BLACKSKY LABS · BLACKSKYUP

LLM 101

Large Language Models for Students

GRADES	5 – 12
TIERS	Middle School · High School
FORMATS	Workshop · Intensive · Weekly · Residency
VERSION	April 2026

Developed by Blacksky Labs · blackskymedia.org · mario@blackskymedia.org

Table of Contents

LLM 101: Large Language Models for Students

A K-12 Curriculum Framework

Developed for: Grades 5–12 Prepared by: Blacksky Media Version: 1.0 | April 2026

About This Curriculum

LLM 101 introduces students to Large Language Models — the technology behind tools like ChatGPT, Claude, Gemini, and others — through age-appropriate, hands-on learning experiences. This framework is organized into two tiers and four delivery formats, allowing educators to deploy the curriculum in settings ranging from a single-day event to a full semester residency.

How to Use This Document

Two Student Tiers: - Tier 1 — Middle School (Grades 5–8): Concept-first, analogy-driven, creative and exploratory - Tier 2 — High School (Grades 9–12): Technically deeper, critically analytical, career-connected

Four Delivery Formats (per tier):

Format	Duration	Best For
Single-Day Workshop	1 day (~6 hrs)	Intro events, STEM fairs, career days, admin demos
Multi-Day Intensive	3–5 consecutive days	Summer programs, school breaks, bootcamps
Weekly Series	6–12 weekly sessions	Semester elective, after-school program

Format	Duration	Best For
Semester Residency	16–18 weeks	Full elective course, co-instructor, portfolio outcomes

Topic Progression by Format

Topics scale with depth and time available. Each format builds on the previous.

Topic	Workshop	Intensive	Weekly	Residency
What is AI? What is an LLM?				
How LLMs process language (tokens, prediction)	Intro			
Prompt engineering basics				
Prompt engineering advanced	—	Intro		
AI ethics & bias	Discussion			
AI in the real world / careers				
Critical thinking & AI-generated content	—			
Training data & how models learn	—	Intro		
Fine-tuning & model customization	—	—	Intro	
AI policy & society	—	—		
Capstone project / student portfolio	—	Mini		

TIER 1 · GRADES 5–8

TIER 1: MIDDLE SCHOOL (Grades 5–8)

Tone: Curious, playful, accessible. Heavy use of analogy, storytelling, and hands-on exploration. Minimal jargon — when technical terms are introduced, they are always connected to something familiar.

DELIVERY FORMAT

Format A — Single-Day Workshop 1 day · ~6 hrs

Duration: 6 hours (with breaks) Ideal Setting: STEM fair, career day, intro event, or administrator buy-in demo Class Size: Up to 30 students Materials: Devices with internet access (1 per 2 students minimum), printed activity sheets, whiteboard or chart paper

---

Learning Objectives

By the end of the day, students will be able to: - Explain what an AI and an LLM is in their own words - Describe at least one way LLMs are already part of their daily lives - Write and improve a prompt to get a better response from an AI tool - Identify one ethical consideration when using AI

Schedule Overview

Time	Block	Activity
9:00–9:20	Welcome	Icebreaker + "What do you already know about AI?" poll
9:20–10:15	Block 1	What Is an LLM?
10:15–10:30	Break	
10:30–11:30	Block 2	Talk to the Machine (Prompt Play)
11:30–12:15	Lunch	

Time	Block	Activity
12:15–1:00	Block 3	Can You Trust It? (Ethics & Critical Thinking)
1:00–1:45	Block 4	The Challenge Round
1:45–2:00	Wrap-Up	Showcase + Reflection

Block 1: What Is an LLM? (55 min)

Core Concepts: - Artificial intelligence vs. human intelligence - What makes something a "language model" - How LLMs predict the next word (autocomplete analogy) - Examples students already use: autocomplete on phones, smart replies, voice assistants

Key Analogy — The Giant Guesser:

> ***"Imagine you've read every book, article, and website ever written. Someone asks you to finish a sentence. You don't think — you just guess the most likely next word based on everything you've ever read. That's basically what an LLM does, billions of times per second."***

Activities: 1. Word prediction warm-up (10 min): Students complete sentences on paper, then compare their answers to the AI's. Discuss: why did the AI guess differently? 2. LLM timeline (15 min): Quick visual tour — from early chatbots (ELIZA, 1966) to today. Students place AI milestones on a class timeline. 3. Discussion (10 min): "What can an LLM do? What can't it do?" — class brainstorm on whiteboard.

Block 2: Talk to the Machine — Prompt Play (60 min)

Core Concepts: - What is a prompt? - How the way you ask changes what you get - Iteration: making prompts better

Prompt Engineering Concepts (Middle School Level): - Be specific: "Write a poem" vs. "Write a 4-line rhyming poem about space for a 6th grader" - Give context: Tell the AI who you are and what you need - Ask for a format: "Make it a list" or "Explain it like I'm 10"

Activities: 1. Guided prompt exploration (20 min): Students try a provided set of prompts with increasing specificity. They record what changed in the response. 2. Prompt makeover (25 min): Each student or pair receives a "bad prompt" card. Their job: rewrite it to get a better answer. Share results with the class. 3. Debrief (15 min): What made a prompt good? Class builds a shared "Prompt Recipe" on the board.

Sample Bad Prompt Cards: - "Tell me about animals" - "Help me with my project" - "Write something funny" - "Explain history"

Block 3: Can You Trust It? (45 min)

Core Concepts: - LLMs can be wrong (hallucinations) - LLMs reflect their training data — including bias - Importance of fact-checking AI output

Activities: 1. Spot the mistake (20 min): Students are given AI-generated paragraphs with intentional errors (wrong dates, made-up facts). They act as "AI fact-checkers." 2. Discussion: Who's responsible? (25 min): Scenario cards — students discuss in small groups and share out.

Sample Scenario Cards: - "You ask the AI to write your book report. You turn it in. The teacher finds out. Whose fault is it?" - "The AI gives someone wrong medical advice. Who is responsible?" - "The AI describes a group of people in a stereotyped way. Where did that come from?"

Block 4: The Challenge Round (45 min)

Format: Team challenge — groups of 4–5 students compete in a fun LLM prompt challenge.

