

Learning and Creativity in the Global Workspace

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

- An account of the **phenomenon of mind** which is
 - ▶ mechanistic
 - ▶ empirically supported
 - ▶ credible from an evolutionary perspective
- Uniform **generalist** approach
 - ▶ cf. approaches that focus on specific problems in isolation
- Account for
 - ▶ reasoning
 - ▶ imagination
 - ▶ **creativity**

- In the cognitive context, we must be able to tell a story of how something could evolve
- **Prediction** and **Learning** are the central concepts here
 - ▶ an organism which can **predict** is better able to manage a dynamic world than one that cannot
 - ▶ an organism which can **learn** is better able to predict a dynamic world than one that cannot
- Consequence: **extreme Ockham!**
 - ▶ always choose simplest mechanism to account for an effect, unless there is a reason not to

Requirements for predictive success


- Learning
 - ▶ perceptual inputs
 - ▶ memory
 - ⊙ representations
- Generation
 - ▶ predictions
 - ⊙ selection mechanism
- Validation/feedback
 - ▶ determine whether predictions were useful
 - ⊙ NB. death or serious damage is **not helpful** as a negative example

- Concomitant requirements (hypothetical)
 - ▶ optimisation of memory vs. representation
 - ▶ episodic segmentation of input
 - ▶ regulation of attention (= processing power)
 - ◎ limited, expensive (in energetic and evolutionary terms)

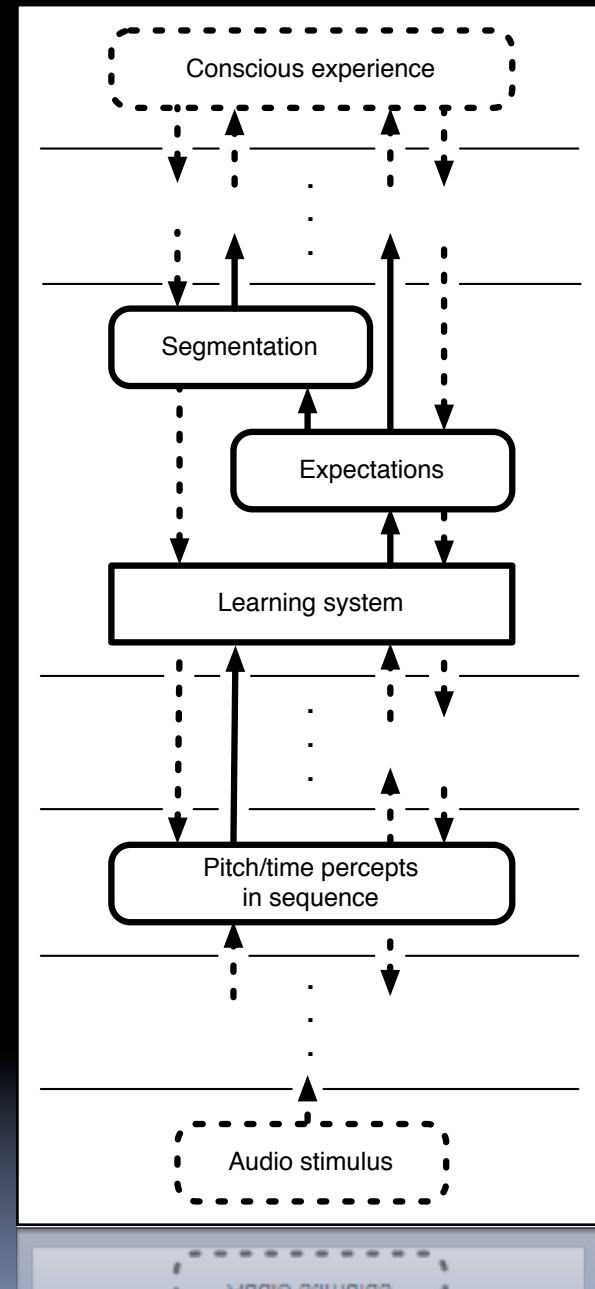
- Consider an organism that has evolved
 - ▶ to predict the next state of the world
 - ▶ to compare the current state of the world with its expectation
- Emotional response to unresolved expectation
 - ▶ unexpected stimulus leads to enhanced arousal
- More generally, tension and, in extremis, anxiety, fear
 - ▶ Anxious organism = cautious organism = safe(r) organism
 - ▶ Anxious organism = organism ready for f(l)ight = safe(r) organism

Expectation in music

- Emotional response to unresolved expectation
 - ▶ musical tension



- Middle layer of cognitive model of conscious musical experience
- Designed by Marcus Pearce (2005)
- 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



- 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:

