Prediction, a Prerequisite to Goal-directed Behavior, and Its Possible Origin in Delay Compensation

INNS Texas/MIND/UTA Conference on Goal-Directed Neural Systems

November 3, 2006

Yoonsuck Choe

Department of Computer Science, Texas A&M University

Joint work with Jaerock Kwon and Heejin Lim

Origin of Goal-Directed Behavior: Recasting the Question

- What is a goal?
  - It is something in the future.
  - To form a goal, one needs to see into the future.
  - That is, **prediction** is necessary.
- How does prediction arise in the nervous system and how does it affect behavior?

Whence Prediction?

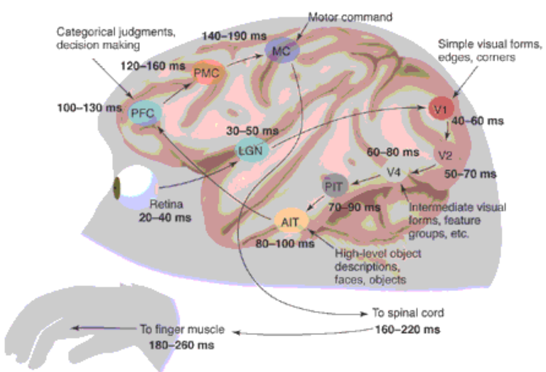

Thorpe and Fabre-Thorpe (2001)

- Due to **neural conduction delay** (couple of 100 ms), we cannot even seem to catch up with the present.
- At best, we will be predicting the **present**, based on the **past**.

Flash Lag Effect: Evidence of Prediction in the Brain?

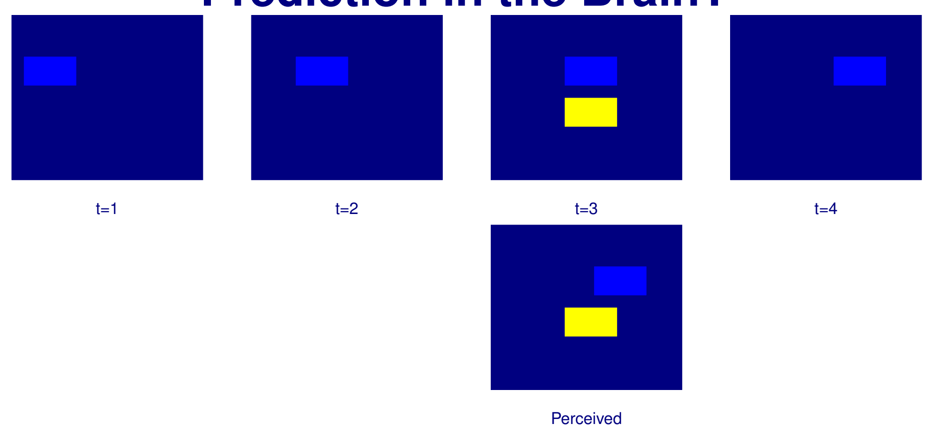

- Flash-Lag Effect (Nijhawan 1994) suggests that the brain may be performing extrapolation to compensate for delay.

Demo: Flash Lag Effect

Implications of FLE

- There may be mechanisms in the brain for **delay compensation** through **extrapolation**.
- The brain may **predict** the present, based on the past.
- Alternative hypotheses: differential latency (Whitney and Murakami 1998), postdiction (Eagleman and Sejnowski 2000), etc.

