AMCC 5160 · Study guide: ten lectures, one line each

Lecture 11 · 10 November 2026. If you can explain each line with an example from your own work, you are ready for the final presentation.

Lecture 01 · Sep 1 · Neural networks to AI filmmaking

A generated video is a learned prediction about change. Read any generator as data → representation → model → sampling.

- L01 · 8: Every generator can be read through four stages — https://hyang.org/amcc5160/lecture-01/#slide-8
- L01 · 44: Compress space and time before the transformer — https://hyang.org/amcc5160/lecture-01/#slide-44
- L01 · 102: Technical literacy becomes creative agency — https://hyang.org/amcc5160/lecture-01/#slide-102

Lecture 02 · Sep 8 · Generative editing

Write the edit contract (change, keep, adapt). Choose a control by the information that is missing. A video edit must survive the next frame.

- L02 · 9: The edit contract — https://hyang.org/amcc5160/lecture-02/#slide-9
- L02 · 29: Different controls carry different information — https://hyang.org/amcc5160/lecture-02/#slide-29
- L02 · 94: The frozen video wins the wrong test — https://hyang.org/amcc5160/lecture-02/#slide-94

Lecture 03 · Sep 15 · Inside the generator

The model predicts; the sampler takes a step. Seed, guidance and steps are different knobs; guidance extrapolates.

- L03 · 21: The model predicts. The sampler takes a step. — https://hyang.org/amcc5160/lecture-03/#slide-21
- L03 · 54: A small guidance calculation — https://hyang.org/amcc5160/lecture-03/#slide-54
- L03 · 64: Some models need only a few steps — https://hyang.org/amcc5160/lecture-03/#slide-64

Lecture 04 · Sep 22 · Video generation and art filmmaking

Frame time is not denoising time. Data filtering shapes what a model can make. The artist's intention comes first.

- L04 · 9: A video has two kinds of time — https://hyang.org/amcc5160/lecture-04/?slide=9#slides
- L04 · 15: Training data shapes the result — https://hyang.org/amcc5160/lecture-04/?slide=15#slides
- L04 · 37: The artist's intention comes first — https://hyang.org/amcc5160/lecture-04/?slide=37#slides

Lecture 05 · Sep 29 · Making the model yours

Textual Inversion, DreamBooth, Custom Diffusion, LoRA, IP-Adapter: what is trained? Memorization and concept ablation. Whose model is it?

- L05 · 27: DreamBooth vs. Textual Inversion — https://hyang.org/amcc5160/lecture-05/?slide=27#slides
- L05 · 61: Low-rank adaptation (LoRA) — https://hyang.org/amcc5160/lecture-05/?slide=61#slides
- L05 · 3: Whose model is it? — https://hyang.org/amcc5160/lecture-05/?slide=3#slides

Lecture 06 · Oct 6 · Quantization for faster model inference (TA talk)

Yaofu Liu's TA talk: quantization stores each number in fewer bits, so the same GPU does more per millisecond.

Lecture 07 · Oct 13 · Music, minutes, and machines

The unit of AI video became the scene. Autoregressive generation, KV cache, distillation, drift and Self Forcing. Keep your sources; pin your versions.

- L07 · 10: Two ways to make a long video — https://hyang.org/amcc5160/lecture-07/?slide=10#slides
- L07 · 21: Under the hood: the train-test gap — https://hyang.org/amcc5160/lecture-07/?slide=21#slides
- L07 · 65: A project folder that survives — https://hyang.org/amcc5160/lecture-07/?slide=65#slides

Lecture 08 · Oct 20 · AI video that talks

Synchresis. Two streams, one clock. Lip-sync and its scoreboard trap. Real-time budgets. Consent and the performer.

- L08 · 19: Inside LTX-2: two streams, one clock — https://hyang.org/amcc5160/lecture-08/?slide=19#slides
- L08 · 44: Reading the scoreboard (and its trap) — https://hyang.org/amcc5160/lecture-08/?slide=44#slides
- L08 · 65: Three tricks you already know — https://hyang.org/amcc5160/lecture-08/?slide=65#slides

Lecture 09 · Oct 27 · Worlds, 3D, and robots

A world model adds an action. Pixels or geometry (Gaussian splats). VLAs, compounding errors, and practice beyond copying.

- L09 · 8: What is a world model? — https://hyang.org/amcc5160/lecture-09/?slide=8#slides
- L09 · 34: Under the hood: only put stuff where stuff is — https://hyang.org/amcc5160/lecture-09/?slide=34#slides
- L09 · 48: Under the hood: small mistakes compound — https://hyang.org/amcc5160/lecture-09/?slide=48#slides

Lecture 10 · Nov 3 · Is it any good?

Four questions: AI? good? fair? allowed? What scores can and cannot see. Consent has layers. C2PA, watermarks, and the credit block.

- L10 · 13: The most useful slide tonight: four questions people mix up — https://hyang.org/amcc5160/lecture-10/?slide=13#slides
- L10 · 31: Worked example: the copycat gets a perfect score — https://hyang.org/amcc5160/lecture-10/?slide=31#slides
- L10 · 73: Useful: a credit block you can copy for A2 and A3 — https://hyang.org/amcc5160/lecture-10/?slide=73#slides

Four threads

1. How a generator works: data → representation → model → sampling; the model predicts, the sampler steps; seed, guidance, steps.
2. Control: the edit contract; choose a control by the missing information; what gets trained when you make a model yours.
3. Time: from the shot to the scene; drift in video, worlds and robots; change less, step less, store less, remember smartly.
4. Judgment: AI? good? fair? allowed?; every score has a trap; consent has layers; sign your work.