Lecture 1: Characterising Computational Creativity

Geraint A. Wiggins
Professor of Computational Creativity
CCLab
Queen Mary University of London

Overview of lectures

Queen Mary University of London

- Lecture 1: Computational Creativity
  - What is it?
  - How can we do it?
  - How can we study it?
  - A Framework for Studying Creativity
  - Examples (from musical CC)

Overview of lectures

Queen Mary
University of London

- Lecture 2 (double): Cognitive Modelling of Musical Creativity
  - How can we begin to study music(al creativity) in an objective way?
  - The case for music as a psychological construct
  - The idea of Cognitive Modelling
  - Statistical Models, which come in (at least) two flavours
  - Implicit musical learning and a statistical model thereof
  - How cognitive modelling can contribute to music analysis
  - How Shannon information theory can apply to a cognitive model
  - How to add evidence for the correctness of a model
  - What does it mean to evaluate "creativity"?
  - What else should we evaluate?
  - How do we evaluate it?

Overview of lectures

Queen Mary
University of London

- Lecture 3: Creativity in the Global Workspace
  - A general cognitive architecture that may account for creative thought

Example: Automated Composition

Queen Mary
University of London

- Computer composition was first suggested by Ada, Lady Lovelace
- First recorded attempt:
  - Illiac Suite for string quartet (Hiller & Isaacson, 1957)
    - stochastic, rule-based generation
    - not very successful, musically (but still impressive)
- Many subsequent attempts
  - often concerned with style replication (Bach...)
  - often concerned with genre replication (jazz...)
  - rarely (almost never) evaluated scientifically

Some computational creativity

Queen Mary
University of London

- A notable success in automated composition is the work of Kemal Ebcioğlu (1980, etc.)
- Ebcioğlu's system CHORAL is capable of harmonising a given chorale theme according to some 350-odd rules and constraints which, it is claimed, capture the style of J S Bach

1. Chorale 48 (Bach)

2. Chorale 48 (CHORAL)

An AI model of creativity

Queen Mary
University of London

- Margaret Boden was the first artificial intelligence (AI) researcher to approach creativity seriously
  - in *Artificial Intelligence and Natural Man*, Boden, 1977
- Her 1990 book, *The Creative Mind*, outlines a broad characterisation of creative behaviour
- However, the characterisation is rather vague, since the discussion is more philosophical than scientific
- The aim here is to cast Boden's characterisation in more precise terms

The conceptual space

- Creative activity is cast as the discovery of *concepts* in a *conceptual space*
- The conceptual space contains all the possible concepts available to the creative agent
- The space is defined/constrained by rules
- *Exploratory creativity* is defined as the action of searching the conceptual space for a new concept
- This is an *abstraction* - no strong claim that it works this way in minds/ brains