W/O Delay Compensation: No FLE

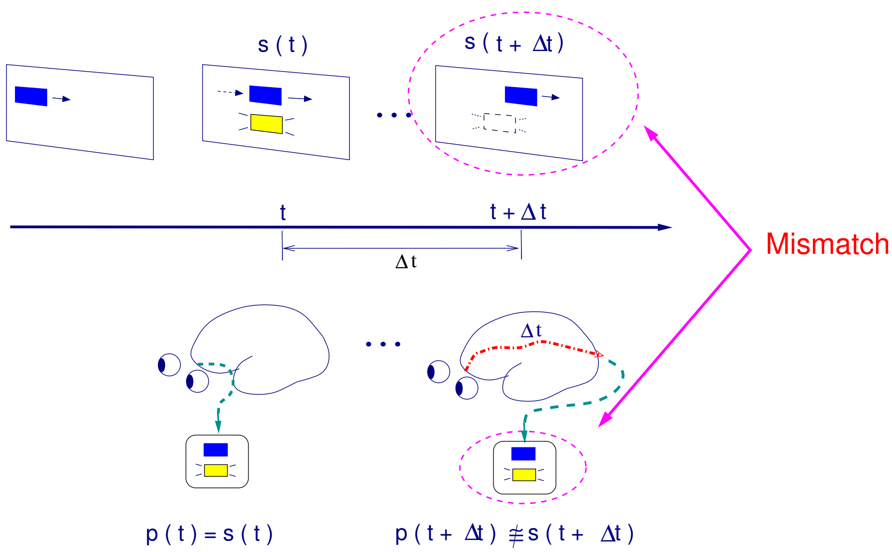

With Delay Compensation: FLE

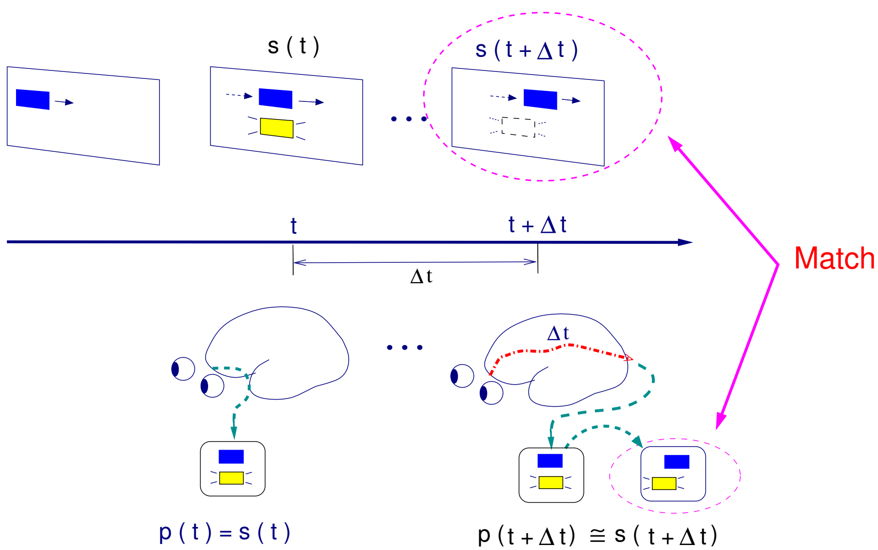

Research Questions

- How can the nervous system compensate for internal delay?
- Are there single-neuron-level mechanisms for that?

Potential Answers

Extrapolation can be used to compensate for delay:

- That can happen at a single-neuron level.
- Facilitatory neural dynamics may be the underlying mechanism.
- FLE may be a side-effect of such a compensatory process.

Approach

Integrate insights from:

1. Psychophysics: Flash-lag effect
2. Neurophysiology: Dynamic synapses
3. Computational theory: Extrapolation

And, potential link to neurology (autism and dyslexia).

Dynamic Synapses

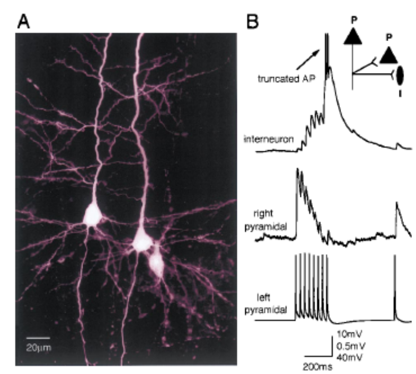

FIG. 2. Differential synaptic facilitation and depression via the same axon innervating two different targets. (*A*) A light microscopic psuedocolor image of three biocytin-filled neurons. The pyramidal neuron on the left innervated the pyramidal neuron on the right and the bipolar interneuron on the right. (*B*) Single trial responses (30 Hz) to same AP train. Failure rate for first EPSP: interneurons, 24%; pyramidal neuron, 0% (60 sweeps). Coefficient of variation (CV; as in ref. 15) for first EPSP: interneuron, 1.12; pyramidal neuron, 0.15. CV for 6th EPSP: interneuron, 0.32; pyramidal neuron, 0.68.

(Markram et al. 1998)

Dynamic Synapses

The effect of synaptic transmission changes dynamically.

