Preference Distributions of Primary Motor Cortex Neurons Reflect Control Solutions Optimized for Limb Biomechanics

Timothy P. Lillicrap and Stephen H. Scott

Neuron, Jan 9 2013

How can we understand neural systems?

The job of the experimentalist

- design a task
- record neural activity
- make observations

How can we understand neural systems?

The job of the experimentalist

- design a task
- record neural activity
- make observations

The job of the theoretician

How can we understand neural systems?

The job of the experimentalist

- design a task
- record neural activity
- make observations

The job of the theoretician

- ~~design~~ learn an artificial network of neurons that can solve the task

How can we understand neural systems?

The job of the experimentalist

- design a task
- record neural activity
- make observations

The job of the theoretician

- ~~design~~ learn an artificial network of neurons that can solve the task
- Need some objective function for learning: task objective

How can we understand neural systems?

The job of the experimentalist

- design a task
- record neural activity
- make observations

The job of the theoretician

- ~~design~~ learn an artificial network of neurons that can solve the task
- Need some objective function for learning: task objective
- add in realistic constraints

How can we understand motor cortex, M1?

The job of the experimentalist

- design a task: center-out reaching, loaded-posture
- record neural activity
- make observations: neurons are tuned to movement/torque direction, but biased

The job of the theoretician (this paper)

- ~~design~~ learn an artificial network of neurons that can solve the task: standard neural network with feedback and threshold nonlinearities
- Need some objective function for learning: task objective reach target, balance torques
- add in realistic constraints: minimize muscle activity, neural activity and synaptic weights

Distribution of orientation tuning in M1 neurons is biased and bimodal

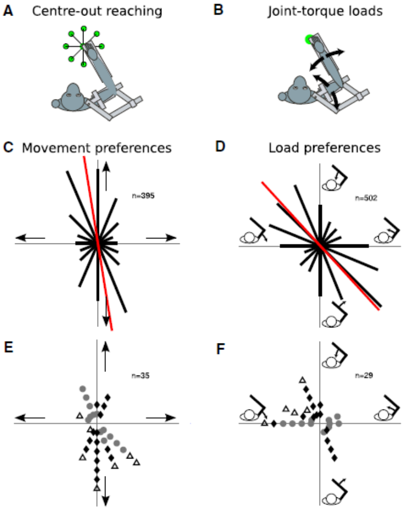

Putting this work in context.

- Unresolved M1 question: do neurons encode goals, arm state or muscle state?
- Is the spinal cord doing something very complicated?
- What is the motor system optimized for?
  - The "natural statistics" of reaching movements? Explains orientation bias, just like in V1.
  - End-point accuracy?
  - Minimize velocities, torques, third derivatives etc?

Putting this work in context. This paper argues.

- Unresolved M1 question: do neurons encode goals, arm state or muscle state? muscle state (torques)
- Is the spinal cord doing something very complicated? no, just linear transformation
- What is the motor system optimized for?
  - The "natural statistics" of reaching movements? Explains orientation bias, just like in V1.
  - End-point accuracy?
  - Minimize velocities, acceleration, third derivatives etc?
  - none of the above, optimized for the biomechanics of the plant.

Putting this work in context. This paper argues.

- Unresolved M1 question: do neurons encode goals, arm state or muscle state? muscle state (torques)
- Is the spinal cord doing something very complicated? no, just linear transformation
- What is the motor system optimized for?
  - The "natural statistics" of reaching movements? Explains orientation bias, just like in V1.
  - End-point accuracy?
  - Minimize velocities, acceleration, third derivatives etc?
  - none of the above, optimized for the biomechanics of the plant.

Follow-up question: what is the plant optimized for?

The proposed model (multiple versions)

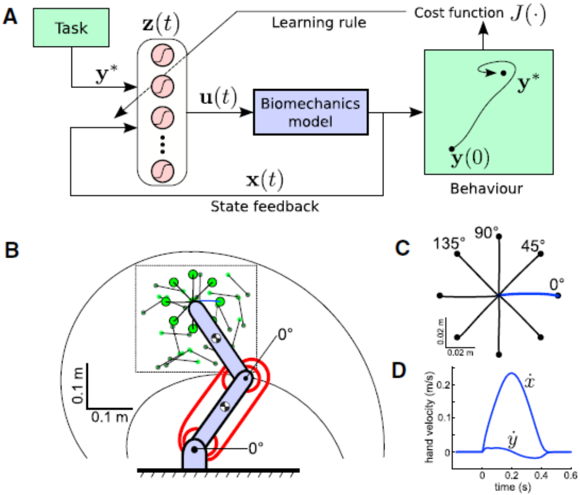

Results - model neurons (basis functions)

Orientation tuning

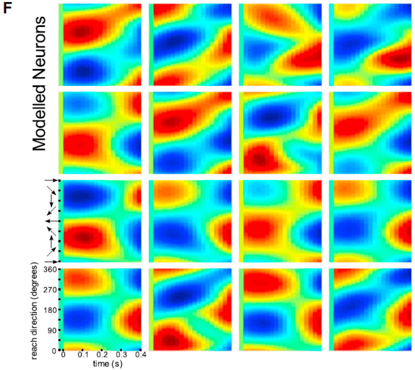

Results - model neurons (basis functions)