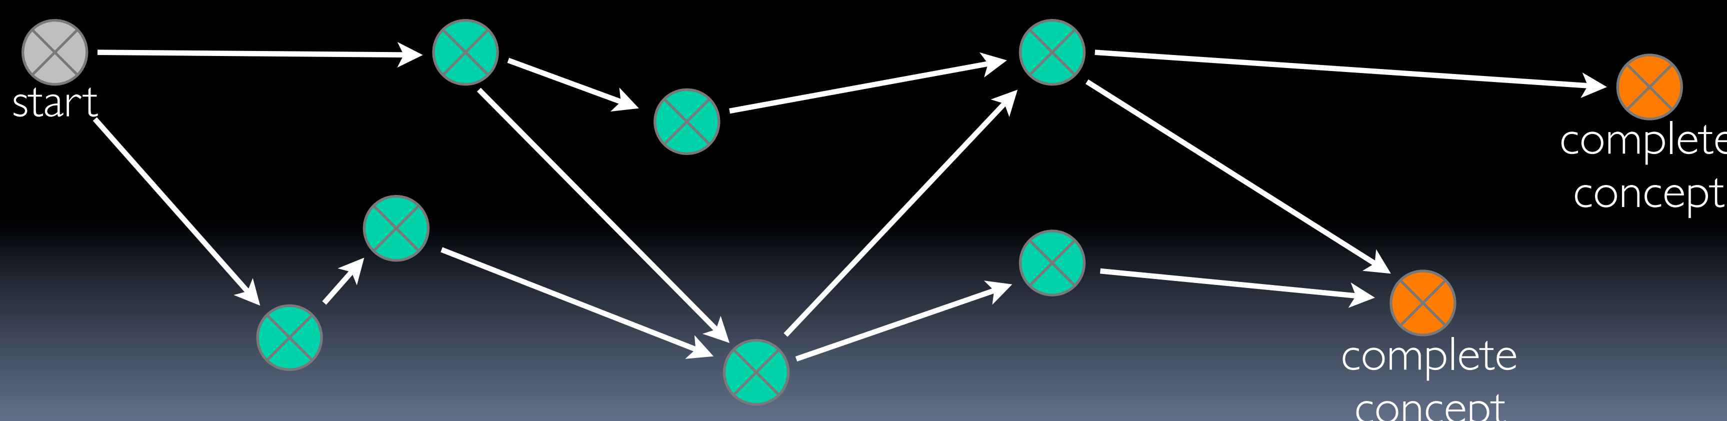

Transformational creativity

Queen Mary
University of London

- An alternative kind of Boden creativity is *transformational creativity*
- This is where the rules defining the conceptual space are changed so as to create a different (but presumably related) space
- Boden suggests that transformational creativity is more significant than exploratory creativity, because it is in a sense “bigger thinking”
- Bundy (1998) and Wiggins (2006b) argue against this, as an overly simple definition

Reasons why not

Queen Mary University of London

- “A symbolic system cannot create new concepts”
  - weighted semantic networks allow us freely to define new concepts in terms of old ones
  - conceptual blending allows us to create new semantic structures directly
  - geometrical representations of meaning allow arbitrary interpolation between concepts (e.g., Gärdenfors, 2000)
    - though we do need to think carefully about what the resulting representations mean!!

Reasons why not

- "A system which is exploring a search space defined by a representation is not being creative"
  - not necessarily true: it depends on the expressive power of the representation
  - creating an artefact by *explicit mechanistic inference* doesn't make doing so any less creative
  - cognitively speaking, creative insight does not "feel" like enumeration
    - but such introspection is misleading

Reasons why not

Queen Mary
University of London

- “Non-symbolic systems generalise via a simple mathematical process, which is not creative”
  - There is no evidence that the human mind does not create in this way
  - There are suggestions (e.g., Kanerva’s sparse distributed memory) that this is exactly how the human mind creates
  - Anyway, interpolation and generalisation may be a perfectly good model of creativity

Formalising Boden's model

Queen Mary
University of London

- Let us represent the conceptual space as a multidimensional (possibly metric) space
- Partial and complete concepts are represented as points in the space
- Each dimension of the space represents a feature of the domain
- (So each point denotes a set of property/value pairs)

Defining a conceptual space

- Suppose now that we have a set of rules, **R**, which defines a conceptual space, **C**
- The existence of transformational creativity implies that there must be a larger set, **U**, containing **C**
- So **R** is a set of rules which picks the elements of **C** from **U**
- $\mathbf{C} \subset \mathbf{U}$

Defining a conceptual space

- In order to give our rules, **R**, we need a language, **L**, and an interpreter for it
- Let $[\![.]\!]$ be an interpreter which maps its argument (a set of rules in **L**) to an effective procedure for selecting elements of **U**
- $\mathbf{C} = [\![\mathbf{R}]\!](\mathbf{U})$
- We also need a null concept, $\top$

Exploring a conceptual space

Queen Mary
University of London

- Let us also allow another set of rules, **T**, describing our creative agent's method for exploring **C**
- One more ingredient of Boden's model remains: it is necessary to be able to choose the better concepts from the less good ones
- We introduce a set of rules, **E**, written in **L**, which may be used to accept or reject concepts in terms of their quality
- We will need a more complex interpreter, 《.,.,.》, which, given three sets of rules in **L**, will return an effective procedure for computing an ordered set of (partial) concepts, $\mathbf{c}_{out}$, from another, $\mathbf{c}_{in}$

$$\mathbf{c}_{out} = 《\mathbf{R},\mathbf{T},\mathbf{E}》(\mathbf{c}_{in})$$

Exploring a conceptual space

- It will be useful to add the operator $^{\diamond}$ which will allow us to compute the set defined by repeated applications of a function

$$F^{\diamond}(X) = \bigcup_{n=0,\infty} F^{n}(X)$$

- We can now define the enumeration of the conceptual space, **C**, by our creative agent:

$$\mathbf{e_C} = \langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\})$$

Exploring a conceptual space

Queen Mary University of London

- Note that $\mathbf{e}_{\mathbf{C}}$ may be a subset of **C**
- This is because a creative agent's exploratory technique, as captured by **T**, need not be strong enough to discover all the concepts which are actually admissible under **R**
- Or $\mathbf{e}_{\mathbf{C}}$ may intersect **C**, producing some acceptable and some unacceptable concepts