Rounds: 1. Shortest prompt, best answer: Who can get the best poem with the fewest words? 2. Break the AI: Who can get the AI to say something strange or unexpected (within school guidelines)? 3. AI as tutor: Use the AI to explain the water cycle to a first grader. Judge: clearest explanation wins. 4. Speed round: First team to get the AI to produce a specific output (haiku about lunch, a 3-step recipe, etc.)

Wrap-Up & Reflection (15 min)

- Exit ticket: "One thing I learned, one thing I'm still wondering, one thing I want to try"
- Share 2–3 student favorites from the day
- Resources handout for students to explore at home

Assessment

- Informal: Observation during activities, class discussions
- Exit ticket collected for facilitator review
- Optional: Photo/screenshot of students' best prompt of the day

---

DELIVERY FORMAT

Format B — Multi-Day Intensive 3–5 days · ~3–4 hrs/day

Duration: 3–5 consecutive days, approximately 3–4 hours per day Ideal Setting: Summer programs, school break camps, workforce bootcamps Class Size: Up to 25 students Materials: Devices (1 per student preferred), journals or notebooks, printed activity packets

Learning Objectives

By the end of the intensive, students will be able to: - Explain what LLMs are, how they learn, and what they can and cannot do - Write effective prompts for multiple real-world use cases - Analyze AI-generated content for accuracy, bias, and usefulness - Articulate at least two ethical considerations surrounding AI - Complete a mini-project using AI as a tool

Day-by-Day Overview

Day 1 — Meet the Machine (~3.5 hrs)

Theme: What is AI and why does it matter?

Morning (90 min): - Welcome, group norms, and icebreaker - Lesson: A brief history of AI — from rule-based systems to neural networks to LLMs - Analogy deep-dive: The Library Brain — LLMs as a compression of human text knowledge - Activity: Student "AI encounter" maps — where have you already met AI today?

Midday (30 min break)

Afternoon (90 min): - Lesson: How does an LLM actually work? Tokens, probability, and prediction (no math required) - Analogy: "Mad Libs" as a simplified model of prediction - Activity: Token scavenger hunt — students break sample sentences into "tokens" using an online tokenizer tool - Reflection journal: "What surprised me about how LLMs actually work?"

Day 2 — Learning to Speak AI (~3.5 hrs)

Theme: Prompt engineering — the art of asking well

Morning (90 min): - Lesson: Anatomy of a great prompt — role, context, task, format, constraints - Demo: Live prompt comparison — same task, three different prompts, three very different results - Activity: "Prompt

Surgery" — students receive weak prompts and operate to improve them

Midday (30 min break)

Afternoon (90 min): - Activity: Prompt Engineering Challenges (stations) - Station 1: Make the AI tell a story in your favorite genre - Station 2: Use the AI to help study for a test (student picks the subject) - Station 3: Ask the AI to help you plan a school event - Station 4: Use the AI to explain something you've always wondered about - Group share-out: Best prompts and what made them work

---

Day 3 — AI Can Be Wrong (and Biased) (~3.5 hrs)

Theme: Critical thinking and AI ethics

Morning (90 min): - Lesson: What is a hallucination? Why does the AI make things up? - Analogy: The overconfident guesser who doesn't know when to say "I don't know" - Activity: Hallucination Hunt — students fact-check AI responses on assigned topics using credible sources - Discussion: How should we respond when AI is wrong?

Midday (30 min break)

Afternoon (90 min): - Lesson: Bias in AI — where does it come from? What does it look like? - Activity: Bias audit — students ask the AI the same question in different ways and look for patterns - E.g., "Describe a nurse," "Describe a CEO," "Describe a criminal" - Ethics fishbowl discussion: "Should AI be used to grade student work? Why or why not?" - Reflection journal: "Something AI does that I think needs to change"

Day 4 — AI in the Real World (~3.5 hrs)

Theme: Applications, careers, and the future

Morning (90 min): - Lesson: Where is AI being used right now? (Healthcare, education, art, science, law, etc.) - Guest speaker or video panel: People working with AI in different careers (can be facilitated via recorded video if in-person not possible) - Activity: "AI job board" — students research and present a real job that involves working with AI

Midday (30 min break)

Afternoon (90 min): - Lesson: What can AI NOT do? What makes humans irreplaceable? - Creativity, empathy, embodied experience, moral judgment - Activity: Human vs. AI creativity challenge — students and AI create something on the same prompt. Class votes on which they prefer (blind). Discussion follows. - Mini-project introduction and planning session

---

Day 5 — Mini-Project & Showcase (~3.5 hrs)

Theme: Build something, share something

Morning (90 min): - Mini-project work time (see options below) - Facilitator circulates for coaching and feedback

Midday (30 min break)

Afternoon (90 min): - Project presentations (2–3 min each) - Peer feedback using "warm/cool/curious" format - Reflection: "What I'll take with me from this week" - Certificates and celebration

Mini-Project Options (student choice): 1. Prompt Portfolio: A curated set of 5+ prompts with reflections on what worked, what didn't, and why — plus screenshots of the best AI responses 2. AI Explainer: A short video, poster, or written piece explaining how LLMs work to a younger student 3. AI Story: A short story co-written with an AI — student writes a reflection on the collaboration experience 4. AI Audit Report: An investigation into bias or errors found in AI tools — with evidence and recommendations

---

Assessment

- Daily reflection journals (formative)
- Mini-project (summative)
- Optional: Peer evaluation form for project showcase

DELIVERY FORMAT

Format C — Weekly Series 6–12 weekly sessions

Duration: 6–12 weekly sessions, approximately 60–75 minutes each Ideal Setting: After-school program, semester elective, enrichment course Class Size: Up to 25 students Materials: Devices, student notebooks/journals, session handouts

Note: This format can be run as a 6-session course (core sessions only) or extended to 12 sessions by adding the enrichment sessions marked with .

Learning Objectives

By the end of the series, students will be able to: - Explain how LLMs work using their own analogies - Write sophisticated, multi-part prompts for a variety of purposes - Critically evaluate AI output for accuracy, bias, and usefulness - Discuss AI ethics with nuance and evidence - Complete and present an independent AI project

---

Session Map

Session 1: Welcome to the Age of AI

Focus: Building a shared foundation and curiosity Key Concepts: What is AI, what is an LLM, brief history Activities: - "What do you already know?" KWL chart (Know / Want to know / Learned) - Speed dating with AI tools — students briefly try 3 different AI tools and note first impressions - Discussion: "What's the same? What's different between the tools?" Homework: Find one place you encountered AI this week (outside school). Be ready to share.

