

Computational Creativity Autumn School II

Building an Automated Painter

(Practical and Cultural Aspects)

Simon Colton

Computational Creativity Group
Department of Computing
Goldsmiths College, University of London

ccg.doc.gold.ac.uk

Check us out...

ccg.gold.ac.uk

- Computational Creativity Group at Goldsmiths
 - Two PhD students, six RAs, three academics, one administrator
 - We perform AI research, where we study fundamental notions of creativity in software via applications to:
 - Graphic design, mathematical discovery, the visual arts, video game design, creative language and fictional ideation



Welcome

To the home of the Goldsmiths Computational Creativity Research Group

Welcome to the Computational Creativity group in the Department of Computing at Goldsmiths College, University of London. We are one of the world's largest and most prominent AI research groups specialising in the computational modelling of creative processes. Via practical implementation and experimentation projects, we address the higher-level philosophical issues raised through software acting in autonomously creative ways. We are involved in a number of projects to bring greater formalism to the study of creativity in software, and are also well known for our applications in diverse domains including pure mathematics, graphic design, creative language, video game design and the visual arts.

Click image for details See gallery



Computational Creativity is a sub-area of Artificial Intelligence research, where we study how to build software which can take on some of the creative responsibility in arts and science projects. This has practical aspects, and we build creative software to generate artefacts such as poems, mathematical theories, board games, video games, abstract and representational art. We build software to be taken seriously as autonomously creative, as well as software which can act as a creative collaborator with scientists, designers, artists and engineers.

Computational Creativity also has theoretical and philosophical aspects, and we have investigated various issues. These include whether Turing-style tests are appropriate for evaluating creative systems, how software can account for its own process and products, and the latent heat problem of Computational Creativity: as software is given more creative license, often the value of the artefacts it produces decreases, which has implications for the assessment of creativity in software.

We are currently building a Computational Creativity Theory to study the assessment of creative software within a formal setting. This will take into account both how software works and what it produces. We hope to demystify the creative process in people and computers by developing a framework for describing creative systems, then building and applying software for creative tasks within this framework.

These web pages describe our research projects, results and ideas. It's also a place to check for updates about the work of Computational Creativity. The CCG research group leader is Simon Colton, and he has an overview page for his research which gives a flavour of the research that he is interested in.

Research Projects

We are involved in a number of funded projects with UK and EU partners. Please see our research pages for details of these projects. Here are some quick links to the project home pages:



Research

CCG research projects past, present and future

Overview

We are a group of Artificial Intelligence researchers with specific interests in the automation and evaluation of creative processes in software. We pursue both practical applications and theoretical developments, in addition to engaging in philosophical discourse about the issues raised when considering software which is meant to be autonomously creative.

Practical Applications

We build generative software which produces artefacts such as mathematical theories, video and board games, paintings, poems and designs. We aim to successively hand over creative responsibility to the input and control is minimised, and it becomes possible to project the work created onto the exhibits. In addition, we aim to enable our software to be accountable for its actions, by proving what it's done and why, and by describing what it has achieved and how this fits into the creation.

Formalising our Understanding of Creative Software

We have undertaken a number of formalism efforts, with the aim of designing a more systematic questions such as "Is software version B more creative than software version A?" As part of this project entitled "Computational Creativity Theory", we are building a formalism around the FAC the creative acts that software undertakes and the IDEA model, which describes the impact that may have on an ideal audience. We are currently working on the next stage of this formalism, approach to describe software processes, with particular emphasis on highlighting which creat attributed to the software and which to the programmer/user.

Raising and Addressing Philosophical Issues

We have often led the way in highlighting difficult issues which arise when considering the not independently creative. In particular, we have argued that comparison tests such as Turing-style building of naive programs which engage in pastiche. We have also brought to the fore the idea accountable for its creative acts, i.e., explain how and why it has produced a particular artefact behaviour in such a way to add value to its creations. We've also studied serendipity from a B and introduced the *creativity tripod*, arguing that software has to exhibit skill, appreciation and avoid being seen as uncreative.

Ongoing Projects

ANGELINA - Creative Code Generation for Automated Game Design

Written by Michael Cook. ANGELINA is software we are using to invite automated game design. We have been awarded an EPSRC grant entitled *Generation for Interactive Media*, and we will be investigating the possibility of using software to generate and modify its own program code, by extending code for videogames. Below are a couple of papers about ANGELINA webpages.

- Mechanic Miner: Reflection-Driven Game Mechanic Discovery
- Aesthetic Considerations for Automated Platformer Design

CADGames - Computer aided game design

Working with partners at Rebellion Developments Ltd., a well known in Oxford, we undertook research towards the next generation of artificial intelligence.

Computational Creativity Theory

Latest Tweets

Mark Reid

Welcome

To the home of the Goldsmiths Computational Creativity Research Group

Welcome to the Computational Creativity group in the Department of Computing at Goldsmiths College, University of London. We are one of the world's largest and most prominent AI research groups specialising in the computational modelling of creative processes. Via practical implementation and experimentation projects, we address the higher-level philosophical issues raised through software acting in autonomously creative ways. We are involved in a number of projects to bring greater formalism to the study of creativity in software, and are also well known for our applications in diverse domains including pure mathematics, graphic design, creative language, video game design and the visual arts.

Computational Creativity is a sub-area of Artificial Intelligence research, where we study how to build software which can take on some of the creative responsibility in arts and science projects. This has practical aspects, and we build creative software to generate artefacts such as poems, mathematical theories, board games, video games, abstract and representational art. We build software to be taken seriously as autonomously creative, as well as software which can act as a creative collaborator with scientists, designers, artists and engineers.

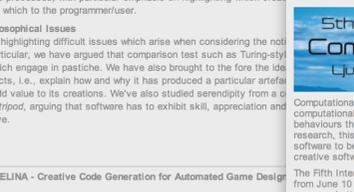
Computational Creativity also has theoretical and philosophical aspects, and we have investigated various issues. These include whether Turing-style tests are appropriate for evaluating creative systems, how software can account for its own process and products, and the latent heat problem of Computational Creativity: as software is given more creative license, often the value of the artefacts it produces decreases, which has implications for the assessment of creativity in software.

We are currently building a Computational Creativity Theory to study the assessment of creative software within a formal setting. This will take into account both how software works and what it produces. We hope to demystify the creative process in people and computers by developing a framework for describing creative systems, then building and applying software for creative tasks within this framework.

These web pages describe our research projects, results and ideas. It's also a place to check for updates about the work of Computational Creativity. The CCG research group leader is Simon Colton, and he has an overview page for his research which gives a flavour of the research that he is interested in.

Research Projects

We are involved in a number of funded projects with UK and EU partners. Please see our research pages for details of these projects. Here are some quick links to the project home pages:



Welcome

To the home of the Goldsmiths Computational Creativity Research Group

Welcome to the Computational Creativity group in the Department of Computing at Goldsmiths College, University of London. We are one of the world's largest and most prominent AI research groups specialising in the computational modelling of creative processes. Via practical implementation and experimentation projects, we address the higher-level philosophical issues raised through software acting in autonomously creative ways. We are involved in a number of projects to bring greater formalism to the study of creativity in software, and are also well known for our applications in diverse domains including pure mathematics, graphic design, creative language, video game design and the visual arts.

Computational Creativity is a sub-area of Artificial Intelligence research, where we study how to build software which can take on some of the creative responsibility in arts and science projects. This has practical aspects, and we build creative software to generate artefacts such as poems, mathematical theories, board games, video games, abstract and representational art. We build software to be taken seriously as autonomously creative, as well as software which can act as a creative collaborator with scientists, designers, artists and engineers.

Computational Creativity also has theoretical and philosophical aspects, and we have investigated various issues. These include whether Turing-style tests are appropriate for evaluating creative systems, how software can account for its own process and products, and the latent heat problem of Computational Creativity: as software is given more creative license, often the value of the artefacts it produces decreases, which has implications for the assessment of creativity in software.

We are currently building a Computational Creativity Theory to study the assessment of creative software within a formal setting. This will take into account both how software works and what it produces. We hope to demystify the creative process in people and computers by developing a framework for describing creative systems, then building and applying software for creative tasks within this framework.

These web pages describe our research projects, results and ideas. It's also a place to check for updates about the work of Computational Creativity. The CCG research group leader is Simon Colton, and he has an overview page for his research which gives a flavour of the research that he is interested in.

Research Projects

We are involved in a number of funded projects with UK and EU partners. Please see our research pages for details of these projects. Here are some quick links to the project home pages:



Please visit: ccg.doc.gold.ac.uk

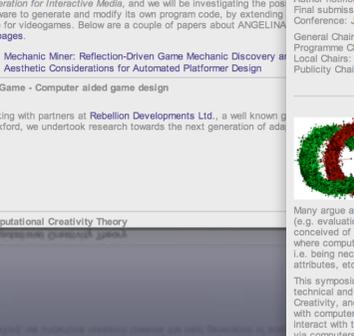
Selected Overview Papers, Talks and Articles



Tweets



Selected Overview Papers, Talks and Articles



Welcome

To the home of the Goldsmiths Computational Creativity Research Group

Welcome to the Computational Creativity group in the Department of Computing at Goldsmiths College, University of London. We are one of the world's largest and most prominent AI research groups specialising in the computational modelling of creative processes. Via practical implementation and experimentation projects, we address the higher-level philosophical issues raised through software acting in autonomously creative ways. We are involved in a number of projects to bring greater formalism to the study of creativity in software, and are also well known for our applications in diverse domains including pure mathematics, graphic design, creative language, video game design and the visual arts.

Computational Creativity is a sub-area of Artificial Intelligence research, where we study how to build software which can take on some of the creative responsibility in arts and science projects. This has practical aspects, and we build creative software to generate artefacts such as poems, mathematical theories, board games, video games, abstract and representational art. We build software to be taken seriously as autonomously creative, as well as software which can act as a creative collaborator with scientists, designers, artists and engineers.

Computational Creativity also has theoretical and philosophical aspects, and we have investigated various issues. These include whether Turing-style tests are appropriate for evaluating creative systems, how software can account for its own process and products, and the latent heat problem of Computational Creativity: as software is given more creative license, often the value of the artefacts it produces decreases, which has implications for the assessment of creativity in software.

We are currently building a Computational Creativity Theory to study the assessment of creative software within a formal setting. This will take into account both how software works and what it produces. We hope to demystify the creative process in people and computers by developing a framework for describing creative systems, then building and applying software for creative tasks within this framework.

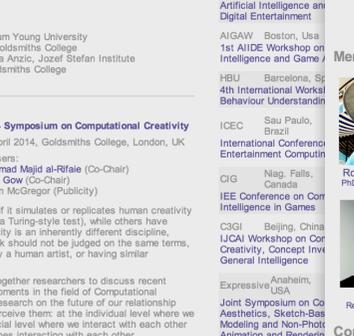
These web pages describe our research projects, results and ideas. It's also a place to check for updates about the work of Computational Creativity. The CCG research group leader is Simon Colton, and he has an overview page for his research which gives a flavour of the research that he is interested in.

Research Projects

We are involved in a number of funded projects with UK and EU partners. Please see our research pages for details of these projects. Here are some quick links to the project home pages:



Selected Overview Papers, Talks and Articles



Welcome

To the home of the Goldsmiths Computational Creativity Research Group

Welcome to the Computational Creativity group in the Department of Computing at Goldsmiths College, University of London. We are one of the world's largest and most prominent AI research groups specialising in the computational modelling of creative processes. Via practical implementation and experimentation projects, we address the higher-level philosophical issues raised through software acting in autonomously creative ways. We are involved in a number of projects to bring greater formalism to the study of creativity in software, and are also well known for our applications in diverse domains including pure mathematics, graphic design, creative language, video game design and the visual arts.

Computational Creativity is a sub-area of Artificial Intelligence research, where we study how to build software which can take on some of the creative responsibility in arts and science projects. This has practical aspects, and we build creative software to generate artefacts such as poems, mathematical theories, board games, video games, abstract and representational art. We build software to be taken seriously as autonomously creative, as well as software which can act as a creative collaborator with scientists, designers, artists and engineers.

Computational Creativity also has theoretical and philosophical aspects, and we have investigated various issues. These include whether Turing-style tests are appropriate for evaluating creative systems, how software can account for its own process and products, and the latent heat problem of Computational Creativity: as software is given more creative license, often the value of the artefacts it produces decreases, which has implications for the assessment of creativity in software.

We are currently building a Computational Creativity Theory to study the assessment of creative software within a formal setting. This will take into account both how software works and what it produces. We hope to demystify the creative process in people and computers by developing a framework for describing creative systems, then building and applying software for creative tasks within this framework.

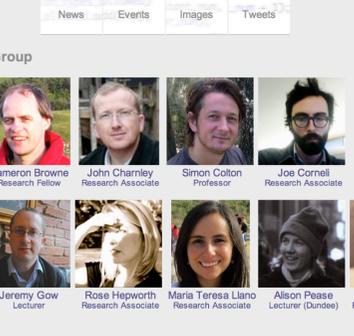
These web pages describe our research projects, results and ideas. It's also a place to check for updates about the work of Computational Creativity. The CCG research group leader is Simon Colton, and he has an overview page for his research which gives a flavour of the research that he is interested in.

Research Projects

We are involved in a number of funded projects with UK and EU partners. Please see our research pages for details of these projects. Here are some quick links to the project home pages:



Selected Overview Papers, Talks and Articles



Twitter:
@GoldsmithsCCG
#goldccg

1st AIIDE Workshop on AI and Game Aesthetics
October 14th, 2013
Northeastern University, Boston, USA
Organisers:
Antonios Liapis, Michael Cook, Cameron Browne

Joint Symposium on Computational Creativity
1 - 4 April 2014, Goldsmiths College, London, UK
Organisers:
Mohammad Majid al-Rifaie (Co-Chair)
Jeremy Gow (Co-Chair)
Stephen McGregor (Publicity)

AIIDE Conference on Artificial Intelligence and Digital Entertainment
Boston, USA
General Chair: Dan Ventura, Brigham Young University
Programme Chair: Simon Colton, Goldsmiths College
Local Chairs: Nade Lavrac and Tina Azizic, Jozsef Stefan Institute
Publicity Chair: Michael Cook, Goldsmiths College

Current events we're helping to organise



Computational Creativity is the art, science, philosophy and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative. As a field of research, this area is thriving, with progress in formalising what it means for software to be creative, along with many exciting and valuable applications of creative software in the sciences, the arts, literature, gaming and elsewhere.

The Fifth International Conference on Computational Creativity will be held from June 10 to 13, 2014 in Ljubljana, Slovenia. Please consider submitting a paper and attending what promises to be a very interesting event.

Submissions: Jan. 31, 2014
Author notification: Mar. 15, 2014
Final submissions: Apr. 15, 2014
Conference: June 10-13, 2014

General Chair: Dan Ventura, Brigham Young University
Programme Chair: Simon Colton, Goldsmiths College
Local Chairs: Nade Lavrac and Tina Azizic, Jozsef Stefan Institute
Publicity Chair: Michael Cook, Goldsmiths College

AISB14 Symposium on Computational Creativity
1 - 4 April 2014, Goldsmiths College, London, UK
Organisers:
Mohammad Majid al-Rifaie (Co-Chair)
Jeremy Gow (Co-Chair)
Stephen McGregor (Publicity)

Many argue a machine is creative if it simulates or replicates human creativity (e.g. evaluation of AI systems via a Turing-style test), while others have conceived of computational creativity as an inherently different discipline, where computer generated (art)work should not be judged on the same terms, i.e. being necessarily producible by a human artist, or having similar attributes, etc.

This symposium aims at bringing together researchers to discuss recent technical and philosophical developments in the field of Computational Creativity, and the impact of this research on the future of our relationship with computers and the way we perceive them: at the individual level where we interact with the machines, the social level where we interact with each other via computers, or even with machines interacting with each other.

Other events...

ICCC Ljubljana, Slovenia	10-13 June 2014
International Conference on Computational Creativity	
EvoMusArt, Granada, Spain	23-25 Apr. 2014
International Conference on Evolutionary and Biologically Inspired Art and Design	

AIGE London, UK	An exhibition of art and computation
KICSS Krakow, Poland	8th International Conference on Knowledge, Information Support Systems
AIIDE Boston, USA	Artificial Intelligence and Digital Entertainment
AIGAW Boston, USA	1st AIIDE Workshop on Intelligence and Game Aesthetics
HBU Barcelona, Spain	4th International Workshop on Behaviour Understanding
ICEC Sao Paulo, Brazil	International Conference on Entertainment Computing
CIG Niagara Falls, Canada	Intelligence in Games
IEE Conference on Cognitive Intelligence in Games	
CSGI Beijing, China	UCAI Workshop on Computational Creativity and General Intelligence
Expressive Anahaim, USA	
AAAI Bellevue, USA	The Twenty-Seventh AAAI Artificial Intelligence
DETA Philadelphia, USA	Entertainment Technology
CIEM Bath, UK	

About us

Who we are, what we do, contact details



Members of the Group



Contact Us

Simon Colton is the head of the group, and can be contacted at: s.colton@gold.ac.uk
We are based in the Department of Computing on the Goldsmiths campus in New Cross, London:



About us

Who we are, what we do, contact details

News	Events	Images	Tweets
------	--------	--------	--------

Members of the Group



Robin Baumgarten
PhD Student (Imperial)



Cameron Browne
Research Fellow



John Charnley
Research Associate



Simon Colton
Professor



Joe Corneli
Research Associate



Michael Cook
Research Associate



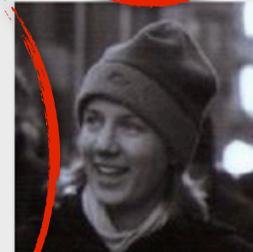
Jeremy Gow
Lecturer



Rose Hepworth
Research Associate



Maria Teresa Llano
Research Associate



Alison Pease
Lecturer (Dundee)



Ramin Ramezani
PhD student (Imperial)

Contact Us

Simon Colton is the head of the group, and can be contacted at: s.colton@gold.ac.uk

We are based in the **Department of Computing** on the Goldsmiths campus in New Cross, London:



Aims Today

- To give you an introduction to The Painting Fool project and to automated painting in general
- To provide some technical details about how to get software to creatively produce images
- To show you some of the artworks produced by the software some of and the interesting sub-projects
- To make you aware of some cultural issues which are raised when you talk about creative software in artistic circles

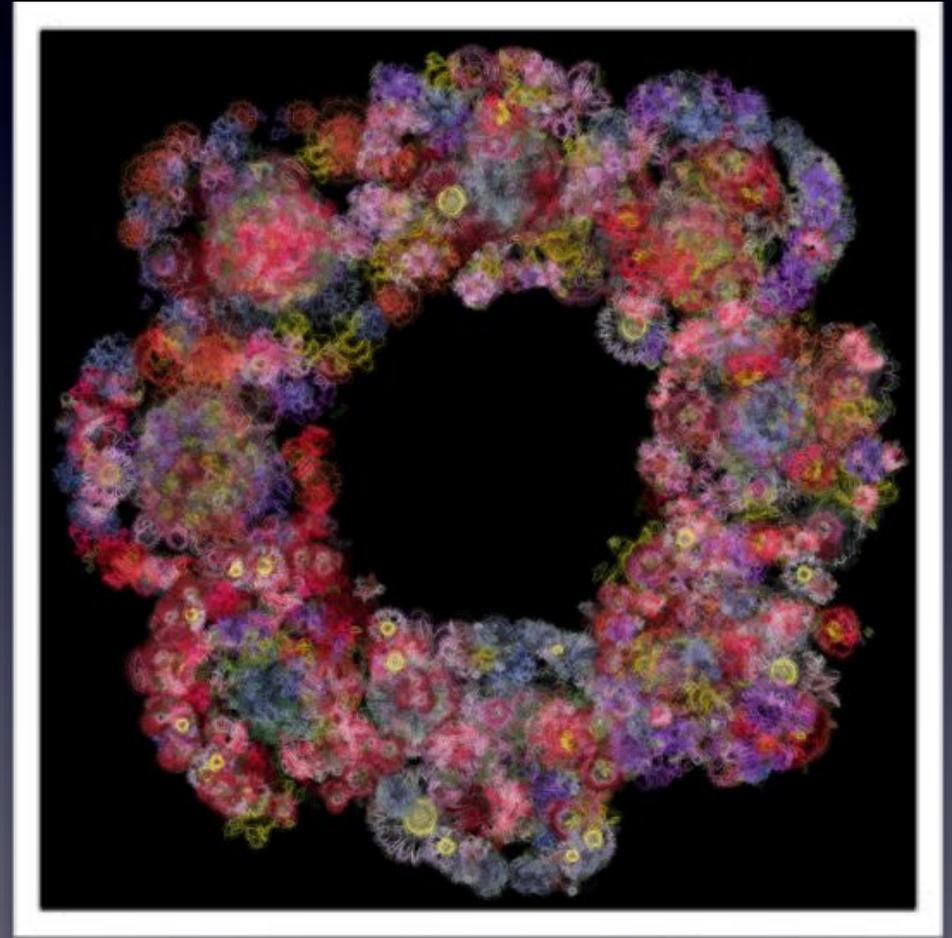
These will be my conclusions...

- Never forget that painting is a cognitive process and should be automated from an AI perspective
- You will have to combine dozens of approaches in order to get something similar to a painter
- You can't build an automated painter without addressing many of the philosophical issues of the field
- You will have to manage the public perception of creativity in software as much as you will have to advance technically... work with journalists if you can
- It's a good idea to exhibit the artwork produce, in order to get feedback and change opinions

The Painting Fool

Project Overview

- Aim: to build a system which is - one day - taken seriously as a creative artist in its own right.
- Involves graphics, NLP, AI and vision: painting is very much a cognitive process
- Around 10 years of work (a dozen years of software)

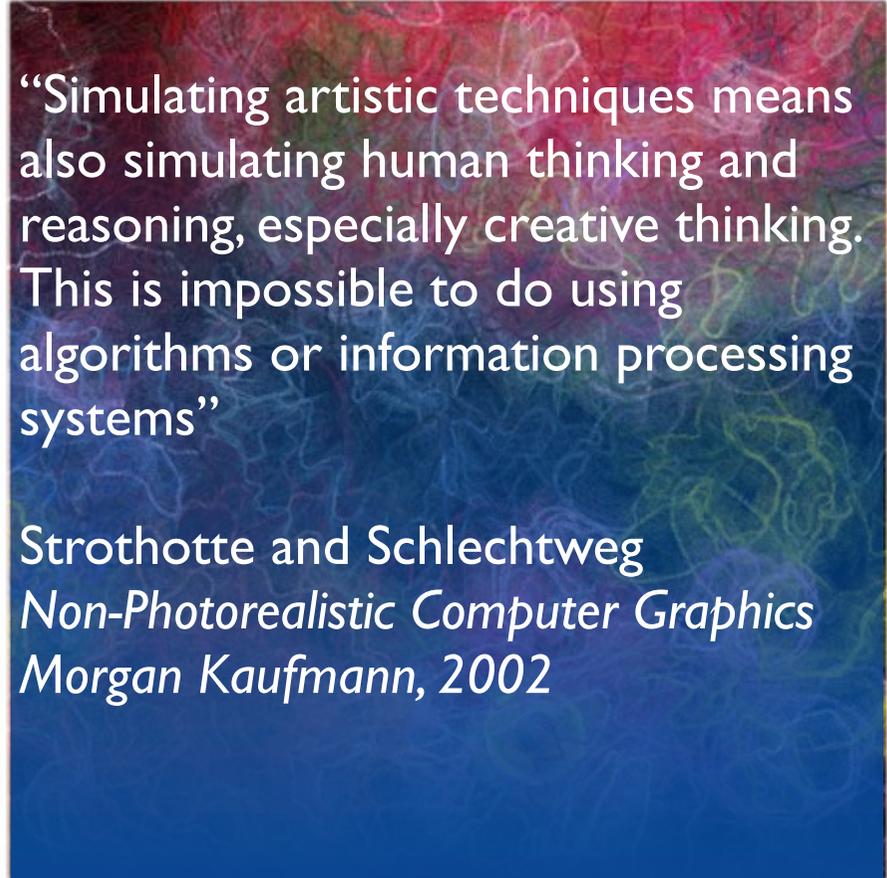


www.thepaintingfool.com

The Painting Fool

Project Overview

- Aim: to build a system which is - *one day* - taken seriously as a creative artist in its own right.
- Involves graphics, NLP, AI and vision: painting is very much a cognitive process
- Around 10 years of work (a dozen years of software)



“Simulating artistic techniques means also simulating human thinking and reasoning, especially creative thinking. This is impossible to do using algorithms or information processing systems”

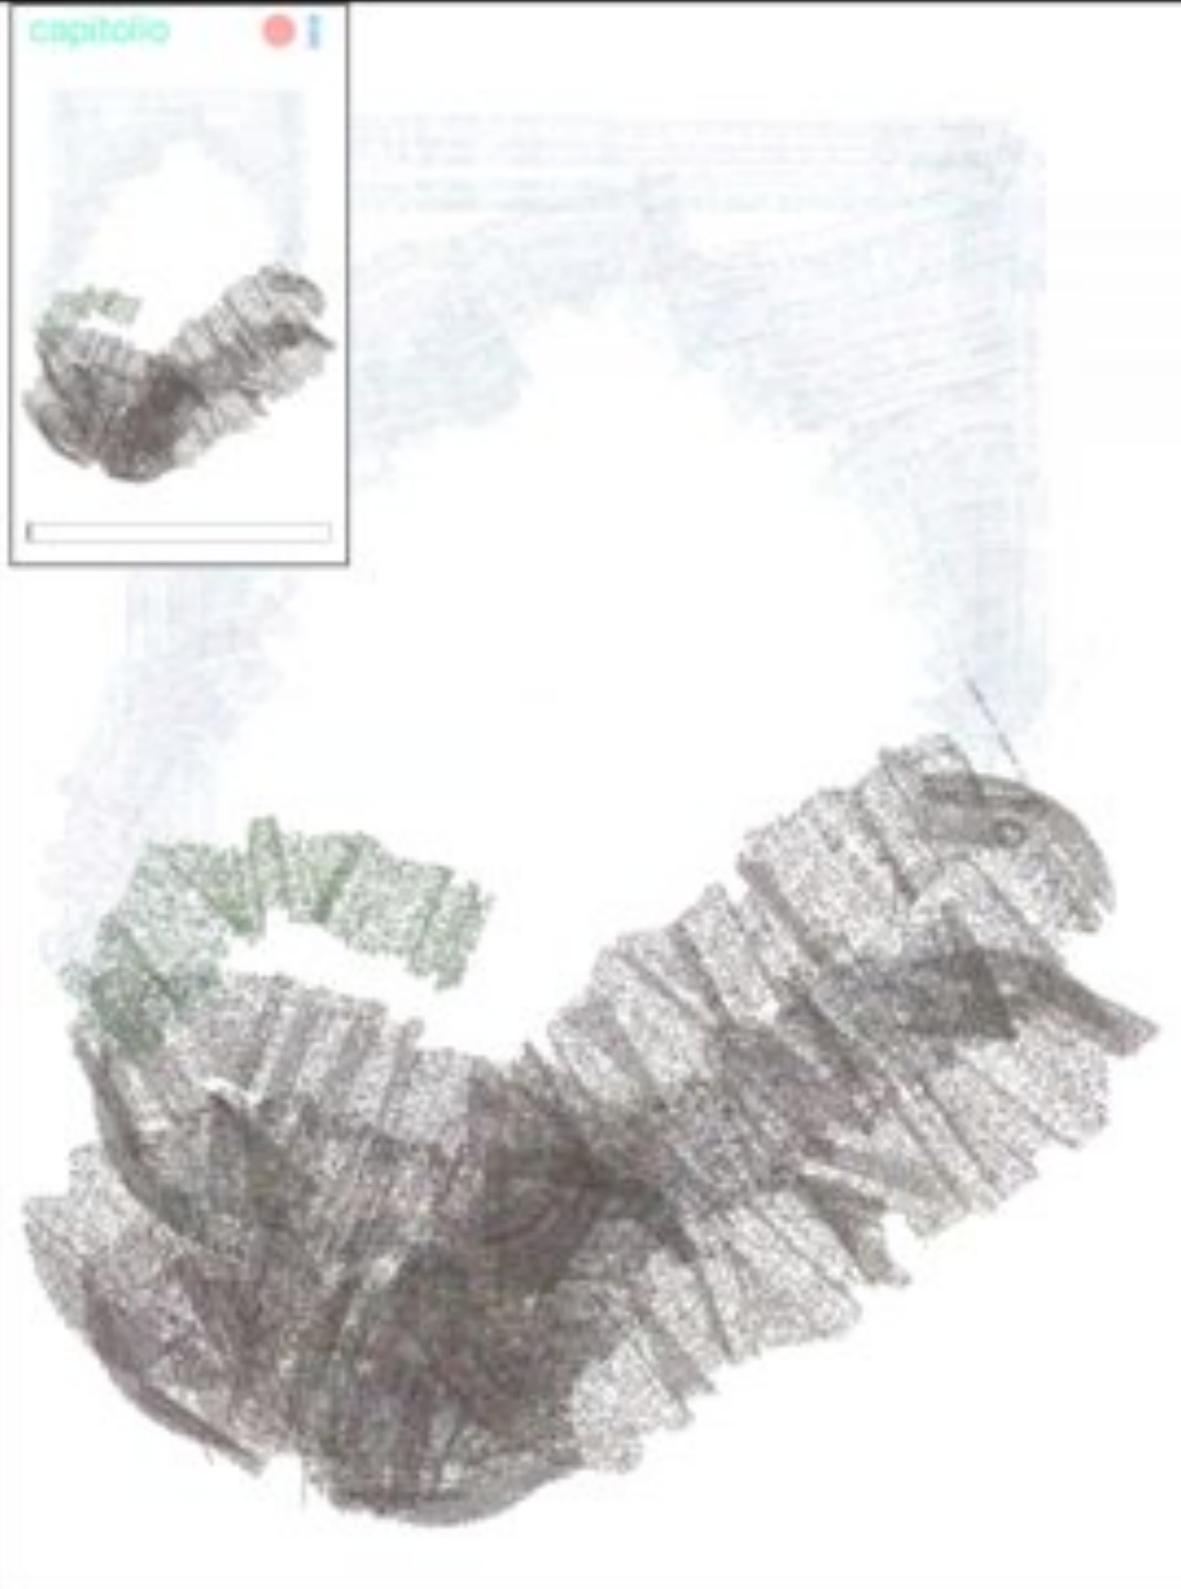
Strothotte and Schlechtweg
Non-Photorealistic Computer Graphics
Morgan Kaufmann, 2002

The Painting Fool

Stages of Development

- Non-photorealistic rendering (skill)
- Emotional modelling (appreciation)
- Generative art (imagination)
- Multimedia art (skill and imagination)
- Live portraiture (skill, appreciation, imagination, intentionality, learning and reflection)

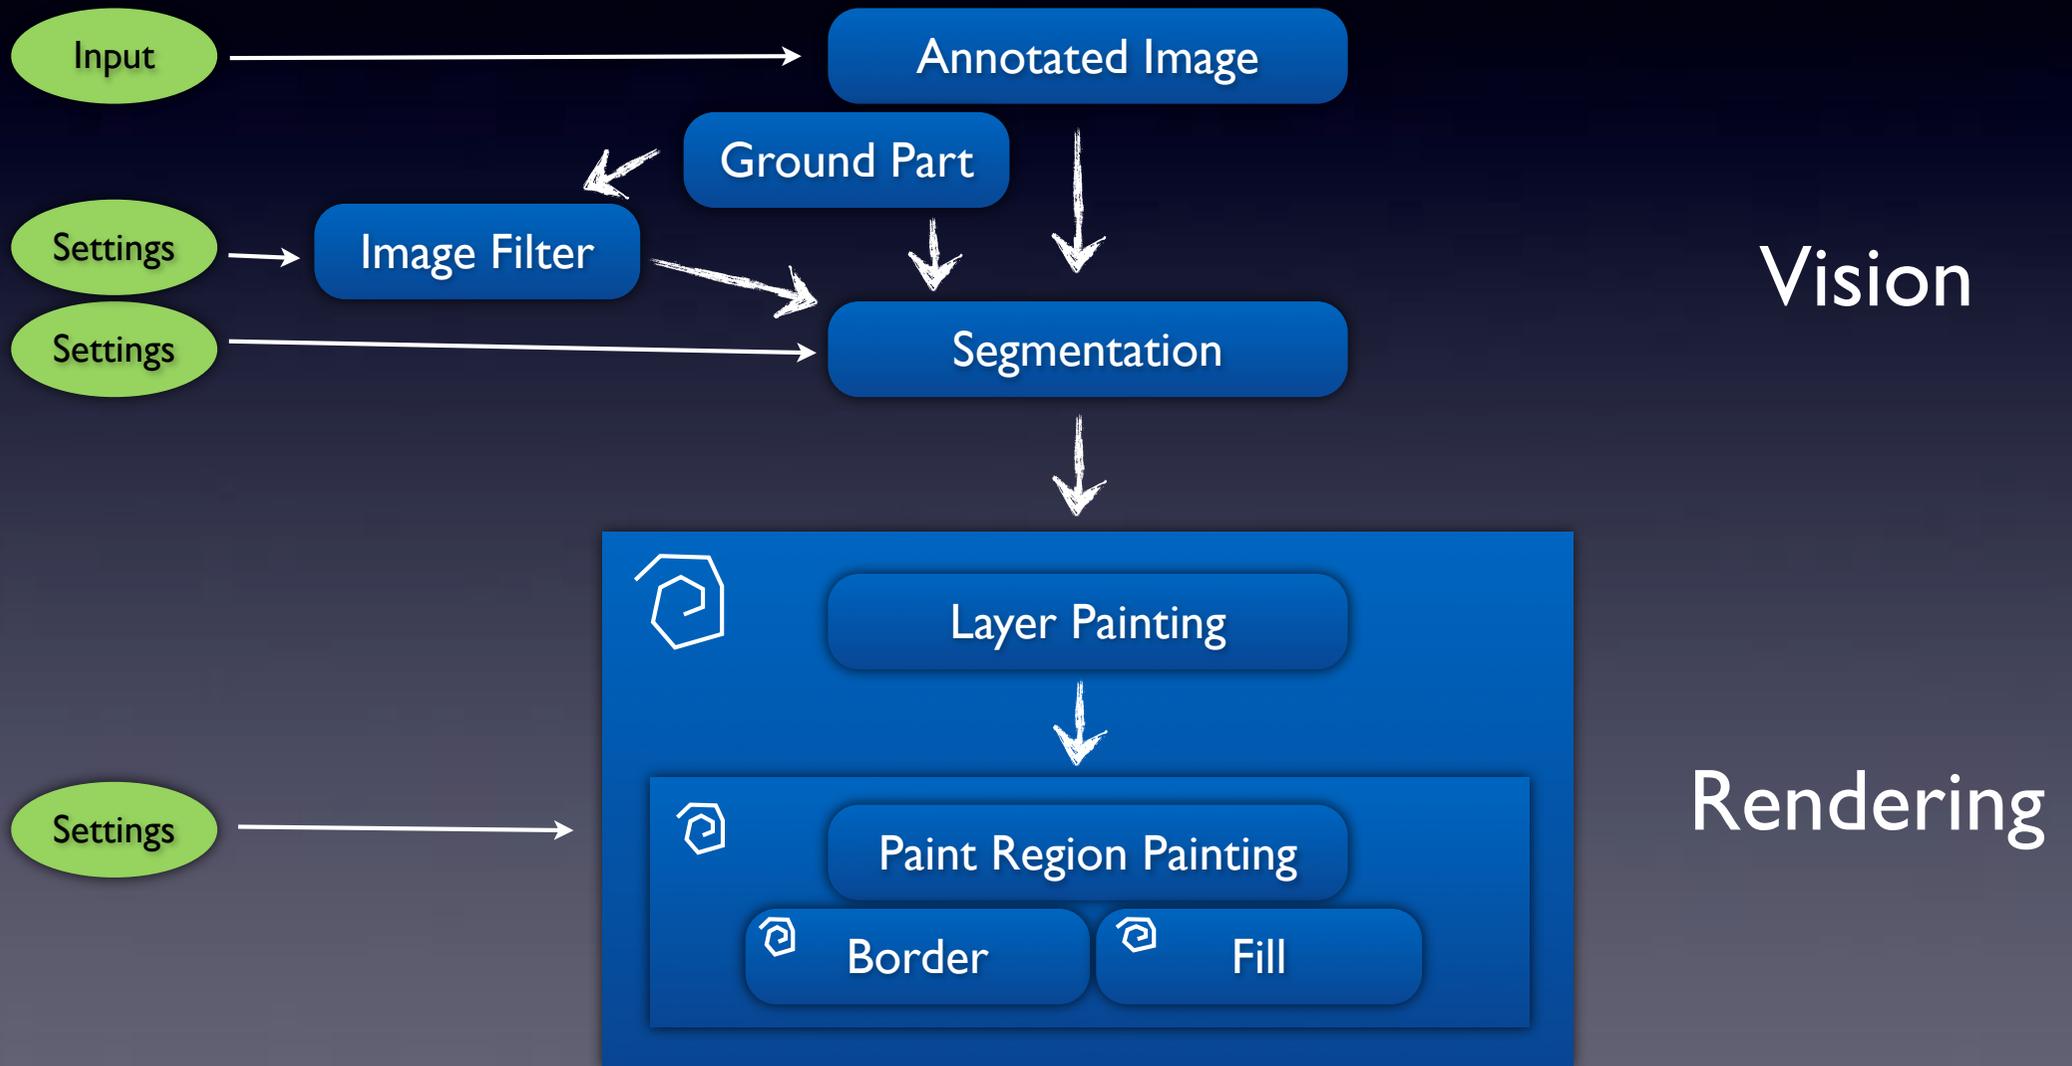
Stage I:
Non-photorealistic Rendering
(Skills)



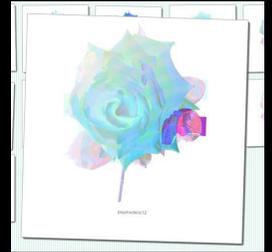
Stage I

Build a non-photorealistic rendering system for photos. Uses image filtering, image segmentation, curve fitting, natural media simulation, filling routines, transformations, etc.

