Lecture 1: Characterising Computational Creativity

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

Overview of lectures

- 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

- 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

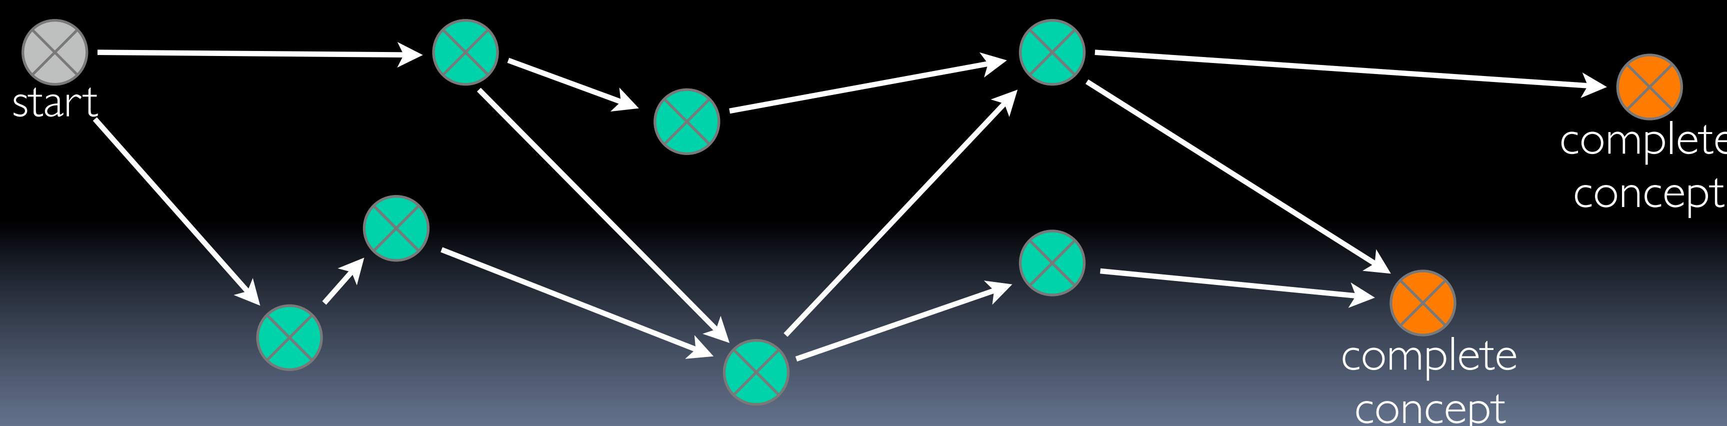

Transformational creativity

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

Queen Mary University of London

- "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, $\mathbf{c}_{in}$

$$\mathbf{c}_{out} = \langle\!\langle \mathbf{R},\mathbf{T},\mathbf{E} \rangle\!\rangle(\mathbf{c}_{in})$$

Exploring a conceptual space

Queen Mary
University of London

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

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

- 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

- Note that $\mathbf{e_C}$ 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 $\mathbf{e_C}$ may intersect **C**, producing some acceptable and some unacceptable concepts

An exploratory creative system

- 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

- 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

- 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

Queen Mary
University of London

- 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 $\langle\!\langle .,.,. \rangle\!\rangle$

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

- *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:

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

- 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

- *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

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

Perfect Aberration

Queen Mary
University of London

- 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

Queen Mary University of London

- 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

Queen Mary University of London

- 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

Queen Mary University of London

- 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 ($\mathbf{c}_{in}$)

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

- 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

Lecture 2: Modelling musical creativity

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

Overview

Queen Mary
University of London

- Aims:
  - to motivate the scientific study of musical creativity
  - to provide an approach to doing so
  - to demonstrate the approach
  - to place the model in the context of the Creative Systems Framework

A Science of Music?

Queen Mary
University of London

- Why would we want to study a science of music?
- Music is
  - an art form
  - a cultural construct
  - precious, even sacrosanct, to many people
- What can science tell us that musicology cannot?
  - Place of music in general cognition
  - Contibution of musical behaviour to human developement
    - individual
    - cultural

A Science of Music?

Queen Mary
University of London

- Music seems to be uniquely a human faculty
  - There is no known human culture without music
  - Music is everywhere in every human culture
  - Music is irresistible to the majority of people
  - No other species has been shown to exhibit musical behaviour in the sense that humans do
- Yet, no known bio-evolutionary advantage is given by music
  - This needs to be explained, and not just wondered at!
- What is more, musical behaviour is a fundamental part of being human
  - So we need to understand music if we want to understand ourselves

Milton Babbitt's representational domains

- Milton Babbitt (1965) proposed three *domains* of music representation

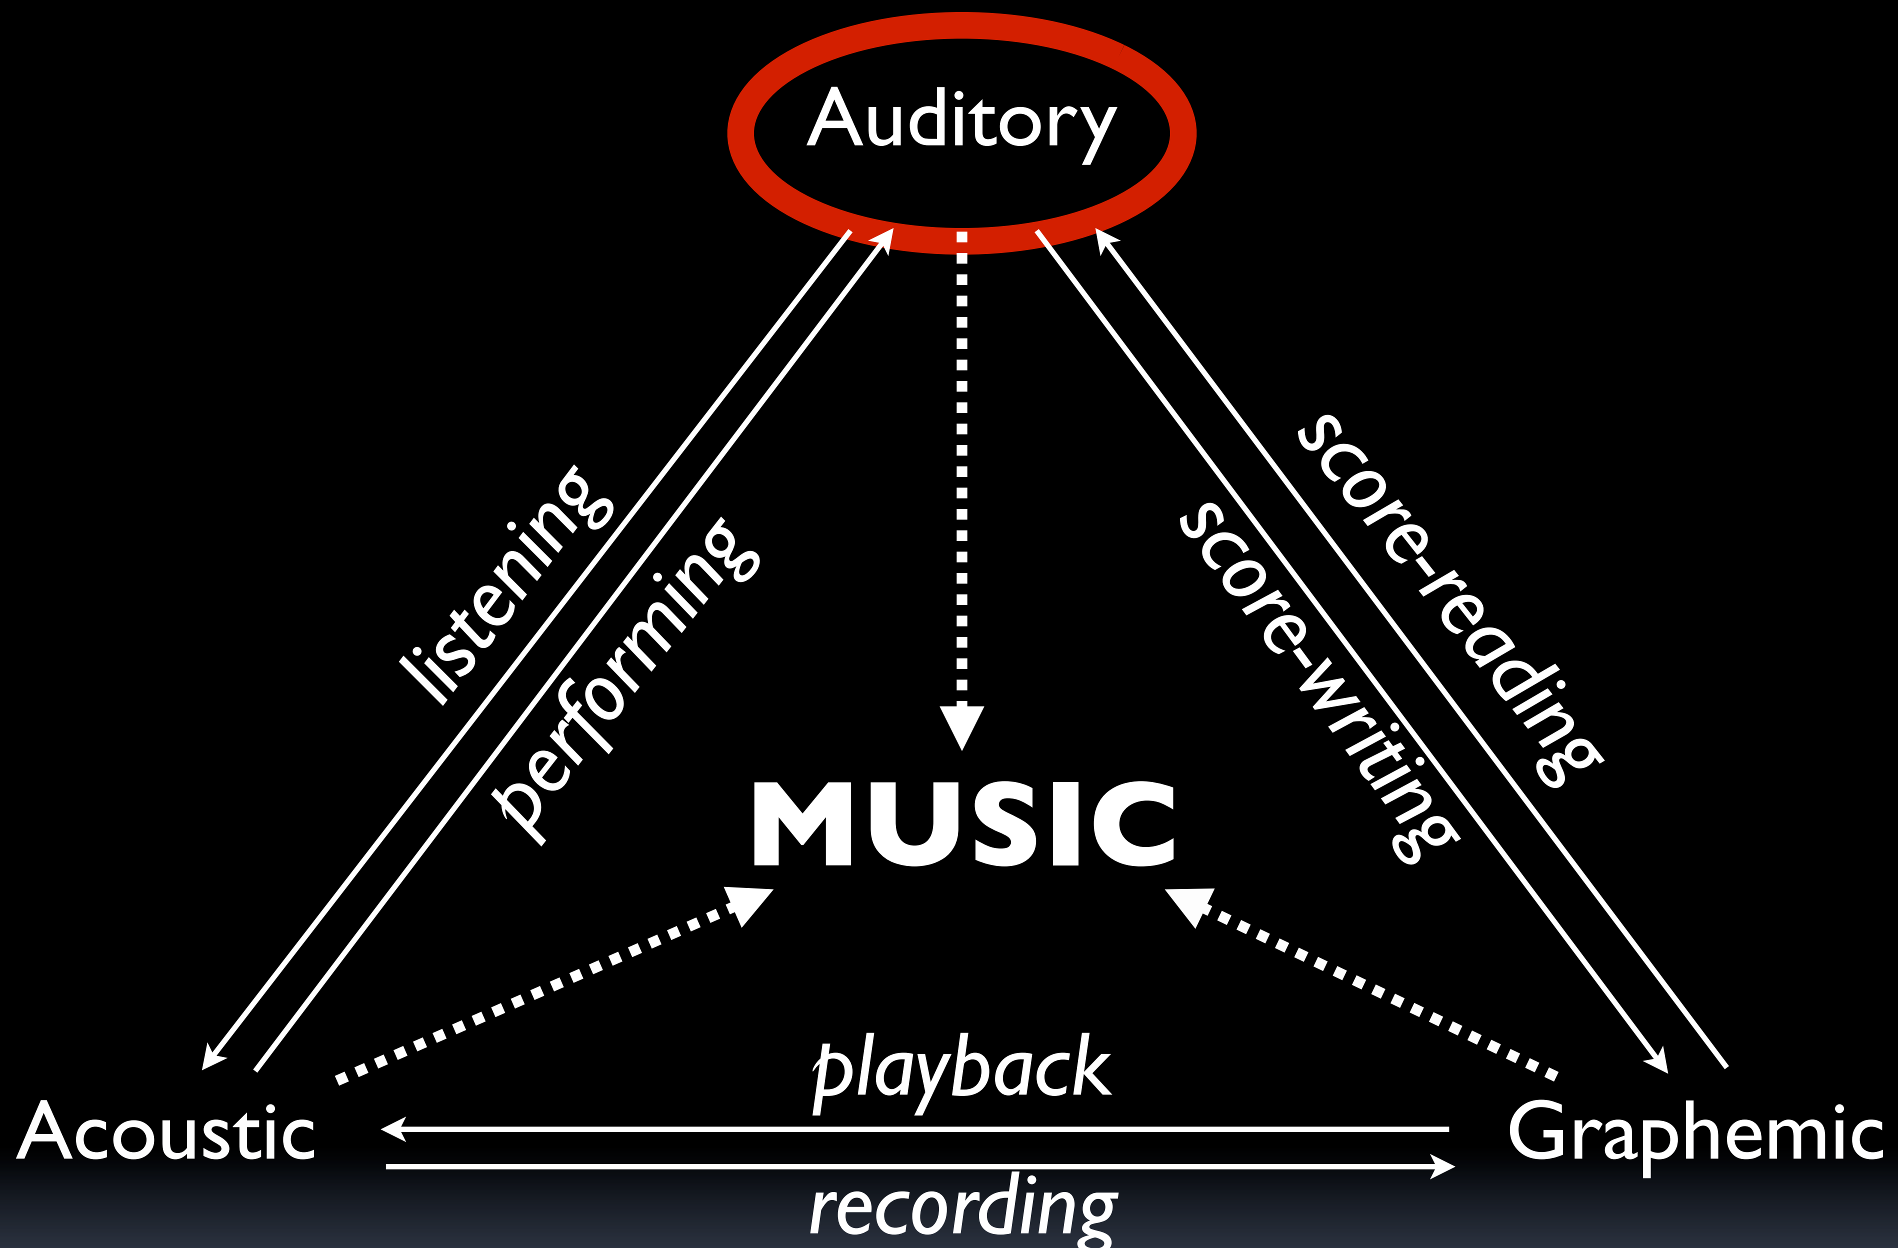

Music and meaning - a bonus for CC

Queen Mary University of London

- Music is all about meaning
  - associations
  - emotional responses
- but it (usually) has no semantic content
  - so a musical creative system has a massively reduced framing problem
  - the frame is syntactic, not semantic
  - ie it’s about style, intrinsic to music, not about an extrinsic world model

