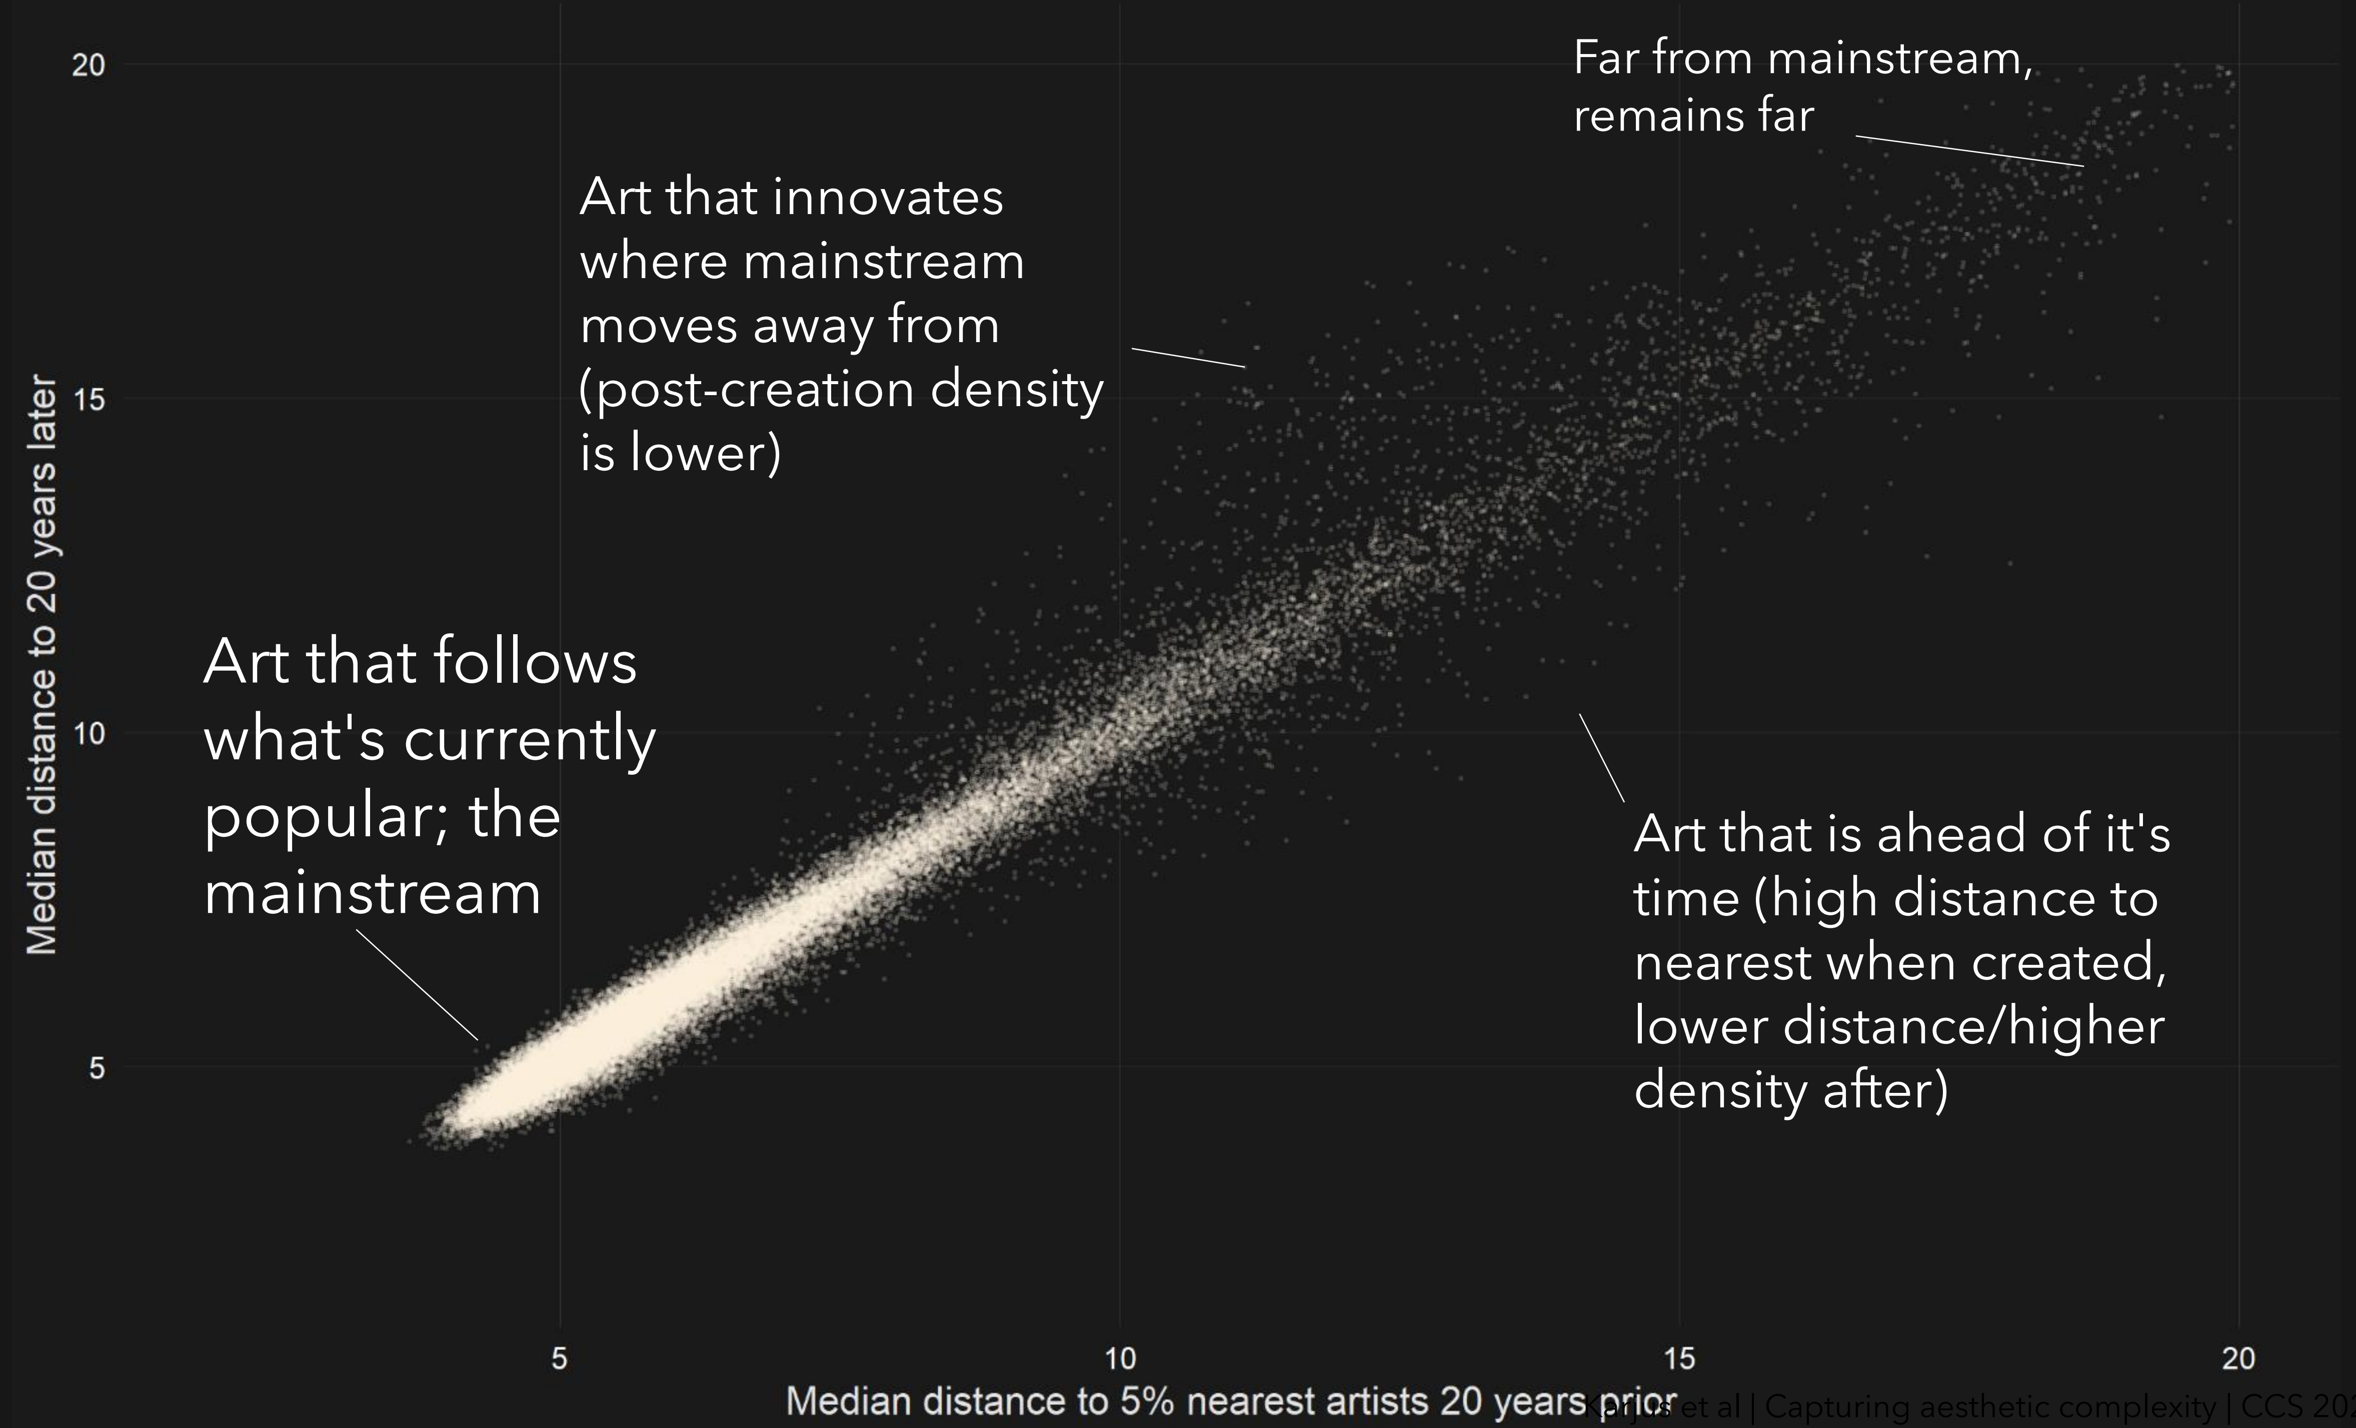
Far from mainstream, remains far
Art that innovates where mainstream moves away from (post-creation density is lower)
Art that follows what's currently popular; the mainstream