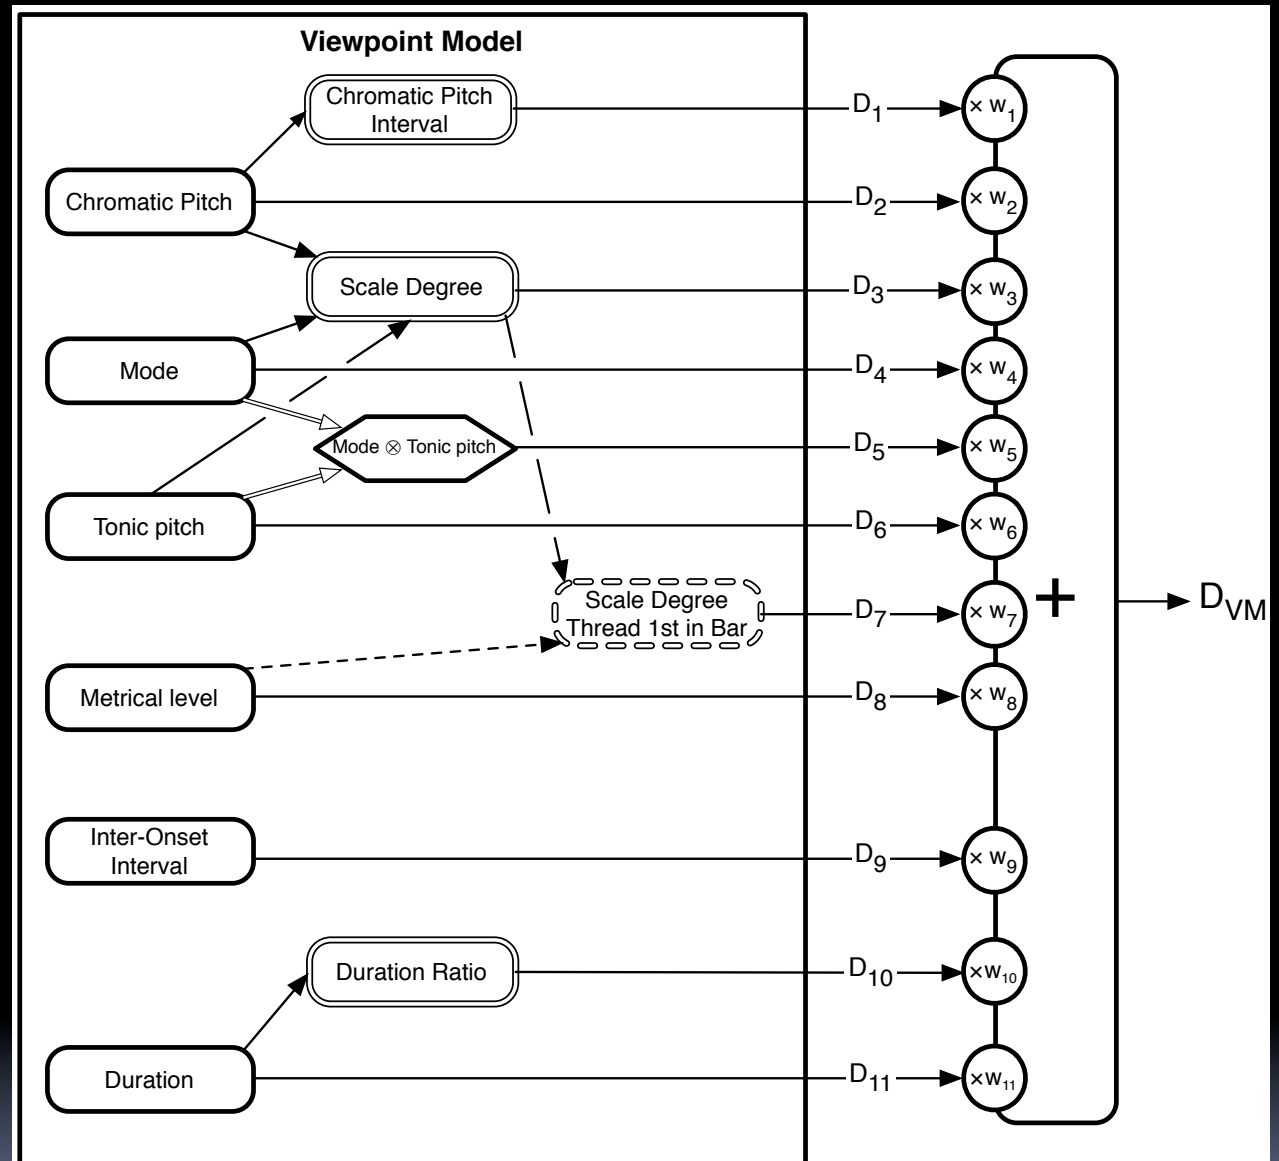
$$h = -\log_2 p_s$$

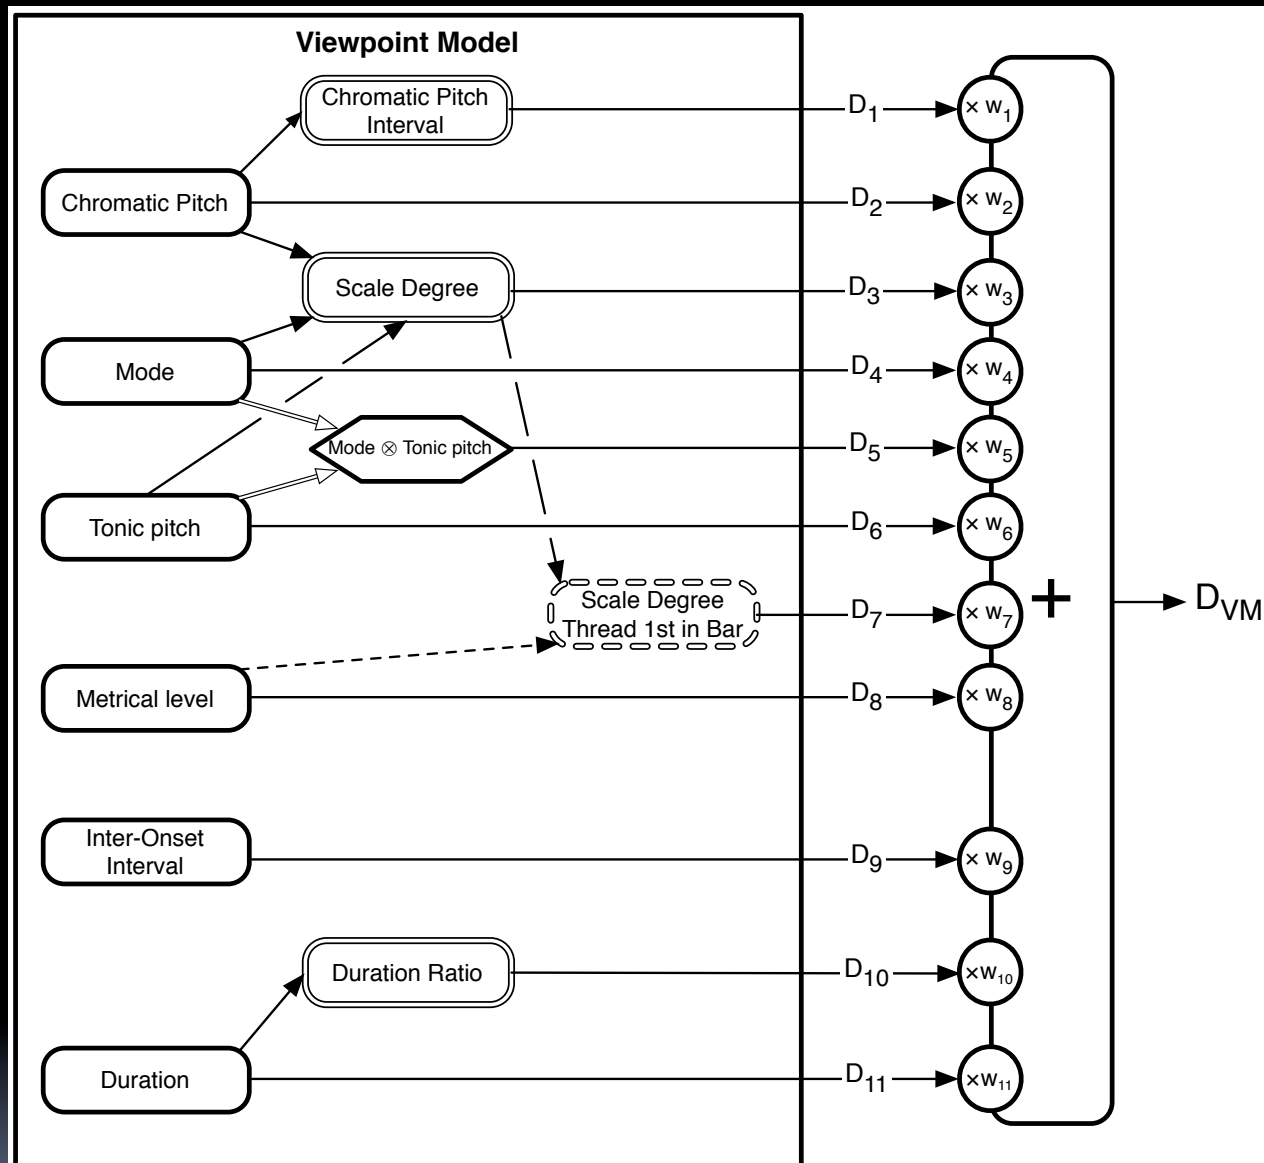
- ◉ models **unexpectedness**
- ▶ *entropy*: expected value of the number of bits required to transmit a symbol from a given distribution, prior to sending/receipt:

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

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

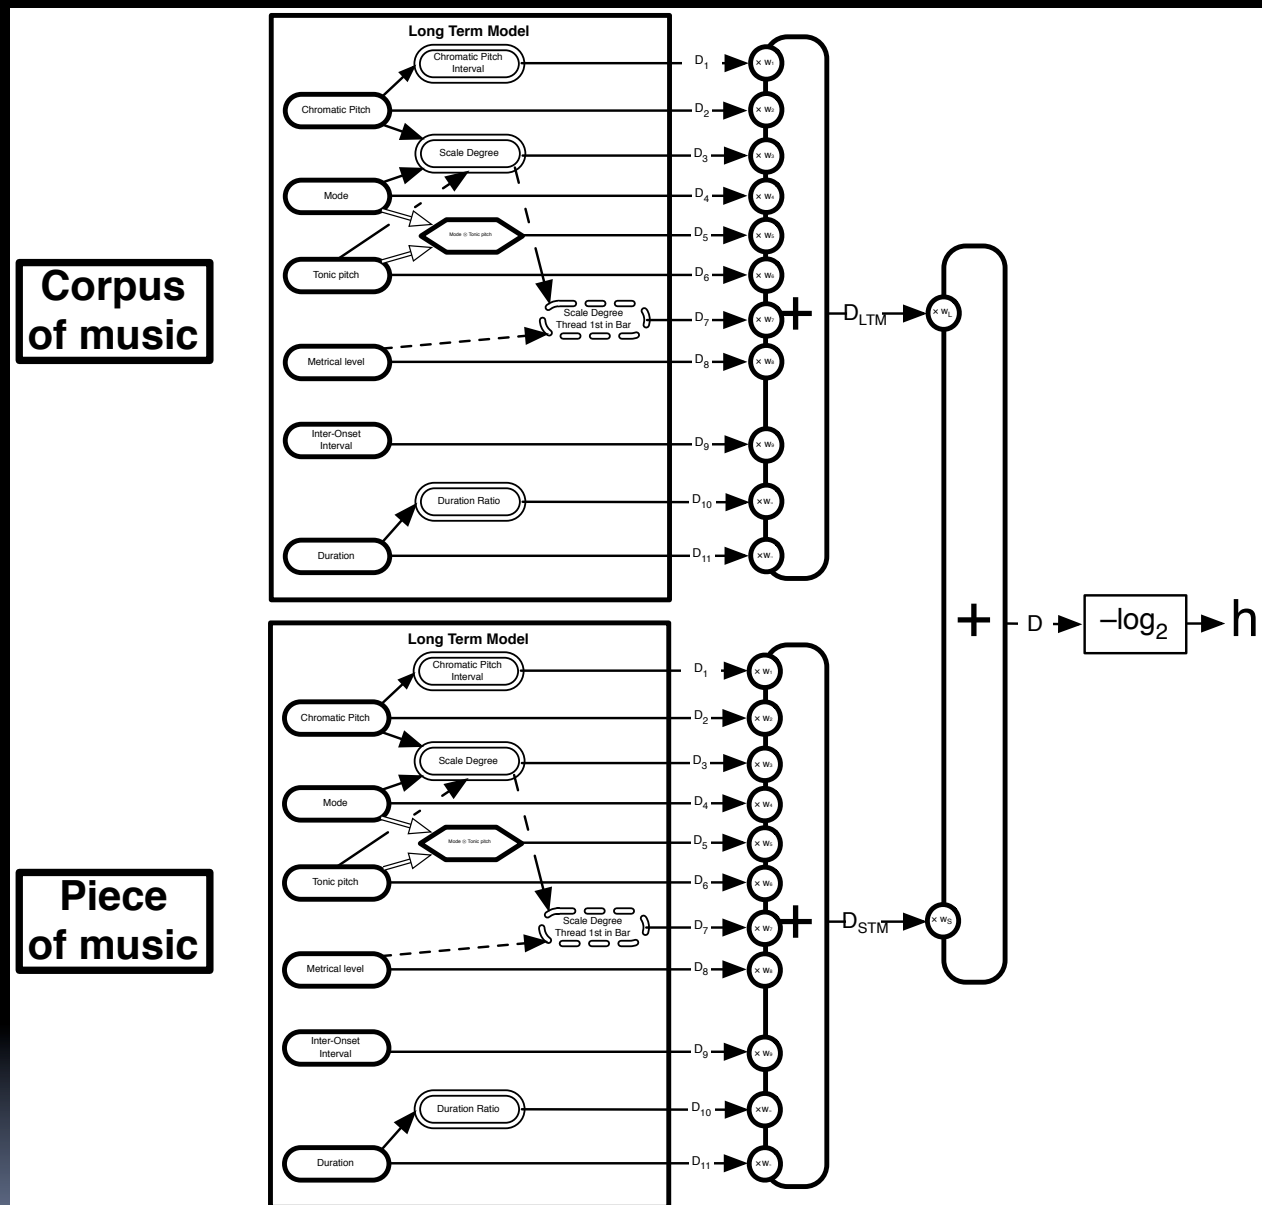
- Efficient implementation of simple Markov chains
 - ▶ but with multi-dimensional symbols
 - ▶ select feature sequences (*viewpoints*)
 - basic
 - derived
 - ▶ calculus of viewpoints
 - differentiation (δ)
 - cross-product (pairing)
 - thread (sub-sequence selection)





