Pre-Assignment for ML Engineer Candidates

Notes

- You may use LLMs.
- In the interview, we will ask questions using test data related to this assignment and will prioritize that over your written answers when evaluating.
- Your answers are meant to help us understand your reasoning. They need not be polished; handwritten answers are acceptable.

Problem 1

Fix integers T, d with $d < T$.

Symbol	Description
$\mathbf{X} \in \mathbb{R}^{T \times d}$	Input sequence (each row is a d-dimensional vector)
$\mathbf{W}^Q, \mathbf{W}^K \in \mathbb{R}^{d \times d}$	Learnable weight matrices

(i) Query and Key generation

$$\mathbf{Q} \;=\; \mathbf{X}\mathbf{W}^Q, \qquad \mathbf{K} \;=\; \mathbf{X}\mathbf{W}^K \quad \big(\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{T \times d}\big).$$

(ii) Score matrix

$$\boxed{\mathbf{S} \;=\; \frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}}$$

Problem 1-1

What is the maximum possible value of

$$\operatorname{rank}\big(\mathbf{S}\big)\ ?$$

Within the scope of $\mathbf{X}$, $\mathbf{W}^Q$, $\mathbf{W}^K$ as defined above, express the maximum possible rank of $\mathbf{S}$ using only $\mathbf{T}$ and $\mathbf{d}$.

Problem 1-2

Consider the scalar function defined by

$$L\big(\mathbf{W}^Q, \mathbf{W}^K\big) \;=\; \frac{1}{2}\,\|\mathbf{S}\|_F^2 \;=\; \frac{1}{2}\,\mathrm{tr}\big(\mathbf{S}^\top \mathbf{S}\big) \qquad \big(\mathbf{S} = \tfrac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\big).$$

Derive closed-form expressions for the gradients

$$\nabla_{\mathbf{W}^Q} L, \qquad \nabla_{\mathbf{W}^K} L$$

in terms of $\mathbf{X}, \mathbf{W}^Q, \mathbf{W}^K, d$ only.

Problem 2

Let $X_0 \sim \mathcal{N}(0,1)$ and $X_1 \sim \mathcal{N}(\mu, \sigma^2)$ be independent real-valued random variables, and fix $t \in (0,1)$. Define

$$X_t \;=\; (1-t)X_0 + tX_1.$$

Problem 2-1

Express

$$v(x,t) \;=\; \mathbb{E}\big[X_1 - X_0 \,\big|\, X_t = x\big]$$

in closed form in terms of μ, σ, t, x.

Problem 3

Background

This problem is based on a past example from one of our tasks; it is not one of our current tasks (which now include, e.g., towel folding), but the data below is real and the question stands on its own. We previously worked on a task of removing a pillow cover as part of the bed-making task. Specifically, as shown in Figure 1, the goal is to open the closed part of the pillow cover and bring it into a state where the pillow can be taken out when the opposite side is lifted. We refer to this task as the RPE (Remove Pillow Envelope) task. The objective was to improve the success rate of this task. In machine learning theory, it is known that generalization error decreases as the amount of data increases. Therefore, with the aim of reducing generalization error and improving success rate by increasing data, we conducted experiments to measure the relationship between data size, generalization error, and success rate.

The method used for the experiment was Behavior Cloning. First, teacher data was collected through human teleoperation. The input data consisted of camera images (with cameras mounted on the robot body, right hand, and left hand, all of which were used as input) and the current joint angles. The output was a sequence of action command values for 2 seconds into the future (50 FPS, totaling 100 frames), with a shape of (100, 16), where 100 represents the number of frames and 16 the action dimensions. This input-output relationship was learned using a Transformer-based model. The loss function employed was the L1 loss.

Experiments are conducted with four different numbers of episodes: 1432, 716, 358, and 179.

For cases with 716 episodes or fewer, episodes are randomly sampled to limit the number of episodes, and then the data is split into train and validation sets.