An exploratory creative system

Queen Mary
University of London

- We are now able to describe an exploratory creative system with the following septuplet:

$$\langle \mathbf{U}, \mathbf{L}, [\![.]\!], \langle\!\langle .,.,. \rangle\!\rangle, \mathbf{R}, \mathbf{T}, \mathbf{E} \rangle$$

U	The universe of all concepts
L	A language for expressing rules and concepts
$[\![.]\!]$	A testing interpreter (for **R**)
$\langle\!\langle .,.,. \rangle\!\rangle$	An enumerating interpreter (for **R**, **T** and **E**)
R	A set of rules defining a conceptual space, **C**, in **U**
T	A set of rules allowing traversal of **U** (around **C**)
E	A set of rules evaluating concepts found using $\langle\!\langle .,.,. \rangle\!\rangle$

Transformational creativity

Queen Mary University of London

- Boden describes *transformational creativity* as changing the rules, **R**, which define the conceptual space
- In our formulation, there are two sets of rules which can be transformed
- Transforming **R** is transforming what is allowed as the output of the creativity process
- Transforming **T** is transforming the creative agent's personal method

Transformational creativity

Queen Mary
University of London

- There is a search space of rule sets, which is itself a conceptual space
- That search space is the power set of the language, **L**: **L***
- So **L*** is now the universe in which we are searching
- We can describe **L** (and **L***) with a metalanguage $\mathbf{L}_\mathbf{L}$

Transformational creativity

- To capture the exploration of the rule space, we need some constraints on what is syntactically well-formed, $\mathbf{R_L}$
- We also need to define the search strategy, $\mathbf{T_L}$
- If we use the metalanguage $\mathbf{L_L}$ as before for these specifications, we can use the same interpreters as before, ⟦.⟧ and 《.,.,.》

Transformational creativity

- The only thing outstanding is the evaluation of the transformation, which can be done with a set of rules, $\mathbf{E_L}$

- We now have another *exploratory* septuple:

$$\langle\ \mathbf{L^*}, \mathbf{L_L}, [\![.]\!], \langle\!\langle .,.,. \rangle\!\rangle, \mathbf{R_L}, \mathbf{T_L}, \mathbf{E_L}\ \rangle$$

- So transformational creativity is exploratory creativity at the meta-level of conceptual spaces

- $\mathbf{E_L}$ may be characterised in terms of $\mathbf{E}$ (see Wiggins, 2006a, for how)

On failing to create...

Queen Mary University of London

- We are now in a position to examine the behaviour of creative systems
- The different components of the descriptions interact, and how they interact can tell us useful information
- Now, we discuss ways in which a system can fail to create
- Therefore, a creative system can introspect about how to improve itself

Uninspiration

Queen Mary University of London

- *Uninspiration* is the inability to produce valued outputs
- There are three kinds of uninspiration:
  - Hopeless
  - Conceptual
  - Generative
- It is useful to know about uninspiration, because it can act as
  - a "well-formedness" check
  - a trigger to transform a creative system in one way or another

Hopeless Uninspiration

Queen Mary University of London

- The simplest case of uninspiration is where there are no valued concepts in the universe:

$$[\![\mathbf{E}]\!](\mathbf{U}) = \varnothing$$

- This means that no creative agent in this universe can ever produce anything valued
- It is a property which we should attempt to disprove of any creative system, *a priori*

Conceptual Uninspiration

- *Conceptual uninspiration* is where there are no valued concepts in a given conceptual space:

$$[\![\mathbf{E}]\!](\mathbf{C}) = [\![\mathbf{E}]\!]([\![\mathbf{R}]\!](\mathbf{U})) = \emptyset$$

- This means that no creative agent exploring this conceptual space can ever produce anything valued
- It is a property which we should attempt to disprove of any exploratory-creative system, *a priori*
- Conceptual uninspiration can be used as a cue to encourage *aberrant* behaviour

Generative Uninspiration

Queen Mary
University of London

- *Generative uninspiration* is where a creative agent's technique, **T**, causes it to miss the valued members of the conceptual space:

$$[\![\mathbf{E}]\!](\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\})) = \emptyset$$

- This means that the agent will never produce anything valued
- It is a property which we should attempt to disprove of any exploratory-creative system, *a priori*
- It can act as a trigger for transformation of **T** (or **R**)