Feeling the beat

Queen Mary
University of London

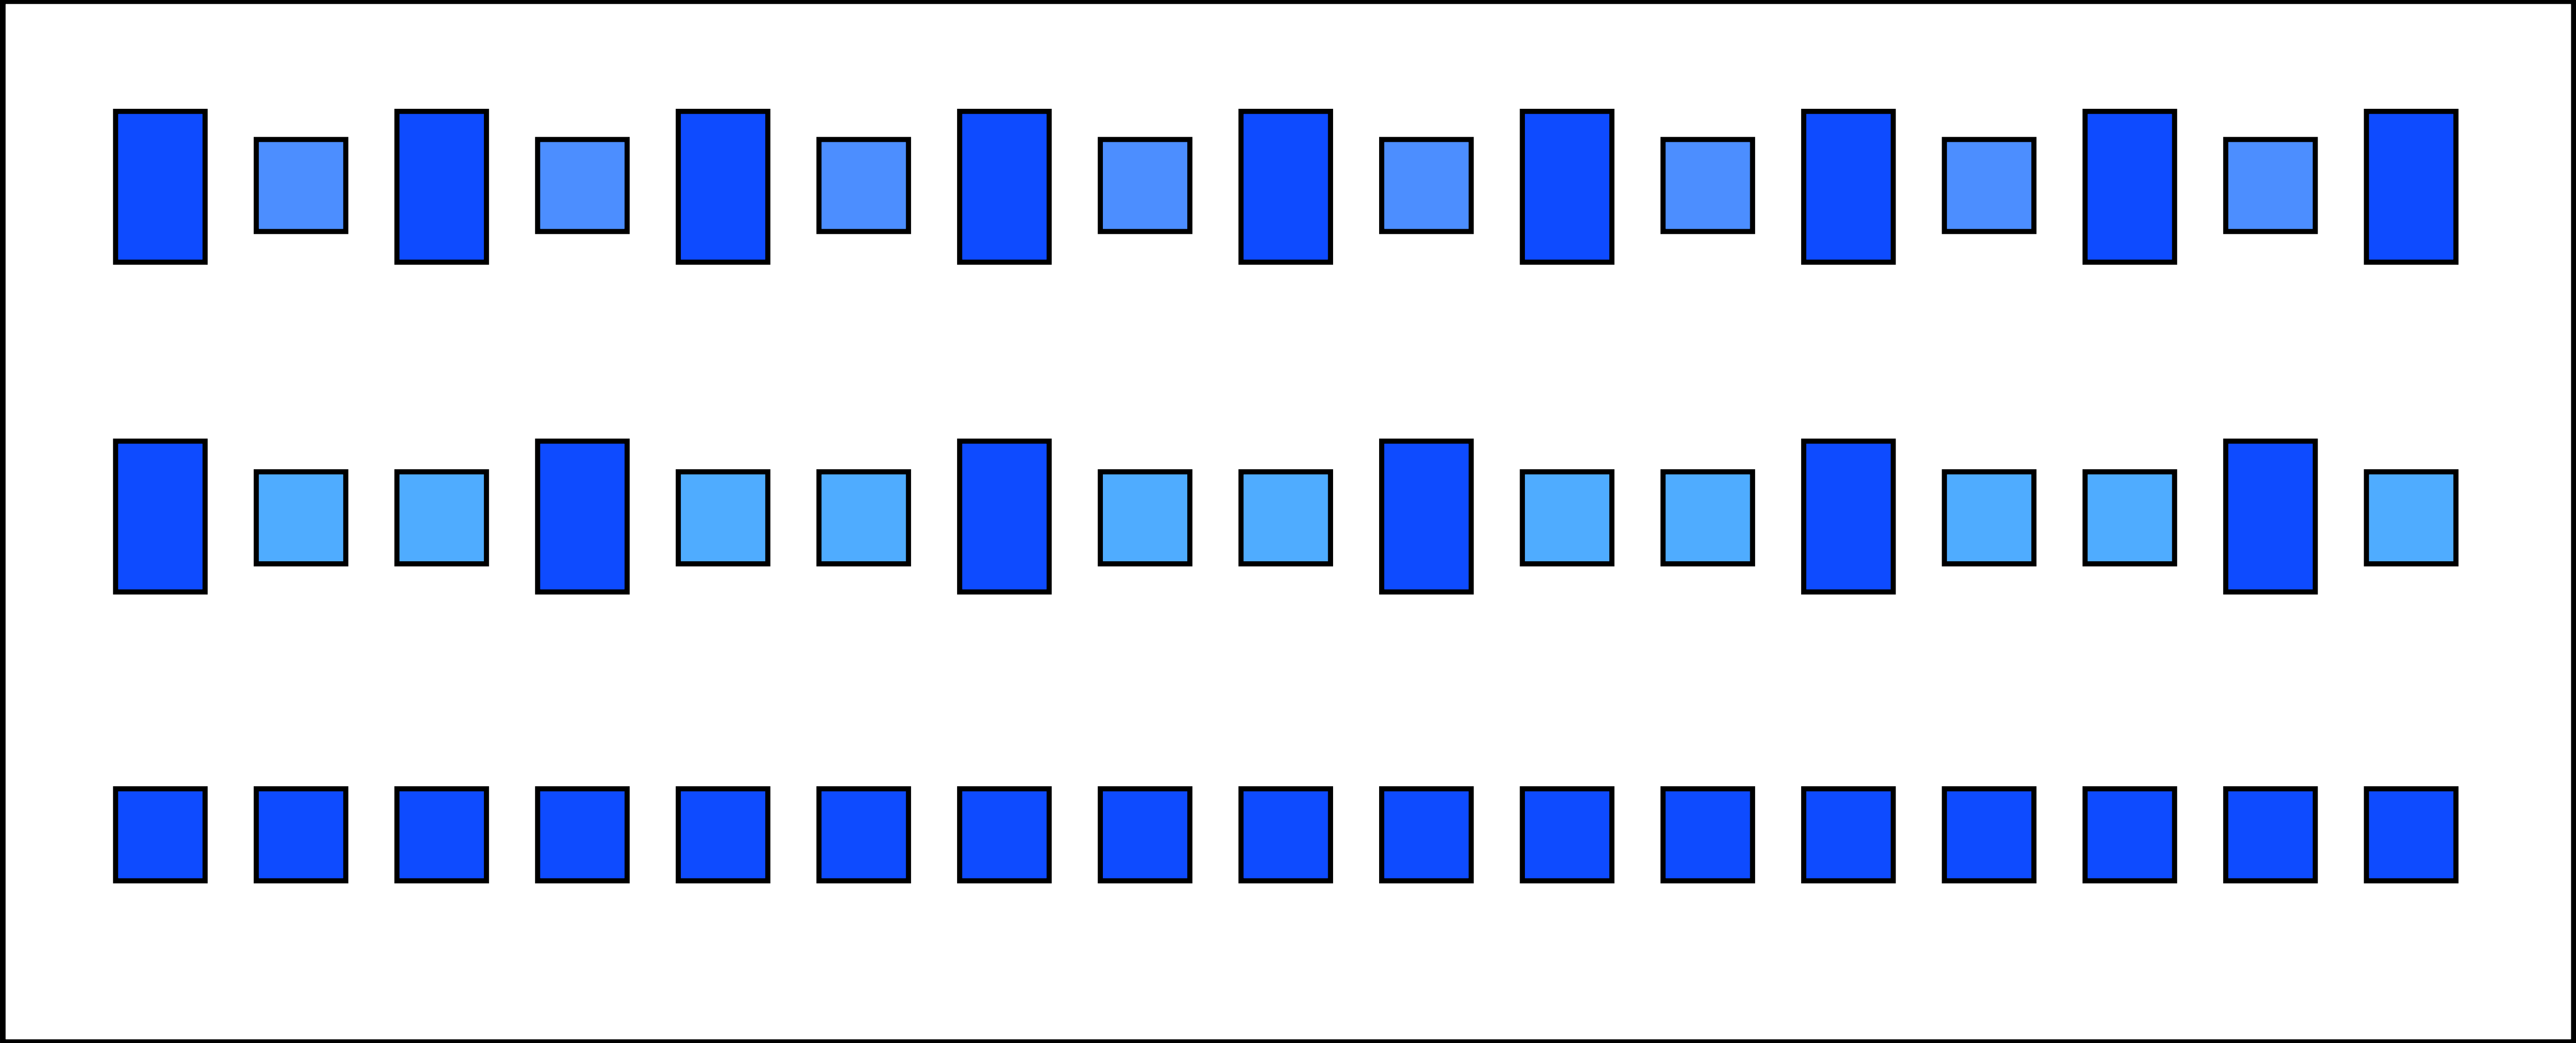

Feeling the beat

Queen Mary University of London

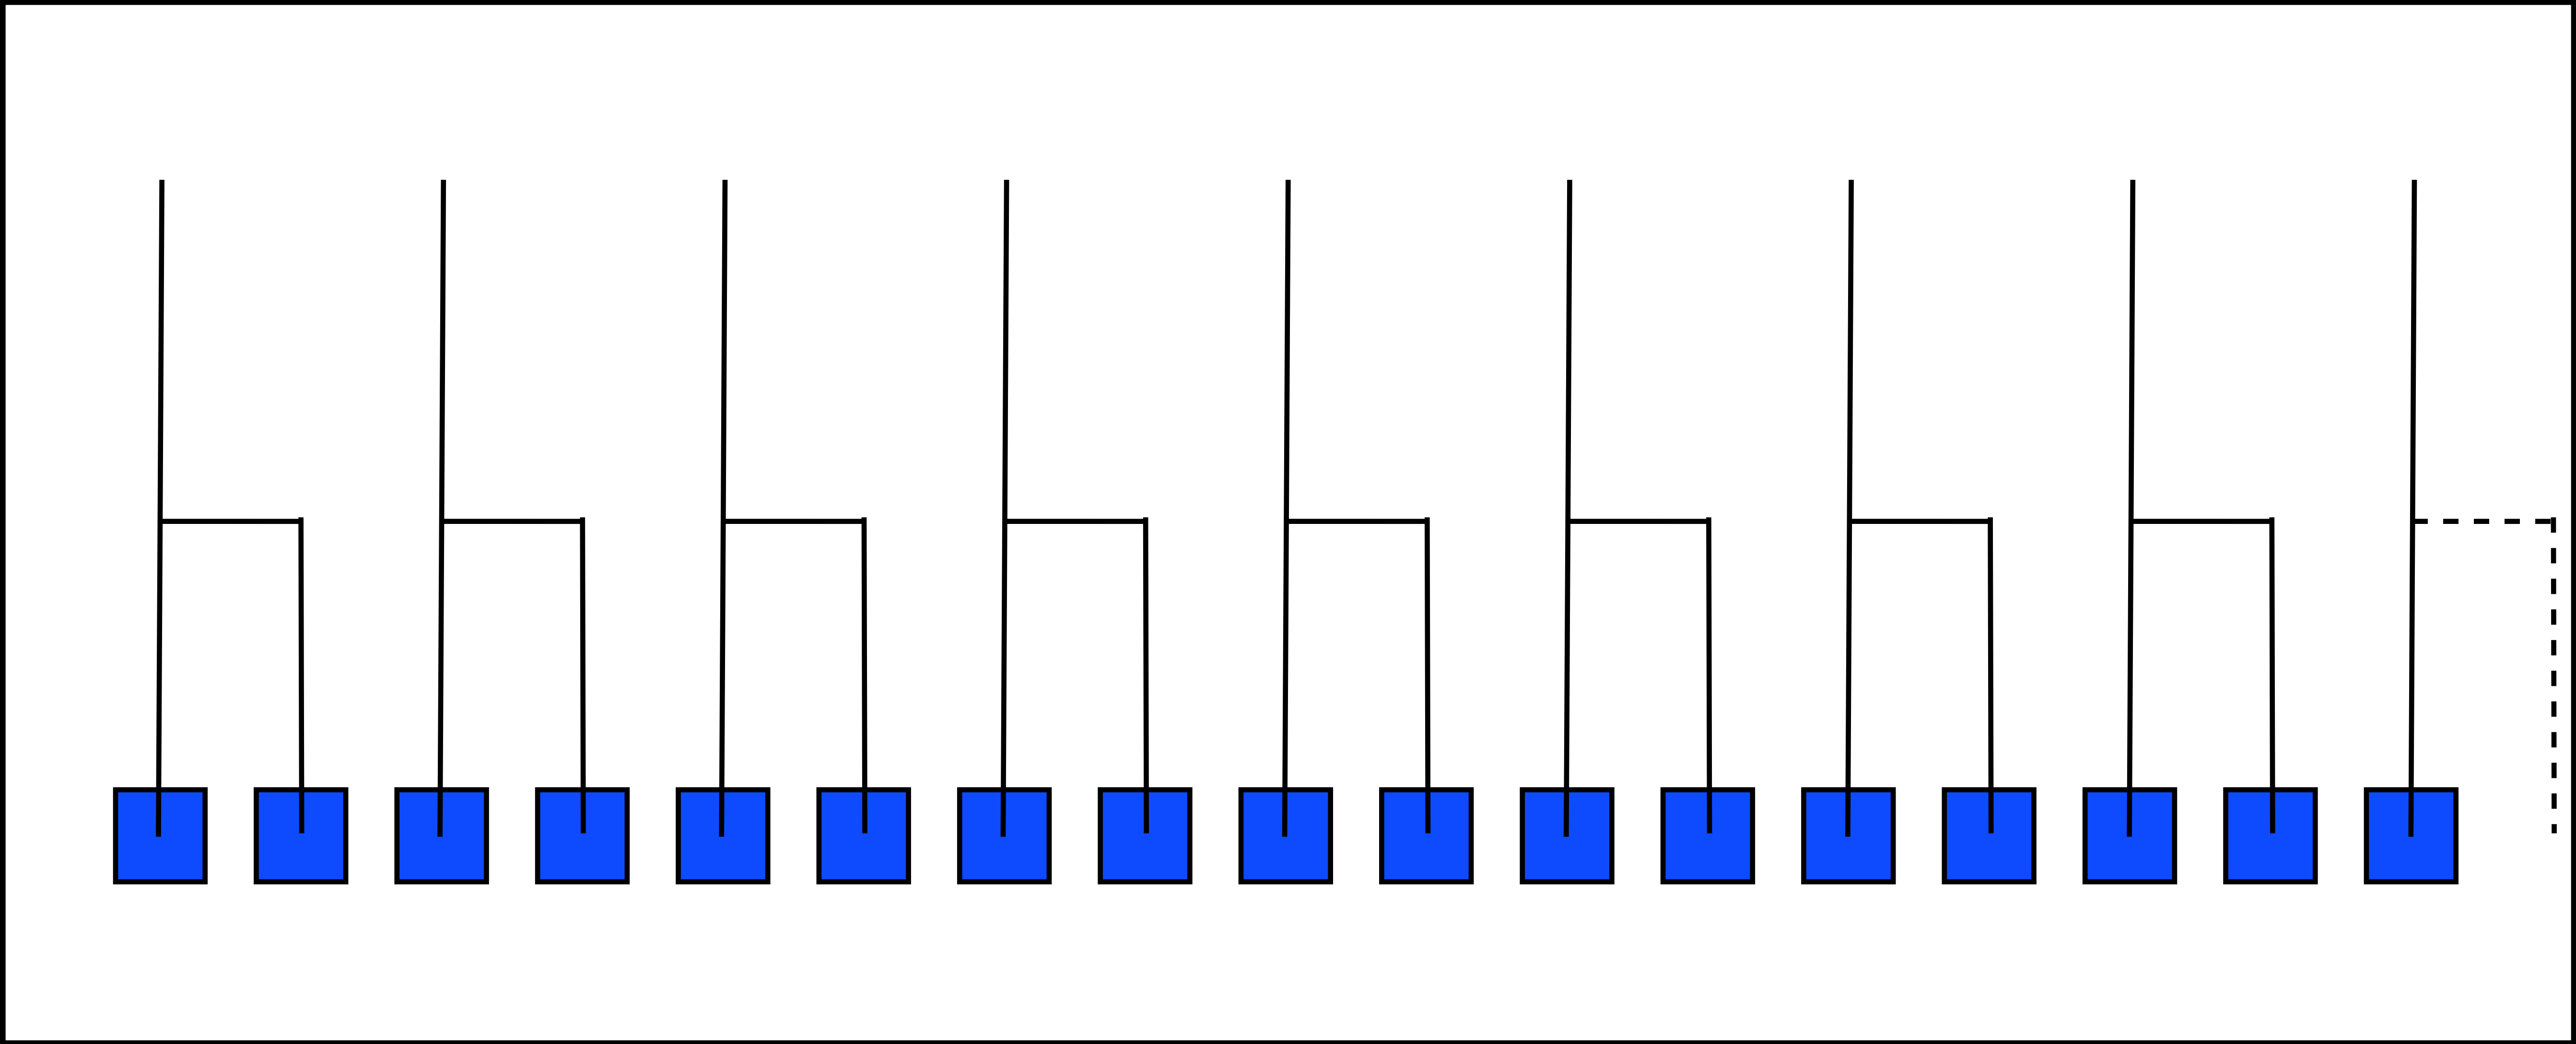

Feeling the beat

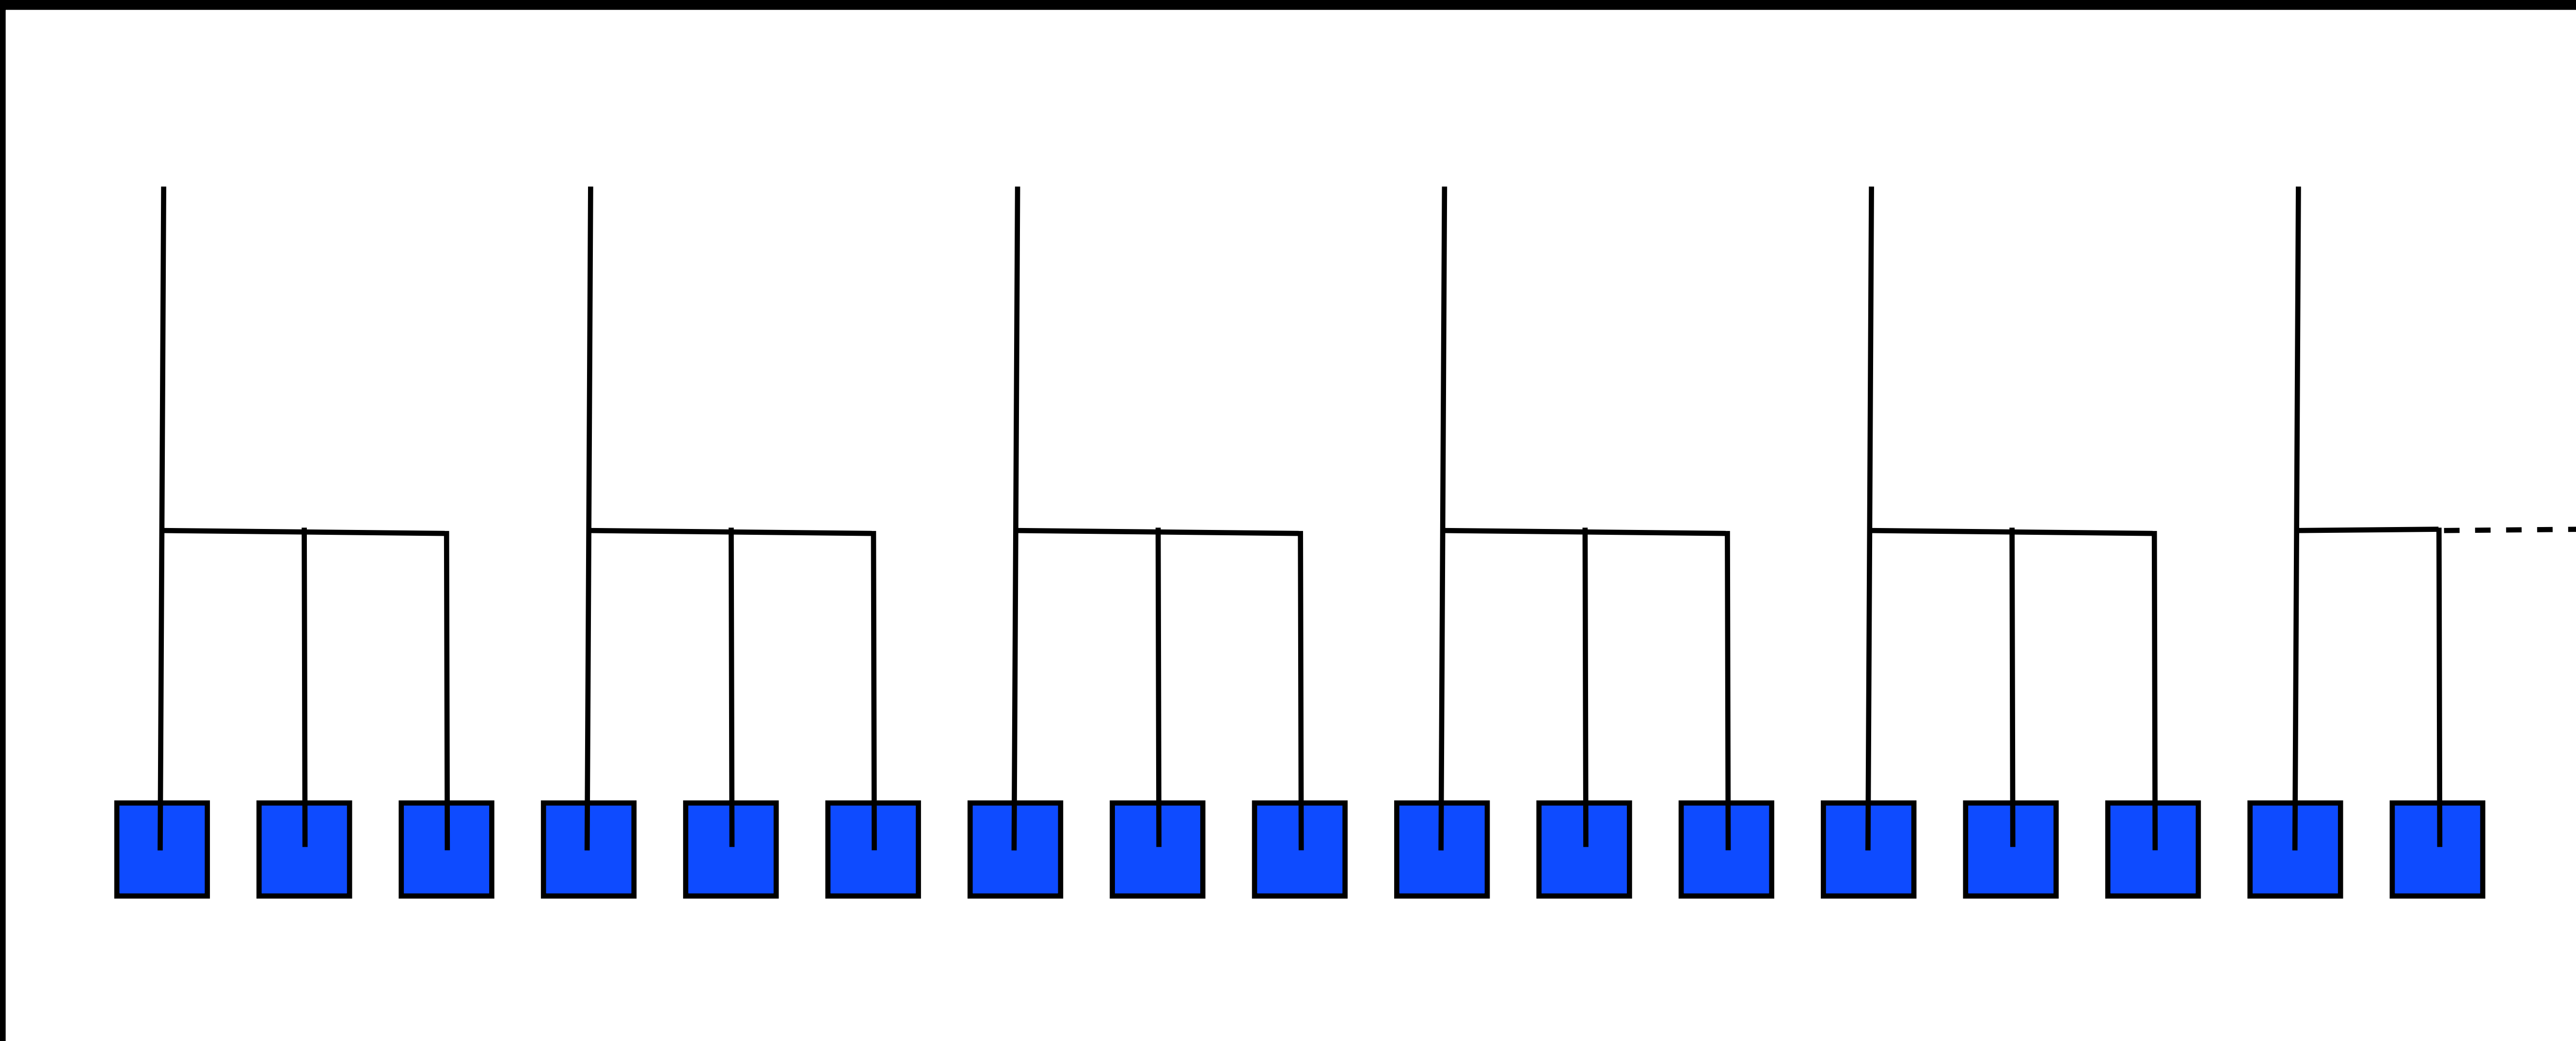

Feeling the beat

- The pulse train demonstration shows that human listeners tend to hear rhythmic structure in sound...

  ...even when it isn't there

- When we know it isn't there, we can manipulate our own perception, to hear either twos or threes

Why is there grouping?

Queen Mary
University of London

- Our senses are subject to a constant barrage of information
- One way to help deal with this information is to break it into more easily manageable lumps
- Broadly speaking, this is what is happening in grouping phenomena (and others)
- “Chunking” is a general phenomenon at all levels throughout cognition

Grouping for music?

- We make smaller and larger scale groupings too
  - Small scale: harmonics *fuse* together to make notes
  - Medium scale: notes fuse together to make chords
  - Larger scale: chords work in sequence to make progressions
- And some authors define music as "Organised Sound"

Grouping for music?

- But why?
- It seems unlikely that we evolved the ability to group sounds because music was useful in evolutionary terms
  - Music does not help find food
  - Music does not protect against predators

Grouping for music?

Queen Mary
University of London

- But the things that contribute to music perception do contribute to evolutionary fitness in other ways:
  - perceiving complex sound as tone colour
    Is that the sound of a lion (lute)?
  - location of sound sources
    Where is that tiger (trombone)?
  - grouping of separated rhythmic sounds
    Do I hear footsteps (fandango)?
  - perceiving expressed emotion in sound
    Is the other hominid pleased to see me?

Entrainment

- Entrainment (in music) is the ability to perceive and duplicate a sequence of events in real time
  - Chimpanzees and other primates are not capable of entrainment
  - Until recently, it was thought that only humans could entrain, but...
- Entrainment is fundamental to musical behaviour
  - Musicians need to anticipate and reciprocate rhythm
- What is more, humans really really like to entrain:
  - tapping a foot to a beat
  - clapping along with a song
  - walking together in step

Hypothesis on entrainment

- If entrainment and the associated affect was an early development (in human history),...

  ...then pairs and groups of humans would have enjoyed entraining together

- In turn, this would be likely to increase social bonding...

  ...which is an evolutionary advantage for weak organisms like hominids...

  ...and which would in turn reinforce the genetic basis of the entraining behaviour

What is computational modelling of cognition?

Queen Mary
University of London

- It is difficult to study minds
  - you can’t see them
  - you can’t stick electrodes in them
  - their relationship with brains is almost completely unclear
  - it is unethical to distort/deform them for testing purposes
  - *etc*
- Before the advent of computers, psychologists had two means of study:
  - look at what happened when things went wrong
  - make predictions from theory about what would happen in certain precise circumstances (hypotheses), and test them (experiments)
- This is very time-consuming (decades, not hours), error-prone, and (in the first case) dependent on chance

What is computational modelling of cognition?

Queen Mary
University of London

- With computers, however, new things become possible
- We can write computer programs which embody theories and then test them to destruction (ethically!)
- We can also make predictions by computer which can then be tested in experiments with humans
- This can be much faster than the human-driven approach
- It is more objective than the human-driven approach (so long as the program is written objectively)

What's the point?

- This is really the only (ethical) way to understand how a cognitive phenomenon actually works
  - duplicate it in an artificial system and test that to destruction
  - if it matches human behaviour in all circumstances, it is a good model
  - it's important to choose and stick to your level of abstraction
- If you can write a program which embodies your theory, then your theory is fully worked through (a Very Good Thing)

How do we build a cognitive model?

Queen Mary University of London

- Apply reductionist science!
  - accept that most phenomena are too complex to understand all at once
  - identify part(s) of the phenomenon that are (as) separable (as possible)
  - be careful to use stimuli (music) that do not go beyond these boundaries
  - remember that the resulting model is probably an oversimplification
  - when you have understood the parts of the phenomenon, put them together, study the interactions between them, and test them in concert
- This is quite different from, and antagonistic to, the holistic view usually taken in the humanities, but it is not incompatible
- Human (musical) behaviour *must* be at the start and end of this process:
  - theories behind the models come from observation of musical behaviour
  - results from models are tested against musical behaviour

What are the limitations of cognitive modelling?

Queen Mary
University of London

- A model is only as good as
  - the theory it embodies
  - the computational implementation
  - the input data
  - the input and output data representation
- We must always question and test (and re-test) results because of these potential sources of error

What are the limitations of cognitive modelling?

Queen Mary
University of London

- We can only take one small step at a time
  - this science is in its infancy: we must not rush ahead and make mistakes
- Therefore, we have to be satisfied with small, focused, isolated results
  - we look at how a given aspect of something changes, given that everything else stays the same – an artificial situation
- The results are only ever approximations
  - we continue to refine models as our understanding improves

What are the requirements of a cognitive model?

Queen Mary University of London

- We must be careful to make the right *abstraction* of our data
  - A representation based on a 12-note octave will not be able to model phenomena related to microtonal music
  - A representation based on a 12-note octave will not be able to model phenomena related to conventional tonal tuning (eg playing into the key)
- And that leads to thinking carefully about the abstraction of the model
  - (NB cf: define universe or define conceptual space?)
- A very good abstraction of Western Common Practice music: the score
  - models categorical pitch and time perception (and tonality if need be)
  - evolved over about 1,000 years to do this *well*
  - not good for everything (eg no means of representing instrumental timbre)
  - but very good at quite a lot!
- Many cognitive models of music use (an equivalent of) score notation

Two kinds of cognitive model

Queen Mary
University of London

- Some models are **descriptive** (Wiggins, 2007, 2011)
  - they say what happens when stimuli are applied in each circumstance
  - the predict results in terms only of the application of rules
  - these rules may be complicated
  - these models do not explain WHY a cognitive effect is the way it is
  - they do explain WHAT the cognitive effect is, at the same level of abstraction as the representation they use
- Some models are **explanatory** (Wiggins, 2007, 2011)
  - they give a **general underlying mechanism** by which a phenomenon occurs
  - they predict results using this mechanism
  - they explain WHY a cognitive effect is the way it is (possibly at some level of abstraction different from the representation)

Example 1: GTTM

Queen Mary University of London

- Generative Theory of Tonal Music (Lerdahl & Jackendoff, 1983)
  - "complete" theory of tonal music (actually not – still being updated)
  - has 4 components, each being a set of rules, written in English
    - grouping
    - metre
    - time-span reduction
    - prolongation
  - within each, there are two kinds of rule
    - fixed rules
    - preference rules
  - "preference" rules without conditions for application mean that GTTM is not a computerisable theory
  - therefore, it is not a rigorously objective model
  - it is only a descriptive model, because no mechanism is given

Towards an explanatory model

Queen Mary
University of London

- Choice of domain
  - We have to apply reductionist methodology (creativity, or even just music, is too complicated to model in one go)
  - We should aim for ecological validity - so stimuli are as natural as possible
  - We should aim to reduce "with the grain" of the domain
- Choice of model
  - Programmed, rule-based - probably not explanatory
  - Learning based - possibly explanatory, depending on how it works
    - Needs to be unsupervised
  - Therefore we need an underlying theory motivating how the model works

Learning-based models of music

Queen Mary
University of London

- There is plenty of evidence (all around us) that the ability to behave musically is universal(ly valued)
- However, different (sub-)cultures have different musics
- Music from other (sub-)cultures is often incomprehensible
  - but it's based mostly on the same constructs
    - rhythm/meter
    - notes
    - patterns
    - repetition
    - timbre (tone colour)

Learning-based models of music

Queen Mary
University of London

- All this suggests a collection of evolved perceptual mechanisms which combined to create music cognition
- There is no reason at all why these things need to be otherwise connected
- Then, we hypothesise, the musical experience is derived from the processing of percepts at this level by some general mechanism
- However, the musical culture is learned *implicitly*
- That is to say that a good model will not require explicit training
    - in other words, we don't tell it what the outputs we expect are

A general perceptual computation

Queen Mary University of London

- Problem: information is everywhere
- There is too much information for even a brain to process
- Possible solution: compress it as it arrives
- We hypothesise that brains use *structural compression* to help manage the continuous information overload
- So, for example, a chair is perceived as a chair, and not as a set comprising a seat, four legs and a back
- Then, when we see many chairs, it is more efficient to represent them all as references to a definition of chairness, rather than as a detailed description of each one individually

Evolutionary perspective

- Key evolutionary points
  - organisms survive better if they can learn
  - organisms survive better if they can anticipate
  - organisms survive better if they can anticipate from what they learn
  - organisms cannot be merely reactive
    - anticipation must be proactive – it must result in embodied action
  - organisms must regulate cognitive resource – attention is expensive
- Need to evolve a mechanism where organisms can learn without damage
  - gives rise to "tension": emotional warning of uncertainty in the organism's world model
  - cf. musical tension; narrative tension

Managing information in time

Queen Mary
University of London

- The same principle applies to things arranged in time as to things arranged in space
- Music is often called "a time-based art" because it has no instantaneous existence (except *perhaps* in the minds of some highly skilled musicans)
- And there are many other situations where the ability to react to sequences of events in time at a perceptual level is evolutionarily useful
- The key property that admits this is...

EXPECTATION

- Expectation allows us to deal with the world
  - there is too much data out there to process in real time
  - we need to manage it by predicting what comes next, so we have a chance to get ahead
- Expectation works in many domains
  - vision
  - movement
  - speech
  - music

EXPECTATION

- There must be a mechanism that
  - learns from data
  - predicts from data
  - generalises from data (so it can deal with data it hasn't seen before)
- It will also be able to enlist cognitive resources
  - so that unexpected things can be dealt with

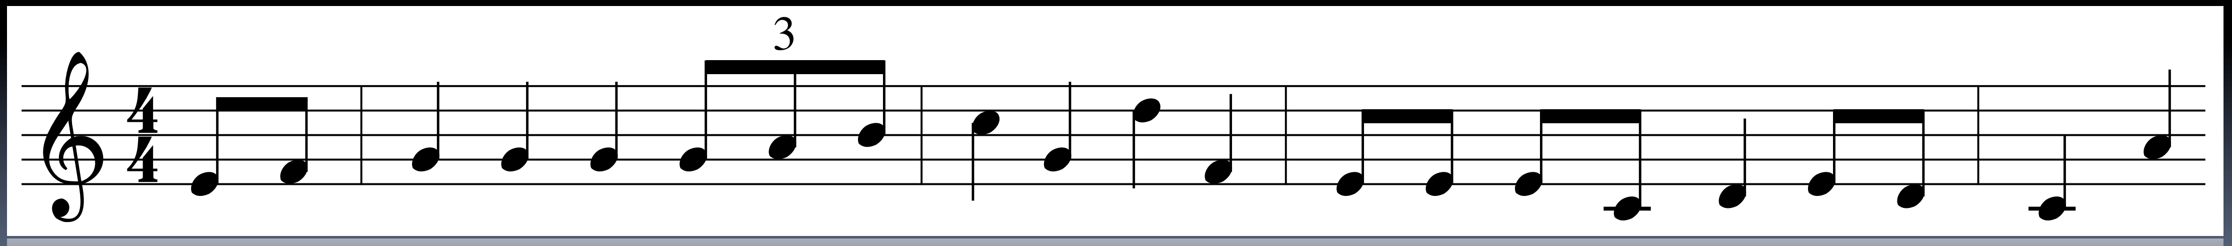

Managing information in time

Queen Mary
University of London

EXPECTATION

- In speech and language understanding
  - **it's easy to wreck a nice beach**

Managing information in time

Queen Mary
University of London

- Knowing things can reduce the amount of information required to transmit information
- Claude Shannon (1948) proposed a mathematical model of this idea, which he called "information theory", in the context of telecommunications engineering
- It turns out that Shannon information theory works very well in a model of human perception based on (advanced) Markov Models

Skip Markov demo

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

a b c b d e	* | a	a | b	b | c	c | b	b | d	d | e
a b c a b e	* | a	a | b	b | c	c | a	a | b	b | e
a b d b d e	* | a	a | b	b | d	d | b	b | d	d | e

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

<table>
<tr><td>a b c b d e</td><td>* | a</td><td>a | b</td><td>b | c</td><td>c | b</td><td>b | d</td><td>d | e</td></tr>
<tr><td>a b c a b e</td><td>* | a</td><td>a | b</td><td>b | c</td><td>c | a</td><td></td><td>b | e</td></tr>
<tr><td>a b d b d e</td><td>* | a</td><td>a | b</td><td>b | d</td><td>d | b</td><td>b | d</td><td>d | e</td></tr>
<tr><td></td><td></td><td>a | b</td><td></td><td></td><td></td><td></td></tr>
</table>

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

<table>
<tr><td>a b c b d e</td><td>* | a</td><td>a | b</td><td>b | c</td><td>c | a</td><td>d | b</td></tr>
<tr><td>a b c a b e</td><td>* | a</td><td>a | b</td><td>b | c</td><td>c | b</td><td>d | e</td></tr>
<tr><td>a b d b d e</td><td>* | a</td><td>a | b</td><td>b | d</td><td></td><td>d | e</td></tr>
<tr><td></td><td></td><td>a | b</td><td>b | d</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td>b | d</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td>b | e</td><td></td><td></td></tr>
</table>

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

1.0

a b c b d e   * | a   a | b   b | c   c | a   d | b

a b c a b e   a | b   b | c   c | b   d | e

a b d b d e   a | b   b | d   d | e

a | b   b | d

b | d

b | e

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

<table>
<tr><td></td><td>1.0</td><td>1.0</td><td></td><td></td><td></td></tr>
<tr><td>a b c b d e</td><td>* | a</td><td>a | b</td><td>b | c</td><td>c | a</td><td>d | b</td></tr>
<tr><td>a b c a b e</td><td></td><td></td><td>b | c</td><td>c | b</td><td>d | e</td></tr>
<tr><td>a b d b d e</td><td></td><td></td><td>b | d</td><td></td><td>d | e</td></tr>
<tr><td></td><td></td><td></td><td>b | d</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td>b | d</td><td></td><td></td></tr>
<tr><td></td><td></td><td></td><td>b | e</td><td></td><td></td></tr>
</table>

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

<table>
<tr><td>a b c b d e</td><td>1.0
* | a</td><td>1.0
a | b</td><td>0.333
b | c</td><td>c | a</td><td>d | b</td></tr>
<tr><td>a b c a b e</td><td></td><td></td><td>0.5
b | d</td><td>c | b</td><td>d | e</td></tr>
<tr><td>a b d b d e</td><td></td><td></td><td>0.167
b | e</td><td></td><td>d | e</td></tr>
</table>