- 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

- 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; $\kappa = 0.58$
 - vs musicologist judgements



Information Dynamics of Music

- IDyOM predicts

- ▶ listener's expectations of next note in melody


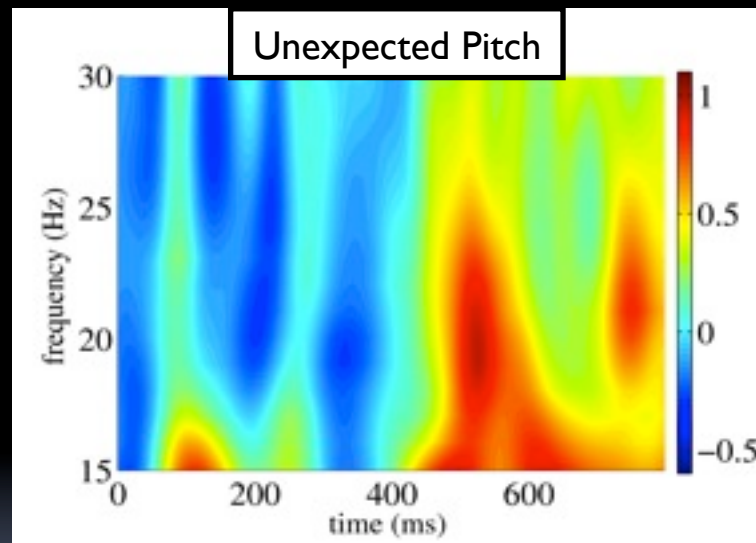
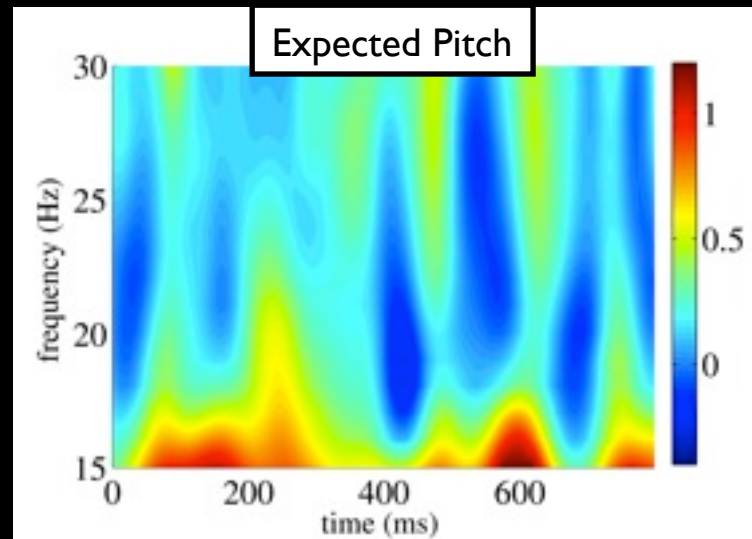
- 4 studies; up to $r=.91$ correlation
- 1 study; very high correlation with musicologists' predictions

