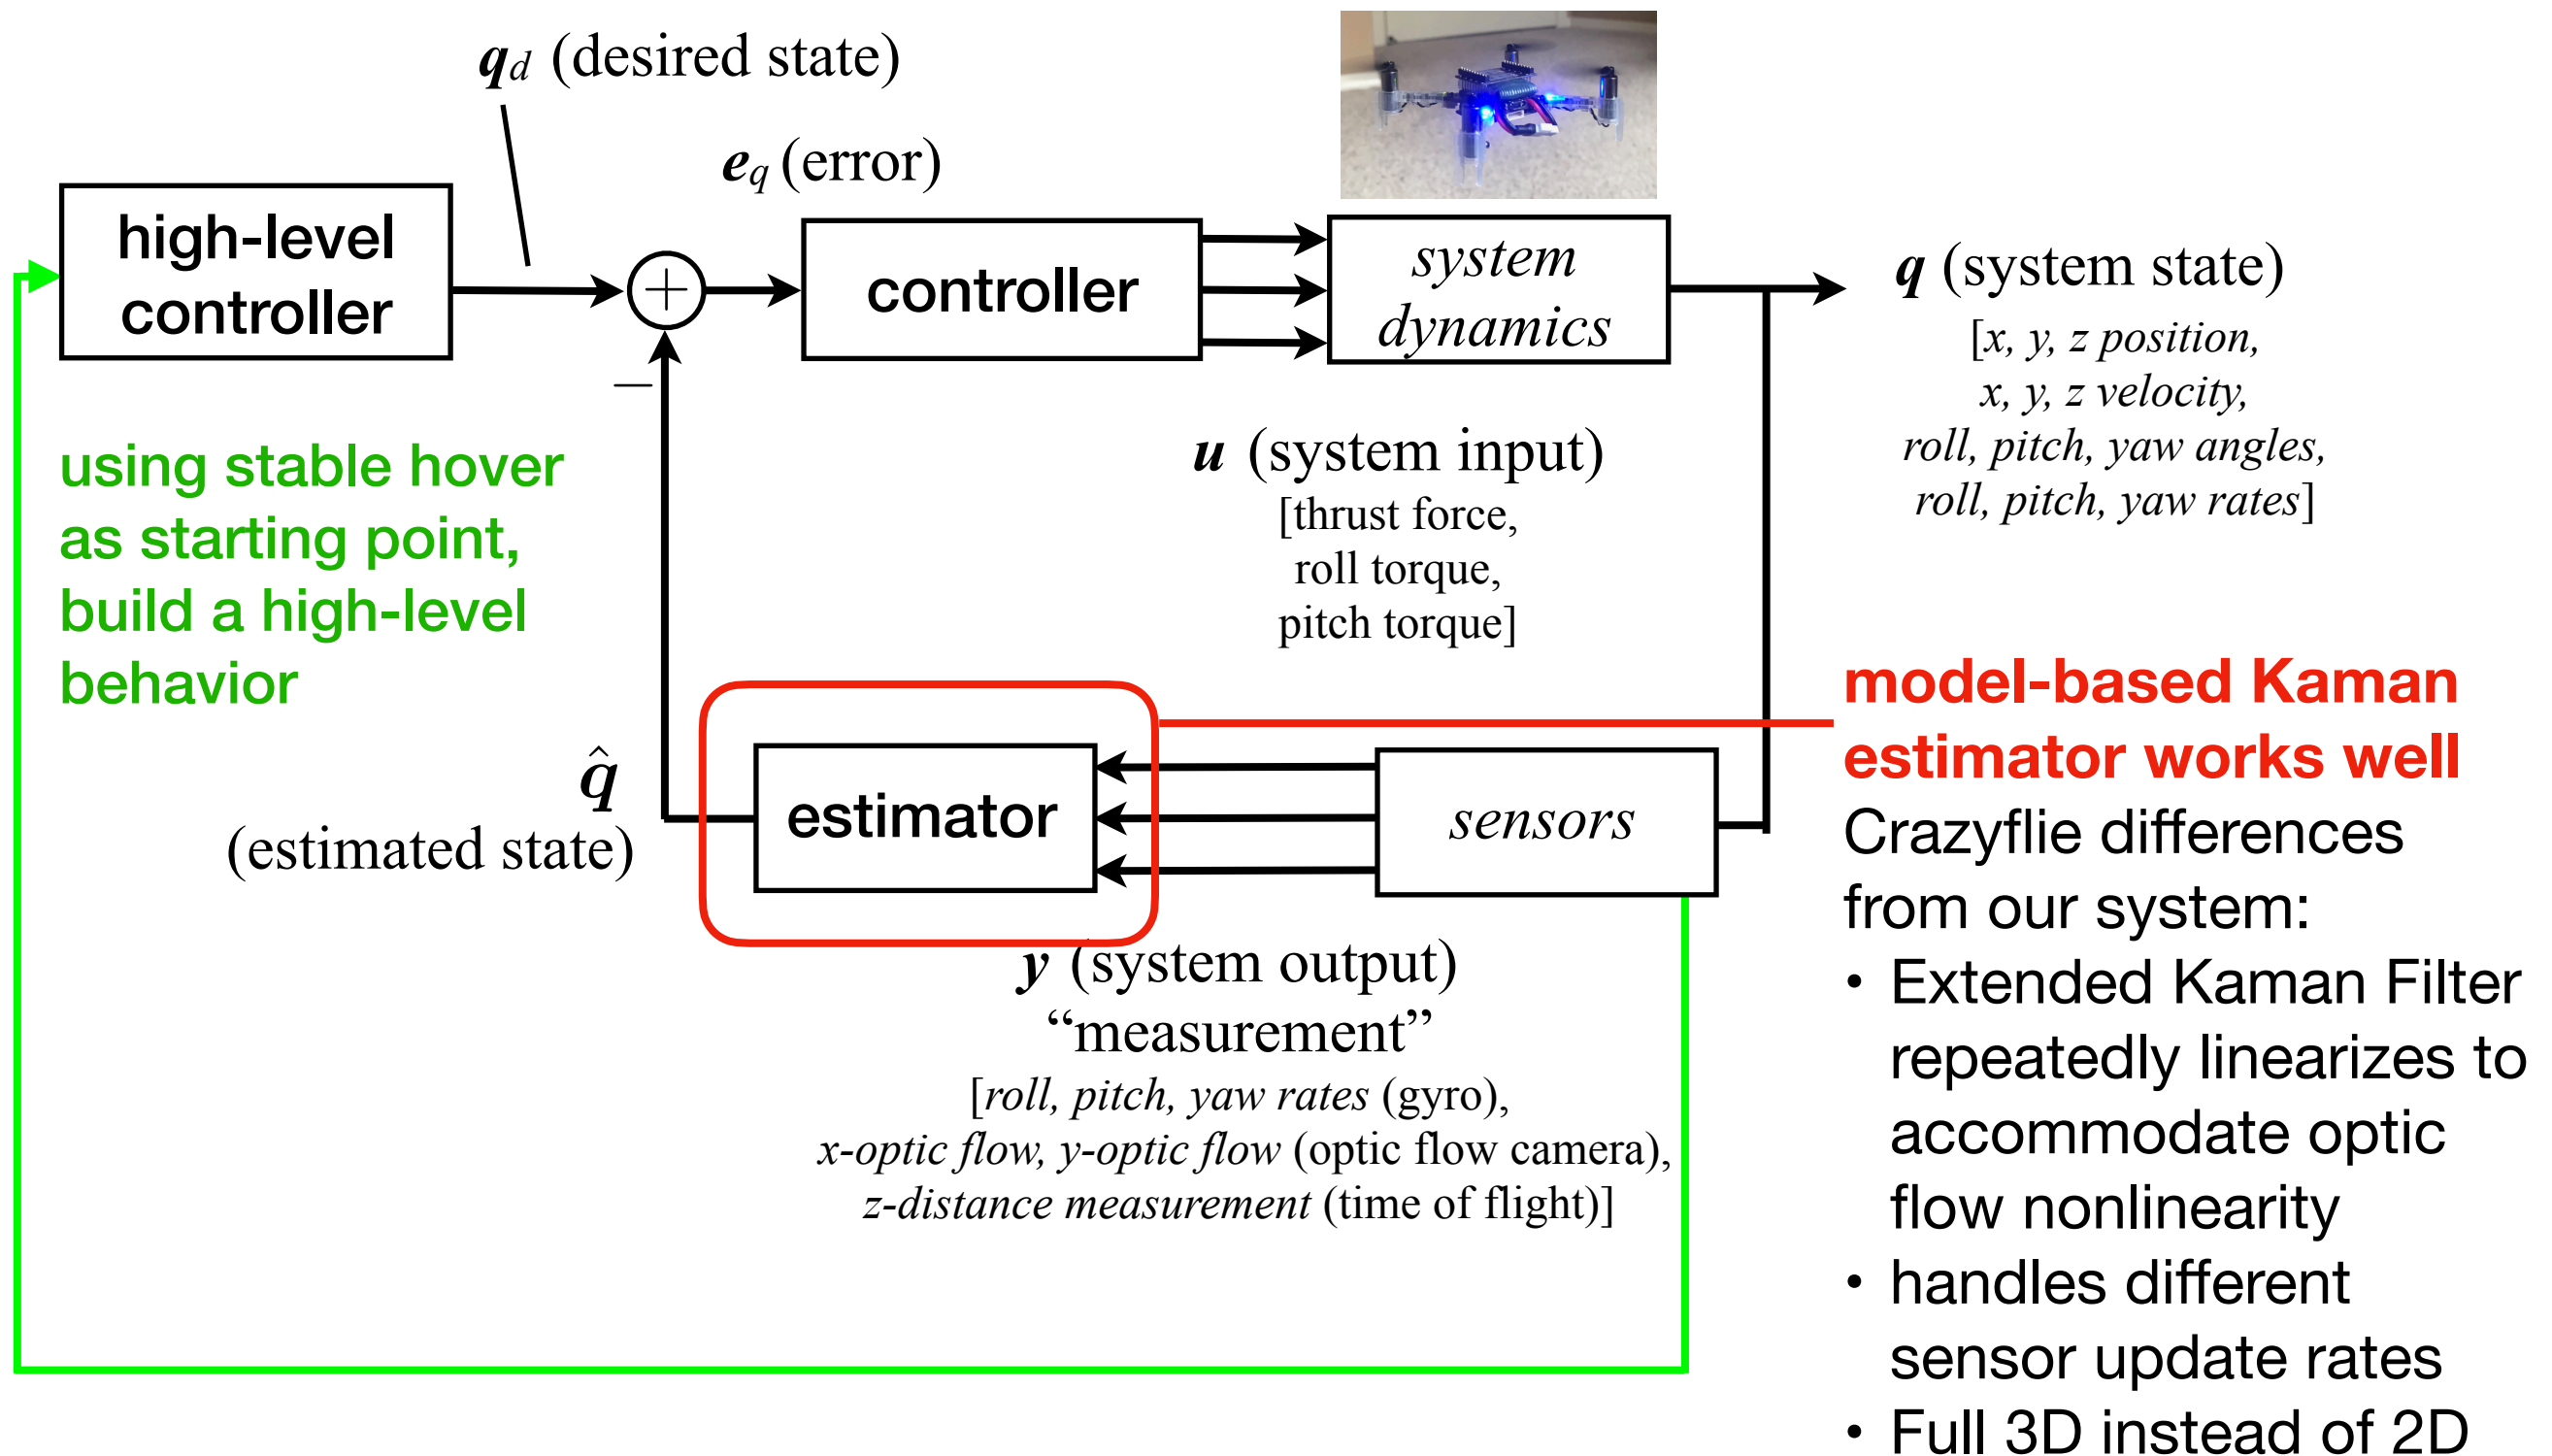