Session 2: How LLMs Actually Work

Focus: Demystifying the technology Key Concepts: Training data, tokens, probability, prediction Activities: - Token activity: manually tokenizing sentences and predicting "next tokens" - Video + discussion: A visual walkthrough of how neural networks learn (3Blue1Brown or similar, age-appropriate) - Analogy building: Students create their own analogy for how LLMs work Homework: Write your LLM analogy and explain it to a family member. Report back.

Session 3: The Art of the Prompt — Part 1

Focus: Introduction to prompt engineering Key Concepts: Specificity, context, format, iteration Activities: - Live demo: Same task, 5 prompts, 5 very different responses - Practice: Prompt improvement workshop (pairs) - Class builds a shared "Prompt Cheat Sheet" Homework: Use an AI tool for something useful this week. Screenshot your best prompt + response.

Session 4: The Art of the Prompt — Part 2

Focus: Advanced prompt strategies Key Concepts: Role assignment, chain of thought, few-shot prompting (concept only) Activities: - Review homework screenshots — class vote on "best prompt of the week" - Lesson: Giving the AI a role ("Act as a coach," "You are a historian") - Challenge: Students use role prompting for a real school task (homework help, research, writing feedback) Homework: Prepare one question for next week's ethics discussion.

Session 5: Can You Trust the AI?

Focus: Accuracy, hallucinations, and critical evaluation Key Concepts: Hallucinations, verification, AI limitations Activities: - Hallucination hunt in teams — fact-check AI-generated content on assigned topics - Discussion: When is AI useful for research? When is it dangerous? - Create a personal "AI Use Policy" — a set of rules for how you'll use AI responsibly Homework: Bring an AI-generated piece of writing to next class for a bias audit.

Session 6: Bias and Ethics in AI

Focus: Where bias comes from and what to do about it Key Concepts: Training data bias, algorithmic bias, AI ethics frameworks Activities: - Bias audit of homework submissions — what patterns do we notice? - Scenario-based ethics discussion (small groups + debrief) - Reflection: "What changes would I make to how AI is built or used?"

For the 6-session format: Begin capstone project planning. End with showcase in session 6.

Session 7: AI and Creativity

Focus: AI as creative collaborator Key Concepts: Generative AI in art, music, and writing; human vs. AI creativity Activities: - Experiment with AI image or music generation tools - Co-create a piece with AI — poem, short story, or visual prompt - Discussion: "Does AI create, or does it remix?"

Session 8: AI in the Real World — Applications

Focus: Real-world use cases across industries Key Concepts: AI in healthcare, education, law, science, media Activities: - "Industry fair" jigsaw — each group researches one industry and presents to peers - Discussion: Benefits, risks, who wins, who loses

---

Session 9: AI, Power, and Society

Focus: Systemic and societal implications Key Concepts: Data privacy, surveillance, environmental cost, digital divide Activities: - Case study analysis: One of the following — facial recognition in schools, AI hiring tools, AI content moderation - Students write a 1-page "policy memo" with a recommendation

---

Session 10: Humans and AI Together

Focus: Collaboration, not replacement Key Concepts: Augmentation vs. automation, future of work, irreplaceable human skills Activities: - Debate: "AI will take our jobs" — for and against - Reflection: "What skills do I want to develop that AI can't replace?"

Session 11: Capstone Project Work Session

Focus: Independent project development Activities: - Project work time with facilitator coaching - Peer feedback sessions (pairs) - Project refinement

---

Session 12: Showcase & Celebration

Focus: Presentations and reflection Activities: - Student project presentations (3–5 min each) - Peer and facilitator feedback - Class reflection: "What changed in how I think about AI?" - Certificates and celebration

Capstone Project Options (same as Multi-Day Intensive, with added options): 1. Prompt Portfolio 2. AI Explainer (video, poster, or written) 3. AI Story (co-written with AI + reflection) 4. AI Audit Report 5. AI Policy Proposal — a written proposal for how AI should be used in your school or community 6. AI Teaching Session — student teaches a younger student or family member about LLMs

---

Assessment

- Session participation and reflections (formative)
- Homework assignments (formative)
- Capstone project (summative)
- Optional: Peer evaluation rubric

DELIVERY FORMAT

Format D — Semester Residency 16–18 weeks

Duration: 16–18 weeks; 2–3 class sessions per week Ideal Setting: Full elective course, co-instructor model, or standalone AI literacy class Class Size: Up to 25 students Materials: Devices, student portfolios (physical or digital), full curriculum packet

---

Learning Objectives

By the end of the residency, students will be able to: - Explain how LLMs work at a conceptual and functional level - Design and execute sophisticated prompt strategies across multiple domains - Conduct an independent analysis of AI output for accuracy, bias, and ethics - Connect AI development to social, historical, and political contexts - Produce a substantial capstone portfolio demonstrating mastery

---

Unit Overview

Unit	Title	Weeks	Focus
1	Introduction to AI and LLMs	1–3	Foundation, history, how it works
2	Prompt Engineering	4–7	Skills-based mastery of prompting
3	AI Applications	8–10	Real-world uses and hands-on exploration
4	Ethics, Bias, and Society	11–14	Critical thinking and social impact
5	Capstone Project	15–18	Independent project + portfolio

Unit 1: Introduction to AI and LLMs (Weeks 1–3)

Essential Questions: - What is intelligence — human or artificial? - How did we get here? What came before LLMs? - How does an LLM actually produce language?

Week 1 — What Is AI? - History of AI: rule-based systems, machine learning, deep learning, LLMs - Turing Test activity: Can you tell human from machine? - KWL charts and unit overview

Week 2 — How LLMs Work - Training data: what is it, where does it come from? - Tokenization, embeddings, and probability (conceptual) - Token lab: students experiment with a tokenizer tool - Analogy project: students create and present their own LLM analogy

Week 3 — LLMs in Practice - Survey of major LLM tools (ChatGPT, Claude, Gemini, Copilot) - Comparative testing: same prompts across multiple tools - Discussion: Why do they give different answers? - Vocabulary quiz and Unit 1 reflection

---

Unit 2: Prompt Engineering (Weeks 4–7)

Essential Questions: - How does the way I communicate change what I get back? - What does it mean to be a skilled AI communicator? - How can I use AI as a genuine learning tool without cheating myself?