Art that is ahead of it's time (high distance to nearest when created, lower distance/higher density after)
Median distance to 20 years later
Median distance to 5% nearest artists 20 years prior
20
15
10
5
5
10
15
20

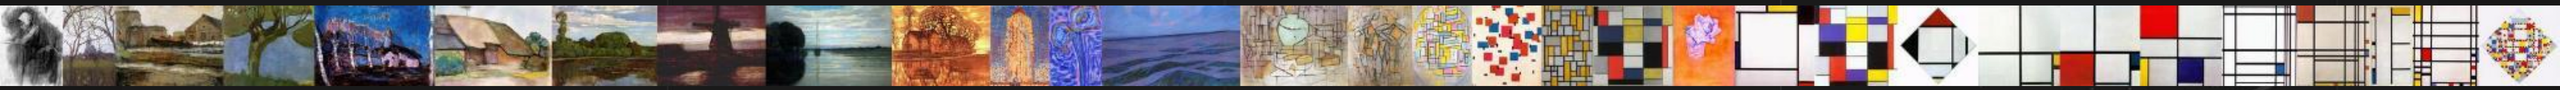

Piet Mondrian
(pink is early career, mainstream; yellow is later career, more experimental

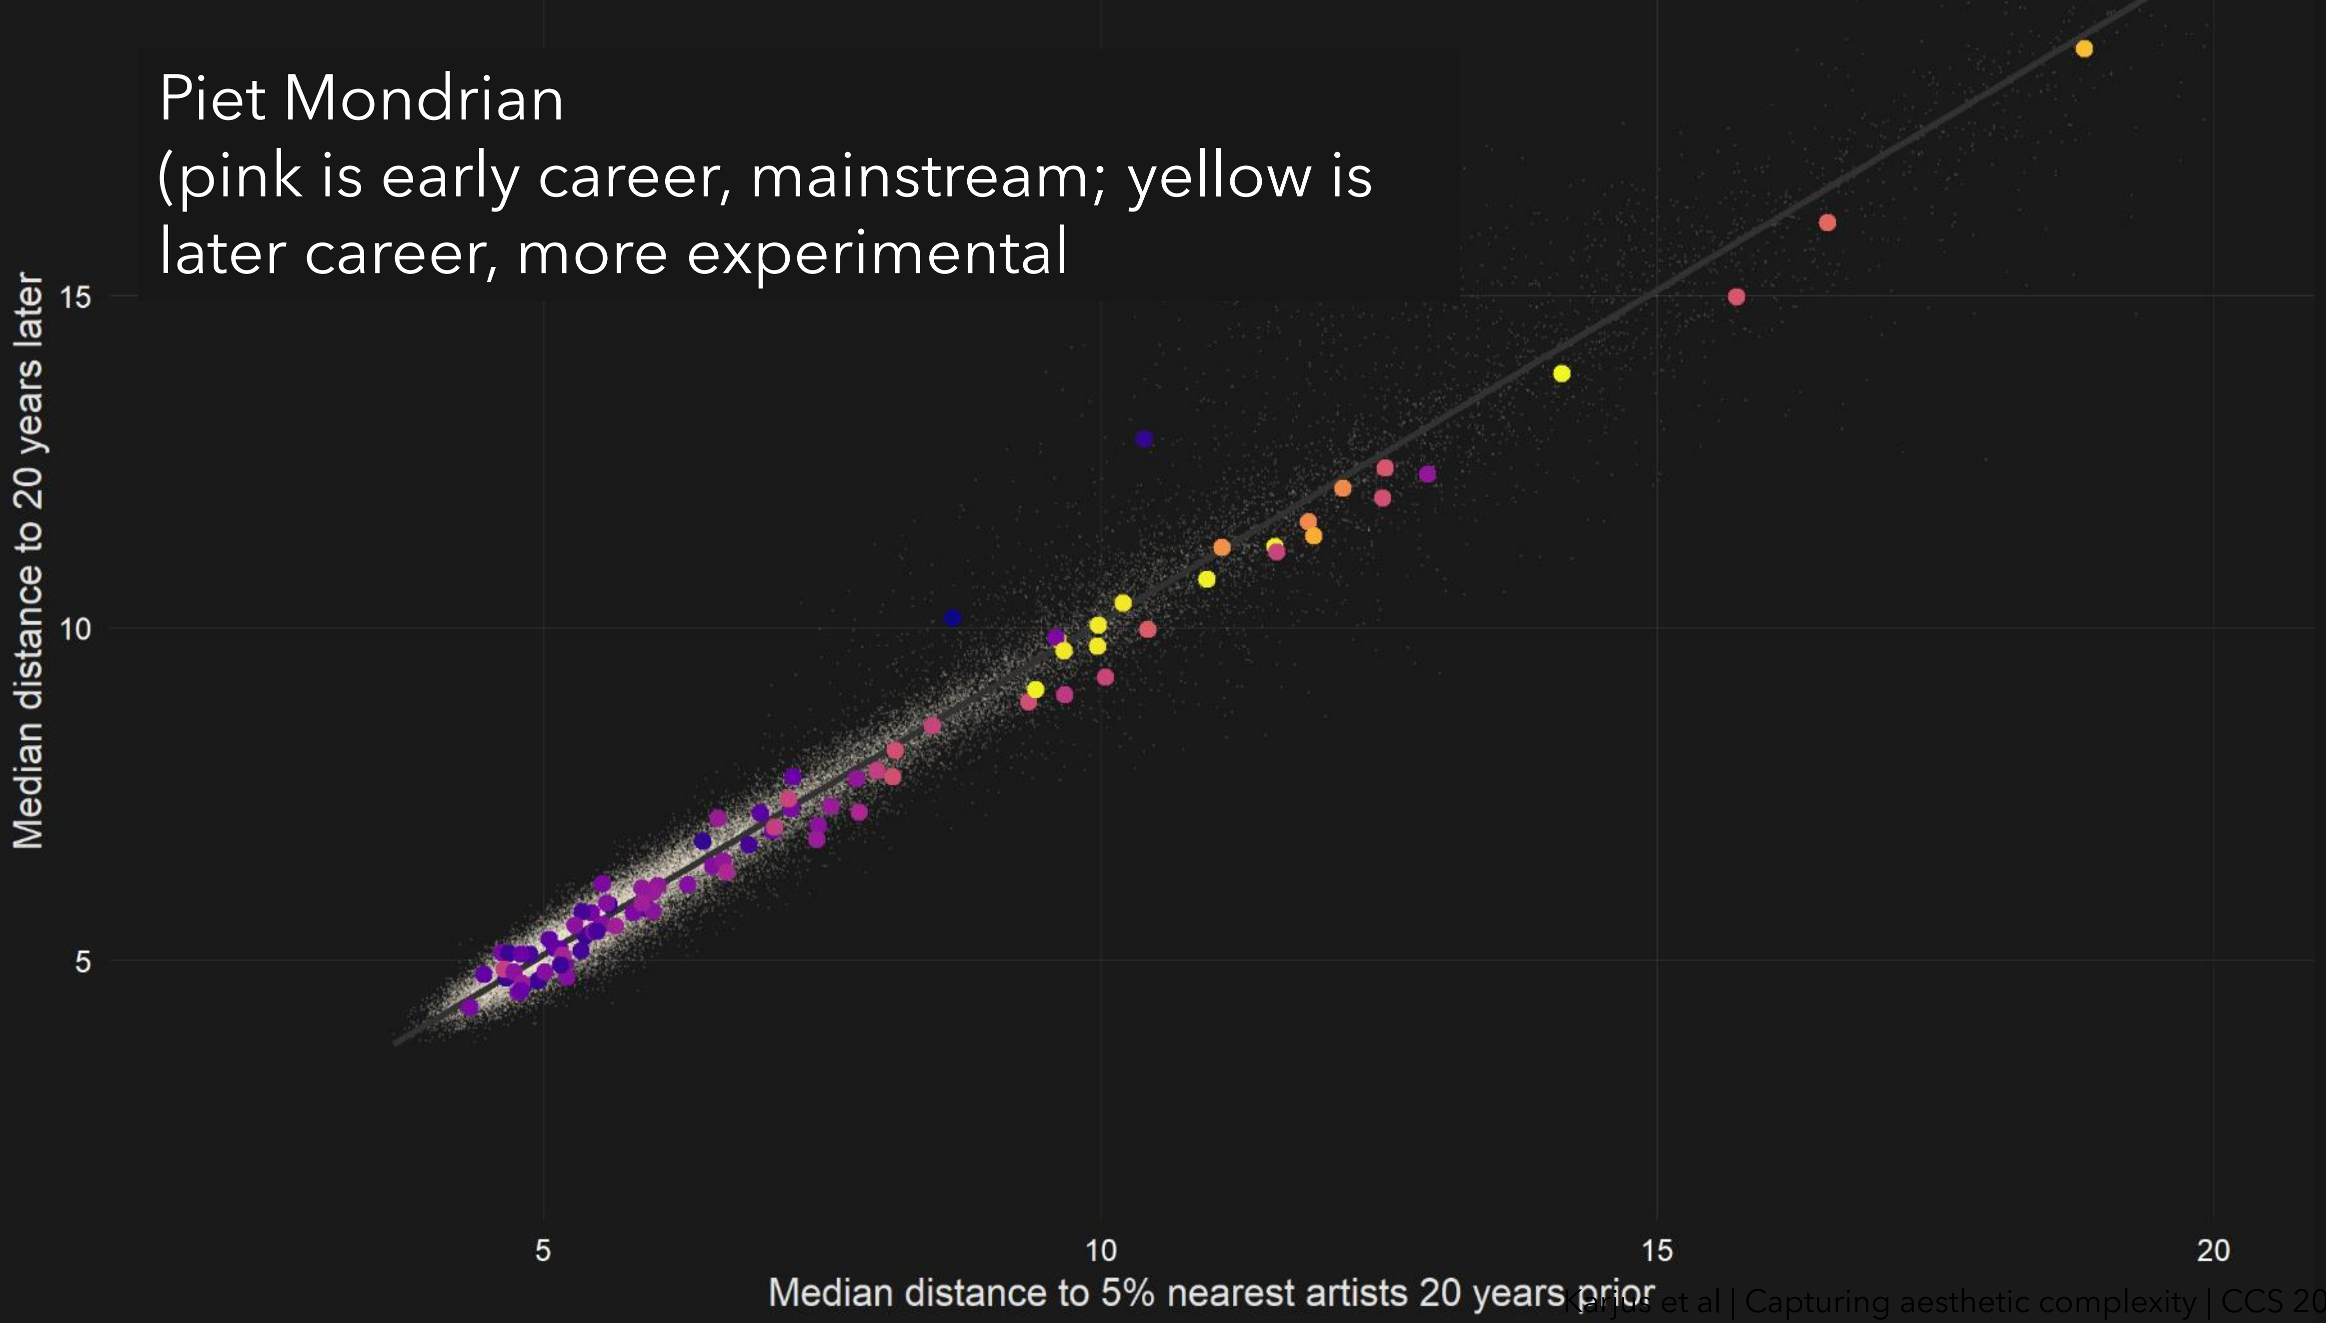

