




NEW YORK UNIVERSITY

Perceptual decision making is supported by a hierarchical cascade in both biological and artificial neural networks


Laura Gwilliams

11th July 2018




The world is an uncertain place


❖ Noise



❖ Ambiguity



AI can categorise, too




- ❖ Artificial intelligence has sought to solve a similar problem in visual processing
- ❖ Deep neural networks (DNNs) can label images very accurately

Three screenshots demonstrating AI image classification results:

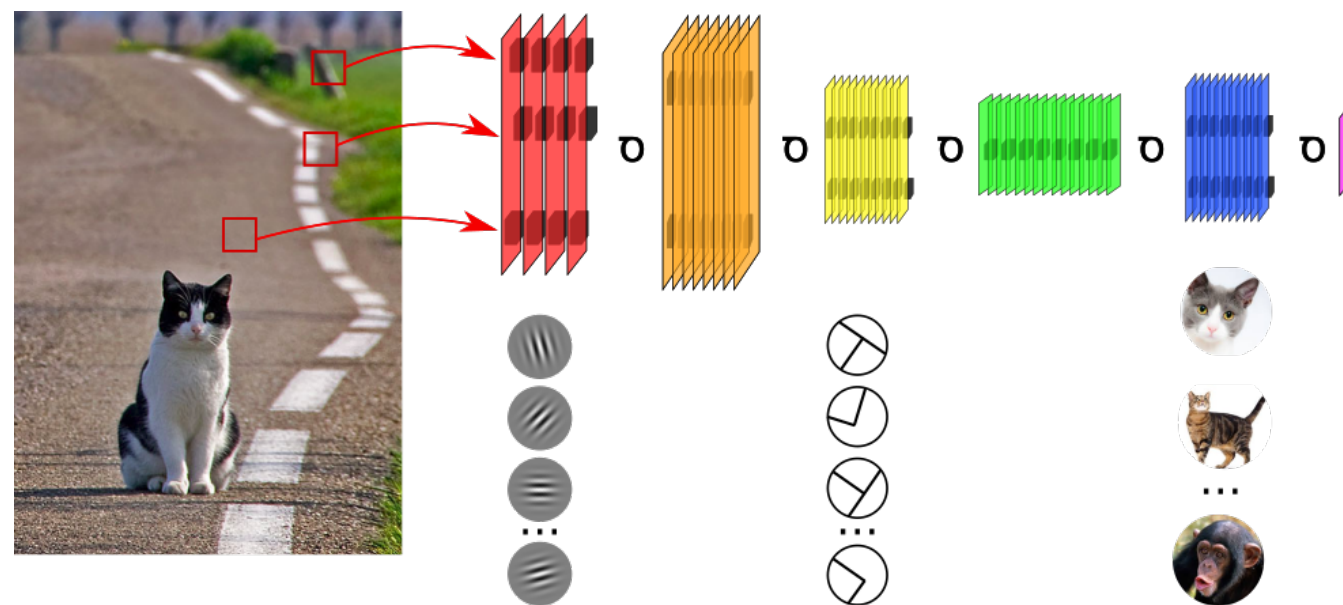

- Soccer Ball:** A white soccer ball with green and yellow patterns. The terminal output shows: `Label: soccer_ball, 93.43%` and a list of other categories: 2. rugby_ball: 6.06%, 3. golf_ball: 0.26%, 4. volleyball: 0.17%, 5. tennis_ball: 0.05%.
- Revolver:** A man holding a revolver. The terminal output shows: `Label: revolver, 69.79%` and a list of other categories: 2. rifle: 7.74%, 3. microphone: 7.31%, 4. assault_rifle: 5.63%, 5. bow: 1.96%.
- Convertible Car:** A white convertible car parked on a beach. The terminal output shows: `Label: convertible, 91.76%` and a list of other categories: 2. sports_car: 4.98%, 3. limousine: 1.06%, 4. car_wheel: 0.75%, 5. beach_wagon: 0.41%.



AI and neural convergence




- ❖ Correspondence has been found in terms of the *representations* employed by brains and DNNs






AI and neural convergence




- ❖ Not so surprising, given that aspects of DNNs are modelled on vision neuroscience
- ❖ There is more to characterising a system than simply knowing the representations it uses:
 - ❖ Architecture
 - ❖ Computation

Research Question

What is the computational architecture of perceptual decision making?




