Goldsmiths
UNIVERSITY OF LONDON

Computational Creativity Group

Computational Creativity Autumn School II

Philosophical Issues

Simon Colton

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

ccg.doc.gold.ac.uk

Aims

- This morning:
  - To raise and discuss some of the issues that affect us all in Computational Creativity research
- This afternoon:
  - To address these philosophical issues with some practical guidance on how to build and assess creative systems

Computational Creativity...

Creative responsibilities

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

Audience participation

Also note the deliberate lack of mention of *value* of generated artefacts (poems, paintings, theorems, etc.) and the lack of mention of comparison with *people*

Overview

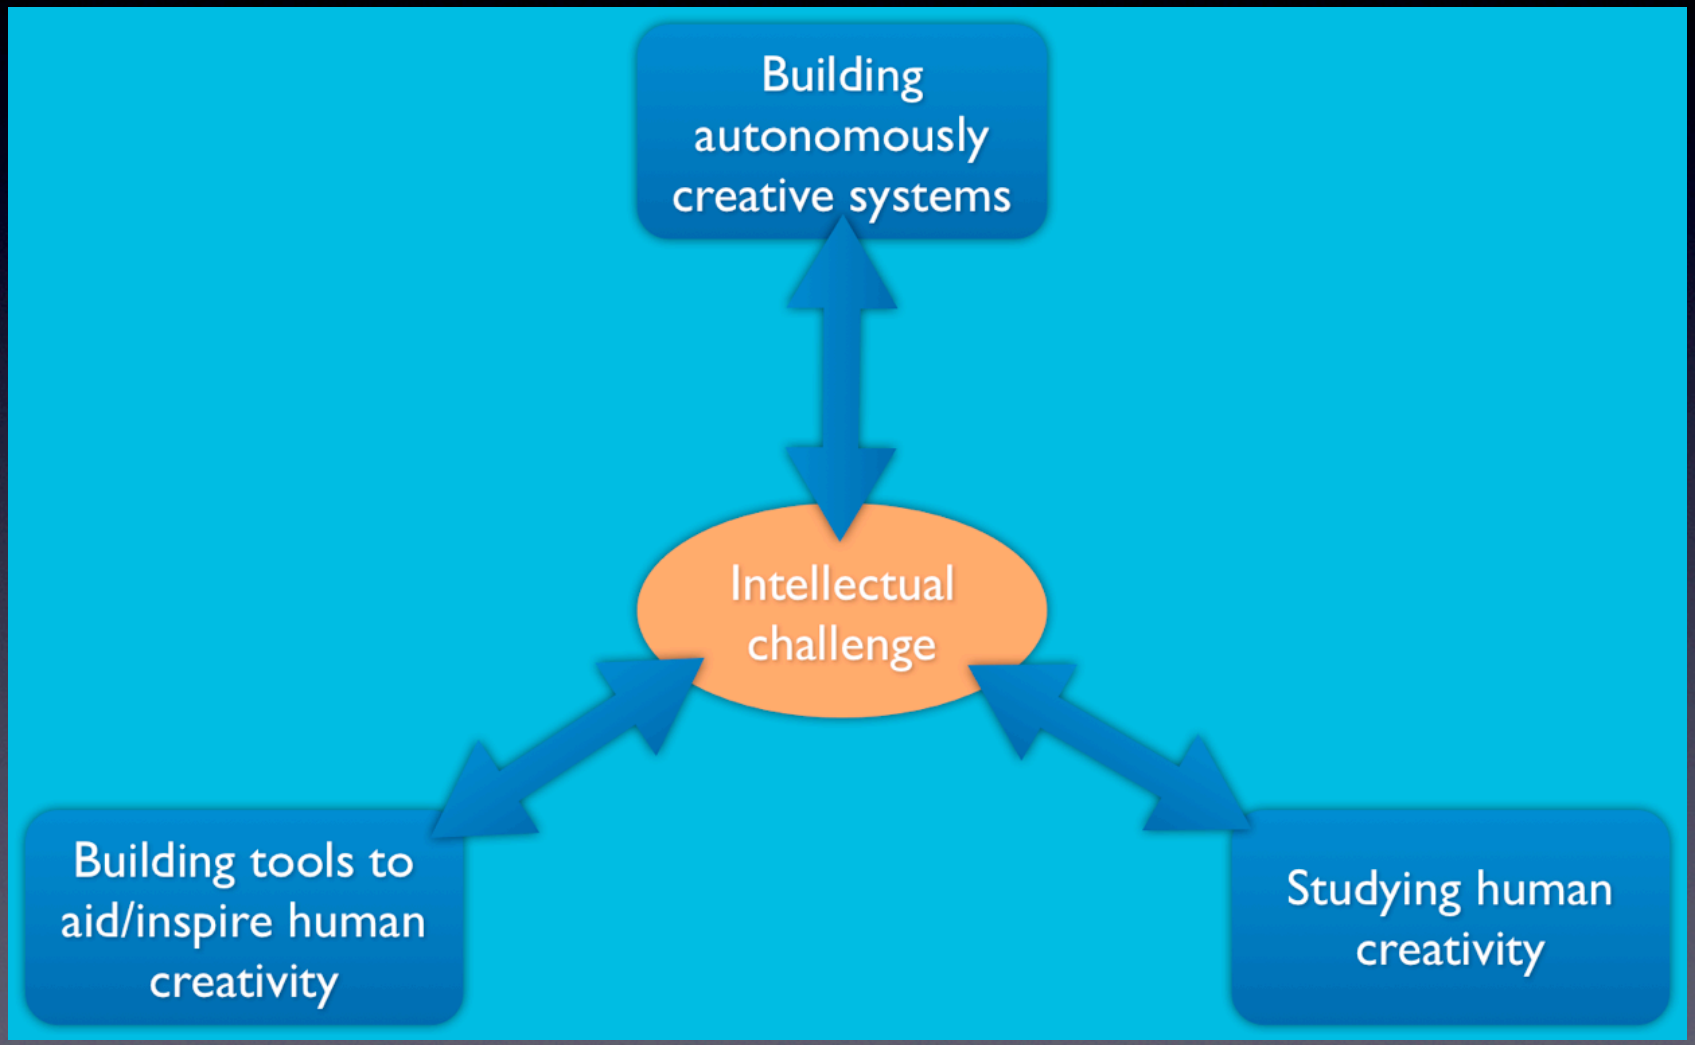
Building autonomously creative systems
Intellectual challenge
Building tools to aid/inspire human creativity
Studying human creativity

Overview

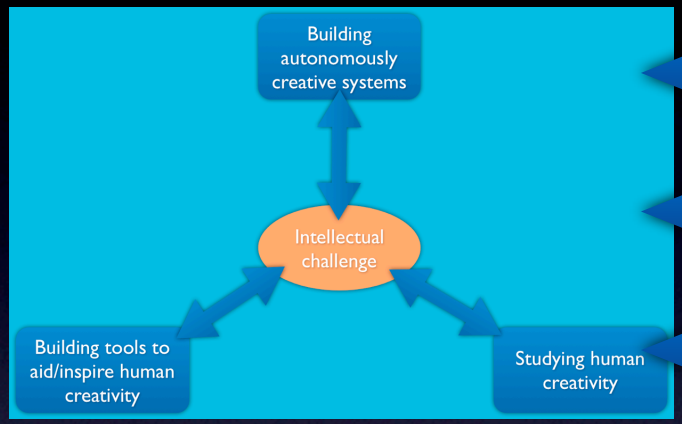

Cognitive sciences, philosophy

AI and other computational sciences

Cultural domains

Music | Art | Design | Maths | Science | Games | Poetry | Stories | Recipes

Cultural Outputs

art exhibitions + prints, scientific publications, concerts, books + anthologies, games, dinner parties

Industrial Involvement

video games companies, art galleries, advertisers (graphics + design firms, music software companies)

Issues in the Field

- With the word ‘creative’
- With handing over creative responsibilities
- With evaluating software which creates
- With software not being human

Some Difficult Notions to Digest...

There's no such thing as creativity

We shouldn't agree on how people perceive creativity

Psychology envy can be a bad thing

Levelling the playing field can go badly wrong

Output quality and autonomy of software
can be inversely proportional

We don't all agree in Computational Creativity!

