R250 / Autoencoders and generative models

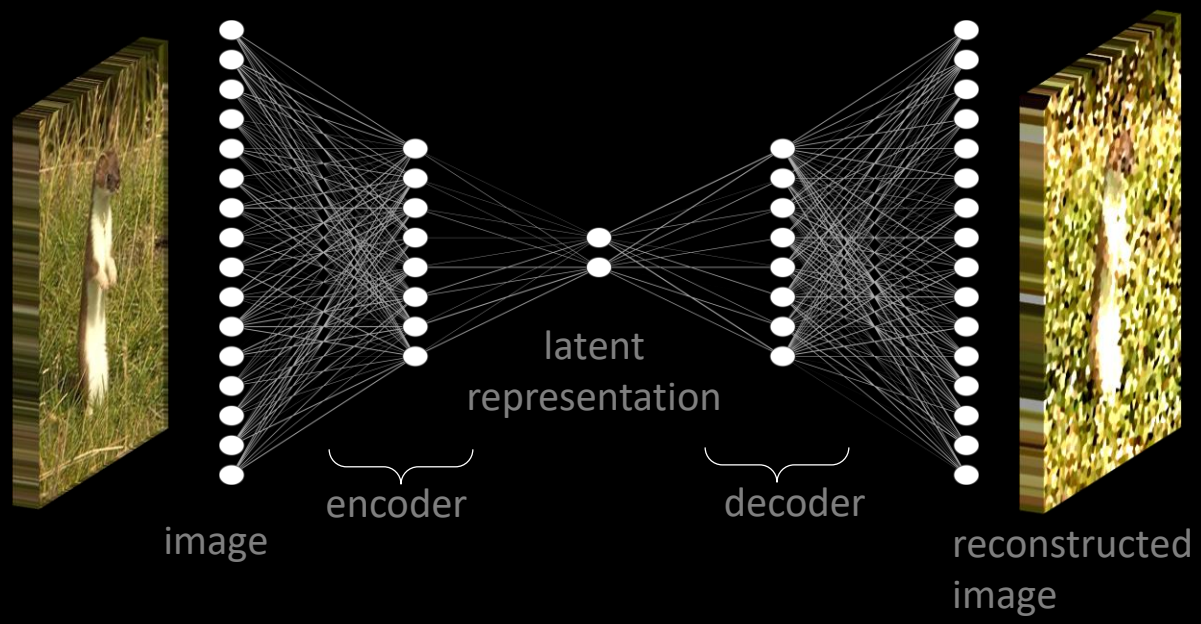

Training Task

given an input, reconstruct it

Constraint

the data must be squeezed through a low-dimension “bottleneck” layer

By training a neural network in this way, we hope

- the neural network learns a meaningful representation of the data
  (since how else could it reconstruct the image from only a few dimensions?)
- we can generate novel outputs
  (by simply feeding in arbitrary values at the bottleneck, and seeing what comes out)

But how do you train and cross validate a network of this sort?
(We can’t just measure prediction accuracy, the gold standard for supervised learning.)

Practical arrangements

- students will each present one paper
- there will be six papers, covering the three aspects
  (meaningful representations, generative models, evaluation)

- students will choose their coursework topic based on what interests them
  (with some suggestions from the lecturer)
  - replicate the results in an early paper,
    using newer methodology
  - analytical literature review,
    bringing together common threads across a range of papers
  - novel research:
    come up with an idea, do some investigation, write up a plan for further development