- Dynamic increase: Facilitating synapse.
- Dynamic decrease: Depressing synapse.
- Time scale: several hundred milliseconds from the onset (Liaw and Berger 1999; Fortune and Rose 2001; Markram 2002)

Alternative Role of Dynamic Synapses

- Previous: **memory** (sensitization and habituation) (Zucker 1989; Fisher et al. 1997).
- Previous: **temporal information processing** (Fuhrmann et al. 2002; Markram et al. 1998; Fortune and Rose 2001).
- Proposed: **extrapolation** (facilitating synapses).

Target Experiment: Luminance FLE

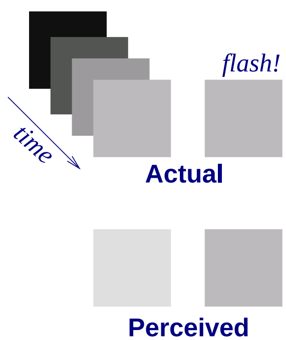

Sheth et al. (2000)

- Works in both directions: increasing or decreasing.
- A single neuron can model the phenomenon.
  - Firing rate represents the perceived luminance.

Available Resource (R) and Synaptic Efficacy (U)

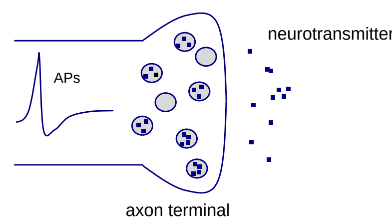

- R: Fraction of recovered neurotransmitters.
- U: Probability of neurotransmitter release.
- Postsynaptic response is dependent on R and U.

Model: Dynamic Synapse

- Synaptic efficacy U (Markram et al. 1998; Fuhrmann et al. 2002):

$$\frac{dU}{dt} = -\frac{U}{\tau_f} + C(1-U)\delta(t-t_s), \qquad (1)$$

where τ_f: time constant for the decay of U; C a constant determining the increase in U due to spikes at t_s; and $\delta(\cdot)$ the Dirac delta function.

- To model extrapolation in the decreasing direction:

$$C = \left(\frac{I(n-1)-I(n)}{|I(n-1)-I(n)|}\right)\left(\frac{I(n-1)}{I(n)}\right) r, \qquad (2)$$

where $I(n)$ is the inter-spike interval.

Model: Membrane Potential

- Postsynaptic current $P(t)$:

$$P(t) = Ee^{-\frac{t}{\tau_P}}, \qquad (3)$$

$$E = AU, \qquad (4)$$

- Membrane potential $V_m(t)$:

$$V_m(t) = V_m(t-1)e^{-\frac{t}{\tau_m}} + P(t)(1-e^{-\frac{t}{\tau_m}}). \qquad (5)$$

- Once V_m exceeds the spike threshold θ, a spike is generated, followed by an absolute refractory period of τ_{refrac}.

Results

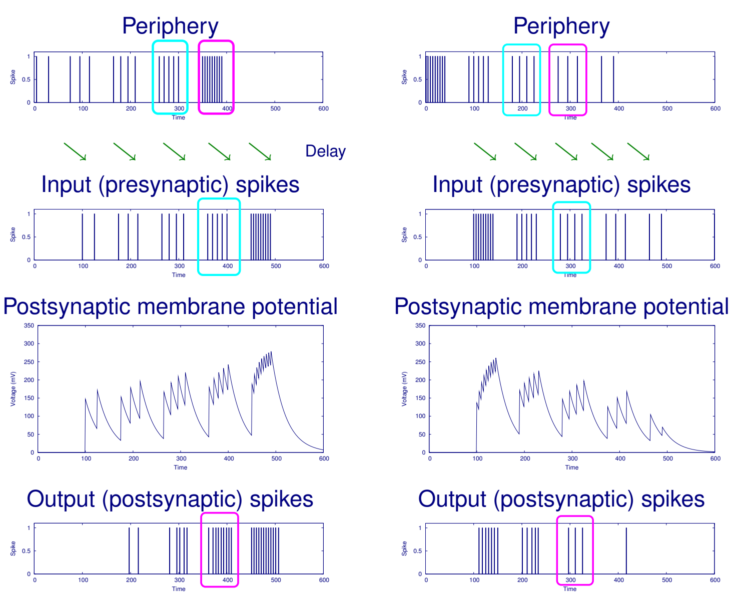

(a) Increasing firing rate (b) Decreasing firing rate

Luminance FLE: Summary

- FLE can be due to delay compensation mechanism.
- Facilitating synapses may be the neural basis of delay compensation.
- Limitations:
  - Cannot explain cross-neuronal facilitation such as orientation FLE