Week 4 — Prompt Basics - Components of a prompt: role, context, task, format, constraints - Bad prompt / good prompt comparison activities - Prompt iteration practice

Week 5 — Prompt Strategies - Role prompting, chain-of-thought prompting, few-shot prompting (concept-level) - Use cases: study help, creative writing, research, planning - Challenge: "Best prompt of the week" competition

Week 6 — Prompting for Learning - Using AI ethically for homework and studying - AI as Socratic tutor: asking AI questions vs. asking for answers - Discussion: Academic integrity and AI - Students draft their personal AI use policy

Week 7 — Prompt Mastery Assessment - Prompt challenge exam: students complete timed prompt tasks with rubric - Reflection: "How has my prompting changed since Week 4?" - Peer feedback on each other's prompt portfolios

---

Unit 3: AI Applications (Weeks 8–10)

Essential Questions: - Where is AI already changing the world? - What are the opportunities and risks of AI in different fields? - What role do humans still play?

Week 8 — AI in Creative Fields - AI in art, music, writing, film - Activity: Create a collaborative piece with AI - Discussion: Authorship, originality, and credit

Week 9 — AI in Science and Society - AI in medicine, climate science, education, law enforcement - Guest speaker or video panel - "Industry research" presentations (student groups)

Week 10 — Humans and AI Together - Augmentation vs. automation: what's the difference? - Skills humans bring that AI cannot replicate - Future of work discussion and career exploration - Unit 3 reflection

---

Unit 4: Ethics, Bias, and Society (Weeks 11–14)

Essential Questions: - Where does AI go wrong? Who gets hurt? - What responsibilities do creators and users of AI have? - What should AI policy look like?

Week 11 — AI Accuracy and Hallucinations - Deep dive on hallucinations: why they happen, real-world examples - Research and verification skills - Hallucination documentation project: students track and analyze AI errors

Week 12 — Bias in AI - Sources of bias: training data, design choices, feedback loops - Bias audit activity: testing AI across demographic categories - Discussion: Whose voices are in the training data? Whose are missing?

Week 13 — AI Ethics Frameworks - Introduction to ethics frameworks: utilitarianism, rights-based thinking, fairness - Applied to AI: how should AI be designed and governed? - Case study analysis: one major AI ethics controversy (facial recognition, deepfakes, hiring algorithms, etc.)

Week 14 — AI Policy and Advocacy - Who regulates AI? What laws exist? - Student policy proposal drafts - Peer review of policy proposals - Unit 4 reflection

---

Unit 5: Capstone Project (Weeks 15–18)

Essential Questions: - What do I know about AI that I want to share with the world? - How can I demonstrate what I've learned in a meaningful way?

Week 15 — Project Planning - Project options overview and selection - Proposal drafts and facilitator approval - Research and resource gathering

Week 16 — Project Development - Work sessions with facilitator coaching - Milestone check-in: rough draft or prototype due

Week 17 — Refinement and Peer Feedback - Peer review sessions using structured feedback protocol - Final revisions - Portfolio assembly

Week 18 — Showcase and Celebration - Public or in-school showcase - Presentations (5–7 min each) - Portfolio submission - Certificates and celebration

Capstone Project Options: 1. Prompt Portfolio (curated, annotated, with reflection essay) 2. AI Explainer Video or Podcast Episode 3. AI Audit Report (research-based, with evidence and recommendations) 4. AI Policy Proposal (2–3 pages with research support) 5. AI-Assisted Creative Work + Artist Statement 6. AI Teaching Curriculum (design a mini-lesson to teach peers or younger students)

Semester Residency Assessment Structure

Component	Weight	Description
Unit Reflections	20%	End-of-unit written reflections
Prompt Portfolio	20%	Ongoing collection of prompts with annotations
Participation & Discussions	15%	Engagement in class discussions and activities
Unit Projects (x2)	20%	Mid-residency assessed projects (bias audit, policy memo)
Capstone Project	25%	Final project + presentation

TIER 2 · GRADES 9–12

TIER 2: HIGH SCHOOL (Grades 9–12)

Tone: Analytically rigorous, career-connected, critically sophisticated. Encourages genuine questioning of AI narratives. Technical depth increases with format length. Students are treated as emerging adults who will shape how AI is used.

DELIVERY FORMAT

Format A — Single-Day Workshop 1 day · ~6 hrs

Duration: 6 hours (with breaks) Ideal Setting: STEM fair, career day, intro event, administrator or community demo Class Size: Up to 30 students Materials: Devices with internet access (1 per student preferred), printed activity sheets

---

Learning Objectives

By the end of the day, students will be able to: - Explain the architecture of LLMs at a conceptual level - Write and iterate on effective prompts for professional and academic contexts - Critically evaluate AI output for accuracy, bias, and reliability - Articulate the social and ethical stakes of AI at a systemic level

---

Schedule Overview

Time	Block	Activity
9:00–9:20	Welcome	Icebreaker + baseline knowledge survey
9:20–10:15	Block 1	How LLMs Actually Work
10:15–10:30	Break	
10:30–11:30	Block 2	Prompt Engineering for Real Life
11:30–12:15	Lunch	

Time	Block	Activity
12:15–1:00	Block 3	AI and Society: Who Benefits? Who Doesn't?
1:00–1:45	Block 4	The Challenge Round (Advanced)
1:45–2:00	Wrap-Up	Reflection + Career Connections

Block 1: How LLMs Actually Work (55 min)

Core Concepts: - Transformer architecture (conceptual overview, no math) - Training data, parameters, and scale - The difference between LLMs and traditional search engines - Why LLMs hallucinate — the statistical prediction problem

Activities: 1. Pre-assessment discussion (10 min): "What's happening when you type a message to ChatGPT?" Students share theories, facilitator annotates. 2. Visual walkthrough (20 min): Animated or visual explanation of: training → tokenization → embedding → attention → output. Students take notes on a provided diagram. 3. Predict-the-token live demo (15 min): Use a token probability visualizer to show students the model's top predicted tokens in real time. 4. Discussion (10 min): "If an LLM is just predicting the next most-likely word, why does it seem so smart? What does that tell us about intelligence?"

