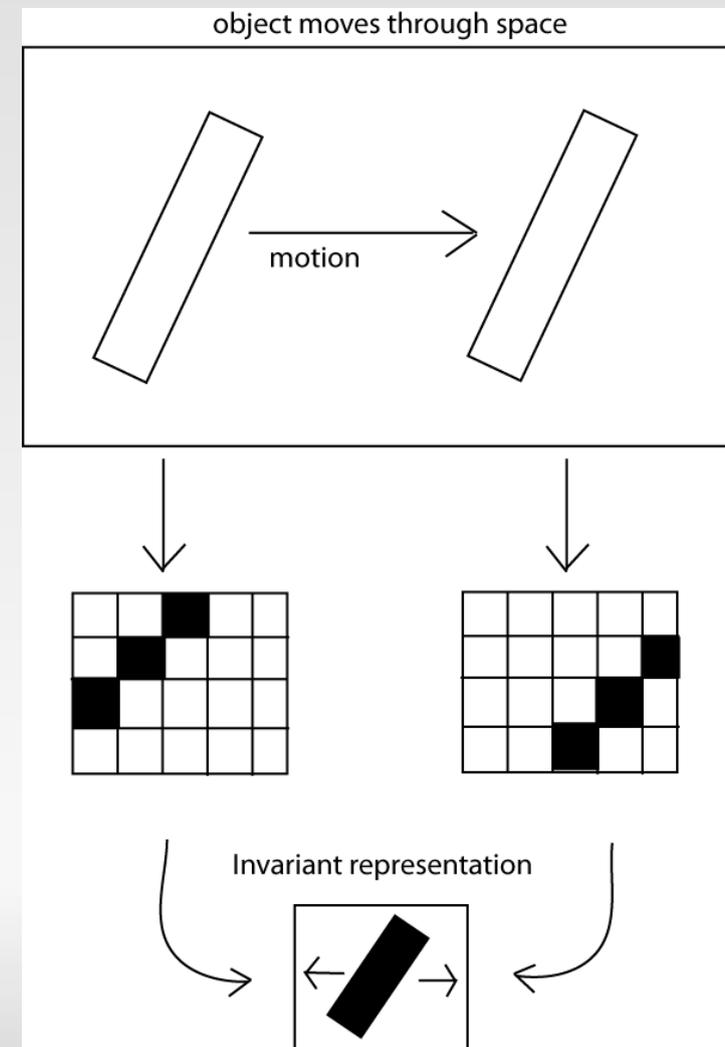


When invariance-learning goes wrong...

A 9.520 project should you choose to accept it.....

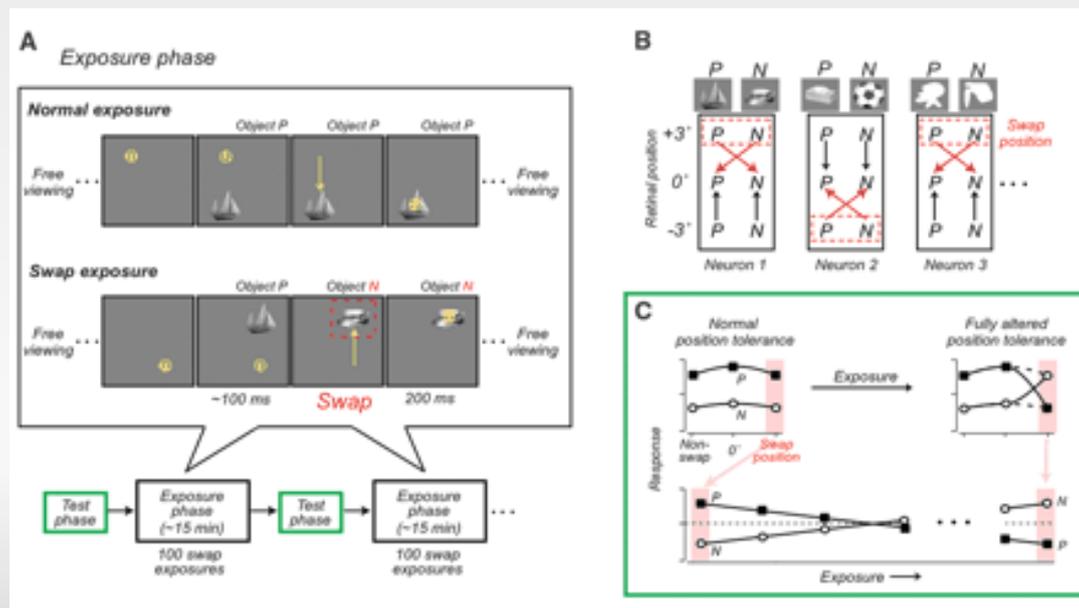
When invariance-learning goes wrong...

- The brain must learn representations that tolerate identity-preserving stimulus transformations.
- Temporal contiguity is a useful cue.



When invariance-learning goes wrong...

- Various algorithms have been proposed for this problem. The brain may be implementing versions of these: The trace rule (Foldiak), Slow feature analysis (Wiskott et al), etc.
- There is both physiological and psychophysical evidence that the brain uses temporal contiguity to build up invariant representations e.g. Cox & DiCarlo 2005, Wallis & Bulthoff 2001, Li & DiCarlo 2008/2009.



When invariance-learning goes wrong...

- If invariant representations are built up using natural experience in this manner, then what happens whenever there are sharp discontinuities in temporal continuity? Presumably this would introduce errors into the resulting representation. *There could be cells that prefer one stimulus at most retinal locations but another stimulus in another location.*
- In this project, you will investigate *how bad these errors are*. How many incorrect associations does it take to ruin translation invariance in a hierarchical model of object recognition?
- I conjecture that these models are rather robust to such errors (as long as they occur *randomly*). This might have neuroscience implications...

When invariance-learning goes wrong...

What you would do:

- Investigate translation invariance for "miswired" hierarchical models. For instance, you could delete/change randomly chosen connections between the S2 and C2 layer of the CBCL (HMAX) ventral stream model.

Resources:

- Software at cbcl.mit.edu
- Myself: jzleibo@mit.edu