1. Our issues with the word ‘creative’

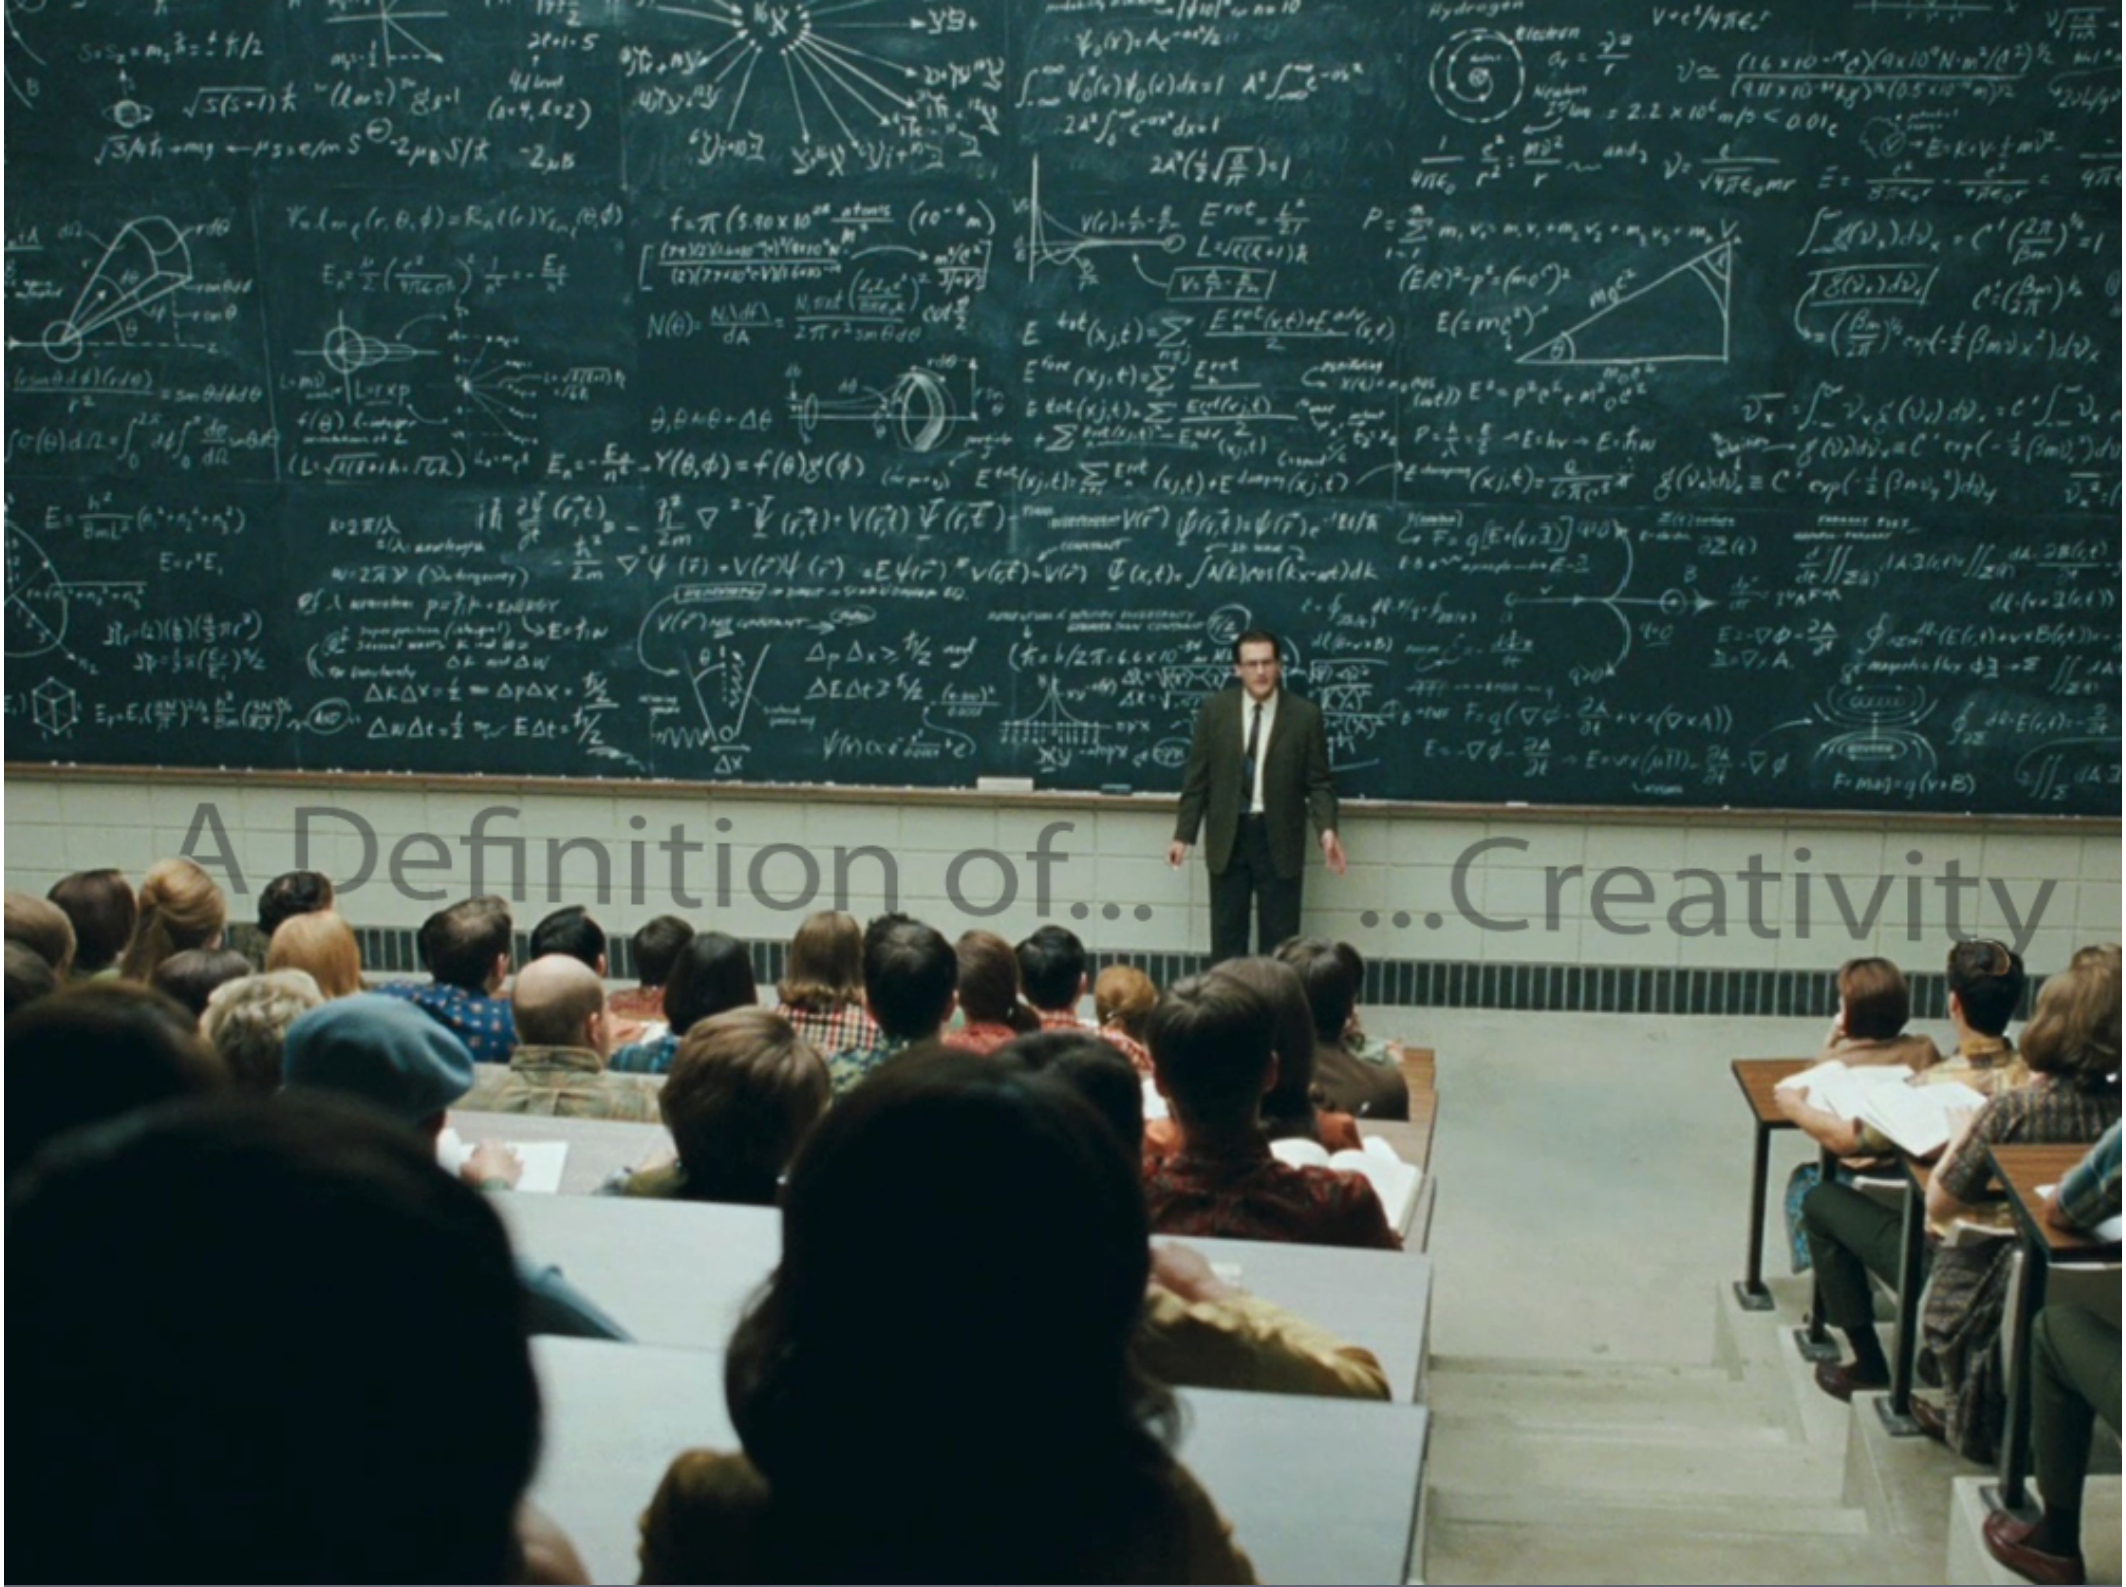
A Definition of...
...Creativity

Essentially contested concept

From Wikipedia, the free encyclopedia

In a paper delivered to the Aristotelian Society on 12 March 1956,[1] Walter Bryce Gallie (1912–1998) introduced the term **essentially contested concept** to facilitate an understanding of the different applications or interpretations of the sorts of abstract, qualitative, and evaluative notions[2] — such as "art" and "social justice" — used in the domains of aesthetics, political philosophy, philosophy of history, and philosophy of religion.

Garver (1978) describes their use as follows:

> The term essentially contested concepts gives a name to a problematic situation that many people recognize: that in certain kinds of talk there is a variety of meanings employed for key terms in an argument, and there is a feeling that dogmatism ("My answer is right and all others are wrong"), skepticism ("All answers are equally true (or false); everyone has a right to his own truth"), and eclecticism ("Each meaning gives a partial view so the more meanings the better") are none of them the appropriate attitude towards that variety of meanings.[3]

Essentially contested concepts involve widespread agreement on a concept (e.g., "fairness"), but not on the best realization thereof.[4]

They are "*concepts the proper use of which inevitably involves endless disputes about their proper uses on the part of their users*",[5] and these disputes "*cannot be settled by appeal to empirical evidence, linguistic usage, or the canons of logic alone*".[6]

Contents [hide]

Essentially contested concept

From Wikipedia, the free encyclopedia

In a paper delivered to the Aristotelian Society on 12 March 1956,[1] Walter Bryce Gallie (1912–1998) introduced the term **essentially contested concept** to facilitate an understanding of the different applications or interpretations of the sorts of abstract, qualitative, and evaluative notions[2] — such as "art" and "social justice" — used in the domains of aesthetics, political philosophy, philosophy of history, and philosophy of religion.

Garver (1978) describes their use as follows:

> The term essentially contested concepts gives a name to a problematic situation that many people recognize: that in certain kinds of talk there is a variety of meanings employed for key terms in an argument, and there is a feeling that dogmatism ("My answer is right and all others are wrong"), skepticism ("All answers are equally true (or false); everyone has a right to his own truth"), and eclecticism ("Each meaning gives a partial view so the more meanings the better") are none of them the appropriate attitude towards that variety of meanings.[3]

Essentially contested concepts involve widespread agreement on a concept (e.g., "fairness"), but not on the best realization thereof.[4]

They are "*concepts the proper use of which inevitably involves endless disputes about their proper uses on the part of their users*",[5] and these disputes "*cannot be settled by appeal to empirical evidence, linguistic usage, or the canons of logic alone*".[6]

Contents [hide]

Credit to Anna Jordanous for highlighting creativity in this context

That Word...