Markov models

- A very neat way to do this is by Markov Modelling
- Each possible symbol in a stream of symbols is recorded, along with each of the contexts in which it is experienced (n-grams; n=2)

a b c b d e	$\overset{1.0}{* \mid a}$	$\overset{1.0}{a \mid b}$	$\overset{0.333}{b \mid c}$	$\overset{0.5}{c \mid a}$	$\overset{0.333}{d \mid b}$
a b c a b e			$\overset{0.5}{b \mid d}$	$\overset{0.5}{c \mid b}$	$\overset{0.667}{d \mid e}$
a b d b d e			$\overset{0.167}{b \mid e}$		

Markov models

<table>
<tr><td>1.0
* | a</td><td>1.0
a | b</td><td>0.333
b | c</td><td>0.5
c | a</td><td>0.333
d | b</td></tr>
<tr><td></td><td></td><td>0.5
b | d</td><td>0.5
c | b</td><td>0.667
d | e</td></tr>
<tr><td></td><td></td><td>0.167
b | e</td><td></td><td></td></tr>
</table>

Markov models

- Now, given a partial leftmost string, we can estimate the probability distribution of the next unseen symbol

1.0	1.0	0.333	0.5	0.333
$* \mid a$	$a \mid b$	$b \mid c$	$c \mid a$	$d \mid b$
		0.5	0.5	0.667
		$b \mid d$	$c \mid b$	$d \mid e$
		0.167		
		$b \mid e$		

Markov models

- Now, given a partial leftmost string, we can estimate the probability distribution of the next unseen symbol

*	a
* a	b
* a b	c or d or e
* a b c	a or b
* a b c b	c or d or e
* a b c b d	b or e
* a b c b d e	

1.0 * | a	1.0 a | b	0.333 b | c	0.5 c | a	0.333 d | b
		0.5 b | d	0.5 c | b	0.667 d | e
		0.167 b | e		

Markov models

- Now, given a partial leftmost string, we can estimate the probability distribution of the next unseen symbol

*	a
* a	b
* a b	c or d or e
* a b c	a or b
* a b c b	c or d or e
* a b c b d	b or e
* a b c b d e	

1.0	1.0	0.333	0.5	0.333
* | a	a | b	b | c	c | a	d | b
		0.5	0.5	0.667
		b | d	c | b	d | e
		0.167		
		b | e		

p(a b c b d e) =
1.0 x 1.0 x 0.333 x 0.5 x 0.5 x 0.667 = 0.111

p(a b c a b e) =
1.0 x 1.0 x 0.333 x 0.5 x 1.0 x 0.167 = 0.028

p(a b d b d e) =
1.0 x 1.0 x 0.5 x 0.333 x 0.5 x 0.667 = 0.022

Markov models

- Now, given a partial leftmost string, we can estimate the probability distribution of the next unseen symbol

*	a
* a	b
* a b	c or d or e
* a b c	a or b
* a b c b	c or d or e
* a b c b d	b or e
* a b c b d e	

1.0 * | a	1.0 a | b	0.333 b | c	0.5 c | a	0.333 d | b
		0.5 b | d	0.5 c | b	0.667 d | e
		0.167 b | e		

p(a b c b d e) =
1.0 x 1.0 x 0.333 x 0.5 x 0.5 x 0.667 = 0.111

p(a b c a b e) =
1.0 x 1.0 x 0.333 x 0.5 x 1.0 x 0.167 = 0.028

p(a b d b d e) =
1.0 x 1.0 x 0.5 x 0.333 x 0.5 x 0.667 = 0.022

p(a b d e) = 1.0 x 1.0 x 0.5 x 0.667 = 0.333

Ingredients for the model

Queen Mary
University of London

- We have
  - an underlying mechanism (Markov models & Shannon information theory)
  - a hypothesis to motivate it
  - a creative subdomain (musical melody)
  - a representation (essentially, the musical score)
- We need
  - data (The Essen Folksong Collection: 907 tonal folk melodies)
  - a computational implementation (Pearce, 2005)

IDyOM: unexpectedness

- Shannon (1948) defines a measure of *entropy*, which has been interpreted in several ways in the literature
- We interpret it in two ways here:
  - the entropy of a given note is a measure of its *unexpectedness in context*
  - the entropy of the distribution of an unseen note is a measure of the model's *uncertainty in context*
- *unexpectedness,* or *information content,* of a note with probability p is defined as

$$-\log_2 p$$

- Does this quantity model human listeners' estimates of their own perception of expectedness when listening?

Information Dynamics of Music

Queen Mary University of London

- Middle layer of cognitive model of conscious musical experience
- Unsupervised, implicit learning
- Inputs are sequences of basic percepts
  - notes, with pitch & time features
  - derived percepts, e.g.,
    - interval
    - tonal centre
- Outputs are
  - distributions of predicted pitches
  - information-theoretic derivatives of distributions

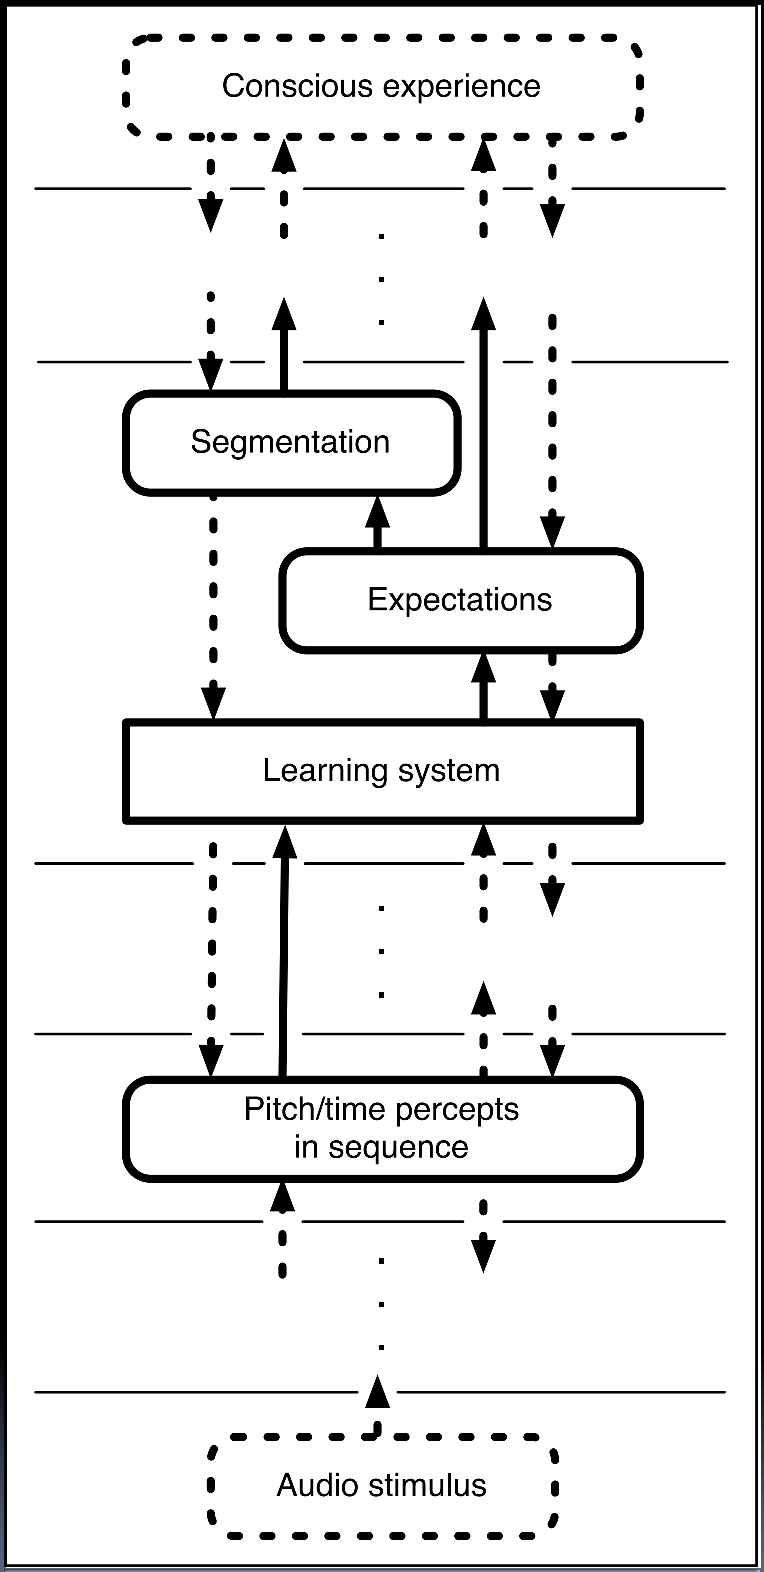

The IDyOM model

Queen Mary
University of London

- IDyOM = Information Dynamics of Music
- Model assembled and evaluated by Marcus Pearce (2005)
- Uses Markov models as a simulation of perception of events in time-sequence
- Clever implementation using suffix trees
- Implicit learning: learns the likelihood of each symbol appearing in a sequence from mere exposure, then predicts from this information
- Estimates the likelihood of unseen symbols (uniform distribution)
- Uses Shannon information theory to weight different components of the model as they contribute to a combined distribution

Information Dynamics of Music

Queen Mary
University of London

- Efficient implementation of simple Markov chains
  - but with multi-dimensional symbols
  - select feature sequences (*viewpoints*)
    - basic
    - derived
  - calculus of viewpoints
    - differentiation (delta)
    - cross-product (pairing)
    - thread (sub-sequence selection)

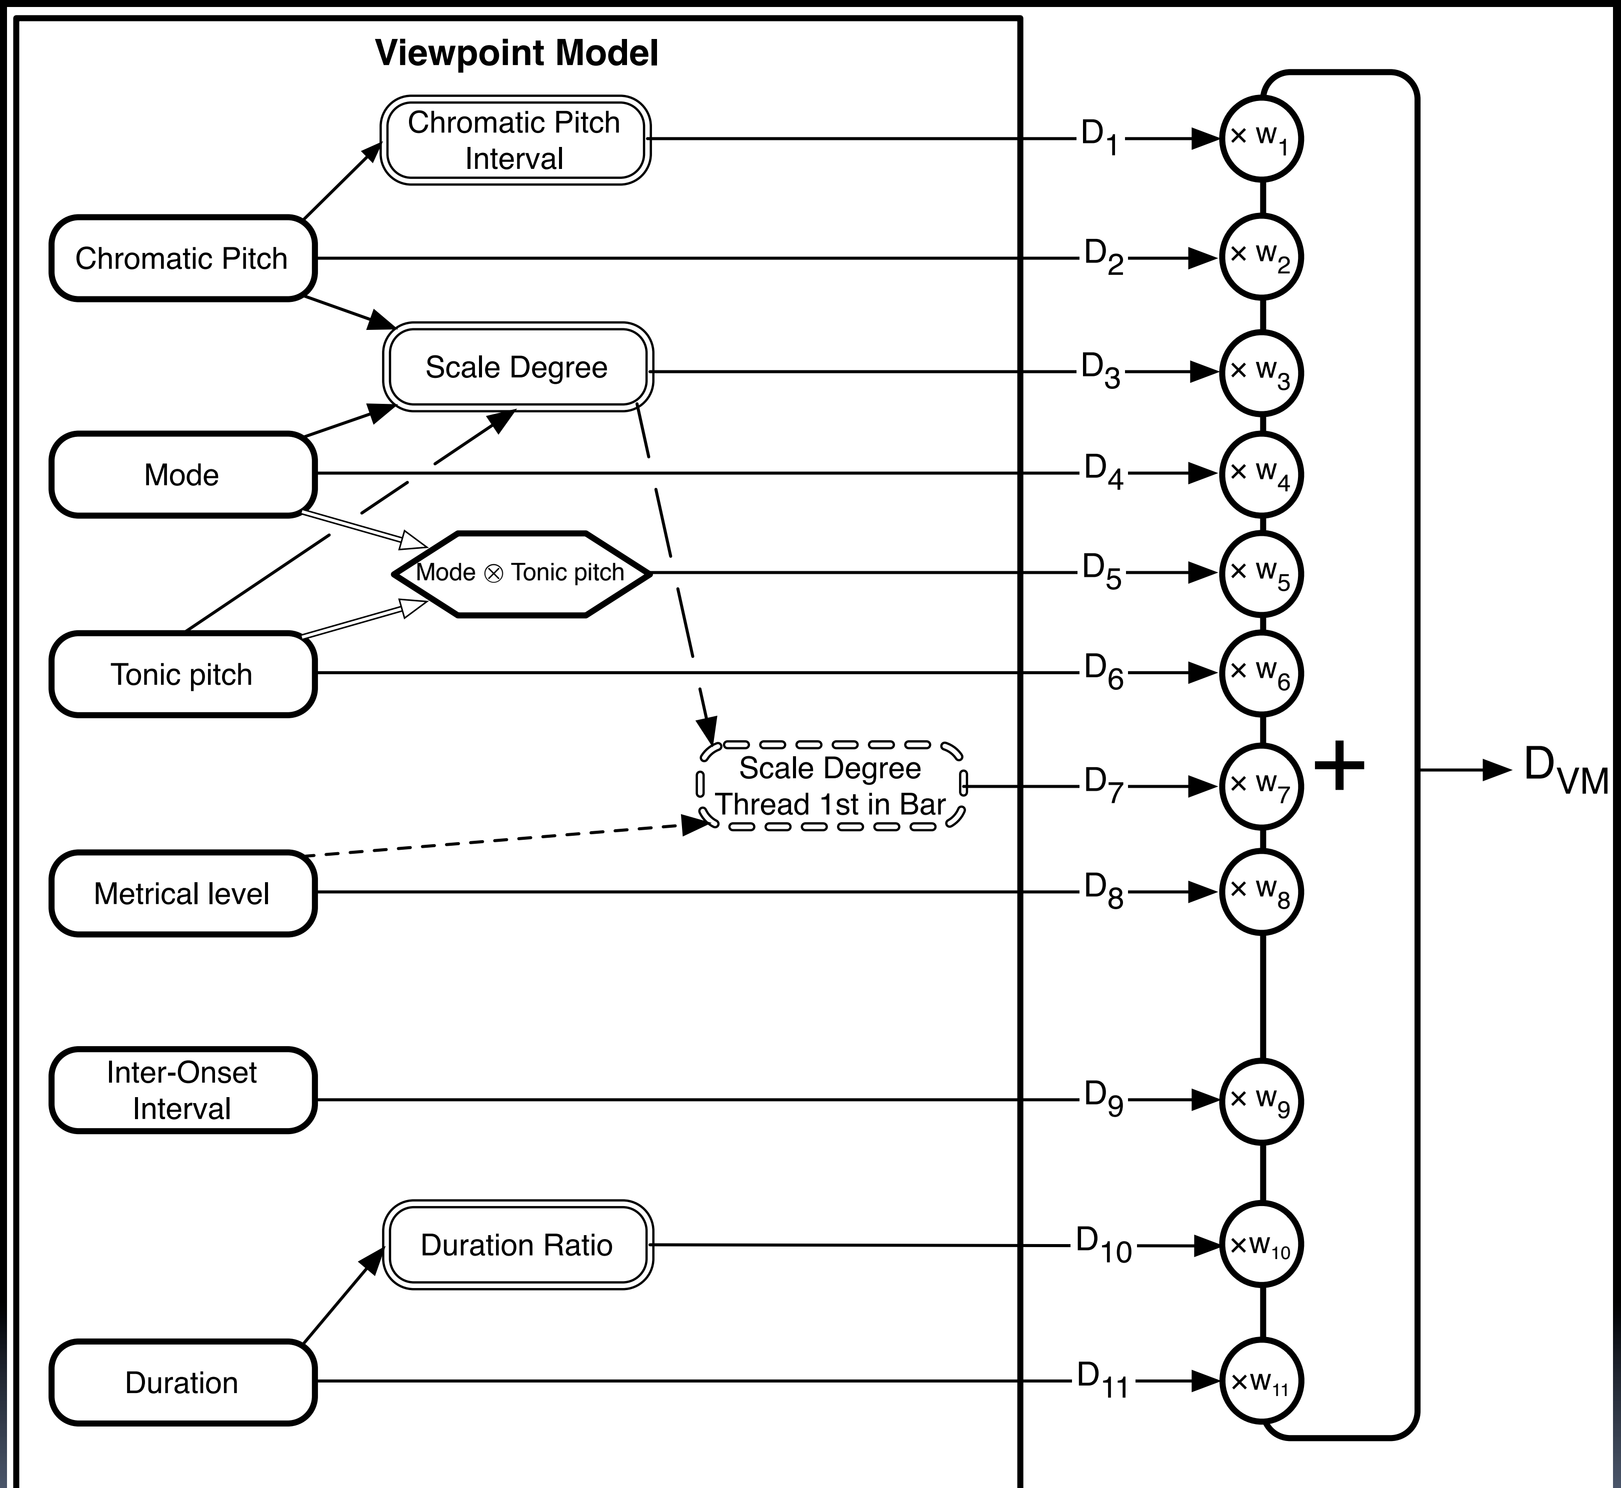

Information Dynamics of Music

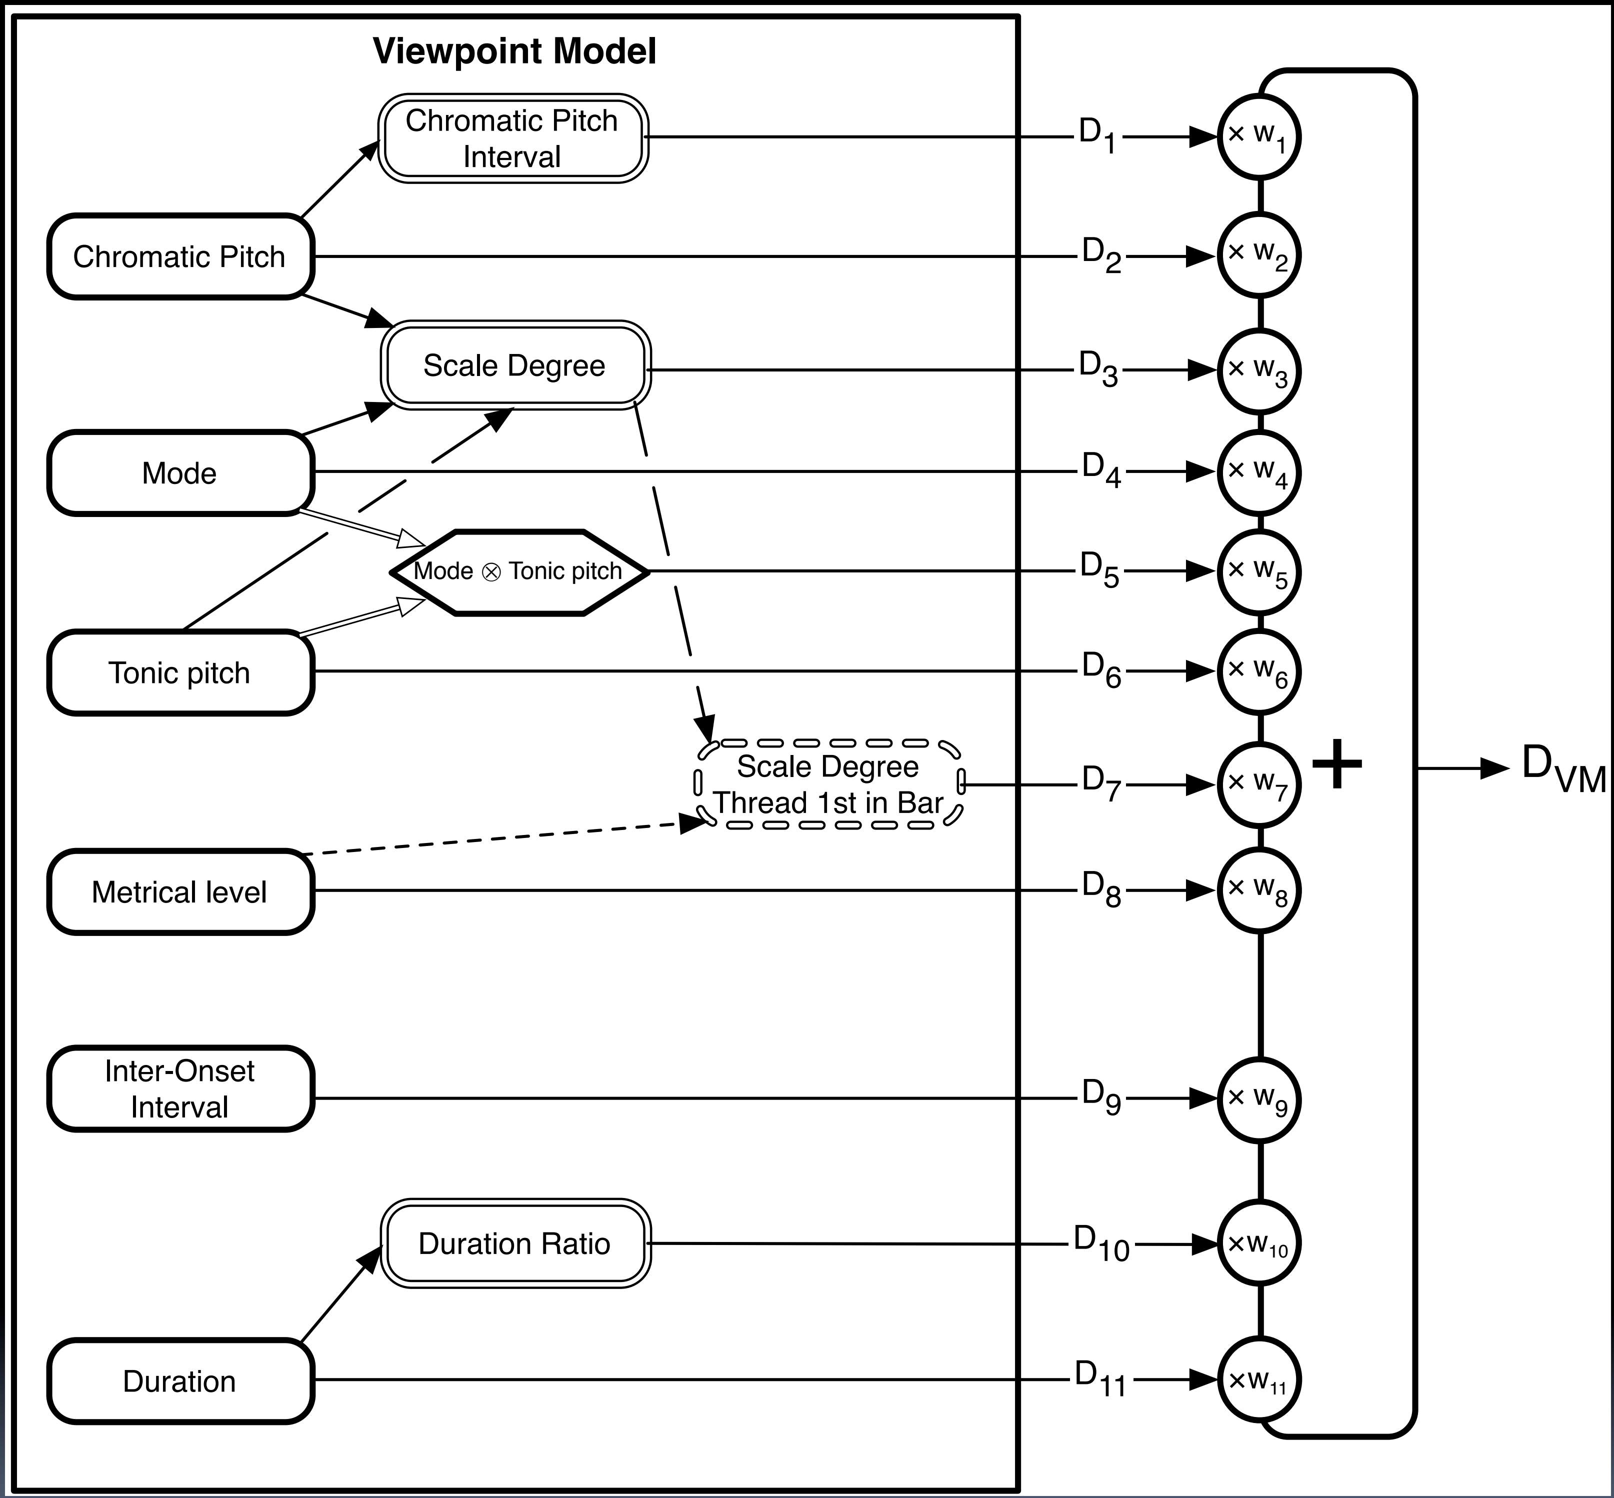

- Predictions made by
  - matching current context with strings in memory
  - all orders between 0 and maximum available
  - all contribute to final distribution
  - Feature predictions combined as linear sum weighted by entropy