- ▶ melodic segmentation


- 2 studies; $\kappa = 0.58$
- vs musicologist judgements

- ▶ neural activation with unexpectedness

- centro-parietal region
- strong sync. in beta-band

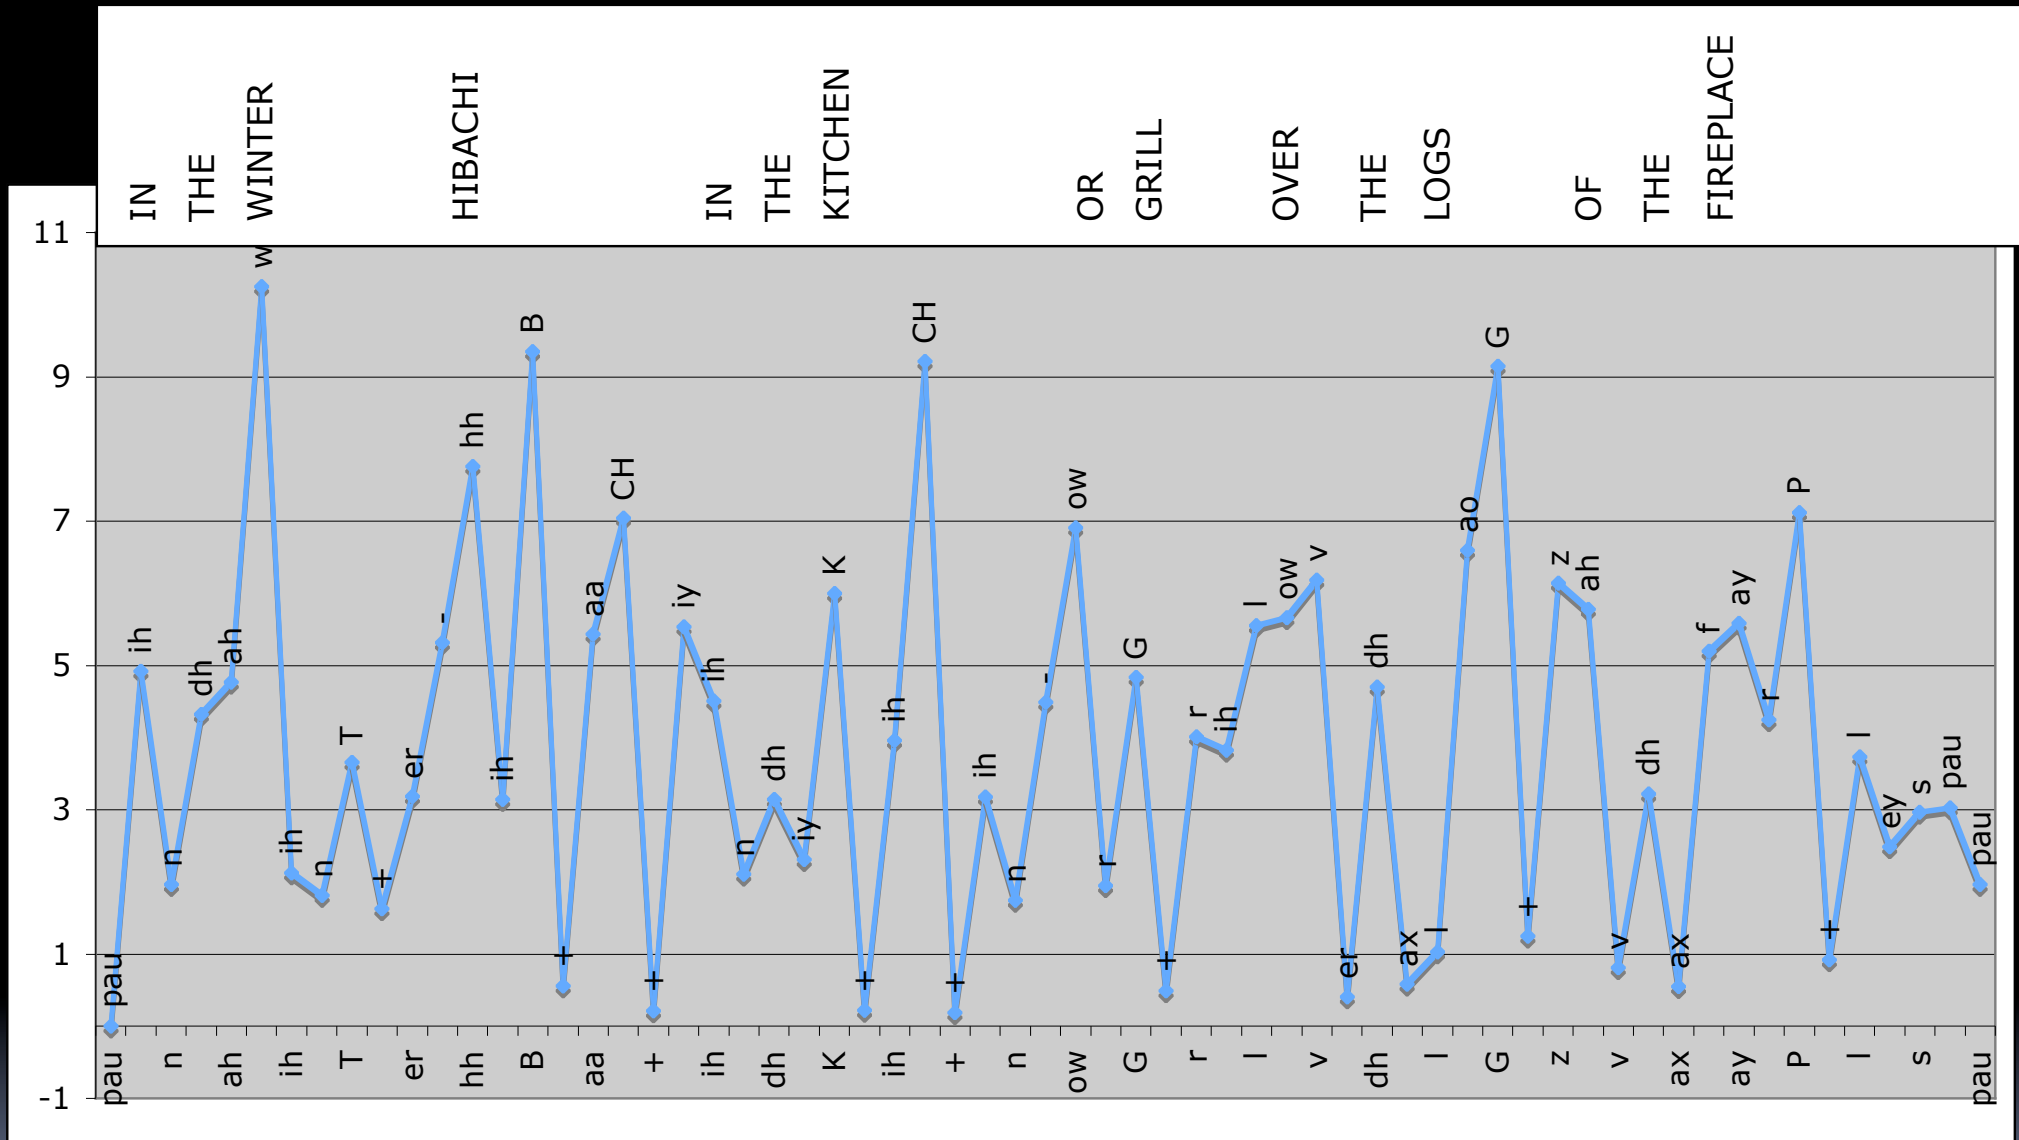


Unexpected – Expected



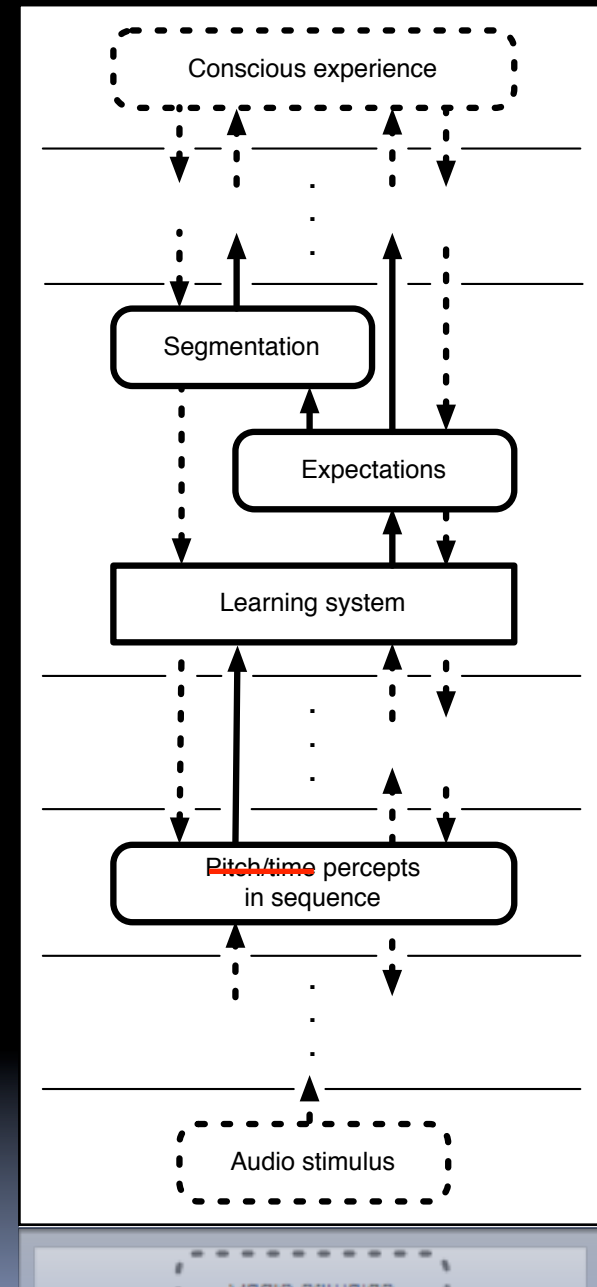
- Replace pitch/time representation with
 - ▶ phonetic symbols (quasi IPA)
 - ▶ stress symbols (none, weak, strong)
 - ▶ try sentence segmentation – expect morphemes or syllables
- Expose model to TIMIT meta-data
 - ▶ intended for audio-speech recognition, but just use comparison data
 - ▶ compare results with syllable segmentation supplied with TIMIT