- 'Creative' is just a word that one person uses to describe another person, like 'funny'. It's a word to describe a *perception* people have, not an inherent property of someone
  - If everyone agrees that they perceive someone as being creative, then it's fair to call them creative, not just 'perceived to be creative'
- It can also be used to describe a process/behaviour... often shorthand for the person
- Artefacts are generally not 'creative'. In common parlance, a 'creative building' or a 'creative metaphor' can mean one of a number of different things
- As scientists, we should be more precise in our usage of this word
  - We might disagree between ourselves, but we should be internally coherent
- If we ask vague questions about the "creativity" of a person, process, or worse, the "creativity" of a building or poem, we should expect to learn zero from the study, except (yet again) that people use the word in different ways
  - We have plenty of other words we could use to help people assess quality of output, e.g., beautiful. And we have plenty of other words to use in domains, e.g., artistic

In Computational Creativity research, when talking about software, we should only use the word ‘creative’ to describe how people perceive what our software actually does

And we shouldn’t ask that question directly

2. Our issues with handing over creative responsibilities

Weak and Strong Computational Creativity Subprojects

- Weak subprojects
  - "I want software to create wonderful artefacts of type X"
  - "I'm more interested in the domain of X, and how we can contribute to that culture than to simulating creativity"
- Strong subprojects
  - "I want to build software which is one day taken seriously as being creative in generating wonderful artefacts of type X"
  - "I'm more interested in the study of creativity in software, and I want to use domain X to further study that"

Some Painting Fool Subprojects

Weak

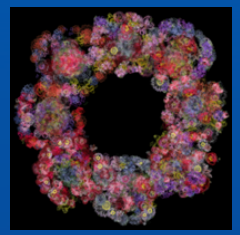

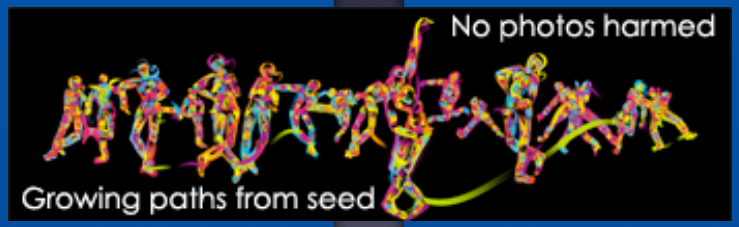

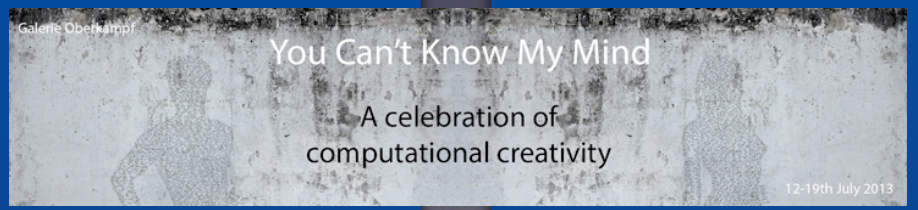

Strong

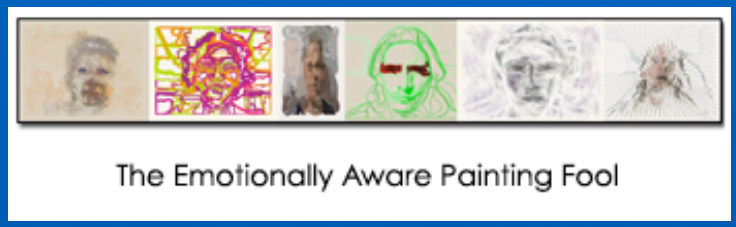

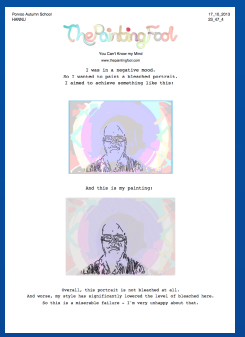

ThePaintingFool

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 portrait is not bleached at all.
And worse, my style has significantly lowered the level of bleached here.
So this is a miserable failure - I'm very unhappy about that.

Being Seen to Be AI

Automated Poetry Generation

- Strongly creative software cannot just produce valuable artefacts (poems, sonatas, theorems, paintings)
  - It has to do so in intelligent and (relatively) difficult to follow ways (not necessarily randomly)
  - And it needs to convince audiences that it has behaved in interesting and creative ways
- Will hopefully turn a vicious circle into a virtuous circle
- Practical implications:
  - Software should produce commentaries, then stories and ultimately be able to answer questions...

Circadian No. 39

[ ] [ ] of a [ ],
[ ] [ ] of a [ ].

In the morning, I am [ ]
Like the [ ] of a [ ].
But the morning grows more [ ]
Than the [ ] of a [ ].
And the mid-day makes me [ ]
Like the [ ] of a [ ].

Active [ ] of a [ ],
[ ] [ ] of a [ ].

In the daytime, I am [ ]
Like the [ ] of a [ ].
But the sunlight grows more [ ]
Than the [ ] of a [ ].
And the day, it makes me [ ]
Like the [ ] of a [ ].

[ ] [ ] of a [ ],
[ ] [ ] of a [ ].

In the [ ], I am [ ]
Like the [ ] of a [ ].
But the evening grows more [ ]
Than the [ ] of a [ ].
And the [ ] makes me [ ]
Like the [ ] of a [ ].

[ ] [ ] of a [ ],
[ ] [ ] of a [ ].

In the night-time, I am [ ]
Like the [ ] of a [ ].
But the moonlight grows more [ ]
Than the [ ] of an [ ].
And the darkness makes me [ ]
Like the [ ] of a [ ].

[ ] [ ] of a [ ],
[ ] [ ] of a [ ].

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

Shiny luster of a diamond,
Homey feeling of a bed.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Active frenzy of a beehive,
Dreary blackness of a cave.

Famous fervor of a poet,
Wily movement of a cat.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

- Repetition
- Frequency
- Stemming
- Metre
- Rhyme
- Non-rhyming
- Sentiment
- Advertising

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

- Repetition
- **Frequency**
- Stemming
- Metre
- Rhyme
- Non-rhyming
- Sentiment
- Advertising

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

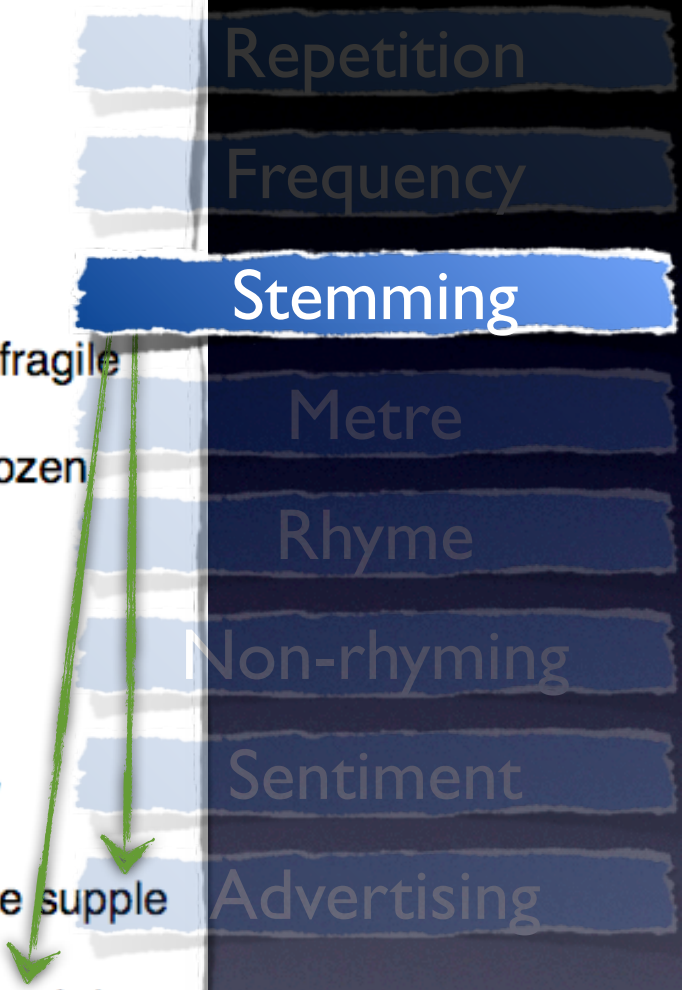

Circadian No. 39

da da da da di di da da
Stealthy swiftness of a leopard,
Happy singing of a bird.
da da da da di di daaa
In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

- Repetition
- Frequency
- Stemming
- **Metre**
- Rhyme
- Non-rhyming
- Sentiment
- Advertising

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

- Repetition
- Frequency
- Stemming
- Metre
- **Rhyme**
- Non-rhyming
- Sentiment
- Advertising

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

Repetition

Frequency

Stemming

Metre

Rhyme

Non-rhyming

Sentiment

Advertising

Circadian No. 39

Stealthy swiftness of a leopard
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more supple
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard
Happy singing of a bird.

- Repetition
- Frequency
- Stemming
- Metre
- Rhyme
- Non-rhyming
- **Sentiment**
- Advertising

Circadian No. 39

Stealthy swiftness of a leopard,
Happy singing of a bird.

In the morning, I am loyal
Like the comfort of a friend.
But the morning grows more lifeless
Than the fabric of a rag.
And the mid-day makes me nervous
Like the spirit of a bride.

Active frenzy of a beehive,
Dreary blackness of a cave.

In the daytime, I am slimy
Like the motion of a snake.
But the sunlight grows more comfy
Than the confines of a couch.
And the day, it makes me tasty
Like the flavor of a coke.

Shiny luster of a diamond,
Homey feeling of a bed.

In the evening, I am solid
Like the haven of a house.
But the evening grows more fragile
Than the mindset of a child.
And the twilight makes me frozen
Like the bosom of a corpse.

Famous fervor of a poet,
Wily movement of a cat.

In the night-time, I am hollow
Like the body of a drum.
But the moonlight grows more su
Than the coating of an eel.
And the darkness makes me subtle
Like the color of a gem.

Stealthy swiftness of a leopard,
Happy singing of a bird.

- Repetition
- Frequency
- Stemming
- Metre
- Rhyme
- Non-rhyming
- Sentiment
- Advertising

Repetition Frequency Stemming Metre

Searching Questions #1

My point is:

There are possibly benefits (in terms of increased value projected onto the artefacts) to be gained by presenting explicit information about the process behind the generation of an artefact

We can provide this information through technical papers, talks and notes, but it would be better for the software to do this itself, because it simulates appreciation and/or reflection, which we value in creative people

- D
- D ease
- D ving
- V

about these constraints, and throw in some more stuff?