Target Experiment: Orientation FLE

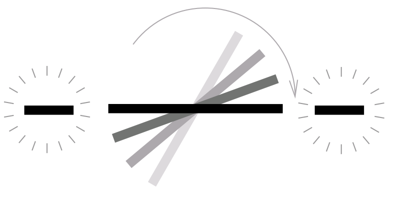

(a) Physical

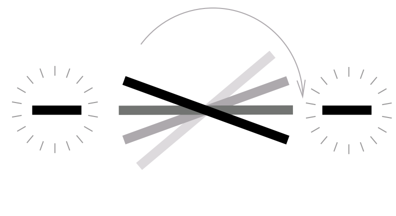

(b) Perceived

- Cannot model with single neuron.
  - V1 orientation-tuned cells have narrow tuning.
- Need network of neurons, with directionally biased weights.

Model: A Ring of Orientation Cells

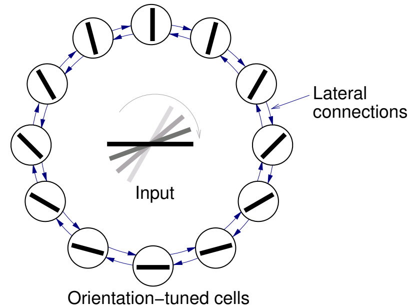

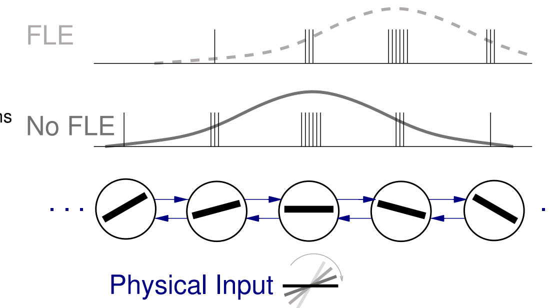

- Shift in firing rate distribution when FLE occurs.
- Needed:
  - Directionally biased connection weights.
  - Facilitating dynamics.

Model: STDP and Facil. Synapses

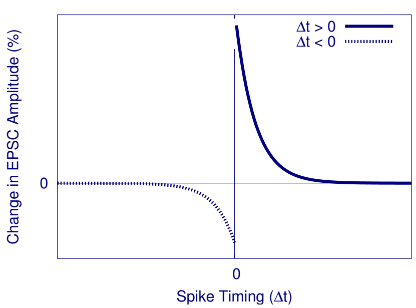

(a) STDP ($\Delta t = t_{\text{post}} - t_{\text{pre}}$)

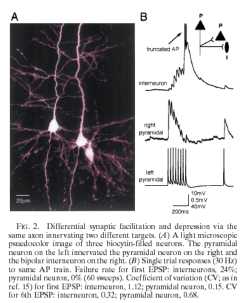

(b) Facil. Synapses

- Spike Timing Dependent Plasticity (Bi and Poo 1998): Set up directionally biased weights.
- Facilitating Synapses: Extrapolation across connections.

Results: Learned Weights

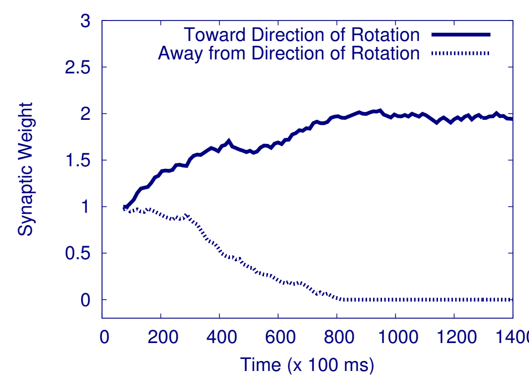

- Weight in the direction of rotation increases.
- Weight in the opposite direction of rotation decreases.