Block 2: Prompt Engineering for Real Life (60 min)

Core Concepts: - Advanced prompt structure: system prompts, user prompts, context windows - Prompt patterns: chain-of-thought, few-shot, role prompting - Professional and academic use cases

Activities: 1. Prompt anatomy lesson (15 min): Breaking down an expert-level prompt — role, background, task, format, constraints, examples 2. Career prompt challenge (30 min): Students choose a career field and design a series of prompts as if they were a professional in that field - E.g., a journalist prompting for research, a nurse prompting for patient education materials, a lawyer prompting for case summaries 3. Peer review (15 min): Students exchange prompts and evaluate using a provided rubric

---

Block 3: AI and Society — Who Benefits? Who Doesn't? (45 min)

Core Concepts: - Who builds AI? Who controls it? - Disparate impact of AI errors - Environmental costs of AI infrastructure - The labor behind AI (data labeling, content moderation)

Activities: 1. Short reading + discussion (25 min): Students read a 1-page case study (provided) on one AI ethics controversy. Small groups discuss: What happened? Who was harmed? What should have been done differently? 2. Structured debate (20 min): "AI development is moving faster than society can handle." Students argue both sides.

Block 4: The Challenge Round — Advanced (45 min)

Format: Individual or pair challenge across 4 stations (10 min each)

Stations: 1. Adversarial prompting defense: Given an AI response that seems convincing but contains errors, students must identify the flaws and verify claims 2. Prompt for impact: Design a prompt that uses AI to address a real problem in your school or community 3. Bias audit sprint: Test the AI on 5 prompts designed to reveal potential bias. Document findings 4. Future of work: Use the AI to identify 3 skills you should develop that AI cannot replicate — then critique the AI's answer

---

Wrap-Up (15 min)

- Exit ticket: "One thing I now understand that I didn't before; one question I'm taking with me"
- Career connection: "These are real jobs where people use AI daily..."
- Resources and next steps handout

Assessment

- Informal observation and discussion participation
- Exit ticket
- Optional: Career prompt submitted for portfolio

DELIVERY FORMAT

Format B — Multi-Day Intensive 3–5 days · ~3–4 hrs/day

Duration: 3–5 consecutive days, approximately 4 hours per day Ideal Setting: Summer programs, workforce readiness bootcamps, school break intensives Class Size: Up to 25 students Materials: Devices (1 per student), notebooks, printed activity packets

---

Learning Objectives

By the end of the intensive, students will be able to: - Explain the technical architecture of LLMs at a conceptual level - Design sophisticated, multi-component prompts for professional and academic contexts - Conduct a structured audit of AI output for accuracy and bias - Analyze an AI ethics issue with evidence and structured reasoning - Produce a mini-project demonstrating applied understanding

---

Day-by-Day Overview

Day 1 — Architecture and Foundations (~4 hrs)

Morning (2 hrs): - Welcome and norms - History of AI: from symbolic systems to statistical models to neural networks to transformers - Technical deep-dive (conceptual): Transformers, attention mechanisms, context windows - Analogy: attention as "knowing which words in a sentence matter most when predicting the next one" - Activity: Reading and annotating a simplified technical diagram of an LLM pipeline

Afternoon (2 hrs): - Training data deep-dive: What goes into an LLM? Web crawls, books, code, conversations - Activity: "What's in the training data?" — students hypothesize what might be over- or under-represented - Discussion: Why does scale matter? What happened between GPT-2 and GPT-4? - Reflection journal: "Three things about LLMs that surprised me"

Day 2 — Prompt Engineering Mastery (~4 hrs)

Morning (2 hrs): - Lesson: The anatomy of a professional-grade prompt - Prompt patterns: zero-shot, one-shot, few-shot, chain-of-thought, tree-of-thought (conceptual) - Live demo: Building a prompt in real time — layering components to get dramatically better results - Activity: Deconstructing expert prompts — students reverse-engineer what makes them work

Afternoon (2 hrs): - Career field prompt design challenge (solo or pair): - Students select a career domain and build a 5-prompt "toolkit" as if they work in that field - Peer critique and refinement - Gallery walk: Best prompts posted, students leave sticky note feedback - Reflection journal: "What does being good at prompting tell us about being good at communication?"

Day 3 — Accuracy, Hallucinations, and Evaluation (~4 hrs)

Morning (2 hrs): - Lesson: Why LLMs hallucinate — the statistical root cause - Real-world examples of AI hallucinations with consequences (legal, medical, journalistic) - Activity: Structured hallucination audit - Students choose a topic, generate AI responses, verify against credible sources, document discrepancies

Afternoon (2 hrs): - Lesson: How to use AI responsibly for research - AI as a starting point, not a source - Triangulation strategies - Activity: Comparative analysis — AI-generated summary vs. original source - Discussion: "Should AI output be cited? How? When?" - Reflection journal: "My personal AI research protocol"

Day 4 — Bias, Ethics, and Power (~4 hrs)

Morning (2 hrs): - Lesson: Sources of bias in AI — data, design, deployment - Activity: Structured bias audit - Students test AI with demographically coded prompts and document the response patterns - Discussion: Who is harmed when AI is biased? Who has the power to change it?

Afternoon (2 hrs): - Lesson: The AI industry — who are the major players, what are their incentives? - Case study deep-dive (student groups each take one): - AI in hiring - Facial recognition and law enforcement - AI content moderation - Deepfakes and misinformation - Group presentations: What happened? Who was affected? What should change? - Reflection journal: "If I had the power to change one thing about how AI is built or governed, it would be..."

Day 5 — Project Showcase and Futures (~4 hrs)

Morning (2 hrs): - Mini-project work time and final prep - Facilitator coaching sessions

Afternoon (2 hrs): - Project presentations - Peer feedback (warm/cool/curious) - Panel discussion: "What does the AI landscape look like in 10 years?" - Reflection and close

Mini-Project Options: 1. Prompt Toolkit (professional domain, 5+ prompts with annotation and reflection) 2. Bias Audit Report (structured, with evidence and recommendations) 3. AI Policy Brief (1–2 pages, addressing a specific AI ethics issue) 4. AI Explainer for a Non-Technical Audience (video, podcast, or article) 5. Technical Deep-Dive (how a specific aspect of LLMs works — written or presented)

Assessment

- Daily reflection journals (formative)
- Mini-project + presentation (summative)
- Optional: Peer evaluation rubric