Rhyme Non-rhyming Sentiment Advertising

Searching Questions #2

My point is:

There may be a virtuous circle (rather than the usual vicious circle) we can get into, if we know that the software has done something difficult and intelligent. People will always read/see/hear things that aren't in artefacts by design. If those people know a little about the intelligence of the system, then they might be prepared to give software the benefit of the doubt

Software could make stuff up (tell a story...)

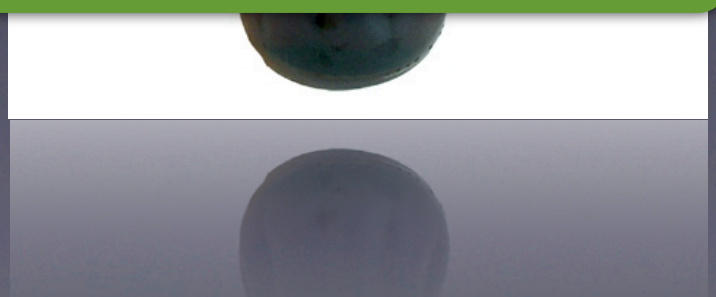

Searching Questions #3

- ”?
- This doesn’t feel like Computational Creativity research yet
- The software needs to take on more creative responsibility ”?

My point is:

This is a clearly a Simon-generated poem

in the production of the poems...

Handing over Creative Responsibility

The FACE Model

- We can judge progress in terms of the types of generative acts that software undertakes (regardless of the value of its output)

- Examples - simile multiplication, phrase overlapping
- Concepts - generating templates
- Aesthetics - inventing measures of value
- Framing - producing a commentary

Handing over Creative Responsibility

Poetry Generation Pipeline

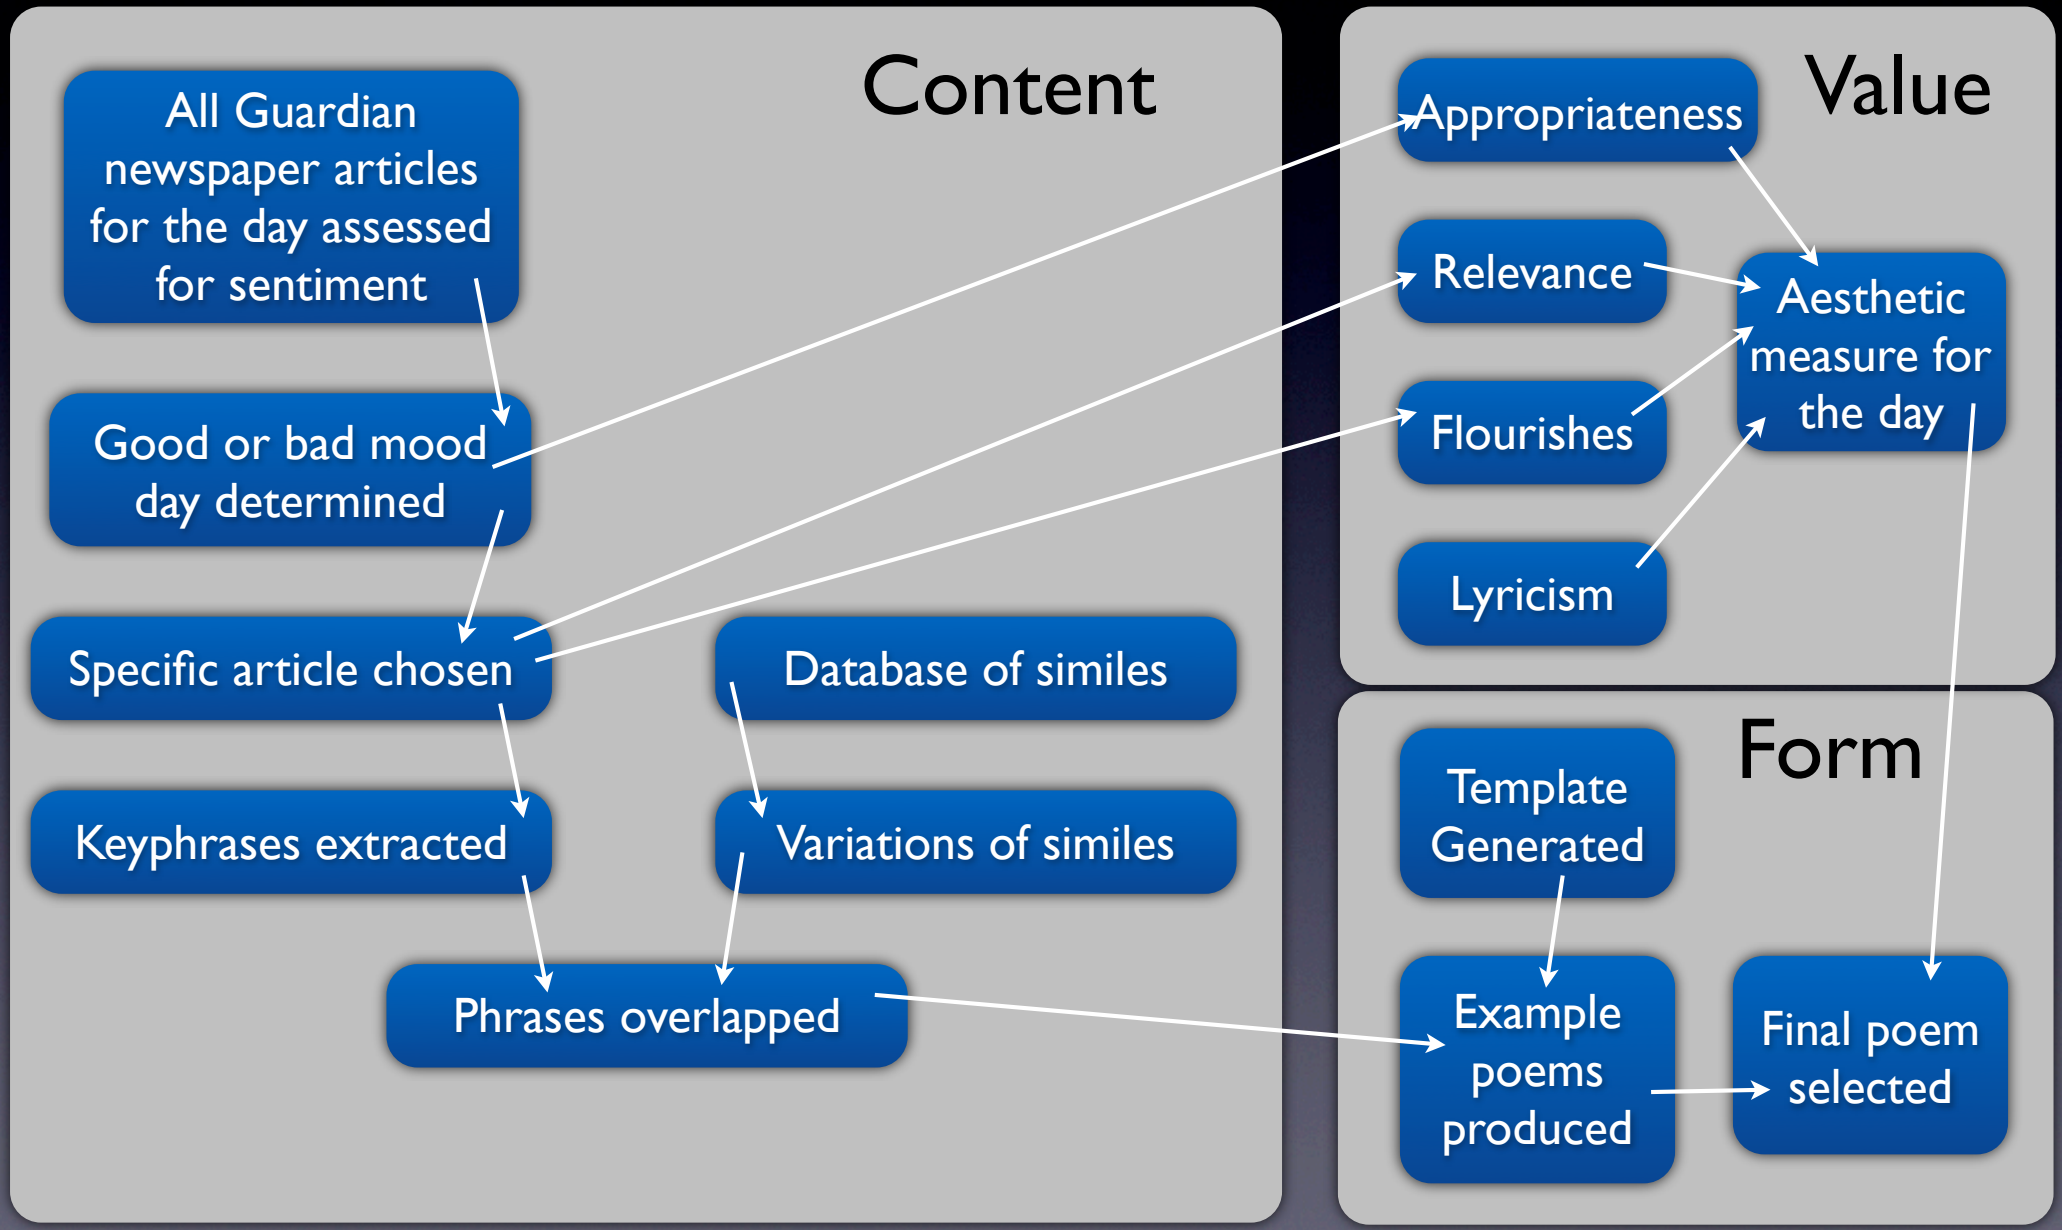

New Poems...

It was generally a good news day. I read a story in the Guardian culture section entitled: "South Africa's ANC celebrates centenary with moment in the sun". It talked of south africans, interfaith prayers and monochrome photos. Apparently, "The heroic struggle against a racist regime was remembered: those thousands who sacrificed their lives in a quest for human rights and democracy that took more than eight decades" and "At midnight he watched with amusement as Zuma lit the centenary flame, at the second attempt, with some help from a man in blue overalls marked 'Explosives'". I wanted to write something highly relevant to the original article. I wrote this poem.

Blue overalls

the repetitive attention of some traditional african chants
a heroic struggle, like the personality of a soldier

an unbearable symbolic timing, like a scream
blue overalls, each like a blueberry
some presidential many selfless leaders

oh! such influential presidents
such great presidents
blueberry-blue overalls

lark-blue overalls
a knight-heroic struggle

Discussion...