Laura Gwinn Jones | NYU



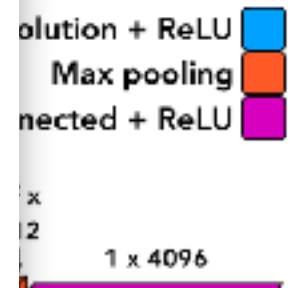
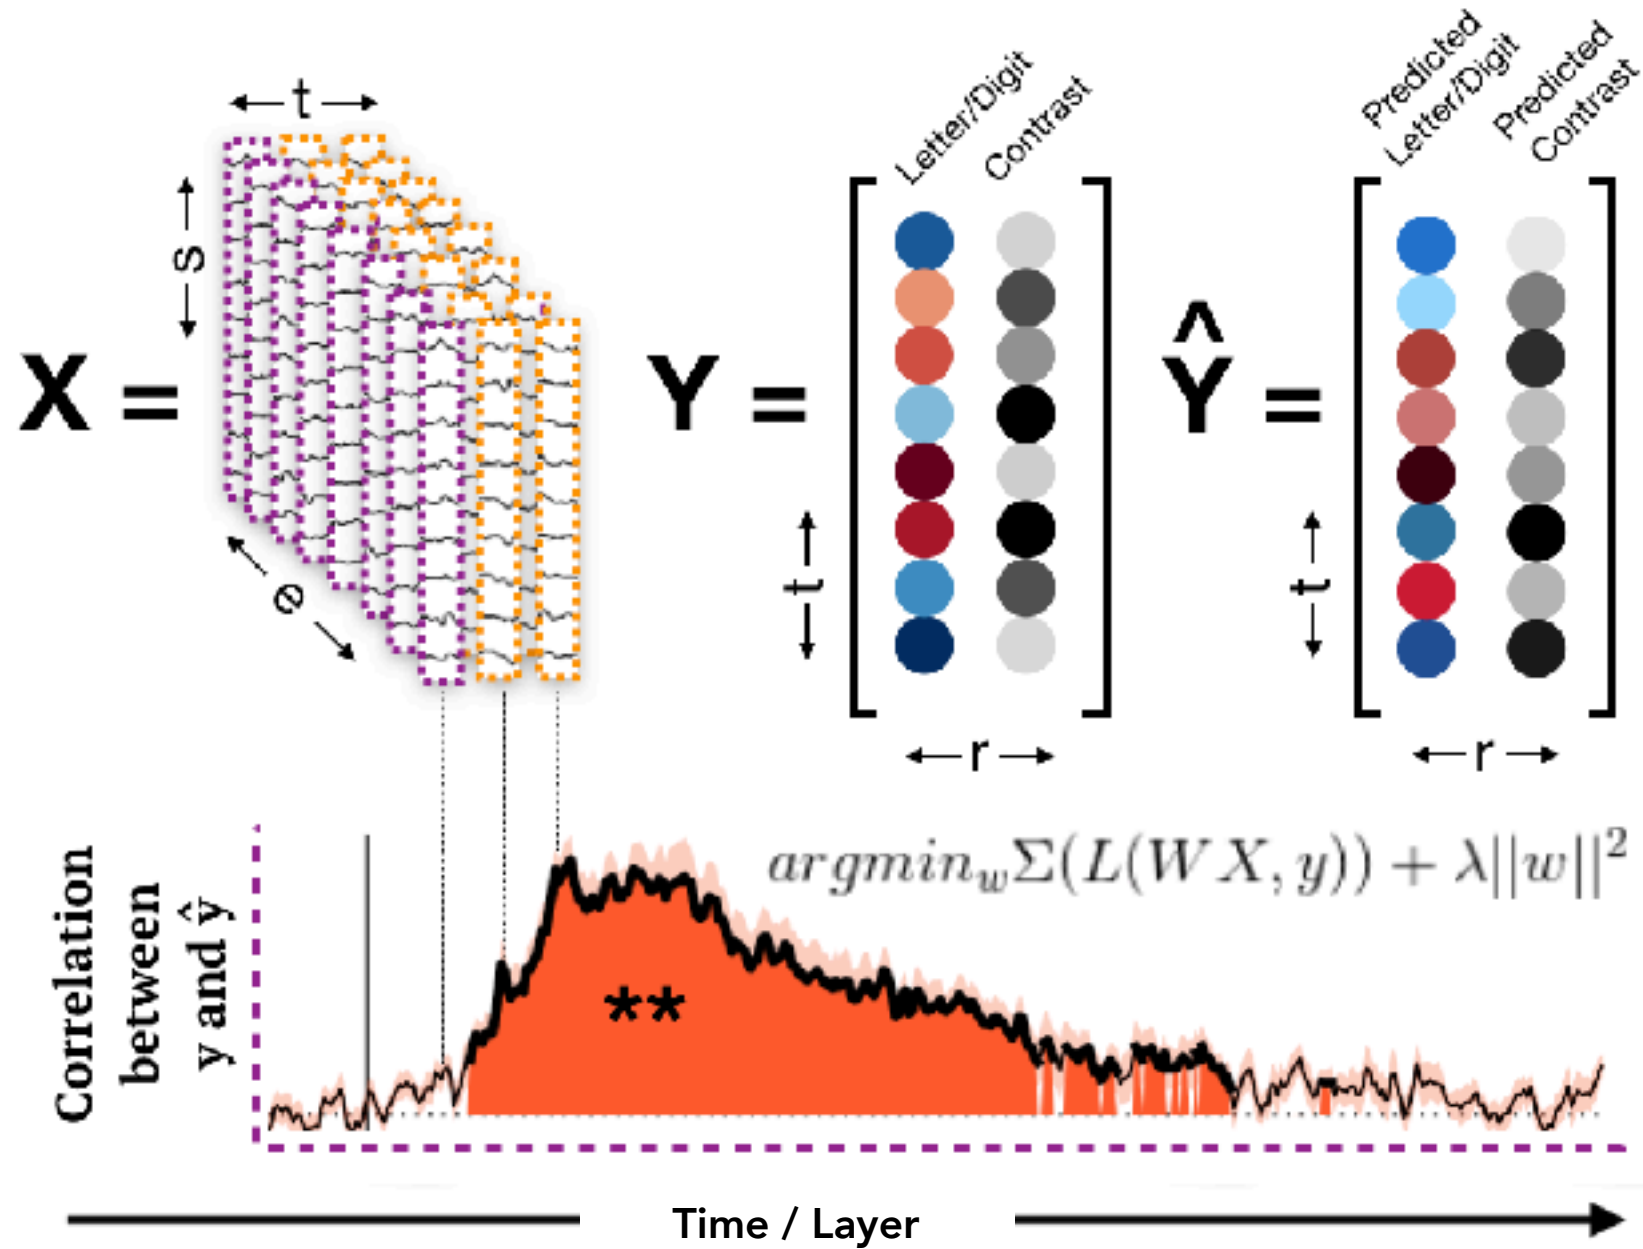