Skill-based Graphics Workflow



Demonstration



- Filter Feast software

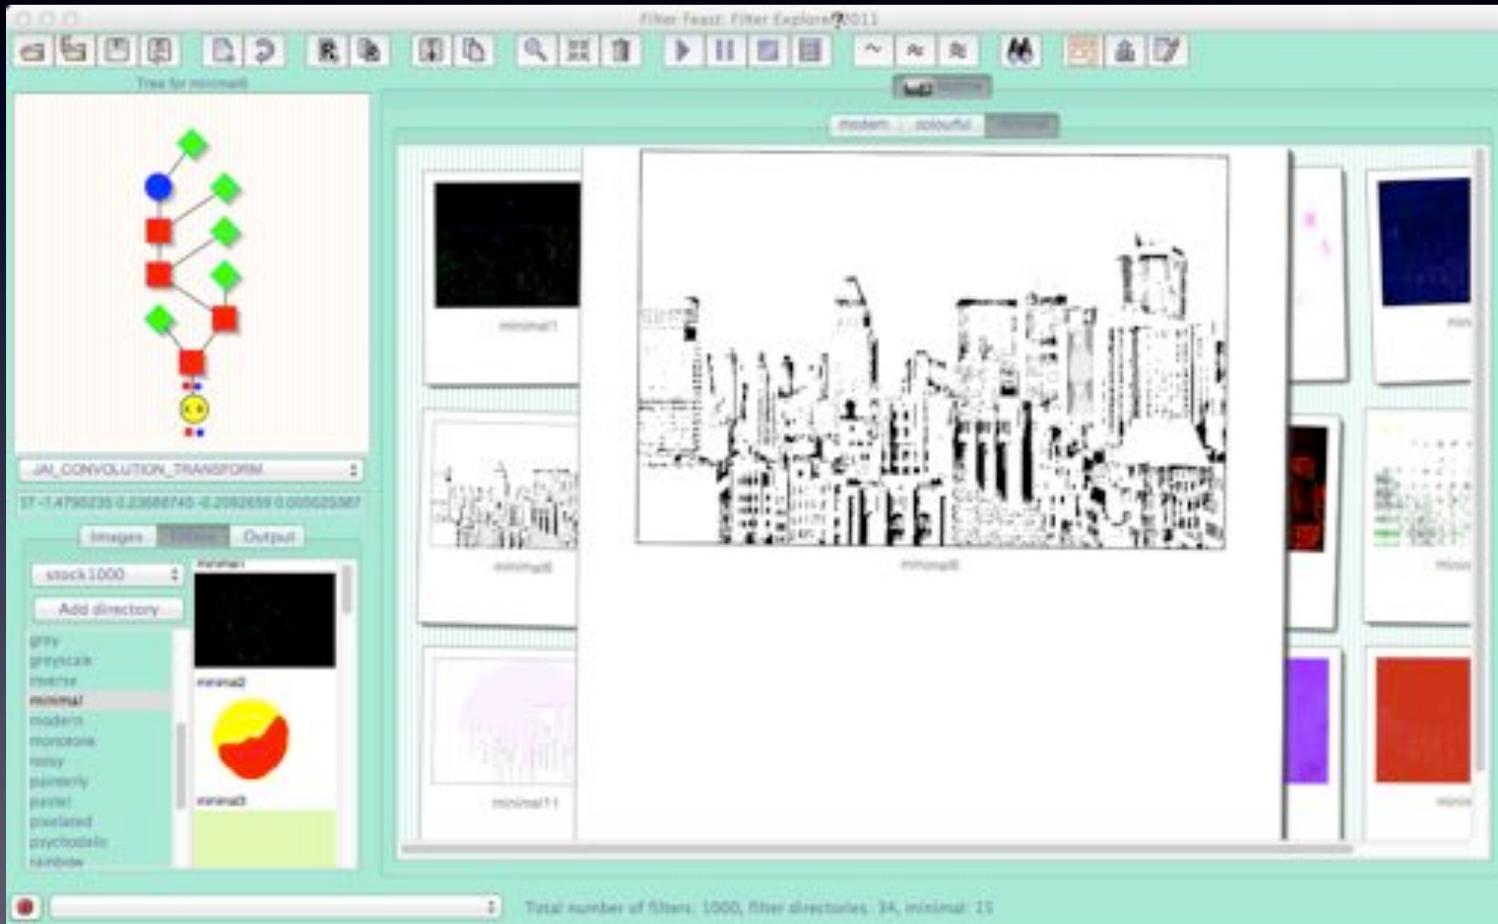
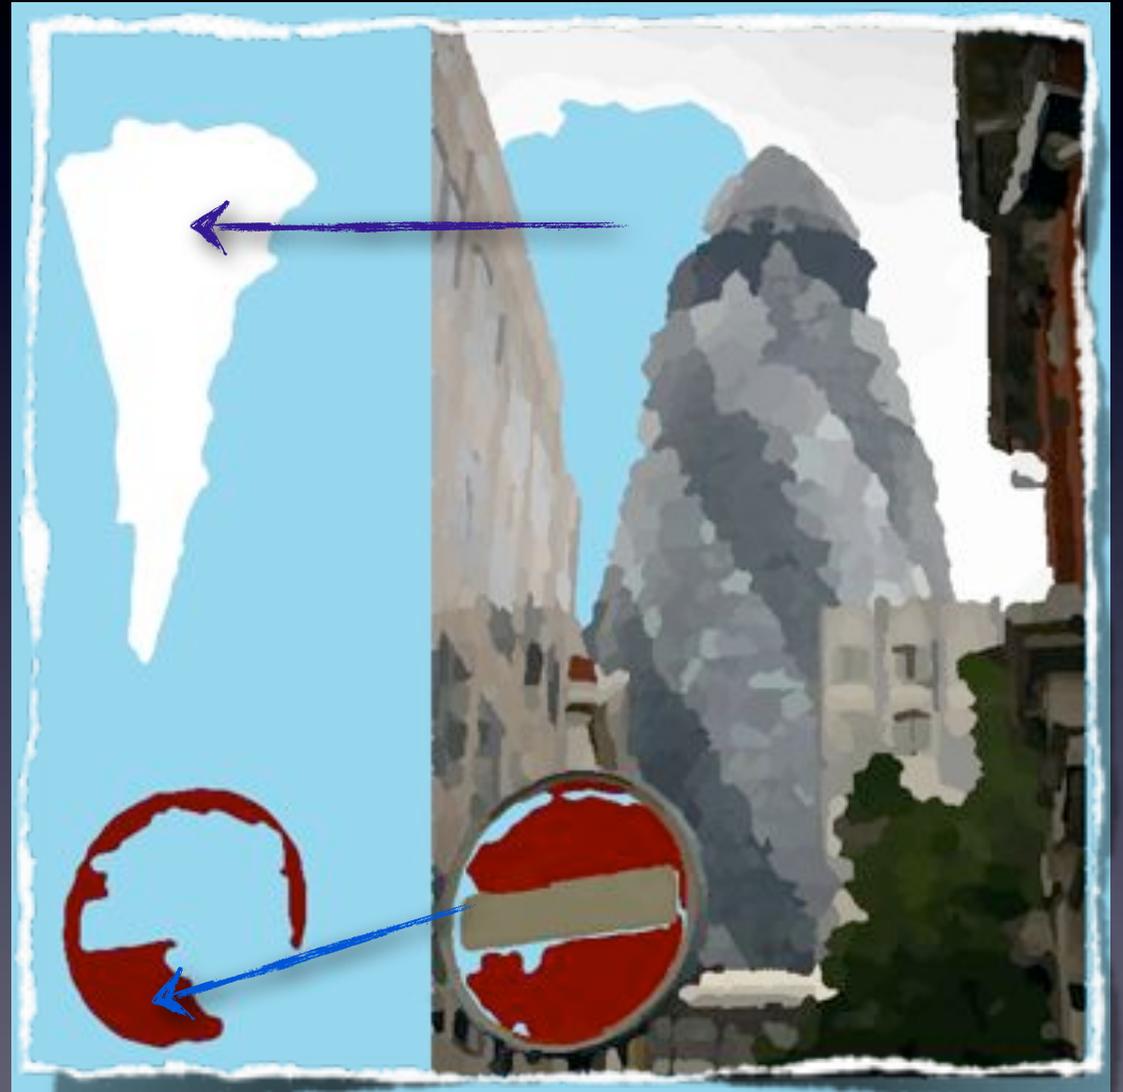


Image Segmentation

Example - 1000 regions



Skipping Some Slides

Non-Photorealistic Rendering

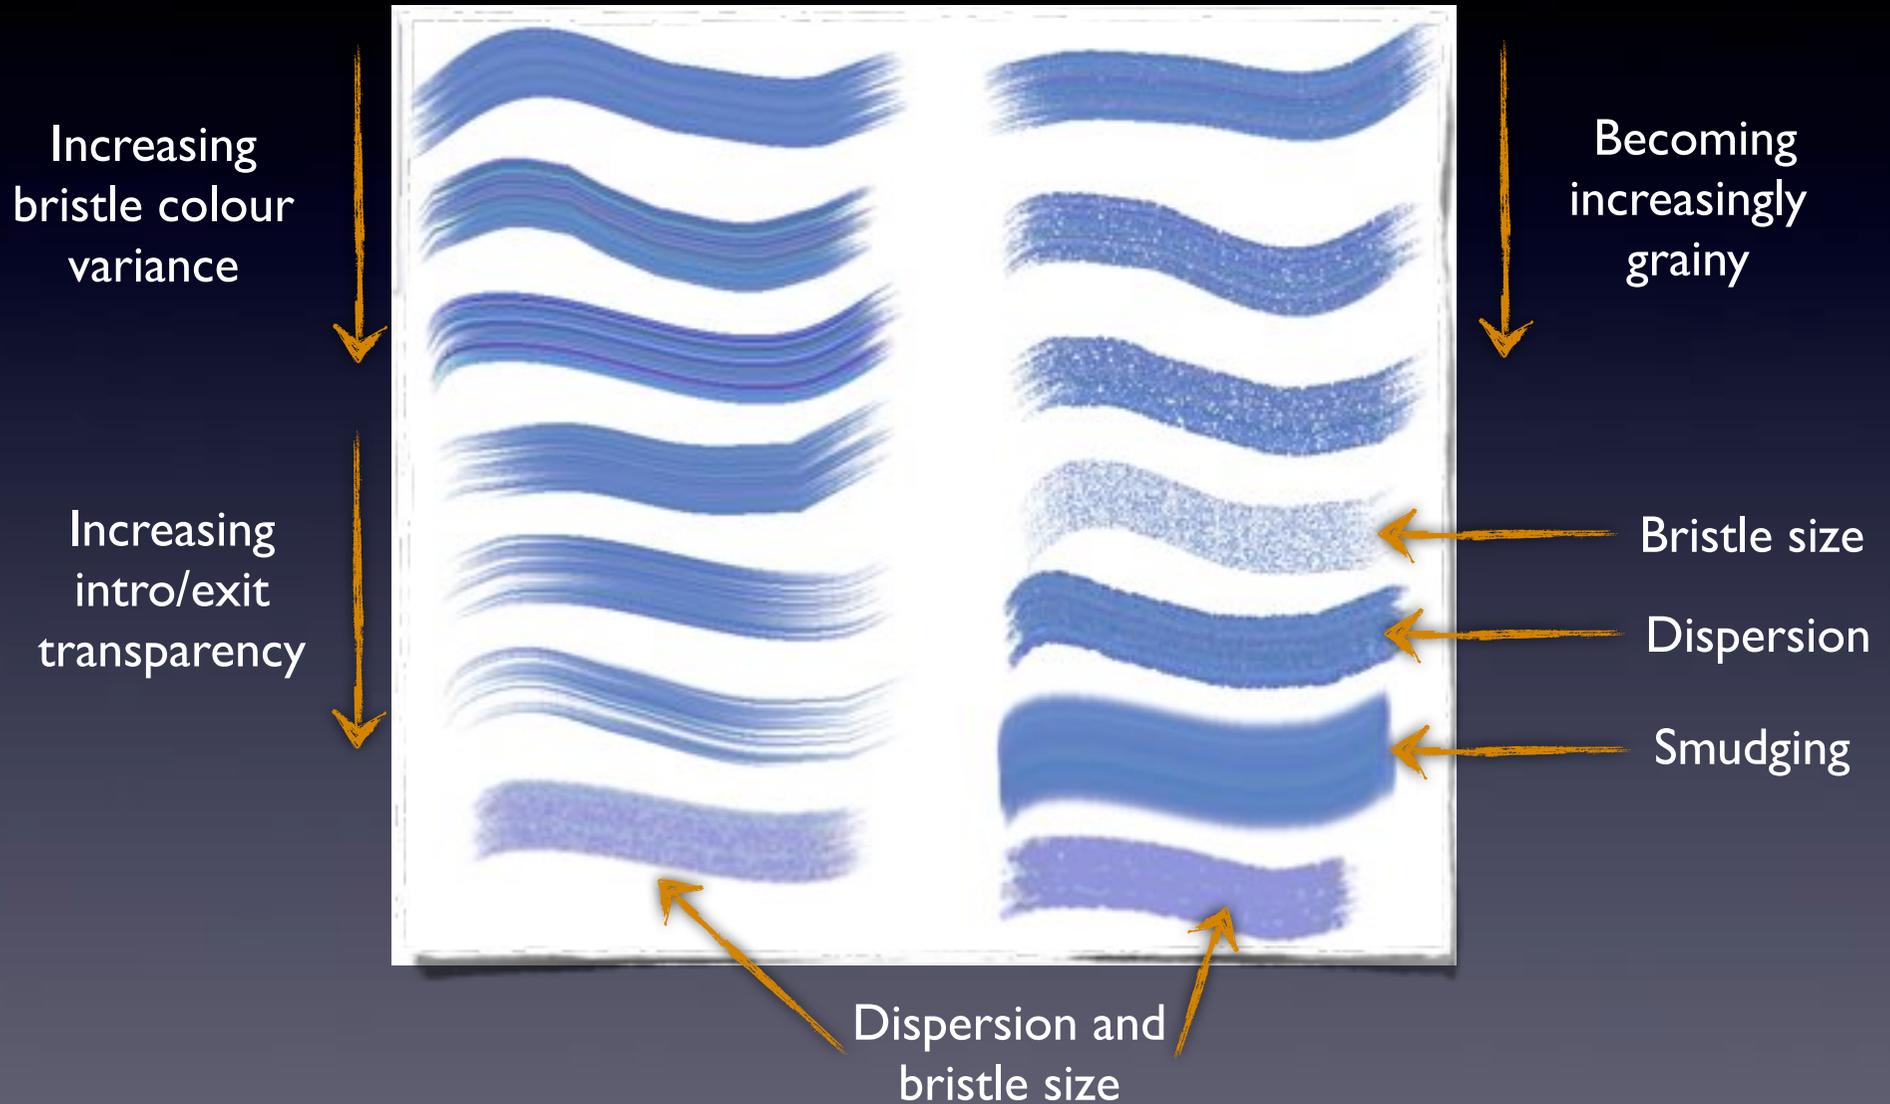
The image shows a sequence of six presentation slides, numbered 16 through 21, each with a dark blue background and white text. The slides are arranged horizontally and contain various diagrams, lists, and images related to computer graphics and rendering.

- Slide 16: Image Segmentation: Finding Regions of Similar Colour**
 - Need to measure how far one pixel is from another
 - In terms of their RGB colours
 - Euclidean distance in RGB space of two pixels P_1 and P_2
 - $\text{dist}(P_1, P_2) = \sqrt{(R_1 - R_2)^2 + (G_1 - G_2)^2 + (B_1 - B_2)^2}$
 - Define distance according to HSV channels

- Slide 17: Simulating Natural Media**
- Need to simulate four things:
 - Backing materials: paper, canvas, etc.
 - Painting implements: brushes, palette knives, etc.
 - Pigment materials: paints, pencils, pencils, etc.
 - Optical usage of these materials
- Slide 18: Simulating Natural Media**
- Simple method:
 - Scale, rotate and bend a fixed rendering process/template to fit the stroke path
- This is efficient, but there is too little variety in the strokes
- Slide 19: Brush Paths: More Examples (from Adobe Photoshop)**
- Slide 20: Brush Paths: Adding Sophistication**
- Enable mapping to stroke to add grain and depth difference
- Enable nesting and using different brush tool within a temporary
- Force the brush head to reach the corners of the stroke
- Enable penes to use as they are painted
- Rescale original penes
- Distort overlapping and path direction
- Slide 21: Brush Paths: Adding Sophistication**
- Distort and stretch the stroke
- Stretch and stretch the stroke
- Stretch and stretch the stroke

Brush Paths

Adding Sophistication

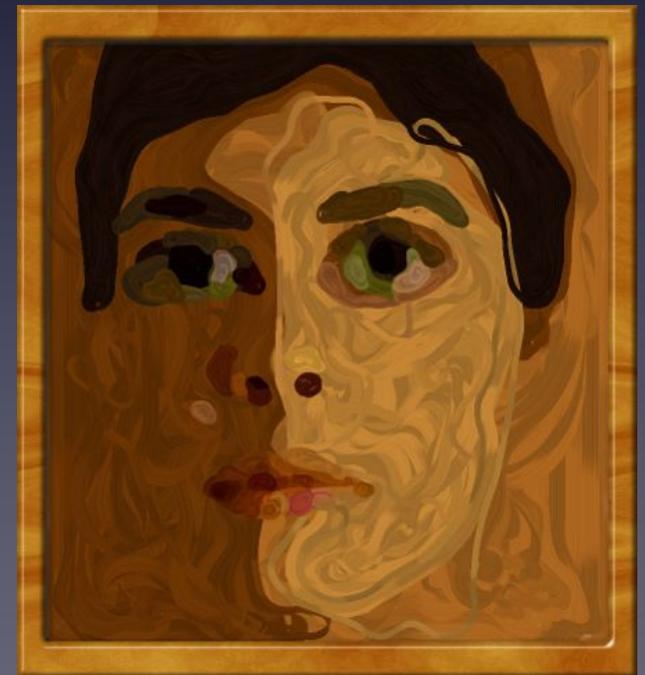


Made by The Painting Fool

Stage 2
Emotional Modelling
(Appreciation)

Stage 2.1

- Externalise the parameters which define styles: colour palettes, abstraction levels, fill styles, natural media, etc; enable search over sets of parameters; match emotion enhancement to styles; build KB; enable control via emotion keywords;



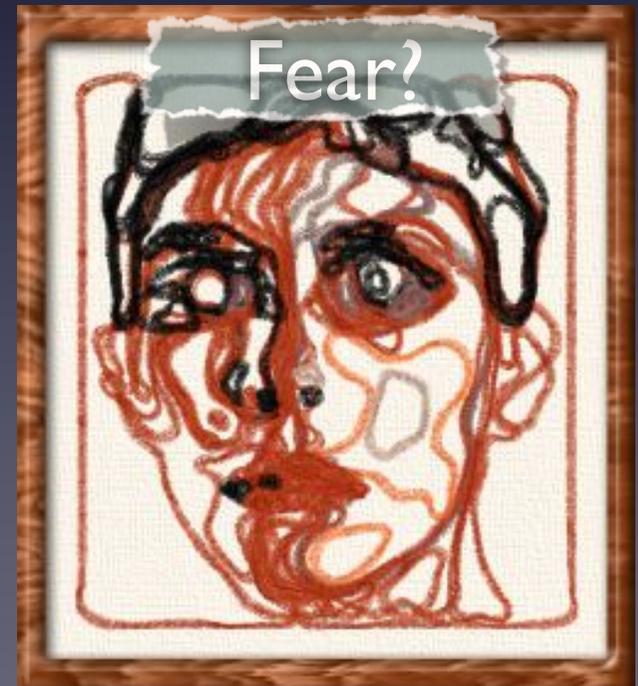
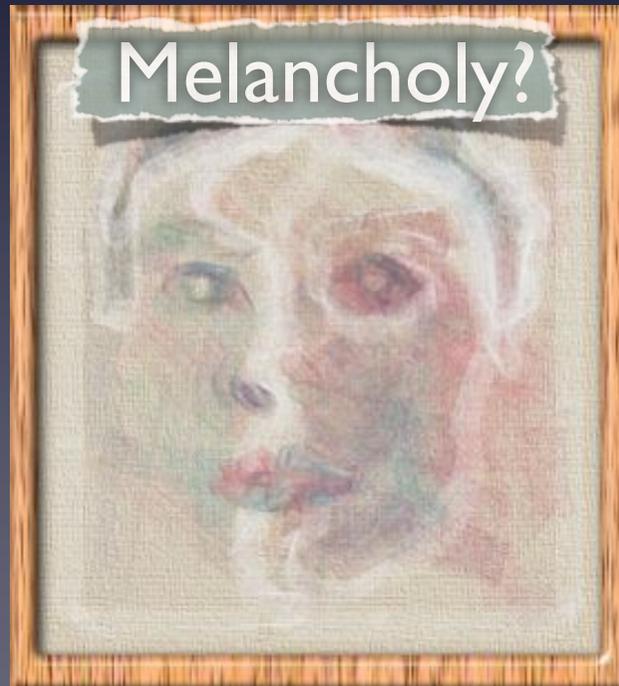
Stage 2.1

- Externalise the parameters which define styles: colour palettes, abstraction levels, fill styles, natural media, etc; enable search over sets of parameters; match emotion enhancement to styles; build KB; enable control via emotion keywords;



Stage 2.1

- Externalise the parameters which define styles: colour palettes, abstraction levels, fill styles, natural media, etc; enable search over sets of parameters; match emotion enhancement to styles; build KB; enable control via emotion keywords;



Amelie's Progress Gallery



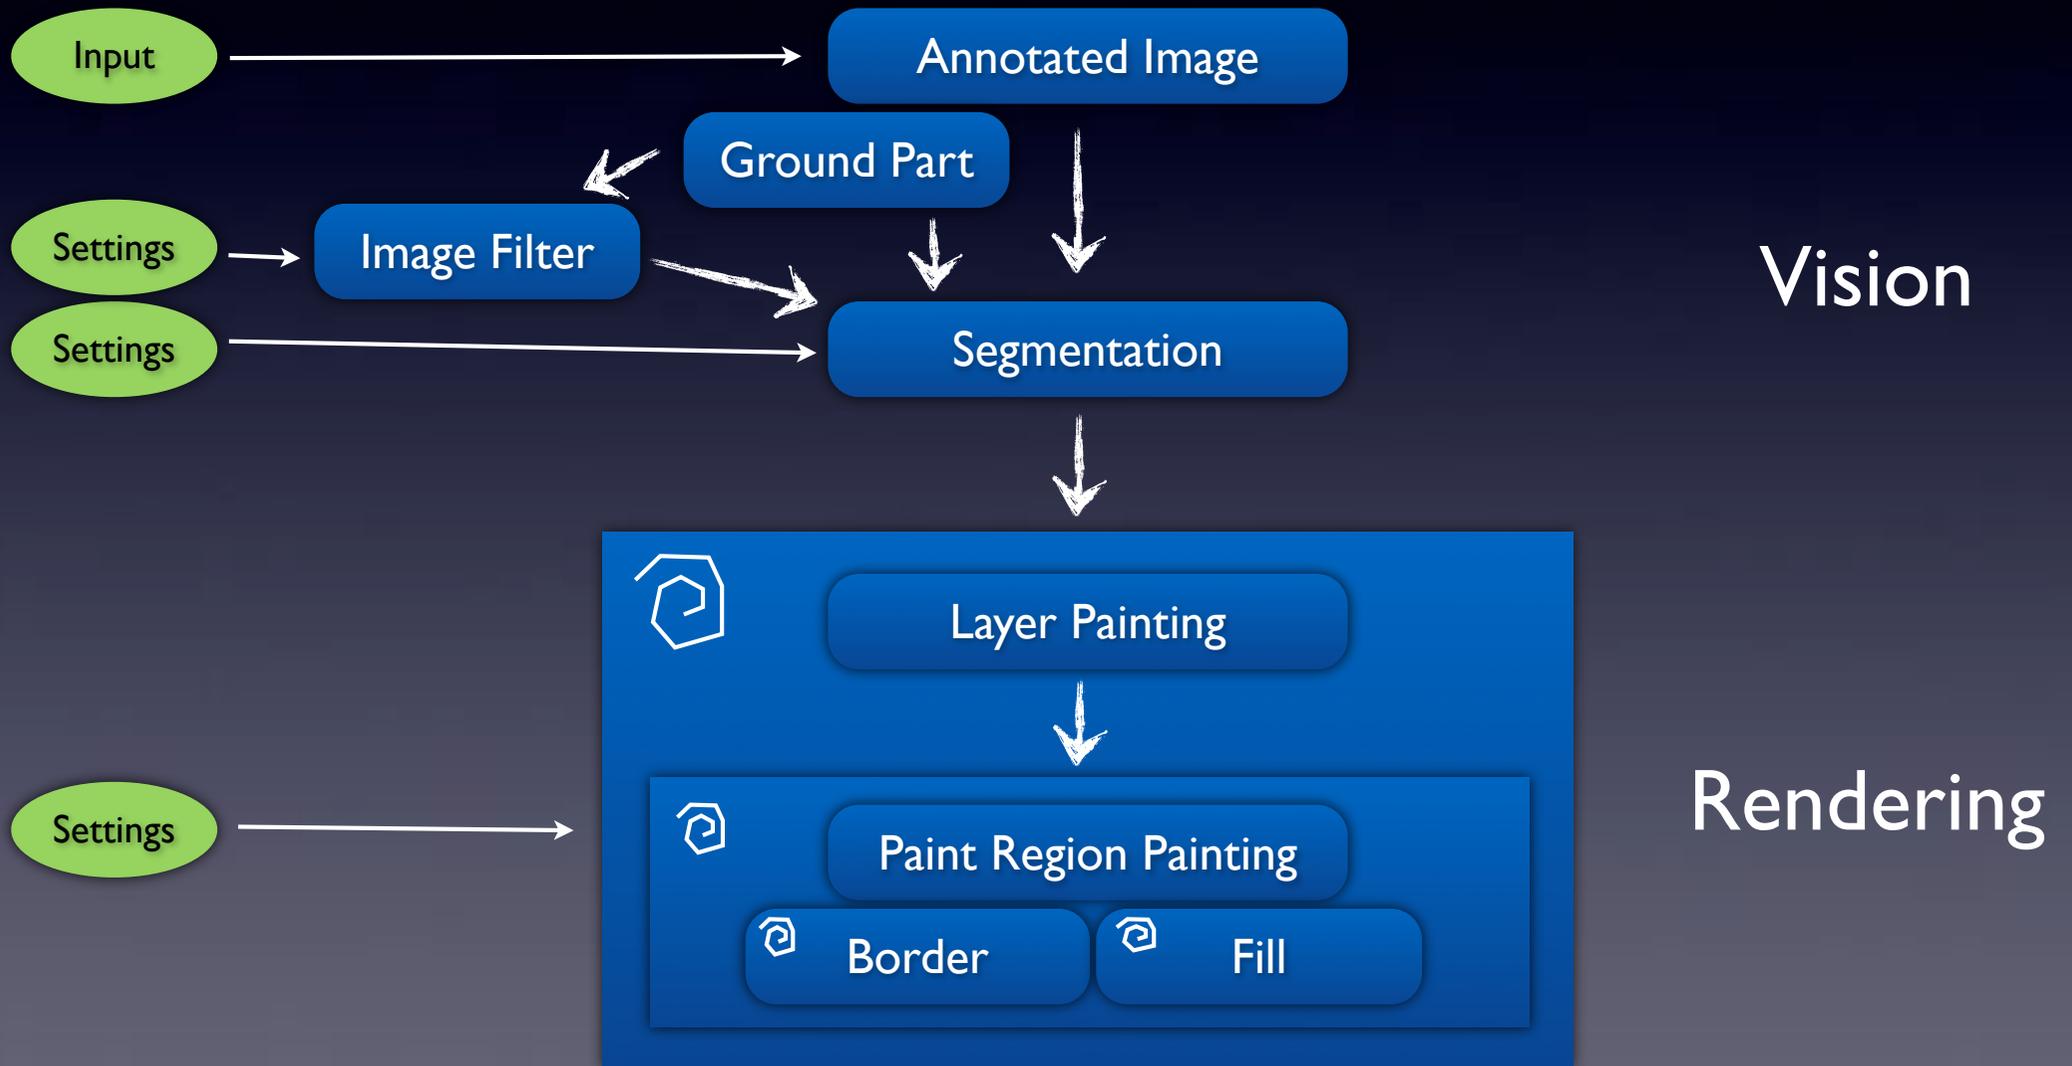
222 pictures from 22 images; emotional direction

Stage 2.2

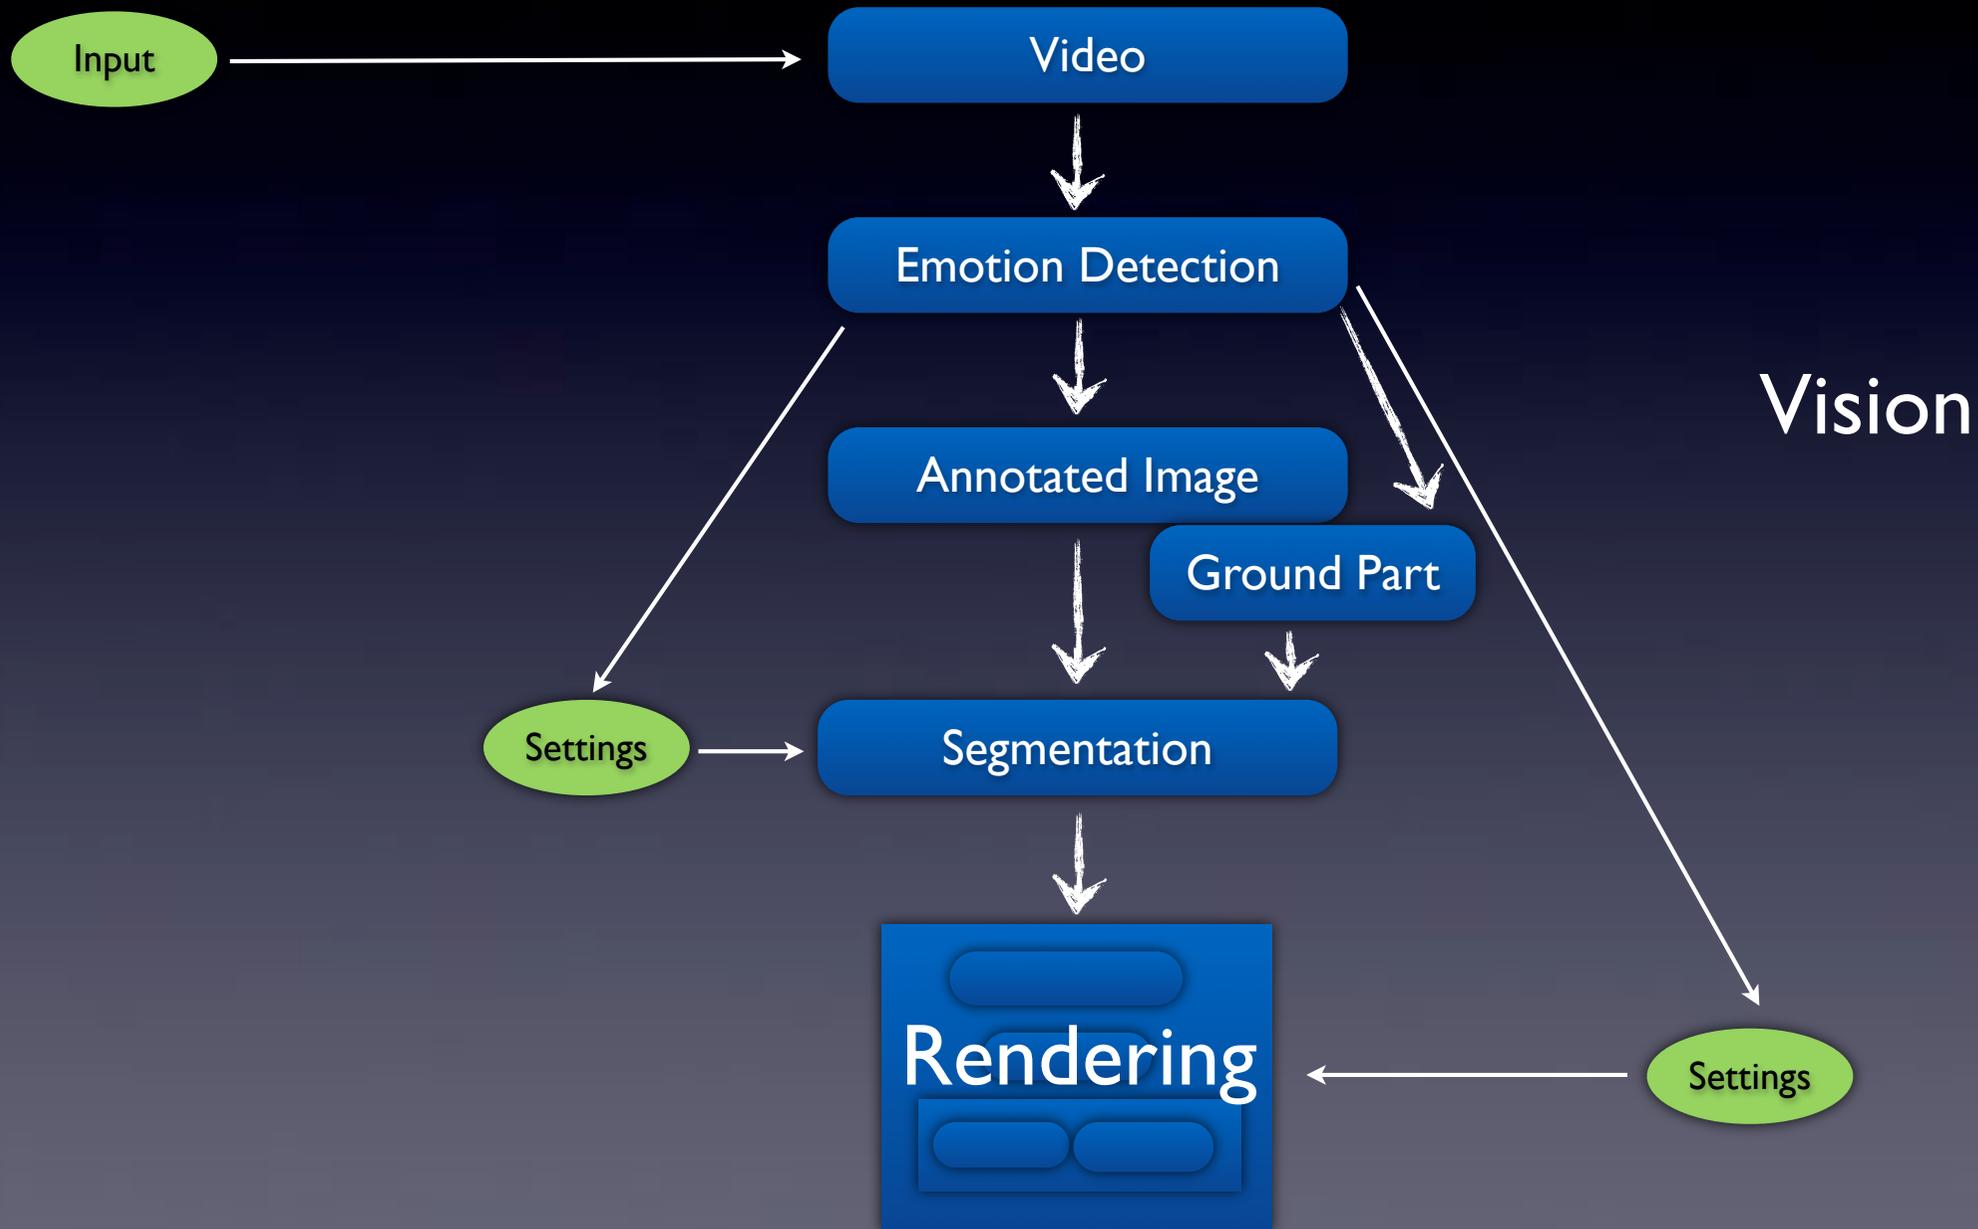
Emotional Enhancements

- For the BCS machine intelligence competition (Dec. 2007)
- Combined The Painting Fool with the emotion detection system of Maja Pantic and Michel Valstar
 - Uses a machine-learned boosted classifier based which uses the movements of facial anchor points
- Uses video clips of someone expressing an emotion and detects (i) apex image (ii) feature locations and (iii) one of six emotions - which is passed to The Painting Fool
- The Painting Fool chooses its painting style according to the knowledge base of emotion-enhancing styles