Information Dynamics of Music

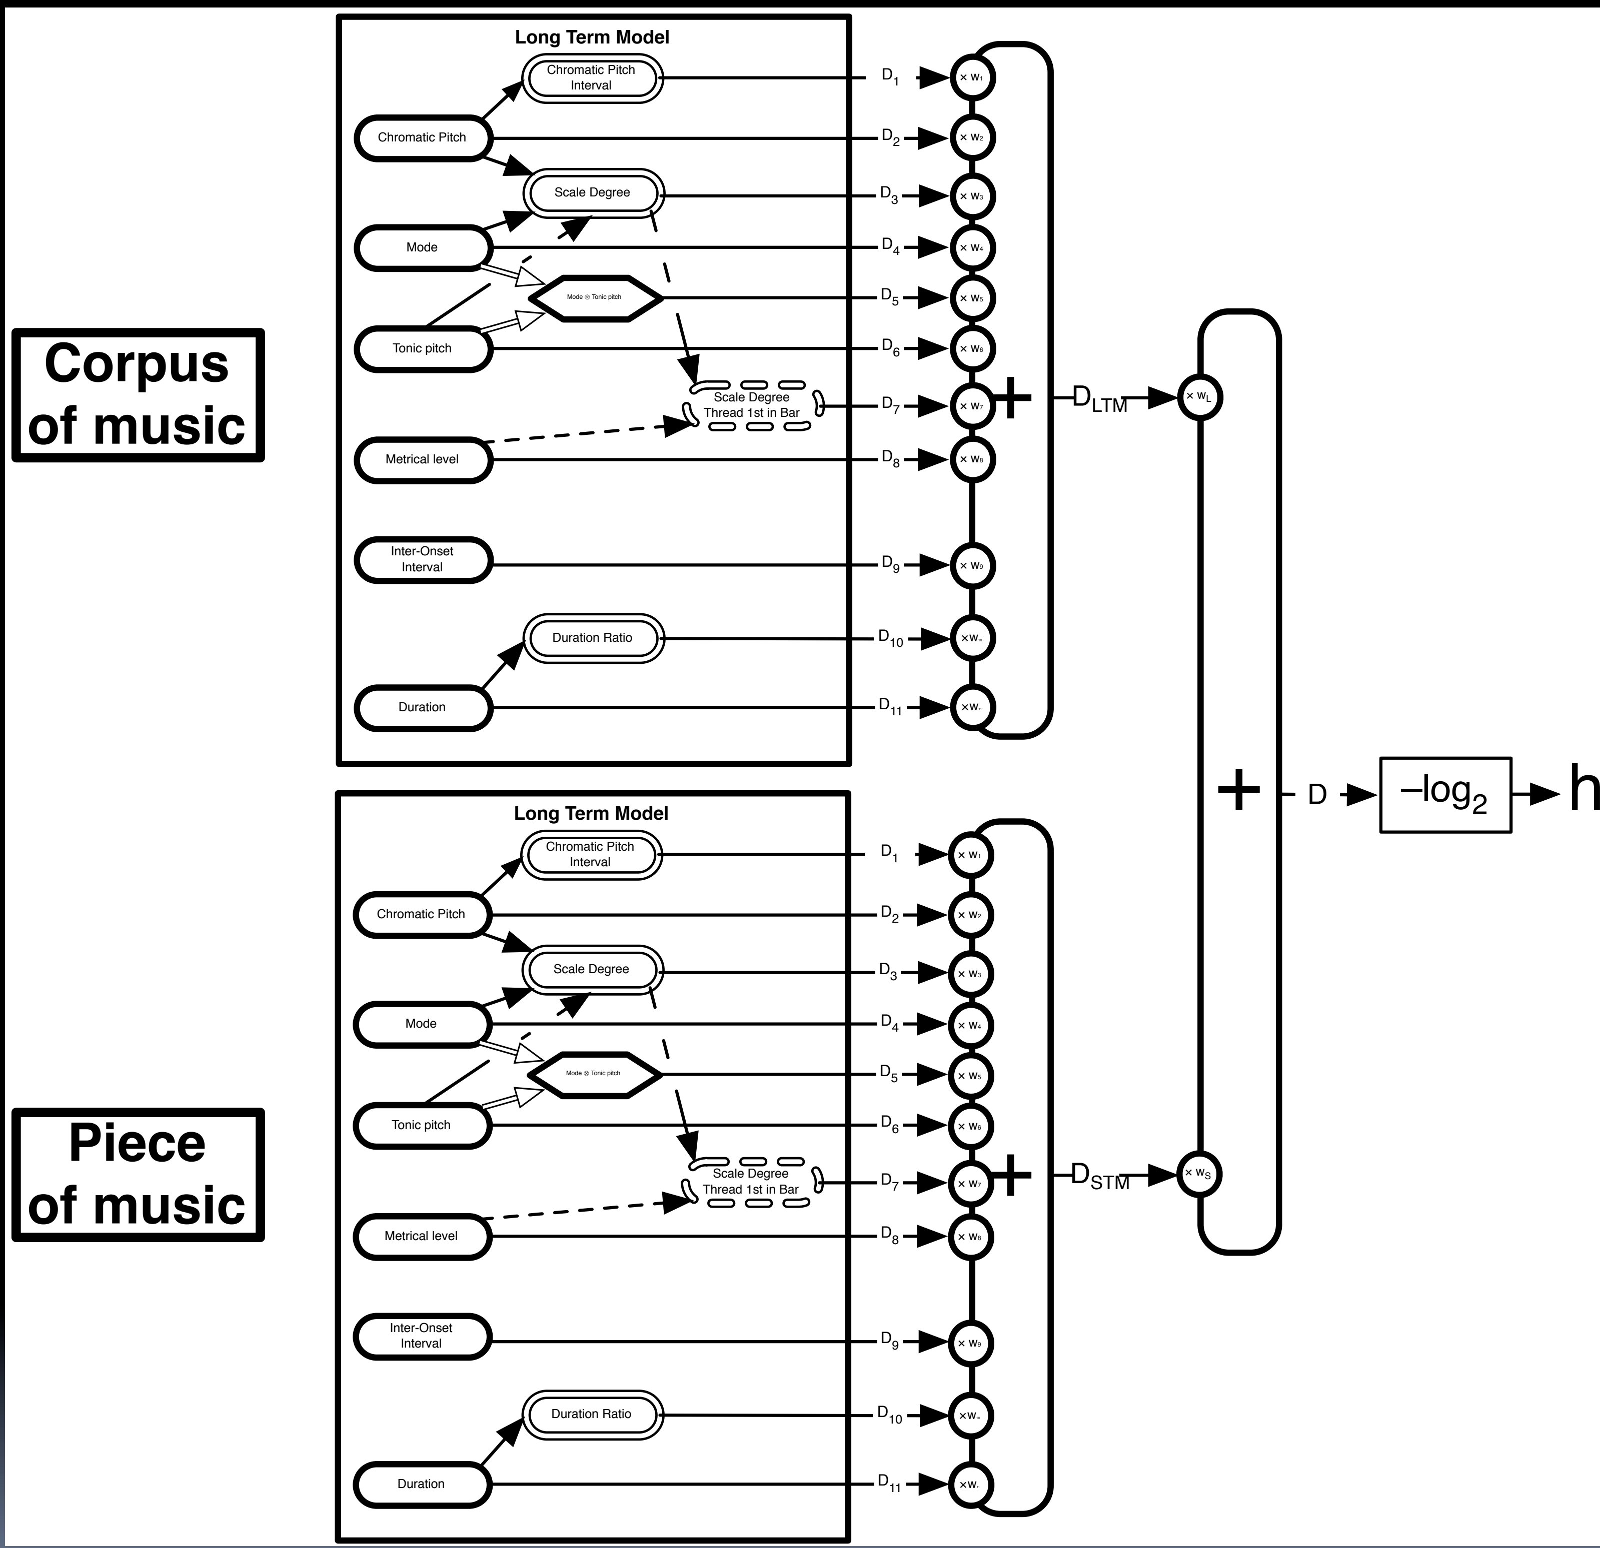

- Combined outputs of two models
  - one exposed to corpus of "enculturation" data
  - one exposed only to current melody
  - Combination is by entropic weighting, as before
- Model is "optimised"
  - inefficient viewpoints are discarded
  - model with lowest average information content is used

Information Dynamics of Music

Queen Mary University of London

- IDyOM predicts
  - listener's expectations of next note in melody
    - 4 studies; up to r=.91 correlation
    - 1 study; very high correlation with musicologists' predictions

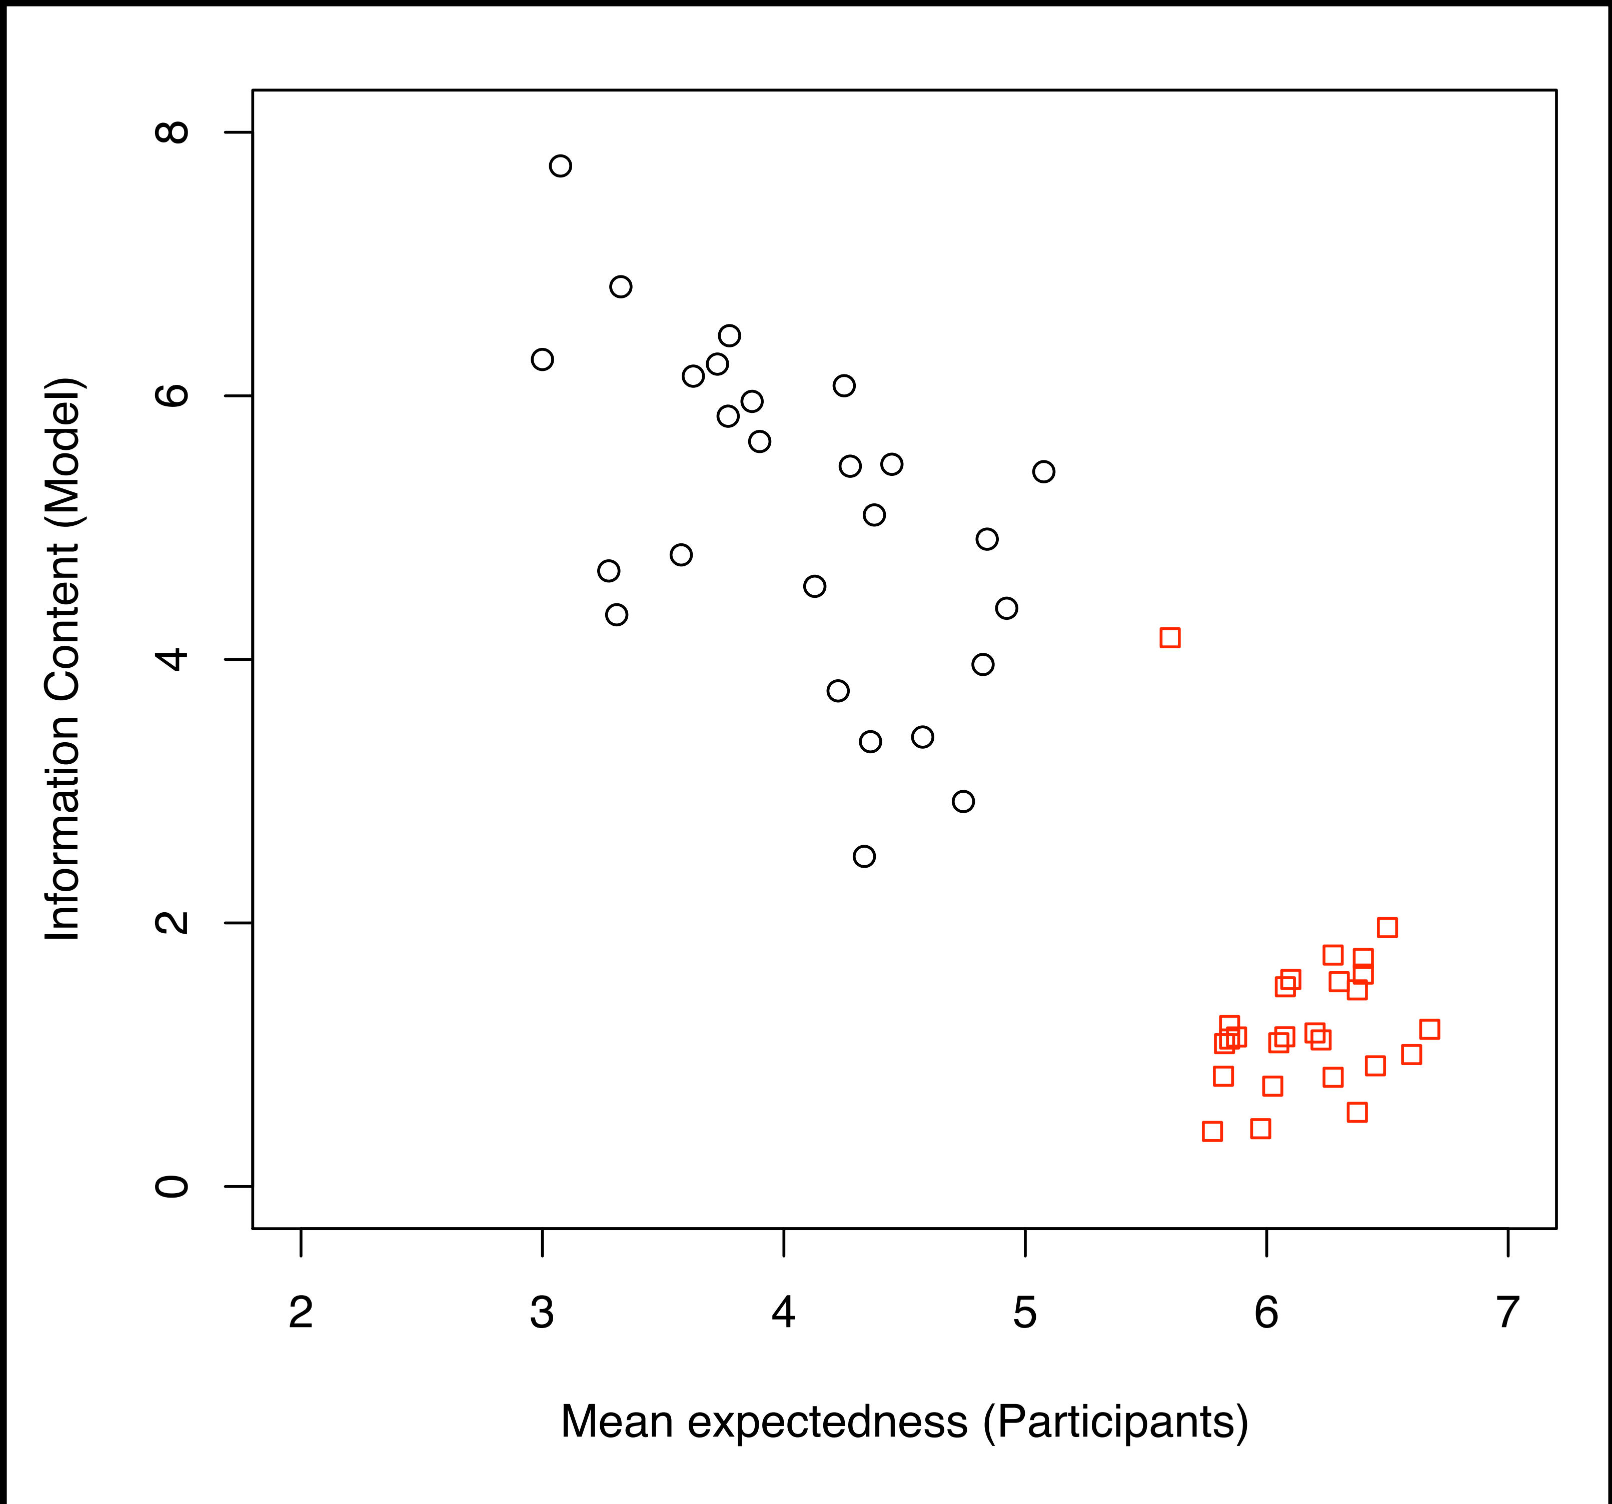

Information Dynamics of Music

Queen Mary University of London

- IDyOM predicts
  - listener's expectations of next note in melody
    - 4 studies; up to r=.91 correlation
    - 1 study; very high correlation with musicologists' predictions
  - melodic segmentation
    - 2 studies; κ = 0.58
    - vs musicologist judgements

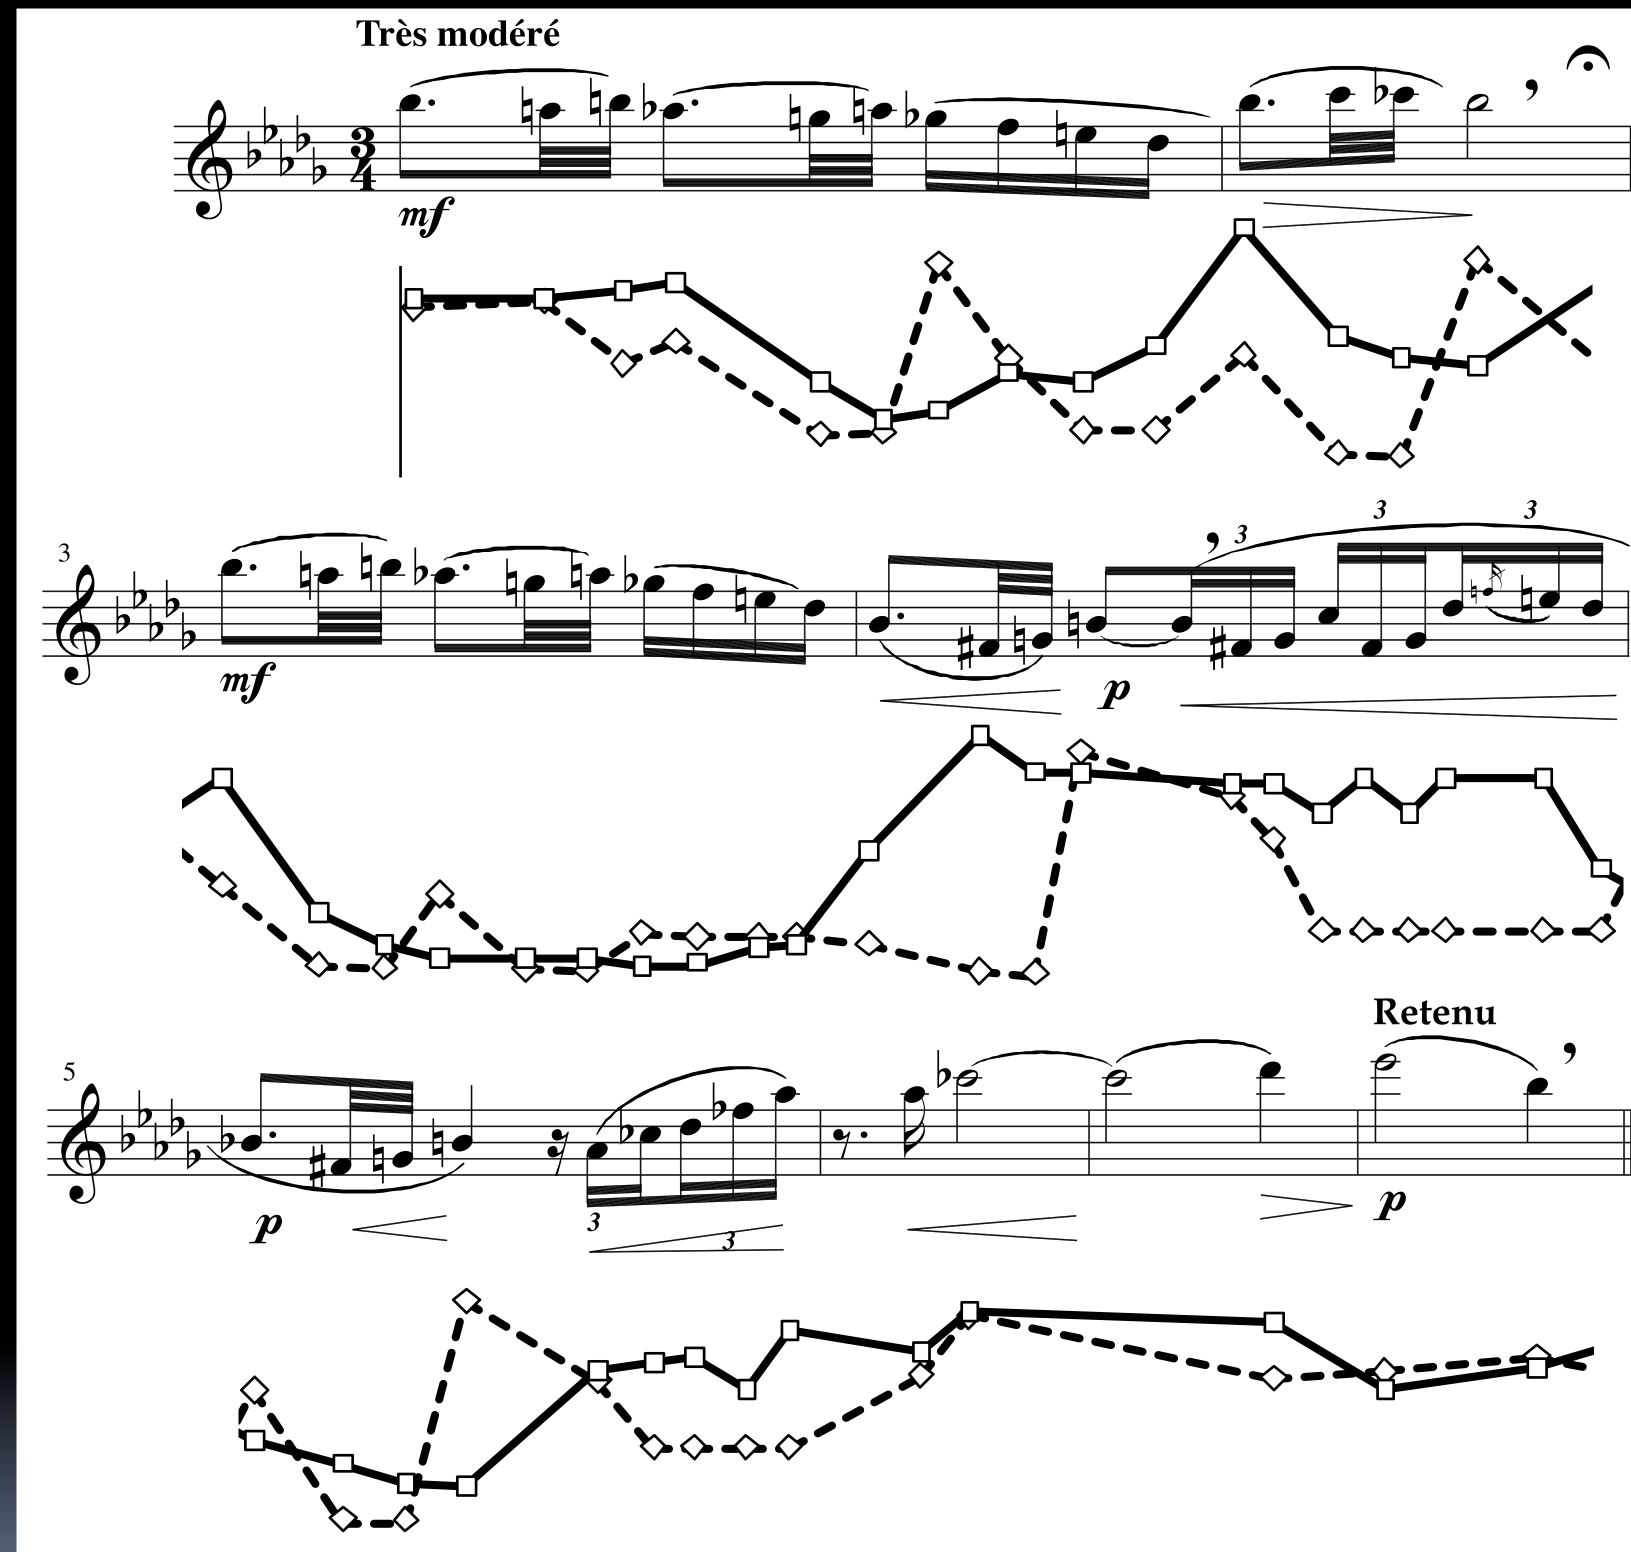

Information Dynamics of Music

- IDyOM predicts
  - listener’s expectations of next note in melody
    - 4 studies; up to r=.91 correlation
    - 1 study; very high correlation with musicologists’ predictions
  - melodic segmentation
    - 2 studies; κ = 0.58
    - vs musicologist judgements
  - neural activation with unexpectedness
    - centro-parietal region
    - strong sync. in beta-band

