Course Syllabus

Grade Level: Suitable for 6-12

Overview

This seven-session curriculum guides students through the full arc of working responsibly with AI. It opens by building foundational fluency – how large language models work and how to prompt them effectively using the BRIEF framework, then immediately instills the habit of critical verification through lateral reading and source-checking. From there, students apply those skills to real research, using NotebookLM to curate sources and develop a project goal connected to a career area of personal interest, then direct Claude as an AI agent to conduct structured comparisons and generate an interactive data visualization. The final three sessions shift students from research into communication and accountability: they build a professional data visualization in Datawrapper, embed it into an AI-generated app mockup designed with UI/UX principles in mind, and close by examining real-world consequences of AI bias and producing formal ethics documentation, a model card or responsible use card, for their own project. Taken together, the curriculum moves students from AI consumer to AI builder to AI evaluator, with student project work, ethical reasoning, and human judgment as consistent throughlines across every session.

Lesson Overview

Lesson 1 – Effective Prompting with BRIEF Students learn how large language models work and practice writing high-quality prompts using the BRIEF framework (Background, Role, Input, End Goal, Format). They compare outputs across multiple LLMs side-by-side and develop critical evaluation habits.

Lesson 2 – Lateral Reading with AI Students learn to verify AI-generated content using lateral reading — the same technique used by professional journalists and
fact-checkers. They identify real vs. fabricated sources and apply verification skills to their own project materials.

Lesson 3 – Tutorials Help You Learn Anything Students use NotebookLM to curate credible sources, generate AI-assisted study aids, and hold structured

conversations with their sources. The lesson culminates in a written project goal statement connecting AI to a career area of personal interest.

Lesson 4 – Using AI Agents for Research and Comparison Students explore how AI agents differ from simple chatbots, then direct Claude to conduct structured research, score a comparison across five criteria, and generate an interactive HTML visualization. Emphasis is placed on iteration, skepticism, and human judgment.

Lesson 5.1 – Data Visualization, Part 1 Students are introduced to data storytelling and explore real-world dashboards in PowerBI. They select data relevant to their own projects, apply ethical data practices, and build an original visualization in Datawrapper.

Lesson 5.2 – Data Visualization, Part 2 Students learn the distinction between UI and UX, then use an AI-assisted design tool (Uizard or UXPilot) to create a mockup of their app or website. They embed their Datawrapper visualization from Part 1 into the mockup and export it for use in their final presentation.

Lesson 6 – Ethics Check Students explore real-world consequences of AI bias, examine model cards from Google, and create either a Model Card or a Responsible Use Card for their own project. The lesson closes with a discussion on explainability and transparency in AI systems. Note: completion of this ethics documentation is required for Samsung Solve for Tomorrow submissions.

How to Familiarize Yourself with Each Lesson

Before teaching each lesson, follow these two steps:

Step 1 – Read the lesson PDF in full. Each lesson document contains the complete session flow with word-for-word facilitation prompts, activity instructions, debrief questions, differentiation suggestions, and "What If" scenarios for common student responses. Reading it fully, not just skimming, is essential, because the facilitation language is intentionally specific.

Step 2 – Review the linked materials for that lesson. Each lesson references external tools and student-facing resources (spreadsheets, templates, example model cards, article lists). Open every linked resource before class so you know what students will see. Where applicable, test the tool yourself by completing the activity as a student would.

A few additional notes:

- Lessons 5.1 and 5.2 are a two-part sequence and should be taught in consecutive sessions. Students carry a file (a Datawrapper PNG) from Part 1 into Part 2.
- Lesson 6 has a prerequisite: students should have an existing AI project (their Samsung Solve for Tomorrow technology) to document. This lesson does not work in isolation.
- Lesson 4 uses Claude; confirm student accounts are active before class. No account setup time is built into those lesson plans.

Materials & Resources Checklist

Student-Facing Documents to Distribute or Copy

- ☐ Prompt Engineering and Comparing LLMs Homework – Lesson 1*
- ☐ LLMs Comparison Spreadsheet – Lesson 1
- ☐ 5 Articles: Real vs. Fake (Claude artifact link) – Lesson 2
- ☐ Prompting AI Agents – Lesson 4
- ☐ Data Visualization In-Class – Lesson 5.1*
- ☐ Data Visualization Pre-Reading – Lesson 5.1*
- ☐ Data Visualization Pre-Reading – Lesson 5.2*
- ☐ Data Visualization Ulzard Student Instructions – Lesson 5.2
- ☐ Data Visualization UXpilot Student Instructions – Lesson 5.2
- ☐ Ethics Check Homework – Lesson 6*
- ☐ Ethics Check Pre-Reading – Lesson 6*
- ☐ Model Card Template – Lesson 6
- ☐ Responsible Use Card Template – Lesson 6
- ☐ Example Responsible Use Card – Lesson 6