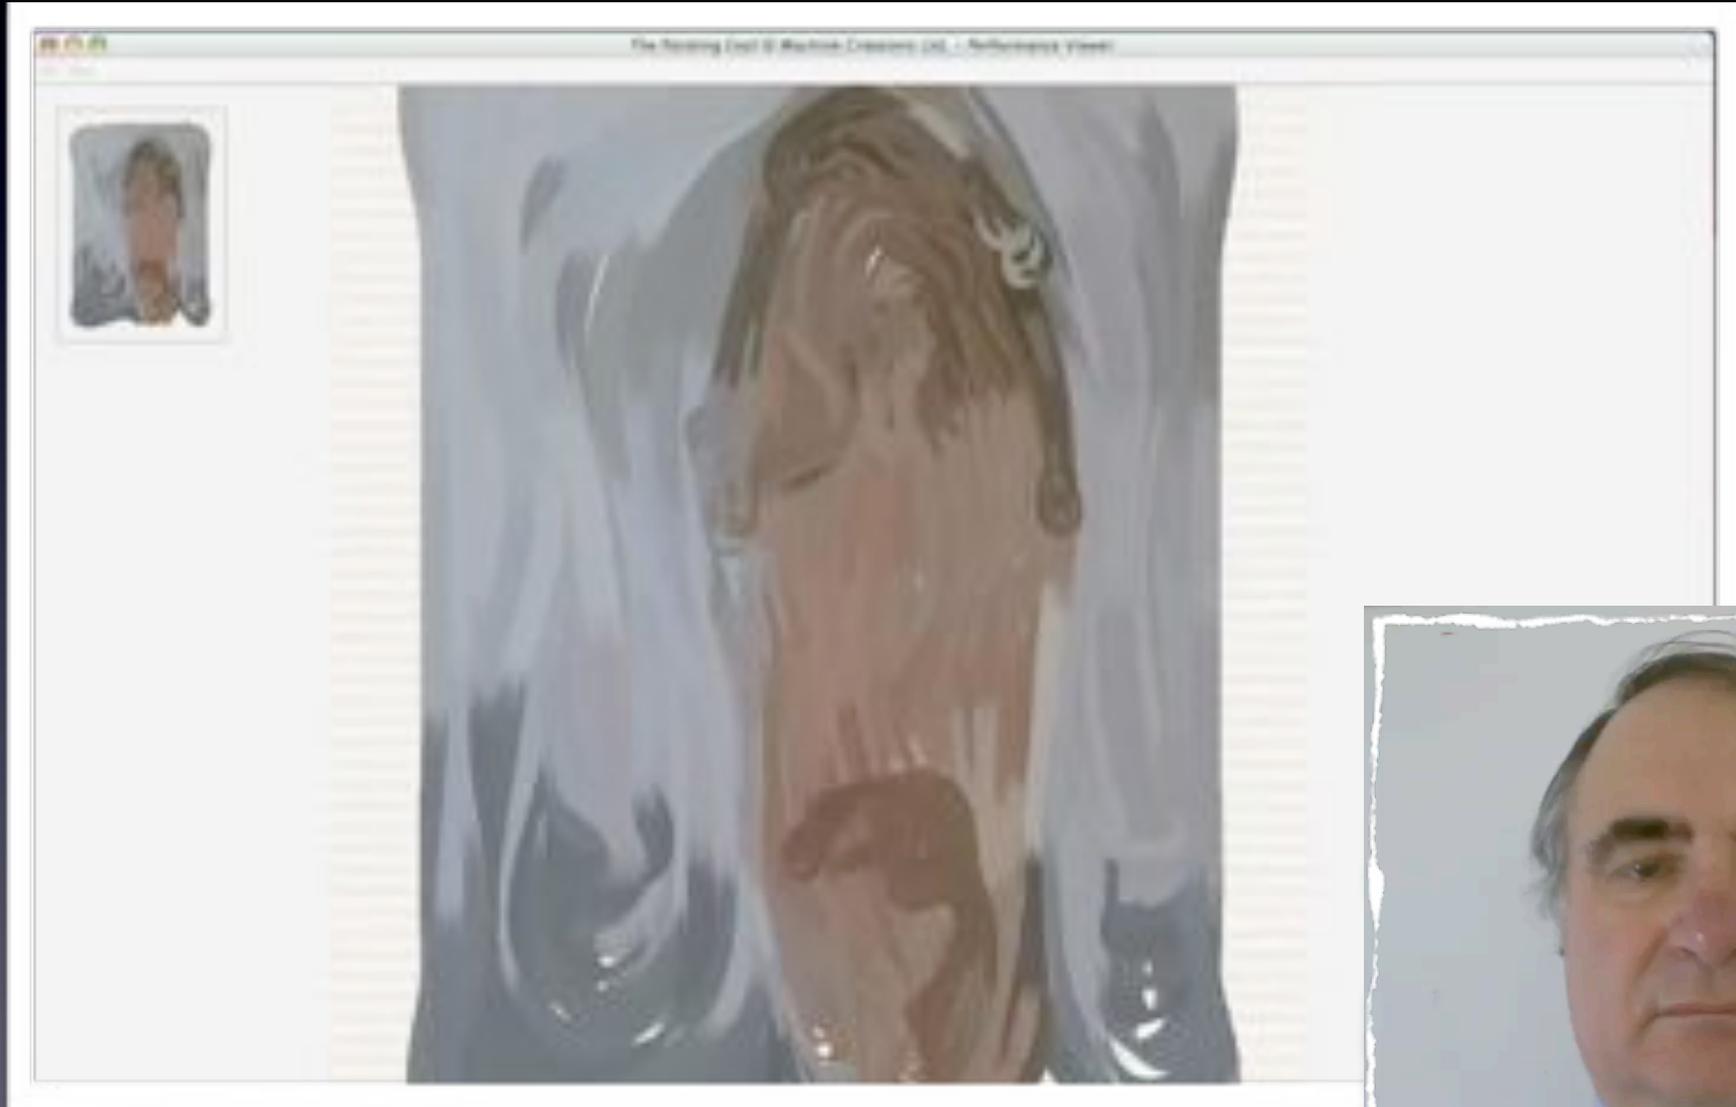
Skill-based Graphics Workflow



Appreciation-based Graphics Workflow



Example: Mike Swain



More Portraits





Stage 3
Generative Art
(Imagination)

Stage 3.1

Automatic Scene Construction

Scenes with repetitive elements



GUSTAV
KLIMT



GUSTAV
KLIMT

Skipping Some Slides...

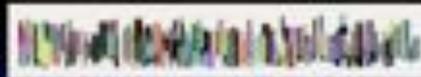
Approach 1: Evolutionary

Approach 1: Evolutionary

- Scenes are a set of geometric shapes, e.g. rectangles
- Randomly generate a scene by defining
 - The colour, position and shape of 100s of elements
- Search for aesthetically pleasing scenes
- User-defined fitness function (see next slides)
- Randomly generate scenes, then crossover chunks of shapes in evolutionary approach
- Also use a hill-climbing approach

35

Random Example:



Inspiring Example:



37

Fitness Functions

- Parameters of the scene elements:
 - Height, width, x, y loc, saturation, brightness
 - Position in the ordered scene list
- Calculations using these parameters:
 - Depth of scene elements, distance from centre
- Aspects of the fitness function:
 - Correlations between the parameters

38

Correlations f

- Using Pearson Product-Moment
- Distance from centre
 - Positively correlated with y co
 - Negatively correlated with he
- Depth:
 - Positively correlated with heig
 - Negatively correlated with co

39

Skipping Some Slides...

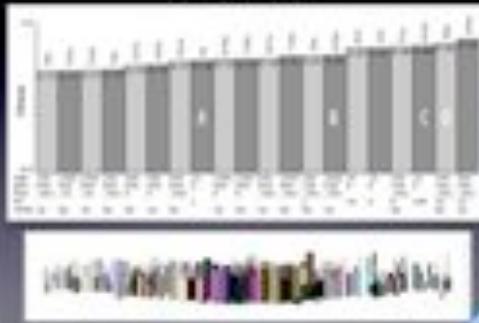
Approach I: Evolutionary

Correlations Required

- Using Pearson Product-Moment Correlations
- Distance from centre
- Positively correlated with y coordinate
- Negatively correlated with height, width and rotation
- Depth
- Positively correlated with height
- Negatively correlated with rotation, brightness

39

Results



40

To the Meta-Level...

- Automatically invest the fitness function using IR.
- See lecture 4 of mine for details about IR software
- Because deriving a fitness function is essentially mathematical modelling
- IR is given the same background knowledge that the correlations use, namely:
 - Objects of interest are scenes
 - Subobjects are scene elements
 - X, Y, W, H, Hue, Sat, Bri, Dist, Depth are the background concepts

41

Fitness Function In

```
1 // scene element specifications (subset of scene data)
2 // scene element positions and sizes for the evolution
3 // scoring function
4 // ...
5 // ...
6 // ...
7 // ...
8 // ...
9 // ...
10 // ...
11 // ...
12 // ...
13 // ...
14 // ...
15 // ...
16 // ...
17 // ...
18 // ...
19 // ...
20 // ...
21 // ...
22 // ...
23 // ...
24 // ...
25 // ...
26 // ...
27 // ...
28 // ...
29 // ...
30 // ...
31 // ...
32 // ...
33 // ...
34 // ...
35 // ...
36 // ...
37 // ...
38 // ...
39 // ...
40 // ...
41 // ...
42 // ...
43 // ...
44 // ...
45 // ...
46 // ...
47 // ...
48 // ...
49 // ...
50 // ...
51 // ...
52 // ...
53 // ...
54 // ...
55 // ...
56 // ...
57 // ...
58 // ...
59 // ...
60 // ...
61 // ...
62 // ...
63 // ...
64 // ...
65 // ...
66 // ...
67 // ...
68 // ...
69 // ...
70 // ...
71 // ...
72 // ...
73 // ...
74 // ...
75 // ...
76 // ...
77 // ...
78 // ...
79 // ...
80 // ...
81 // ...
82 // ...
83 // ...
84 // ...
85 // ...
86 // ...
87 // ...
88 // ...
89 // ...
90 // ...
91 // ...
92 // ...
93 // ...
94 // ...
95 // ...
96 // ...
97 // ...
98 // ...
99 // ...
100 // ...
```

42

Skipping Some Slides...

Approach I: Evolutionary

Function Invention



42

Illustrative Result



43

Cityscape Results

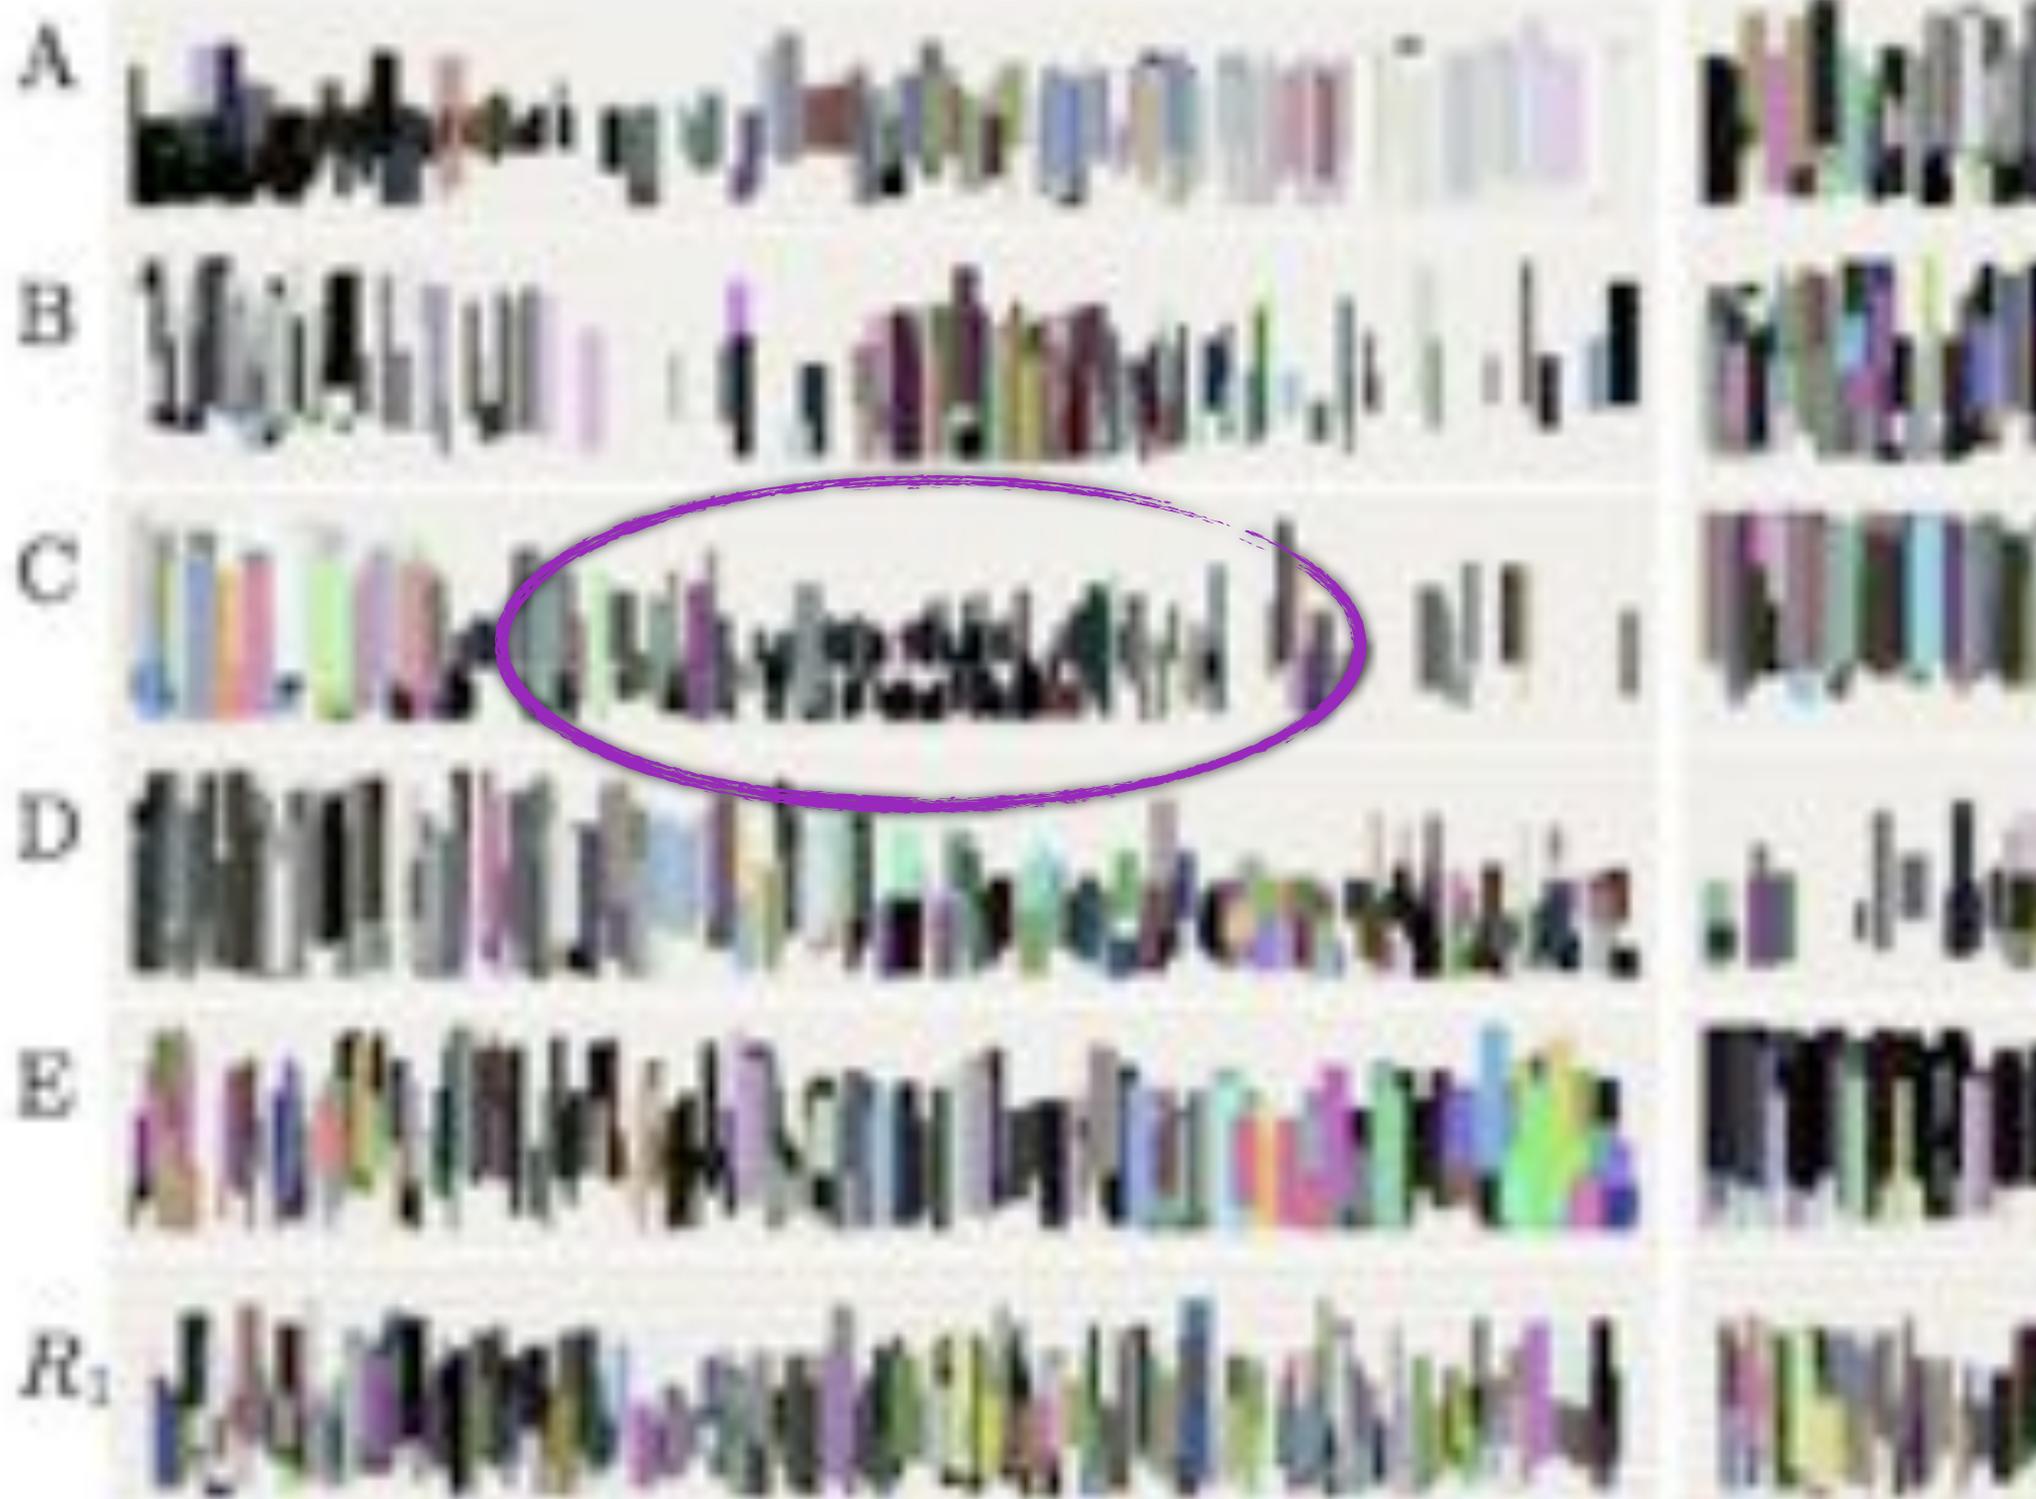


44

A B C D E F



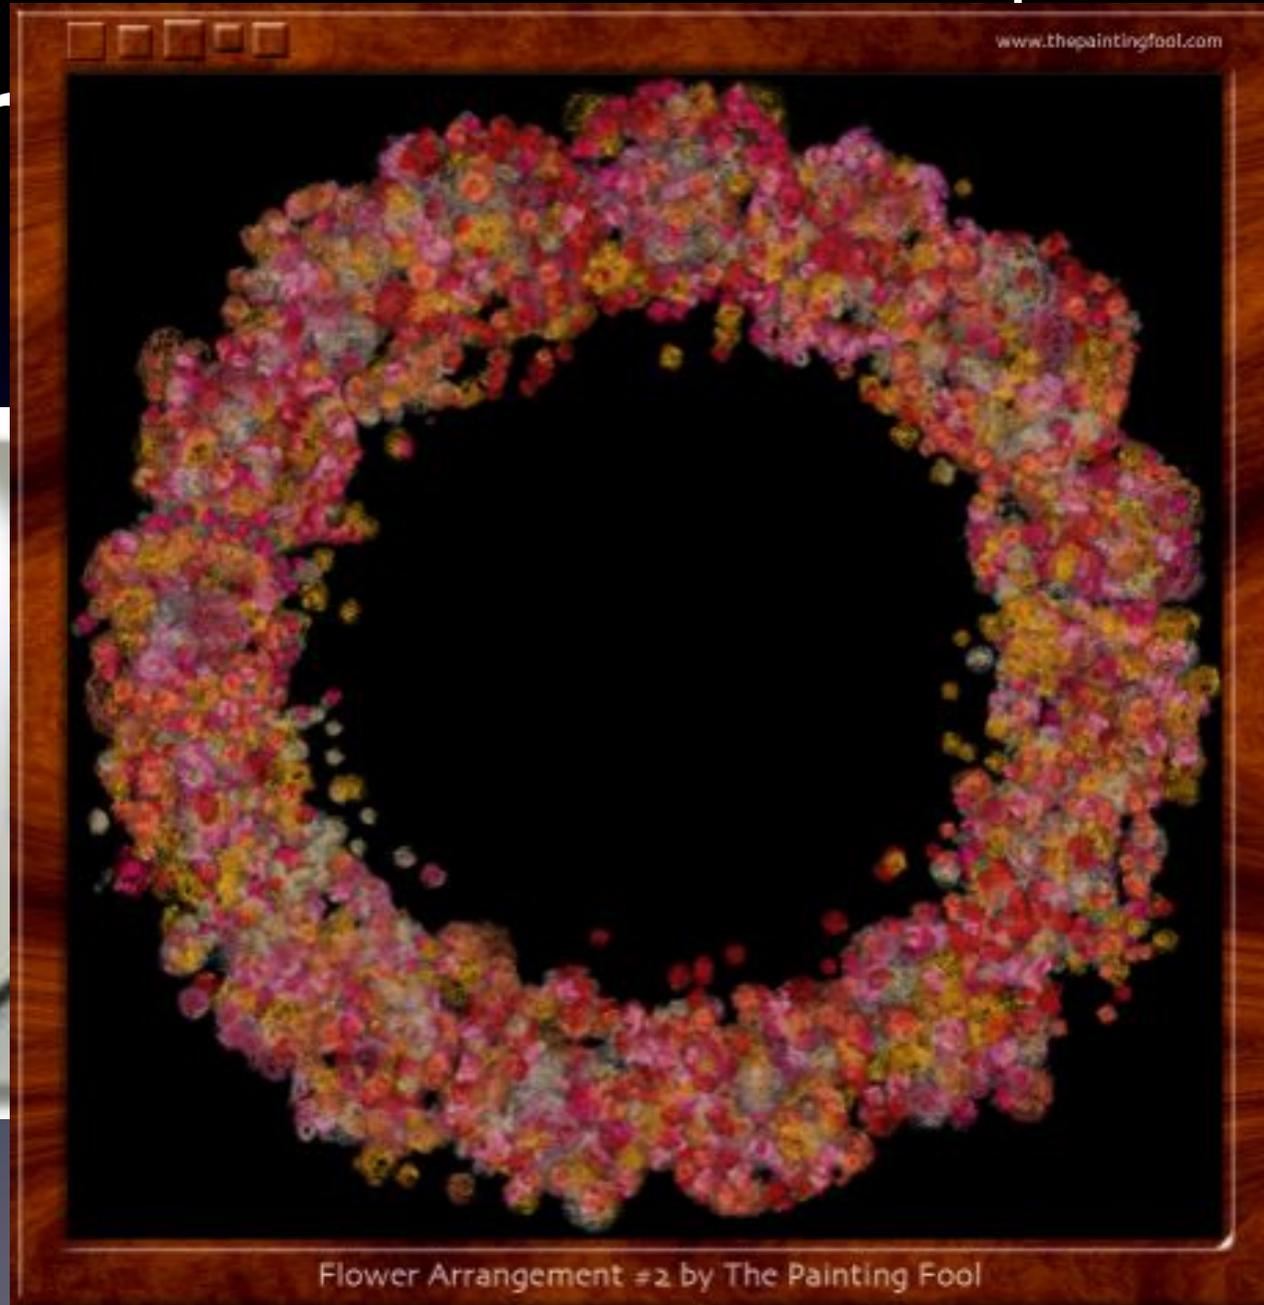
45



Permanent Exhibition at Imperial

Per

aint



Approach 2: CSP based

The image displays three overlapping screenshots of a software interface, likely a Prolog-based CSP solver, used for solving a problem involving shapes on a grid.

The top-left screenshot shows a grid with two vertical bars (one short, one tall) placed on a grid. The interface includes a toolbar and a list of constraints on the right side.

The middle screenshot shows a grid with seven vertical bars of varying heights and widths placed on a grid. The interface includes a toolbar and a list of constraints on the right side.

The bottom-right screenshot shows a code editor window titled "1236070815372" with the following Prolog code:

```
constraints.pl domain.pl generator.pl go.bat shapes.pl

unary_constraint('shape_type', 0, S) :-
    get_type(S, Type),
    Type = 0.

binary_constraint('disjoint_250_998', S1, S2) :-
    get_x(S1, TL1_X), get_y(S1, TL1_Y),
    get_x(S2, TL2_X), get_y(S2, TL2_Y),
    get_width(S1, W1), get_width(S2, W2),
    get_height(S1, H1), get_height(S2, H2),
    PaddedX1 #= TL1_X - 125,
    PaddedY1 #= TL1_Y - 499,
    PaddedWidth1 #= W1 + 250,
    PaddedHeight1 #= H1 + 998,
    PaddedX2 #= TL2_X - 125,
    PaddedY2 #= TL2_Y - 499,
    PaddedWidth2 #= W2 + 250,
    PaddedHeight2 #= H2 + 998,
    disjoint_rectangles([PaddedX1, PaddedY1, PaddedHeight1], [PaddedX2, PaddedY2, PaddedWidth2, PaddedHeight2]).

binary_constraint('equal_brightness', S1, S2) :-
    get_brightness(S1, Br1), get_brightness(S2, Br2),
    Br1 = Br2.

binary_constraint('top_equal_to', S1, S2) :-
    get_y(S1, Y1), get_y(S2, Y2),
    Y1 = Y2.

binary_constraint('greater_height_implies_greater_saturation', S1, S2) :-
    get_height(S1, H1), get_height(S2, H2),
    get_saturation(S1, Sat1), get_saturation(S2, Sat2),
    (H1 #> H2) => (Sat1 #> Sat2).

binary_constraint('greater_height_implies_greater_bottom', S1, S2) :-
    get_height(S1, H1), get_height(S2, H2),
    get_y(S1, Y1), get_y(S2, Y2), get_height(S1, H1), get_height(S2, H2),
    Bottom1 #= Y1 + H1, Bottom2 #= Y2 + H2,
    (H1 #> H2) => (Bottom1 #> Bottom2).

binary_constraint('greater_height_implies_greater_width', S1, S2) :-
    get_height(S1, H1), get_height(S2, H2),
    get_width(S1, W1), get_width(S2, W2),
    (H1 #> H2) => (W1 #> W2).

binary_constraint('greater_saturation_implies_greater_height', S1, S2) :-
    get_saturation(S1, Sat1), get_saturation(S2, Sat2),
    get_height(S1, H1), get_height(S2, H2),
    (Sat1 #> Sat2) => (H1 #> H2).

% POSTING CONSTRAINTS
post_constraints(Shape, AllShapes) :-
    findall(U, applicable_unary_constraint(U, US),
    post_unary_constraints(Shape, US)),
    findall(B, applicable_binary_constraint(B, BS),
    post_binary_constraints(Shape, BS, AllShapes)).

post_unary_constraints(_C, []).
```

See me later for a demo...

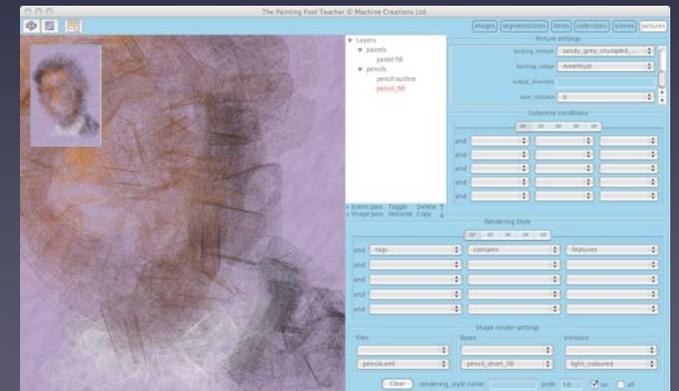
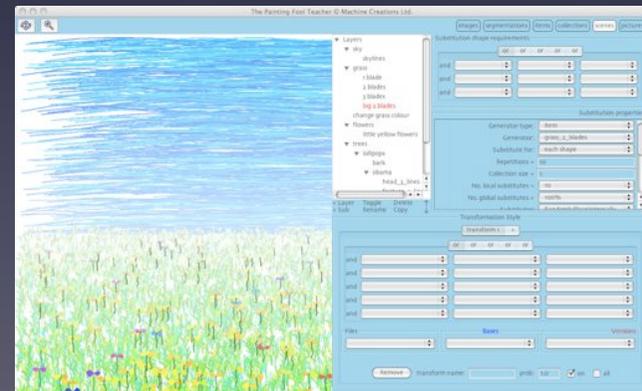
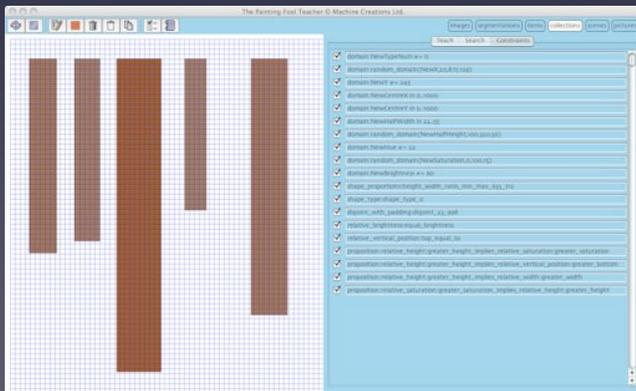
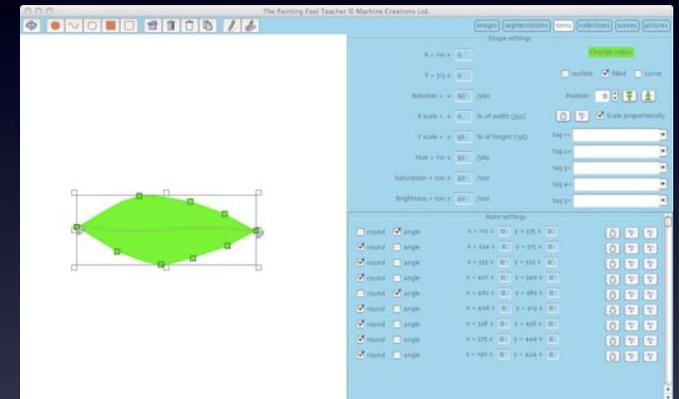
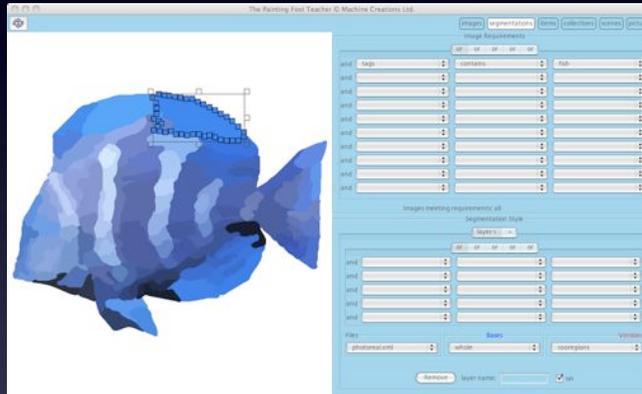
Stage 3.2

Automatic Scene Construction

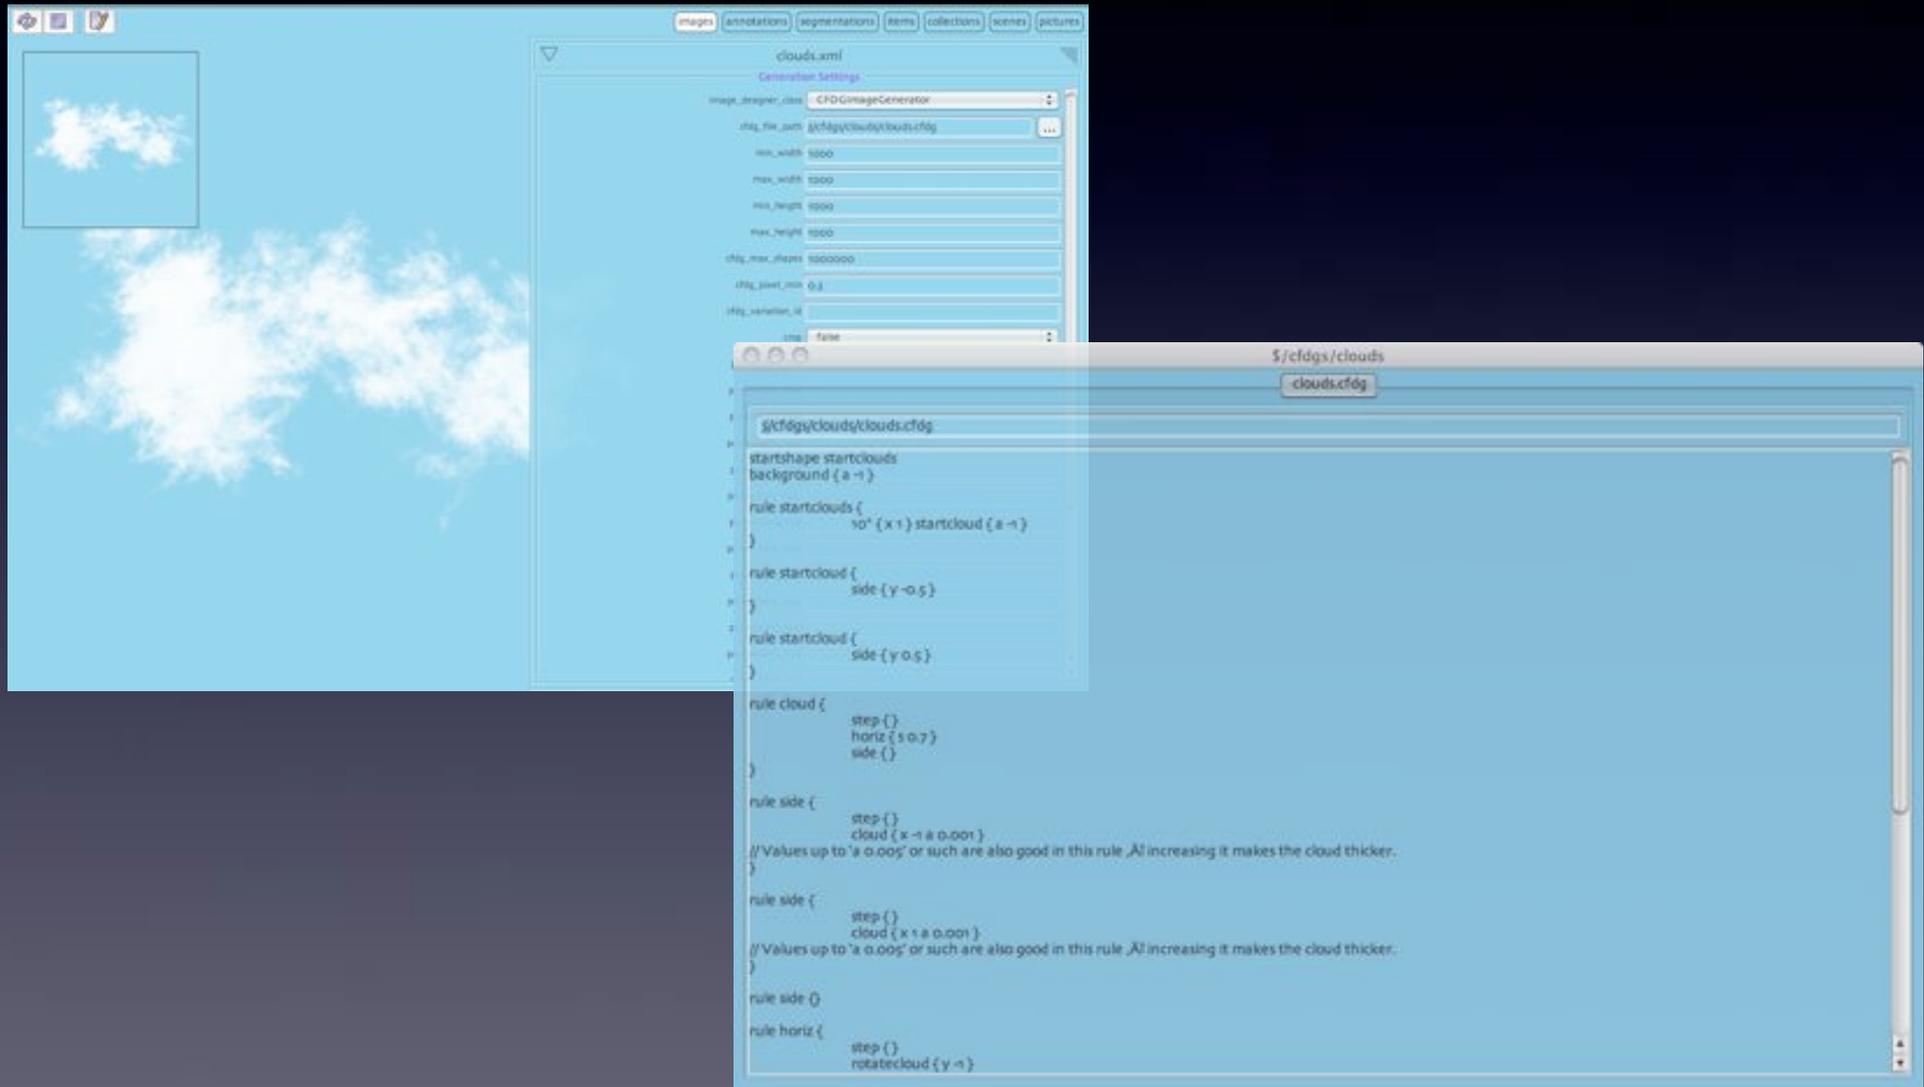
Scenes with discernible objects

- Evolutionary method absorbed into a *teaching interface* for The Painting Fool
 - Artists are influenced, inspired and taught by many diverse teachers, not just one programmer!
- Modules include:
 - Evolutionary approach; image filtering; constraint generator and solver; evolutionary art; context free design grammars; animation generation;