Results

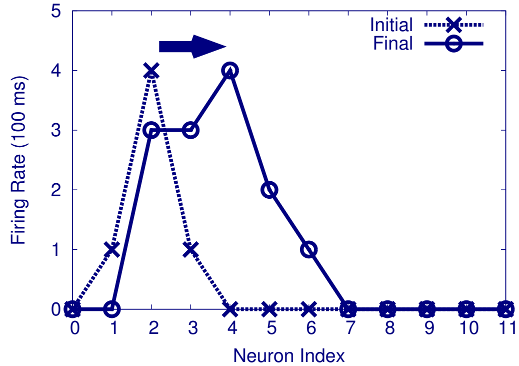

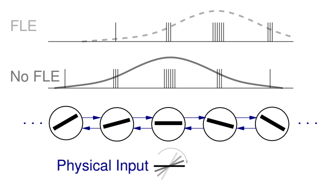

- Peak firing neuron shifts in the direction of rotation.

Results: STDP or Facil. Synapse Alone

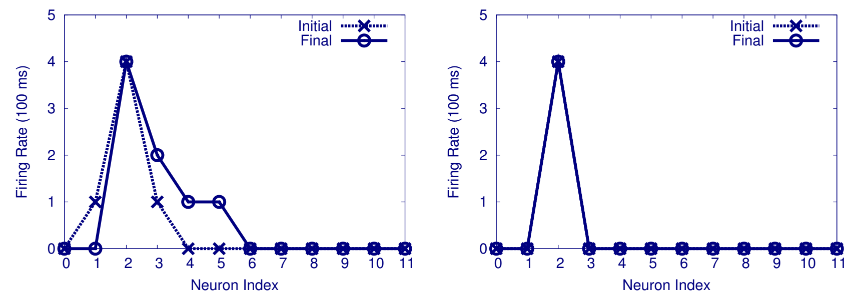

- STDP or facilitating synapses alone was insufficient.

Orientation FLE: Summary

For cross-neuronal facilitation, both

- STDP
- Facilitating synpases

are needed.

Application: Pole Balancing

Modified Pole-Balancing Problem

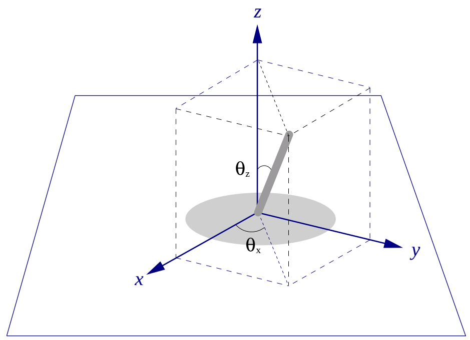

- 2D pole balancing problem.
- **Delay introduced** in input (position and pole angle).

Neuroevolution of Recurrent Neural Network Controller

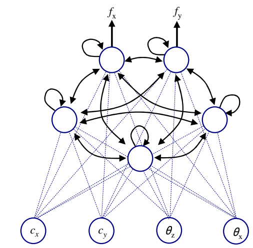

- Fully recurrent neural network controller.
- Trained through neuroevolution (ESP by Gomez and Miikkulainen 1998, 1999; Gomez 2003).

ESP Activation

- Neuron state is determined by **instantaneous** weighted sum of activity:

$$X_i(t) = g\left(\sum_{j \in N_i} w_{ij} X_j(t)\right),$$

where $g(\cdot)$ is a nonlinear activation function, N_i the set of neurons sending activation to neuron i, and w_{ij} the connection weight from neuron j to neuron i.

Approach: Add Dynamics to Neuron Activation

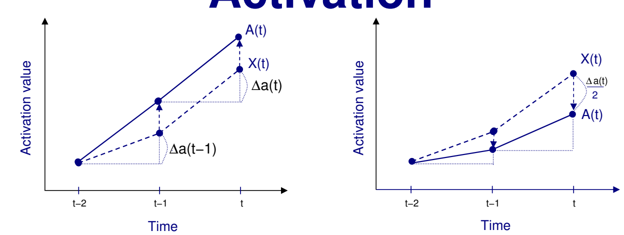

- Facilitatory activity (left):

$$A(t) = X(t) + (X(t) - A(t-1))r,$$

$A(t)$: facilitated activation level at t; $X(t)$: instantaneous activation; r: facilitation rate ($0 \leq r \leq 1$).

- Decaying activity (right): $A(t) = A(t-1)r + X(t)(1-r)$.

Encoding r

- ESP was modified to use the facilitating or decaying dynamics.
- The rate parameter r was encoded in the chromosome so that it can evolve.

Experiment

Compare task performance under three types of dynamics:

- Control: Basic ESP implementation.
- FAN: Facilitatory Activation Network.
- DAN: Decaying Activation Network.