Aberration

Queen Mary
University of London

- *Aberration* is the production of new concepts which are not in the existing conceptual space (that is, deviation from the expected)
- There are three kinds of aberration:
  - Perfect
  - Productive
  - Pointless

Aberration

- Aberration happens when a creative agent finds concepts which are valued, but which are not in the conceptual space
- This is why value (**E**) needs to be represented distinctly from acceptability (**R**)
- In the CSF, this means that

$$\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\}) \setminus [\![\mathbf{R}]\!](\mathbf{U}) \neq \varnothing$$

Perfect Aberration

- Perfect aberration is the case where

$$\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\}) \setminus [\![\mathbf{R}]\!](\mathbf{U}) = [\![\mathbf{E}]\!](\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\}) \setminus [\![\mathbf{R}]\!](\mathbf{U}))$$

  that is, where all the aberrant concepts are valued

- This, in most cases, will be a cue to transform **R** so that it includes the new concepts

Productive Aberration

- Productive aberration is the case when

$$[\![\mathbf{E}]\!](\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\}) \setminus [\![\mathbf{R}]\!](\mathbf{U})) \neq \varnothing$$

  that is, where some aberrant concepts are valued

- This, in many cases, may be a cue to transform **R** or **T** or both

Pointless Aberration

- Pointless aberration is characterised by

$$[\![\mathbf{E}]\!](\langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle^{\diamond}(\{\top\}) \setminus [\![\mathbf{R}]\!](\mathbf{U})) = \varnothing$$

  that is, where no aberrant concepts are valued

- This is a cue to transform **T** but not **R**

Reflection and transformational creativity

- These ideas pave the way towards creative agents which can reason about their own performance, in terms of both value and productivity
- In particular, these analyses, which were not possible in Boden's original framework, allow a system which is essentially exploratory to cue occasional transformational behaviour
- Is this what artists/musicians/scientists do when they (eg) consciously change style?
- Just because we can use the CSF to model creative systems, it doesn't mean that all creative systems have to work by search
- We can usefully conceptualise/model a process as a search mechanism in the abstract even if that is not how it actually works

An important question

Queen Mary
University of London

- What is the difference between Good Old-Fashioned AI Search and Computational Creativity based on the Boden/Wiggins model?

GOFAI Search

- Given an agenda **S** (a sequence of states):
  1. If **head**(**S**) is a solution, stop.
  2. Remove **head**(**S**) from **S** giving remainder **S'**
  3. **expand**(**head**(**S**)) giving **S"**
  4. **merge**(**S"**,**S'**) giving (new) **S**
  5. Repeat from 1
- For Depth-First Search, **merge** = **prepend**
- For Breadth-First Search, **merge** = **append**
- For Best-First Search, Hill-climbing, A, A*, **merge** = **append+sort**

GOFAI Search

Queen Mary
University of London

- Key Features:
  - Representation: can represent all and only output configurations of problem (closed world)
  - Solution detector: Boolean test for (a representation of) a solution
  - Heuristics allow control of search for best one(s)
    - calculate "quality" of solutions
    - calculate "distance" from nearest solution
    - combination of these

Similarities

Queen Mary
University of London

- GOFAI search vs. CSF
  - Representation syntax ≃ Rules of **R**
  - Search space ≃ Conceptual space
  - Algorithmic framework ≃ Algorithmic framework
  - Heuristics ≃ Traversal (**T**) and/or Value (**E**) Rules
  - Agenda (**S**) ≃ Current expansion of space ($\mathbf{c}_{\mathbf{in}}$)

Differences

Queen Mary University of London

- Representation: closed vs. open world (**C** vs **U**)
  - admits "discovery" of solutions not envisaged by system designer
- Algorithmic framework: single vs. multiple operands
  - admits more complex (powerful?) search algorithms, e.g., GA, blending

CSF > GOFAI Search

Queen Mary
University of London

- GOFAI search can be implemented in the CSF
- The CSF cannot be implemented as GOFAI search
  - (unless, in both cases, we disingenuously jump to a meta-level)
  - The CSF is therefore more expressive than the GOFAI search framework
  - So Boden's notion of creativity is not "just AI search"

Summary

Queen Mary
University of London

- Introduced Creative Systems Framework
  - Conceptual Space and Rule Set **R**
  - Traversal of Space to find Concepts and Rule Set **T**
  - Evaluation and Rule Set **E**
- Transformational Creativity is Exploratory Creativity at the meta-level
- The CSF is more expressive than the standard search framework of AI
- We can use the CSF to help conceptualise creative systems...
- ...and that's what we'll do in Lectures 2 and 3