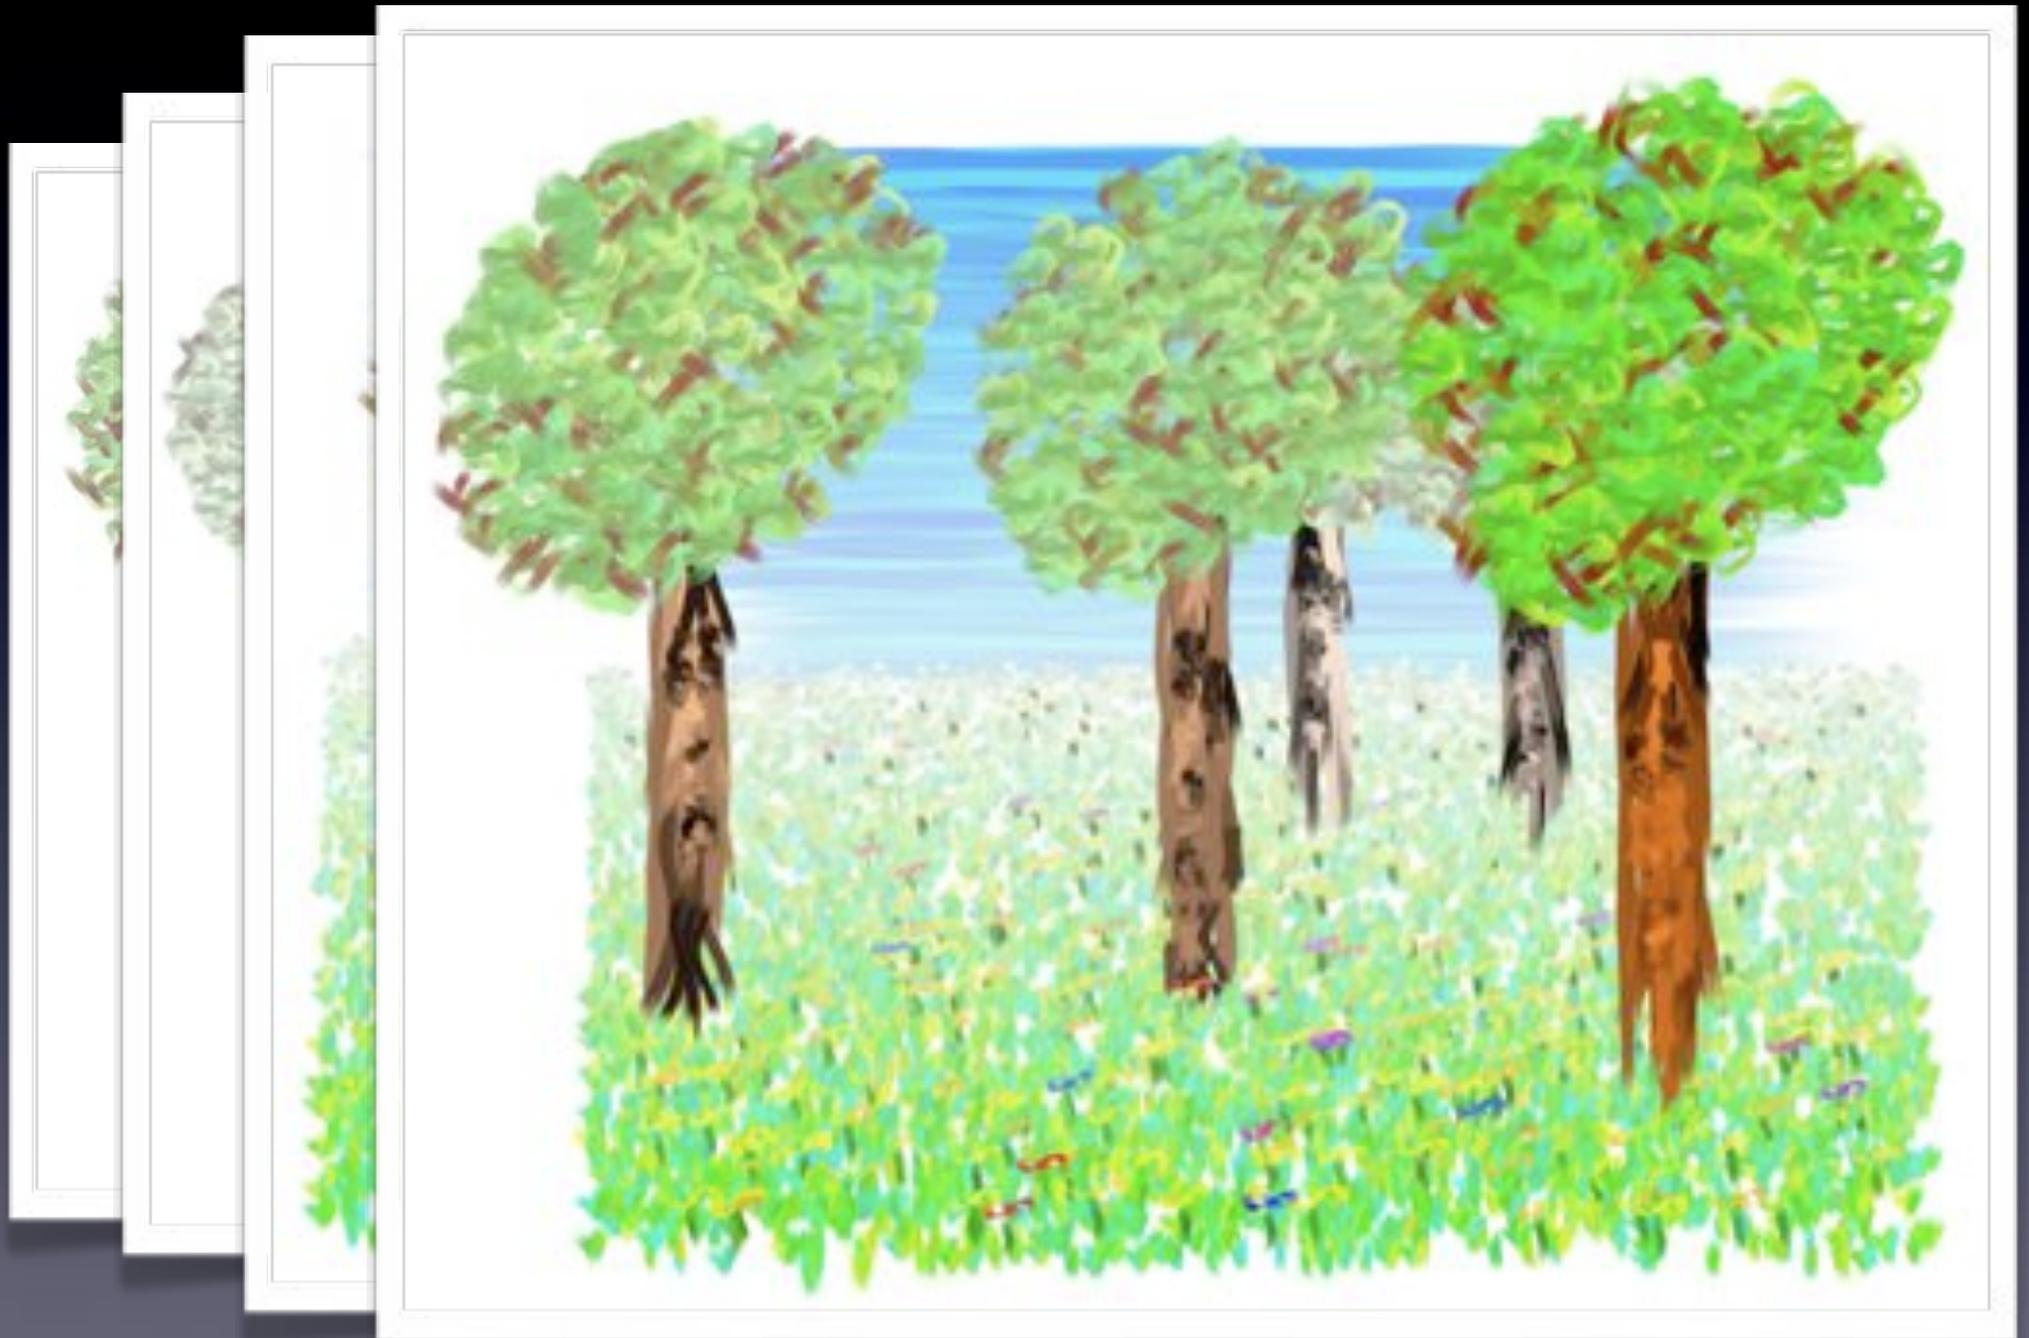
A Teaching Interface

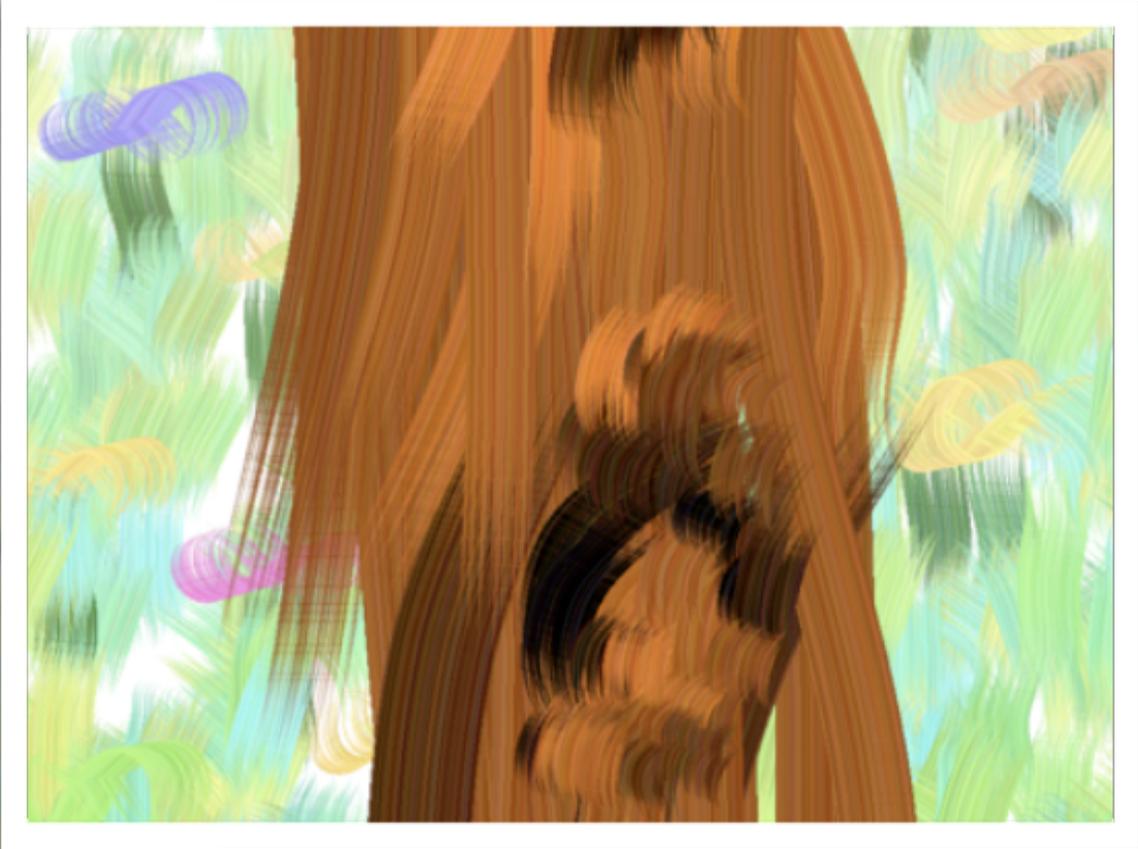


CFDG Module



PresidENTS Series





Fish Fingers Series



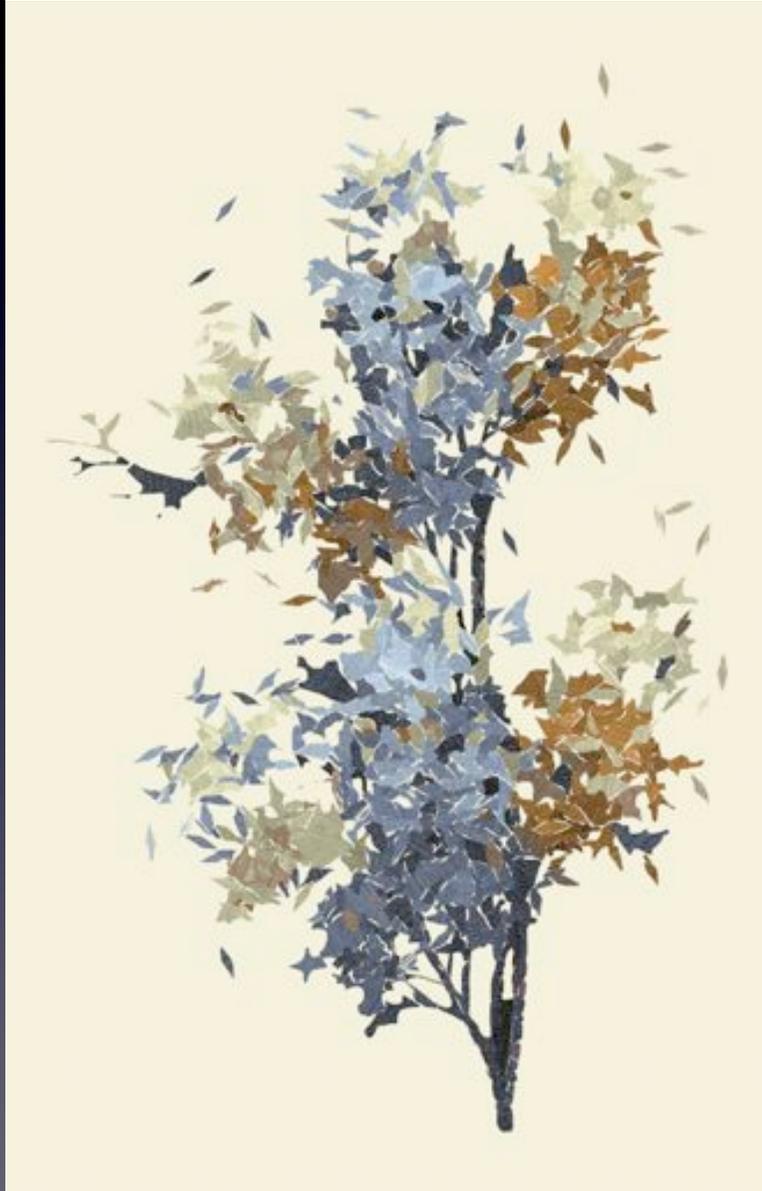
After AARON Series



Countryside Series



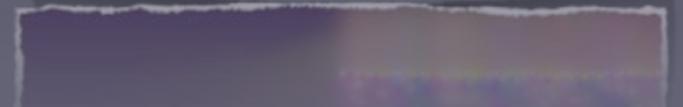
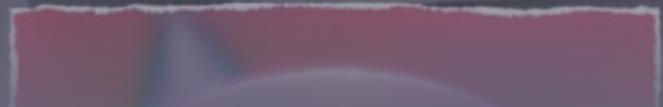
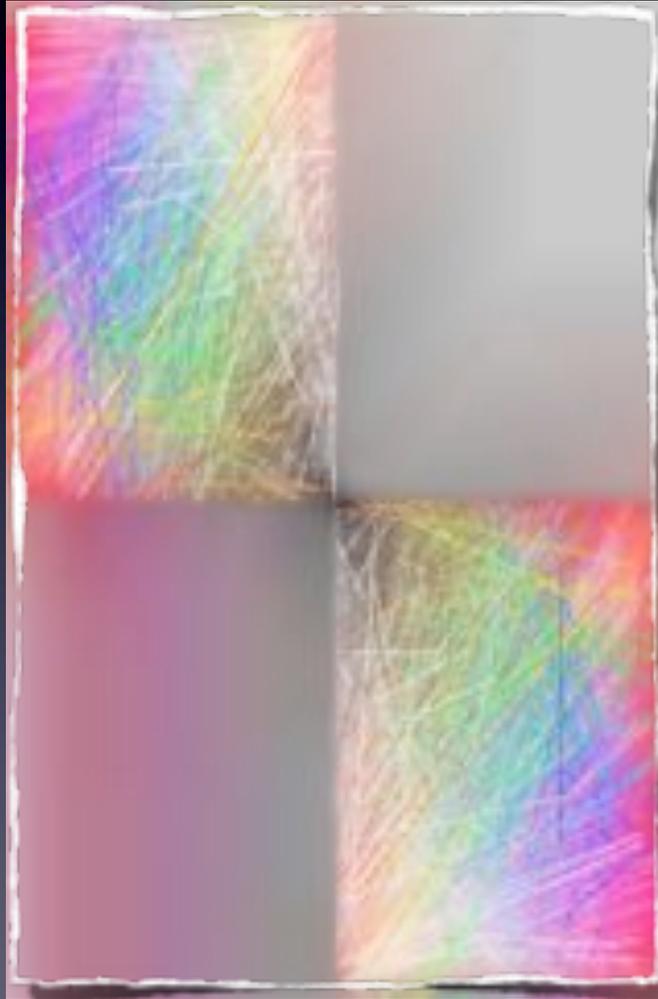
Countryside Series



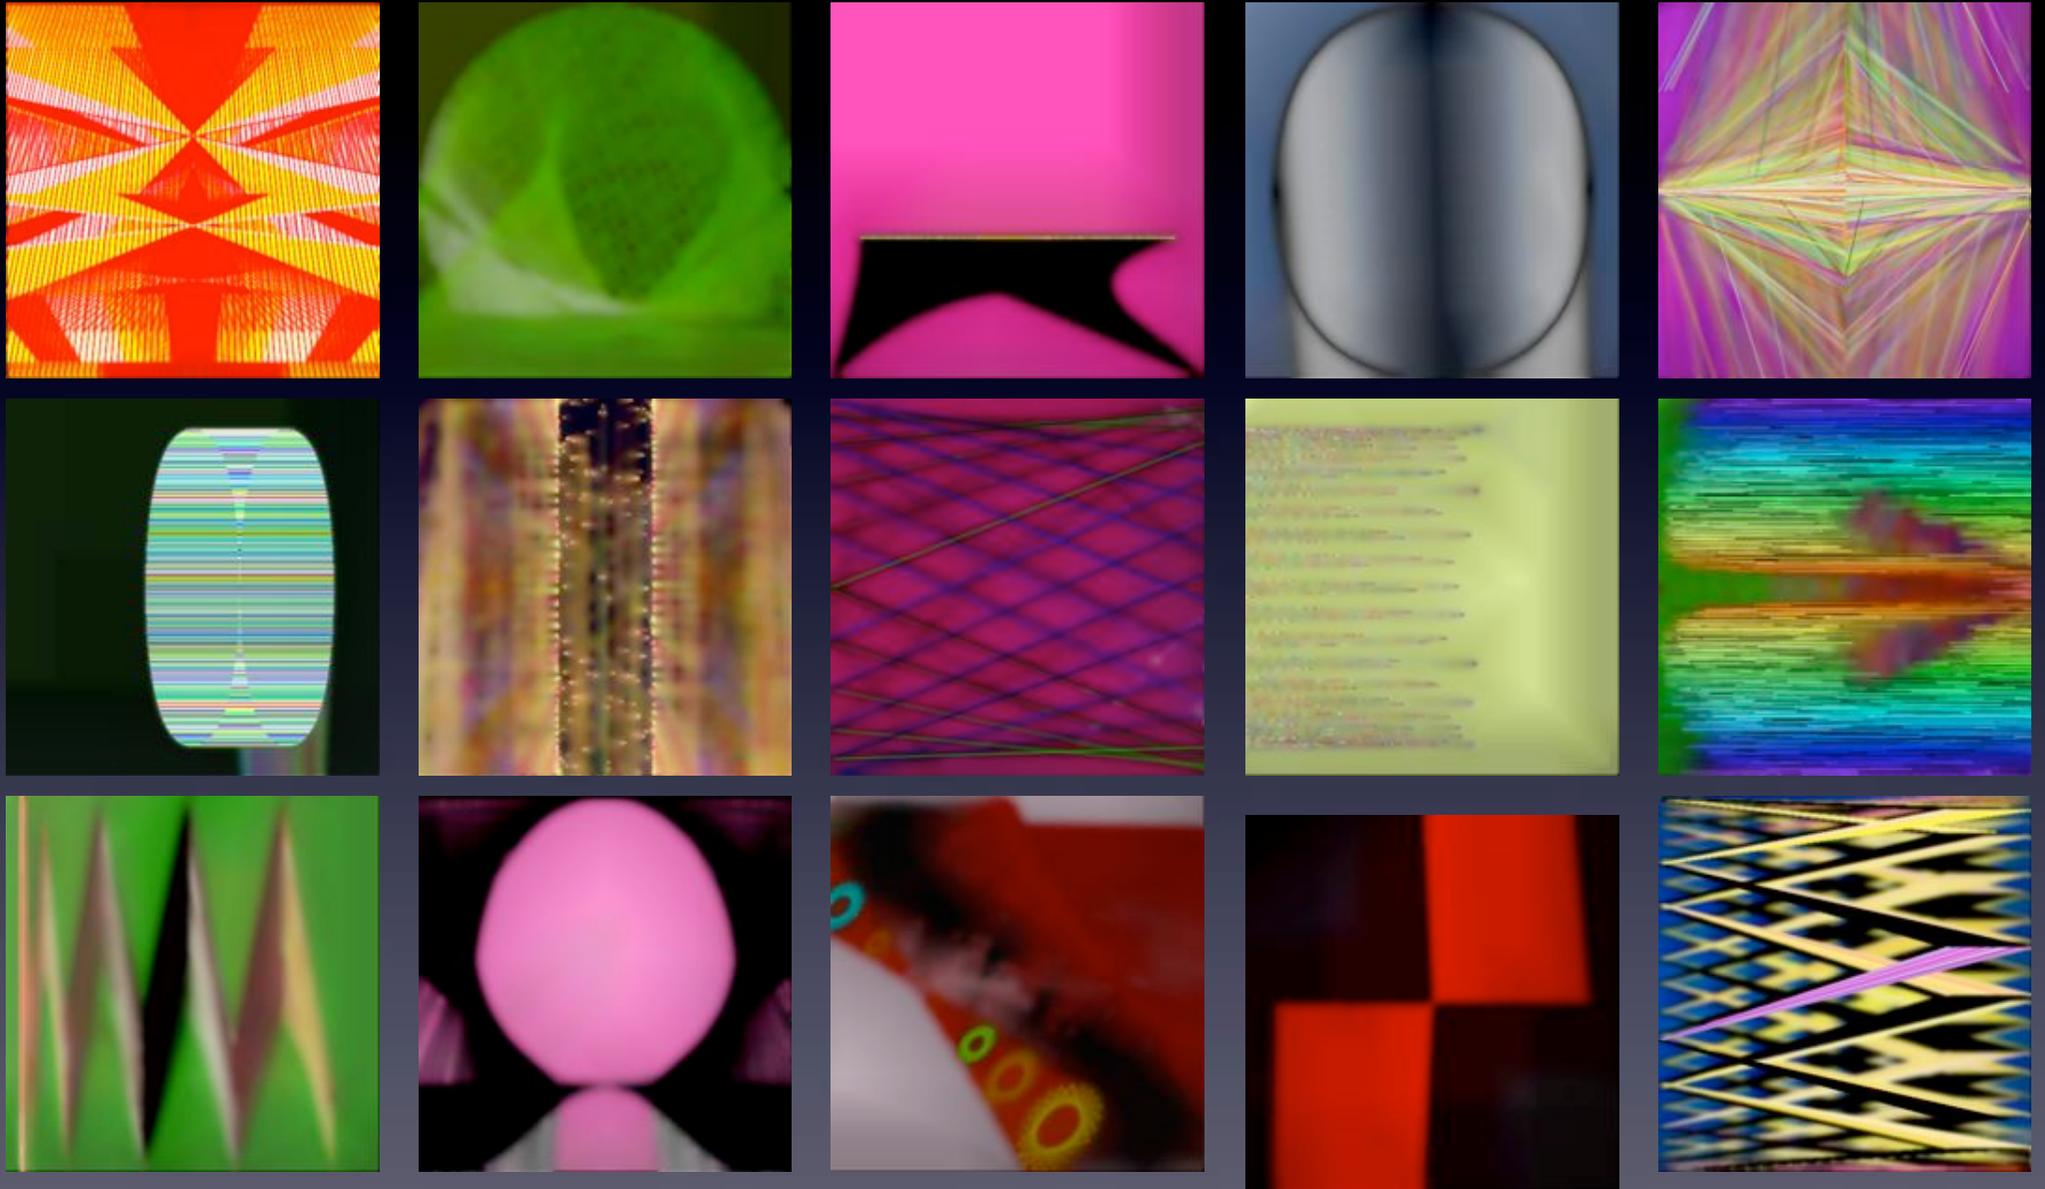


Evolutionary Art Module

Ask me for details about ELVIRA



A Particle-Based Wrapper



There is amazing visual variation in this wrapper

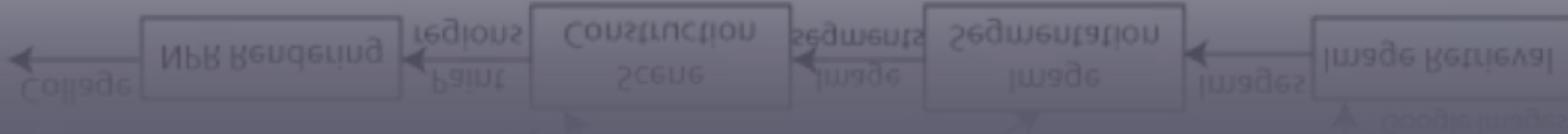
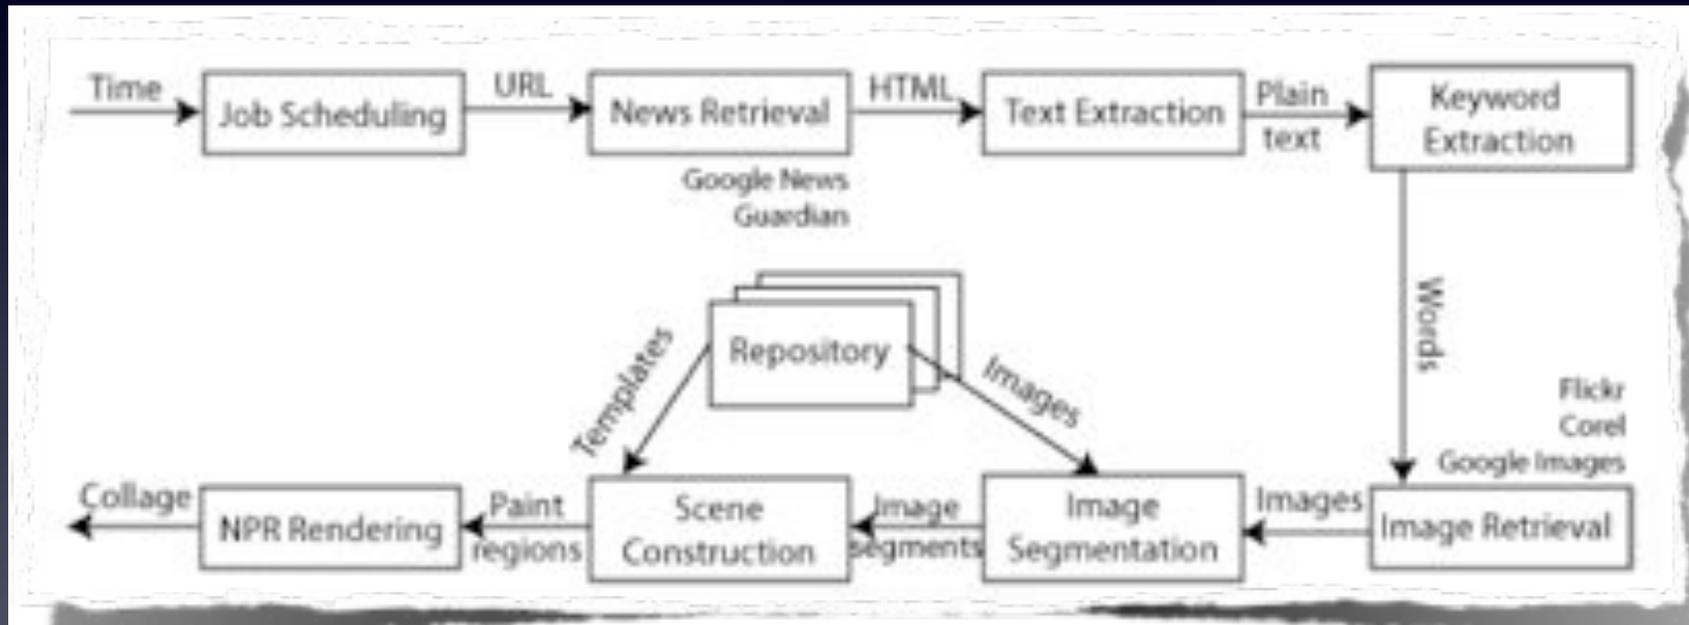
Stage 3.3

Automatic Scene Construction

Constructing scenes for a purpose

- 2009 MSc. Project of Anna Krzeczowska
- Joint with Jad El-Hage and Stephen Clark (CompLab, Cambridge)
- Deals (partially) with the question of *intent* in creative systems:
 - Usually the user supplies all the intentionality in creative projects, software acts as a (possibly intelligent) tool
- Idea is to react to news stories to paint pictures which make people think more about certain issues

Automated Collage Generation

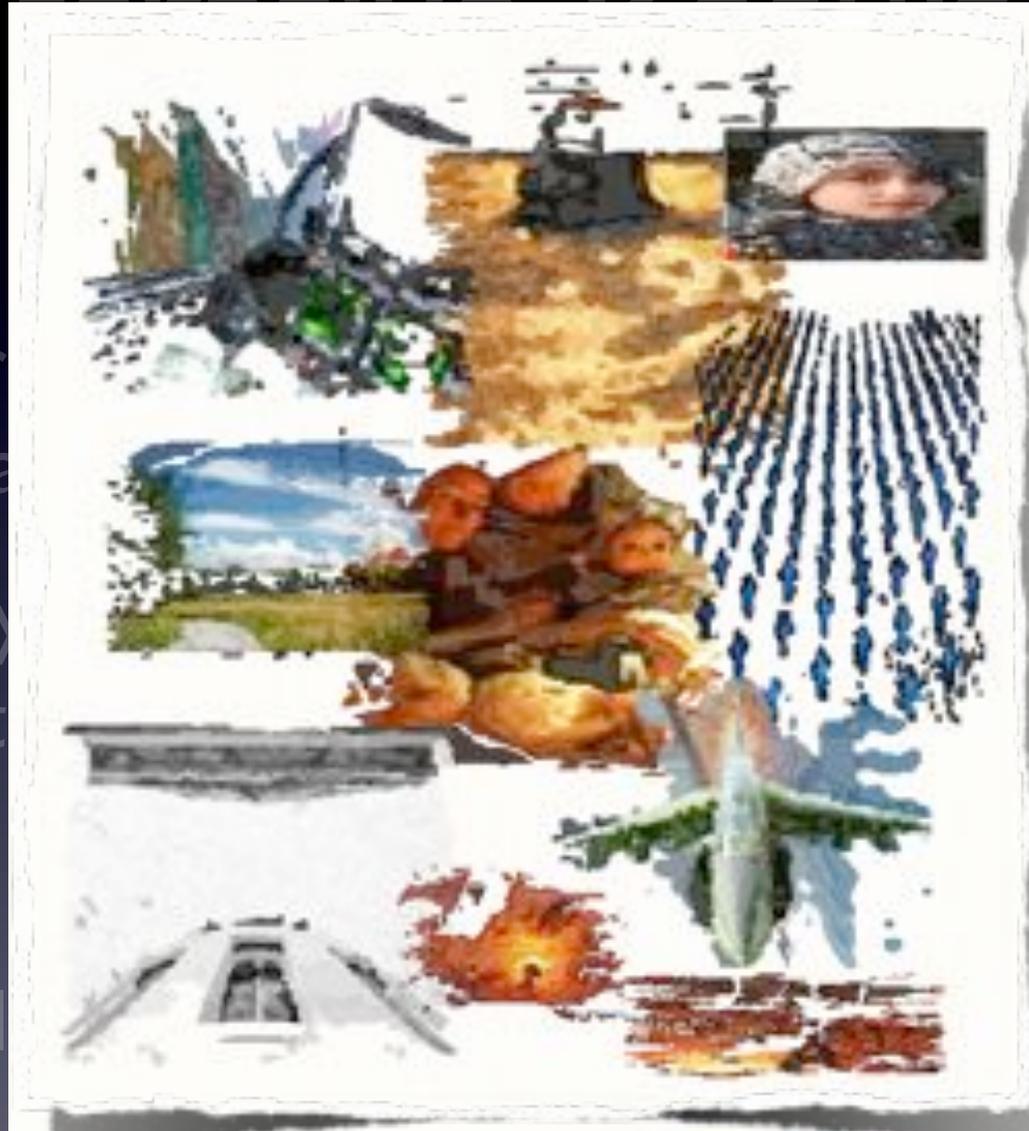


Illustrative Result

- Guardian news story on the war in Afghanistan
 - Keywords: *afghanistan, brown, forces, troops, nato, british, speech, country, more and afghan*
 - Images retrieved from Flickr
- Highlights the potential for poignancy