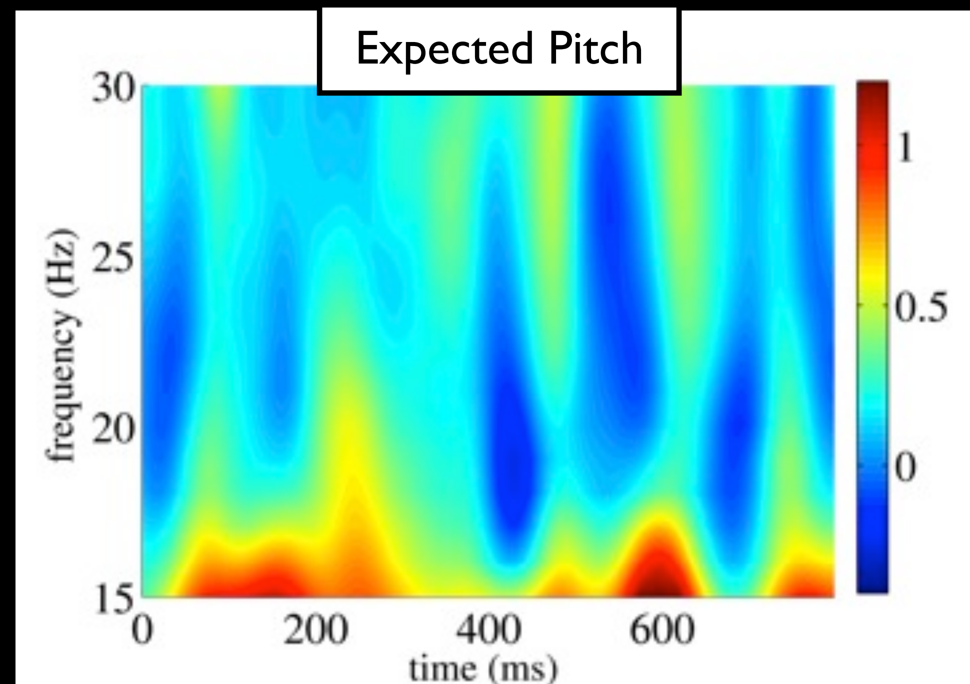

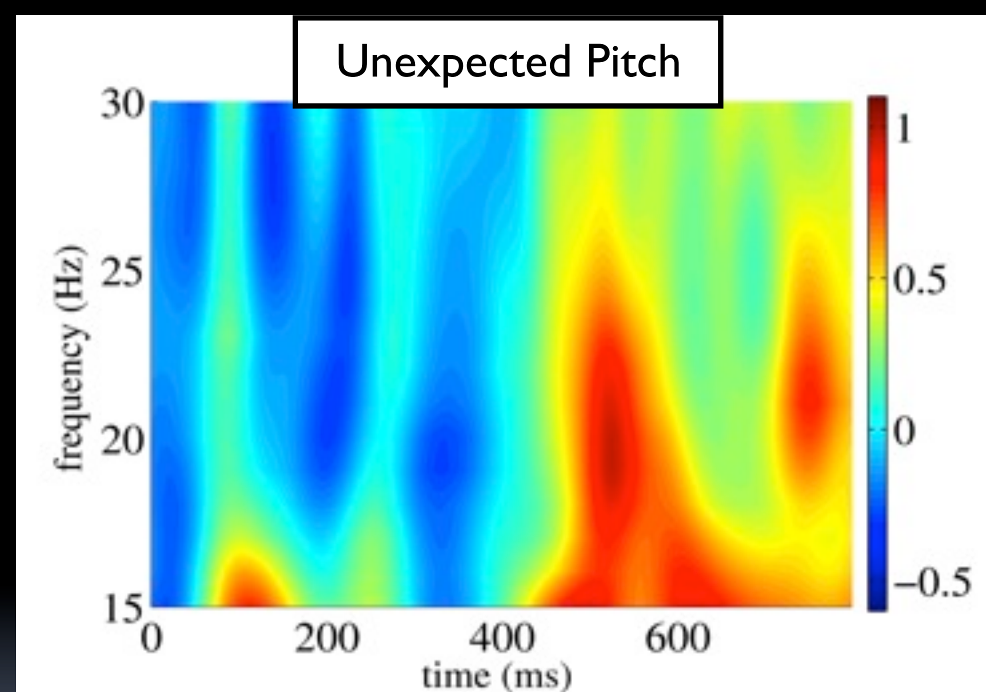

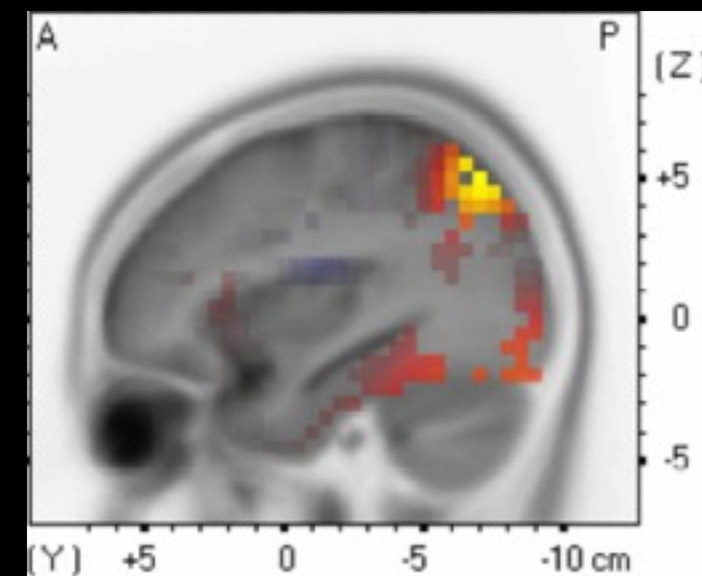

Unexpected – Expected

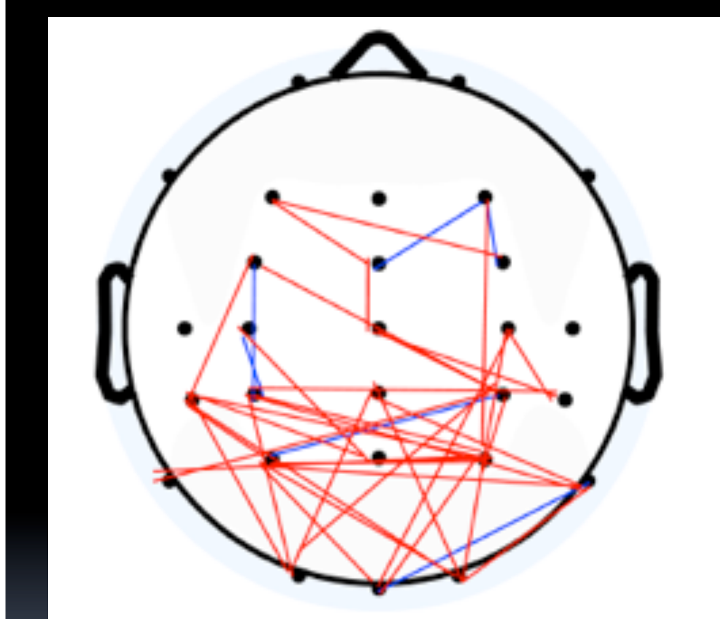

Interim conclusions

Queen Mary University of London

- IDyOM accounts for the vast majority of the data
  - with no programmed musical knowledge
  - with no training
  - without advanced musical concepts, like harmony
- This is evidence not only for IDyOM as a model of pitch perception but also for the Markovian idea of statistical perception in general
  - there is similar evidence in computational linguistics
- But it is possible to give further, stronger evidence

Meta-model

- We introduce the concept of a *meta-model*
- This is a model
  - which *reuses* an existing model
  - *without* changing it
  - *to do something related but different*
  - (meta = "beyond")
- Such a meta-model constitutes strong evidence that the original model is doing something veridical
  - the model can do more than one thing
  - if these things are related, then this evidence that the model has captured the process involved in a strong way

IDyOM: boundary detection

Queen Mary
University of London

- Narmour (1990) hypothesises that musical phrasing is derived from the expectations generated and then realised or denied as a melody proceeds
- As a phrase ends, notes become more expected; when a new phrase starts, there are only weak expectations of what is to come next
- So peak-picking in unexpectedness should allow us to predict boundaries

IDyOM: boundary detection

	Precision (1 - specificity)	Recall (sensitivity)	F1	d’	*K*
Grouper	0.67	0.87	0.76	2.94	0.73
GPR2a	0.95	0.55	0.70	2.93	0.68
LBDM2001	0.86	0.57	0.69	2.61	0.67
IDyOM	0.60	0.63	0.61	2.15	0.58
SimpleSeg.	0.25	0.35	0.29	0.99	0.22
GPR3a	0.16	0.37	0.22	0.69	0.12
always	0.08	1	0.15	-	-0.86
GPR2b	0.13	0.19	0.15	0.38	0.07
saffran.p.pitch	0.10	0.04	0.06	0.14	0.01
never	0	0	-	-	-0.04

Language segmentation

Queen Mary
University of London

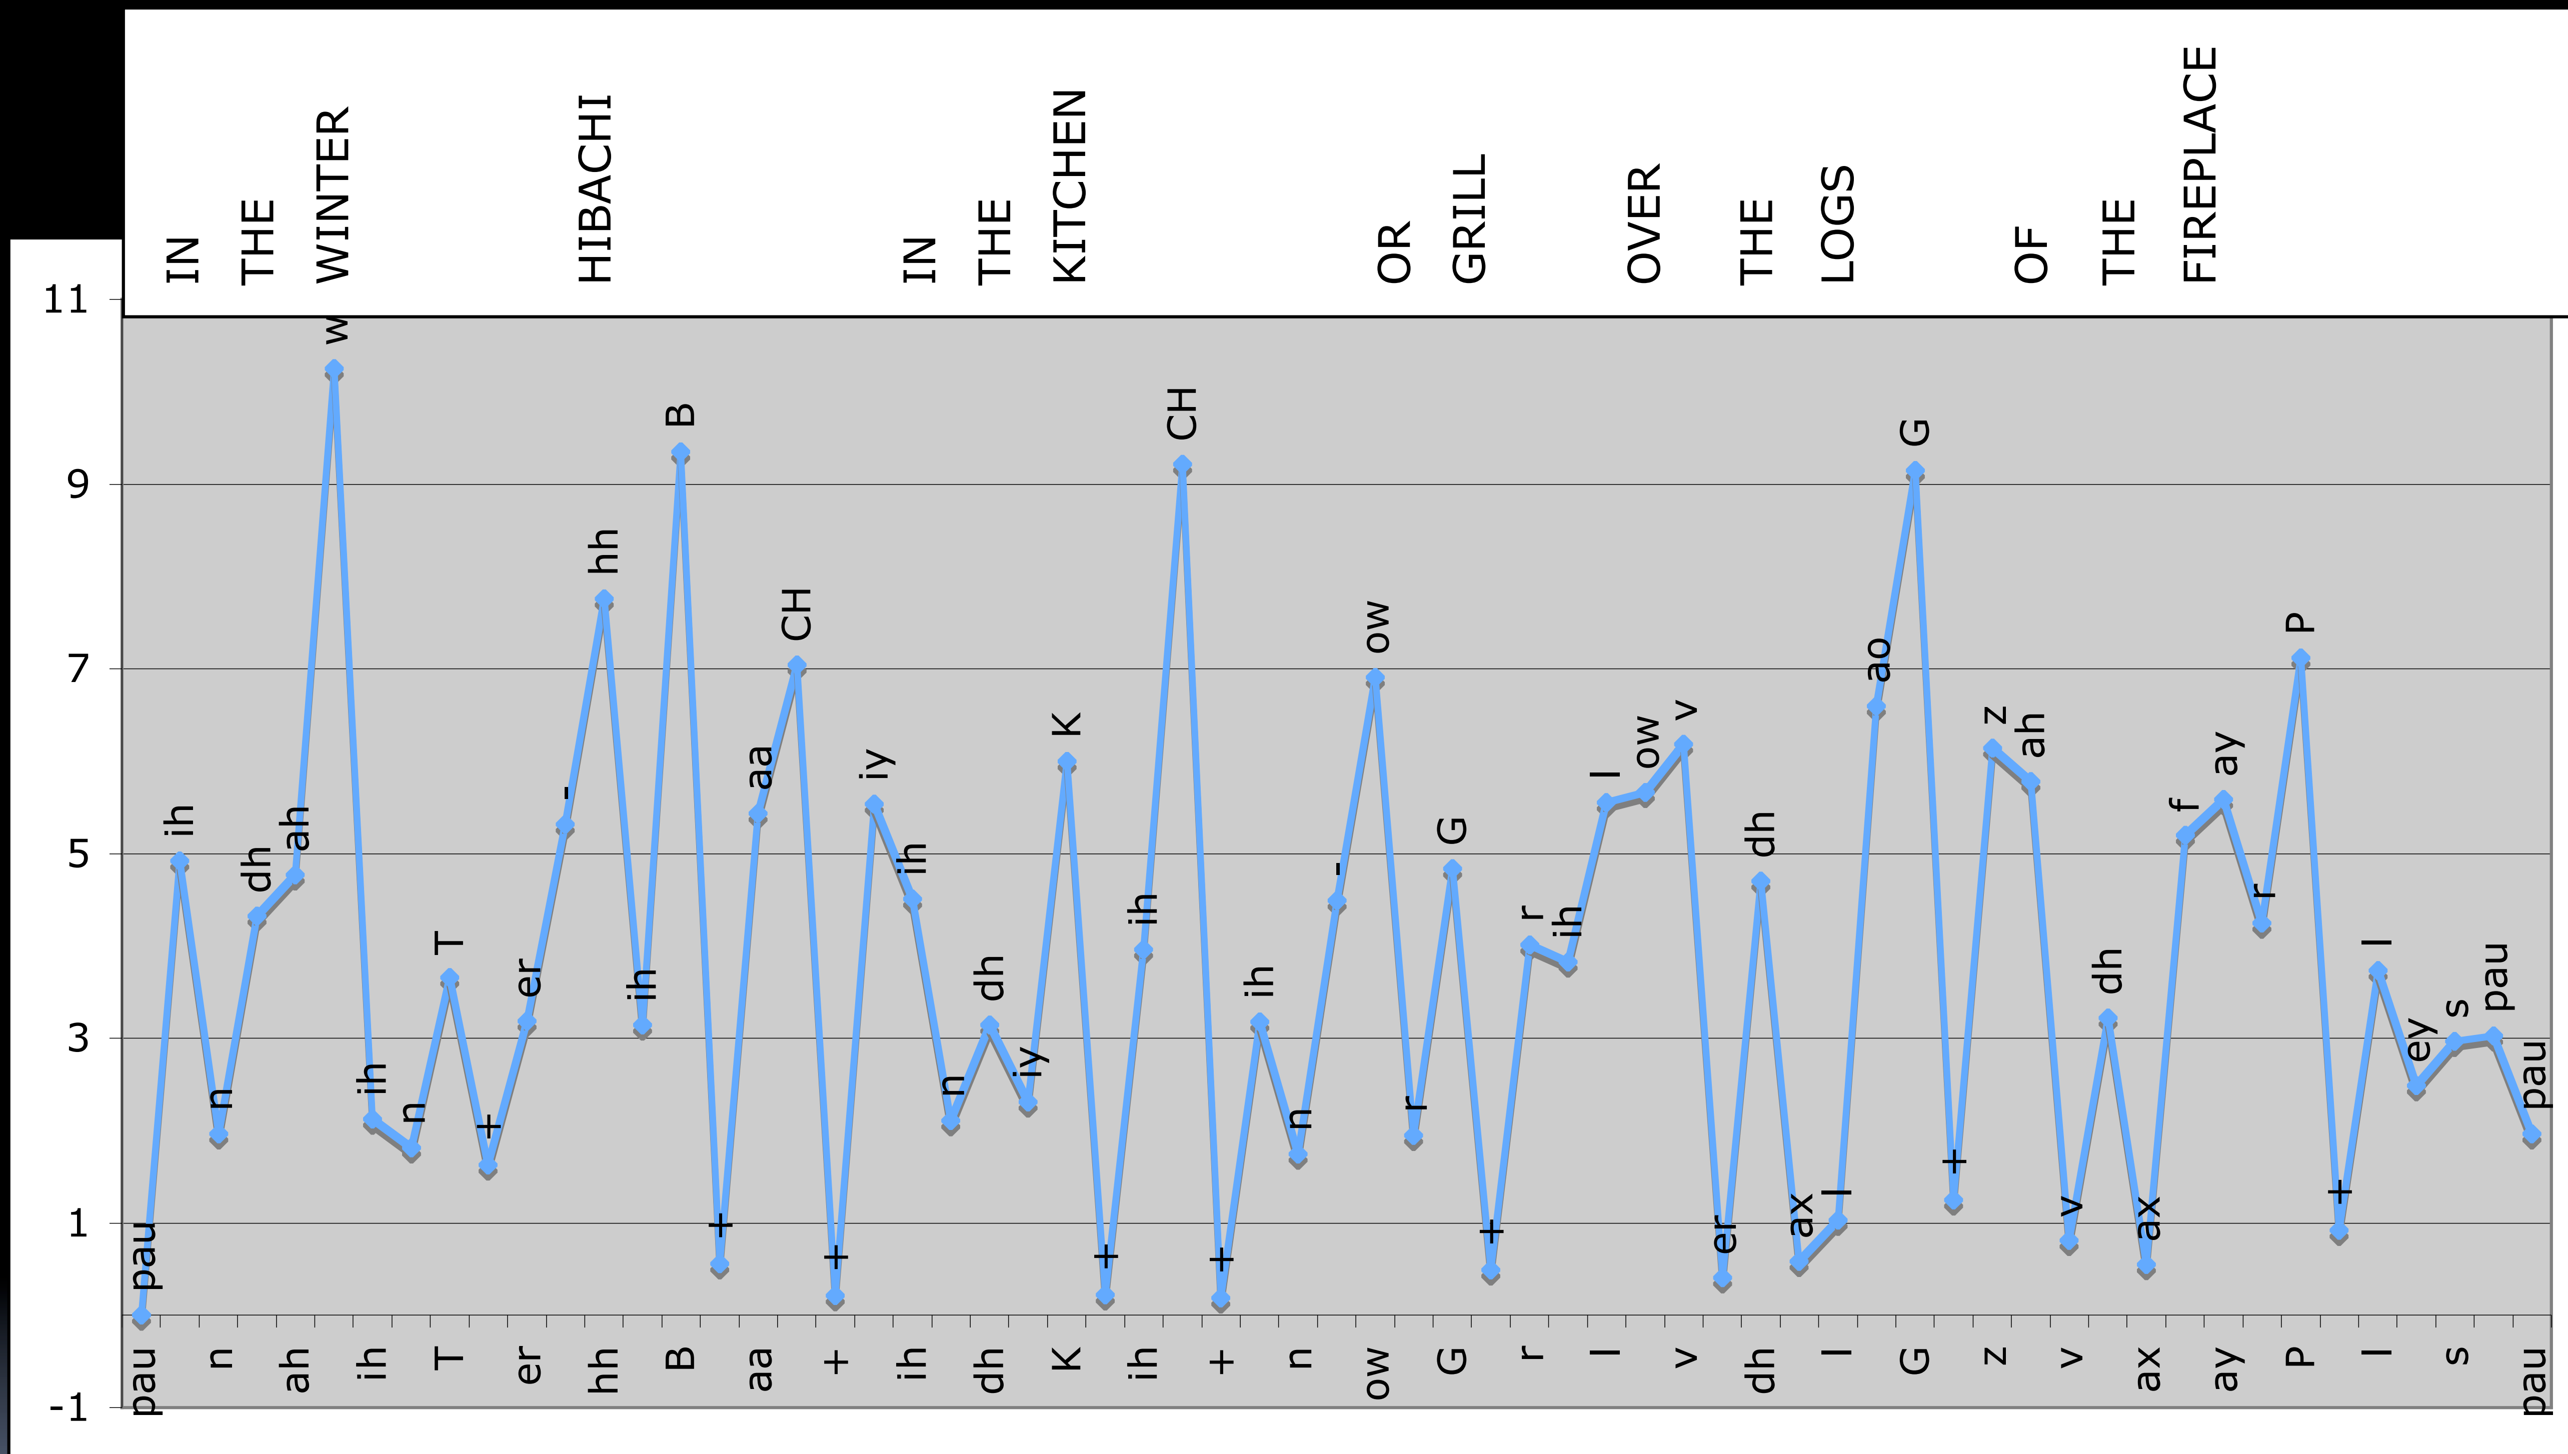

Language segmen

Queen Mary
University of London

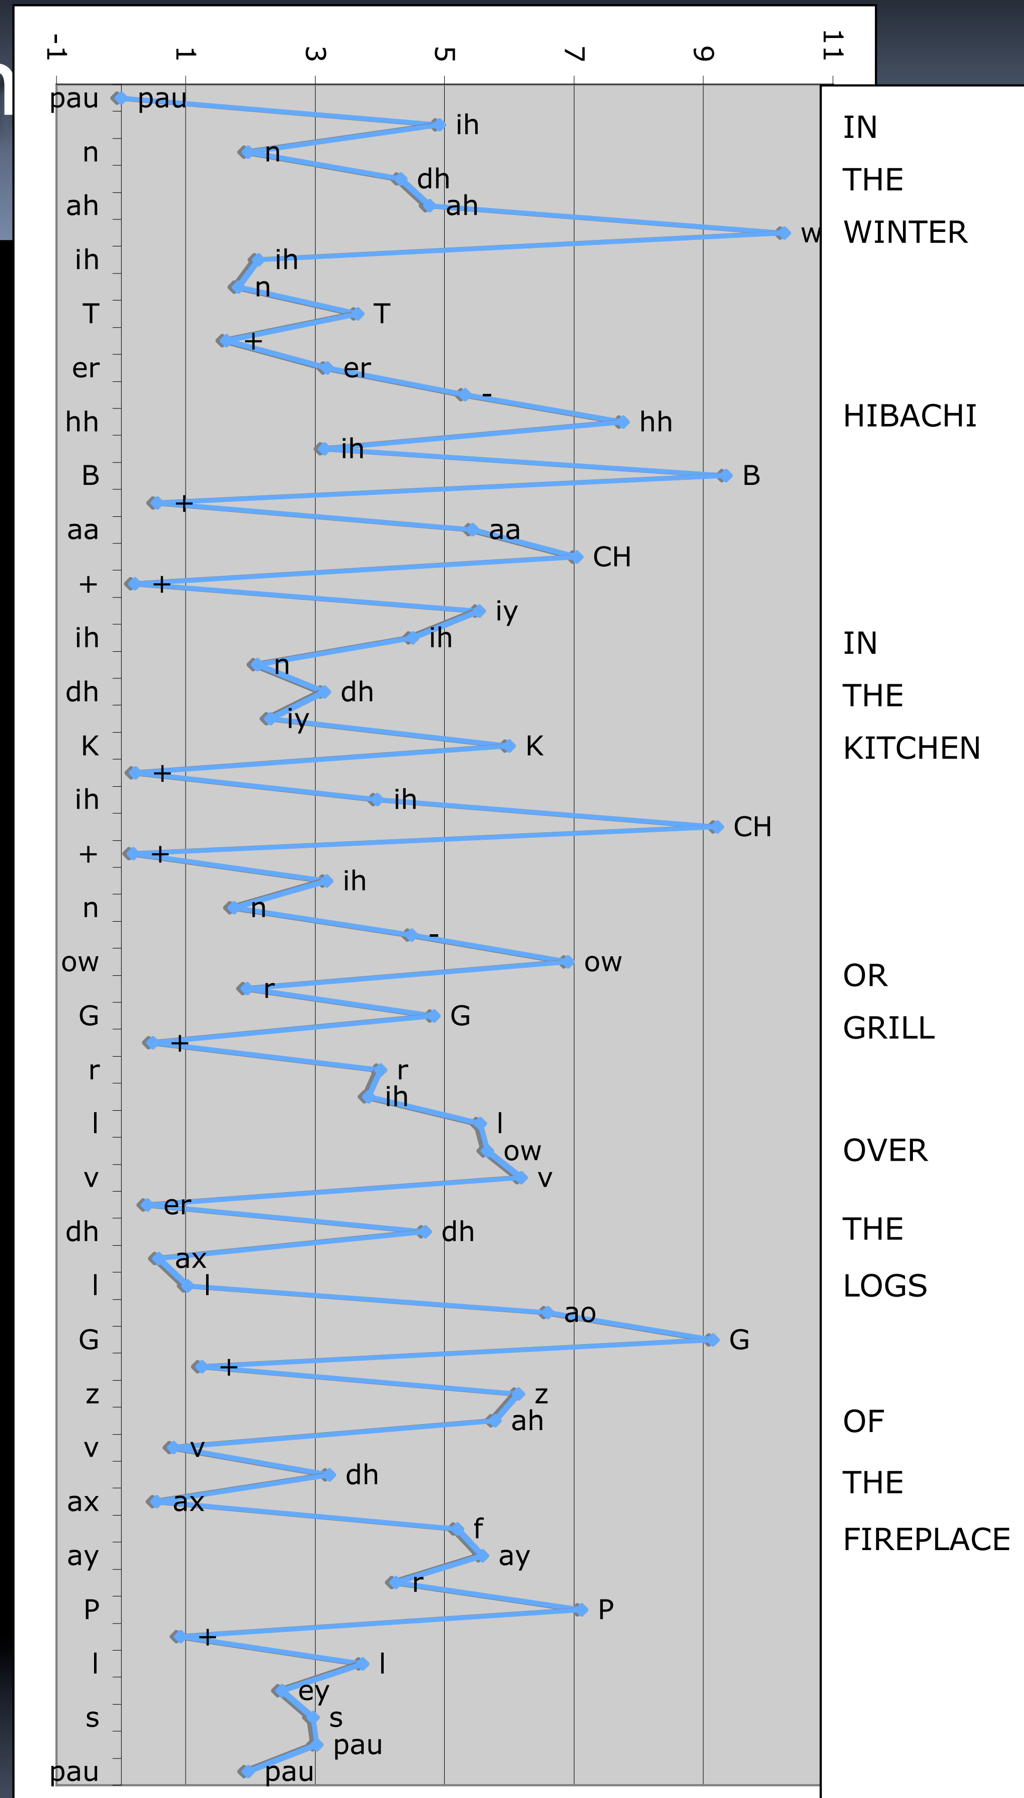

Summary

Summary

- Learned model and sequence retrieval method form **R**
- Metropolis sampling forms **T**
- Can perceptual models be expected reliably only to generate the things they perceive? (Relationship between **R** and **T**)
  - No
- How can the quality of such generation be rigorously evaluated?
- NB. Difference between
  - our evaluation of the scientific contribution
  - the system's evaluation of its own outputs (**E**)

Context

- The models we are studying focus on the conceptual space, defined by the notional rule set **R**
  - in IDyOM this is learned, and stored in the LTM
  - the advanced features of the retrieval system enhance the basic likelihoods with a kind of generalisation
- This work does not attempt to address **E**, the capacity of a creative system to introspect on its own output quality
- The traversal strategy, **T**, not really addressed in IDyOM
  - uses a standard statistical optimisation method, metropolis sampling

Evaluation

Queen Mary
University of London

- A study of **IDyOM**'s melody generation
  - **Motivation**: Evaluating outputs of creative systems
  - **Method**: Consensual Assessment Technique
  - **Question 1**: Can a model of expectation be reliably used as a generative model?
  - **Question 2**: How can quality of output be improved?

Motivation

Queen Mary
University of London

- How can we objectively evaluate the outputs of creative systems?

  “**Evaluation**

  We listened to a large number of the generated tunes, and they sounded quite good.” (********, 1998)

- We need to be much more rigorous than this

Motivation

Queen Mary University of London

- Aesthetic evaluation is (often) primary
- But we can evaluate scientifically too, and if we can, we should
- We can evaluate our generative systems in terms of engineering:
  - are they reactive enough?
  - are they reliable?
- But what does it mean to evaluate an aesthetic output scientifically?

Motivation

Queen Mary
University of London

- We need to say what we mean by “good”
- For example:
  - how well does the music generator keep in time (if that’s what it’s meant to do)
  - how well does the music generator match the implied harmony of the input melody (if that’s what it’s meant to do)

Motivation

- This means that we need to understand
  - the music we are generating
  - its context
  - its purpose
- Just like mono-disciplinary science, we need to formulate precise research questions (which may be as much aesthetic as scientific) and precise tests to evaluate them
- An artefact's (e.g., music) not being precisely specified does **not** mean that we can't ask precise questions about it

Motivation

- But how?
- Creative judgements are (desirably)
  - subjective
  - context-dependent
- How, then, can we be rigorous in evaluating the success or otherwise of creative systems output?
- The Consensual Assessment Technique (Amabile, 1998)

Consensual Assessment

Queen Mary
University of London

- Technique originally used for assessing the creative content of outputs
- Used here to assess stylistic success of outputs
- The task must be open-ended enough to permit considerable flexibility and novelty in the response
- Response must be an observable product which can be rated by judges

Consensual Assessment

- Judges must
  - be experienced in the relevant domain;
  - make independent assessments;
  - assess other aspects of the products, such as technical accomplishment, aesthetic appeal or originality;
  - make relative judgements of each product in relation to the rest of the stimuli;
  - be presented with stimuli and provide ratings in orders randomised differently for each judge.