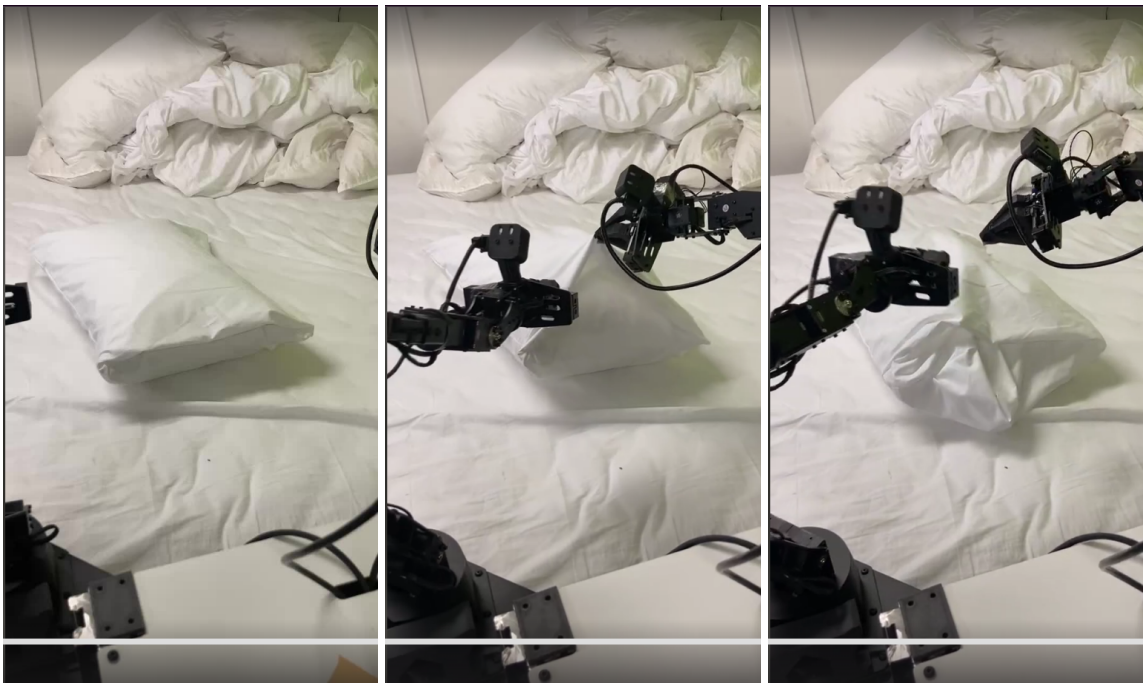

Figure 1: RPE task

Result

The measurement results of the success rate for each data size are shown below.

Episodes	Success rate
179	49%
358	14%
716	6%
1432	2%

Table 1: Number of episodes and corresponding success rate

Figure 2 shows a graph of L1 loss plotted against the number of data samples. The plain L1 represents the L1 loss on the training data, while val_l1 represents the L1 loss on the evaluation data. The training data and evaluation data are split on an episode basis.

Problem 3-1 Interpret the experimental results and discuss what is likely happening. Furthermore, based on the interpretation of these results, discuss what kind of research you would pursue in order to improve the success rate of the RPE task. For example, comprehensively describe the experiments you would like to conduct and the methods you would like to try. In addition, clearly indicate the priority order of these proposed research actions and explain the rationale for that prioritization.

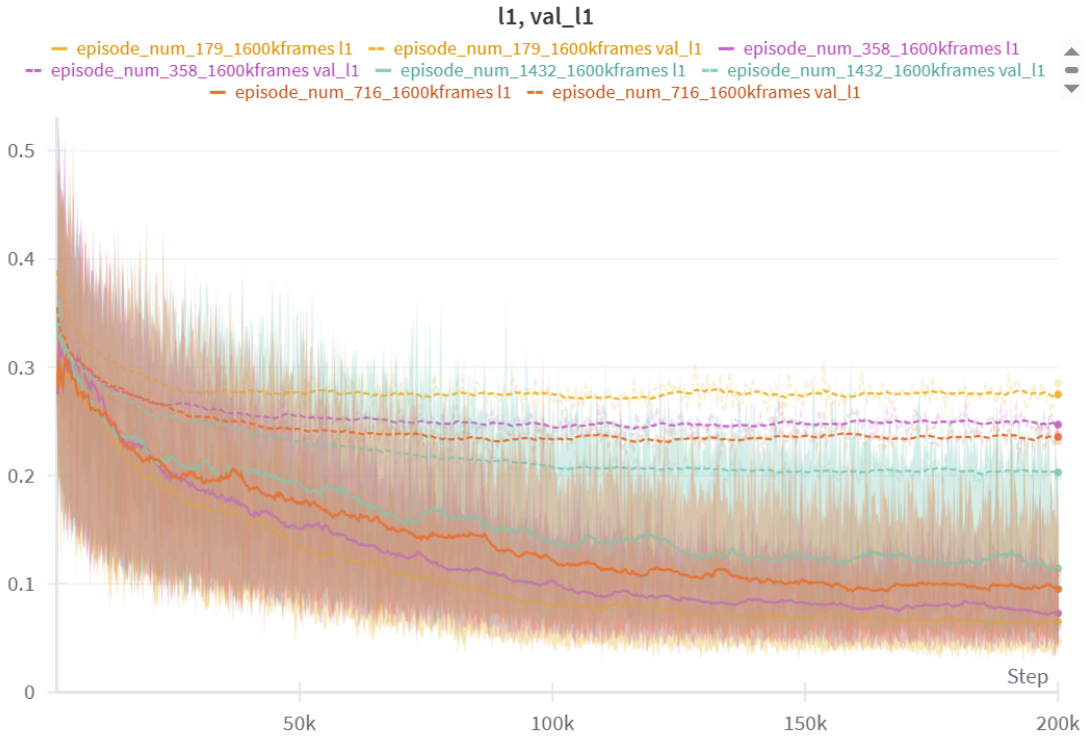

Figure 2: Episode number and l1 loss