Speech segmentation



Predictive cognition in general

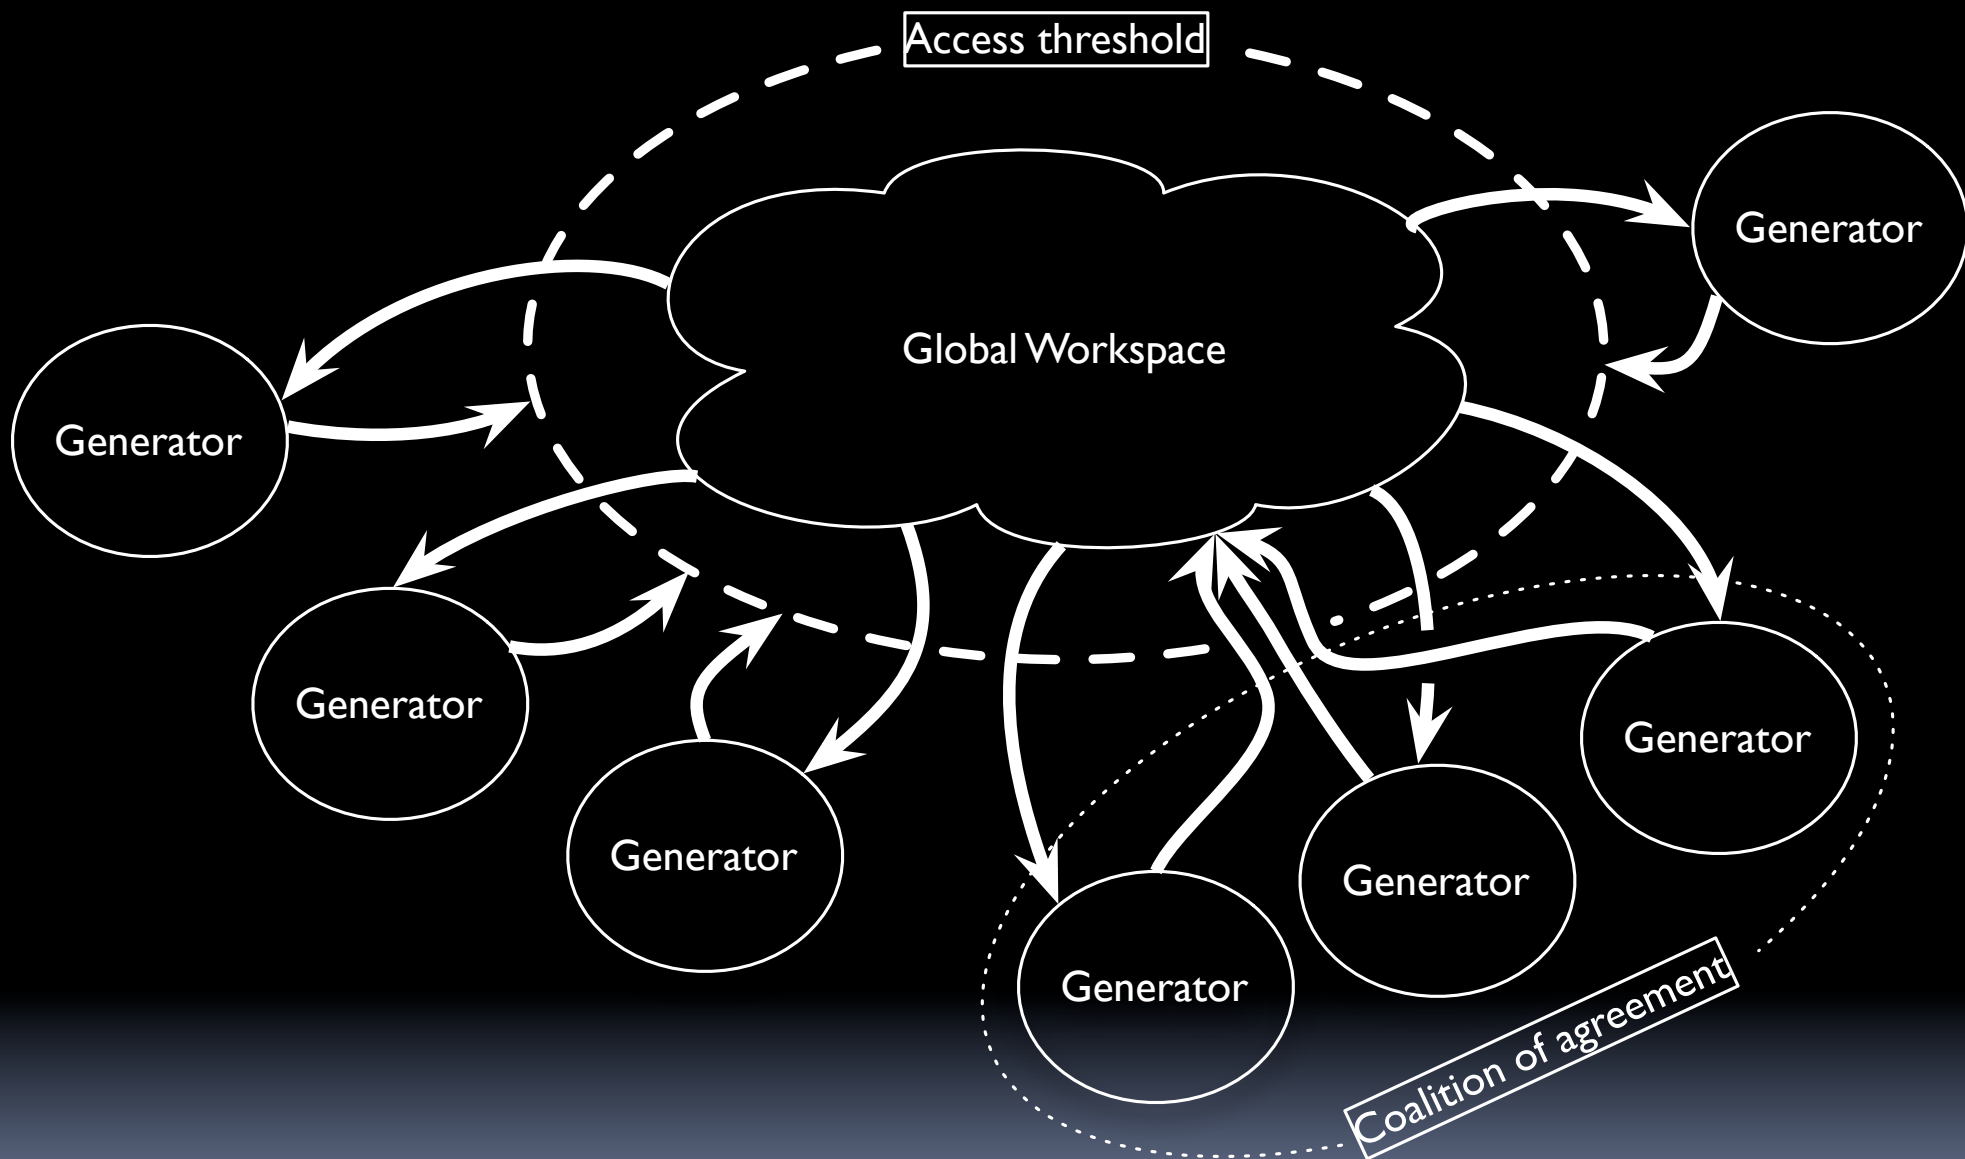
- Information theoretic properties of data determine perceptual segmentation
 - ▶ future work: *exactly* how?
- Information theoretic properties of data determine cognitive representation
 - ▶ future work: *exactly* how?
- Need a mechanism by which predictions are made and regulated
 - ▶ **Global Workspace Theory**
 - ▶ cf Hippolyte Taine (**1871!!**)