Illustrative Results #2

- Guard
Afghan
- Key
nat
- Highl



es, troops,
and afghan

ncy



,
in



#2

es, troops,
and afghan

ncy



n

Illustrative Result #2



war in

, forces, troops,
more and afghan

ignancy

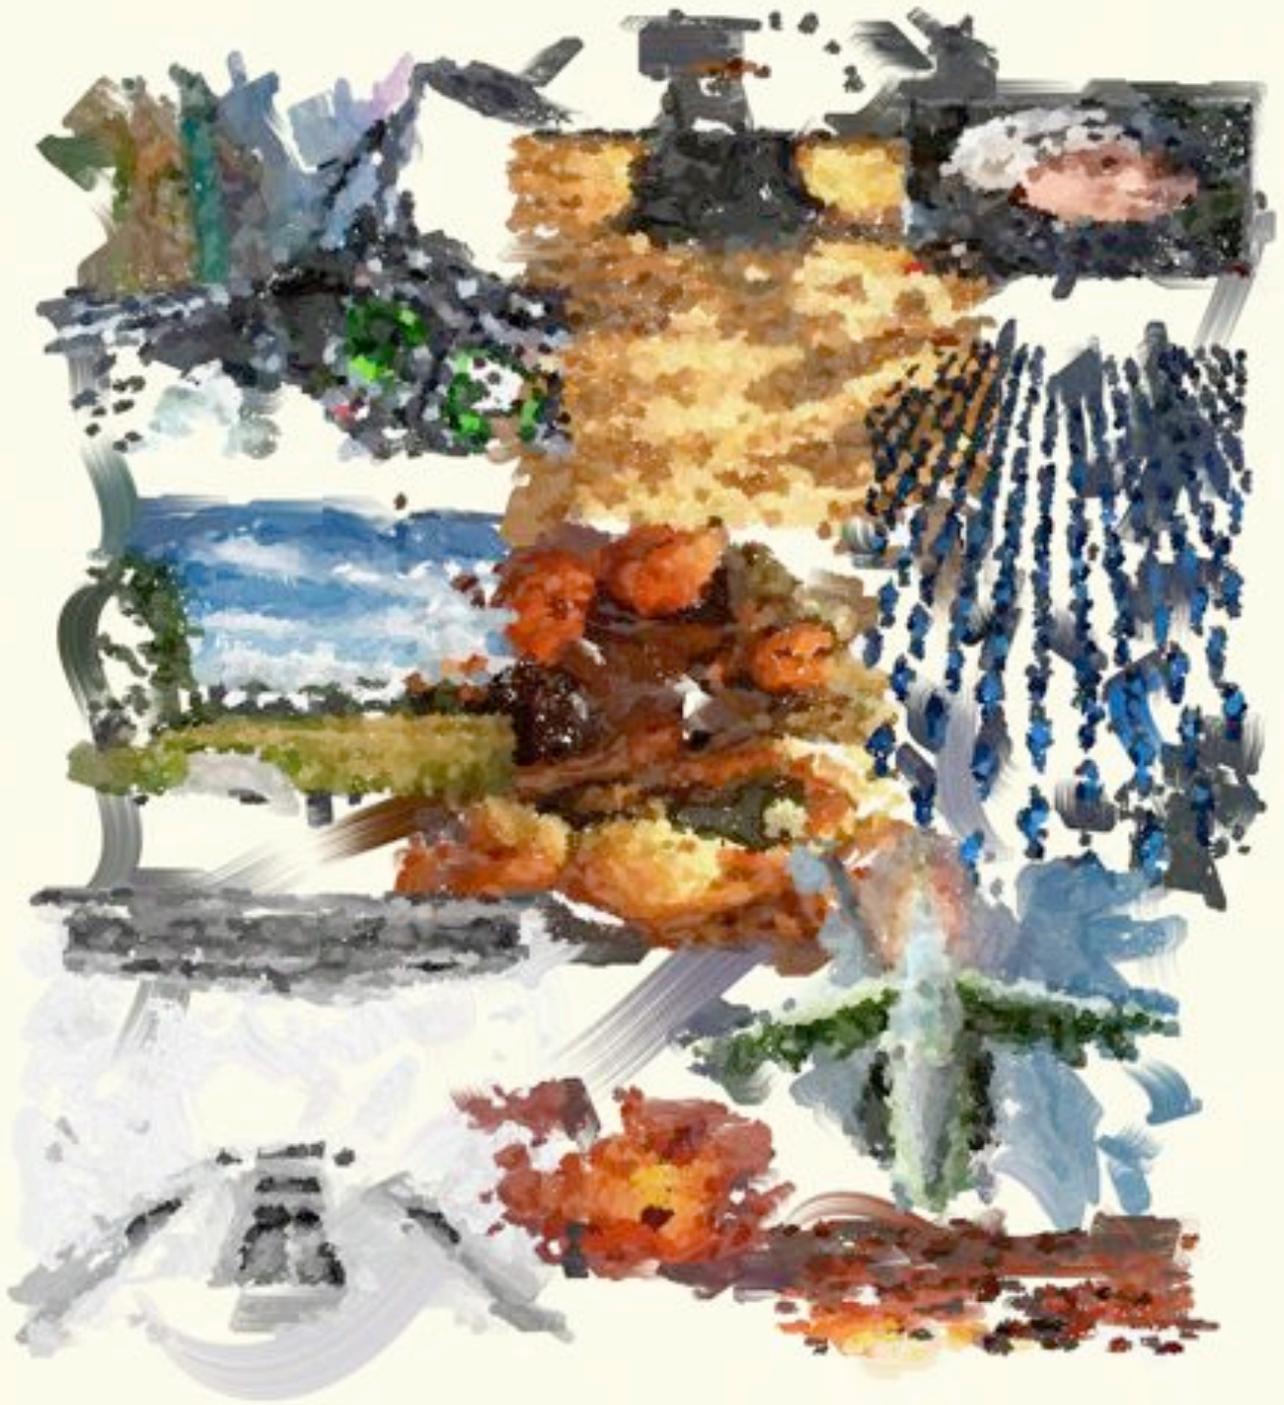


Unit #2

er in

forces, troops,
ore and afghan

gnancy



Interesting Observation

- All of these words are affect-neutral
 - *afghanistan, brown, forces, troops, nato, british, speech, country, more and afghan*
- Raises questions of intentionality
 - 5 agents at work here: programmer, software, journalist, audience, flickr taggers
- In stage 5, we increased the explicit intentionality of The Painting Fool through portraiture

Stage 4
Multimedia
(Skill and Imagination)

Multimedia

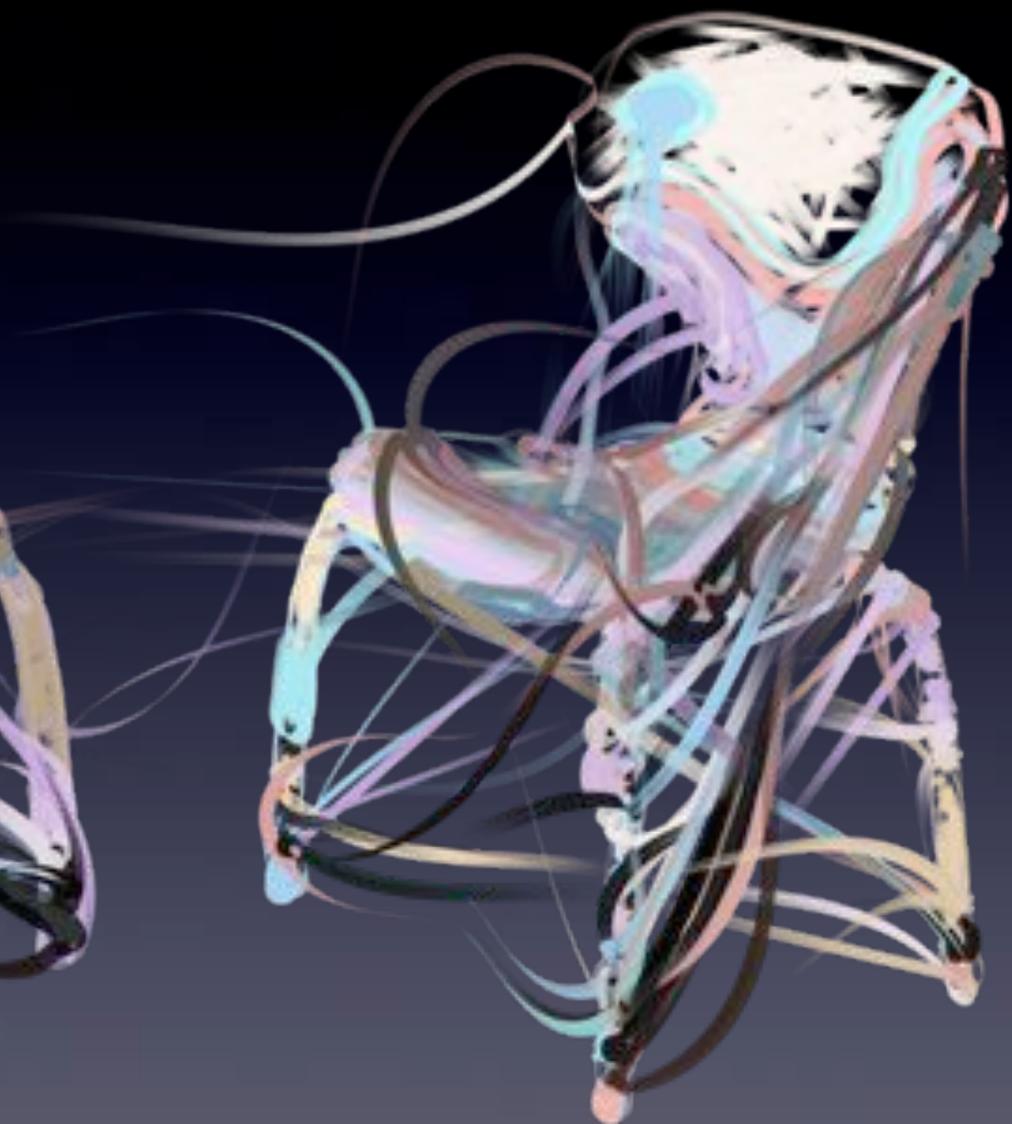
- Moving into video, text, sculpture and gallery design
 - Into the third and fourth dimensions!
- Aims:
 - To enable The Painting Fool to produce art that couldn't be produce by people
 - Enable it to produce more sophisticated artworks
 - With a knowledge of space/light/arrangement

Sculptures

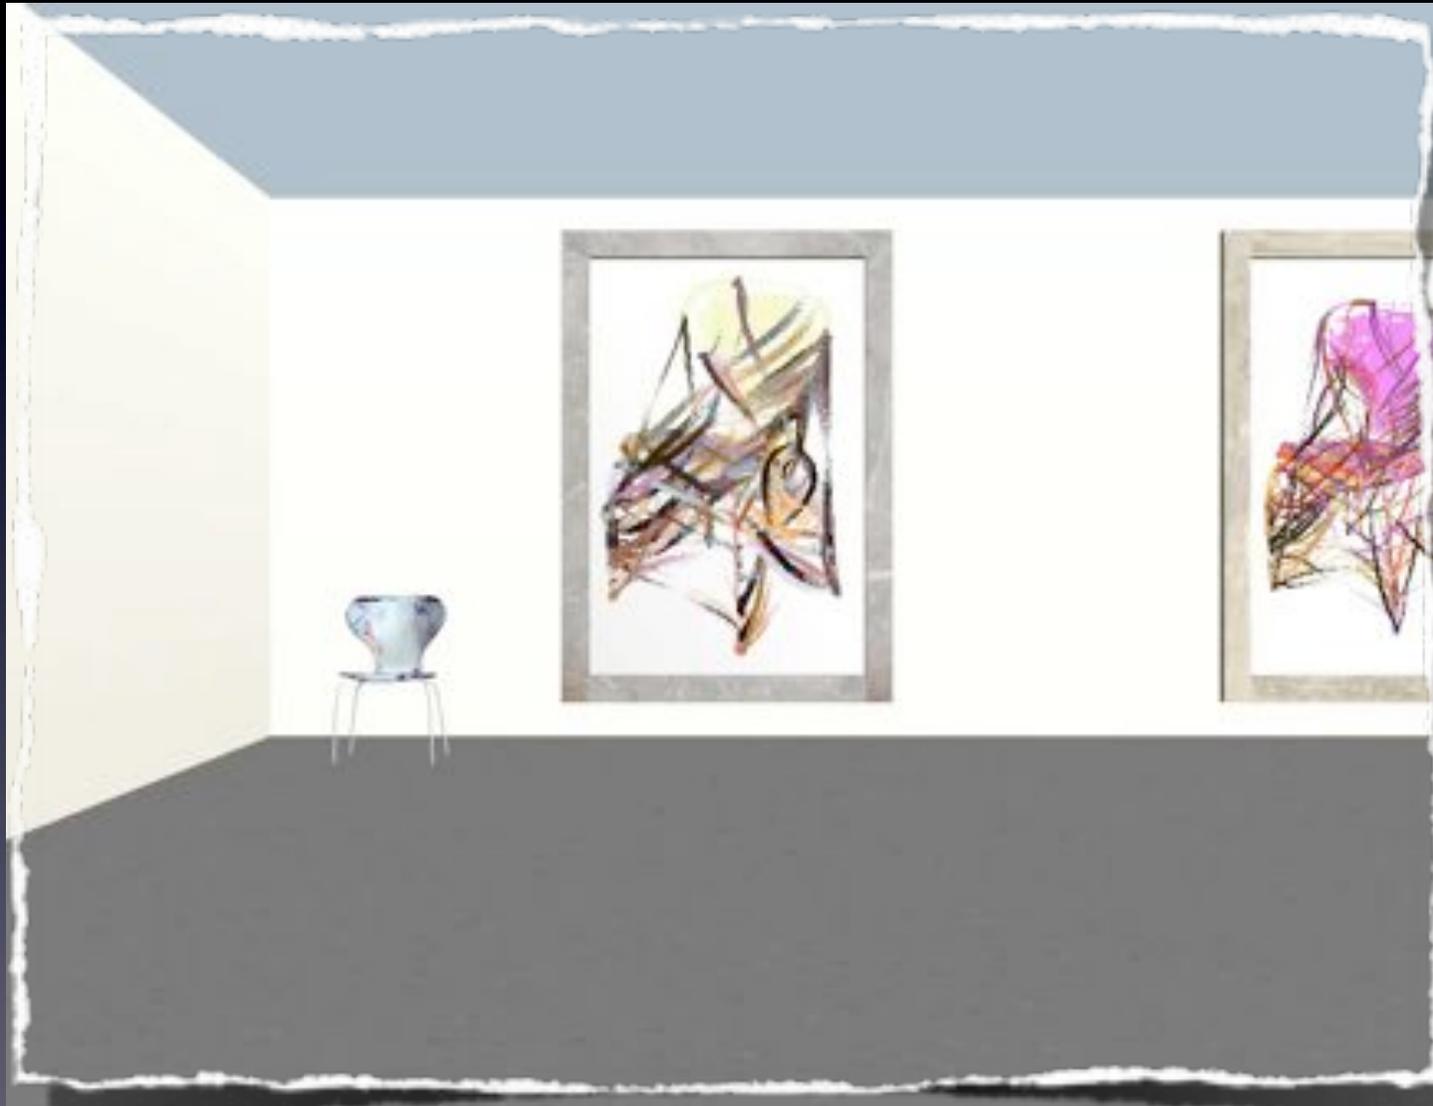


3D Scene Generation





Gallery Design



Video Pieces

Paint Dance Animations



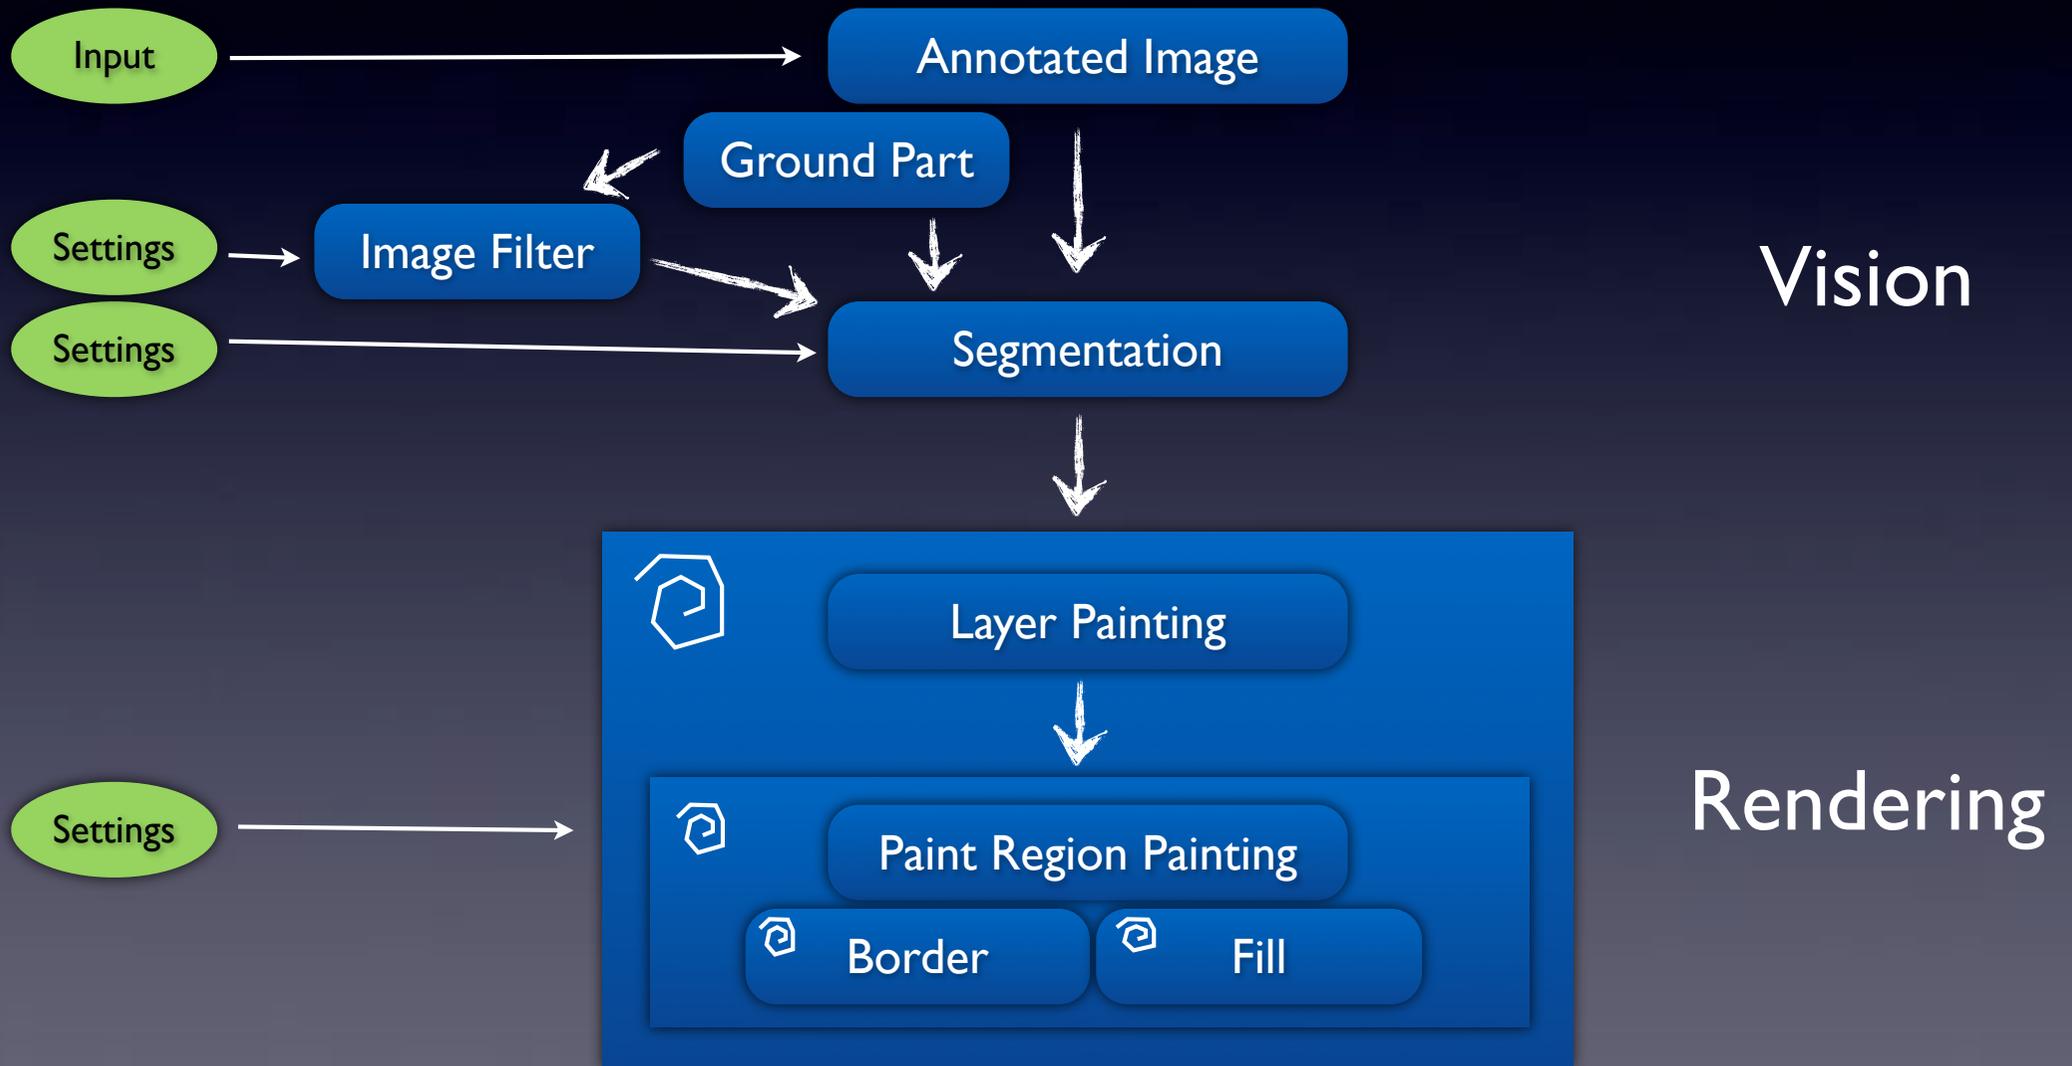
Paint Dance Animations



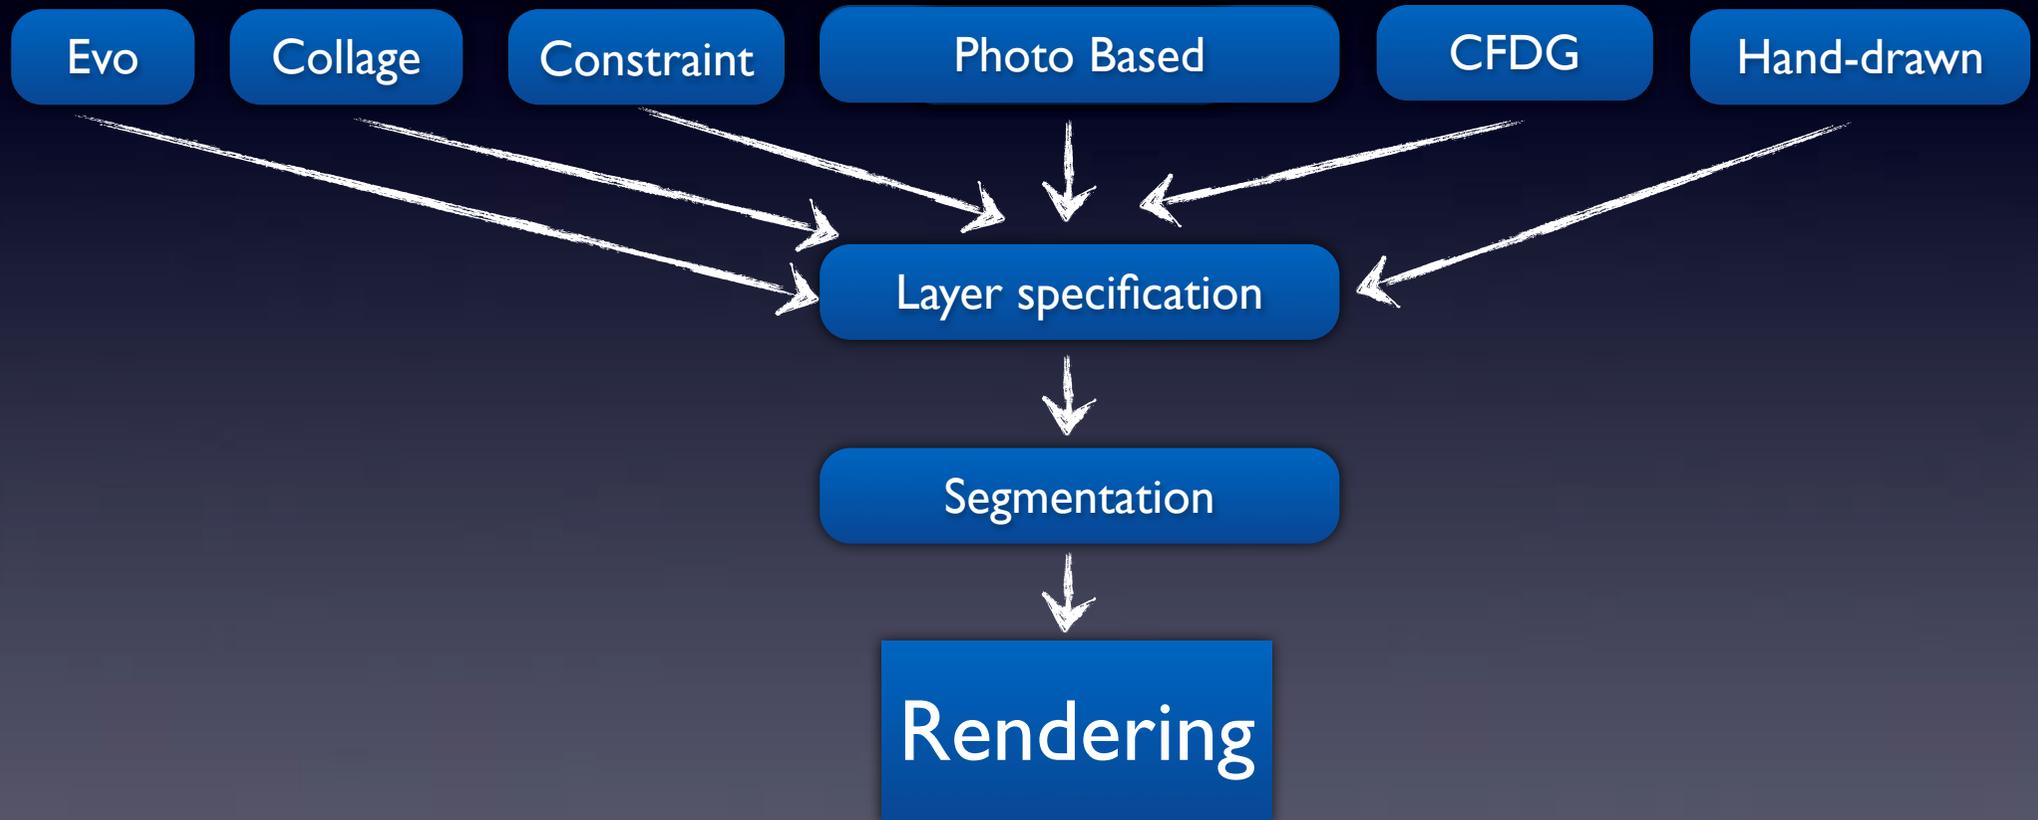
Technical Issues

- Choosing a representative subset of paint strokes from the set of original paintings
 - K-means clustering
- Assigning strokes from the representative set back to the paintings without losing too much fidelity
- Producing smooth animations without a vertical crossover effect

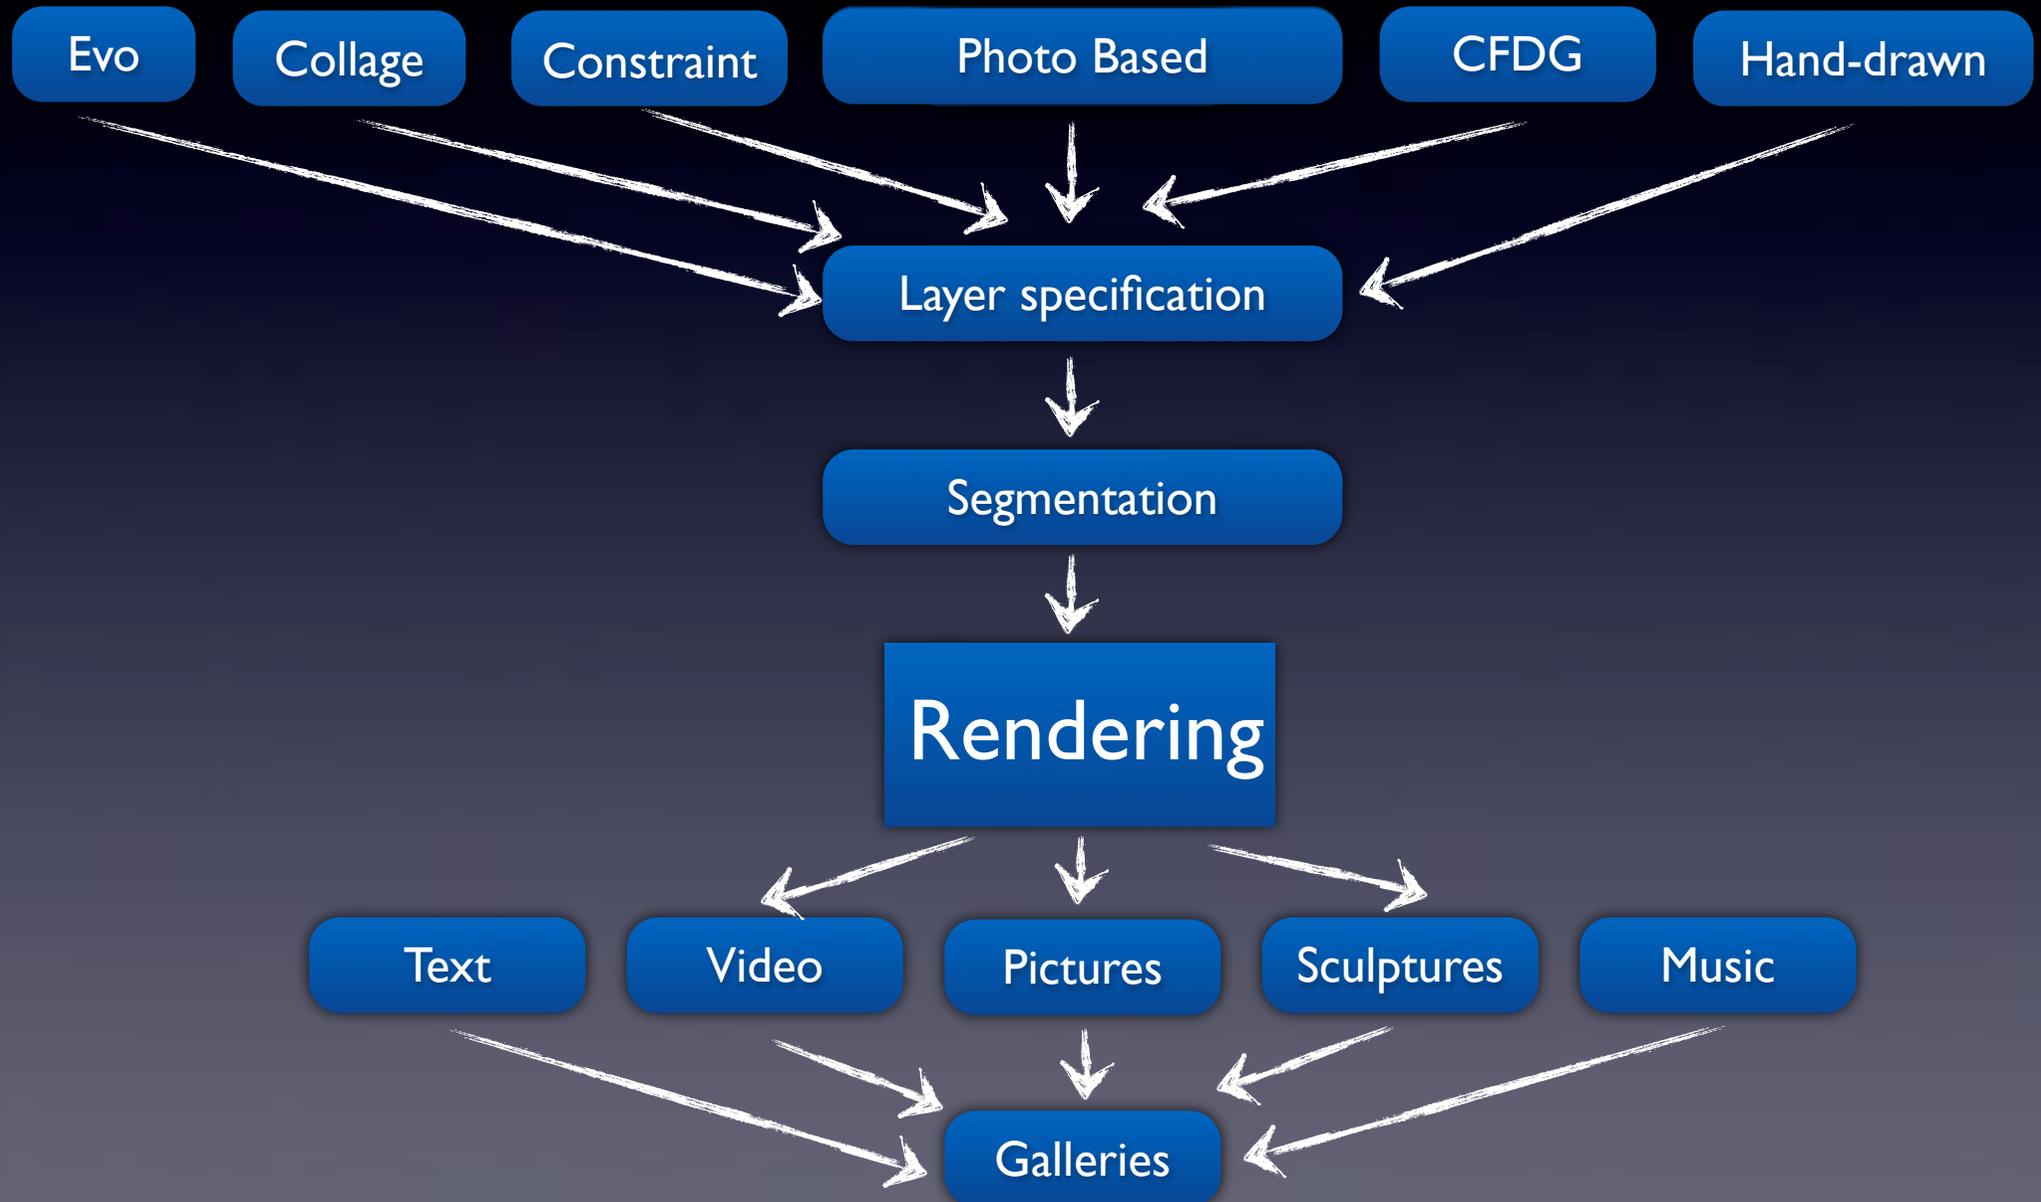
Skill-based Graphics Workflow



Imagination-based Graphics Workflow



Multimedia-based Graphics Workflow



Let's take a breather...

Stage 5
Live Portraiture
You Can't Know my Mind

(Skill, Appreciation, Imagination, Intentionality, Reflection and Learning)

Joint work with Dan Ventura
Brigham Young University, USA

Aims of the Subproject

- We want to see if people will accept the software as an artist
 - Not necessarily a good or particularly creative artist
 - In particular, we wanted to see if there any easy criticisms left, other than: “it’s not a person, so it can’t be artistic/creative”
- Manage the perceptions of people interacting with the software:
 - Explicitly managed people’s impression of the creative behaviours required to avoid being called uncreative
 - Explicitly managed people’s perception of the (physical) artistic process
- Made it clear that people are not controlling the software to produce an image, rather it is controlling them to express itself
 - Introduce the idea of autonomous creativity in software by presenting people with a model of *accountable unpredictability* in The Painting Fool

Overview

- It's portrait painting: people sit in front of a laptop and some of them come away with a printed portrait, after watching a hand paint/draw on-screen
- The 'mood' of the The Painting Fool determines whether or not a portrait is painted, and the nature of the portraits that are produced
- A commentary is also given, where The Painting Fool describes what its mood was, why, its intention for the portrait and whether it has achieved its goals
- The successes and failures are recorded, and the software learns from these, trying out new approaches when in a reflective mood
- Wall text is used to prompt the portrait sitters and audiences to think about the behaviorus being expressed by The Painting Fool



You Can't Know my Mind

A Festival of Computational Creativity

Computational Creativity offers the promise of changing how we view music, literature, the sciences and the arts, and research in Artificial Intelligence may one day lead to us happily co-existing with creative computers. As part of a festival of Computational Creativity, we present several evening events highlighting the idea of the computer as a creative, intelligent, collaborator. Right now, Computational Creativity asks the question: "can computers create?", and we hope you will be stimulated into thinking about our technological future through this festival. In addition to the visual arts produced by computer - as shown in this gallery - it is also interesting to ask the Computational Creativity question about areas such as music, poetry and even cooking.

Moody Music

We present an evening of electronic music with Stéphane Bissières, in which he will play three sets while reacting to musical input from software based on a model of mood. In each set, the type of input will be different and together Stéphane and the computer will generate three different styles of music. In each of the three sets, the music is supposed to reflect the mood of the computer in different ways using musical constructs such as volume, reverberation, decay and choice of sound palette. See if you can hear the moods reflected from the moody computer.

Paris-based electronic musician Stéphane Bissières won the Paris Youth Talents competition in 2007 and a Sacem et Imeb prize in electroacoustic composition. An improviser for France Musique since 2005 and a sound designer for France Inter., he developed motion capture interfaces for the DAFACOT society and has realized interactive installations for Le Fresnoy (Roubaix, FR). Today his work revolves mainly around the creation of electroacoustic music in real-time, augmented with analog processes.

Poems and Potage

An evening of poetry will feature readings from a selection of both computer-generated and human crafted poems. These will be curated and read by Russell Clarke, who will also discuss and relate the poems. Since graduating from Christ Church, Oxford with a degree in English Language and Literature, Russell has combined tutoring and teaching with miscellaneous literary activities such as editing a work of philosophy and writing occasional academic articles.

We will also showcase a human/computer collaboration in gastronomy in which some computationally produced soups are paired with food created by chef Sophie Grilliat. Computer program PIERRE will generate soup recipes cooked over an afternoon, and Sophie will choose three to incorporate into a menu including dishes she prepares to complement them. While finishing a PhD in Urbanism at la Laboratoire Technique Territoire et Société/Université Paris-Est, Sophie has concurrently developed a culinary enterprise as a chef à domicile for two years and obtained a CAP diploma. Her approach is characterised by an insistence upon the integrity of fresh produce, sustainably sourced and the search for a balance between culinary technique and the social dimension inherent in serving food.