Orientation tuning

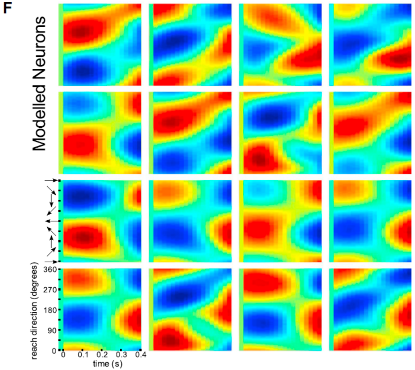

Tuning distributions over the population

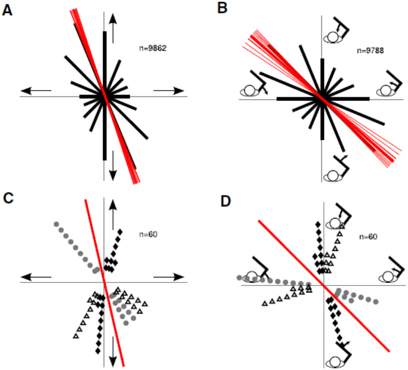

Incremental biomechanics

What constraints of the model explain the results?

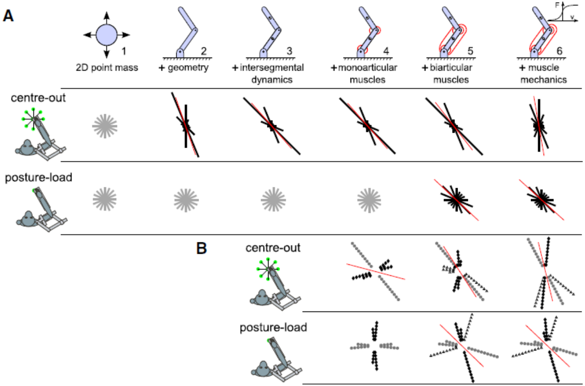

Overall, distribution is biased in the direction of most muscle force (torque).

Incremental biomechanics

What constraints of the model explain the results?

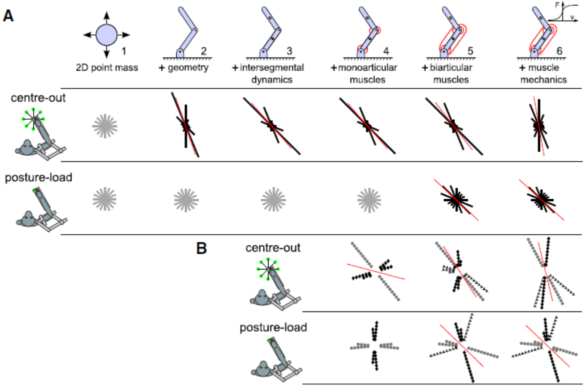

Overall, distribution is biased in the direction of most muscle force (torque).
Results are probably explained by the L2 prior on the neural activities.

(biased) natural statistics of movement do not result in biased orientation tuning

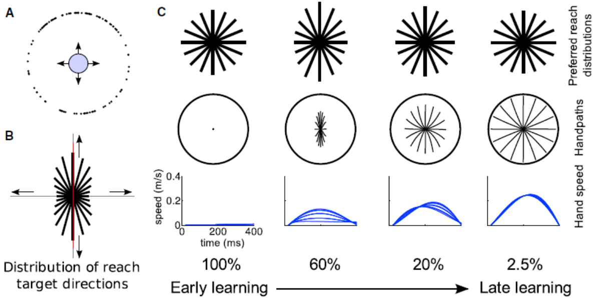

(biased) natural statistics of movement do not result in biased orientation tuning

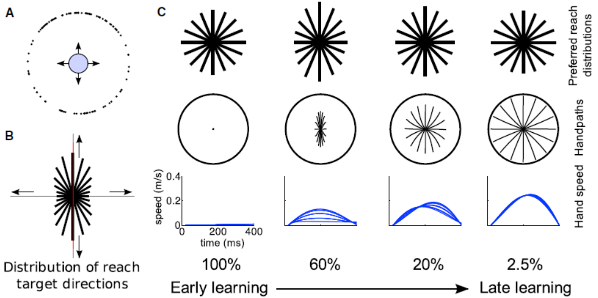

Opposite of orientation tuning in V1, but L1 prior on neural activities used there.

(biased) natural statistics of movement do not result in biased orientation tuning

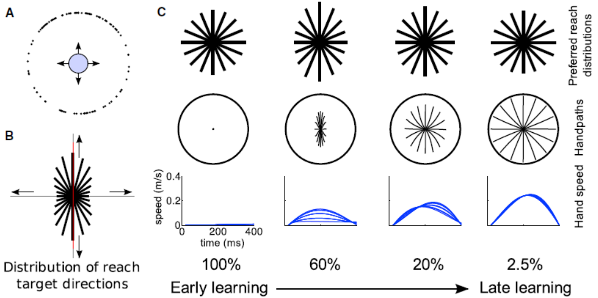

Opposite of orientation tuning in V1, but L1 prior on neural activities used there.
Optimization of this model is very similar to an autoencoder, except for horrendous nonlinearities in the motor plant.

Conclusions

Two results

- Network optimized to control realistic limb learns same direction biases as M1 neurons.
- Network optimized for natural statistics of movement does not learn biases.

Conclusions

Two results

- Network optimized to control realistic limb learns same direction biases as M1 neurons.
- Network optimized for natural statistics of movement does not learn biases.
- Why should it be different from vision?

Conclusions

Two results

- Network optimized to control realistic limb learns same direction biases as M1 neurons.
- Network optimized for natural statistics of movement does not learn biases.

- Why should it be different from vision?
- I believe the L2 vs L1 prior on neural activities explains the difference.