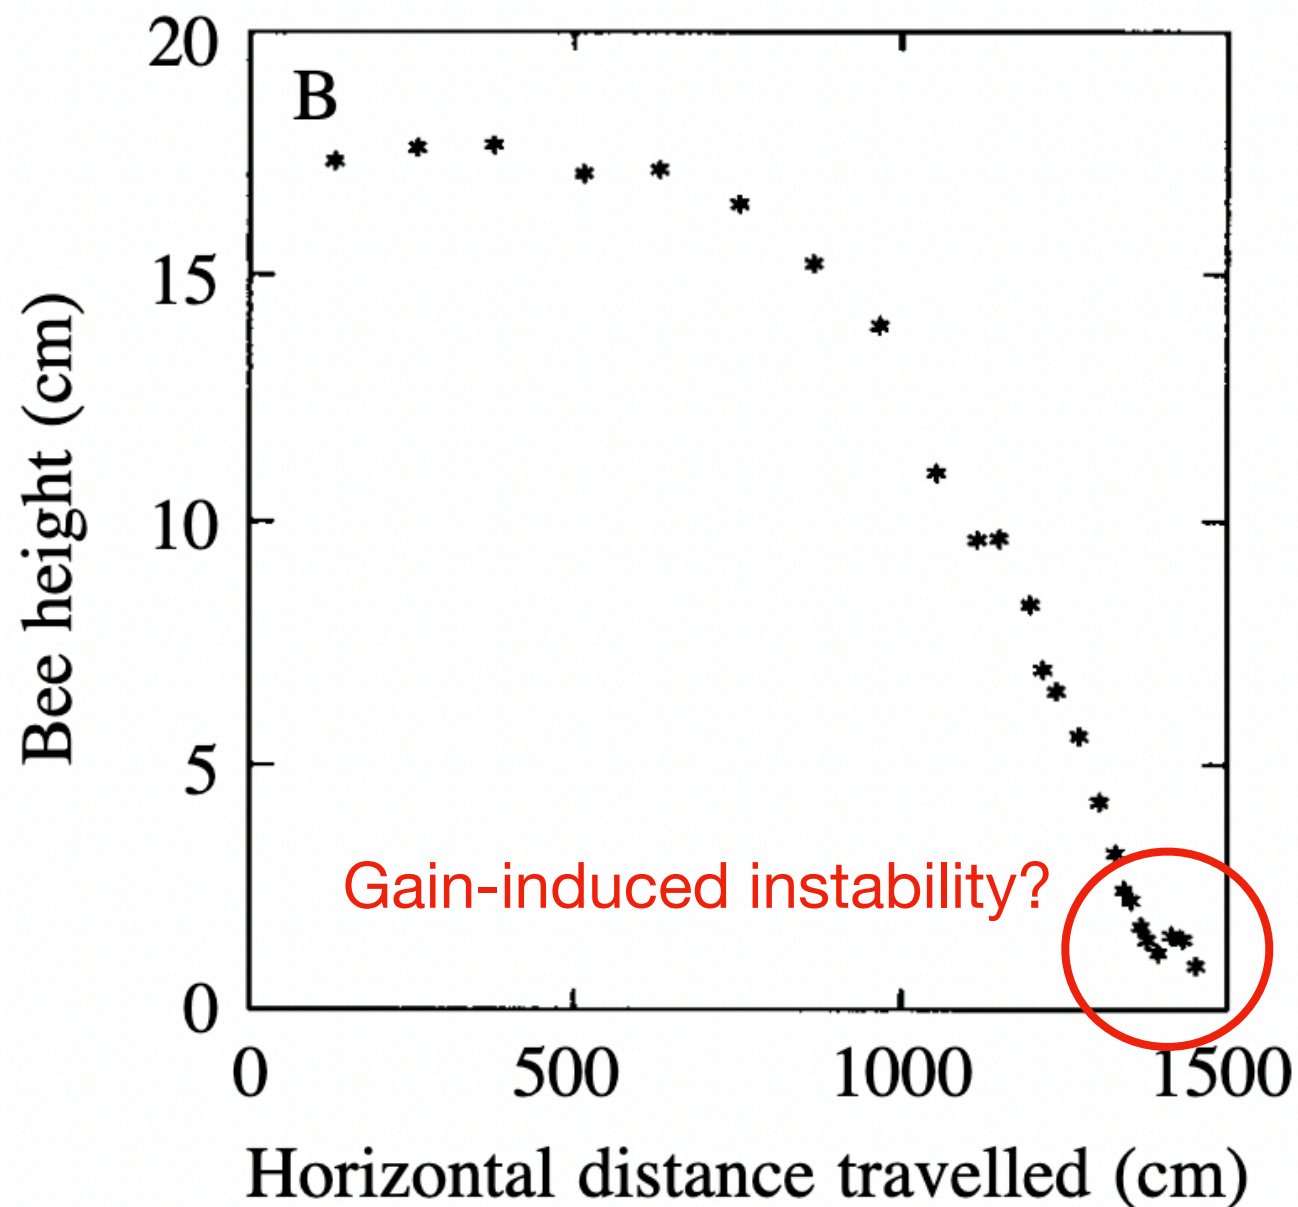


