L173 Pre-reading response for session 4

D.J. Wischik

Read *Auto-encoding variational Bayes*, Kingma and Welling 2013, §§1–3 and figure 4 in appendix A, and *Visualizing data using t-SNE*, van der Maaten and Hinton 2008, §§1–2 and figure 2. Write a short response, about half a page pdf, answering these questions. LLM usage is permitted.

1. When we train a VAE, we learn an encoder; later at inference time we use our previously-trained encoder to compute the encoding of a new unseen datapoint. The t-SNE paper doesn't describe an equivalent inference-time version of its encoding; if it did, what would it be?

2. The goal in training a VAE is to fit a generative model, comprising both a decoder $p_\theta(x|z)$ and also a latent variable Z with its prior; this generative model has an associated encoder. Suggest a generative model that might be associated with your t-SNE encoder.

The questions are mostly about judgement, and ask you to read between the lines or think about the paper's unstated premise, or explain an argument that is implied rather than directly stated. You get the marks for interpretation and engaging critically with the paper, not for simply repeating what it says.

Be concrete. For example, point to a specific passage, equation, or figure, say why it matters and what it does. Use the course's notation, not the paper's. Answer from the papers and the lectures; facts brought in from elsewhere earn no marks, however true. Engaging critically means setting the paper against the lectures, not against the later literature.