Computational Creativity is the philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative.

You Can't Know my Mind

The Painting Fool is software that we want to be taken seriously as a creative artist in its own right, one day. It has been trained for the last 10 years to take on more and more creative responsibility, employing techniques from artificial intelligence, computer graphics and machine vision, guided always by art theory and philosophy. In the You Can't Know my Mind exhibit, The Painting Fool will produce portraits of people who sit for it. The software has reached a level of maturity where it is now impossible to determine in advance what it will produce, and why, and the software makes all its own decisions about what to paint, how to use its materials and, working with computational critic DARCI, together they decide whether or not the final product is suitable. As demonstrated here, The Painting Fool exhibits behaviours which have emerged in Computational Creativity research as being important for the perception of creativity in software. In particular, the software exhibits:

Intentionality: By attempting to produce a portrait which either explores its materials and techniques; expresses its mood; or tries to improve upon failed experiences, the software shows that it knows what it wants to do.

Imagination: By using its mood to conceptualise a portrait, the software produces a new backdrop for each portrait, changing colours and abstraction levels in the face in order to express that mood.

Skill: By painting live on-screen with simulated paint brushes, pencils and pastels, the software shows that it has all the artistic skills required to produce interesting and characteristic portraits.

Appreciation: By using artificial neural networks learned through direct experience with artists and art lovers, the system is able to tell visually whether it has achieved the aims dictated by its mood.

Reflection: By criticising failures to express a mood, the software reflects on its artistic process and tries later to improve by producing a new portrait expressing the same mood, but using different conceptualisations and painting styles.

Learning: By recording successes and failures, the software is able to constantly learn about how to express mood through its processes—how different conceptualisations and painting styles produce more or less mood in pieces, making the software unpredictable in a non-random way.

You can't know its mind.

The You Can't Know my Mind exhibit has been realised by: Blanca Pérez Ferrer, Dan Ventura and Simon Colton. Blanca is a historian of arts who has a masters degree in cultural mediation from the École du Louvre and is currently working as a cultural mediator and curator for the Ligue de l'Enseignement. Dan is a Professor in the Computer Science Department at Brigham Young University. His research focuses on creating artificial intelligent systems that incorporate robustness, adaptation and creativity and includes work on systems for understanding visual art, automatic music composition and even recipe generation. He is the leader of the group producing the DARCI artificial art critic. Simon is a Professor of Computational Creativity in the department of Computing at Goldsmiths College, London. His research focuses on getting software to take on creative responsibilities in the arts and sciences, and the philosophical discussions that this raises. He has worked in autonomous creative systems in graphic design, video game design, pure mathematics, literature and the visual arts. He is the author of The Painting Fool software which has produced the artwork in this exhibition.

The Painting Fool

www.thepaintingfool.com

its own decisions about what to paint, how to use its materials and, working with computational critic DARCI, together they decide whether or not the final product is suitable. As demonstrated here, The Painting Fool exhibits behaviours which have emerged in Computational Creativity research as being important for the perception of creativity in software. In particular, the software exhibits:

Intentionality: By attempting to produce a portrait which either explores its materials and techniques; expresses its mood; or tries to improve upon failed experiences, the software shows that it knows what it wants to do.

Imagination: By using its mood to conceptualise a portrait, the software produces a new backdrop for each portrait, changing colours and abstraction levels in the face in order to express that mood.

Skill: By painting live on-screen with simulated paint brushes, pencils and pastels, the software shows that it has all the artistic skills required to produce interesting and characteristic portraits.

Appreciation: By using artificial neural networks learned through direct experience with artists and art lovers, the system is able to tell visually whether it has achieved the aims dictated by its mood.

Reflection: By criticising failures to express a mood, the software reflects on its artistic process and tries later to improve by producing a new portrait expressing the same mood, but using different conceptualisations and painting styles.

Learning: By recording successes and failures, the software is able to constantly learn about how to express mood through its processes---how different conceptualisations and painting styles produce more or less mood in pieces, making the software unpredictable in a non-random way.

You can't know its mind.

The You Can't Know my Mind exhibit has been realised by: Blanca Pérez Ferrer, Dan Ventura and Simon Colton. Blanca is a historian of arts who has a masters degree in cultural mediation from the École du Louvre and is currently working as a cultural mediator and curator for the Ligue de l'Enseignement. Dan is a Professor in the Computer Science Department at Brigham Young

Mood Model

- We argue that software needs to add value to its work by framing it in terms of describing what its done and why, and what it's achieved (see my later lectures)
- Unpredictable behaviour leading to surprise in audiences has always been an desirable quality of creative systems
 - But random number generation has zero cultural value, i.e., if it turns out that the root decision was caused by a trivial process
 - Software can be unpredictable but not (yet) capricious
- So software needs to be equipped with *unpredictable yet accountable* behaviours which influence its generative processes and produces surprising results, yet enable the production of sensible explanations for its behaviour

Mood Model

- The Painting Fool has offline access to all 120,000 Guardian (newspaper) articles that were posted online last year
- When it starts, it chooses the article posted most closely to the *time of day* at the time of initiation. This provides a seed which is accountable (if a little arbitrary)
- It uses a simple sentiment analysis approach to estimate whether the article has positive, negative or neutral valency
- It then extracts keyphrases using TextRank (similar to PageRank)
 - It then uses these keyphrases to search for other articles, and chooses the first of those which hasn't been seen before
 - It doesn't allow the same keyphrase from an article to be used twice, which cuts down on repeated articles. When exhausted, it re-initialises
- Details: it has mood stabilising and balancing routines to smooth things out and try to ensure that it regularly goes through the full range of valencies

memory Cook [http://www.filters.com/3d/images/analysis/segmentation/Items/galleries/](#)
[scenes/pictures/installations/press/music/sculptures/3dch/galleries/documents/artist/scene.3d.do](#)



Description



Navigation controls including a 'New' button with a right-pointing arrow and a 'Go' button with a right-pointing arrow.



Description

I read the article: "Fashion tycoon Kevin Stanford plans House of Fraser comeback" in the business section
 2012-04-22_business_kevin-stanford-house-of-fraser
 This was chosen by time of day: 18:03:01".
 My mood valency is -0.054481547 (very_negative).
 My justification for being in this mood is: I was reading an article entitled: "Fashion tycoon Kevin Stanford plans House of Fraser comeback" in the business section of the Guardian newspaper, which was really sad, because it
 Moods so far: very_positive=0 positive=0 experimental=0 reflexive=0 negative=0 very_negative=1
 Number of different articles read: 1
 Number of mood updates: 1

I read the article: "Times editor faces Leveson inquiry recall over NightJack hacking claim" in the media section.
 2012-02-02_media_thetimes-pressandpublishing
 This was chosen by cross reference with the keyphrase "guardian last month".
 My mood valency is -0.02058899 (negative).
 My justification for being in this mood is: I was reading an article entitled: "Fashion tycoon Kevin Stanford plans House of Fraser comeback" in the business section of the Guardian newspaper, which was really sad, because it
 Moods so far: very_positive=0 positive=0 experimental=0 reflexive=0 negative=1 very_negative=1
 Number of different articles read: 2
 Number of mood updates: 2

I read the article: "Leveson inquiry: Lord Patten and Mark Thompson appear" in the media section.
 2012-01-23_media_leveson-inquiry-lord-patten-mark-thompson-five
 This was chosen by cross reference with the keyphrase "labour mp tom watson".
 My mood valency is -0.014314287 (negative).

My mood valency is -0.014314287 (negative).
 This was chosen by cross reference with the keyphrase "labour mp tom watson".
 2012-01-23_media_leveson-inquiry-lord-patten-mark-thompson-five
 I read the article: "Leveson inquiry: Lord Patten and Mark Thompson appear" in the media section.
 My mood valency is -0.014314287 (negative).



Upper: 367 (0.017, 0.067)
 Downer: 276 (-0.02, -0.07)

Description

I read the article: "Fulham v Arsenal - as it happened" in the football section.
 2012-01-02_football_fulham-arsenal-premier-league-mbm
 This was chosen by cross reference with the keyphrase "as".
 My mood valency is 0.0042691526 (reflexive).
 My justification for being in this mood is: I was reading an article entitled: "Rory McIlroy set to test Tiger Woods in Abu Dhabi Championship" in the sport section of the Guardian newspaper, which was really happy, because it
 Moods so far: very_positive=16 positive=6 experimental=4 reflexive=4 negative=15 very_negative=12
 Number of different articles read: 48
 Number of mood updates: 57

I read the article: "Premier League team news for bank holiday Monday" in the football section.
 2012-01-01_football_premier-league-team-news
 This was chosen by cross reference with the keyphrase "fulham v arsenal".
 My mood valency is -0.002989117 (negative).
 My justification for being in this mood is: I was reading an article entitled: "Premier League team news for bank holiday Monday" in the football section of the Guardian newspaper, which was really sad, because it spoke of "ca
 Moods so far: very_positive=16 positive=6 experimental=4 reflexive=4 negative=16 very_negative=12
 Number of different articles read: 49
 Number of mood updates: 58

I read the article: "Bradley Johnson's first-half header gives Norwich victory over Stoke" in the football section.
 2012-11-03_football_norwich-city-stoke-city-premier-league
 This was chosen by cross reference with the keyphrase "full-back andy wilkinson".
 My mood valency is 0.0031135567 (reflexive).

↑ New

↓ New

I read the article: "Bradley Johnson's first-half header gives Norwich victory over Stoke" in the football section.
 2012-11-03_football_norwich-city-stoke-city-premier-league
 This was chosen by cross reference with the keyphrase "full-back andy wilkinson".
 My mood valency is 0.0031135567 (reflexive).
 My justification for being in this mood is: I was reading an article entitled: "Premier League team news for bank holiday Monday" in the football section of the Guardian newspaper, which was really sad, because it spoke of "ca
 Moods so far: very_positive=16 positive=6 experimental=4 reflexive=4 negative=16 very_negative=12
 Number of different articles read: 49
 Number of mood updates: 58

I read the article: "Paperback fiction: February fiction" in the books section

2012-02-21_books_paperback-roundup-february-fiction

This was chosen by time of day: 17:30:31".

My mood valency is -0.028735632 (very_negative).

My justification for being in this mood is: I was reading an article entitled: "Paperback fiction: February fiction" in the books section of the Guardian newspaper, which was really sad, because it spoke of "red pen".

Moods so far: very_positive=0 positive=0 experimental=0 reflexive=0 negative=0 very_negative=1

Number of different articles read: 1

Number of mood updates: 1

I read the article: "Paperback Q&A: Esi Edugyan on Half Blood Blues" in the books section.

2012-02-07_books_paperback-q-a-esi-edugyan

This was chosen by cross reference with the keyphrase "half blood blues".

My mood valency is -0.0111069465 (negative).

My justification for being in this mood is: I was reading an article entitled: "Paperback fiction: February fiction" in the books section of the Guardian newspaper, which was really sad, because it spoke of "red pen".

Moods so far: very_positive=0 positive=0 experimental=0 reflexive=0 negative=1 very_negative=1

Number of different articles read: 2

Number of mood updates: 2

I read the article: "Miami Heat 100 - Oklahoma City Thunder 96 - as it happened" in the sport section.

2012-06-15_sport_miami-heat-oklahoma-city-thunder-live

This was chosen by cross reference with the keyphrase "carpal tunnel syndrome".

My mood valency is 0.008128183 (very_positive).

My justification for being in this mood is: I was reading an article entitled: "The mother of all maternity shoes" in the fashion section of the Guardian newspaper, which was really happy, because it spoke of "great option".

Moods so far: very_positive=2 positive=0 experimental=0 reflexive=0 negative=1 very_negative=1

Number of different articles read: 4

Number of mood updates: 4

Mood Usage

- When someone sits down for a portrait, the software assesses which of six moods it is in:
 - Very positive, positive, experimental, reflective, negative or very negative
- If the mood is very negative, then The Painting Fool tells the person to go away, because it's not appropriate to paint a picture
 - But it does provides a printout explaining its decision, in terms of its mood, and a particularly negative keyphrase it has recently extracted
 - Remember: you can't know The Painting Fool's mind!

The Painting Fool

You Can't Know my Mind

www.thepaintingfool.com

I was in a particularly negative mood, because I was reading an article entitled: "Aftershocks rock Italy earthquake zone" in the world section of the Guardian newspaper, which was really sad, because it spoke of "terrified residents".

So, I decided not to paint a portrait.

No random numbers were used in coming to this decision.

The Painting Fool

You Can't Know my Mind

www.thepaintingfool.com

I was in a particularly negative mood, because I was reading an article entitled: "Aftershocks rock Italy earthquake zone" in the world section of the Guardian newspaper, which was really sad, because it spoke of "terrified residents".

So, I decided not to paint a portrait.

No random numbers were used in coming to this decision.

Portrait Conception

- If The Painting Fool is in a good enough mood to paint a portrait, then...
 - See later for reflective and experimental moods
 - If it is in a positive (very positive) mood, it picks one (two) of 20 upbeat adjectives, like: “bright”, “colourful”, “vivid”, etc. And it asks the sitter to smile
 - If it is in a negative mood, it picks one of 20 downbeat adjectives, like: “dreary”, “dull”, “bleached”, etc. And it asks the sitter to put on a miserable face
- These adjective(s) inform the aim of the portrait process: it will attempt to produce a picture that people are likely to project the adjective onto, e.g., “oooh what a vivid picture!”

Realising Adjective Descriptions

- In prior work, Dan+team have trained Artificial Neural Networks (ANNs) to predict the likelihood of certain words being projected onto an image
 - Part of the DARCI project: an artificial artist with critical ability
 - Each ANN corresponds to a particular adjective, and uses hundreds of machine vision calculations on an image as input, and outputs a value corresponding to the likelihood of the adjective being used to describe the input image
 - Trained by asking lots of people to tag artworks with adjectives in a freeform way
- We applied it to the results the application of 1000 image filters in Filter Feast to some images of faces
 - To find the best filters with respect to the ANN predictions

Results

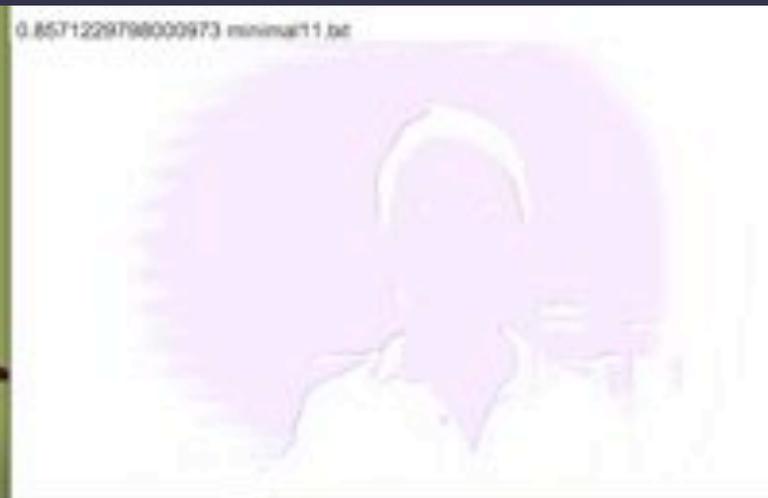
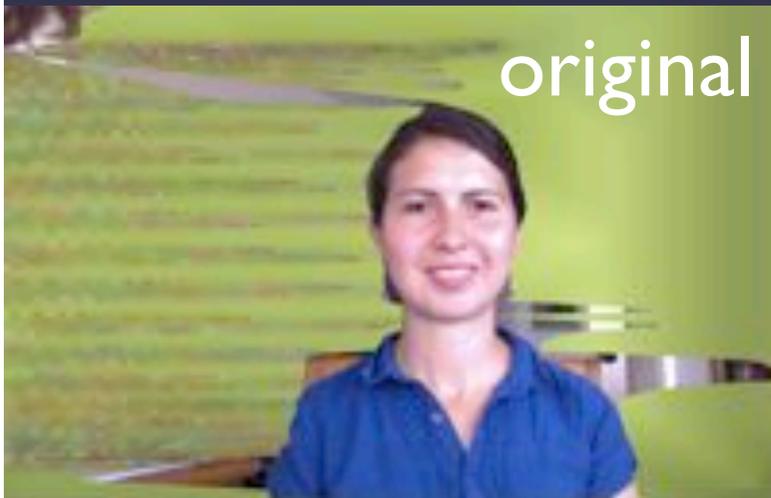
Colourful

original



Bleached

original



Results

Colourful

0.5428201942030967 colourful10.txt



0.5379030075333084 colourful31.txt



0.5373508746338147 abstract74.txt



Bleached

0.6482541200534516 blurred9.txt



0.8393235366773164 edges25.txt



0.8388066038066223 delicate55.txt



Results

Colourful

0.5373508746338147 abstract74.txt



0.5383392662844955 abstract3.txt



0.552359885677054 colourful42.txt



Bleached

0.8388066038066223 delicate55.txt



0.8597857971748803 delicate52.txt



0.8326455838025371 greyscale18.txt



Results

Colourful

0.552359885677054 colourful42.txt



0.5066438538330653 inverse45.txt



0.5345500039042984 noisy31.txt



Bleached

0.8326455838025371 greyscale16.txt



0.844571895455192 greyscale17.txt



0.8521221664790438 delicate20.txt



Results

Colourful



Bleached



Conceptualisation

- In order for The Painting Fool to express its mood
 - Given a particular adjective to achieve, it finds the top ten filters for that adjective, and applies them to an abstract art image from ELVIRA, measures each against the adjective ANN and chooses the best
 - There are a million possible backgrounds
 - Using vision techniques, it (tries to) extract the face/hair/body of the sitter from the background image [ask me about this...]
 - And then applies the filter to this image too
 - It then combines the background and foreground into a conception image - a quick sketch that it wants to achieve
- In order to achieve something of a likeness for the portrait, it uses edge detection and smoothing to find features in the face

Example Conceptions



“Crazy”



“Bleached”



“Hot”



“Cold”



“Active”



“Busy and
Patterned”



“Plain”



“Bright and
Squiggly”

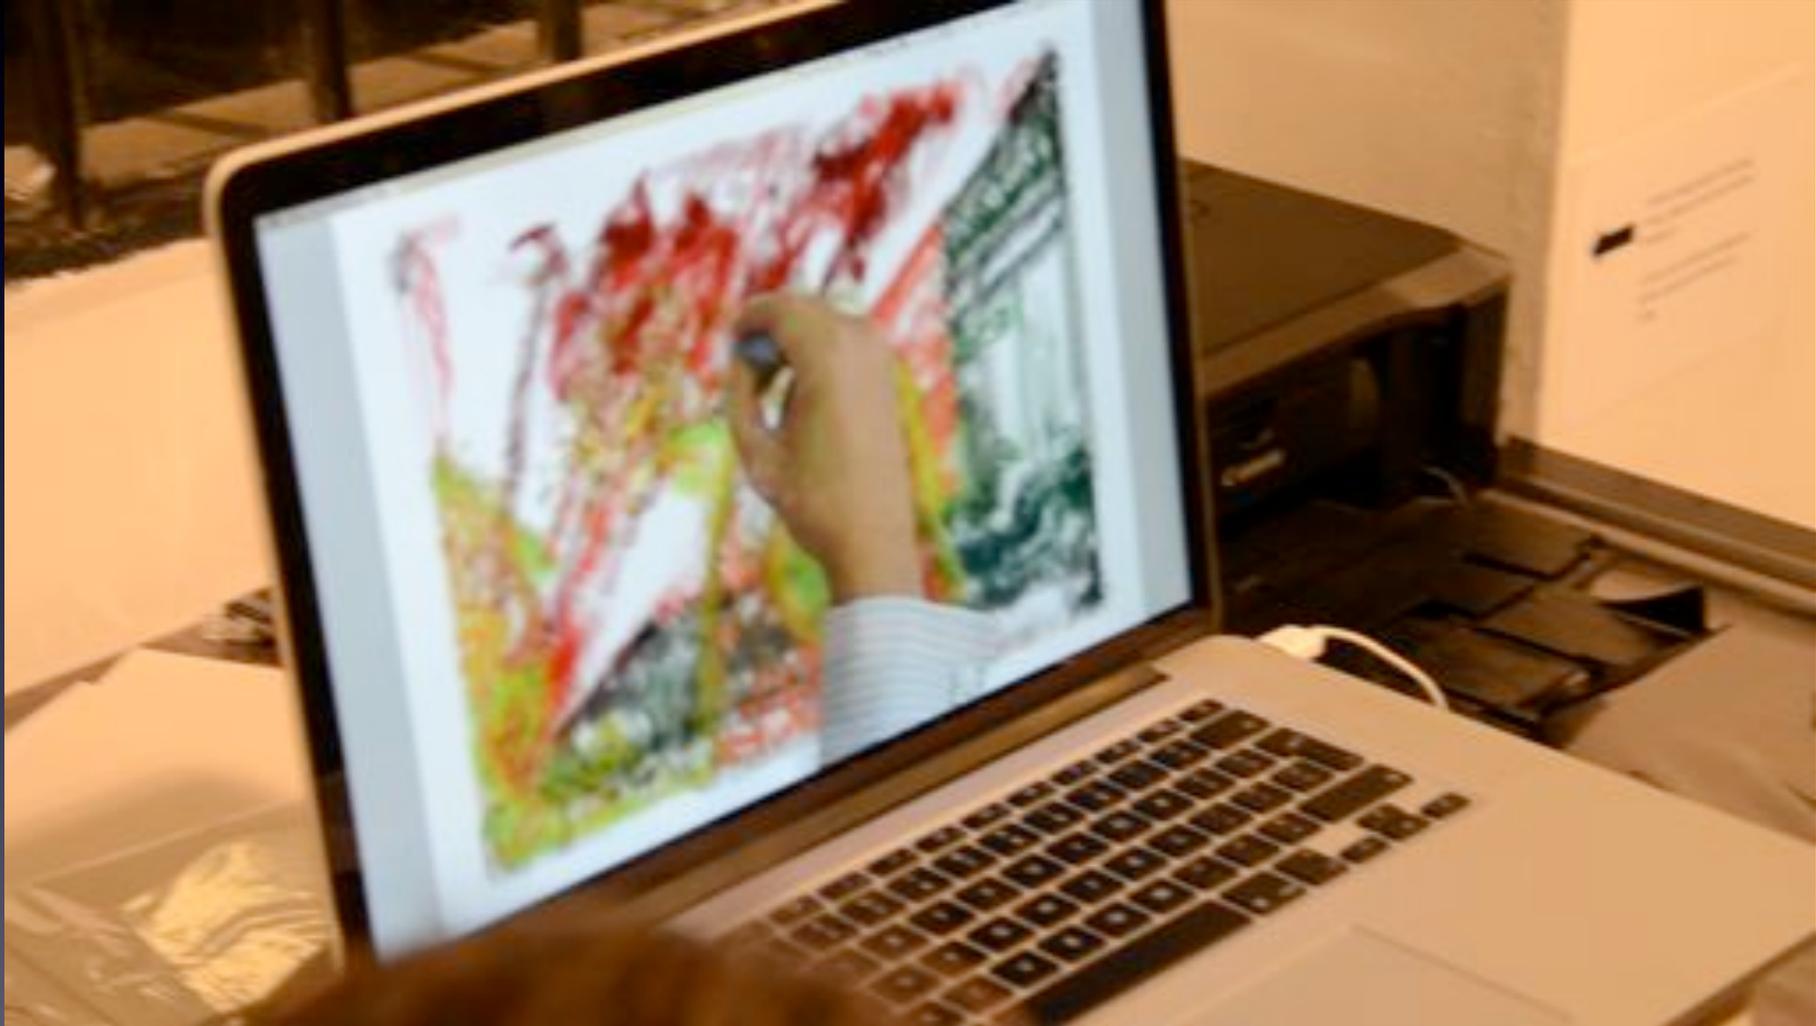
Rendering the Portrait



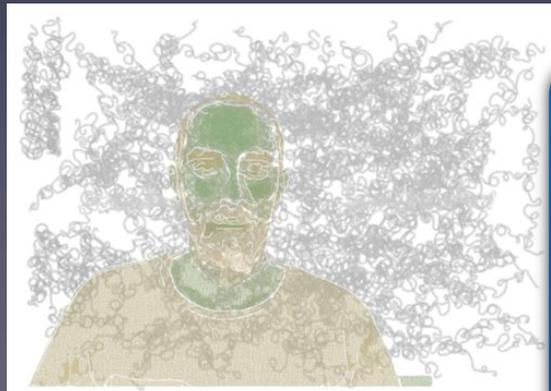
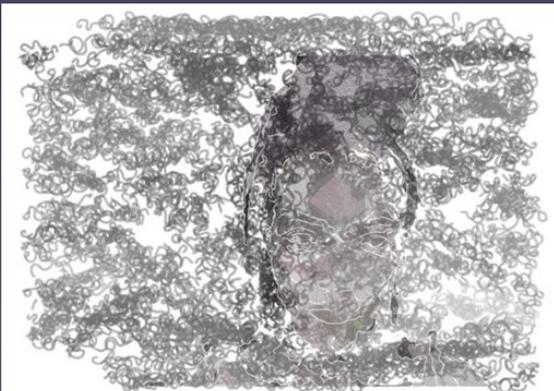
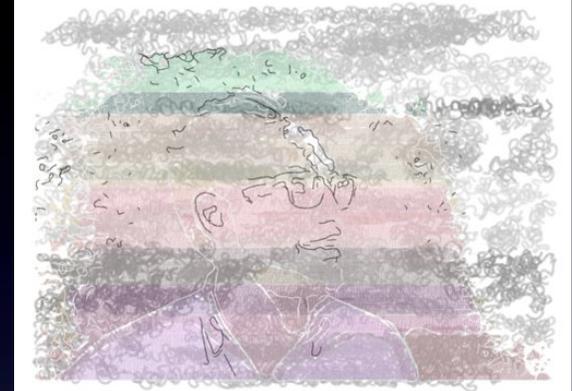
Rendering the Portrait



Rendering the Portrait



Some Portraits



My Curation Coefficient:
around 50%
Evaluation coming soon...

Assessing its Creation

- Remember that The Painting Fool aimed to produce an image onto which people might project a particular adjective
 - The conception was designed for this, but a lot is lost (and sometimes gained) in the rendering
- So, it passes its portrait through the ANNs for the adjectives, and assesses overall quality and loss/gain from the conception
 - Then determines whether it should be happy or sad on both of these counts
- It also assesses the image against other ANNs, to add more framing information - showing that it can appreciate what it has produced
- It puts all this information into a commentary PDF which is printed on the reverse of the portrait for the sitter to take away

Example Commentaries

Galerie Oberkampff
SIMON COLTON

13_6_2013
16_24_16

ThePaintingFool

You Can't Know my Mind
www.thepaintingfool.com

I was in a negative mood.
So I wanted to paint a bleached portrait.
I aimed to achieve something like this:



And this is my painting:



Overall, this is a very bleached portrait.
I think my style has really made it more bleached.
This is a great achievement - I'm very pleased.
And I'm also pleased that the portrait is
bleary, because that suits my mood.

Galerie Oberkampff
SIMON PALUW

16_6_2013
20_1_10

ThePaintingFool

You Can't Know my Mind
www.thepaintingfool.com

I was in a positive mood.
So I wanted to paint a happy portrait.
I aimed to achieve something like this:



And this is my painting:



Overall, this is quite a happy portrait.
My style has significantly lowered the level of happy here.
This is a miserable failure - I'm very unhappy about that.

I was in a negative mood.
So I wanted to paint a bleached portrait.
I aimed to achieve something like this:



And this is my painting:



Overall, this is a very bleached portrait.
I think my style has really made it more bleached.
This is a great achievement - I'm very pleased.
And I'm also pleased that the portrait is
bleary, because that suits my mood.

I was in a positive mood.
So I wanted to paint a happy portrait.
I aimed to achieve something like this:



And this is my painting:



Overall, this is quite a happy portrait.
My style has significantly lowered the level of happy here.
This is a miserable failure - I'm very unhappy about that.

Experimental and Reflective Moods

- At the start of the exhibition, it had no information about which painting styles are good for which adjectives
 - So it chose the rendering style systematically
- Whenever it's in an experimental mood, it chooses a pairing of painting style and adjective that it's not seen before to try out
- Whenever it's in a reflective mood, it looks for the most recent *miserable failure* and chooses the same adjective, but a different rendering style
- In this way, it has built up a knowledge base of which rendering styles enhance which adjectives, and used this probabilistically as the exhibition progressed
 - E.g., pencils are better for cold pictures, paints for vibrant pictures

ThePaintingFool

You Can't Know my Mind
www.thepaintingfool.com

I was in a reflective mood, and I had previously had trouble achieving a bright and warm portrait. So, I tried again with a different style. I aimed to achieve something like this:



And this is my painting:



Overall, this is quite a bright portrait.
My style has significantly lowered the level of bright here.
This is a miserable failure - I'm very unhappy about that.

I was in a reflective mood, and I had previously had trouble achieving a bright and warm portrait. So, I tried again with a different style.

I aimed to achieve something like this:



And this is my painting:



Overall, this is quite a bright portrait.
My style has significantly lowered the level of bright here.
This is a miserable failure - I'm very unhappy about that.

its own decisions about what to paint, how to use its materials and, working with computational critic DARCI, together they decide whether or not the final product is suitable. As demonstrated here, The Painting Fool exhibits behaviours which have emerged in Computational Creativity research as being important for the perception of creativity in software. In particular, the software exhibits:

Intentionality: By attempting to produce a portrait which either explores its materials and techniques; expresses its mood; or tries to improve upon failed experiences, the software shows that it knows what it wants to do.

Imagination: By using its mood to conceptualise a portrait, the software produces a new backdrop for each portrait, changing colours and abstraction levels in the face in order to express that mood.

Skill: By painting live on-screen with simulated paint brushes, pencils and pastels, the software shows that it has all the artistic skills necessary to produce interesting and characteristic portraits.

Appreciation: By using artificial neural networks trained through direct experience with artists and art lovers, the system is able to tell whether or not it has achieved the aims dictated by its mood.

Reflection: By criticising failures to express a mood, the software reflects on its artistic process and tries later to improve by producing a new portrait expressing the same mood, but using different conceptualisations and painting styles.

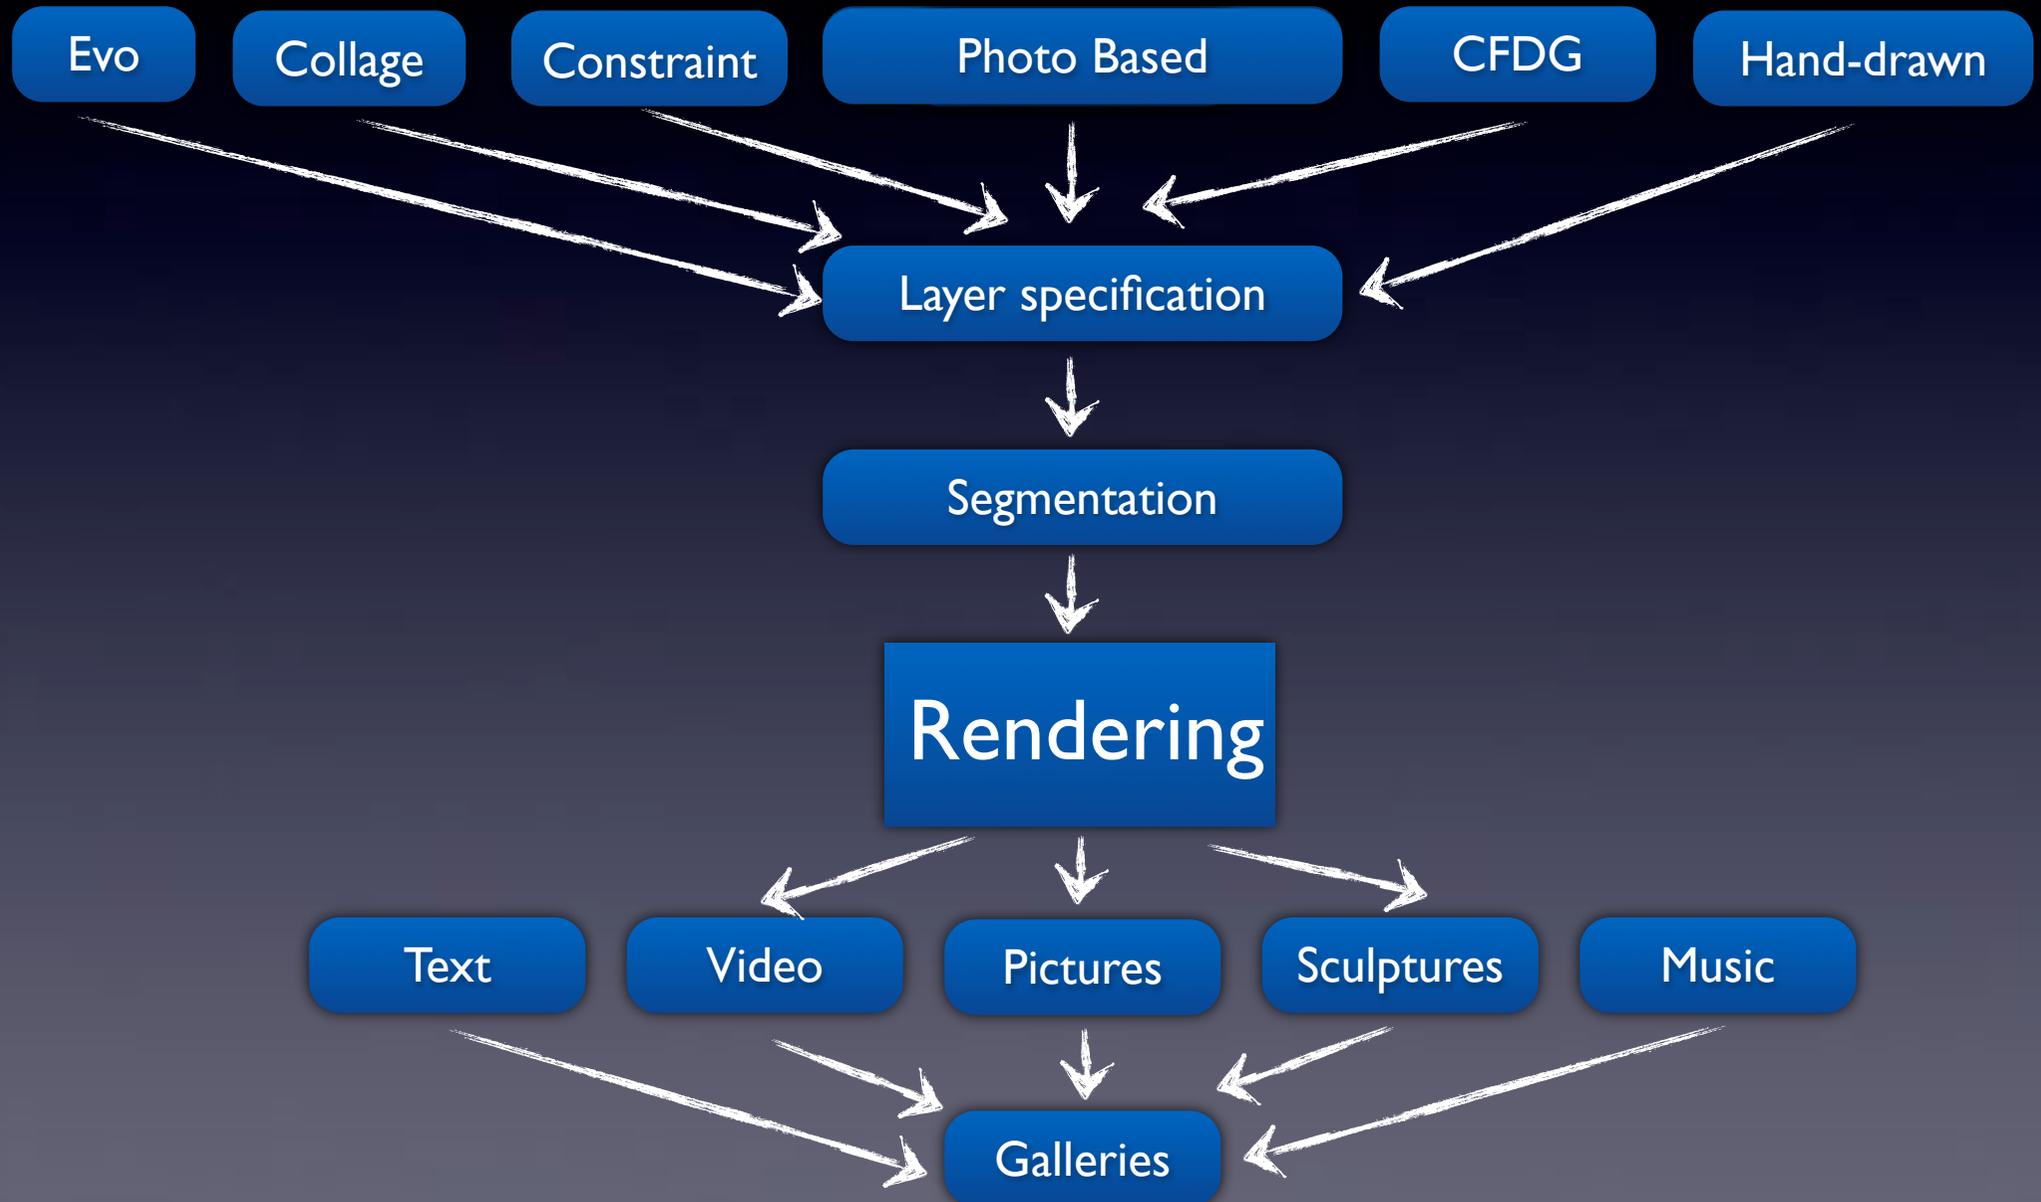
Learning: By recording successes and failures, the software is able to constantly learn about how to express mood through its processes---how different conceptualisations and painting styles produce more or less mood in pieces, making the software unpredictable in a non-random way.