Overview so far




Problem Set 5

Inspired by Decroon2016*, which was inspired by observation of honeybee landing in Srinivasan96*



Nice example of robotics: not just dynamics, not just vision, but the *confluence of both*



Translation-induced optic flow increases as altitude decreases:


$$\Omega_T = -\frac{\dot{x}'}{r}$$

Feedback law demo: $v_{xd} = K\Omega_T$

*available on course webpage

Paper 4 intro. First, a history of neural nets

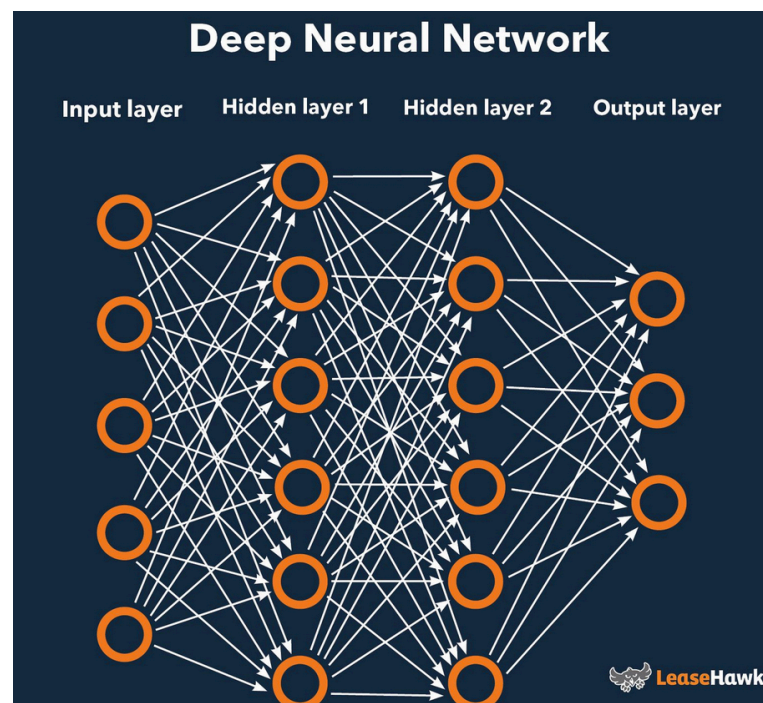
LeCun, Y., et al. (1989). Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4), 541-551.



Your instructor took a course on neural networks.
Why would anybody do that?

???

Krizhevsky, A., Sutskever, I. & Hinton, G. ImageNet classification with deep convolutional neural networks. *NIPS* 25 1090–1098 (2012).




Paper 4 intro

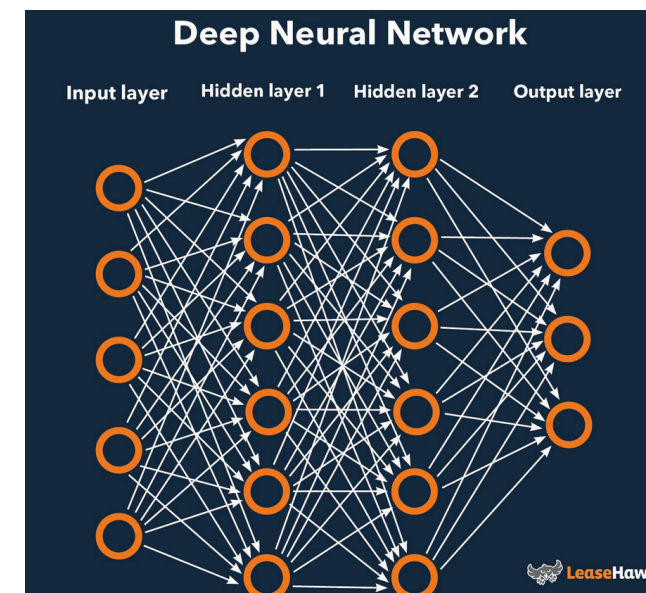
doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih^{1*}, Koray Kavukcuoglu^{1*}, David Silver^{1*}, Andrei A. Rusu¹, Joel Veness¹, Marc G. Bellemare¹, Alex Graves¹, Martin Riedmiller¹, Andreas K. Fidjeland¹, Georg Ostrovski¹, Stig Petersen¹, Charles Beattie¹, Amir Sadik¹, Ioannis Antonoglou¹, Helen King¹, Dharshan Kumaran¹, Daan Wierstra¹, Shane Legg¹ & Demis Hassabis¹



+



State space s : which cell
Action space a : which arrow direction
Learned “Q function” encodes best path

A deep network approximates
“Q function”

Paper 4 is about invention of “Deep reinforcement learning” to play 70’s games.
 s is state of pixels on the screen, a is joystick action, reward comes from the score