DELIVERY FORMAT

Format C — Weekly Series 6–12 weekly sessions

Duration: 6–12 weekly sessions, approximately 75–90 minutes each Ideal Setting: After-school program, elective, workforce development series Class Size: Up to 25 students

---

Learning Objectives

By the end of the series, students will be able to: - Explain LLM architecture and functioning with technical accuracy - Design and execute sophisticated prompt strategies across professional contexts - Conduct structured audits of AI for accuracy, bias, and ethical implications - Analyze systemic social impacts of AI with evidence-based reasoning - Produce and present a capstone project demonstrating applied and critical mastery

---

Session Map

Session 1: The Architecture of Intelligence

Focus: Technical foundations — how do LLMs actually work? Key Concepts: Neural networks, training, parameters, transformers, attention Activities: - Pre-survey: What students already know and believe about AI - Visual walkthrough of the transformer architecture (conceptual) - Live demo: Token probability visualization - Discussion: "If LLMs are pattern-matchers, what does that mean for how we use them?" Homework: Read a provided 2-page overview of a major moment in AI history. Annotate.

Session 2: Training Data and the World Inside the Model

Focus: What went into building these models? Key Concepts: Training corpora, web crawls, RLHF, fine-tuning, scale Activities: - Discussion: "What's in the training data? What's missing?" - Activity: Students hypothesize what biases might result from known data sources - Lesson: RLHF — how human feedback shapes model behavior - Discussion: Who are the humans providing feedback? What are their incentives? Homework: Find and bring in one news article about AI training data. Be ready to discuss.

Session 3: Prompt Engineering — Professional Level

Focus: Mastering the full anatomy of expert prompts Key Concepts: System prompts, few-shot, chain-of-thought, context windows Activities: - Homework article discussion - Prompt anatomy masterclass: deconstructing expert-level prompts - Practice: Students build a professional-grade prompt for a career of their choice - Peer critique using rubric Homework: Build a 3-prompt "AI workflow" for something you actually need to do (college essay, job application, research project, etc.). Screenshot and annotate.

Session 4: Advanced Prompting Strategies

Focus: Sophisticated prompt design patterns Key Concepts: Tree-of-thought, iterative refinement, meta-prompting, adversarial testing Activities: - Homework share-out: "What worked? What didn't?" - Lesson: Advanced prompting strategies with live demos - Challenge: Students redesign their homework prompts using at least two advanced strategies - Discussion: "What does skilled prompting have in common with good writing or teaching?" Homework: Use AI to help with one real academic or personal task. Reflect: was it helpful? How? Did it have limitations?

Session 5: Hallucinations, Reliability, and Research

Focus: Critical evaluation of AI output Key Concepts: Hallucinations, confabulation, statistical confidence, verification strategies Activities: - Homework reflection discussion - Lesson: Why hallucinations happen (mechanism, not just definition) - Structured hallucination audit activity - Activity: Building a personal "AI Reliability Checklist" - Discussion: "Is an AI that lies confidently better or worse than one that says 'I don't know'?" Homework: Fact-check an AI response on a topic you know well. Document errors and reasoning.

Session 6: Bias — Mechanisms, Manifestations, and Stakes

Focus: Understanding and auditing bias in AI Key Concepts: Training data bias, feedback bias, deployment bias, disparate impact Activities: - Homework share: bias documentation reports - Lesson: A taxonomy of AI bias — where it enters, how it compounds - Structured bias audit: demographic testing of AI responses - Case study: A real-world example of harmful AI bias with measurable consequences - Discussion: "Who has the responsibility to fix this? The builder? The deployer? The user? The government?"

For 6-session format: Students begin capstone planning and present at a modified showcase.

Session 7: AI Ethics Frameworks

Focus: Applying ethical reasoning to AI Key Concepts: Utilitarianism, rights-based ethics, fairness, care ethics Activities: - Introduction to ethical frameworks (brief, applied) - Apply each framework to one AI ethics scenario - Students write a structured ethical argument about a chosen AI issue - Peer review and discussion

---

Session 8: AI, Power, and the Political Economy of Tech

Focus: Who controls AI and why it matters Key Concepts: AI concentration, OpenAI/Anthropic/Google dynamics, global AI competition, labor Activities: - Reading and discussion: The AI industry landscape - Discussion: "What happens when a few companies control the most powerful AI systems?" - Activity: Map the AI power landscape — who funds it, who benefits, who's left out - Discussion: AI and national security, AI and democracy

Session 9: AI Policy and Governance

Focus: What rules should govern AI? Key Concepts: AI regulation globally (EU AI Act, US executive orders), responsible AI principles Activities: - Overview of major AI policy initiatives - Simulation: "You're an AI regulator" — students draft a policy for a specific AI use case - Debate: "AI should be regulated like a utility" — for and against

Session 10: AI, Labor, and the Future of Work

Focus: Economic and social impact Key Concepts: Automation, augmentation, job displacement, digital divide Activities: - Data analysis: Which jobs are most affected by AI automation? - Discussion: "How do we build an economy where AI benefits everyone?" - Reflection: Skills I want to develop; how AI changes my career planning

Session 11: Capstone Project Work Session

Activities: - Structured project work time - Facilitator one-on-one coaching - Peer draft review

Session 12: Capstone Showcase

Activities: - Presentations (5 min + Q&A) - Peer feedback - Certificates and celebration

Capstone Project Options (HS Level): 1. Prompt Toolkit (professional domain, annotated, with methodology reflection) 2. Bias Audit Report (structured research report with evidence, analysis, and recommendations) 3. AI Policy Brief (3–5 pages, research-backed, addressing a specific AI governance question) 4. AI Ethics Paper (analytical essay applying an ethics framework to a real AI controversy) 5. Technical Explainer (deep-dive on one technical concept for a non-expert audience) 6. AI Impact Case Study (researched analysis of AI's effect on a specific community or industry)

Assessment

- Session participation and discussion quality (formative)
- Homework submissions (formative)
- Capstone project + presentation (summative)
- Optional: Peer evaluation, self-assessment rubric

DELIVERY FORMAT

Format D — Semester Residency 16–18 weeks

Duration: 16–18 weeks; 2–3 class sessions per week Ideal Setting: Full elective course, embedded co-instructor model, standalone AI literacy class Class Size: Up to 25 students Materials: Devices, digital portfolios, full curriculum packet, access to research databases