Results: Activation Pattern

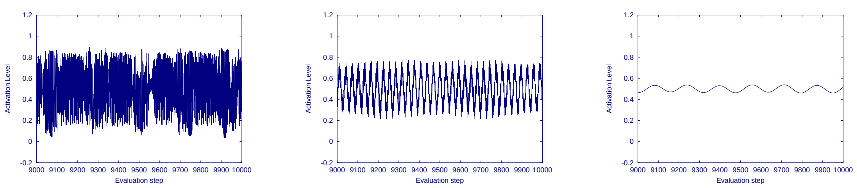

(a) DAN (b) Control (c) FAN

- Last 1000 steps in successful balancing trials.
- 1-step delay, from iteration 50 to 150.
- FAN shows smoother, low-amplitude oscillation.

Results: Cart Trajectory

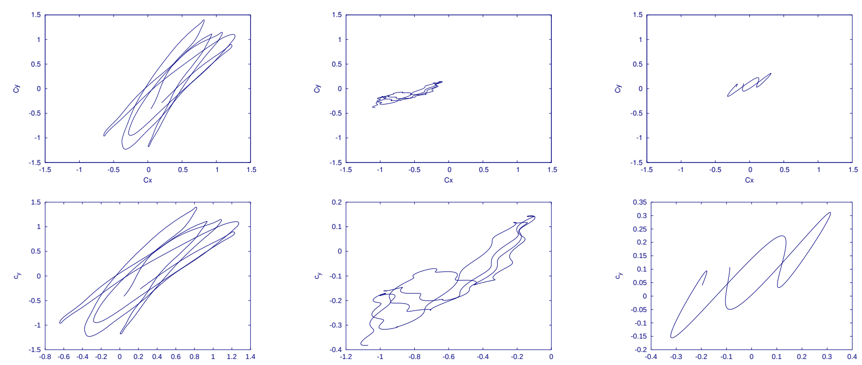

(a) DAN (b) Control (c) FAN

- Last 1000 steps in successful balancing trials.
- 1-step delay, from iteration 50 to 150.
- FAN shows a smooth trajectory with a much smaller footprint.

Results: Success Rate

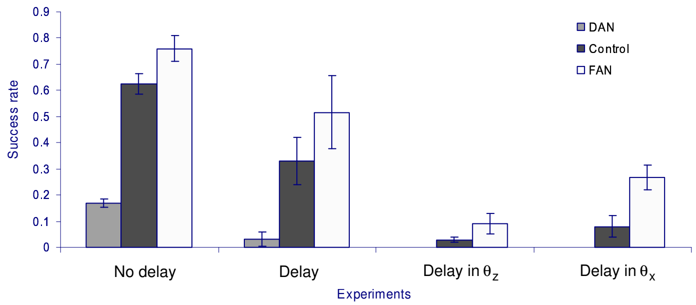

- Different delay conditions were tested.
- FAN showed best performance under all conditions (t-test, $p < 0.005$, $n = 250$).

Results: Speed of Learning

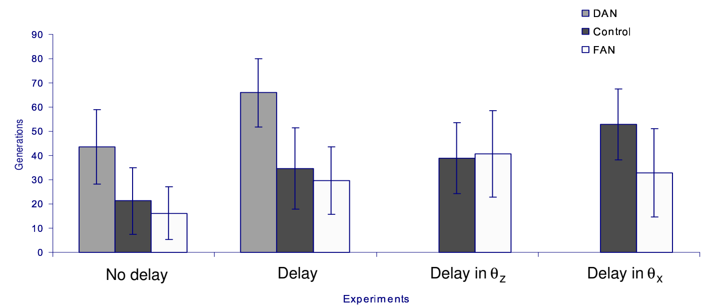

- Different delay conditions were tested (same as above).
- FAN showed best performance under all conditions (t-test, $p < 0.0002$, $n = 250$), except for the θ_z-delay case ($p = 0.84$, i.e., no difference).

