AMCC 5160 · Quiz 11 (Review)

Four technique questions and one reflection, drawn from across the semester. 4 points each; 20 points total. Answer in your own words (1–3 sentences). Submit online at https://hyang.org/amcc5160/quiz-11/ before the end of the semester.

1. Technique: Drift in long video (Lecture 07), a world model that forgets the red door (Lecture 09), and a robot whose small errors compound (Lecture 09) share one cause. Name it, and describe one technique from the course that reduces it. (4 points)

2. Technique: Your main character's face changes from shot to shot in your A2 or A3. Choose two methods from Lectures 02, 05 or 08 (for example a reference image or IP-Adapter, LoRA, DreamBooth, or a motion latent like VASA-1). Say what each one supplies or trains, and give one trade-off between them. (4 points)

3. Technique: A 1024×1024 RGB image becomes a 64×64×64 latent. How many values is that before and after, and why is it not 16 times fewer? Then explain why 45 video frames become 12 latent frames. (Lecture 02) (4 points)

4. Technique: For a 4096×4096 layer, how many parameters does a full update have, and how many does a rank-16 LoRA have? Name one thing a small LoRA file does not guarantee. (Lectures 02 and 05) (4 points)

5. Reflection: Which single idea from this semester most changed how you make or look at work? Name the lecture, and give one concrete example from your own project. (4 points)

Where to look: the Lecture 11 slides and study guide link to every earlier slide you need: https://hyang.org/amcc5160/lecture-11/