---

Learning Objectives

By the end of the residency, students will be able to: - Explain LLM architecture and functioning at a technical and conceptual level - Design and critically evaluate sophisticated prompts for professional and academic use - Conduct independent research-grade audits of AI systems for accuracy, bias, and ethics - Analyze AI governance and policy with evidence-based, structured reasoning - Produce a portfolio-ready capstone project demonstrating mastery across all dimensions

---

Unit Overview

Unit	Title	Weeks	Focus
1	Technical Foundations	1–3	How LLMs work, training, architecture
2	Prompt Engineering Mastery	4–7	Expert prompting across professional domains
3	Critical Evaluation	8–10	Hallucinations, bias, reliability
4	Ethics, Policy, and Society	11–14	Systemic analysis, governance, advocacy
5	Capstone Portfolio	15–18	Independent research + presentation

---

Unit 1: Technical Foundations (Weeks 1–3)

Essential Questions: - What is actually happening when an LLM generates text? - How did we arrive at this moment in AI history? - What are the fundamental limitations of the statistical approach to language?

Week 1 — History and Context - History of AI: symbolic AI, expert systems, machine learning, deep learning - The transformer moment (2017): what changed and why - Timeline activity + discussion: "What surprises you about how recent this all is?"

Week 2 — How LLMs Work - Tokenization, embeddings, attention mechanisms, context windows - Scale: from GPT-2 to GPT-4 — what did more parameters actually change? - Technical reading: excerpt from a technical blog post (accessible level) - Lab: Token visualizer, context window experiments

Week 3 — Training, Fine-Tuning, and RLHF - Pre-training, fine-tuning, RLHF — how a model becomes "helpful and harmless" - Discussion: Who provides the human feedback? What are their incentives? - The alignment problem: introduction - Unit 1 assessment: Annotated technical diagram + short explanation essay

---

Unit 2: Prompt Engineering Mastery (Weeks 4–7)

Essential Questions: - What distinguishes an expert prompt from a novice one? - How does skilled prompting relate to broader communication and critical thinking skills? - What are the ethical boundaries of prompt engineering?

Week 4 — Foundations and Patterns - Anatomy of a professional prompt - Prompt patterns: zero-shot, few-shot, chain-of-thought, role prompting - Lab: Prompt engineering practice across 4 domains

Week 5 — Advanced Strategies - Meta-prompting, iterative refinement, tree-of-thought (conceptual) - Adversarial prompting: how to test AI weaknesses - Professional toolkit project begins

Week 6 — Applied Prompt Engineering - Real-world application: students build full AI workflows for academic tasks - Academic integrity discussion and personal AI use policy development - Peer review of workflows

Week 7 — Prompt Mastery Capstone - Professional domain prompt toolkit (final) - Presentation: Students present their toolkit and methodology - Reflection: "How has thinking about prompting changed how I think about communication?"

Unit 3: Critical Evaluation (Weeks 8–10)

Essential Questions: - When should I trust AI, and when shouldn't I? - What systematic methods can I use to evaluate AI reliability? - What are the consequences when AI gets it wrong?

Week 8 — Hallucinations and Reliability - Technical root causes of hallucination - Taxonomy of AI errors: factual, logical, contextual, temporal - Lab: Hallucination documentation project (ongoing through Week 10)

Week 9 — Bias: Mechanisms and Manifestations - Sources of bias in training data and RLHF - Types of bias: demographic, ideological, temporal, cultural - Structured bias audit lab

Week 10 — Building an Evaluation Framework - Students synthesize what they've learned into a personal AI Evaluation Framework - Peer review of frameworks - Unit 3 assessment: Reliability + Bias audit report on a specific AI tool

---

Unit 4: Ethics, Policy, and Society (Weeks 11–14)

Essential Questions: - Who is harmed by AI, and who has the power to change that? - What ethical frameworks are most useful for thinking about AI? - What does responsible AI governance look like, and who should design it?

Week 11 — AI Ethics Frameworks - Ethical frameworks applied to AI: utilitarian, deontological, care ethics, justice - Case study analysis: one major AI ethics controversy - Applied ethics writing: structured ethical argument

Week 12 — AI and Power - Industry concentration: who controls foundation models? - Labor: data labeling, content moderation, and the invisible workforce - Environmental costs of AI infrastructure - Research-based discussion

Week 13 — AI Governance and Policy - Global AI regulation: EU AI Act, US executive orders, China's approach - What responsible AI principles look like in practice - Simulation: students draft AI policy for a specific context

Week 14 — AI, Society, and Advocacy - AI and democracy, AI and inequality, AI and the future of education - Student-led research presentations: AI's impact on a community or sector of their choice - Discussion: "What's the most important AI issue that isn't getting enough attention?"

Unit 5: Capstone Portfolio (Weeks 15–18)

Essential Questions: - What do I know about AI that matters, and to whom? - How can I communicate complex ideas clearly and persuasively? - What is my position on AI, and how do I support it?

Week 15 — Project Selection and Planning - Portfolio options overview - Proposal writing and facilitator approval - Research design and source gathering

Week 16 — Research and Development - Intensive project work - Facilitator coaching sessions - Mid-point milestone: rough draft or prototype

Week 17 — Refinement and Review - Peer review sessions (structured feedback protocol) - Final revisions - Portfolio assembly: all artifacts + reflective introduction essay

Week 18 — Showcase and Portfolio Submission - Public or in-school showcase - Presentations (7–10 min + Q&A) - Portfolio submission - Celebration and close

Capstone Options (HS Residency Level): 1. Research Paper (8–10 pages, original analysis on an AI ethics or policy question) 2. Policy Brief + Testimony (written brief + mock public testimony) 3. Bias Audit Research Report (systematic study of AI bias with original data collection) 4. AI Technical Explainer Series (3-part educational content for a general audience) 5. AI Ethics Case Study (in-depth analysis of one controversy, journal-article style) 6. Original Curriculum Design (design and pilot a 2-session LLM 101 lesson for younger students)

Semester Residency Assessment Structure

Component	Weight	Description
Unit 1 Assessment	10%	Annotated diagram + explanation essay
Prompt Toolkit Presentation	15%	Professional domain prompt portfolio
Reliability + Bias Audit Report	15%	Unit 3 research deliverable
AI Ethics Paper	15%	Applied ethical analysis
Portfolio Reflective Essay	10%	Integrating reflection across all units
Capstone Project + Presentation	35%	Final major deliverable

APPENDICES

Appendix A: Core Vocabulary

Artificial Intelligence (AI): Computer systems designed to perform tasks that typically require human intelligence, such as understanding language, recognizing images, or making decisions.