Consensual Assessment

Queen Mary
University of London

- Most importantly, in analysing the collected data, the inter-judge reliability of the subjective rating scales must be determined
- If and only if reliability is high, we may correlate creativity ratings with other objective or subjective features of creative products
- In our version, we replace "creativity" ratings with "stylistic success" ratings

Study - Background

Queen Mary University of London

- Can this perceptual model generate music, as some researchers (eg Baroni, 1999) suggest it should?
- Model's primary data is pitch, not rhythm, so
  - create new melodies in existing rhythmic frameworks
- Use human experts to rate outputs
  - work in a well-established, formally-studied style
  - analyse success of generation in terms of stylistic success
- Use Consensual Assessment Technique to produce reliable consensus judgements

Study - Hypotheses

Queen Mary
University of London

- Three generation systems:

  A. pitch only;

  B. interval from 1st note of piece; scale degree x note duration; 1st note in phrase

  C. complex representation, including harmonic implications of melody

- Three null hypotheses:
  - Each system can generate melodies rated as equally stylistically successful in the target style as existing, human-composed melodies.

Study - Participants

Queen Mary
University of London

- Judges were 16 music researchers or students
- 5 were male and 11 female
- Their age range was 20–46 years (mean 25.9, SD 6.5)
- They had been formally musically trained for 2–40 years (mean 13.8, SD 9.4)
- 7 judges reported high familiarity with the chorale genre and nine were moderately familiar
- All judges received a nominal payment, and worked for approximately an hour

Study - Stimuli

Queen Mary
University of London

J. S. Bach: *Jesu, meiner Seelen Wonne* (BWV 359)

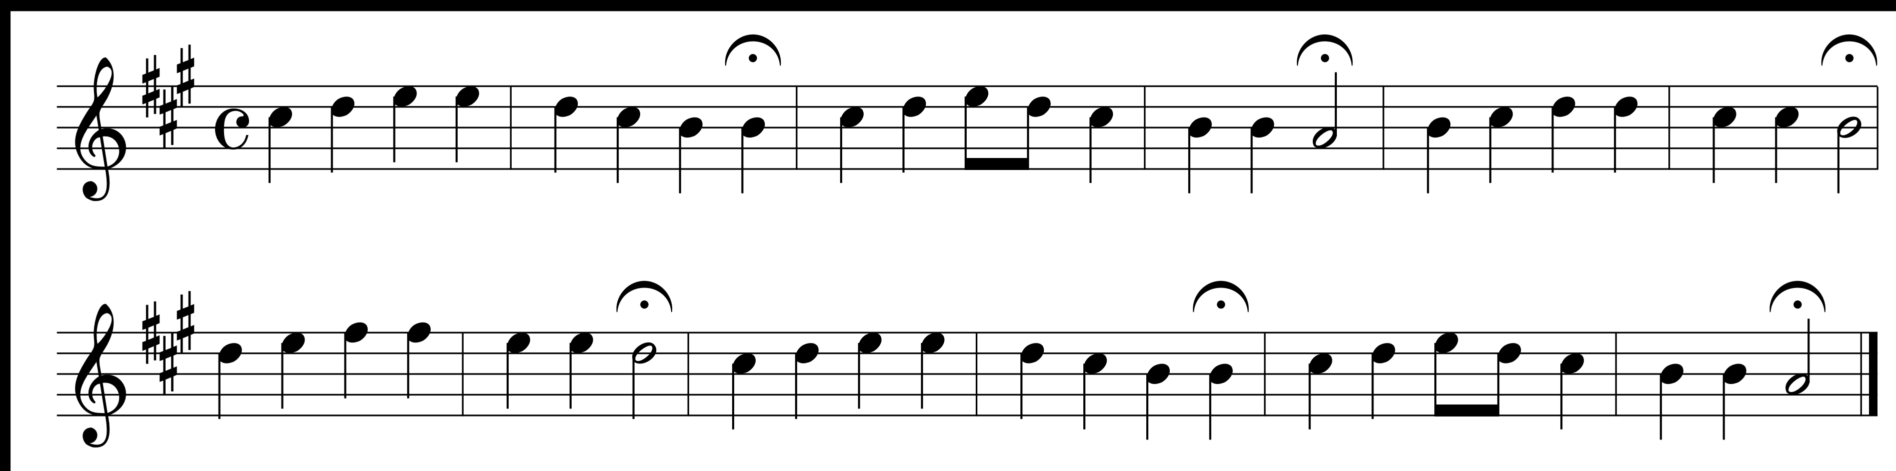

System A: *Jesu, meiner Seelen Wonne*

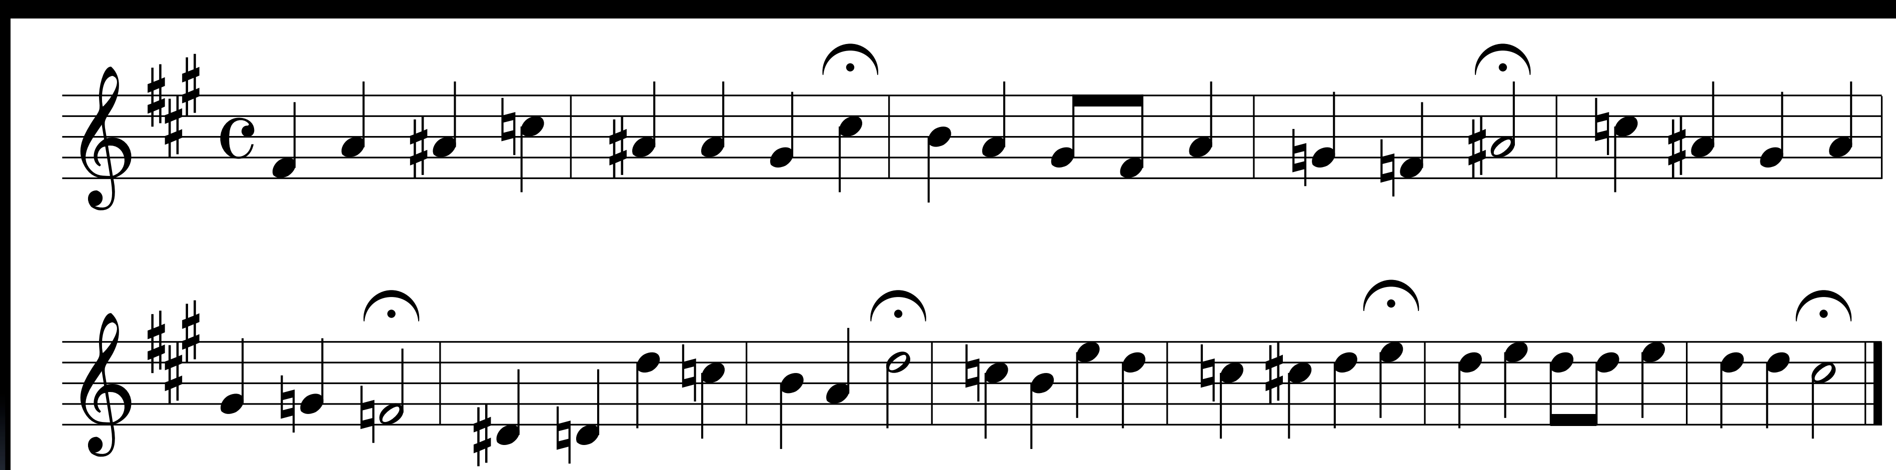

Study - Questions

Queen Mary University of London

- Questions:
  - “How successful is this as a chorale melody?”
  - Judges were advised to reflect such factors as
    - conformity to important stylistic features;
    - tonal organisation;
    - melodic shape and interval structure;
    - melodic form.
  - and required to explain their answers
  - “Do you recognise the melody?”
- NB: participants not told that a computer was involved
  - Moffat & Kelly (2006) demonstrate systematic bias against computers

Study - Results

- All but 2 of the 120 pairwise correlations between judges were significant ($p<0.05$) with a mean coefficient of $r(26)=0.65$ ($p<0.01$)
- This high consistency warrants averaging the ratings for each stimulus across individual judges in subsequent analyses

Study - Results

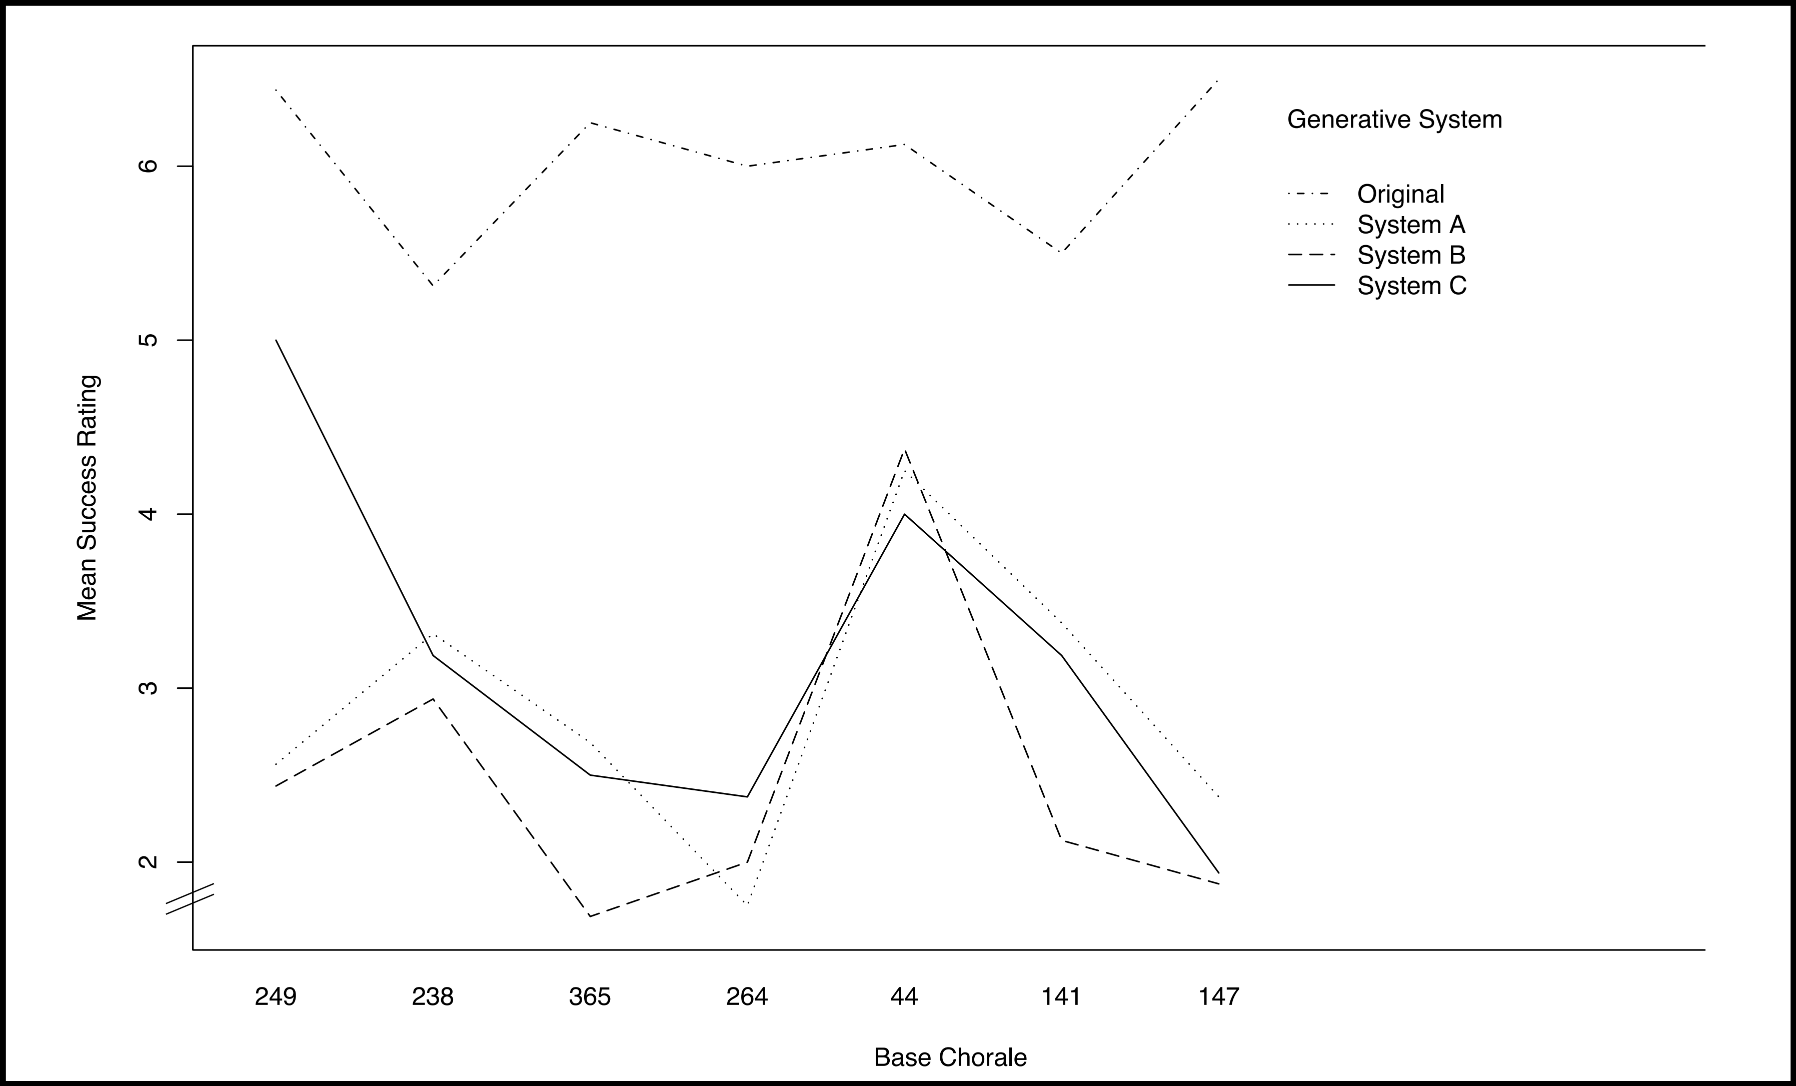

- Friedman's Test shows significant intrasubject effect of system on ratings ($X^2(3)=32$, $p < 0.01$)
- Wilcoxon rank sum test with Holm's Bonferroni correction show "original" ratings differ significantly from IDyOM's ($p < 0.03$)

Using the results

Queen Mary
University of London

- We examined the judges' explanations of their judgements
- We developed a corresponding set of objective descriptors
- We applied the descriptors in a multiple regression analysis
  - dependent variable: the rating scheme averaged across stimuli

Using the results

Queen Mary
University of London

- Objective features:
  - Pitch Range
  - Melodic Structure
  - Tonal Structure
  - Phrase Structure
  - Rhythmic Structure
- Redundant descriptors removed by backwards stepwise elimination

Using the results

System D: *Jesu, meiner Seelen Wonne*

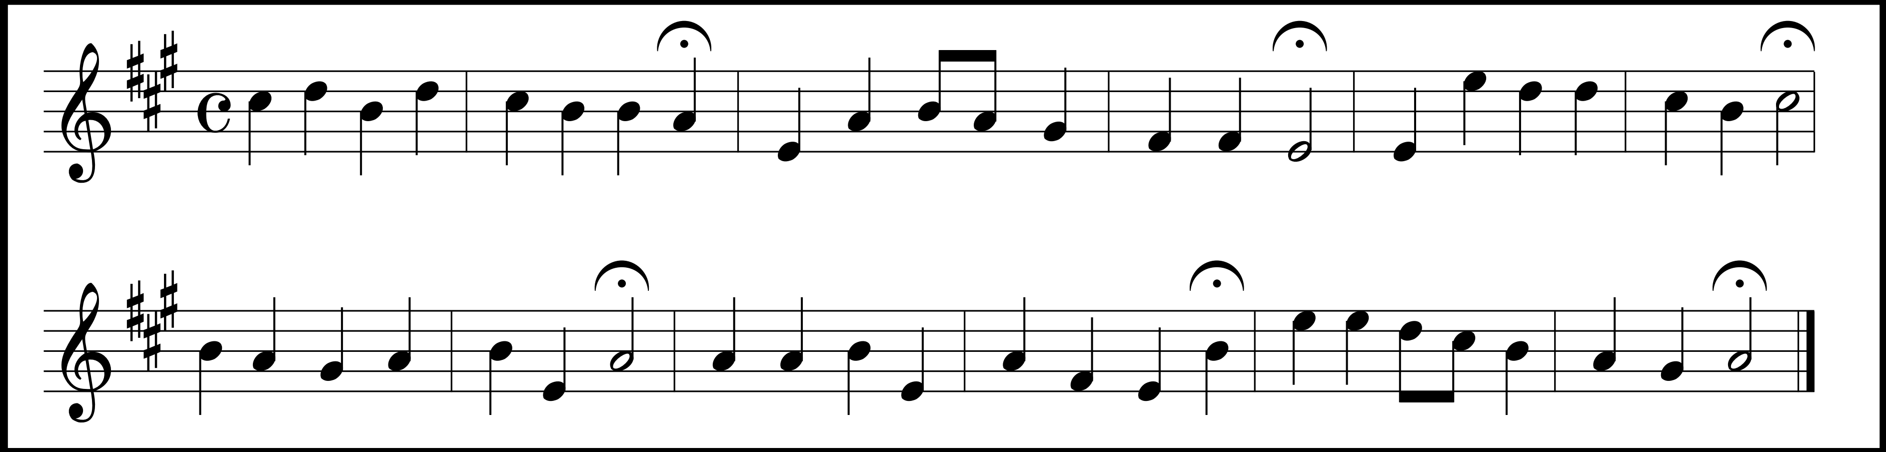

- Much more coherent example than earlier systems
- But System D must be rigorously analysed in further cycles
- What is going on here is transformational creativity:
  - the experimental results are being used to modify **R**

Summary of Results

Queen Mary
University of London

- Systems A-D all fail to compare with original melodies
- System D is, however, significantly better than Systems A-C
- The CAT provides a rigorous way to assess the stylistic success of outputs of computational creative systems
- Statistically significant agreement between expert judges is crucial in supporting the conclusions of the work

Skip to Summary

Summary

Queen Mary
University of London

- Two methods for establishing the quality of automated musical composition systems
  - expert judgements on style
  - non-expert judgements on meaning
- Both are precisely specified (not just “good”)
- Two different techniques
  - focused on particular knowledge, with articulate feedback on what is not right
  - hidden amongst many different possible dimensions of quality, so as not to prime participants
- It is up to researchers to establish what the dimensions of quality are for their own systems and their own work and to use them rigorously
- In this way, perhaps we can also reach a more objective view of aesthetics

The Future

Queen Mary University of London

- Many more models and many more studies are needed
- Studying **T** is paramount (see, e.g., Guilford, Chikszentmihalyi for starting points)
- Models that learn their rules are likely to be more credibly creative than programmed models
- Music is a good domain for this study because it is uncluttered by semantics (in the language sense) so it can be looked at on the perceptual/cognitive level only

Summary

- In this lecture we have looked at
  - a way of characterising creativity that is amenable to computational systems
  - a reason for, and a way of, studying music scientifically and objectively
  - a way of doing that study computationally, and getting more out of it in consequence
  - ways of evaluating what we do rigorously
- We have now looked at
  - **R**: IDyOM perceptual model
  - **E**: Human evaluators; feedback in form of revised perceptual model
- We have used a temporary make-do
  - **T**: Metropolis sampling

Acknowledgements

- This work is funded by UK Engineering and Physical Sciences Research Council grants
  - GR/S82220:
    "Techniques and Algorithms for Understanding the Information Dynamics of Music"
  - EP/H01294X:

    "Information and neural dynamics in the perception of musical structure"
  - EP/D038855:

    "Modelling Musical Memory and the Perception of Melodic Similarity"

References

Lerdahl, F. & Jackendoff, R. (1983) **A Generative Theory of Tonal Music.** Cambridge, MA: MIT Press

Narmour, E. (1990) **The Analysis and Cognition of Basic Melodic Structures: The Implication-realisation Model**. Chicago: University of Chicago Press

Müllensiefen, D., Pearce, M.T., Wiggins, G.A. & Frieler, K. (2007) **Segmenting Pop Melodies: A Model Comparison Approach**. Proceedings of SMPC'07, Montreal, Canada

Pearce, M.T. & Wiggins, G.A. (2006) **Expectancy in melody: The influence of context and learning.** Music Perception, 23(5), 377–405

Potter, K., Wiggins, G.A. & Pearce, M.T. (2007) **Towards Greater Objectivity in Music Theory: Information-Dynamic Analysis of Minimalist Music.** Musicae Scientiae, 11(2), 295–324

Shannon, C. (1948) **A mathematical theory of communication.** Bell System Technical Journal, 27, 379-423, 623-56

Wiggins, G.A. (2007) **Models of Musical Similarity.** Musicae Scientiae, Discussion Forum 4a, 315–338

Pearce, M.T. & Wiggins, G.A. (2012) **Auditory Expectation: The Information Dynamics of Music Perception and Cognition**. Topics in Cognitive Science, in press.

Wiggins, G.A. (2012) **"I let the music speak": cross-domain application of a cognitive model of musical learning**. In Rebuschat, P. and Williams, J. (Eds.), Statistical Learning and Language Acquisition, Mouton De Gruyter, Amsterdam, NL

Creativity before consciousness

a mechanism admitting spontaneous creativity in Baars' Global Workspace

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

Outline

- What I mean by "spontaneous creativity"
- Background
  - Taine's Theatre of Consciousness, the Society of Mind and Global Workspace Theory
    - The Threshold Paradox
  - Statistical models of cognitive process
  - Information theory
- A hypothetical model of cognitive selection that accounts for spontaneous creativity
- Evaluation – a difficult problem
- Motivation: **WHERE DO (MUSICAL) IDEAS COME FROM?**

Two kinds of creativity

Queen Mary
University of London

- One aspect of creativity is **SPONTANEOUS**
  - ideas appear, spontaneously, in consciousness
  - cf. Mozart (Holmes, 2009, p. 317)
    - When I am, as it were, completely myself, entirely alone, and of good cheer – say traveling in a carriage, or walking after a good meal, or during the night when I cannot sleep; it is on such occasions that my ideas flow best and most abundantly.
- Compare with the composer working to build (e.g.) a new version of a TV theme, on schedule, and with constraints on "acceptable style"
  - this is a different kind of activity: **CREATIVE REASONING**
- Most creative acts of any size are a **mixture of both**
- Here, I focus on **SPONTANEOUS CREATIVITY** only

Background: societies of mind

- Hippolyte Taine (1871) proposed the first (?) multi-agent theory of mind, based in a **Theatre of Consciousness**
  - narrow theatre stage, with actors appearing, disappearing, and planning off-set
- Marvin Minsky (1987) proposed the **Society of Mind**
  - computational knowledge-rich agents, communicating & collaborating hierarchically to achieve goals
- Bernard Baars (1988) proposed the **Global Workspace Theory**
  - agents, generating cognitive structures, communicating via a shared blackboard
  - agnostic as to nature of agent-generators
- The three theories are not incompatible
  - Baar's agents/representations are underspecified, and don't contradict Minsky's
  - The key difference is in the communication mechanism
    - but even that may not be contradictory...