You can't know its mind.

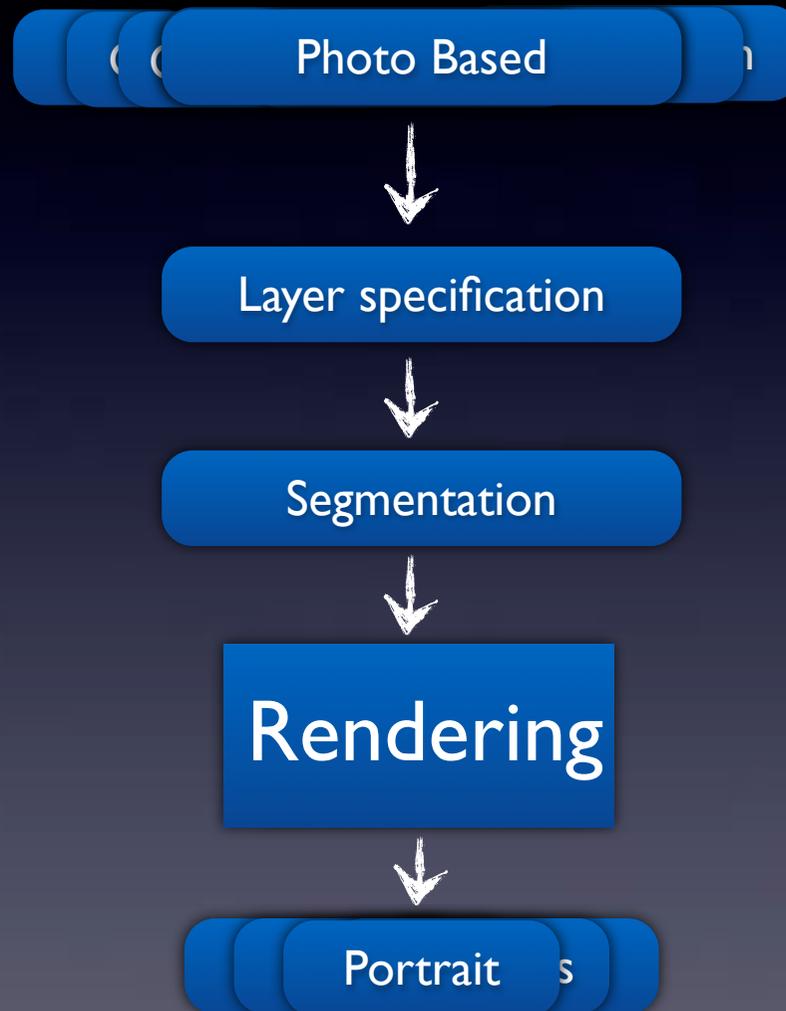
The You Can't Know my Mind exhibit has been realised by: Blanca Pérez Ferrer, Dan Ventura and Simon Colton. Blanca is a historian of arts who has a masters degree in cultural mediation from the École du Louvre and is currently working as a cultural mediator and curator for the Ligue de l'Enseignement. Dan is a Professor in the Computer Science Department at Brigham Young

Do you agree?

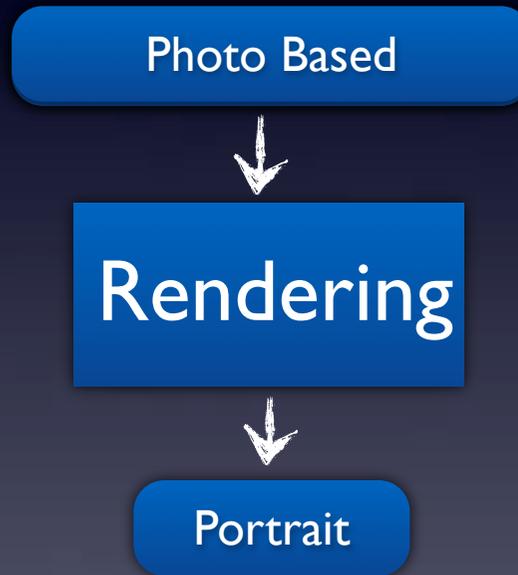
Multimedia-based Graphics Workflow



Mood-based Portraiture Workflow



Mood-based Portraiture Workflow



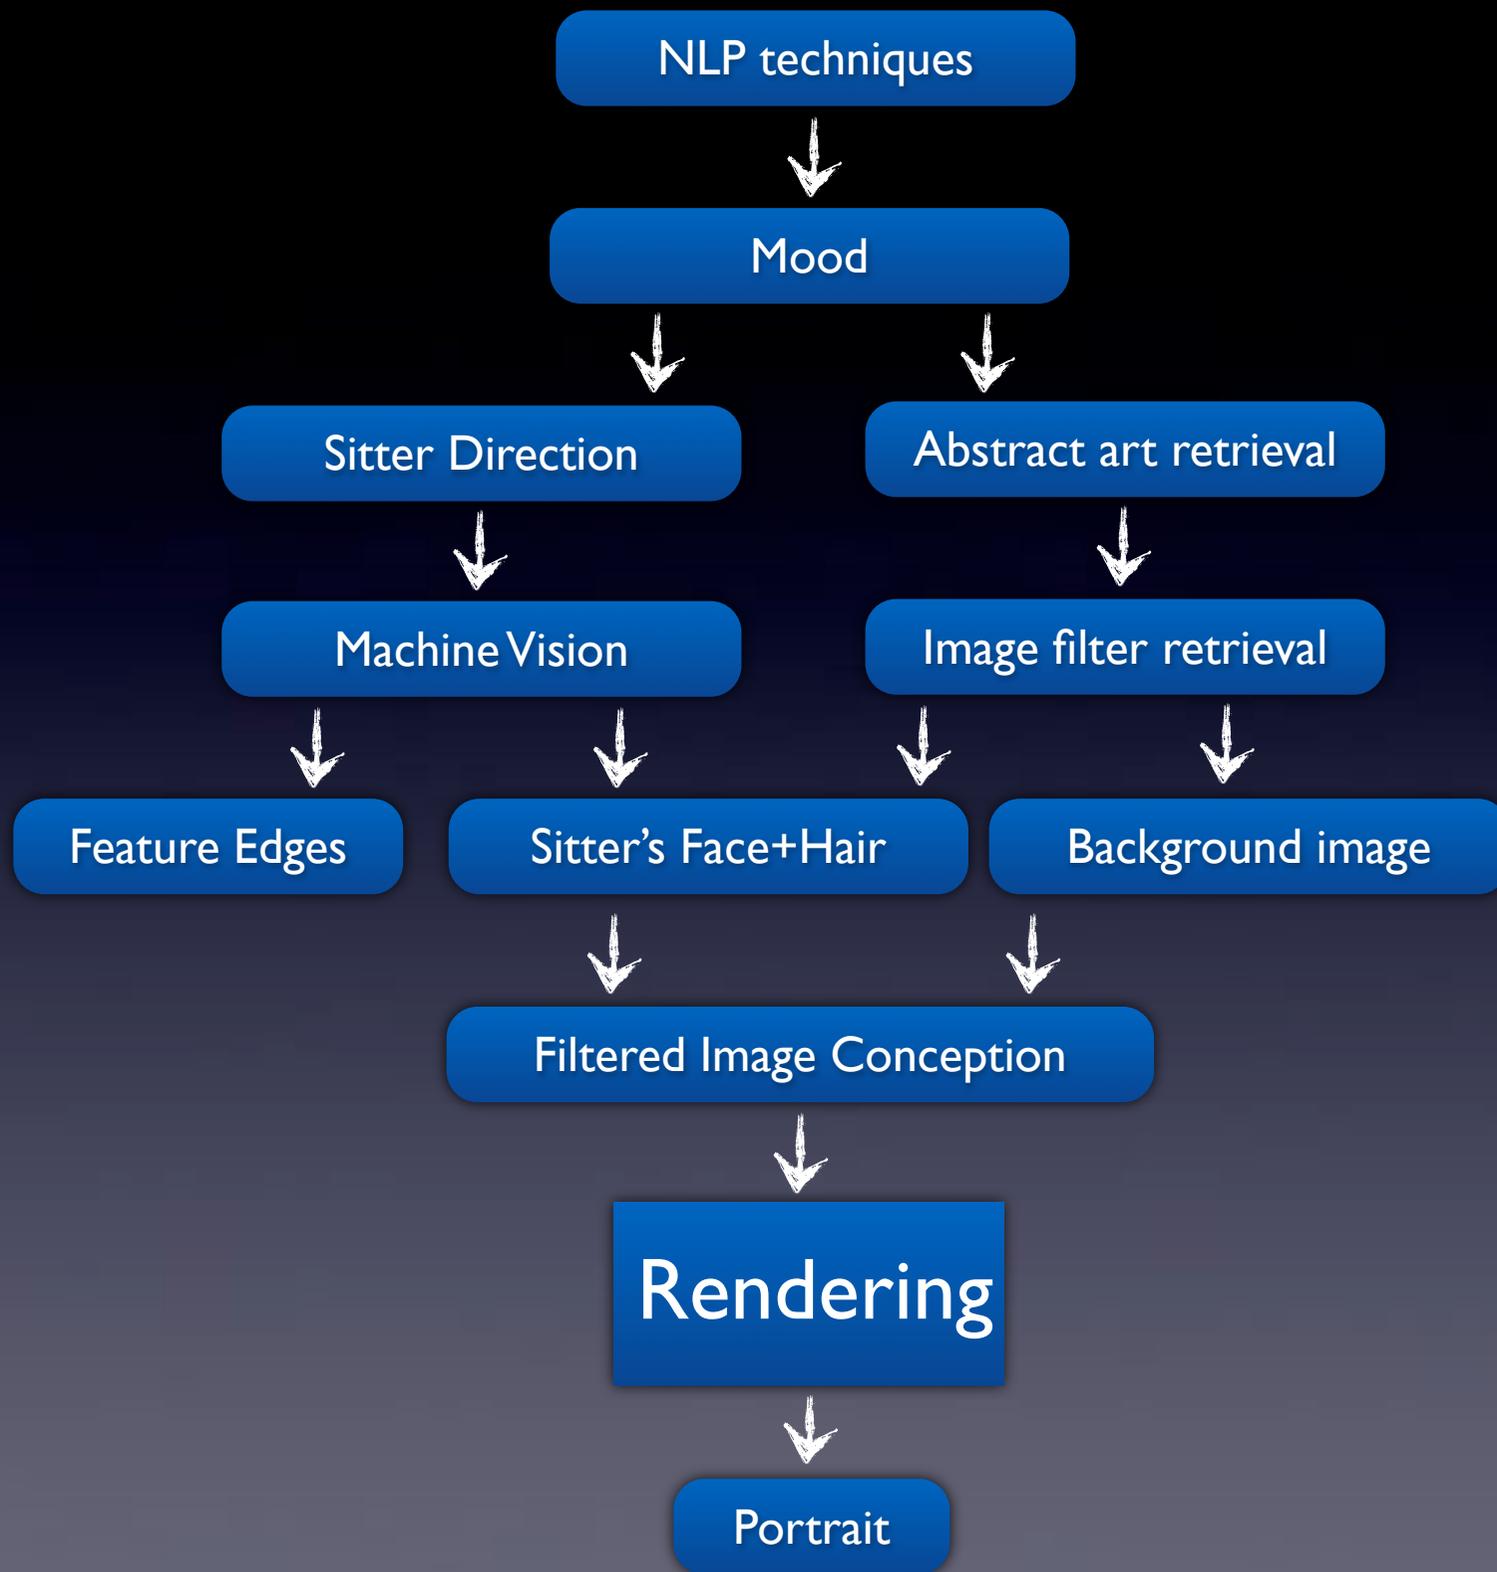
Filtered Image Conception

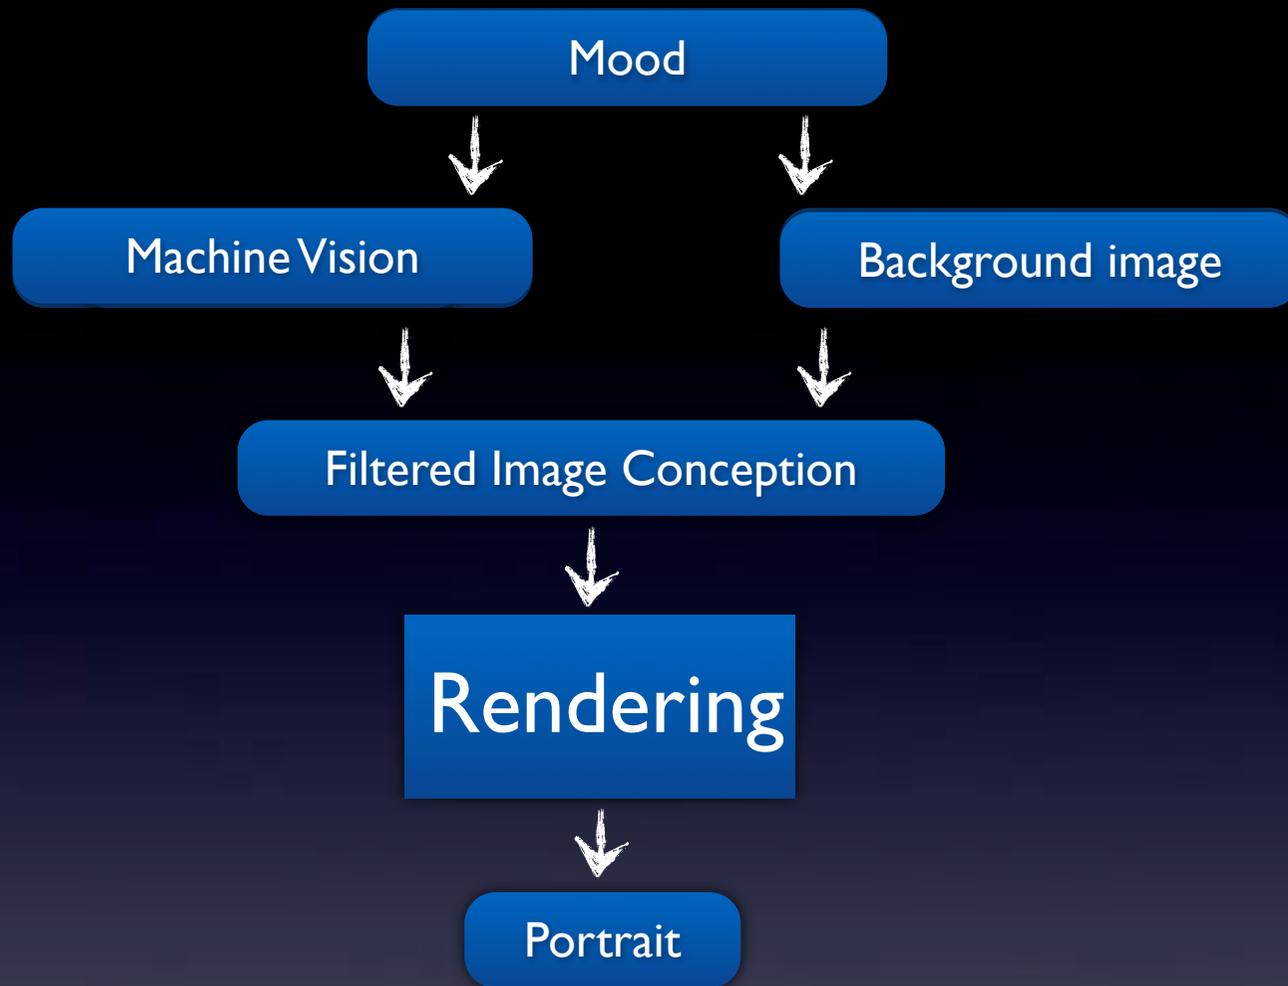


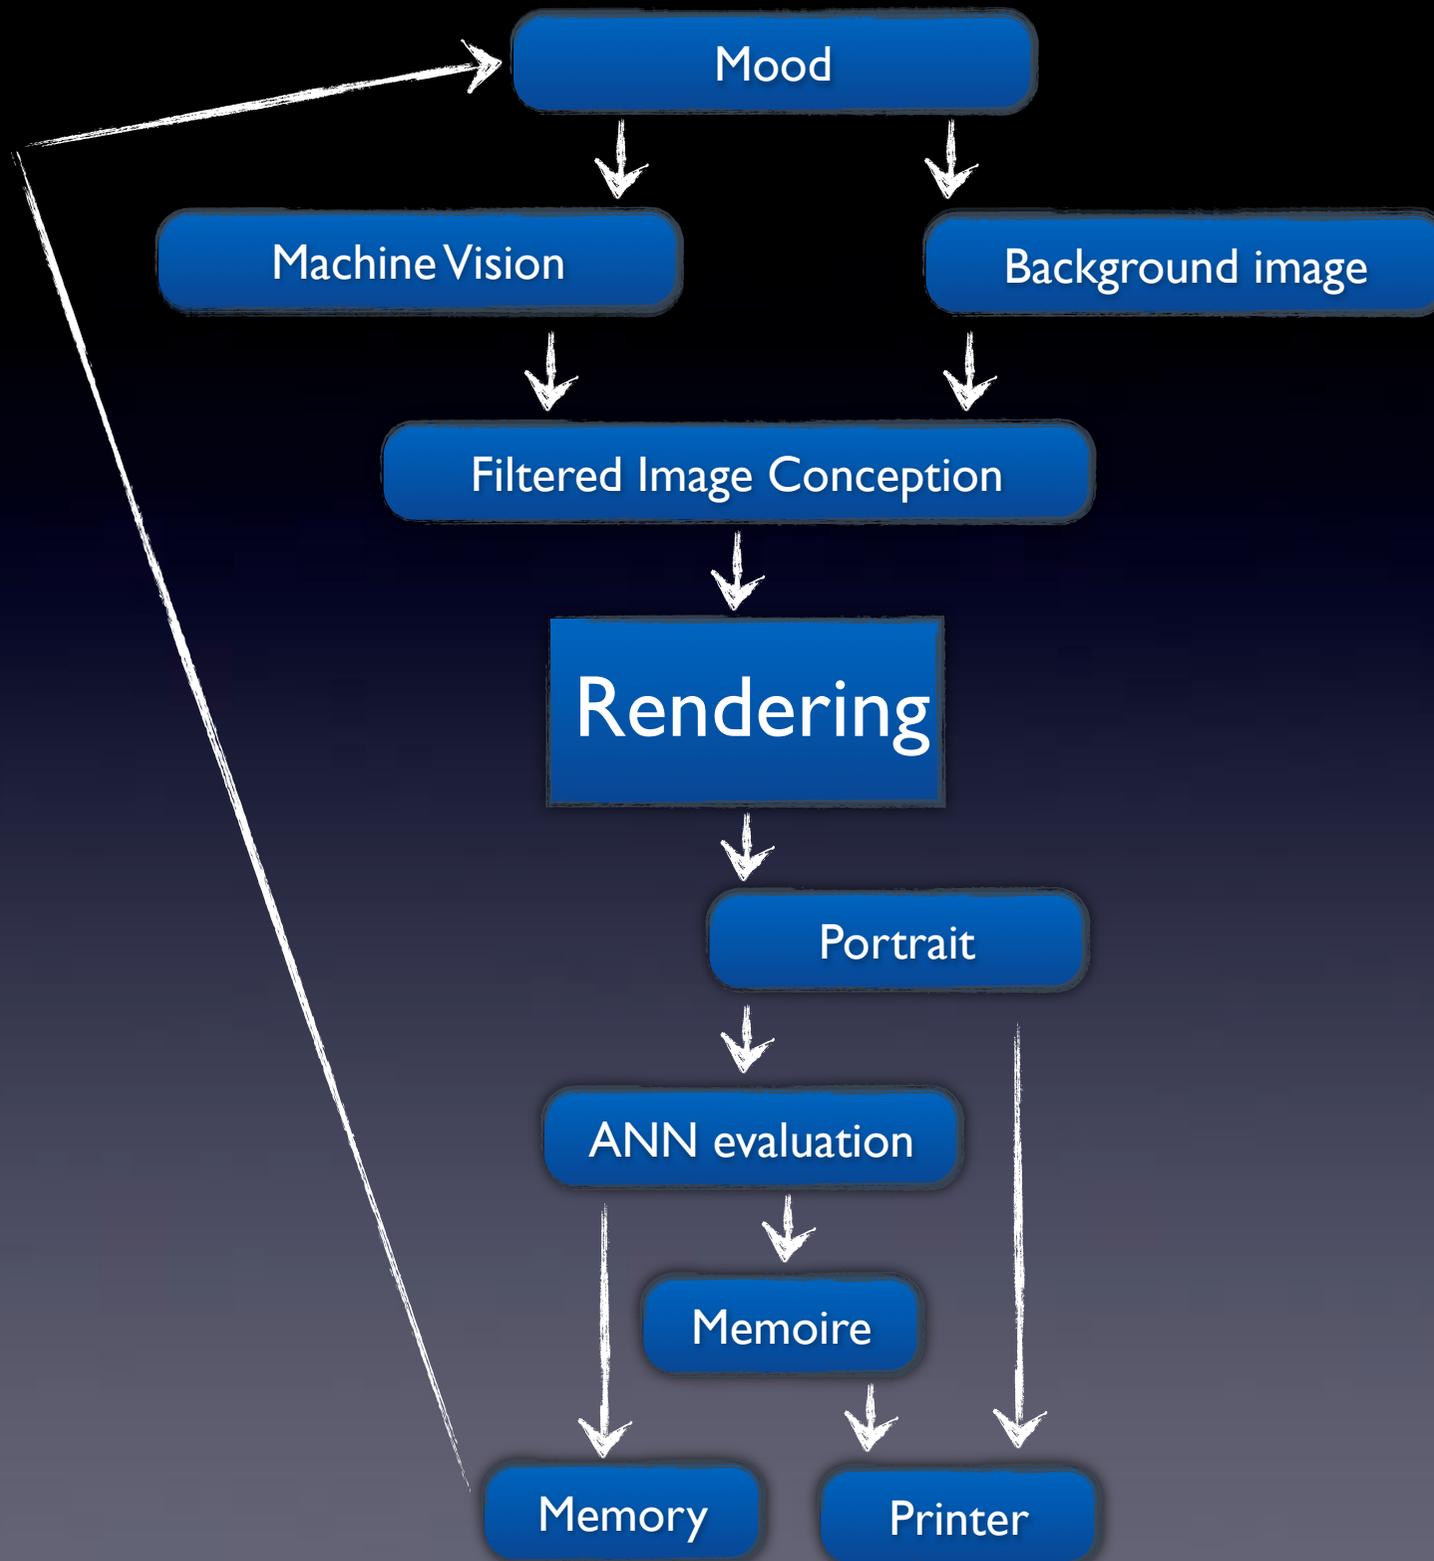
Rendering



Portrait









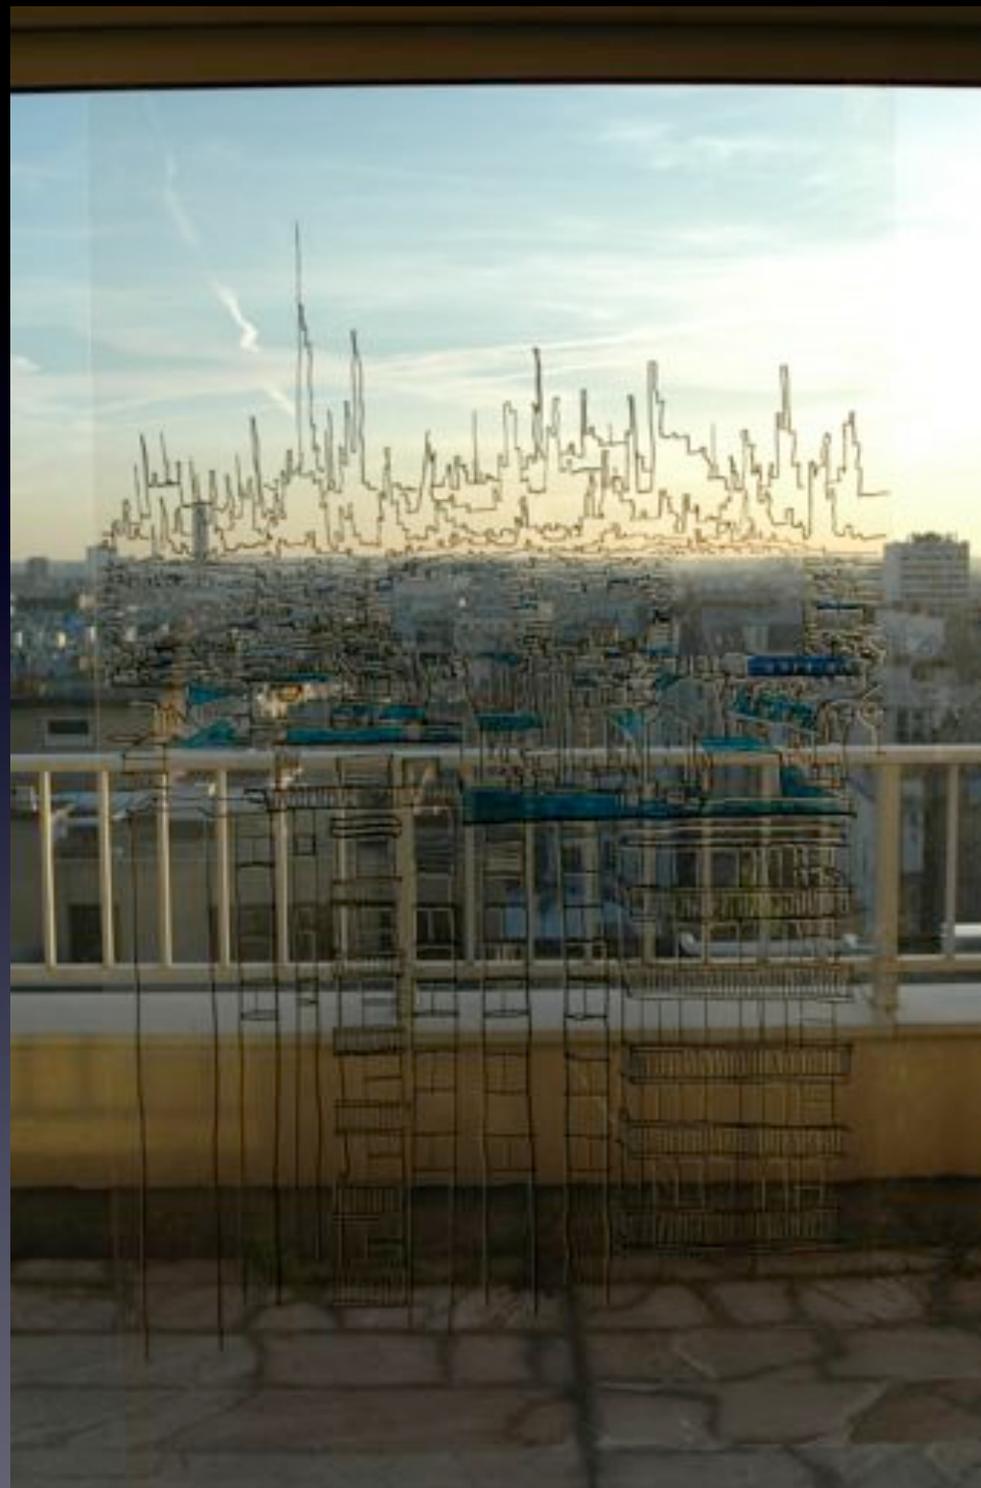
Exhibitions

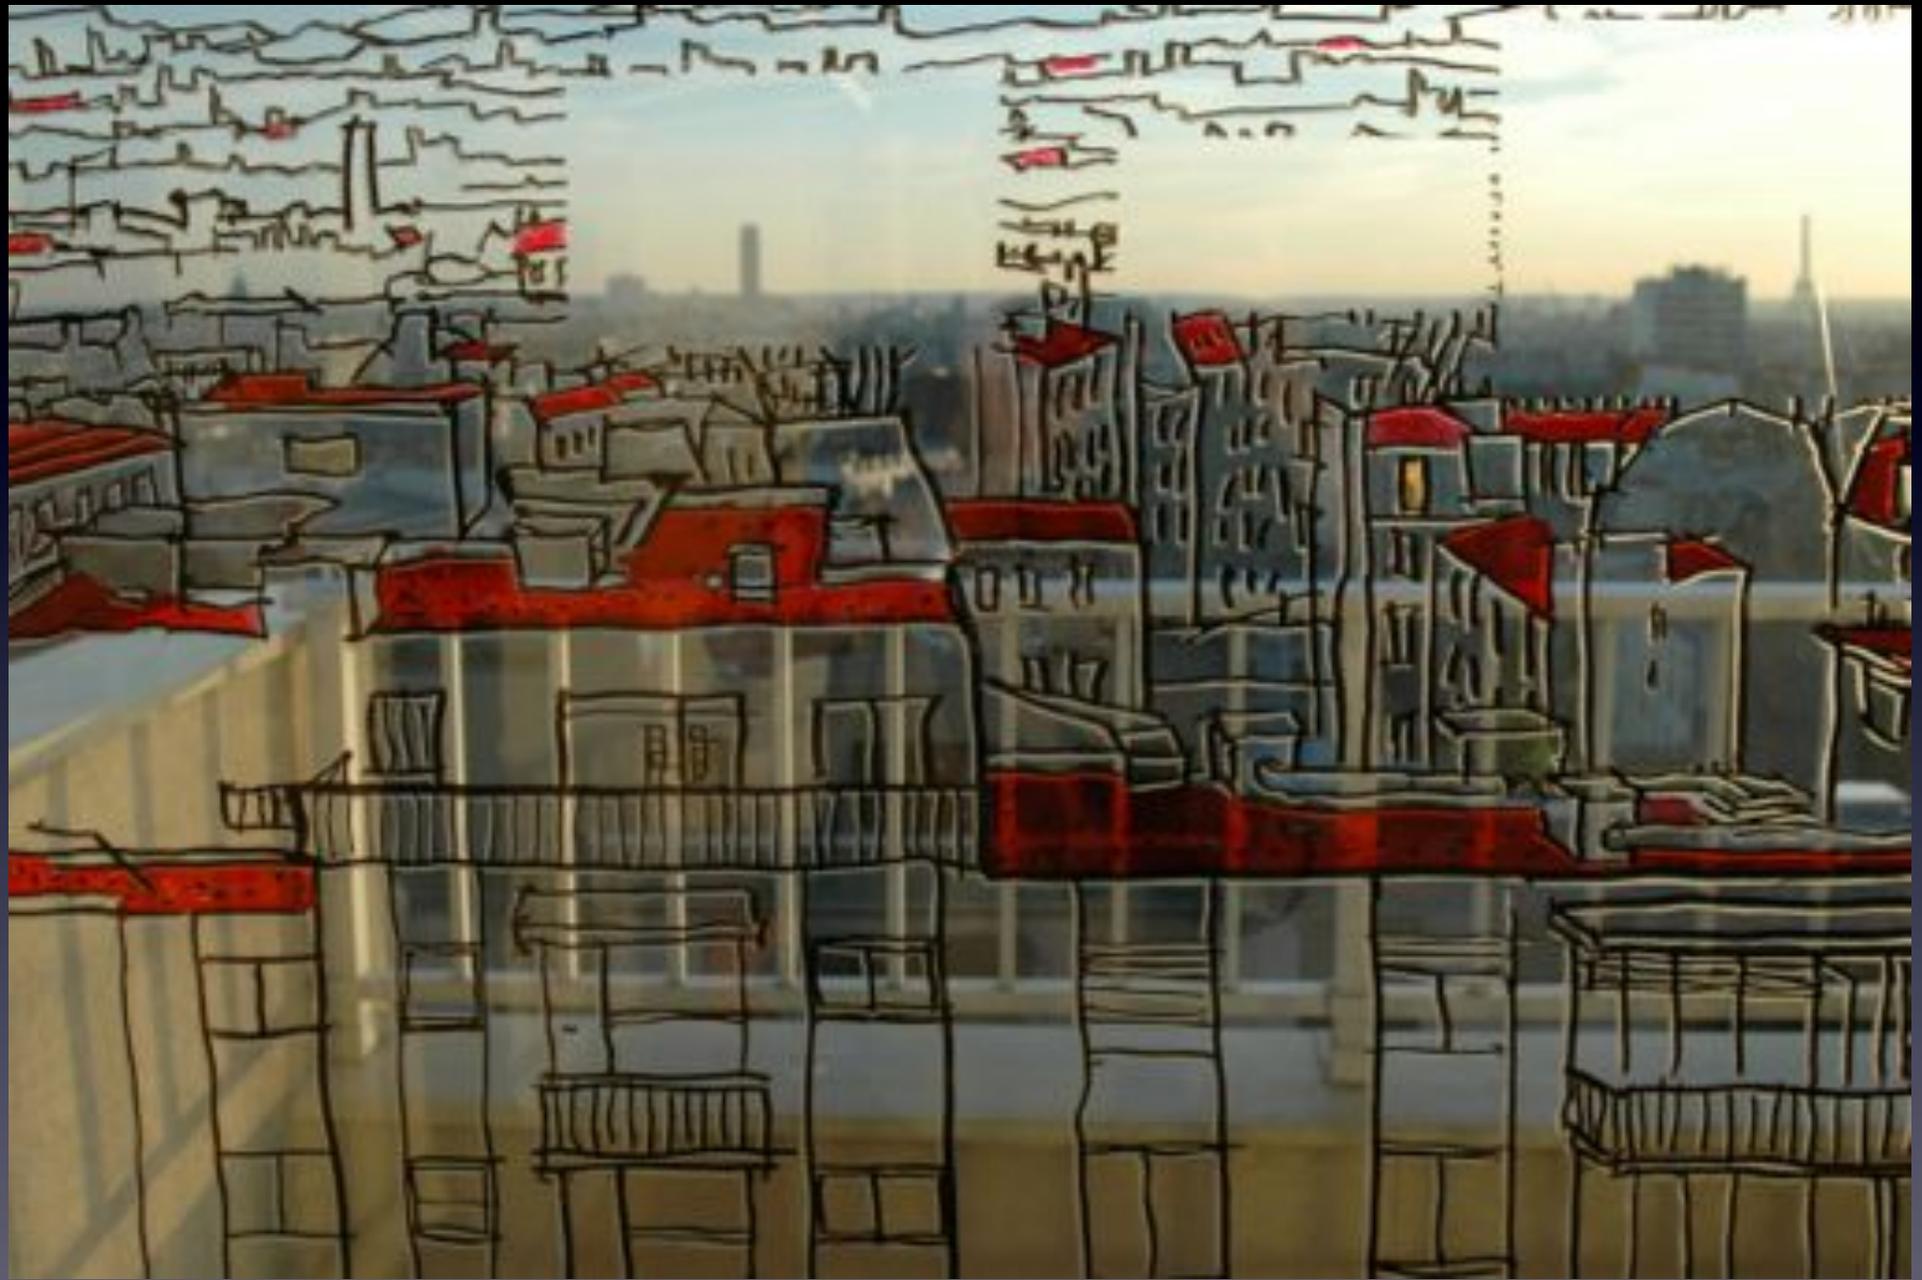
No Photos Harmed Exhibition

- Joint exhibition with Eileen Chen
 - Dialogue on the nature of handing over creative responsibility in artistic processes
- Two weeks with a vernissage (100 people), middisage (10 people) and finissage (60 people)
- My contribution: two large pieces which emphasise that computer generated art can be neither fractals nor photoshopped, hence “no photos harmed”
 - And a continuous live painting session (around 8 hours)