Large Language Model (LLM): A type of AI trained on massive amounts of text data to understand and generate human language. Examples include GPT-4, Claude, and Gemini.

Token: The basic unit an LLM processes. A token is roughly a word or part of a word. "Unbelievable" might be split into ["un", "believ", "able"].

Training Data: The massive collection of text (books, websites, code, etc.) used to teach an LLM. The model learns patterns from this data.

Prompt: The input — the question, instruction, or context — that a user gives to an LLM. The quality of the prompt strongly influences the quality of the response.

Prompt Engineering: The skill of crafting prompts to get more useful, accurate, or creative outputs from an AI system.

Hallucination: When an LLM generates information that sounds confident and plausible but is factually wrong. This happens because LLMs predict likely next words, not verified facts.

Bias: Systematic patterns in AI output that reflect unequal treatment of different groups, often inherited from imbalanced training data or design choices.

Fine-tuning: Additional training of a pre-trained model on a specific dataset or with human feedback to improve performance for a particular task.

RLHF (Reinforcement Learning from Human Feedback): A training technique where human raters evaluate AI responses, and the model is updated to produce responses more like the highly-rated ones.

Context Window: The amount of text an LLM can "see" at once. Everything in the conversation, including instructions and history, must fit within this window.

Transformer: The neural network architecture behind modern LLMs, introduced in 2017. The "attention mechanism" allows transformers to understand relationships between words across long distances.

Attention Mechanism: The component of a transformer that allows the model to weigh the importance of different words when generating each new word.

Embedding: A numerical representation of a word or concept that captures its meaning and relationships to other words.

Parameters: The internal numerical settings of a model, adjusted during training. More parameters generally mean a more capable (but more expensive) model. GPT-4 is estimated to have hundreds of billions.

Generative AI: AI systems that create new content — text, images, audio, code — rather than just classifying or retrieving existing content.

Foundation Model: A large, general-purpose AI model trained on broad data that can be adapted for many different tasks.

AI Alignment: The challenge of ensuring AI systems behave in accordance with human values and intentions.

Deepfake: AI-generated media (video, audio, images) designed to realistically portray someone doing or saying something they never did.

Appendix B: Materials and Technology Checklist

For All Formats

- [ ] Student devices with internet access (1:1 preferred; 1:2 minimum)
- [ ] Projector or display screen for facilitation
- [ ] Whiteboard or chart paper + markers
- [ ] Printed activity packets (or digital equivalent)
- [ ] Student notebooks or journals

Recommended AI Tools for Classroom Use

Note: Schools should verify account and age requirements before use. Most platforms require accounts for full functionality.

- Claude (claude.ai): Strong for nuanced text tasks, excellent for education
- ChatGPT (chat.openai.com): Widely known, good general-purpose tool
- Google Gemini (gemini.google.com): Google account integration, good for research workflows
- Microsoft Copilot (copilot.microsoft.com): Good for Microsoft-integrated schools
- Khan Academy Khanmigo: Designed for education, recommended for Tier 1

Recommended Free Tools

- Tiktokenizer / OpenAI Tokenizer: Visualize how text is tokenized
- Anthropic Claude: Often accessible without account for limited use
- Google's Teachable Machine: Introduction to ML concepts without LLMs
- MIT's Moral Machine: Ethics decision-making (good for Unit 4)

Appendix C: Facilitator Notes

Setting Up for Success

- Run all tools yourself before the session and have backup plans for access issues
- Create example prompts and outputs in advance in case live AI access is unavailable
- Establish community norms around AI use and academic integrity in the first session

- Expect and welcome students who are already expert users — leverage their knowledge

Managing Student Access

- Use school-approved platforms whenever possible
- Review each platform's terms of service and age requirements
- For younger students (Grades 5–7), facilitate as a whole class on a projector when possible rather than individual unsupervised access
- Remind students that AI conversations may be logged by the platform

Addressing Common Misconceptions

- "AI is thinking" → AI is predicting. It has no understanding, intent, or experience
- "AI knows everything" → LLMs have a training cutoff date and confidently fabricate
- "AI is neutral/objective" → LLMs reflect the biases in their training data
- "AI will take all jobs" → More nuanced — some tasks will be automated, human roles will shift
- "AI understands me" → LLMs produce text that matches patterns; they do not "understand"

Handling Difficult Moments

- If students get inappropriate AI outputs: redirect, don't shame — use it as a teaching moment about AI limitations and guardrails
- If students are using AI to cheat: address directly with school policy; also explore why and what learning gaps exist
- If students challenge your knowledge of the technology: be honest about limitations and model good epistemic humility

---

Appendix D: Student Reflection Prompts

Use throughout any format to build metacognitive habits.

- What did I think about AI before today? How has that changed?
- What surprised me most about how LLMs actually work?
- Describe a time I got a bad AI response. What might have caused it?
- What's the difference between a mediocre prompt and a great one?
- If I could fix one thing about how AI is built, what would it be?
- What skills do I have that AI cannot replicate? How do I know?
- Who do I think benefits most from AI? Who do I think is most harmed?

- What question about AI do I most want answered that nobody seems to know yet?
- How has this changed how I plan to use AI in my own life?
- What's something I want to learn more about that we didn't cover?

Appendix E: Family and Community Engagement

What Families Should Know

- Their student is learning to think critically about AI — a skill that will matter for every future career
- AI tools are not banned in this curriculum; students learn to use them wisely and ethically
- Families are encouraged to have conversations about AI at home — what tools do you use? What concerns do you have?

Homework to Share With Family

- The LLM Analogy Project (Tier 1, Session 2): Explain how an LLM works to someone at home
- The AI Encounter Map (Tier 1, Day 1 Intensive): Where did your family encounter AI this week?
- The Personal AI Policy (any format): Share your rules for responsible AI use and get family feedback

---

LLM 101 Curriculum — Blacksky Media | April 2026 For educational use. May be adapted with attribution.

"I am because we are."

UBUNTU · THE PHILOSOPHY BEHIND EVERYTHING BLACKSKY BUILDS

BlackskyUp · Blacksky Labs · blackskymedia.org · mario@blackskymedia.org