Background: societies of mind

Queen Mary
University of London

- Society of Mind uses a hierarchical structure of control
  - agents recruit other agents according to task
  - communication passes up and down hierarchy
  - locus of consciousness is explicitly excluded
- Global Workspace Theory uses a central communication exchange, the Global Workspace
  - corresponds with Taine's "theatre" of consciousness
  - can hold one item at a time (some researchers suggest this should be 2 or 2.5)
  - all agents have read-access to Global Workspace
  - in later developments, Baars proposes a hierarchical system of "local" workspaces feeding into the Global Workspace, reducing information overload
  - there is a "threshold" to be "crossed" to get write-access to the GW
  - granting access can be viewed as **assignment of conscious attention**

Background: the threshold paradox

Queen Mary
University of London

- Baars writes (somewhat metaphorically) about agents "recruiting" others to support a given cognitive structure
  - when enough agents support the structure it is "loud" enough to pass the threshold and enter consciousness
    - I'll use this analogy of "volume" later; Baars proposes synchrony as the implementing mechanism and Shanahan (2010) identifies the necessary neural substrate
- However, there is a problem: **The Threshold Paradox**
  - To communicate in the global workspace, an agent needs to recruit supporters
  - To recruit supporters, an agent must communicate in the global workspace
    - because that is the only medium of communication
- This talk is about an alternative view of access to the Global Workspace

Background: the threshold paradox

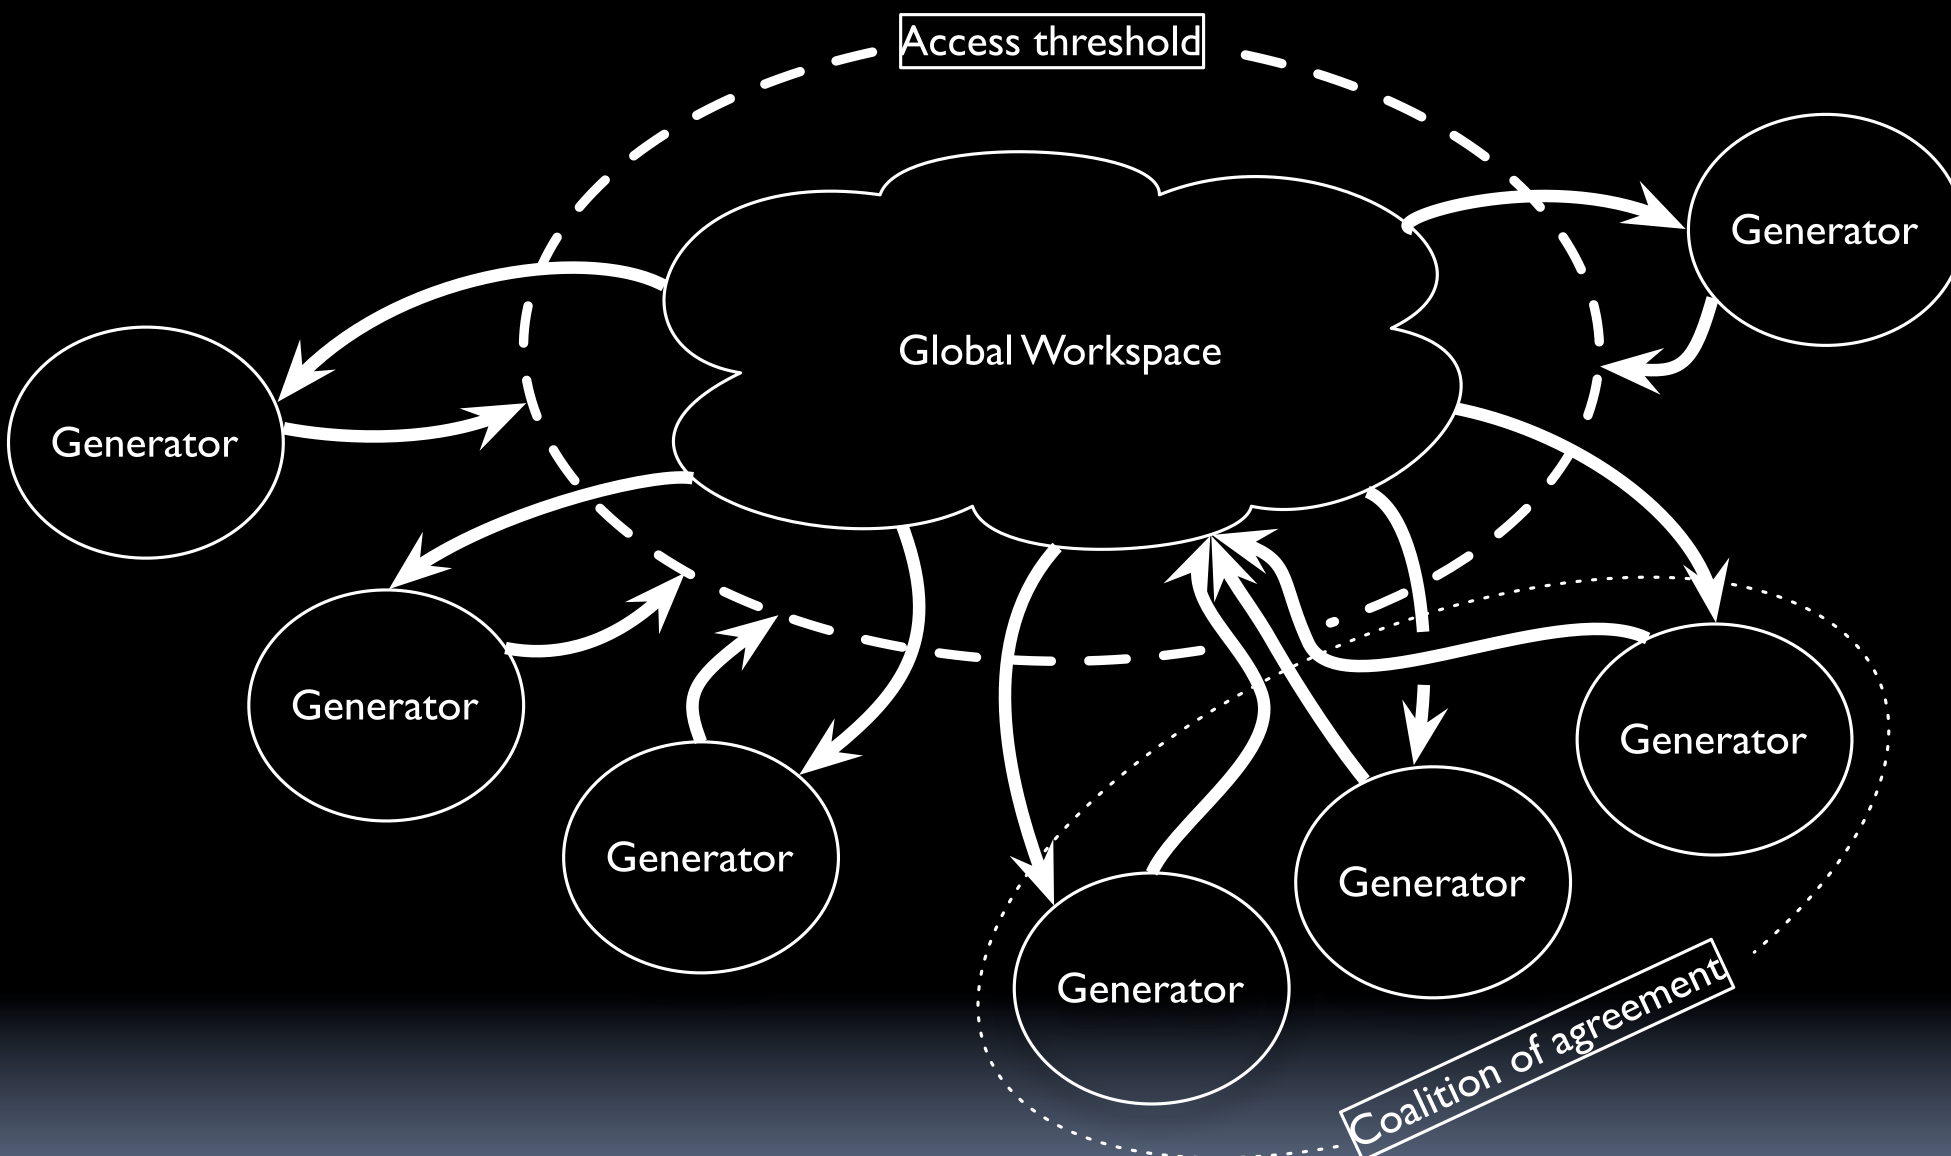

Background: information theory

Queen Mary University of London

- I use two versions of Shannon's *entropy* measure (MacKay, 2003)
  - the number of bits required to transmit data between a hearer and a listener given a shared data model
  - *information content*: estimated number of bits required to transmit a given symbol as it is received:

$$h = -\log_2 p_s$$

  - models **unexpectedness**
  - *entropy*: expected value of the number of bits required to transmit a symbol from a given distribution, prior to sending/receipt:

$$H = -\sum_i p_i \log_2 p_i$$

  - models **uncertainty**
  - p_s, p_i are probabilities of symbols; i ranges over all symbols in the alphabet

Background: statistical cognitive models

Queen Mary
University of London

- Organisms need to be able to **anticipate** the world
  - use (mental) models to predict what is coming next
  - use learned models, trained by observed likelihood
  - use temporal association (implication/consequence)
  - use co-occurrence (conjunction)
- Can model music and language (and other things) in this way
  - currently using IDyOM model (Pearce, 2005; Pearce & Wiggins, 2006)
    - predicts human melodic expectation (R^2=.81; Pearce & Wiggins, 2006)
    - predicts human melodic segmentation (F_1=.61; Pearce, Müllensiefen & Wiggins, 2010)
    - predicts language (phoneme) segmentation (F_1=.67; Wiggins, 2011)
- Claim is that mental process is literally statistical
    - statistical nature means we can apply information theory (Shannon, 1948)

Instantiating the Global Workspace

Queen Mary University of London

- Agent generators (not specified by Baars; simpler than Minsky's?)
  - statistical samplers predicting next in sequence from shared learned models of perceptual and other domains
  - many agents, working in massive parallel
    - at all times, the likelihood of a given prediction is proportional to the number of generators producing it *(this isn't in Baars' theory, but it will be important later)*
  - receive perceptual input from sensory systems
    - continually compare previous predictions with current world state
  - continually predict next world state from current matched predictions
    - sensory input does not enter memory directly
    - the expectation that matches best is recorded
  - consider state n (current) and state n+1 (next)
    - at state n, we can calculate h_n, H_n, and H_{n+1} (but not h_{n+1}, because it hasn't happened yet)