Georgia O'Keeffe
(career oscillates between being a bit closer to mainstream and more experimental)

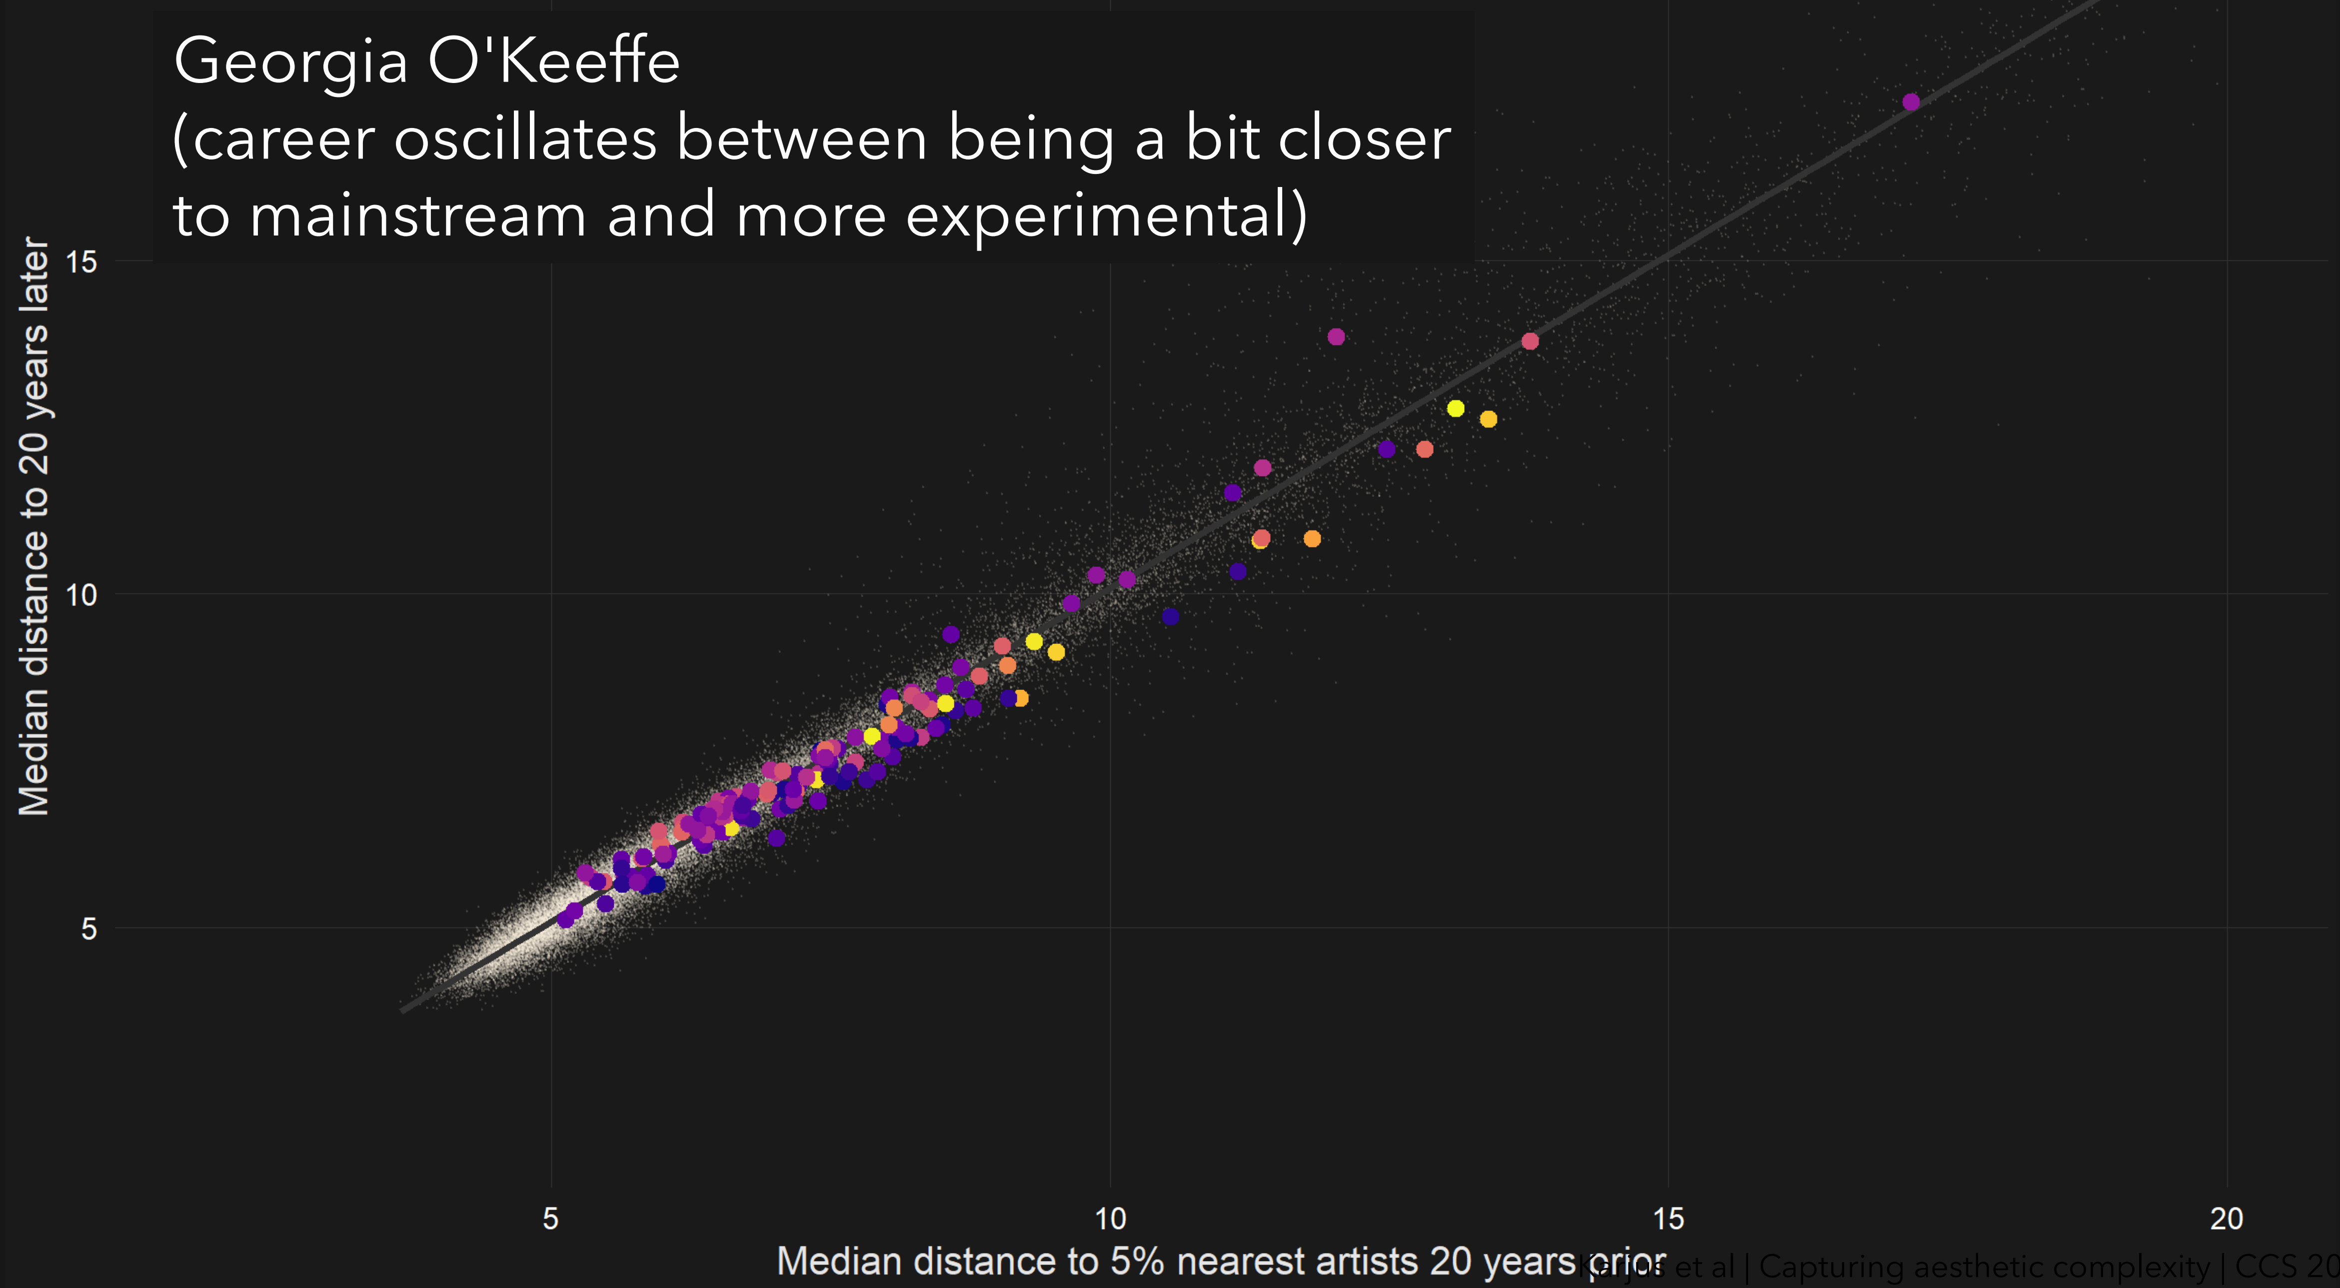

Work in progress: an online art complexity explorer app