Results: Effect of Increased Delay

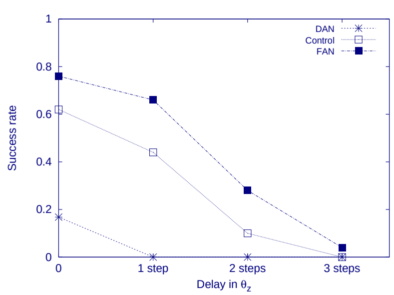

(a) Increased delay

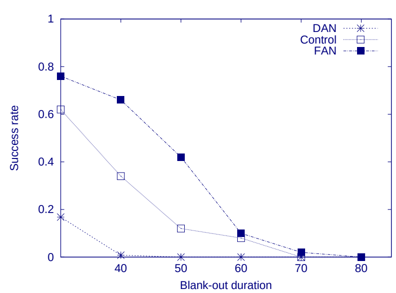

(b) Increased blank-out duration

- Performance under increased delay and input blank-out period.
- In all conditions, FAN performed the best.

Blank-Out as External Delay

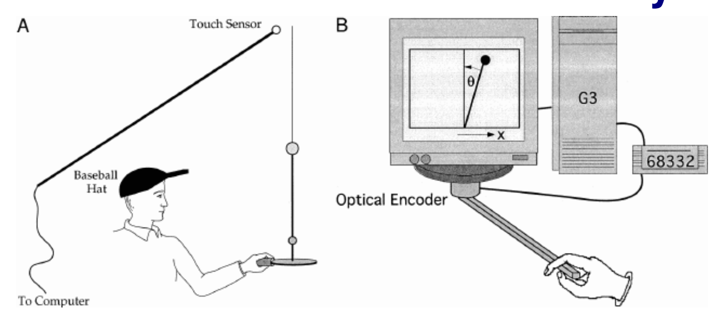

Mehta and Schaal (2002)

- Input feed cut off for $40 \sim 500$ ms while balancing a virtual pole.
- Humans are good at dealing with input blank-out.
- FAM shows similar robustness.

Analysis: Evolution of r

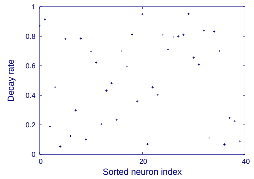

(a) DAN: Initial state

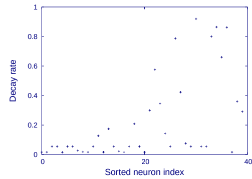

(b) DAN: Final state

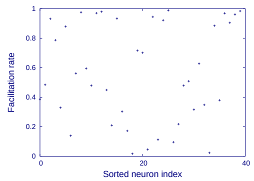

(c) FAN: Initial state

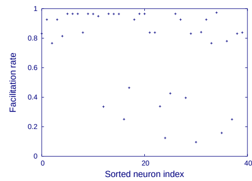

(d) FAN: Final state

- FAN: best neurons had high r
- DAN: best neurons had low r

Summary: Pole Balancing

- Facilitatory dynamics help alleviate debilitating effects of delay in the input.
- Facilitatory dynamics can help in delay in **external** environment as well (potential for real prediction?).
- Decaying dynamics make things worse.

Discussion

- Can predictive traces be found in cortical/cerebellar dynamics? (cf. Harter; Kozma and Freeman; Principe; Werner; Voicu)
- Role of prediction in decision making? (cf. Levine; Wunsch)
- Nonlinear control with delay? (cf. Lewis)
- Facilitating (afferent) and nonfacilitating (associative) synapses in the olfactory system (cf. Gutierrez-Osuna)
- Predicted future state (and goal) as a moving target to be optimized against? (cf. Werbos)
- Use of delay in simulated agents to facilitate the evolution of predictive capabilities (cf. Miikkulainen)
- Differential role of prediction in differentiation-oriented vs. synthesis-oriented cultures (Perelovsky).

Future Directions

Autism:

- Problem in coherent motion detection (Milne et al. 2002).
- Problem with processing moderately rapid motion (Gepner et al. 2001; Gepner 2002).

Dyslexia:

- Difficulty with processing rapidly changing stimulus (Hari and Renvall 2001)

Predictions:

- Autistics and dyslexics may not perceive FLE.
- Abnormal growth in brain size may have outgrown built-in delay compensation mechanisms.

Conclusions

- Facilitatory (extrapolatory) dynamics at a single-neuron level can help compensate for neural delay.
- Facilitatory synapses may be implementing such a function: They are not just for memory!
- Such mechanisms may have evolved into predictive mechanisms providing access to estimated future states.

Question: Why Do We Have a Brain?