- Taken out of context, e.g., in an evaluation test or a Turing-style test, poem #1 would probably score higher as a "poem shaped object" than poem #2
- But, when we read about *how* the software produced the poems, it's likely that people will project more creativity onto the software producing poem #2
- So, the more sophisticated software represents a backward step in automated poetry generation, but an advance in computational creativity
  - The Latent Heat effect in Computational Creativity

The Latent Heat Effect

in Computational Creativity

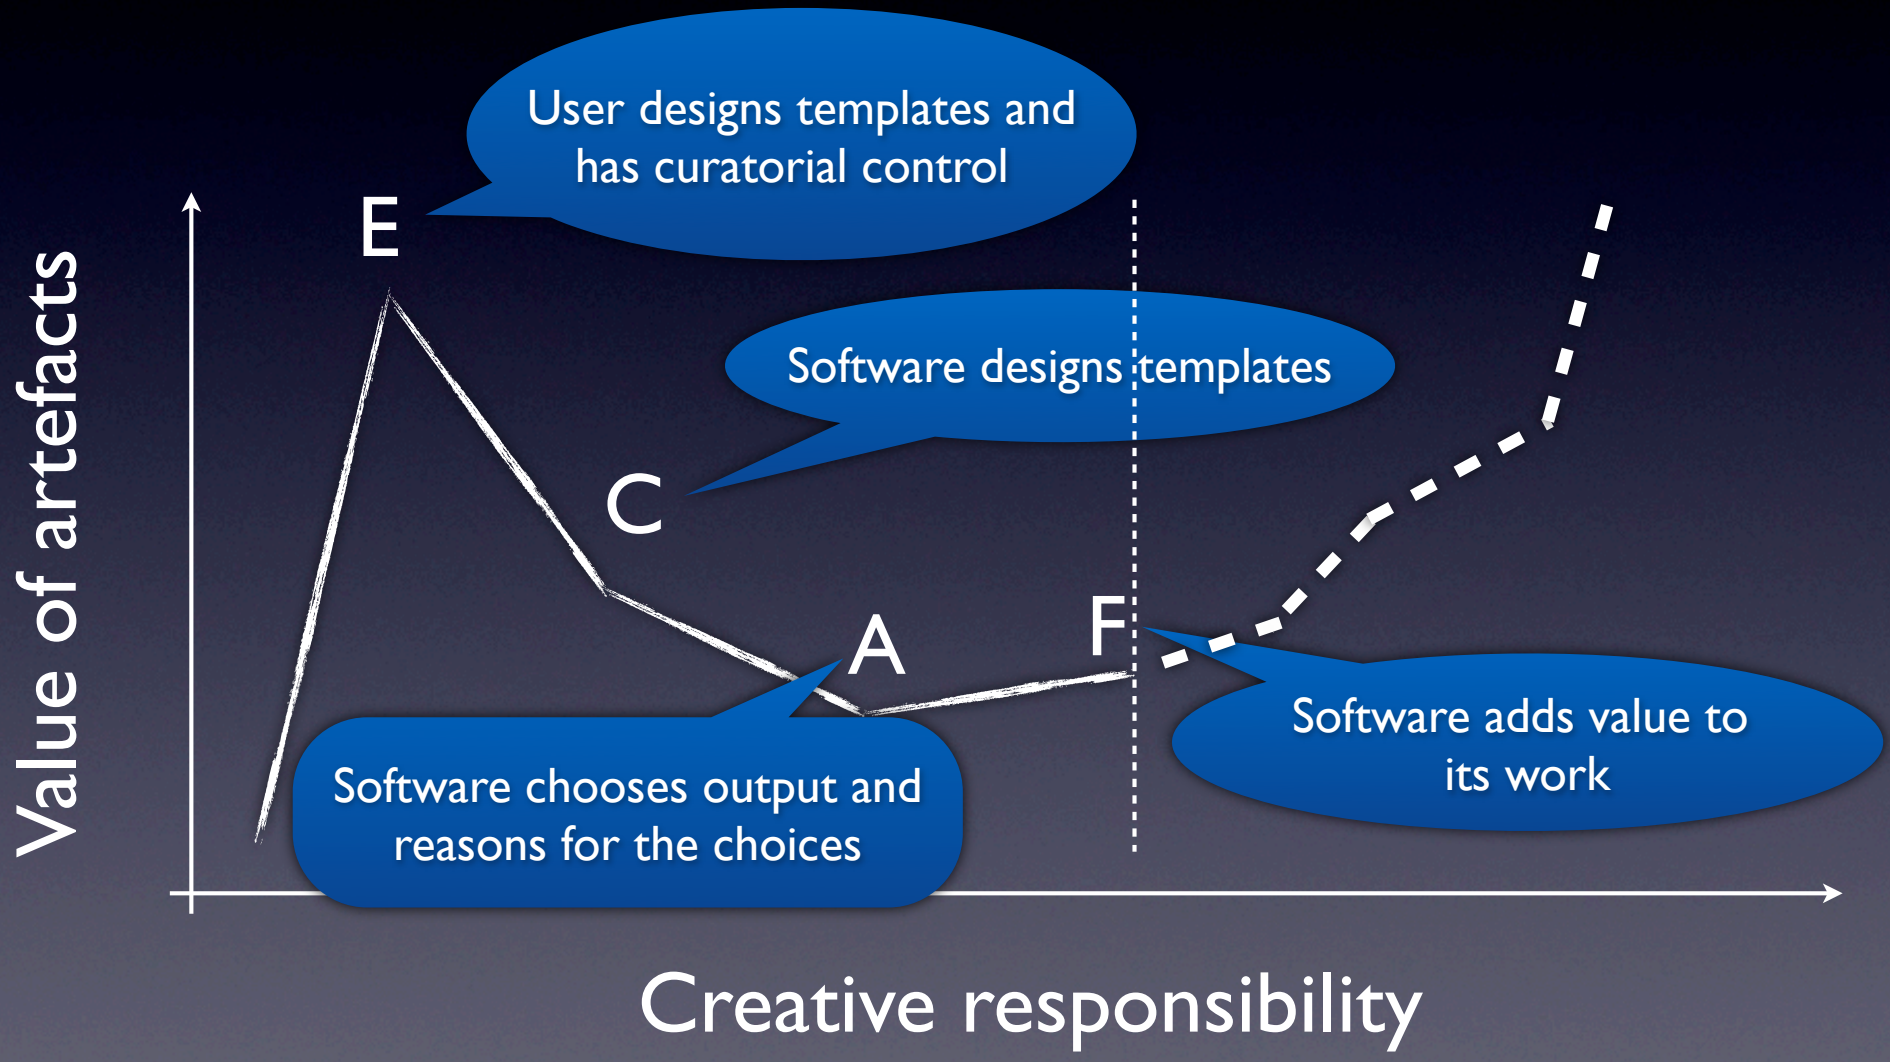

3. Our issues with evaluating software which creates

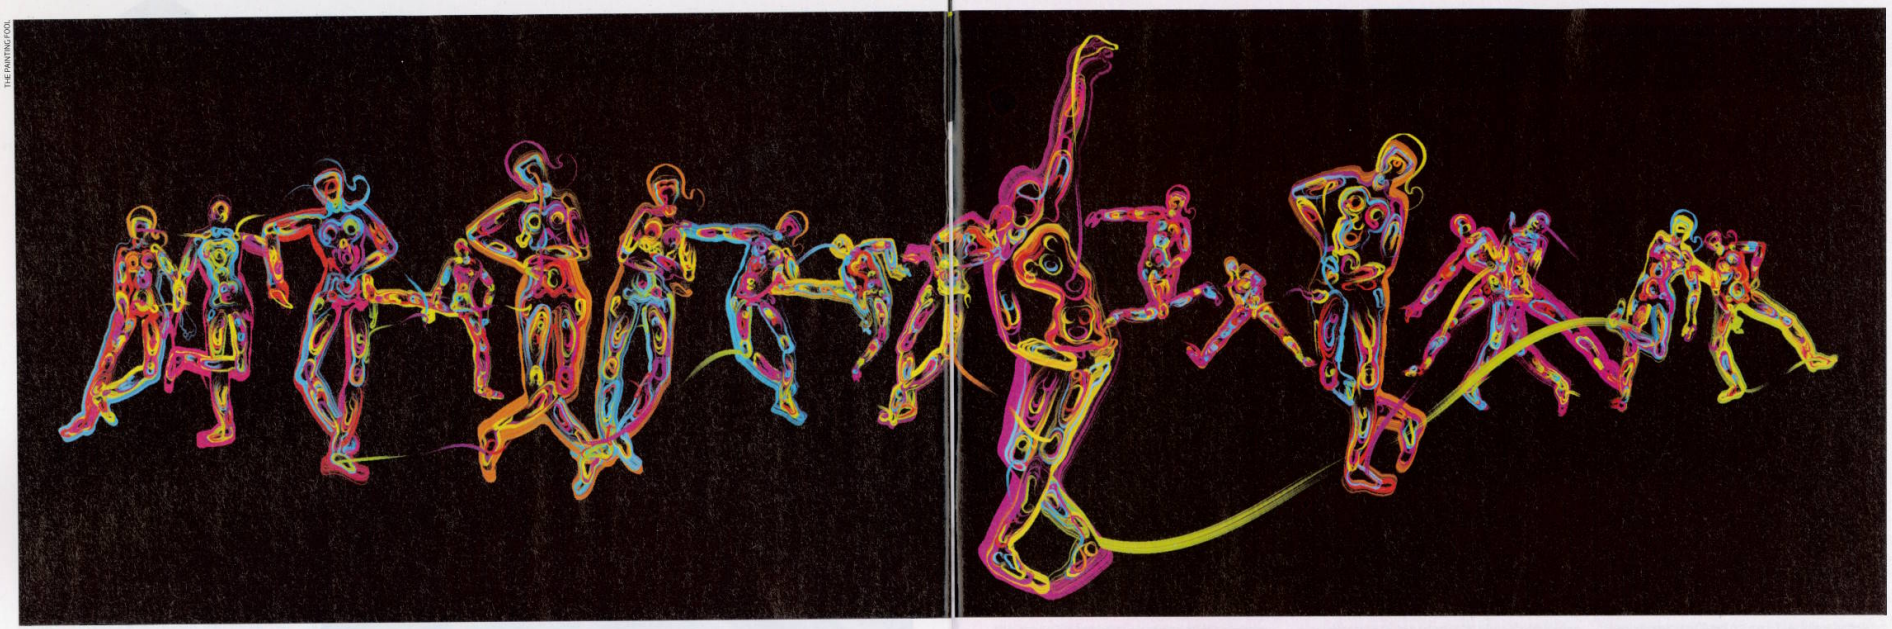

THE PAINTING FOOL

Creative sparks

Does it really take a human to make a masterpiece? **Catherine de Lange** sizes up the artificial artists forcing us to change our ideas about creativity

IN A loft overlooking the rooftops of one of the buzzing artistic neighbourhoods of Paris, France, Simon Colton is carefully unfurling one giant painting after another. I have waited some time to see these, and am unsure what to expect. To dislike them would be a disappointment, but easy. If I think these paintings are any good, however, then I might have to reconsider my own creative talents. In fact, they might even challenge my understanding of what it means to be human.