Baars' (1988) Global Workspace Theory



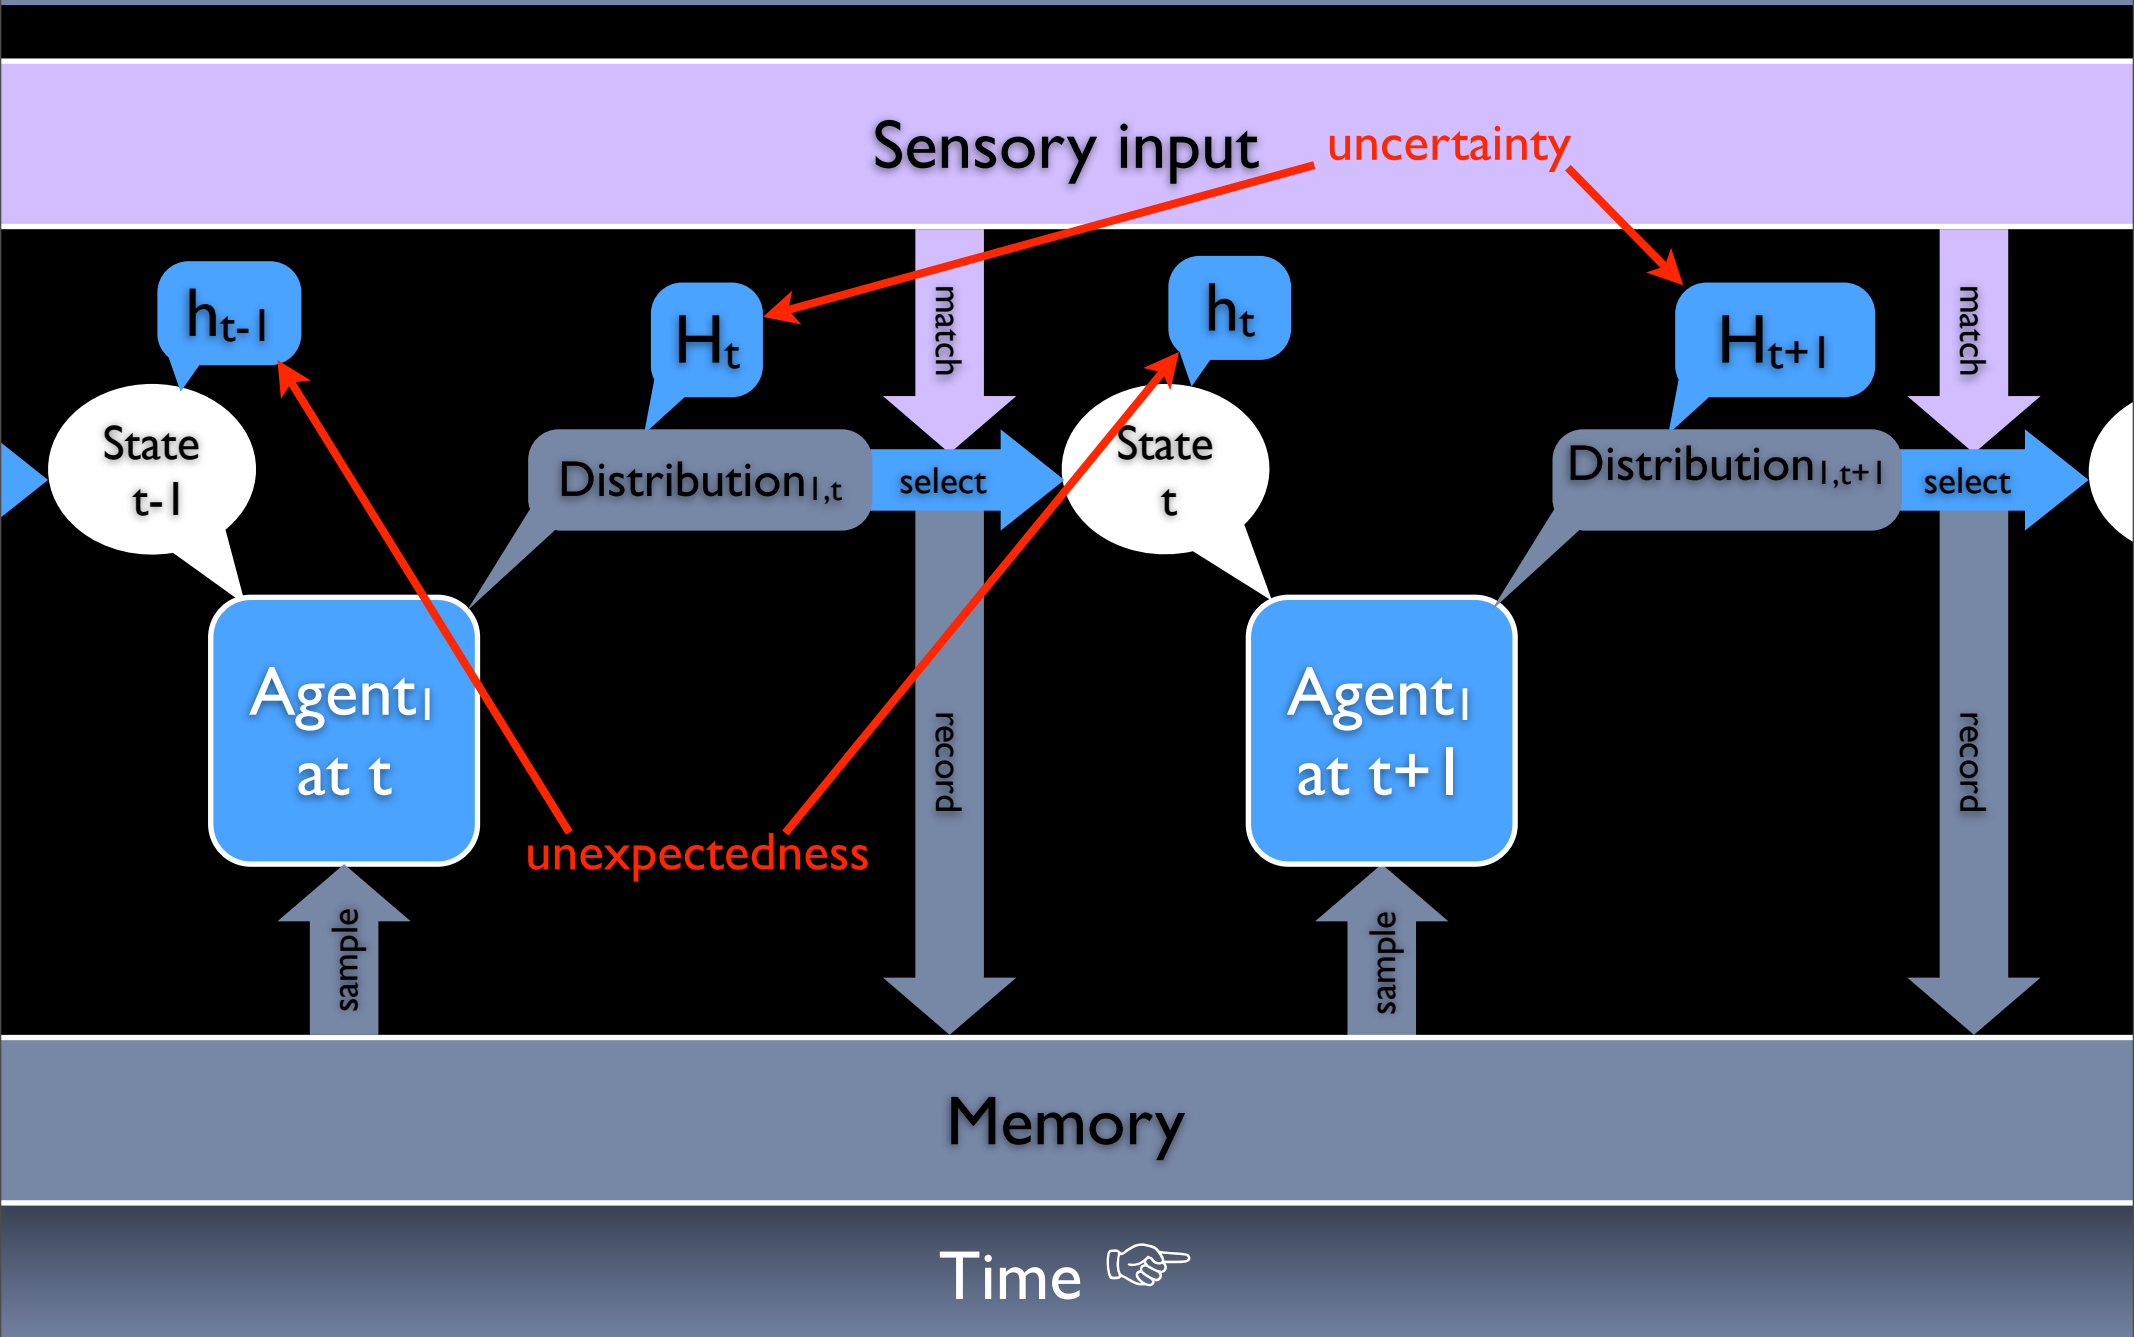
Queen Mary
University of London



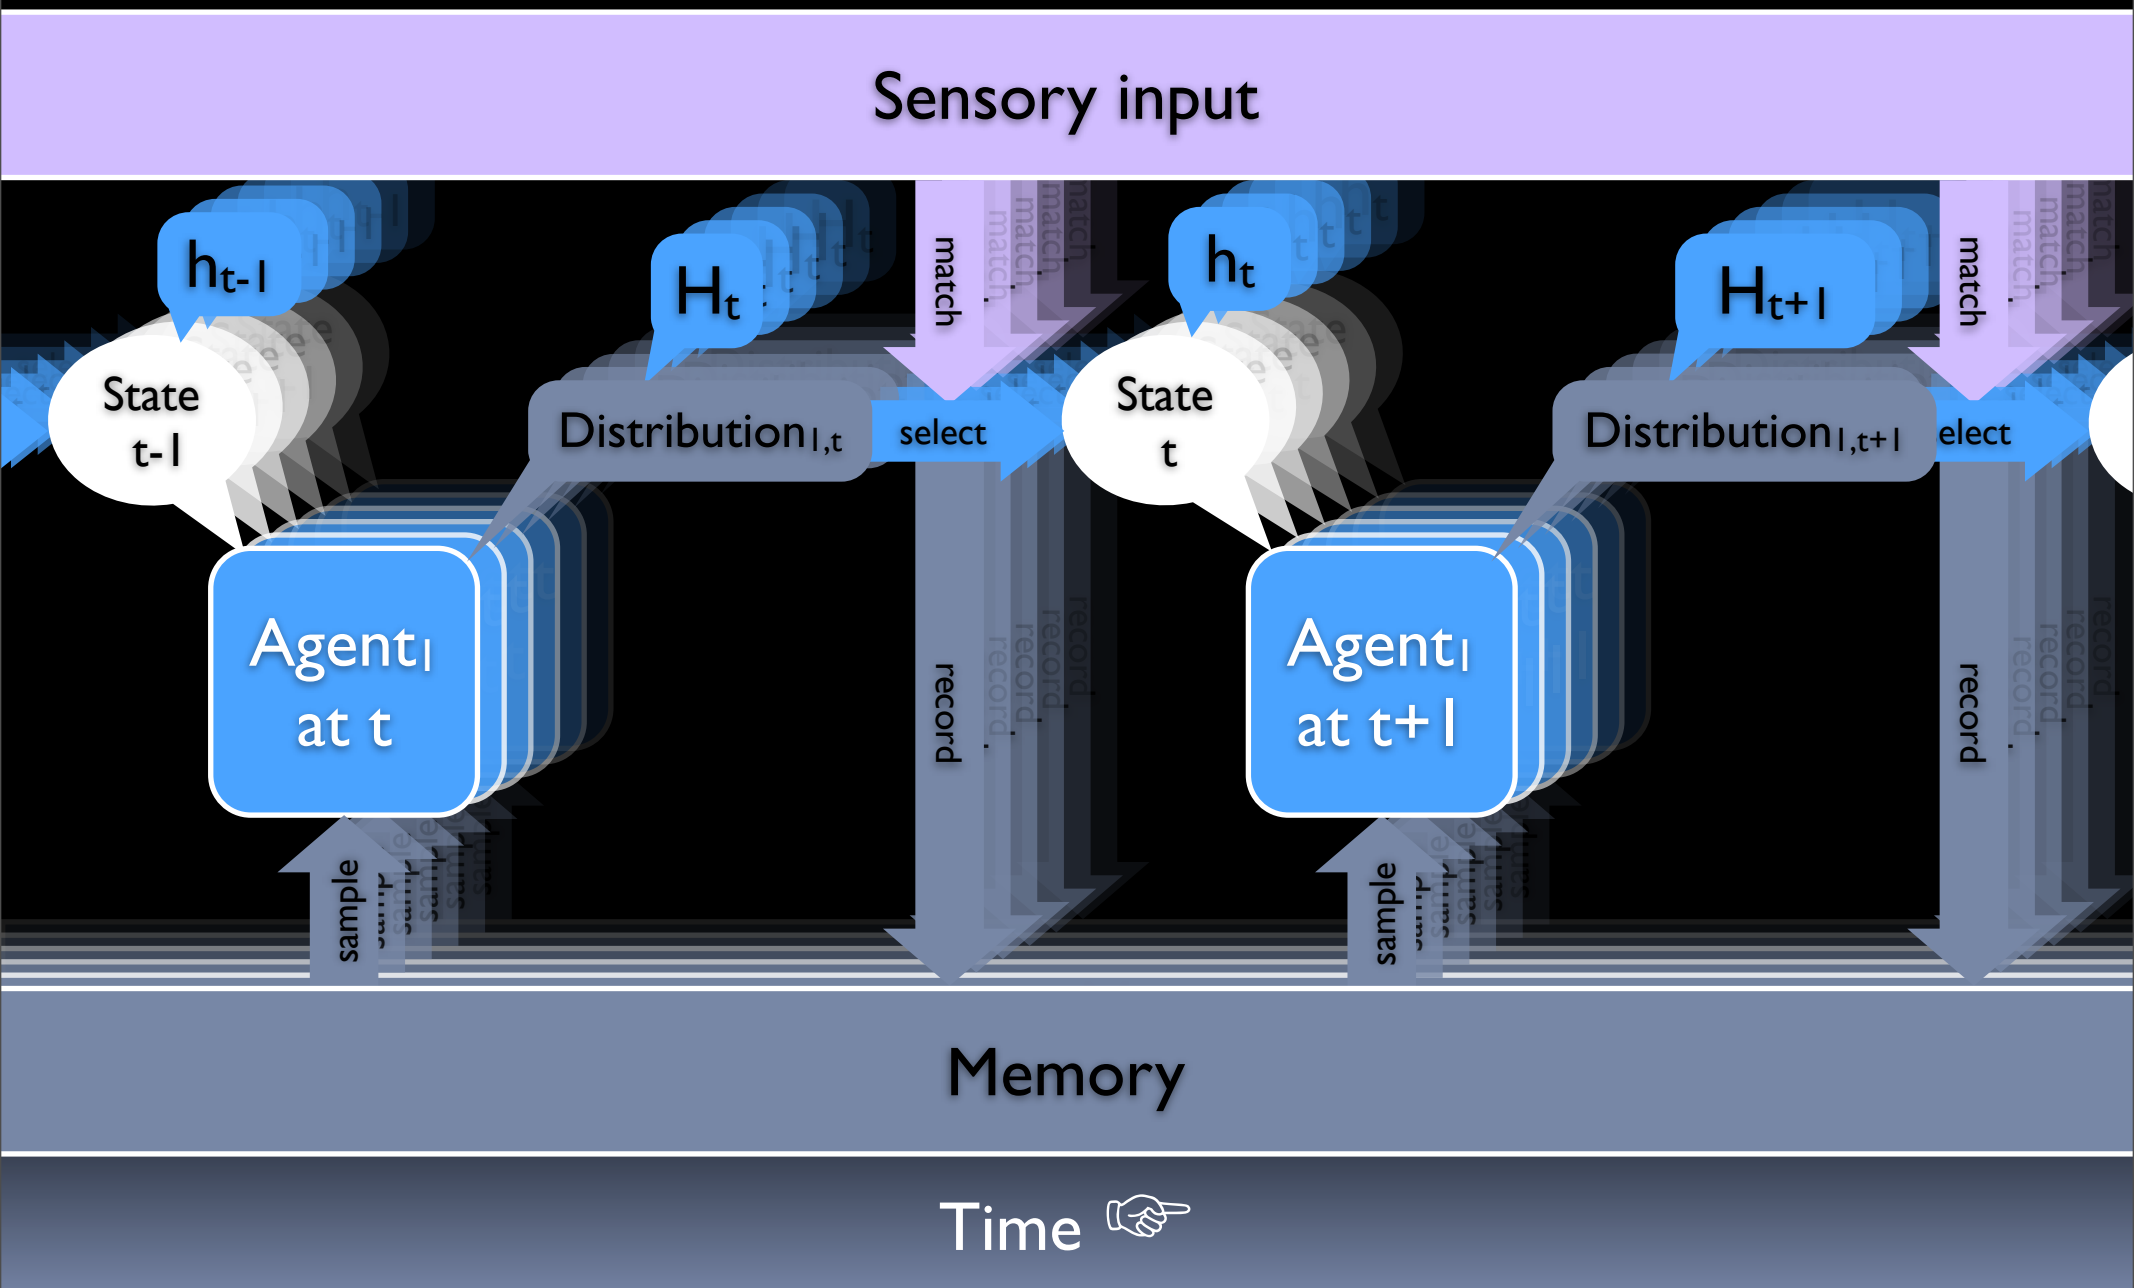


- Compare:
 - ▶ Multiple-agent system communicating via blackboard
 - ▶ Genetic algorithm (breeding replaces blackboard)
 - ▶ Minsky-type Society of Mind (hierarchical management replaced blackboard)
 - ▶ others...
- Neuroscientific evidence for multiple-generator idea
 - ▶ “idling noise” in brain
 - ▶ freewheeling generation when consciousness disengaged
 - ◎ sleeping = dreams
 - ◎ high/pathological = hallucinations