Baars' (1988) Global Workspace Theory

Queen Mary
University of London

- "Aha" moment = passage into consciousness
- Threshold paradox

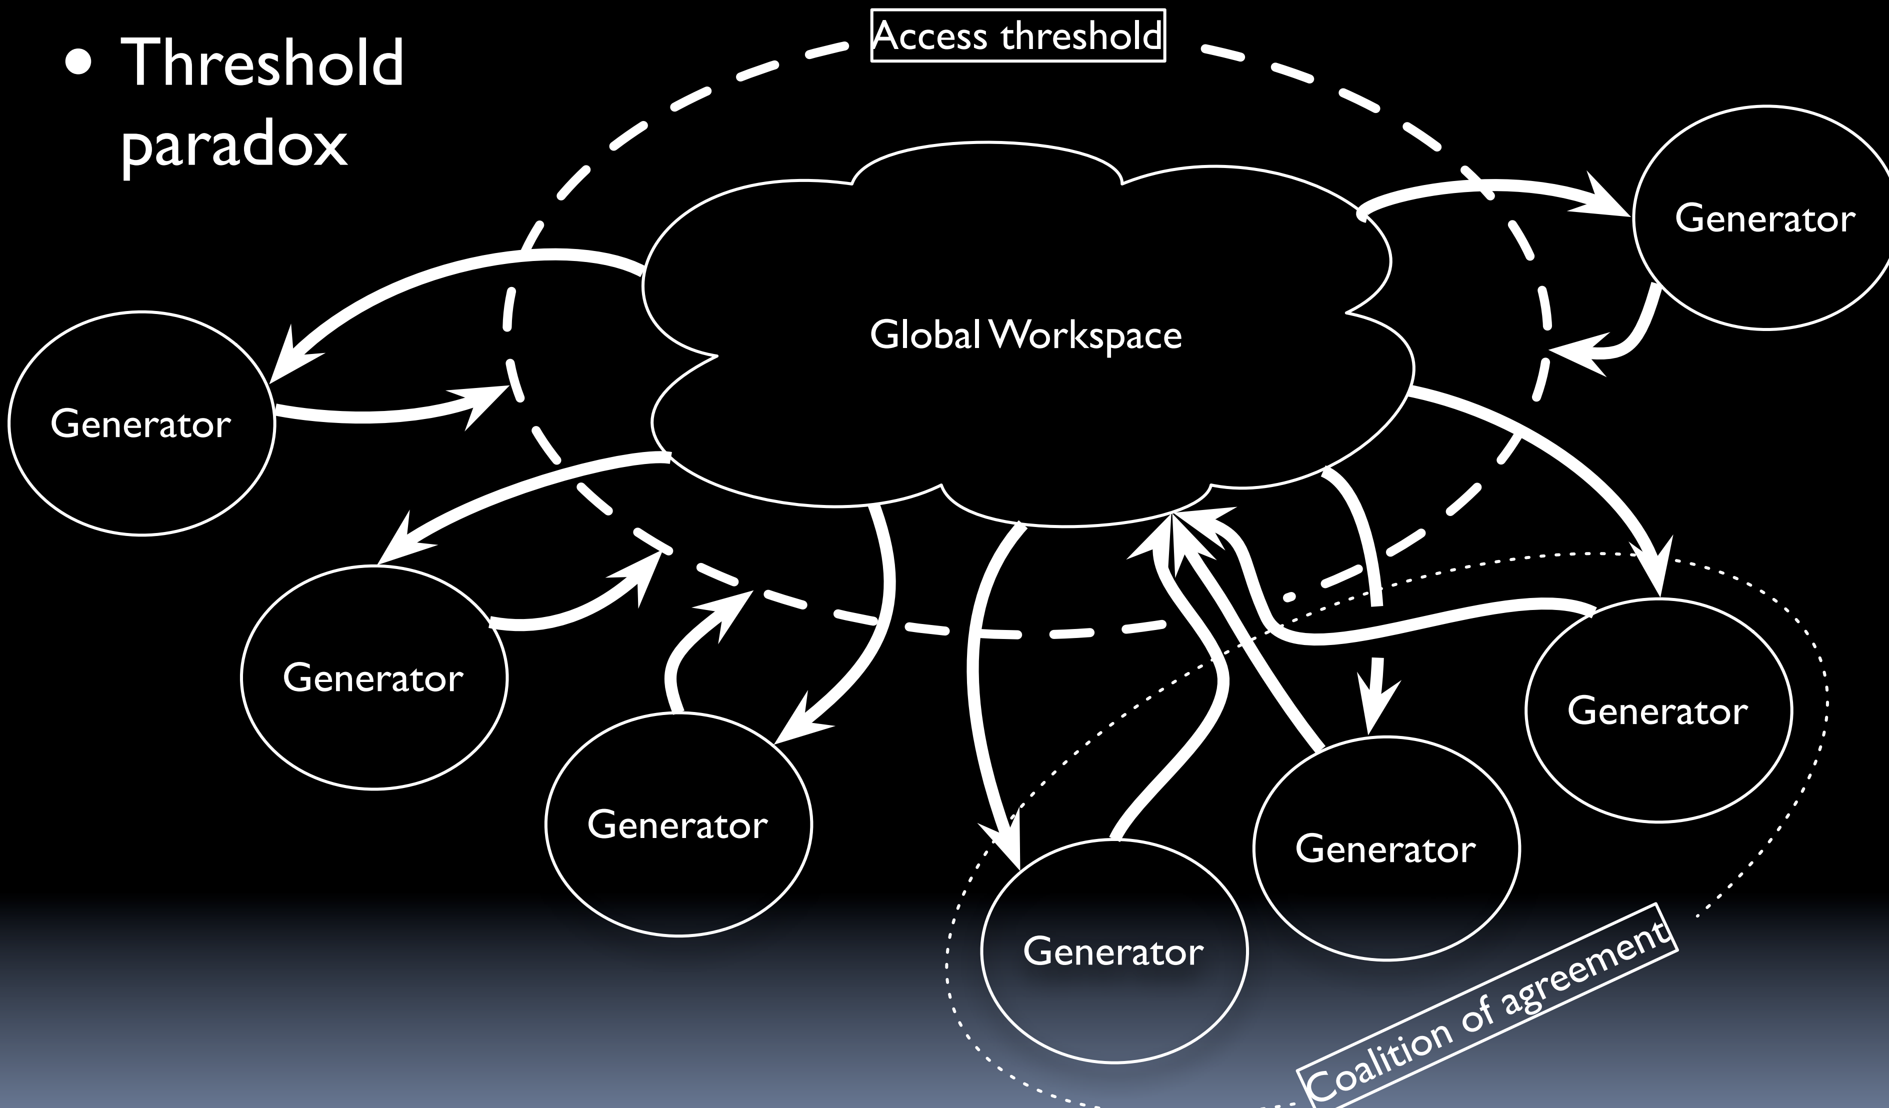

Baars' (1988) Global Workspace Theory

Queen Mary University of London

- "Aha" moment = passage into consciousness
- Threshold paradox

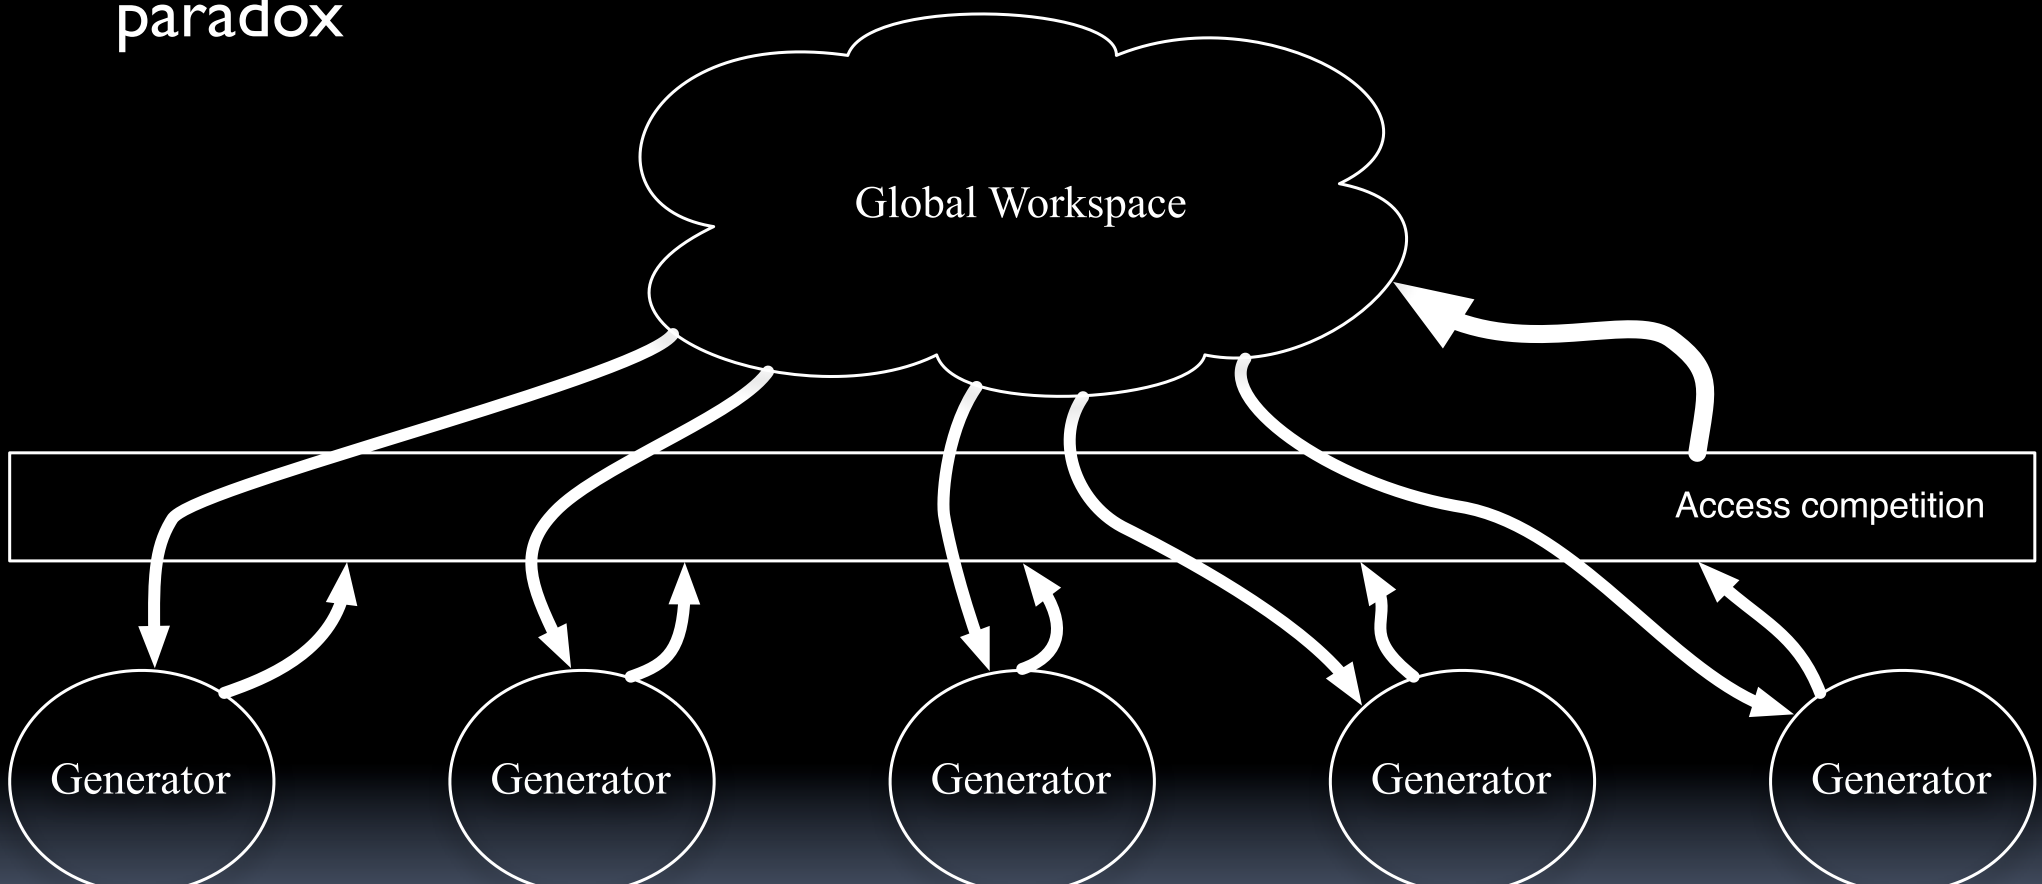

- **Big question: what does "recruitment" mean???**

Anticipatory agent

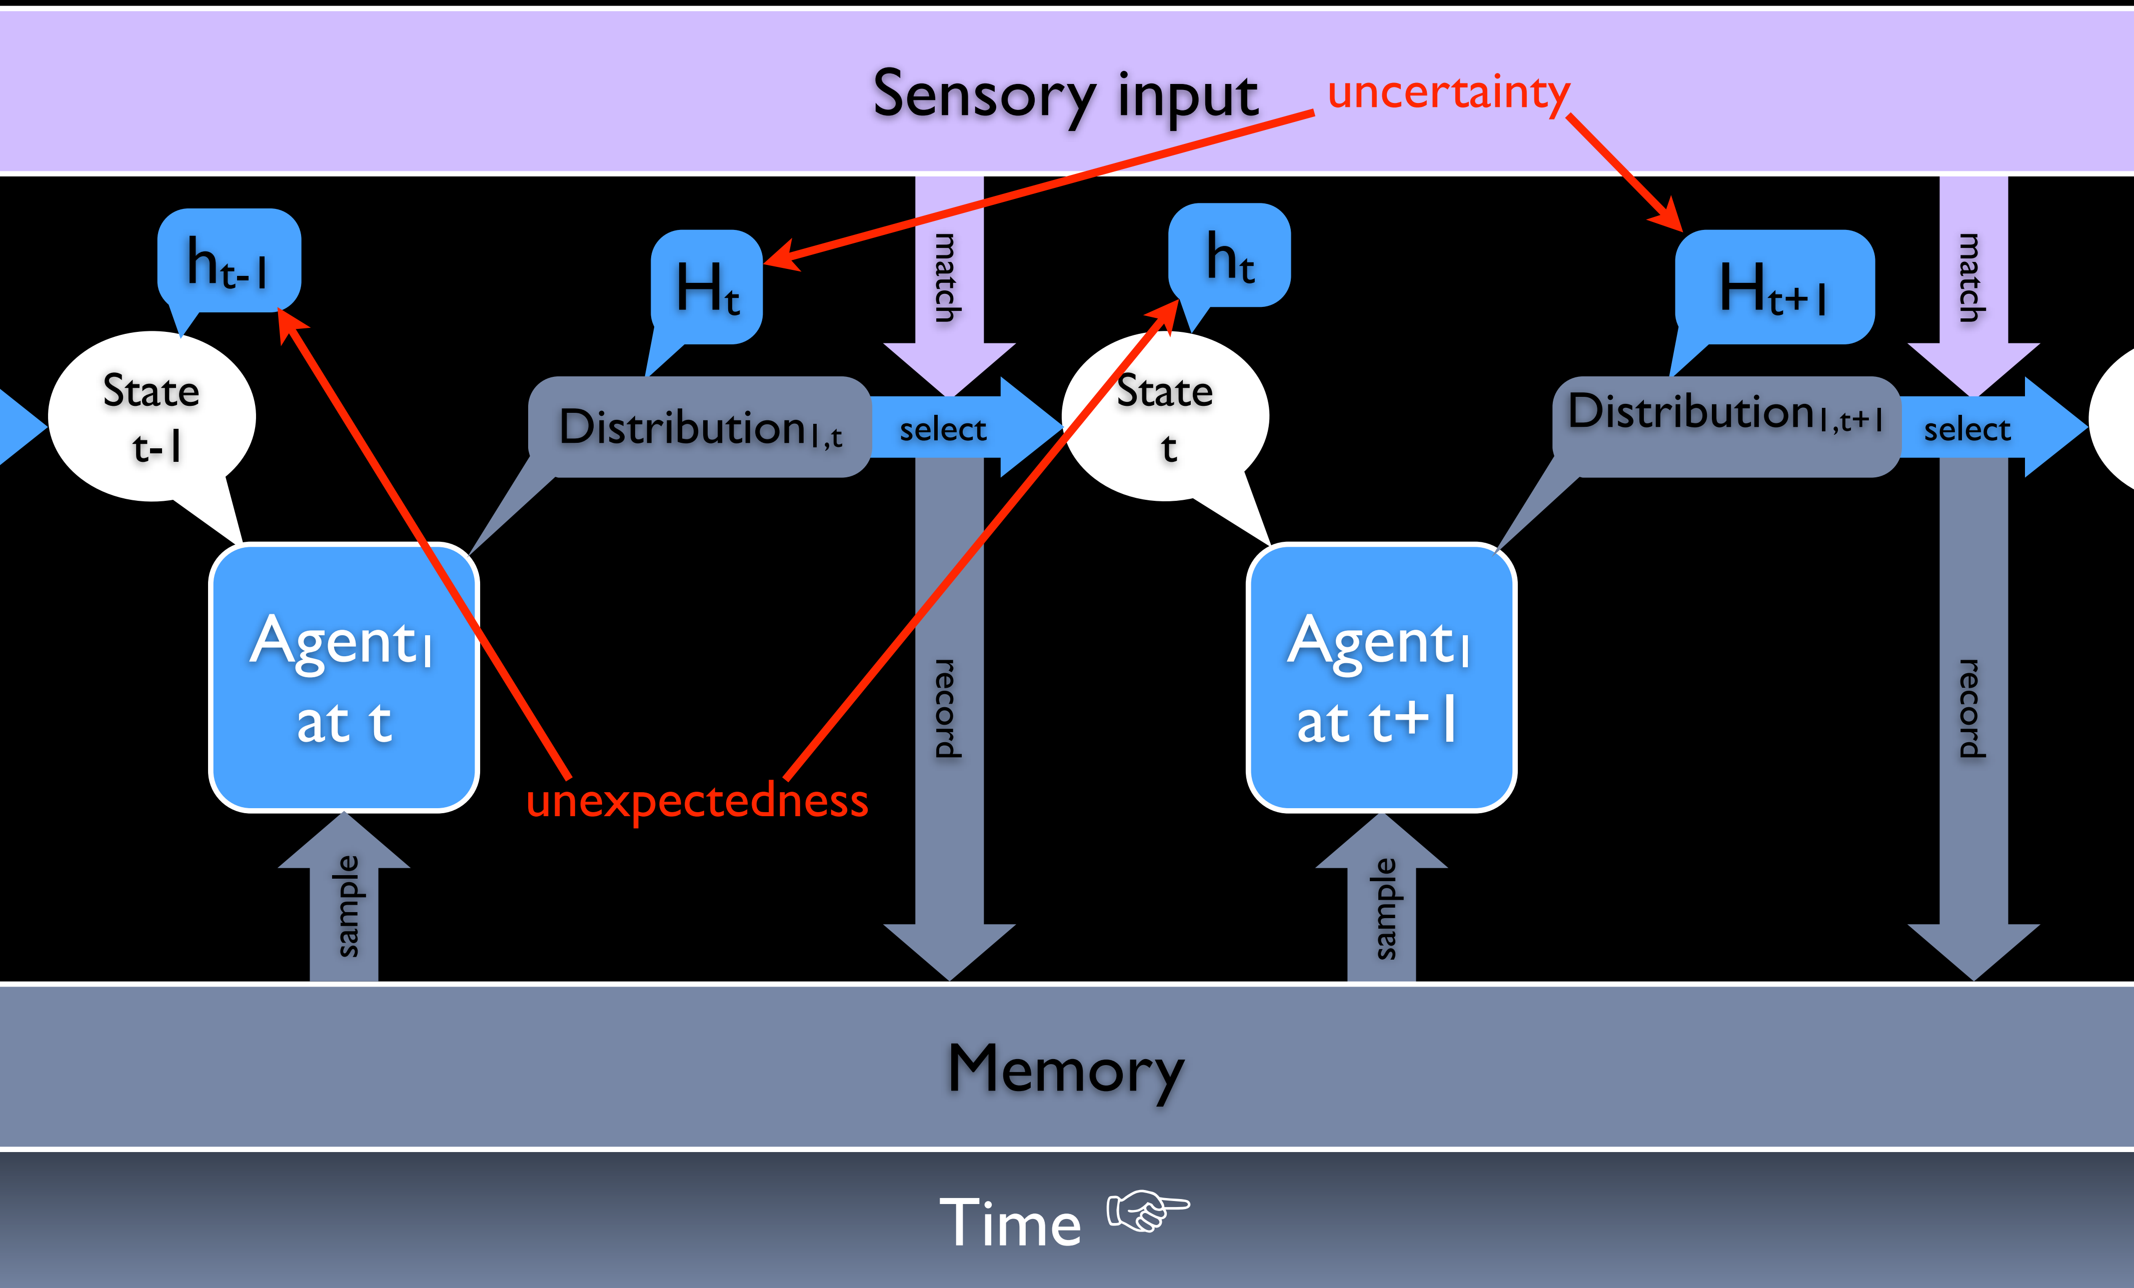

Anticipatory agents

Queen Mary
University of London

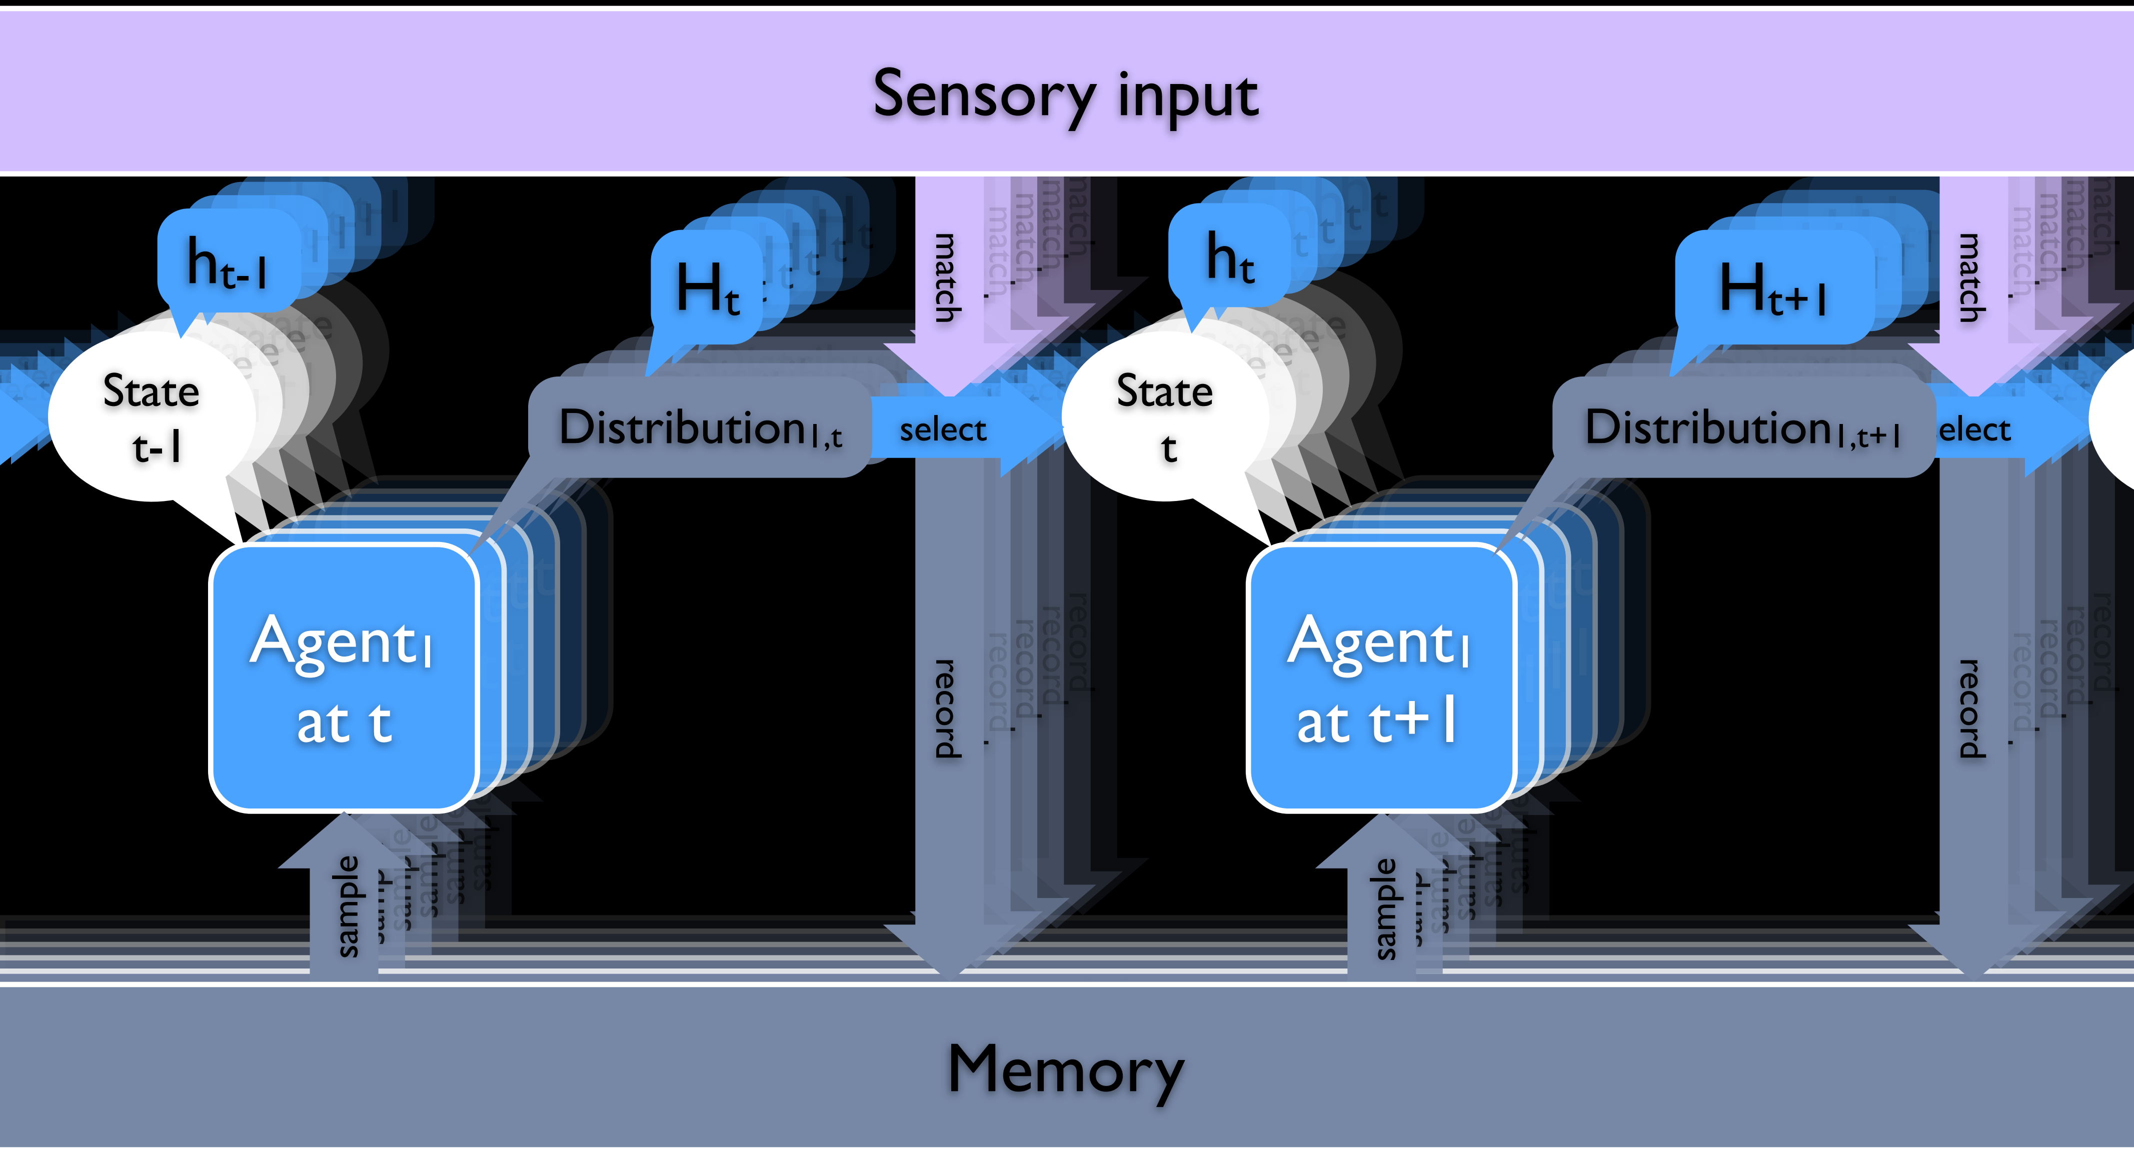

Time ☞

Anticipatory agents in competition

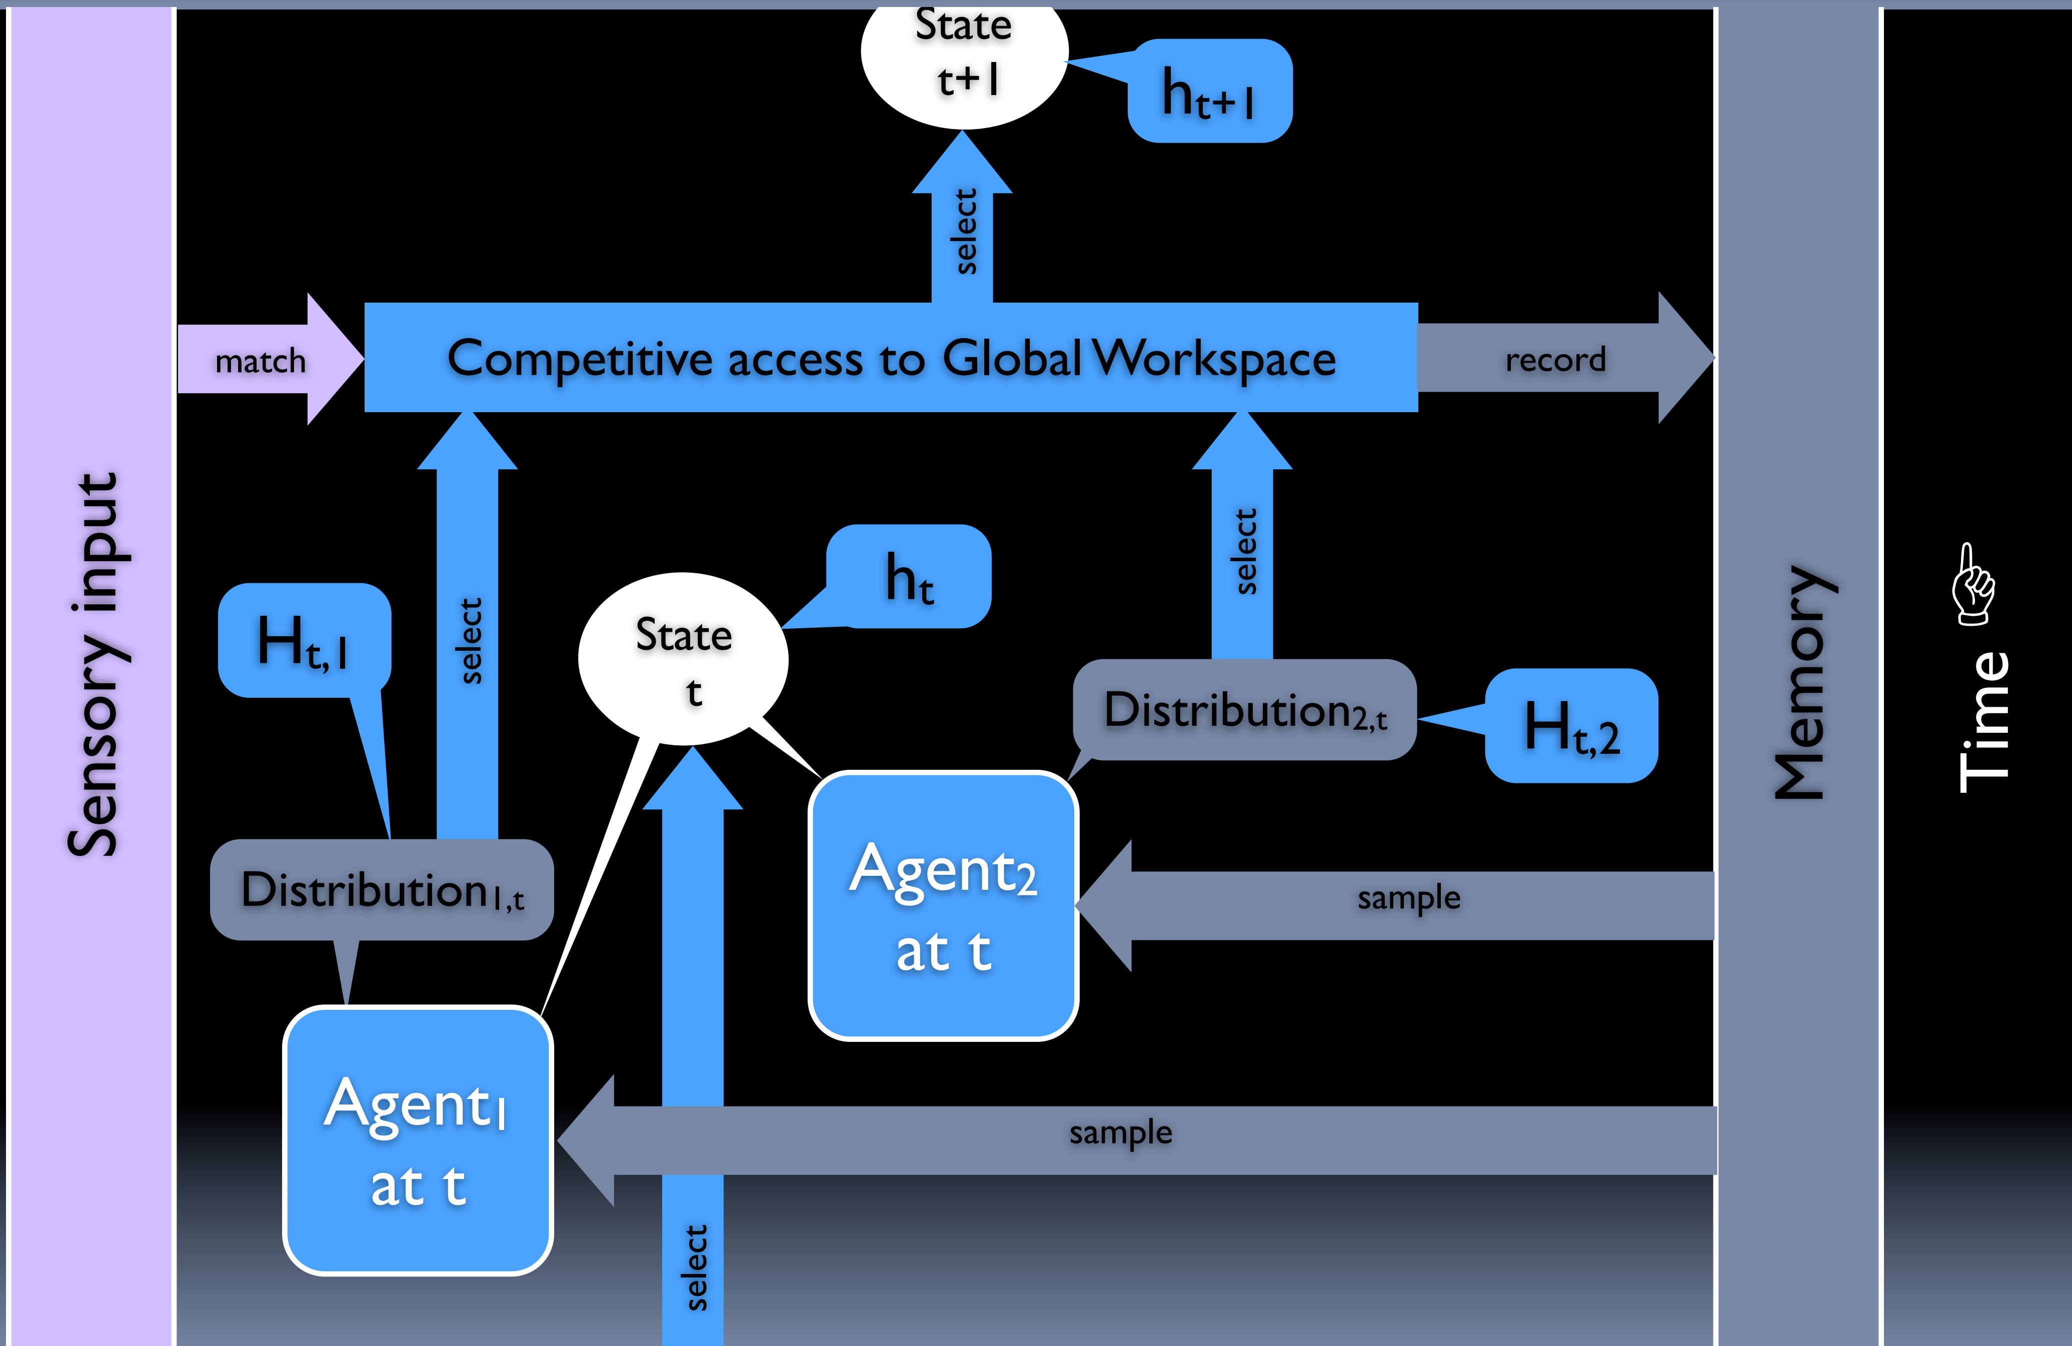

Selecting agent outputs

Competitive access to Global Workspace

- Agents produce (musical) structure representations
- Probability of structure (in learned model) increases "volume"
  - likely structures are generated more often
  - multiple identical predictions are "additive"
  - avoid "recruitment" question in model
    - need fewer agents?
- Unexpectedness increases "volume"
  - information content predicts unexpectedness
- Uncertainty decreases "volume"
  - entropy predicts uncertainty

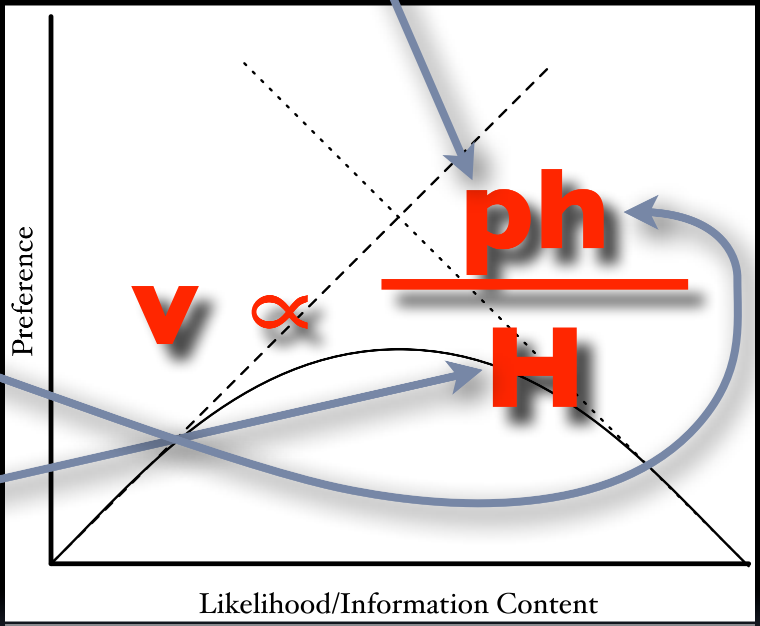

Selecting agent outputs

Competitive access to Global Workspace

- Agents produce (musical) structure representations
- Probability of structure (in learned model) increases “volume”
  - likely structures are generated more often
  - multiple identical predictions are “additive”
  - avoid “recruitment” question in model
    - need fewer agents?

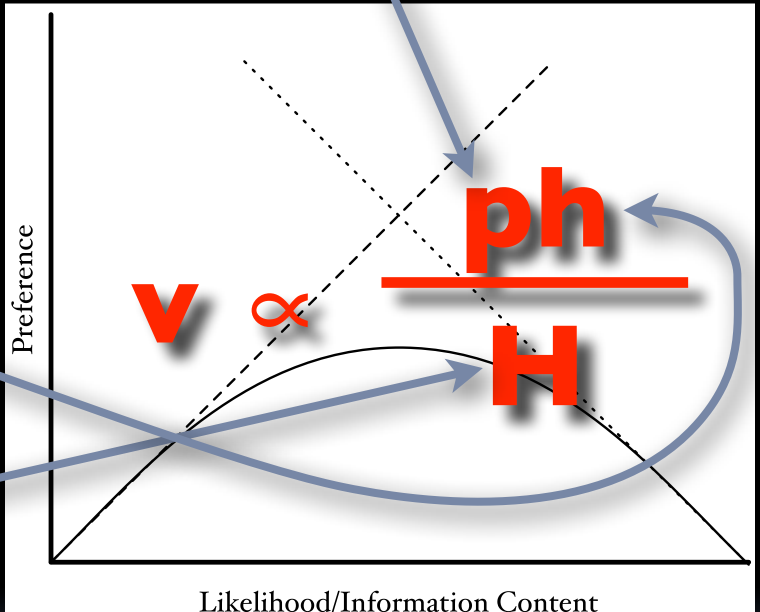

$$v \propto \frac{ph}{H}$$

- Unexpectedness increases “volume”
  - information content predicts unexpectedness
- Uncertainty decreases “volume”
  - entropy predicts uncertainty
- Predictions matched with sensory input, but can compete without it

Spontaneous creativity to order

- In the absence of distracting perceptual input, generators freewheel
- Predictions are produced from memory, spontaneously
- Some may be prioritised enough to enter consciousness as “ideas”
  - cf. Wallas (1926) “illumination”
  - the “Aha!” moment
- Such ideas can be selected...

Where to find more

Queen Mary
University of London

- Full (long) paper:
  - Wiggins, G. (2012) The Mind's Chorus: Creativity before Consciousness. Cognitive Computation. Special issue on Computational Creativity, Intelligence and Autonomy, 4(3):306–319

Final example

Queen Mary
University of London

- Example: harmony by Raymond Whorley's autonomous composer
  - NB statistical model alone - no GW, no feedback, no deep learning

Mozart's explanation (Holmes, 2009)

Queen Mary
University of London

When I am, as it were, completely myself, entirely alone, and of good cheer – say traveling in a carriage, or walking after a good meal, or during the night when I cannot sleep; it is on such occasions that my ideas flow best and most abundantly. Whence and how they come, I know not; nor can I force them. Those ideas that please me I retain in memory, and am accustomed, as I have been told, to hum them to myself.

All this fires my soul, and provided I am not disturbed, my subject enlarges itself, becomes methodized and defined, and the whole, though it be long, stands almost completed and finished in my mind, so that I can survey it, like a fine picture or a beautiful statue, at a glance. Nor do I hear in my imagination the parts successively, but I hear them, as it were, all at once. What a delight this is I cannot tell! All this inventing, this producing takes place in a pleasing lively dream. Still the actual hearing of the toutensemble is after all the best. What has been thus produced I do not easily forget, and this is perhaps the best gift I have my Divine Maker to thank for.

Mozart's explanation (Holmes, 2009)

Queen Mary
University of London

When I am, as it were, completely myself, entirely alone, and of good cheer – say traveling in a carriage, or walking after a good meal, or during the night when I cannot sleep; it is on such occasions that my ideas flow best and most abundantly. Whence and how they come, I know not; nor can I force them. Those ideas that please me I retain in memory, and am accustomed, as I have been told, to hum them to myself.

All this fires my soul, and provided I am not disturbed, my subject enlarges itself, becomes methodized and defined, and the whole, though it be long, stands almost completed and finished in my mind, so that I can survey it, like a fine picture or a beautiful statue, at a glance. Nor do I hear in my imagination the parts successively, but I hear them, as it were, all at once. What a delight this is I cannot tell! All this inventing, this producing takes place in a pleasing lively dream. Still the actual hearing of the toutensemble is after all the best. What has been thus produced I do not easily forget, and this is perhaps the best gift I have my Divine Maker to thank for.