The thing is, these paintings are not the work of an ordinary artist. Nor of Colton, who is a computer scientist based at Imperial College London. Instead, they have been created by a piece of software that can seek artistic inspiration and, arguably, has a rudimentary imagination. Called the Painting Fool, it may have been designed by Colton, but its artwork is its own.

It sounds unlikely that any computer, unguided by the human hand and eye, could create artwork with any feeling or resonance. How could it be creative without having experienced the world? Now, as I take a first glimpse at the paintings, will I be forced to reconsider? Could software, which has no shared experience with my own, create a painting that touches me?

The Painting Fool is one of a growing number of computers which, so their makers claim, possess creative talents. Classical music by an artificial composer has had audiences enraptured, and even tricked them into believing a human was behind the score. Artworks painted by a robot have sold for thousands of dollars and been hung in prestigious galleries. And software has been built which creates art that could not have been imagined by the programmer. "It scares a lot of people," says Geraint Wiggins, a computational creativity researcher at Goldsmiths, University of London. "They are worried it is taking away something special from what it means to be human."

While some animals such as crows and monkeys have displayed traits that could be labelled as limited creativity, we are the only species to perform sophisticated creative acts regularly. If we can break this process down into computer code, where does that leave human creativity? "This is a question at the very core of humanity," says Wiggins.

To some extent, we are all familiar with computerised art. Software that is used to create or manipulate art is ubiquitous, but these are mere tools for a human artist. The question is: where does the work of a person stop and the creativity of the computer begin?

Consider one of the oldest machine artists, Aaron, a robot that has had paintings exhibited in London's Tate Modern and the San Francisco Museum of Modern Art. In some respects, then, Aaron passes some kind of creative Turing test – its works are good enough to be exhibited alongside some of the best human art and people spend good money on them.

Aaron can pick up a paintbrush with its robotic arm and paint on canvas on its own. Impressive, perhaps, but it can never break free from the tightly controlled rules it has been given by its programmer, the artist and founder of machine fine art, Harold Cohen. So I remain unconvinced that Aaron is much more than a tool to realise Cohen's own creative ideas. Colton also dismisses the machine as "rather limited" because "it still only creates one kind of artwork: people in a room with pot plants".

Colton is keen to make sure the Painting Fool doesn't fall foul of the same criticism,

> "The machine will wake up in the morning and look at newspaper headlines for source material"

Composed and painted by a computer – can it be enjoyed and lauded by a human being?

and so has sought to give it as much autonomy as possible. Although the software does not physically apply paint to canvas, it simulates many styles digitally, from collage to paint strokes. One of the first paintings Colton shows me is a touching portrait of a young, fragile-looking girl with porcelain skin and long brown hair. I am impressed that a computer could capture such subtleties, until Colton tells me the software just applied its own painting style to photographs of the girl.

That sounds like cheating to me, but Colton assures me that was an early work. Today, the Painting Fool only needs minimal direction and can come up with its own concepts by going online for source material. "I don't even give it the notion of a person or a topic," says Colton. "It will wake up in the morning and look at the newspaper headlines." The software runs its own web searches and trawls through social media websites such as Twitter and Flickr. The idea is that this approach will let it produce art that is meaningful to the audience, because it is essentially drawing on the human experience as we act, feel and argue on the web.

In 2009, Colton and graduate student Anna Krzeczkowska asked the Painting Fool to ›

produce its own interpretation of the war in Afghanistan, based on a news story. The result is a striking juxtaposition of Afghan citizens, explosions and war graves. "This piece struck a chord with me, and shows the potential for the software to add poignancy and intentionality to its paintings," says Colton.

The Painting Fool is now beginning to display a kind of imagination too, creating pictures from scratch. One of its original works, part of a series that Colton called *Four Seasons*, depicts fuzzy panels of simple landscapes (see image, below). I think it looks rather mechanical.

Comparable to Bach

Having said that, I am swayed by Colton's argument that my reaction arises from my double standards towards software-produced and human-produced art. After all, he says, consider that the Painting Fool painted the landscapes without referring to a photo. "If a child painted a new scene from its head, you'd say it has a certain level of imagination, even if it's just a little bit," he points out. "The same should be true of a machine."

Software bugs can also lead to unexpected results. I see this for myself when Colton shows me some paintings of a chair, which came out black and white thanks to a glitch. It gives the work an eerie, ghostlike quality. Human artists like Ellsworth Kelly are lauded for limiting their colour palate – so why should computers be any different? Nonetheless, these mechanical steps towards creating new styles are barely comparable to the talents of, say, Picasso or Mozart. Or are they?

Researchers like Colton don't believe it is right to compare machine creativity directly to that of humans, who "have had millennia to develop our skills". Others, though, are fascinated by the prospect that a computer might create something as original, emotional and subtle as our best artists. So far, only one has come close.

> "After discovering the truth, one music lover told Cope he had 'killed music' and tried to punch him"

One day, David Cope was suffering from "composer's block". He had been commissioned to write an opera, but was struggling to come up with the goods. If only a computer could understand his style, he thought, and help him write new material. That idea was the starting point for what was to become one of the most controversial pieces of creative software to date. Cope came up with a program called Experiments in Musical Intelligence, or EMI. He fed in musical scores and out popped new material in the composer's style. Not only did EMI create compositions in his style, but also that of the most revered classical composers, including Bach and Mozart.

To my untrained ear, it sounds like any other classical music. I found the purported Chopin, in particular, to be rich and emotional. Audiences who heard the music have been moved to tears, and EMI even fooled classical music experts into thinking they were hearing genuine Bach. If ever there were a successful Turing test for computational creativity, that had to be it.

THE PAINTING FOOL

If a child produced these scenes, you might say they were imaginative

Not everyone was impressed, however. Some critics, such as Wiggins, have blasted Cope's work as pseudoscience, saying his explanation of how the software works is "smoke and mirrors" leaving others unable to reproduce the results (*Literary and Linguist Computing*, vol 23, p 109). Douglas Hofstadter, at Indiana University, Bloomington, says Cope merely scratches at the surface of creativity, using superficial elements of an artist's work to create replicas, which still rely on the original artist's creative impulses.

Nonetheless, for others EMI's ability to mimic Bach or Chopin has serious implications. If it is so easy to break down the style of some of the world's most original composers into computer code, that means some of the best human artists are more machine-like than we would like to think. Indeed, when audiences found out the truth about EMI they were often outraged – one music lover allegedly told Cope he had "killed music" and tried to punch him. Amid such controversy, in 2004 Cope decided that EMI's time was up, and destroyed its vital databases.

But why did so many people love the music, yet recoil when they discovered what composed it? A study by David Moffat, a computer scientist at Glasgow Caledonian University in the UK, provides a clue. He asked both expert musicians and non-experts to assess the creative worth of six compositions. The participants weren't told beforehand whether the tunes were composed by humans or computers, but were asked to guess, and then rate how much they liked each one. Perhaps unsurprisingly, people who thought the composer was a computer tended to dislike that piece more than those who believed it was human. This was true even among experts, who you might think would be more objective in their analysis of musical quality.

Where does this prejudice come from? Psychologist Paul Bloom of Yale University has a suggestion: he reckons part of the pleasure we get from art comes from our perception of the creative process behind it. This can give it an "irresistible essence", says Bloom. This idea explains why a painting loses its value if exposed as a fake, even though we might have loved it when we thought it was an original. Indeed, experiments by psychologist Justin Kruger of New York University have shown that people's enjoyment of an artwork increases if they think more time and effort was needed to create it (*Journal of Experimental Social Psychology*, vol 40, p 91).

Similarly, Colton thinks that when people experience art, they engage in a discourse with the artist. We wonder what the artist might have been thinking, or ponder what they are trying to tell us. With computers producing art, this speculation is cut short – there's nothing to explore. But as the software becomes increasingly complex, finding those greater depths in the art may become possible. That's why Colton asks the Painting Fool to tap into online social networks for its inspiration: hopefully this way it will choose themes that will already mean something to us.

THOMAS STRUTH

Our enjoyment of art is influenced by the time and effort it required

Unconscious creativity

Hofstadter thinks the more complex machines become, the more easily we will accept their art – especially if they can interact more with the physical world. If robots bumped into things and had goals, successes and failures, then that might be enough. "They would be sort of pathetic and laughable and once in a while heroic," he says. "I don't think people would be uncomfortable with creatures like that writing an essay or composing a piece of music or painting a picture."

Yet the fact that machines now lack this kind of self-awareness is perhaps the most irksome element of computational creativity. How can you be creative without even being conscious? Surprisingly, it is not a computer scientist who talks me out of this reaction, but a neuroscientist. Our brains work creatively even when we aren't consciously thinking about it, says Arne Dietrich at the American University of Beirut in Lebanon. Just think back to a time when the solution to a problem you had forgotten about just popped into your head. "We know that there are several different types of creativity – some of them are conscious, some of them unconscious," he says. "Creativity can happen when you try effortfully, or it can happen in your sleep."

In any case, Dietrich believes that the creative brain might work much like software. Neuroscientists suspect that creativity is essentially about discovery, rather than anything mystical. "It's a mechanical process in the brain that generates possible solutions and then eliminates them systematically," Dietrich says. He believes our tendency to dismiss computational creativity as inferior to our own comes from an ingrained dualism in human culture: "We are over-evaluating ourselves and underestimating them. As a neuroscientist, I tackle the brain as a machine, and I don't see machine creativity as different." Suddenly, the idea that the human brain has a unique claim to creative talents seems a limited perspective.

Back in Paris, Colton continues to show me painting after painting, all signed by the Painting Fool. Some of the work genuinely speaks to me. One of my favourites, called *The Dancing Salesman Problem*, features colourful human figures dancing on a black background (see main image, page 42). Again, the software did not base its composition on existing pictures. The dancers are painted in long, flowing strokes, so they appear full of movement: they contort into beautiful poses, and the bright colours bring the scene to life. The work could never be to everyone's taste, but I would have stopped to look at it in a gallery, and I don't mind that it was created by a machine.