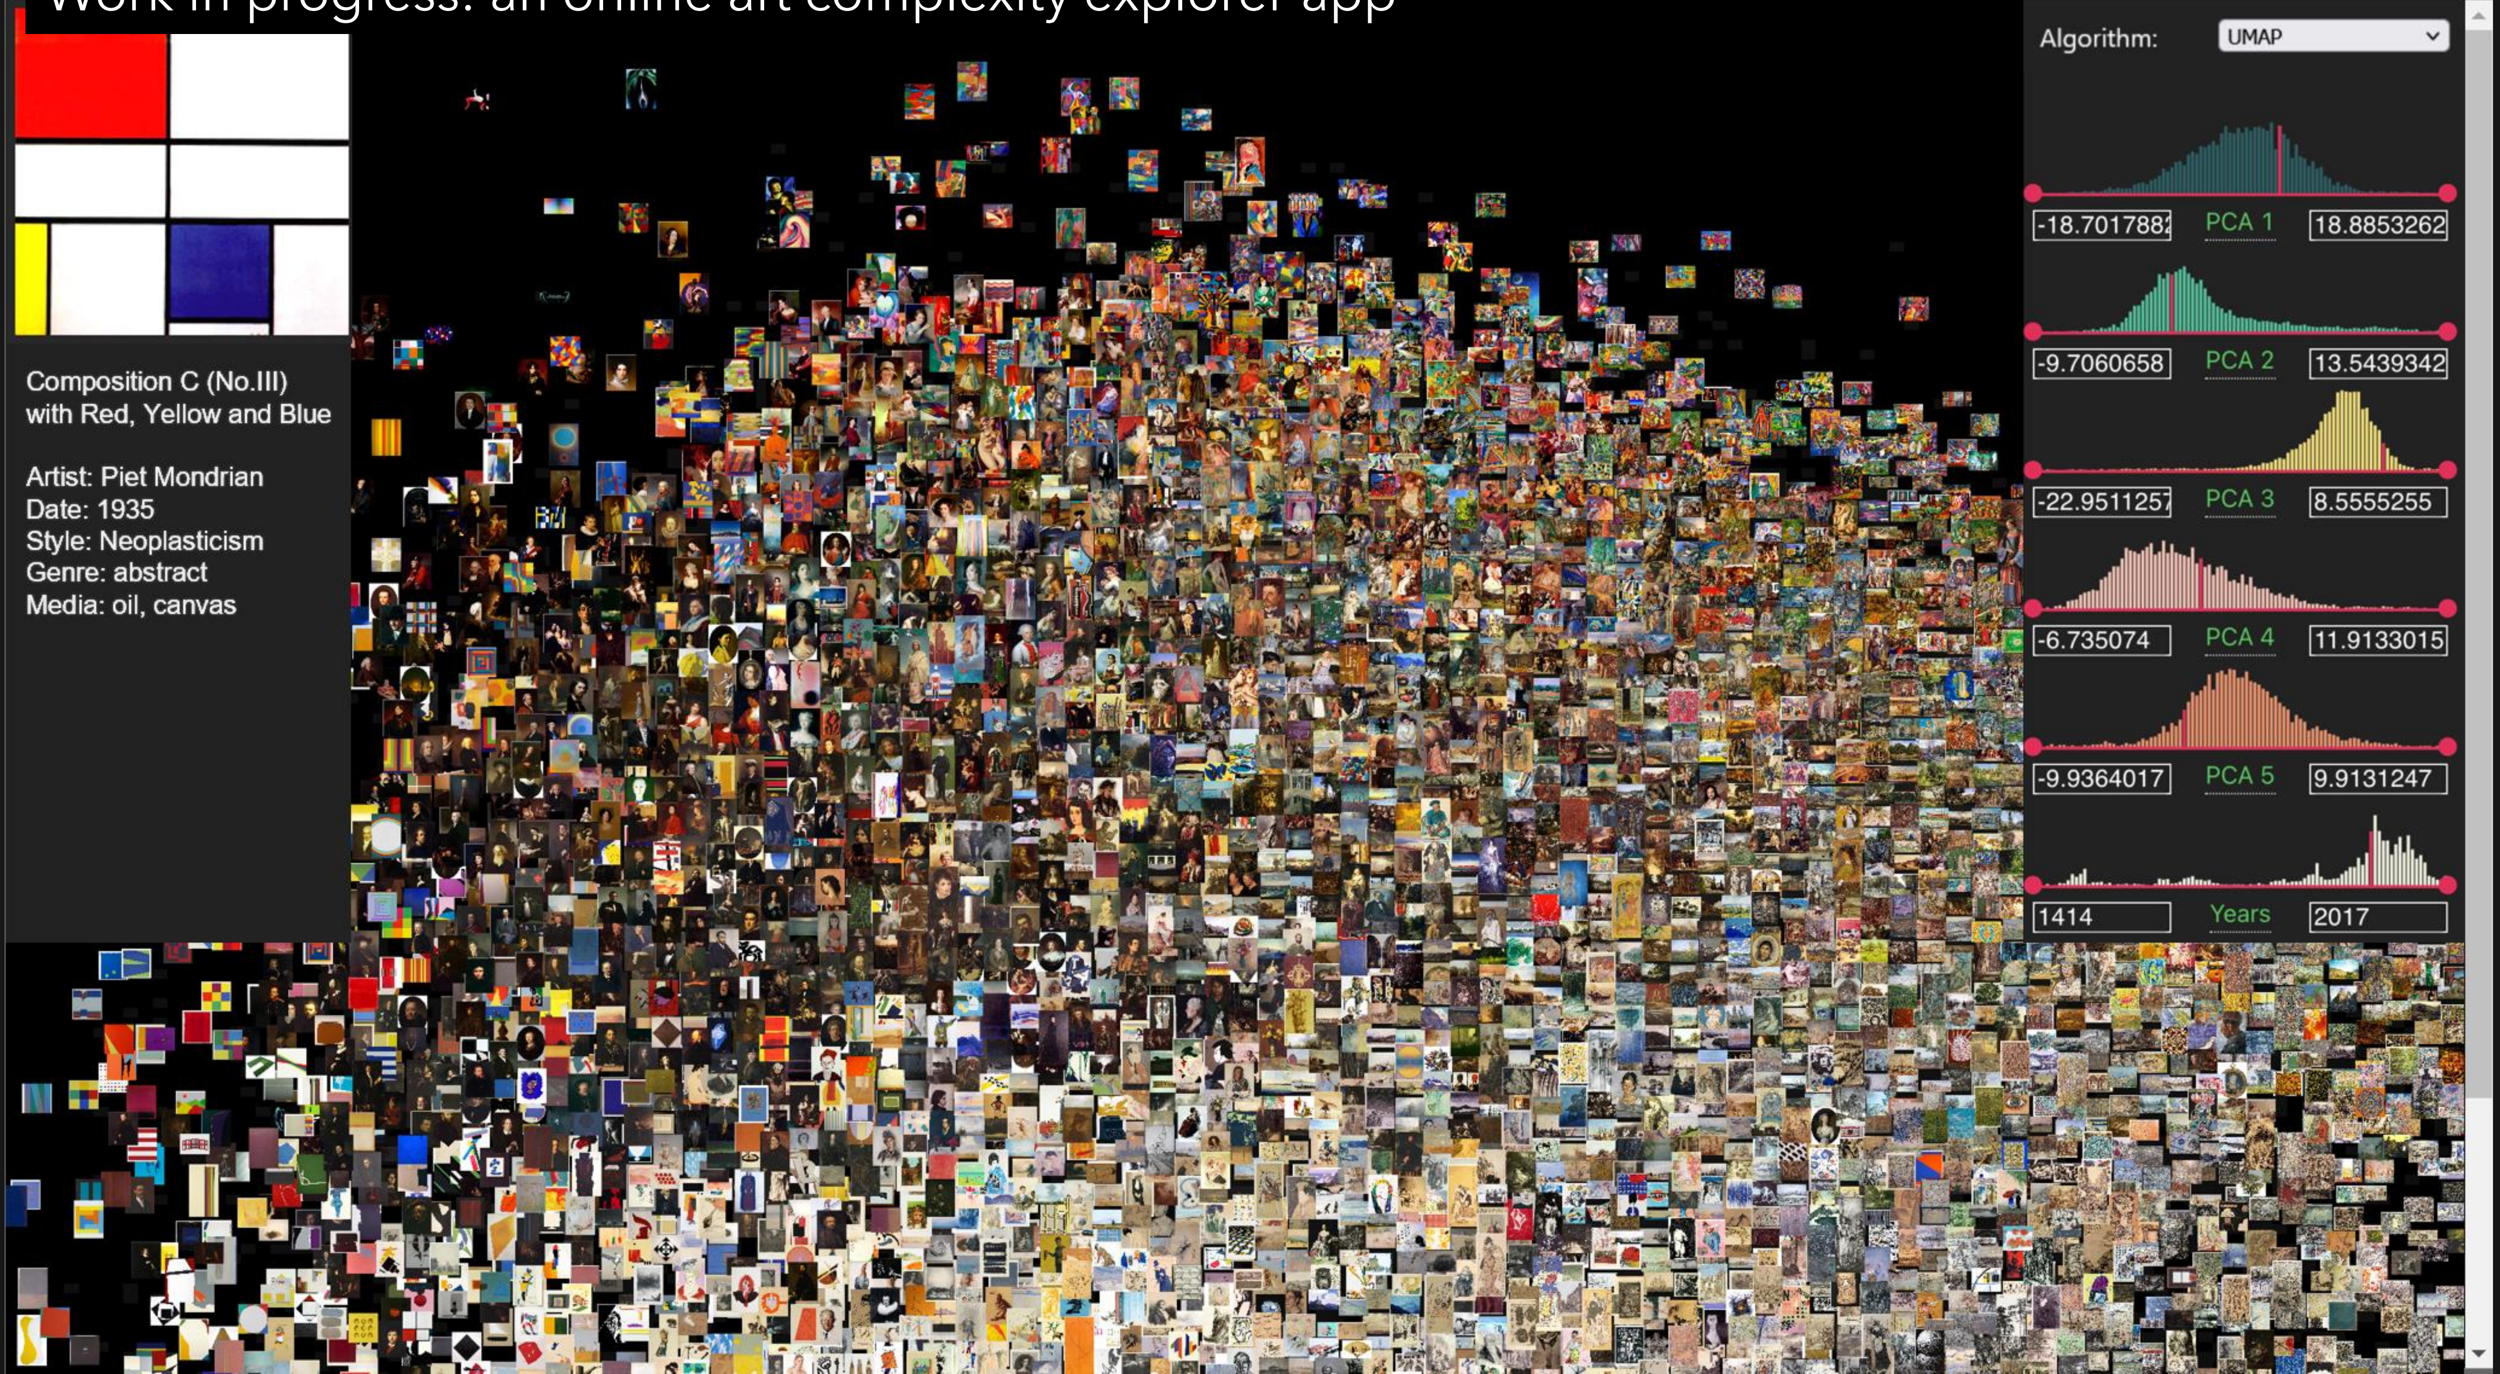

Conclusions

- A novel ensemble approach for quantifying visual complexity
- Cognitively plausible, outperforms the single-compression approach; also captures similarities and differences between artists and artworks
- Here art historical questions, but applicable to any visuals
- General ensemble approach should be applicable to any domain – instead of trying to find the best estimator, use all the estimators
- Future work: quantifying art evolution in the complexity space, comparing and reasoning about artists' trajectories; application to other domains. Complexity explorer app and a paper in the works, stay tuned!
- More questions? → twitter.com/AndresKarjus
- Slides and contact: andreskarjus.github.io

Supported by European Union Horizon 2020, Project No. 810961; see cudan.tlu.ee