

MLPR Lecture 1

Lots of information here.

tinyurl.com/edmlpr

i.murray@ed.ac.uk "Iain" (ee-an)

In lectures:

- Ask questions
- Point out mistakes

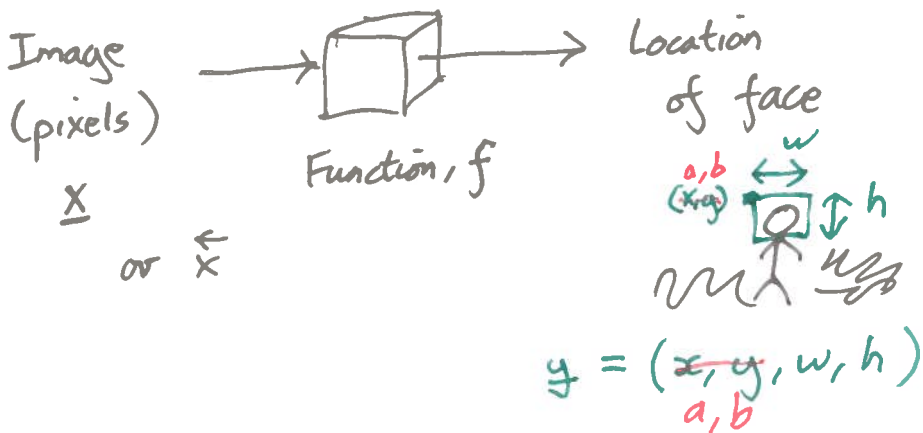
After class questions:

- here briefly
- outside ... or Inf. Forum
- office hour
- hypothesis
- email

← Please sign up and try it out.

Instructions on website.

Cartoon view of Machine Learning



Email (text) \rightarrow Label $f(\underline{x}) \in \{\text{spam}, \text{ok}, \text{phishing}\}$

Write f by hand?

if "Ray-Ban" in \underline{x} : spam $+ = 10$

if "IT Help Desk" in \underline{x} : phishing $+ = 100$

if "Bayesian" in \underline{x} : ok $+ = 10^6$


parameters $\underline{\theta}$ or \underline{w} "weights"

Auto-encoders



Goal is $f(\underline{x}) \approx \underline{x}$

Could have a bottleneck



Or

- Adding noise
- Or sparse A.C.

For the papers I showed you,
and the Acton rant, see the course notes.