I have come to appreciate that computers can create subtle and original artwork. Will others accept that idea? The trick, says Colton, is to stop trying to compare computer artists to human ones. If we can embrace computer creativity for what it is, and stop trying to make it look human, not only will computers teach us new things about our own creative talents, but they might become creative in ways that we cannot begin to imagine. They are creating a whole new form of art with the potential to delight, challenge and surprise us.

Will that take something away from being human? "It's not taking away anything at all," says Wiggins. "It's helping us to understand how things work. And when you understand how things work they tend to become more amazing, not less so." ■

Catherine de Lange is a writer based in London.

To see more examples of software-produced artwork, visit newscientist.com/gallery/painting-fool

e Painting Fool painted the
out referring to a photo. "If a
new scene from its head, you'd
in level of imagination, even
bit," he points out. "The same
f a machine."
s can also lead to unexpected
s for myself when Colton
paintings of a chair, which
and white thanks to a glitch.
an eerie, ghostlike quality.
ike Ellsworth Kelly are lauded
ir colour palate – so why should
ny different? Nonetheless,
al steps towards creating new
comparable to the talents of,
Mozart. Or are they?
ike Colton don't believe it is
e machine creativity directly

starting point for what was to become one of the most controversial pieces of creative software to date. Cope came up with a program called Experiments in Musical Intelligence, or EMI. He fed in musical scores and out popped new material in the composer's style. Not only did EMI create compositions in his style, but also that of the most revered classical composers, including Bach and Mozart.

To my untrained ear, it sounds like any other classical music. I found the purported Chopin, in particular, to be rich and emotional. Audiences who heard the music have been moved to tears, and EMI even fooled classical music experts into thinking they were hearing genuine Bach. If ever there were a successful Turing test for computational creativity, that had to be it.

If a child produced these scenes, you might say they were imaginative

music lover allegedly told Cop
music" and tried to punch hir
controversy, in 2004 Cope de
time was u troyed it
But pl
ye
b d n
orth of si
The participants weren't told
whether the tunes were comp
or computers, but were asked
then rate how much they like
Perhaps unsurprisingly, peop
the composer was a compute
dislike that piece more than th
it was human. This was true e
experts, who you might think
objective in their analysis of
Where does this prejudice
Psychologist Paul Bloom of Y
a suggestion: he reckons part
we get from art comes from c
of the creative process behin
it an "irresistible essence", sa
idea explains why a painting

Hmmmmm.
If ever there was a Turing test for pastiche generation, that had to be it

A Turing Test for Being Creative

- Probably because of our longing for objectivity, it's easy to mix up the process of *generating high quality artefacts*, with *being creative*
  - Pastiches can be high quality, yet it's unlikely that you would call the artist/writer/musician particularly creative, as there's no new dialogue arising from it
- You might argue: "What does it matter, as long as high quality output is produced?"
  - Until you realise that as a society, we really value our creative individuals - perhaps more than what they produce
  - And that when we celebrate an artefact (poem/picture/sonata), we are actually (also) celebrating the creative act that led to it
  - And for strong Computational Creativity, creative behaviour is paramount

On impact and evaluation in Computational Creativity: A discussion of the Turing Test and an alternative proposal

Alison Pease[1] and **Simon Colton**[2]

Abstract. Computational Creativity is the AI subfield in which we study how to build computational models of creative thought in science and the arts. From an engineering perspective, it is desirable to have concrete measures for assessing the progress made from one version of a program to another, or for comparing and contrasting different software systems for the same creative task. We describe the Turing Test and versions of it which have been used in order to measure progress in Computational Creativity. We show that the versions proposed thus far lack the important aspect of interaction, without which much of the power of the Turing Test is lost. We argue that the Turing Test is largely inappropriate for the purposes of evaluation in Computational Creativity, since it attempts to homogenise creativity into a single (human) style, does not take into account the importance of background and contextual information for a creative act, encourages superficial, uninteresting advances in front-ends, and rewards creativity which adheres to a certain style over that which creates something which is genuinely novel. We further argue that although there may be some place for Turing-style tests for Computational Creativity at some point in the future, it is currently untenable to apply any defensible version of the Turing Test.

As an alternative to Turing-style tests, we introduce two descriptive models for evaluating creative software, the FACE model which describes creative acts performed by software in terms of tuples of generative acts, and the IDEA model which describes how such creative acts can have an impact upon an ideal audience, given ideal information about background knowledge and the software development process. While these models require further study and elaboration, we believe that they can be usefully applied to current systems as well as guiding further development of creative systems.

1 The Turing Test and Computational Creativity

The Turing Test (TT), in which a computer and human are interrogated, with the computer considered intelligent if the human interrogator is unable to distinguish between them, is principally a philo-

Computational Creativity (CC) is a subfield of AI, in which researchers aim to model creative thought by building programs which can produce ideas and artefacts which are novel, surprising and valuable, either autonomously or in conjunction with humans. There are three main motivations for the study of Computational Creativity:

- to provide a computational perspective on human creativity, in order to help us to understand it (cognitive science);
- to enable machines to be creative, in order to enhance our lives in some way (engineering); and
- to produce tools which enhance human creativity (aids for creative individuals).

Creativity can be subdivided into everyday problem-solving, and the sort of creativity reserved for the truly great, in which a problem is solved or an object created that has a major impact on other people. These are respectively known as "little-c" (mundane) and "big-C" (eminent) creativity [2]. Boden [3] draws a similar distinction in her view of creativity as search within a conceptual space, where "exploratory creativity" searches within the space, and "transformational creativity" involves expanding the space by breaking one or more of the defining characteristics and creating a new conceptual space. Boden sees transformational creativity as more surprising, since, according to the defining rules of the conceptual space, ideas within this space could not have been found before.

There are two notions of evaluation in CC: (*i*) judgements which determine whether an idea or artefact is valuable or not (an essential criterion for creativity) – these judgements may be made internally by whoever produced the idea, or externally, by someone else and (*ii*) judgements to determine whether a system is acting creatively or not. In the following discussion, by evaluation, we mean the latter judgement. Finding measures of evaluation of CC is an active area of research, both influenced by, and influencing, practical and theoretical aspects of CC. It is a particularly important area, since such measures suggest ways of defining progress in the field,[3] as well as

Turing-Style Tests

- Style 1: A dialogue where the point of the exercise is to prove that it would be fair to call your software intelligent
  - Closest to what Turing had in mind
- Style 2: A dialogue where the point of the exercise is to prove that people can't tell the difference to talking to a person and talking to your software
  - So, we implement software which often says unintelligent things
- Style 3: A comparison test with no dialogue, where the point of the exercise is to prove that the output of your software is of a similar (or higher) value to that produced by people
  - This has often been applied in Computational Creativity research

Comparison Tests

- It is certainly a milestone in the development of generative software (and for the field as a whole) if the output can be easily confused with that of people. This is because we can refer to the default position that people act creatively when they produce, and hence it is only fair to describe software similarly
- And it allows objective comparison, enabling us to show progress in implementations. Importantly, we can be seen to be scientific in our evaluation methodology
- And journalists love setting up Turing-style tests, as it both informs and worries the general public, which helps to sell newspapers...
  - New Scientist and BBC Horizon

However...

- Imagine a comparison test where the tester performs the *reveal:*
- "So, these paintings were painted by recent Royal College of Art graduates"
- "And these ones were painted by.....

  a mass murderer!"

- Wouldn't your value judgements change?

Problems 1 and 2

- Turing-style comparison tests set the computer up for a fall
  - The implicit assumption is that software should be very grateful if it is mistaken occasionally for a human
  - So, human level output becomes seen as the only goal of Computational Creativity research
- Software is NOT human!
- So, we end up missing out on possibilities where the software creates valuable, interesting artefacts in non-human ways
  - We should instead be loud and proud about the generative system being computer based, and help people to appreciate the value of computer generated creative acts

Problems 3 and 4

- Turing-style comparison tests massively underestimate the importance of process in certain domains
  - This can lead to alienation of people, certainly in the visual art world, where art theory is all about process
- Turing-style comparison tests answer the wrong question, e.g., which would you prefer, if you had to make up your mind without knowing fully how they were produced
  - Whereas in (commercial/artistic/scientific) reality, we will have full/partial disclosure of practice as well as product
  - Or should we go through this charade with our software for the rest of our lives?

Problems 5 and 6

- There are no right or wrongs in the visual arts. However, critics can severely inflict pain by saying that your work is “naive” and/or a “pastiche”
- Turing-style comparison tests might encourage software to act unintelligently, to make it seem more human, hence it could be criticised as naive
- Turing-style comparison tests definitely encourage the generation of pastiche pieces, as the measure of success is whether you have successfully imitated something which isn’t you
- Would art graduates be happy if you said their pieces all looked like Monet pictures?

Well put by Alison...

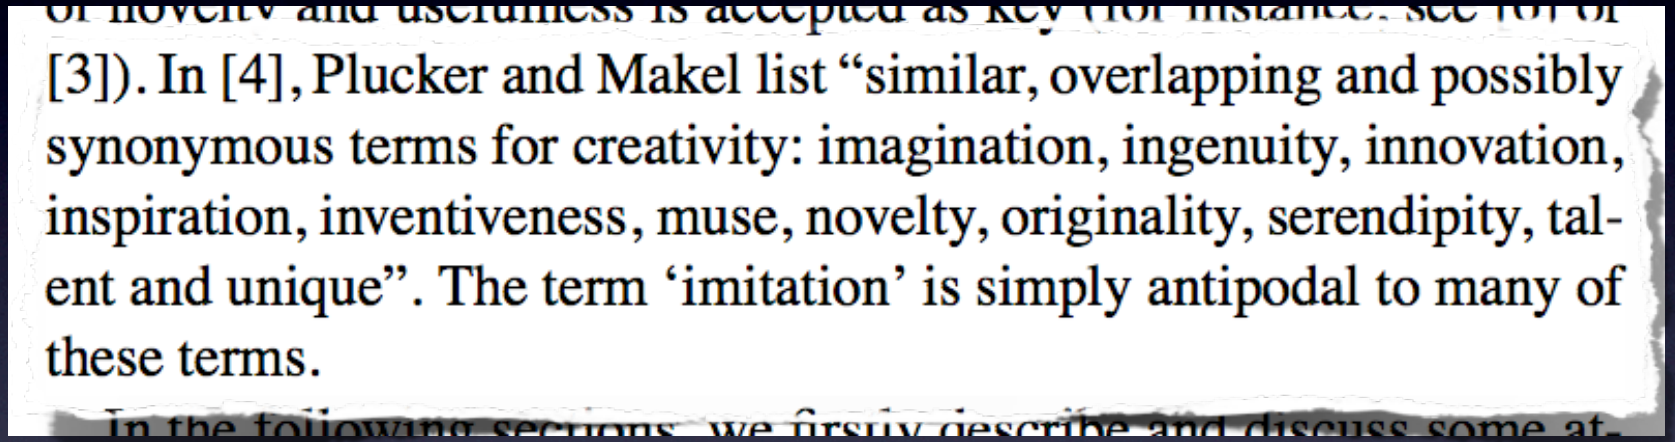

of novelty and usefulness is accepted as key (for instance: see [6] or [3]). In [4], Plucker and Makel list "similar, overlapping and possibly synonymous terms for creativity: imagination, ingenuity, innovation, inspiration, inventiveness, muse, novelty, originality, serendipity, talent and unique". The term 'imitation' is simply antipodal to many of these terms.