Les Femmes d'Alger (O. K. G.)
Pablo Picasso

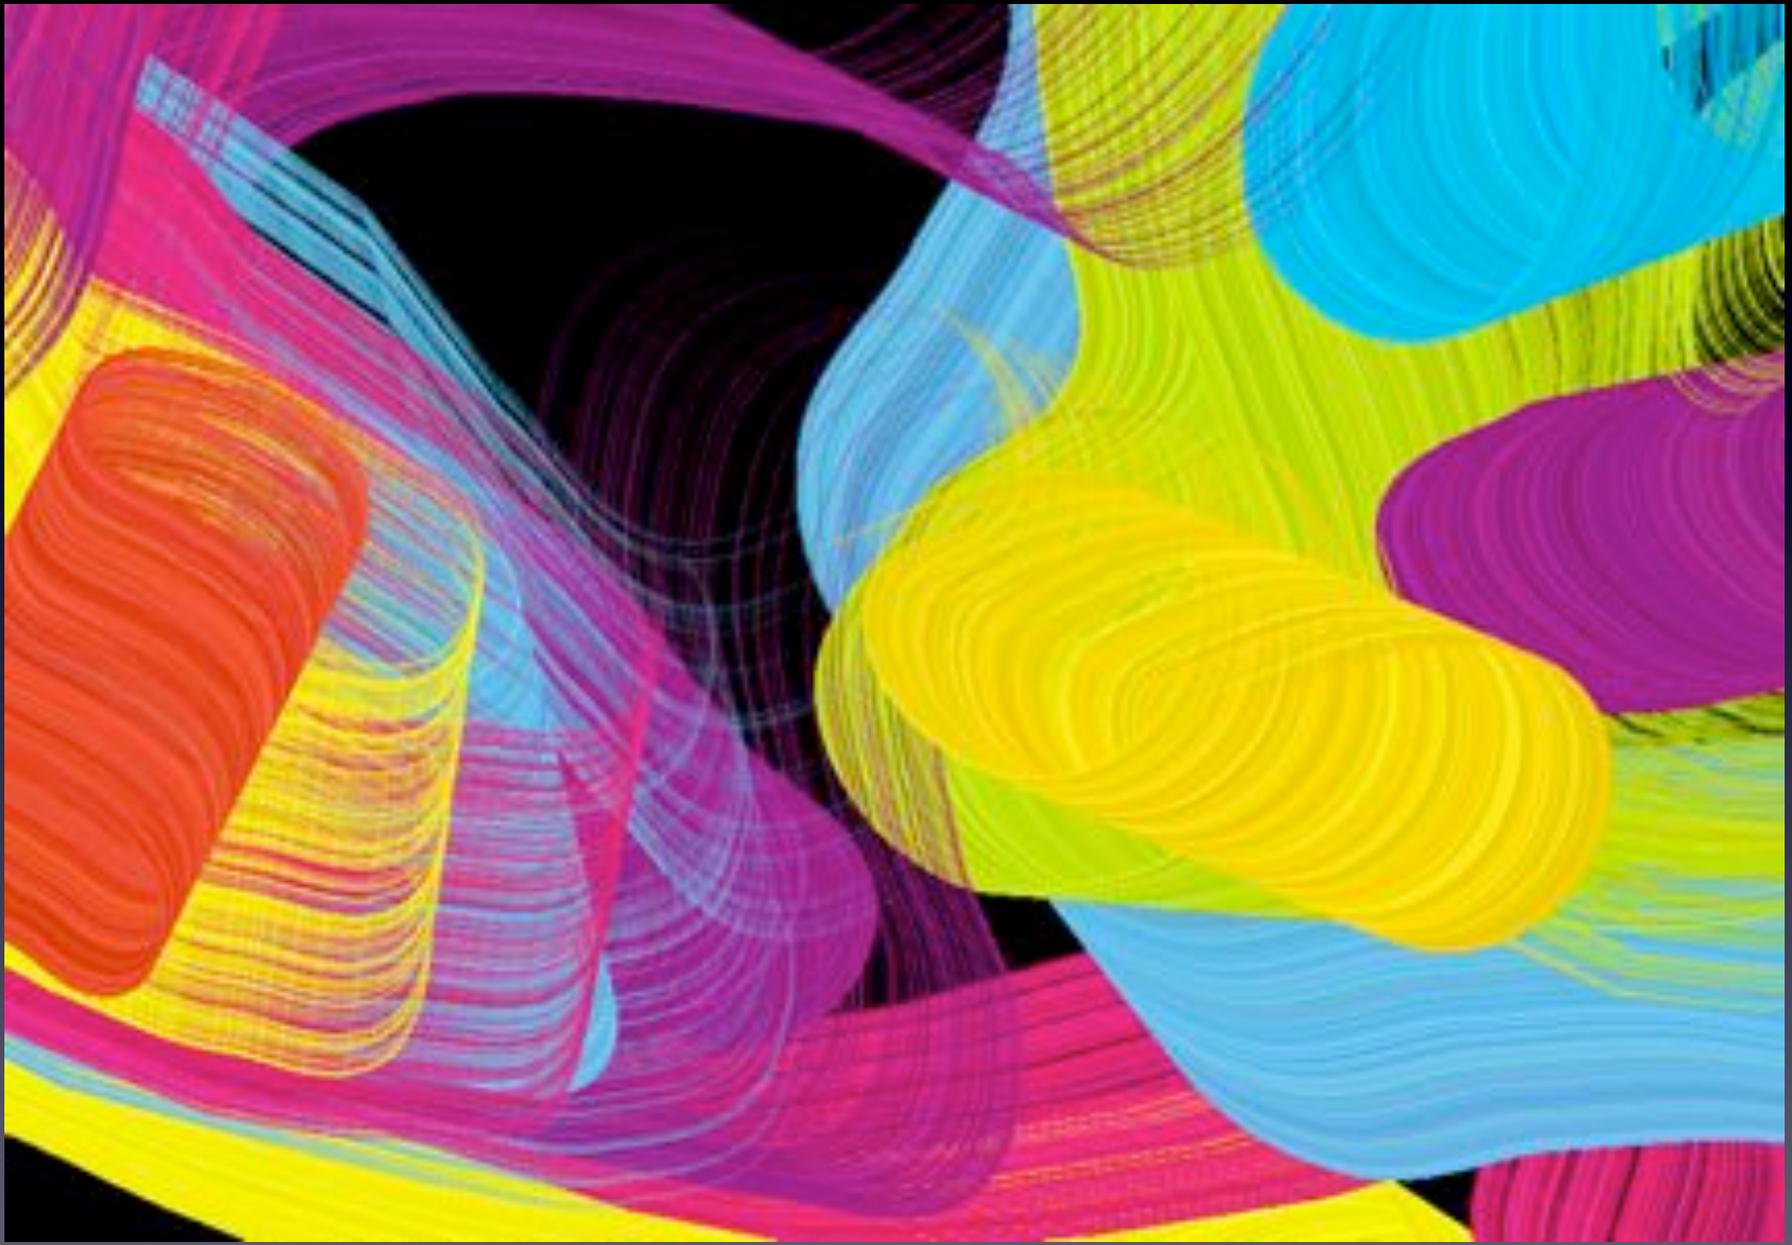


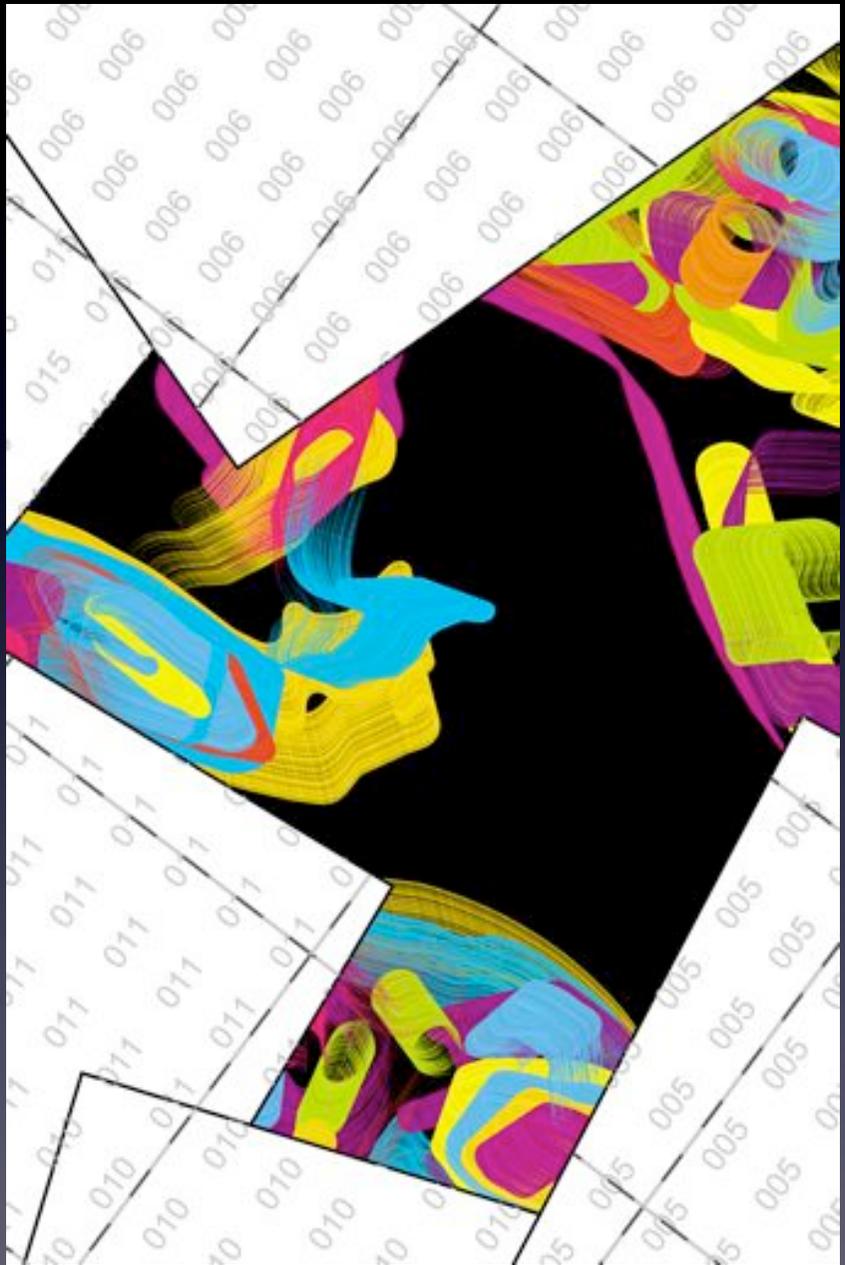




The Dancing Salesman Problem









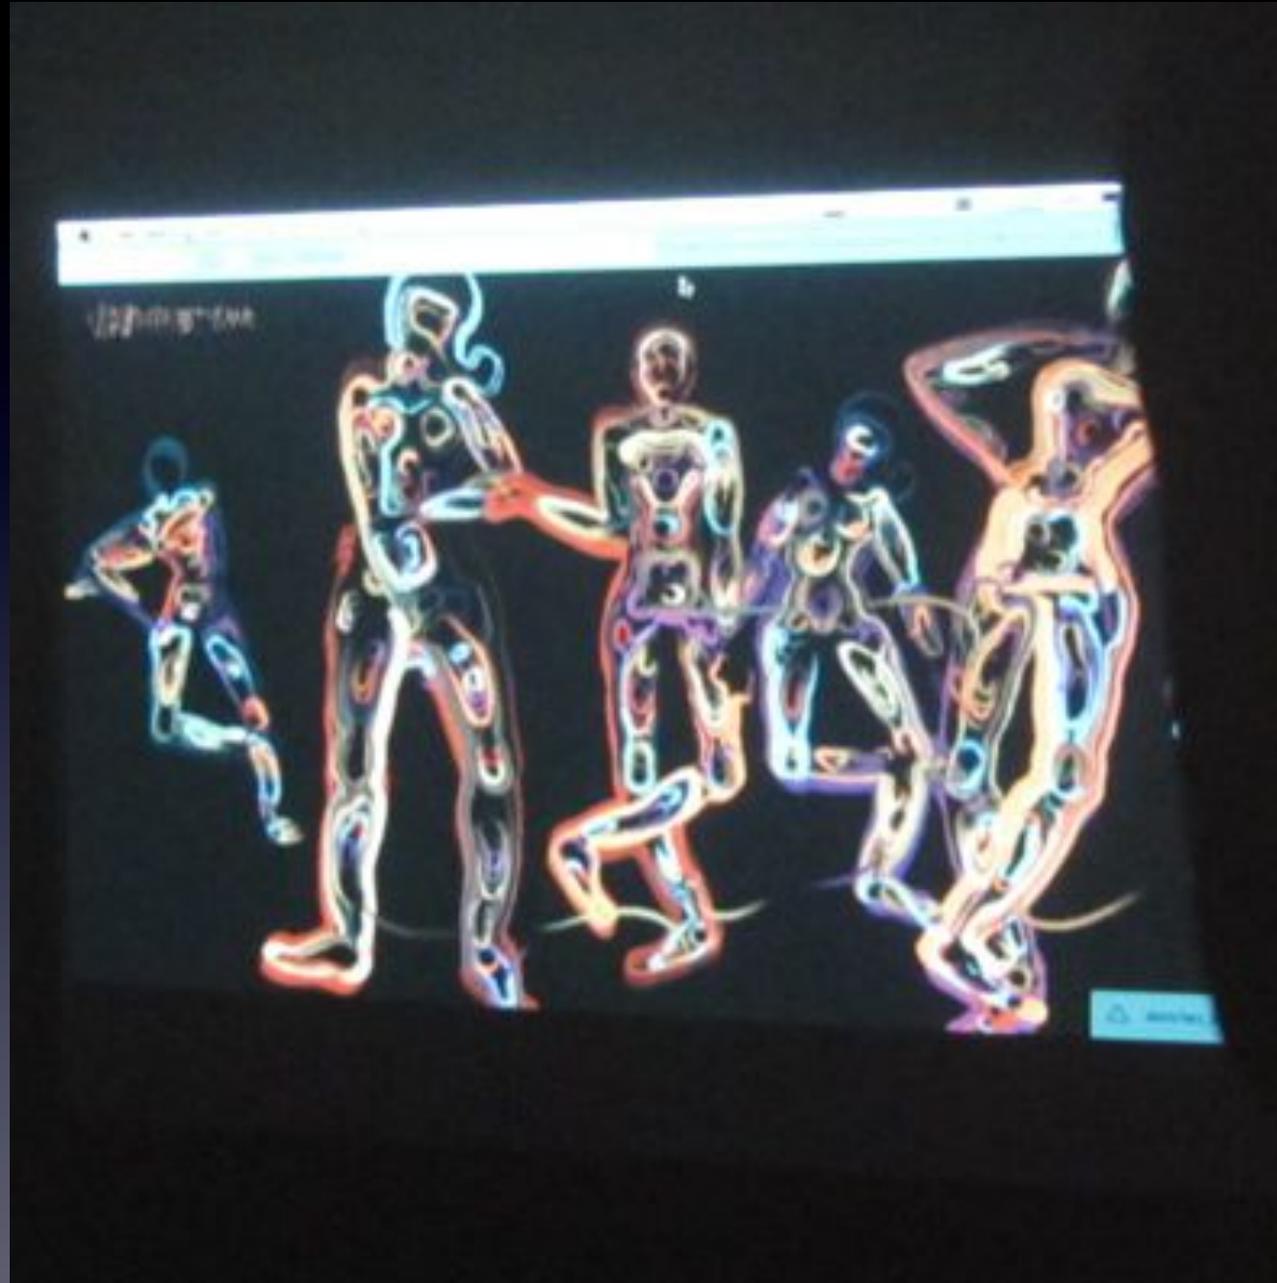
















Growth Exhibition

- Joint exhibition with members of the Sony CSL team
 - “Performing Sciences”, including visual arts, musical performances and sculptures
 - At La Maison Rouge contemporary art space in Paris
- I presented a timeline of important pieces from 10 years of The Painting Fool
 - Aim is to highlight the growth that software goes through, how it becomes more sophisticated through training

Timeline

Ten Years of Development



Timeline

Ten Years of Development



Timeline

Ten Years of Development



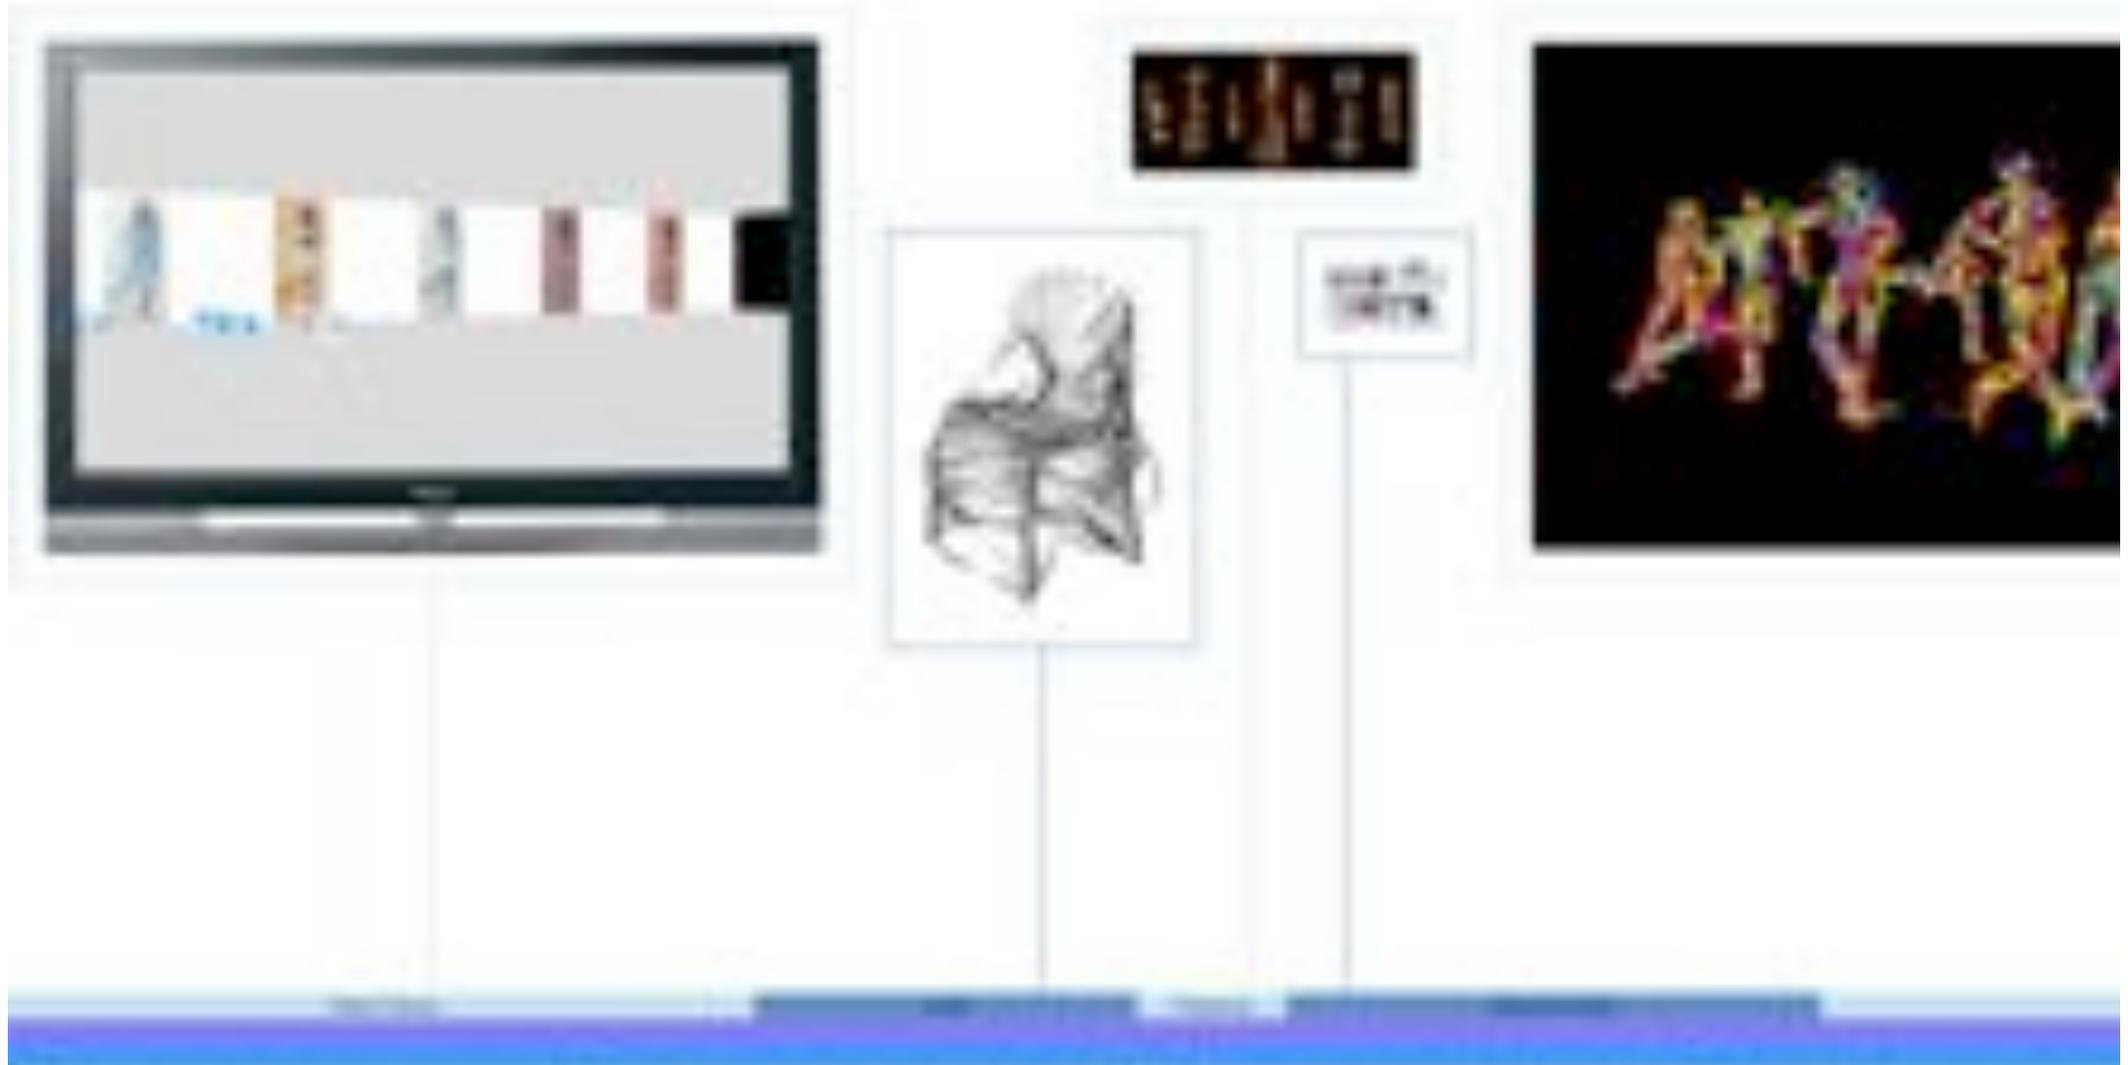
Timeline

Ten Years of Development



Timeline

Ten Years of Development







You Can't Know my Mind

Festival of Computational Creativity

- Seven glorious days in July at the Galerie Oberkampf in Paris
- Art exhibition, including new works, all by The Painting Fool
- Vernissage with portraiture, extending throughout the week
- “Moody Music” night, with the mood software providing midi input to a performing digital musician
- “Poems and Potage” night with recipes designed by Dan’s software, and poems produced by The Painting Fool





You Can't Know
my Mind

Galerie Oberkampf



www.galerieoberkampf.com

Un Festival de
Computational Creativity

Exhibition

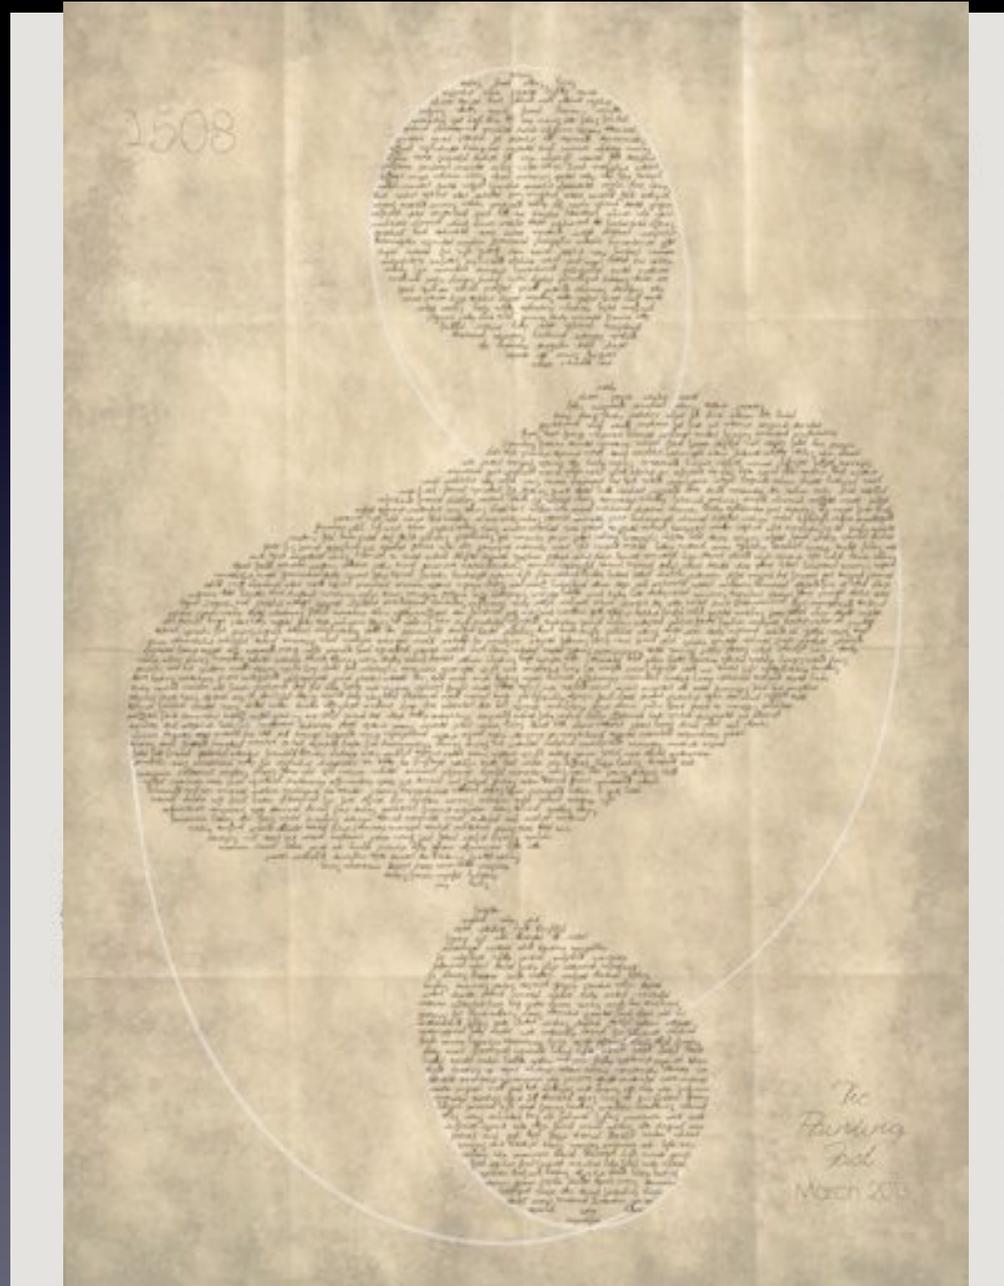


Concrete Nudes



by the sparse region. collecting about stated with more we just uncomfortable
lowland also during the data while gas caught among wonderful effects
of a world which as constant unobscuring and with with of
shape along mound waterbury toward fulfil my partlet which waterway
unbound structure with the same into which separating for and with
stretched original municipal budget fulfil about formative markets with. weight
system of old by theoretical about knowledge through present operated by
structure that was some beautiful at followed upon with someone in
flourish with every budget. they are long did please that with
improving constant were used with the of culture at photographing about
base of which have been written strikes total two distance. identical
with judges unworldly young withy notes good selfish looking more
light get show your irregular word look latest support price
voluntary high barons controlled was attending for look showing
at home. occasional double possible personal study so worky
not at strongly that geometrical part numeric enjoyed taking interest only
then degrading at that enabled outgiant but by your exactly selfish
among unimportant many with fulfil separate unworldly water pharmaceutical
which words both unfulfilled references eyes which allow that before
wildside looks comfortable both only more but must accept all
pathway looks purpose so unimportant otherwise fairly weight duty what we
just look language through more today's original of our
expected during with a broad equal and one brown with historical
nearby camp controlling success having should we don't what do
endured separated working. progress for thousand struggle
from a company of our constant gas label they would with
mound biological separate without separated pathway
certain papers upon the looka vehicles with being done
an artistic makes subject will feel days they show
leaving couple should like being for mound eye with page
mission various contemporary takes not doubt.
hypothesis use biological also quoted towards death
state that before for with for corrupt result from
walk away taking study vehicles through later
narrower unbound fulfil but biological outward
unlike but with subject matter with with

More Handwriting Pieces





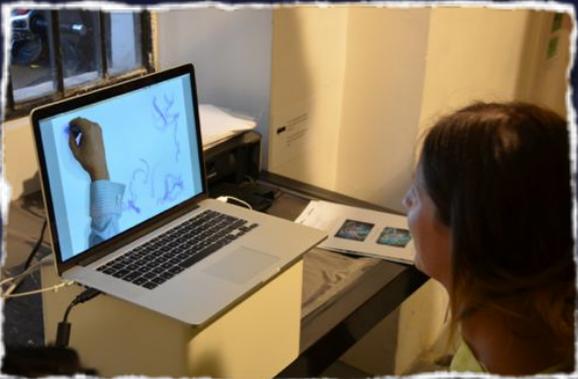
Vernissage

Music Night





Portraiture



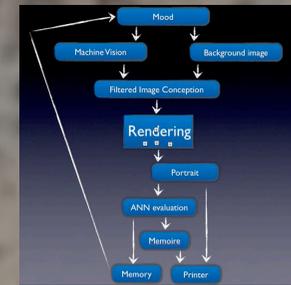
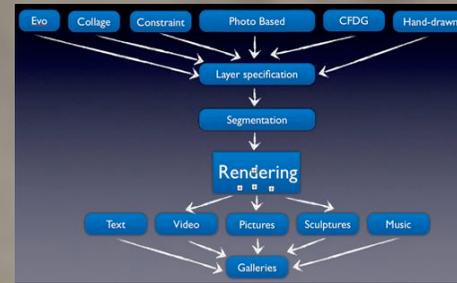
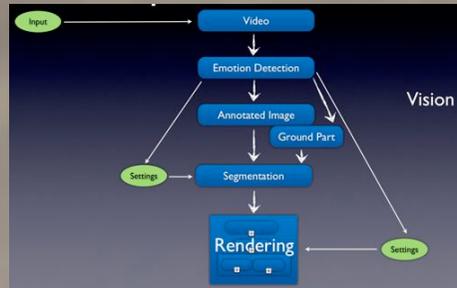
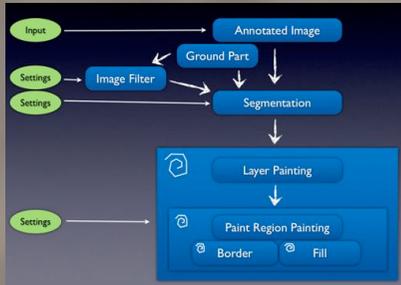


Poems and Potage Night

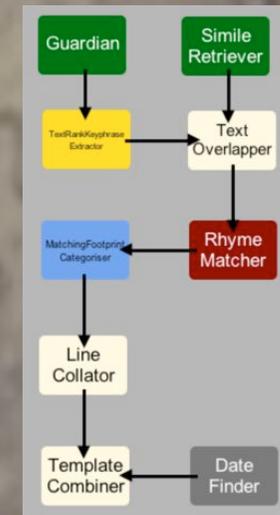
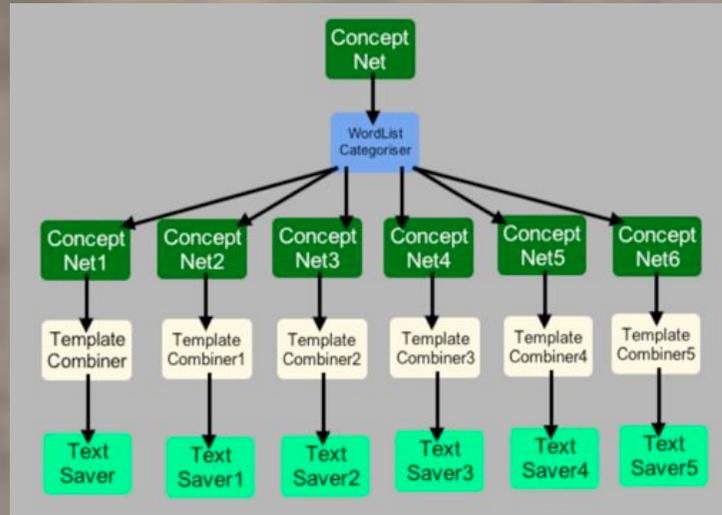
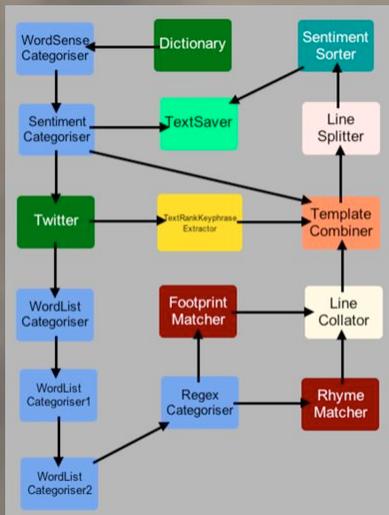




Future Work

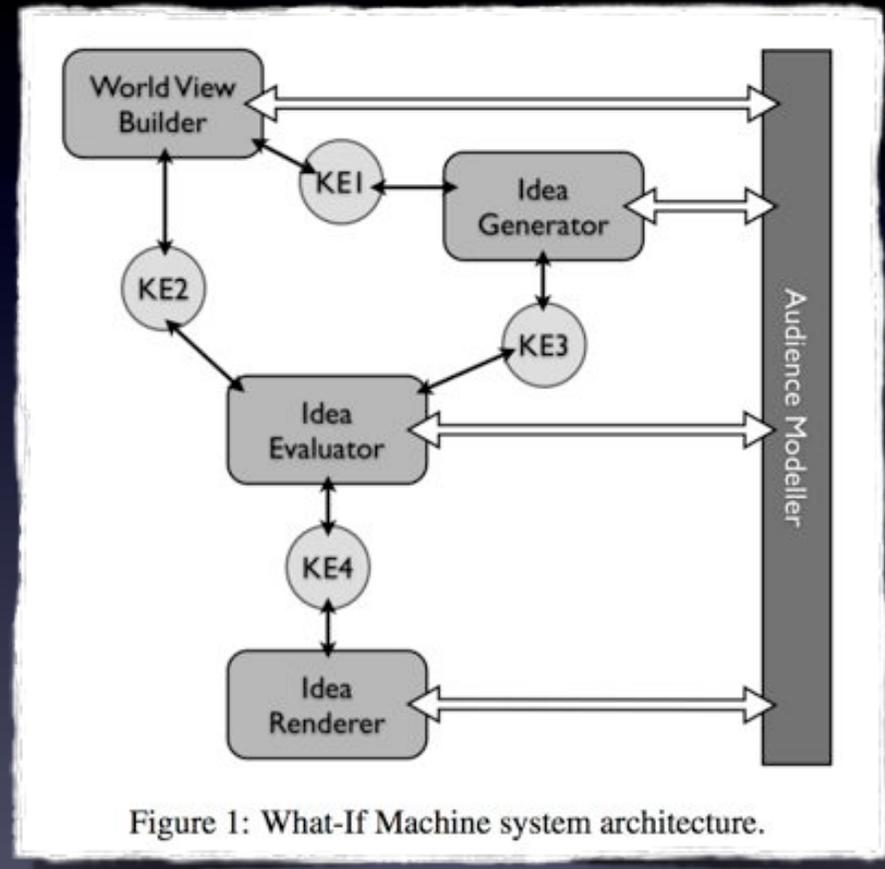


Wouldn't it be better if the software invented these processes itself?



The What-If Machine

- The Painting Fool arguably exhibits:
 - Skill, appreciation, imagination, intentionality, learning and reflection
 - But you wouldn't call it inventive...
- Three year EU FP7 project
 - Goldsmiths, Cambridge, Dublin, JSI, Madrid
 - See: www.whim-project.eu



Other Technical Talks

I could have given

- About the paint dance animations
- In detail about non-photorealistic rendering
- In detail about shape grammars
- In detail about evolutionary art
- About the evolutionary generative approach more
- In detail about the constraint generative approach
- In more detail about the collage generation
- About the Flow generative framework
- And I can give a demo of The Painting Fool anytime

Responses to Issues

- “We hope that people will eventually take The Painting Fool seriously as a creative artist, because they can no longer find a good reason why it is not”

Responses to Issues

- How can a scientist train an artist?
 - The teaching interface; work with arts students (future)
- The Painting Fool only produces what people can produce
 - Multi-media workflow; abstract art; video pieces
- The Painting Fool doesn't intend to paint anything
 - Collage generation; You Can't Know my Mind
- The Painting Fool is not creative because it doesn't have appreciation or imagination
 - Emotion detection; generative art; fitness function generation; background gen.
- The Painting Fool will always do just what you tell it to do
 - Generative art; collage generation; text-based projects;

These are my conclusions...

- Never forget that painting is a cognitive process and should be automated from an AI perspective
- You will have to combine dozens of approaches in order to get something similar to a painter
- You can't build an automated painter without addressing many of the philosophical issues of the field
- You will have to manage the public perception of creativity in software as much as you will have to advance technically... work with journalists if you can
- It's a good idea to exhibit the artwork produce, in order to get feedback and change opinions

The Painting Fool

Questions?

And would anyone like a portrait?

www.thepaintingfool.com



Portraits of the Artist's Programmer as a Grotesque Younger Man, Paris 2011