Tree
(no Brain)

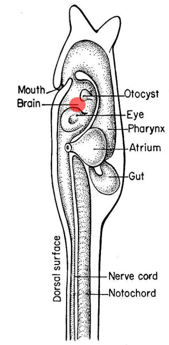

A

Tunicate
Free-floating
(Has brain)

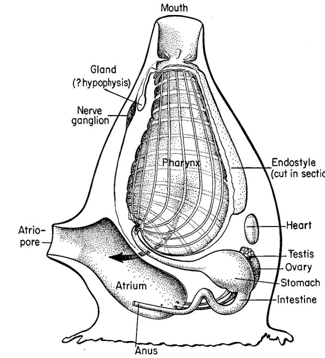

B

Tunicate
Settled
(Brain digested)

- Not just for survival/reproduction, but for **prediction**, to support motion (Llinás 2001).

Sources: http://homepages.inf.ed.ac.uk/jbednar/ and http://bill.srnr.arizona.edu/classes/182/Lecture-9.htm

References

Bi, G.-Q., and Poo, M.-M. (1998). Activity-induced synaptic modifications in hippocampal culture: Dependence on spike timing, synaptic strength and cell type. *Journal of Neuroscience*, 18:10464–10472.

Eagleman, D., and Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science*, 287:2036–2038.

Fisher, S. A., Fisher, T. M., and Carew, T. J. (1997). Multiple overlapping processes underlying short-term synaptic enhancement. *Trends in Neurosciences*, 20:170–177.

Fortune, E. S., and Rose, G. J. (2001). Short-term synaptic plasticity as a temporal filter. *Trends in Neurosciences*, 24:381–385.

Fuhrmann, G., Segev, I., Markram, H., and Tsodyks, M. (2002). Coding of temporal information by activity-dependent synapses. *Journal of Neurophysiology*, 87:140–148.

Gepner, B. (2002). Rapid visual-motion integration deficit in autism. *Trends in Cognitive Sciences*, 6:455.

Gepner, B., Deruelle, C., and Grynfeltt, S. (2001). Motion and emotion: A novel approach to the study of face processing by young autistic children. *Journal of Autism and Developmental Disorders*, 31:37–45.

Gomez, F. (2003). *Robust Non-Linear Control Through Neuroevolution.* PhD thesis, Department of Computer Science, The University of Texas at Austin, Austin, TX. Technical Report AI03-303.

Gomez, F., and Miikkulainen, R. (1998). 2-D pole balancing with recurrent evolutionary networks. In *Proceeding of the International Conference on Artificial Neural Networks*. New York: Elsevier.

Gomez, F., and Miikkulainen, R. (1999). Solving non-markovian control tasks with neuroevolution. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufmann.

Hari, R., and Renvall, H. (2001). Impaired processing of rapid stimulus sequences in dyslexia. *Trends in Cognitive Sciences*, 5:525–532.

Liaw, J., and Berger, T. W. (1999). Dynamic synapse: Harnessing the computing power of synaptic dynamics. *Neurocomputing*, 26-27:199–206.

Llinás, R. R. (2001). *I of the Vortex.* Cambridge, MA: MIT Press.

Markram, H. (2002). Elementary principles of nonlinear synaptic transmission. In Hecht-Nielsen, R., and McKenna, T., editors, *Computational Models for Neuroscience: Human Cortical Information Processing*, chapter 5, 125–169. London: Springer.

Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences*, 95.

Mehta, B., and Schaal, S. (2002). Forward models in visuomotor control. *Journal of Neurophysiology*, 88:942–953.

Milne, E., Swettenham, J., Hansen, P., Campbell, R., Jefferies, H., and Plaisted, K. (2002). Visual motion coherence thresholds in people with autism. *Journal of Child Psychology and Psychiatry*, 43:255–263.

Nijhawan, R. (1994). Motion extrapolation in catching. *Nature*, 370:256–257.

Sheth, B., Nijhawan, R., and Shimojo, S. (2000). Changing objects lead breifly flashed ones. *Nature Neuroscience*, 3:489–495.

Thorpe, S. J., and Fabre-Thorpe, M. (2001). Seeking categories in the brain. *Science*, 291:260–263.

Whitney, D., and Murakami, I. (1998). Latency difference, not spatial extrapolation. *Nature Neuroscience*, 1:656–657.

Zucker, R. S. (1989). Short-term synaptic plasticity. *Annual Review of Neuroscience*, 12:13–31.