

Week 1 Optional Weekend Assignment

This weekend, choose Quick Review or Extended Review if you want additional practice. You may also choose no homework without penalty.

Optional and ungraded: Complete one assignment route only if it is useful to you. If you complete a route, use the form below to create a PDF and submit it in eCampus for formative feedback. Choosing no homework does not affect your grade, Quiz 1 eligibility, or participation.

Choose Your Weekend Route

Quick Review · 8-Minute Video + One Question

3Blue1Brown: Large Language Models explained briefly

<https://www.youtube.com/watch?v=LPZh9BOjkQs>

Optional assignment task: watch the video, then complete the Quick Review Response below by writing one question you still have.

Extended Review · 83 Minutes of Video + Writing

Andrej Karpathy: Deep Dive into LLMs like ChatGPT

<https://www.youtube.com/watch?v=7xTGNNLPyMI>

Optional assignment task: watch 00:01–00:31, 00:59–01:42, and 03:21–03:31, then answer the six Extended Review prompts below. The 83-minute label covers video only; allow additional time for writing.

No Homework · No Penalty

You may stop after the Week 1 class activities. Bring your ungraded Week 1 evidence card back only if it helps your recap.

Choosing this route does not affect your course grade or participation in Week 2.

Deep Dive · Optional Enrichment

Watch the complete 3-hour-31-minute Karpathy lecture and use all 18 timestamped checkpoints.

This route is beyond the optional weekend assignment. The time label covers video only; allow additional time for checkpoints. You may split the work across sittings or change playback speed.

Assignment boundary: Quick Review means the 8-minute video plus one question. Extended Review means the three selected Karpathy segments plus six responses. The baseline, glossary, Deep Dive checkpoints, and self-checks are optional supports rather than additional assignment requirements.

Optional reference: Karpathy's MicroGPT article shows the same GPT ingredients in a tiny setting: documents, tokenizer, model, training loop, and inference loop. Use it only if you want a code-level preview after the class concepts.

<https://karpathy.github.io/2026/02/12/microgpt/>

Optional Baseline

In one sentence, what do you think ChatGPT is doing when it answers a question?

Tiny Glossary

Term	Meaning for this course
Token	A chunk of text the model processes. It can be a word, part of a word, punctuation, or code fragment.
Context window	The current working memory: the text the model can see while answering.
Pretraining	Learning broad patterns from massive text and code data.
SFT	Supervised fine-tuning: training from example conversations or desired outputs.
RLHF	Reinforcement learning from human feedback: using human preferences to shape model behavior.
Inference	Running the trained model to generate output, usually one token at a time.
Hallucination	A fluent answer that is unsupported, misleading, or false.
Tool	An external action the system lets the model request, such as search, calculation, or code execution.
Guardrail	A rule, approval step, or check that limits risky behavior.

Deep-Dive Route: 18 Optional Checkpoints

Timestamps are based on the current YouTube chapters for "Deep Dive into LLMs like ChatGPT." Use them as anchors, not as exact quiz boundaries.

00:01:00 Pretraining data (internet)

What kinds of data help train an LLM? Why does data quality matter?

00:07:47 Tokenization

What is a token? Give one example where tokenization might surprise a user.

00:14:27 Neural network input/output

What goes into the model, and what comes out?

00:20:11 Neural network internals

You do not need the math. What is the useful mental model for what the network has learned?

00:26:01 Inference

Why does the answer appear one piece at a time?

00:31:09 GPT-2: training and inference

What does a base model do? How is that different from a helpful chat assistant?

00:42:52 Llama 3.1 base model inference

What did you notice about interacting with a base model?

00:59:23 Pretraining to post-training

Why is pretraining not enough if we want a useful assistant?

01:01:06 Post-training data (conversations)

How do conversation examples change the behavior of a model?

01:20:32 Hallucinations, tool use, knowledge, working memory

Why can a model sound confident and still be wrong? How can tools help?

01:41:46 Knowledge of self

Why might a model be confused about what it is, what version it is, or what it can access?

01:46:56 Models need tokens to think

Why can writing intermediate reasoning or scratch work help a model perform better?

02:01:11 Tokenization revisited

Why can spelling, counting letters, or manipulating strings be surprisingly hard?

02:04:53 Jagged intelligence

What does it mean for a model to be brilliant at one task and weak at another?

02:07:28 SFT to reinforcement learning

What changes when the system moves beyond examples and starts using feedback or rewards?

02:14:42 Reinforcement learning

Explain reinforcement learning in one sentence using your own words.

02:48:26 RLHF

Why do human preferences matter when building assistant-like models?

03:21:46 Grand summary

Write three takeaways you want to remember before learning agents.

Bridge to Agents

Use this course definition:

Practical agent = LLM + goal + context + tools + loop + guardrails

Classify each example as **chatbot answer**, **tool call**, or **agent loop**.

Example	Your classification	Why?
ChatGPT explains the difference between pretraining and post-training.	Choose one v	
A model asks a calculator tool to compute $438 * 27$.	Choose one v	
A coding system reads files, proposes a plan, edits code, runs tests, and asks before deploying.	Choose one v	

Verification Habit

Pick one important claim from the video or from a chatbot answer. How would you verify it?

Exit Ticket

In 4-6 sentences, explain why an agent needs guardrails.

Optional Self-Check

- ☐ I can explain tokens, context, and next-token generation without copying the glossary.
- ☐ I can explain why fluent output still needs verification.
- ☐ I can distinguish a chatbot answer, a tool call, and an agent loop.

☐ I have one question or example I may bring to Week 2.

Save Your Full Worksheet

This creates a PDF of **every** answer area on this page, including Optional Baseline, all 18 Deep-Dive checkpoints, Bridge to Agents, Verification Habit, Exit Ticket, and Self-Check. Unanswered areas are printed as *(not answered)* so nothing you typed is left out. Keep this for your own records; the eCampus submission still uses the Quick or Extended PDF above.

Create full worksheet PDF

Every answer area on this page saves automatically on this device.

Feedback boundary: If you complete a route, submit the generated PDF in eCampus for formative feedback. No route is scored, recorded as extra credit, required for Quiz 1, or required for full participation in Week 2.