In the following sections, we firstly describe and discuss some at-

- Turing-style comparison tests are inappropriate for testing aspects of creative intelligence in software
  - See paper for other arguments

Boden's
"A Turing Test for Artistic Creativity"

In [11], Boden discusses the Turing Test and artistic creativity. She provides an interpretation of the Turing Test which is specifically designed for computer art systems:

> "I will take it that for an 'artistic' program to pass the TT would be for it to produce artwork which was:
>
> 1. indistinguishable from one produced by a human being; and/or
> 2. was seen as having as much aesthetic value as one produced by a human being." [11, p. 409]

467–491. Basic Books, 1993. Epilogue in [25].

[11] M. A. Boden. The Turing test and artistic creativity. *Kybernetes*, 39(3):409–413, 2010.

[12] P. McCorduck. *AARON's Code: Meta-Art, Artificial Intelligence, and*

Boden and Edmund's

"Turing Test for Artistic Creativity"

Boden describes several systems which produce art or music, which she considers to be either non-interactive or unpredictably interactive (such as a piece of art which responds to audience members or participants in ways they do not understand). She discusses comparisons with both mediocre human art, in this case pastiches of given styles (perhaps comparable to work by an art student exploring a given style), as well as examples which match world class human art, of interest as an artwork in itself (comparable to work done by a practising artist). She argues that the following systems all pass (her version of) the TT:

- Richard Brown's Starfish
- Harold Cohen's AARON
- Art by Boden and Edmunds
- David Cope's EMI

Boden's

"Turing Test for Artistic Creativity"

In particular, Boden argues that "If being exhibited alongside Rothko, in a 'diamond jubilee' celebration of these famous artists, does not count as passing the Turing Test, then I do not know what would." [11, p. 410].

Our Objections...

- It's an interpretation of Turing's test which bears little resemblance to the original idea
  - There is no dialogue or interaction of any kind with the system as part of the test
  - The test can be passed without comparison to human intelligence, or even human output
- So, it's possible to pass the originally conceived Turing test (testably achieving human-level intelligence), yet not pass Boden's test
  - Yet - as evidenced by the Starfish and by Boden and Edmunds' art - it's possible to pass Boden's test without exhibiting any higher level cognitive functions

The Starfish...

4. Our issues with software not being human

Human Level Creativity

- There are four main reasons to study human creativity with Computational Creativity research:

✓ • Because the artefacts produced are for human consumption (product)

✓ • It's ultimately human recognition of our software being creative that we seek (process)

✓ • People take computationally created artefacts and are creative with them (interpretation, etc)

✖ • We can program software to be (perceived to be) more creative by understanding well human creative processes

Psychology Envy

- Psychology research is great! And we should rightly be envious of their achievements
  - But its value to Computational Creativity research is not as high as you might imagine, and it's not a problem to realise this
- Problem 1: the results are often too vague to be turned into computational approaches, although we might get some general motivation or ideas
- Problem 2: because of the volume of knowledge about human creativity, it's too tempting to apply it to computational approaches to creativity
  - Software isn't human, e.g., it doesn't store and process information in the same way
  - It's often wholly inappropriate to analyse software in psychological terms, or pretend that software can be compared to people in meaningful ways
- Problem 3: there are some psychology experiments with methodologies that involve vaguely defined concepts. Our envy might lead us to overlook this and copy their flawed methodologies

Humanity Envy

- Treating our software as human blinds us to the fact that people's perception of software in society is hugely different to their perception of other people

- Me: "Should my software explain how it has produced a poem?"
- Tony: "...but people don't do that when they write poems!"
- Me: "But my software isn't a person"
- Tony: "It'll be seen as juvenile"
- Me: "But my software is juvenile!"...
- Me later: "Hmmm. What does juvenile even mean when we are talking about software?"

I believe that, even if they have exactly the same words in exactly the same order, a computer generated poem should be seen as a *fundamentally different* type of artefact to a poem penned by a person

This is because the effect that each poem has on a reader is fundamentally different, due to the differences in how they were produced and what/who produced them

What do People Know about Human Poets?

They're like us in many ways

They think hard about their poems

They're motivated to tell their story

They want to express their emotions

They stay up all night writing and drinking coffee

They live in Paris and sell poems for meals

They became a poet after a bad break up

They're writing about their dead father

What do People Know about Human Poets?

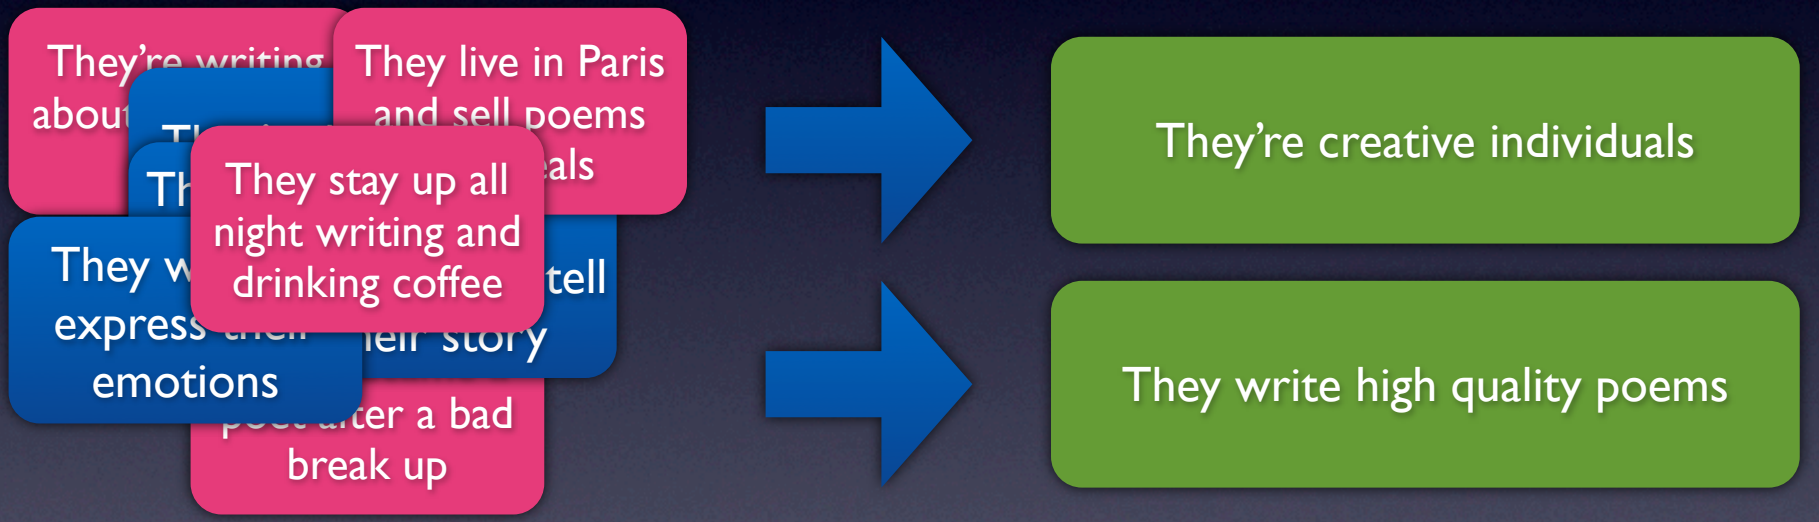

What do People Know about ~~Human~~ Poets?

Computer

- They're dislike us in sooo many ways
- They perform trivial random processing
- They have no story to tell
- They are devoid of personality

PACIFIC★STANDARD

Rise of the Robot Artist

Everyone knows computers are good at processing information and humans are good at creativity. What if it's the other way around?

The Painting Fool is an artist and a piece of software written by the British computer scientist Simon Colton. Having learned the "visual grammar" of the human body, the Fool generated this scene. (IMAGE: COURTESY OF SIMON COLTON)

—audience members were quizzed to see if they could tell the difference. Most couldn't. In later, more advanced experiments, Cope tried crossing, say, the style of Joplin with the style of Chopin. And in the culmination of his artificial intelligence work, he designed a piece of software—with the distinctly unrobotic name Emily Howell—that composes contemporary classical music in a style all its own.

The day is not far off, Cope says, when everyone will be able to carry a personal composer in their pocket, one that can generate new songs in the style of those they already have on their playlist—or in new styles altogether. In fact, "that's almost trivial to imagine," Cope says. "I'm surprised it doesn't already exist."

The main hurdles for computational creativity are not technical but psychological. "Every time a machine can do something, people turn around and say, 'Well, that's not really intelligent,'" says Michael Spranger, an artificial intelligence researcher for Sony Computer Science Laboratories in Tokyo. Colton calls this prejudice against software "silicon bias"—and he has seen convincing evidence that it's pervasive. "If you show people two rows of paintings, get their feedback, and then tell them the second row was painted by a mass murderer, those rankings fall precipitously," he said. "The same thing can happen when you tell them it's painted by a computer."

As the proud father of the Painting Fool—a piece of software with no known copy, and hence a kind of individuality—Colton takes silicon bias more than a little personally. And while he is no doubt himself biased in favor of his own creations, the talents of the Painting Fool do sometimes make him forget it is a set of logical commands atop a knot of silicon and metal. "Every morning when I wake up," he says, "I have to repeat to myself ten times: Software is not human, software is not human, software is not human…."

Some Difficult Notions to Digest...

There's no such thing as creativity

We shouldn't agree on how people perceive creativity

Psychology envy can be a bad thing

Levelling the playing field can go badly wrong

Output quality and autonomy of software
can be inversely proportional

We still don't all agree in Computational Creativity!

It's not all Bad!

- In this afternoon's lecture...
  - Filling the humanity gap
  - Managing people's perception of creativity with the creativity tripod (now spider!)
  - Formalising progress in terms of the processes that software undertakes

Was there anything right about your talk..?