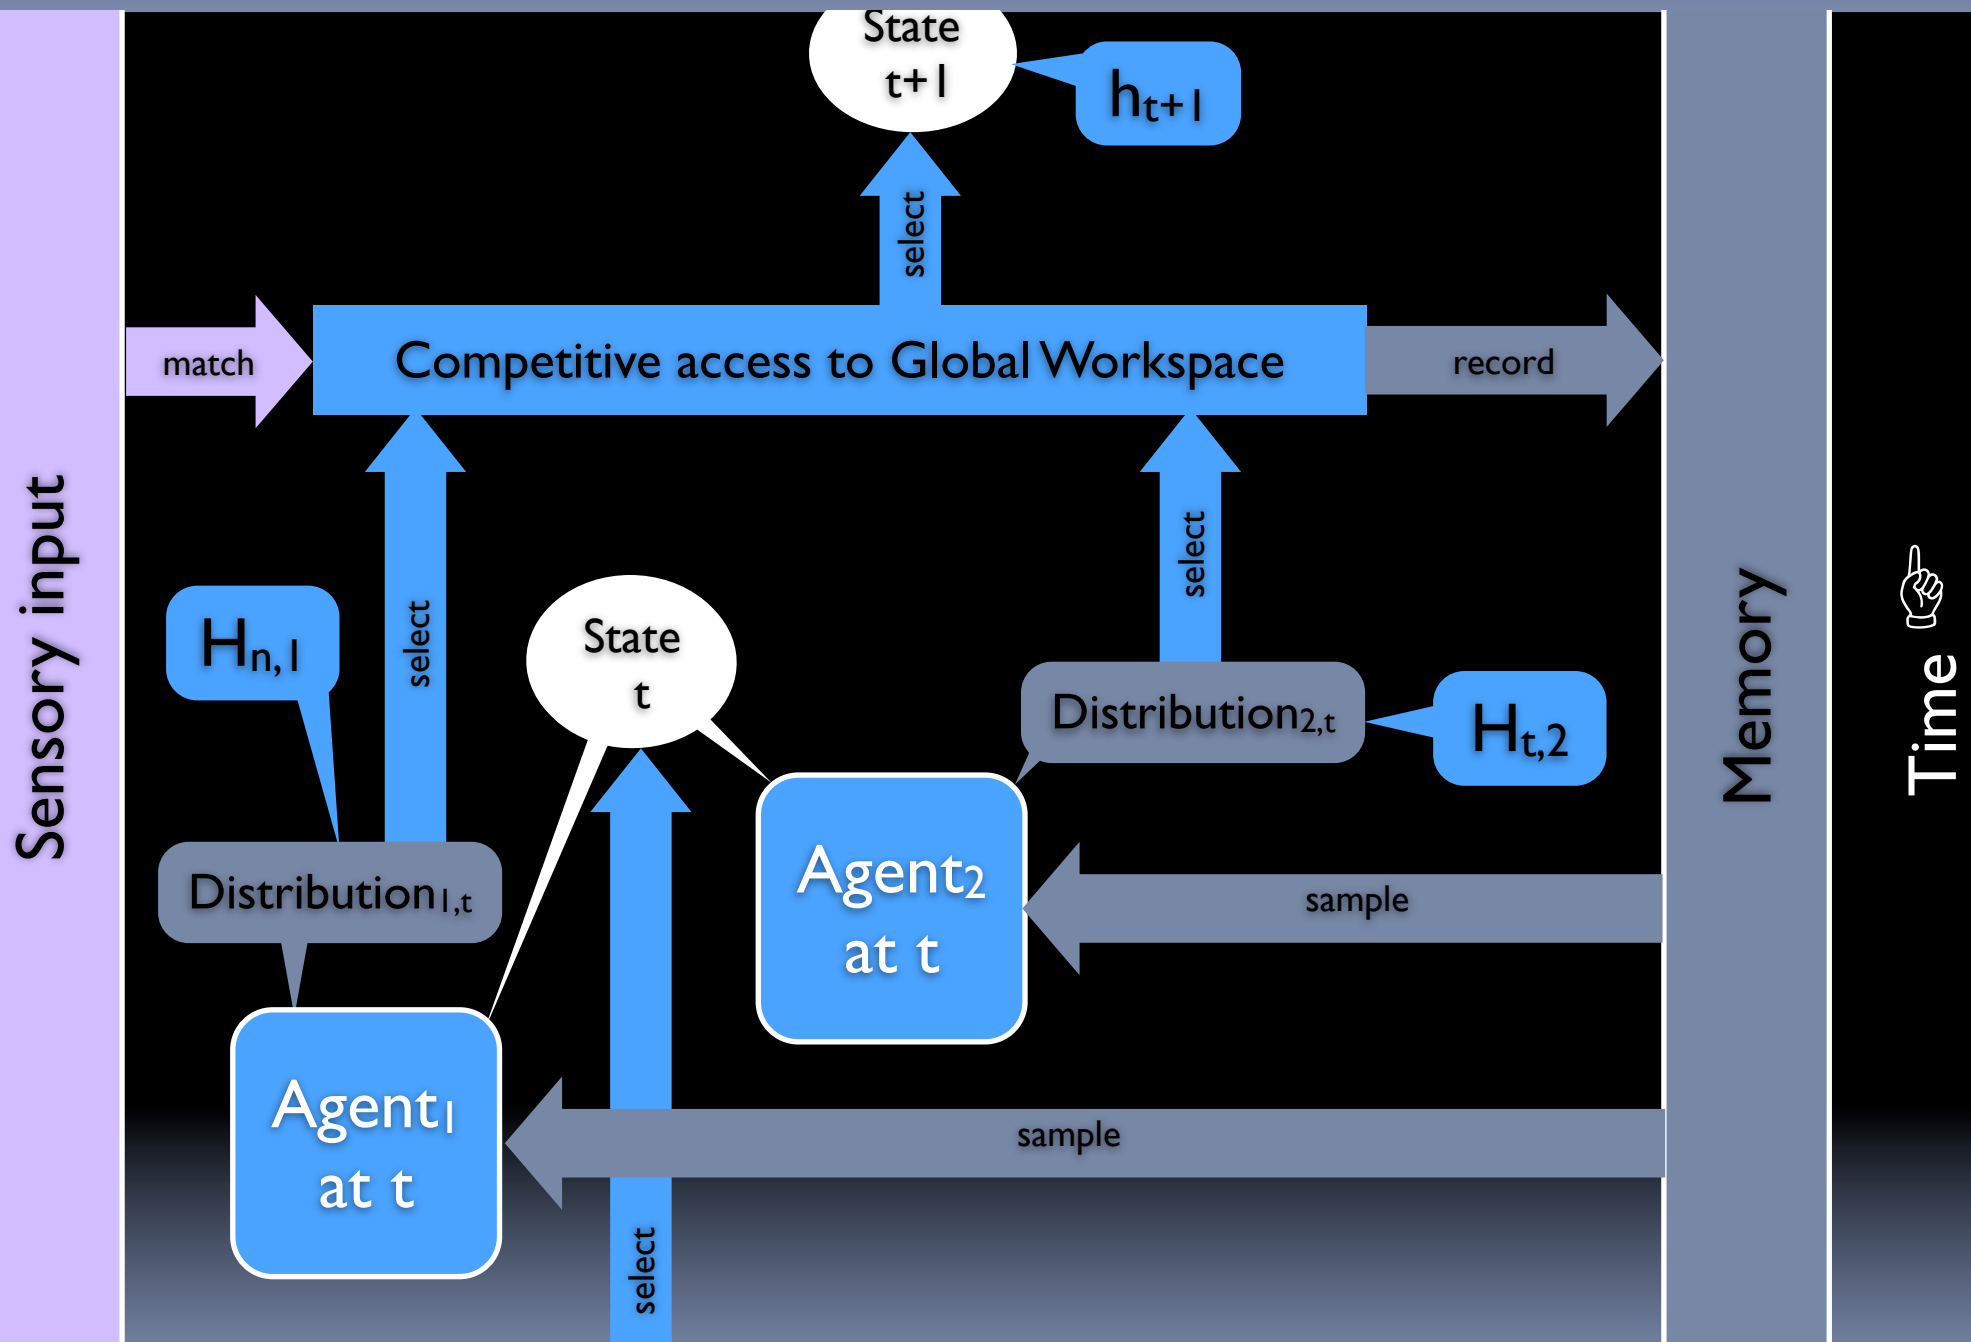
Anticipatory agent



Anticipatory agents




Anticipatory agents in competition



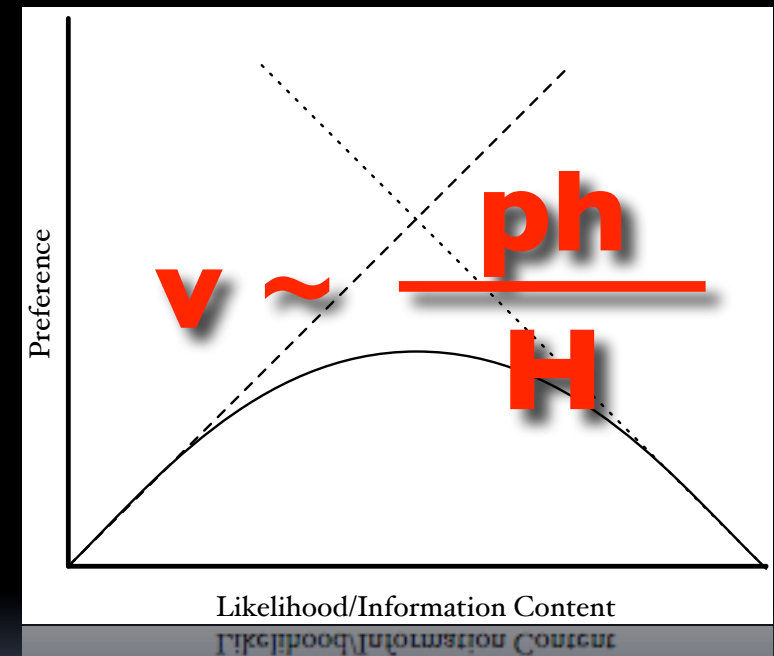
Competitive access to Global Workspace

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




Competitive access to Global Workspace

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



Competitive access to Global Workspace

- Agents produce (musical) structure representations
- Probability of structure (in learned model) increases priority
 - ▶ likely structures are generated more often
 - ▶ multiple identical predictions are “additive”
 - ▶ avoid “recruitment” question in model
 - ◉ need fewer agents?
- Unexpectedness increases priority
 - ▶ information content predicts unexpectedness
- Uncertainty decreases priority
 - ▶ entropy predicts uncertainty
- Predictions matched with sensory input, but can compete without it



- 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...

Mozart's explanation (Holmes, 2009)

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.

- Creativity is a slippery concept in humans
 - ▶ how can we evaluate the model?
- Doing this with music is in a sense easier than with language or other kinds of knowledge
 - ▶ no real-world inference necessary
- **Build the beast and see what it does!**
 - ▶ does it produce novel and interesting (musical) ideas?
 - ▶ does its behaviour match human behaviours,
 - ◎ directly?
 - ◎ tangentially?
 - ▶ use Amabile's (1999) Consensual Assessment Technique to assess creativity and quality of outputs

Where to find more

- Full (long) paper published this week:
 - ▶ Wiggins, G.A. (2012) The Mind's Chorus: Creativity before Consciousness. *Cognitive Computation*, 4(3), 306–319.

- International Conference on Computational Creativity
- Sydney, Australia
- 12-14 June 2013
- computationalcreativity.net