- ❖ 17 healthy adults
- ❖ 306 channel MEG




- ❖ VGG19
- ❖ 19-layer CNN

Parallel Analysis





Question 1

How are processing stages organised — what is the architecture?


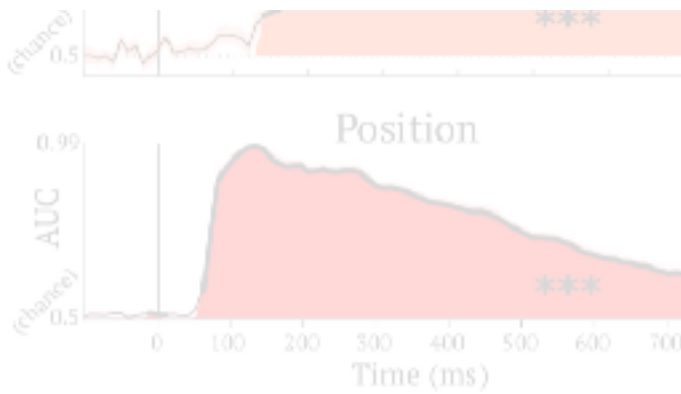




MEG Decoding Scores







MEG Decoding Scores






MEG Decoding Scores

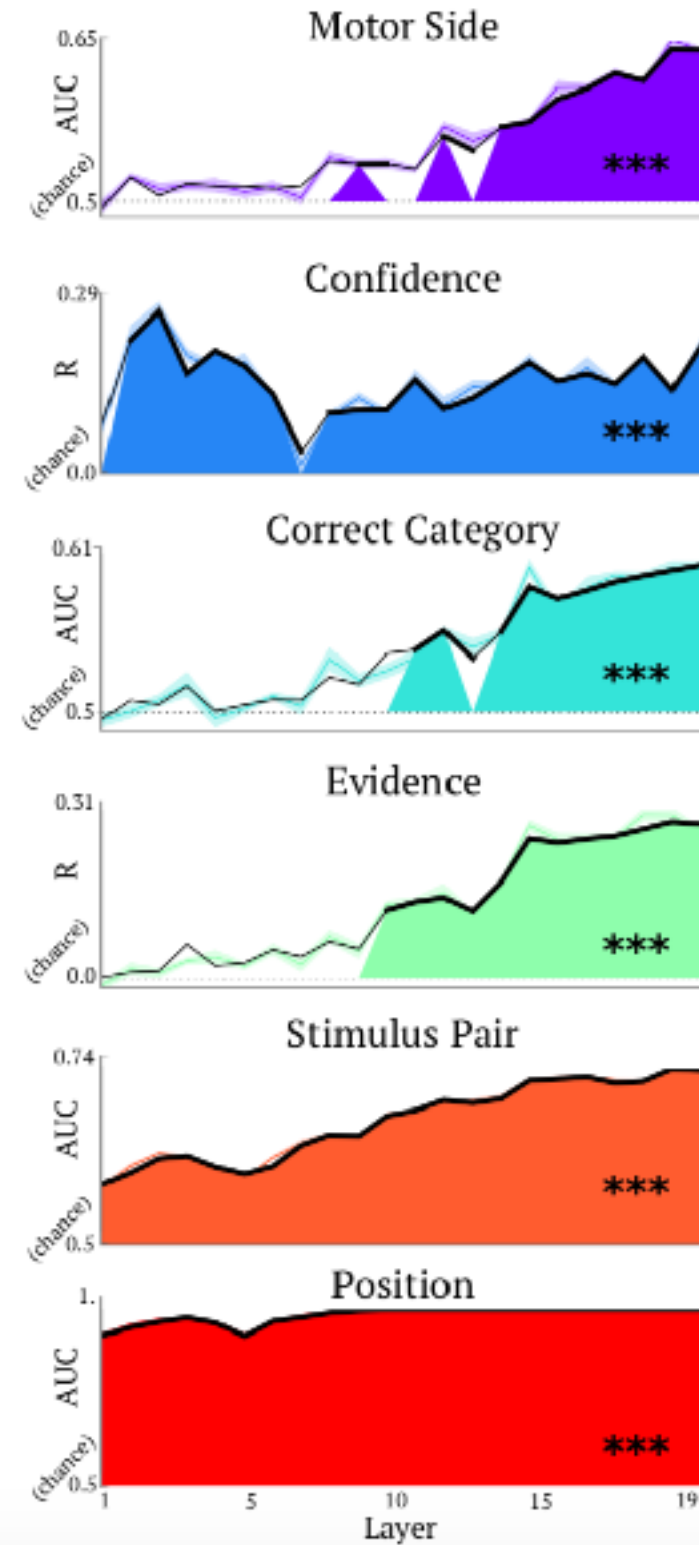







MEG Decoding Scores



DNN Decoding Scores



Question 2

**What are the underlying
computations?**



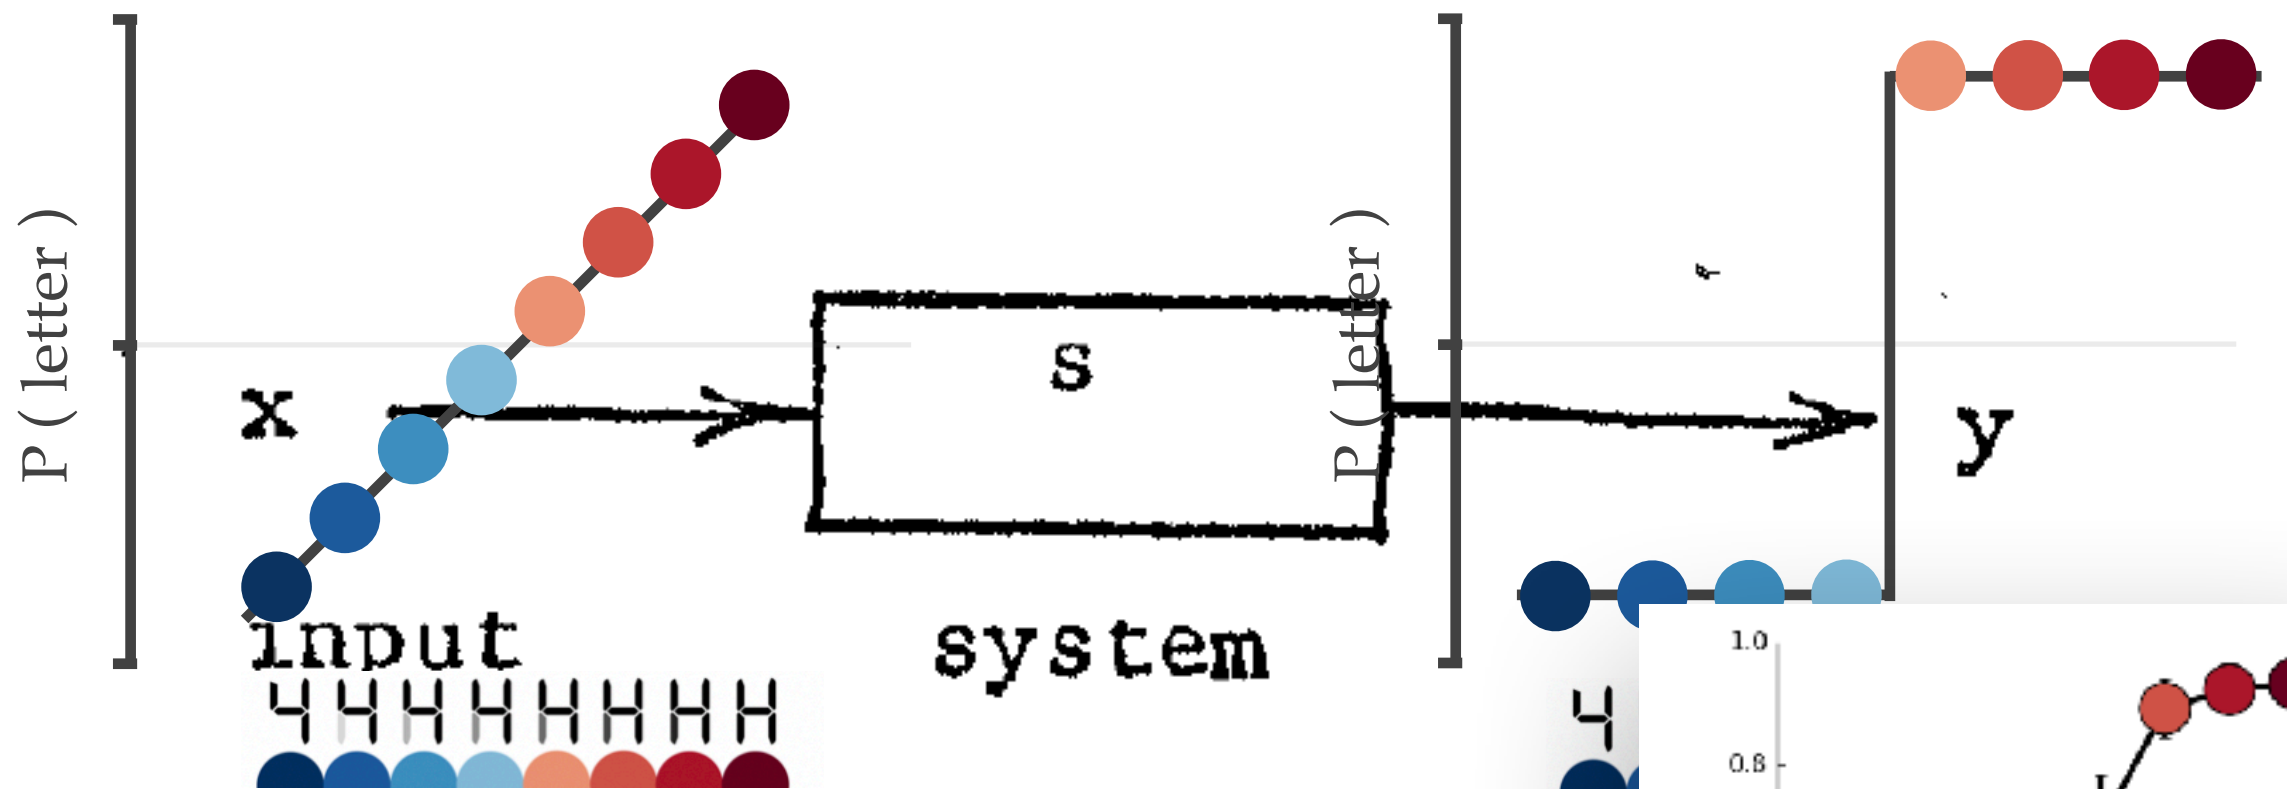


What are the underlying computations?



Linear Evidence

Categorical Percept






What are the underlying computations?




Time: -0.0 s







Categorical
Linear

4 H H H H H H H






What are the underlying computations?





Conclusions



- ❖ We replicate the finding that hierarchical representations converge between biological and artificial NNs
- ❖ **Performance-optimised models only partially predict neural responses during perceptual decision making**


With big thanks to:

@GwilliamsL

- Collaborator **Jean-Rémi King**



- My supervisors, **Alec Marantz** and **David Poeppel**, and everyone in the **Neuroscience of Language Lab** and **Poeppel Lab!**



Laura Gwilliams | NYU



Funding: G1001 Abu Dhabi Institute