*Links with * are optional student documents for shorter class periods.*

Hardware/Logistics

- ☐ Laptops preferred over phones or tablets (every lesson)
- ☐ Projector or smartboard for teacher modeling (especially Lessons 1, 3, 5.1)
- ☐ Whitelist (unblock) all websites on school network before Day 1

Whitelist/Unblock Websites in Advance

- ☐ LMArena (lmarena.ai) – Lesson 1
- ☐ Image for Activity #1 – Lesson 1
- ☐ 5 Claude Articles: Fake vs Real - Lesson 2
- ☐ NotebookLM (notebooklm.google.com) – Lesson 3
- ☐ Claude (claude.ai) – Lesson 4
- ☐ Datasets CDC, census.gov, Google Dataset Search, and Kaggle - Lesson 5.1
- ☐ PowerBI Examples - Lesson 5.1
  - ☐ Arts and Entertainment
  - ☐ Business and Entrepreneurship
  - ☐ Computer Science
  - ☐ Ecology and Food Security
  - ☐ Healthcare
  - ☐ Sports Science
- ☐ Datawrapper (datawrapper.de) – Lesson 5.1
- ☐ Uizard (uizard.io) and UXPilot (uxpilot.ai) – Lesson 5.2
- ☐ Example Model Cards - Lesson 6
  - ☐ Gemma Model Card
  - ☐ MediaPipe FaceMesh Model Card
  - ☐ Object Detection Model Card

Optional Extension Resources

- ☐ News Literacy Project (newslit.org) – Lesson 2
- ☐ Snopes, Google Fact Check Explorer, AllSides – Lesson 2
- ☐ Datawrapper River and Academy – Lesson 5.1
- ☐ Survival of the Best Fit (survivalofthebestfit.com) – Lesson 6
- ☐ The Moral Singularity (themoralsingularity.org) – Lesson 6

Suggested Schedule

This curriculum is designed to be taught as a coherent unit. The ideal schedule runs over seven consecutive class periods, but could be arranged in any format, such as seven Fridays. It is recommended, though, that lessons 5.1 and 5.2 are taught consecutively to form one two-part block.

Flexibility note: Lessons 1–4 can each stand alone if the full unit is not feasible. However, Lessons 5.1 → 5.2 → 6 are most effectively taught together and in order, as students carry artifacts and project work forward from one lesson to the next.

Curriculum Progression

This curriculum is structured as a deliberate learning arc, moving students from AI consumer to AI builder to AI evaluator.

Phase 1 – Understanding & Using AI (Lessons 1–2) Students develop foundational fluency: what LLMs are, how they work, and why they require human oversight. The BRIEF framework gives students a repeatable structure for getting useful outputs.
Lateral reading establishes the critical habit of verification before that habit is needed in later, higher-stakes lessons.

Phase 2 – Applying AI to Real Work (Lessons 3–4) Students move from understanding AI to directing it. In Lesson 3, they use a research tool grounded in their own curated sources to begin forming a project vision. In Lesson 4, they deploy Claude as an agent – writing structured prompts, evaluating outputs skeptically, and iterating toward a usable artifact. Both lessons reinforce that the human's judgment is what gives AI output its value.

Phase 3 – Communicating with Data (Lessons 5.1–5.2) Students shift from research to communication. They learn to turn data into a visual argument (Part 1) and then embed that argument into a designed product mockup (Part 2). These lessons build presentation-ready assets while reinforcing ethical data practices introduced earlier.

Phase 4 – Taking Responsibility (Lesson 6) The final lesson asks students to account for the AI choices they've made throughout the unit. By examining real-world consequences of AI bias and producing formal ethics documentation for their own projects, students internalize accountability as a non-optional part of building with AI – not an afterthought.

Throughlines across all lessons:

- The BRIEF framework is introduced in Lesson 1 and actively referenced in Lessons 3, 4, and implicitly throughout

- Hallucination and verification are introduced in Lesson 1, deepened in Lesson 2, and revisited in Lessons 3 and 4
- Ethical AI considerations appear in every lesson and culminate formally in Lesson 6
- Student project work (Samsung Solve for Tomorrow) threads through Lessons 3, 5.1, 5.2, and 6

Key Vocabulary

Accountability Taking responsibility for the outcomes an AI system produces, including unintended consequences.

AI Agent Software that plans steps, uses tools, and adjusts its approach to complete a multi-step task – as opposed to a simple chatbot that responds to one message at a time.

AI Hallucination When a large language model generates citations, statistics, or facts that appear real but are completely fabricated. The model does not know it is doing this – it is predicting what a convincing answer looks like.

Artifact In the context of this curriculum, the interactive HTML visualization generated by Claude in Lesson 4. More broadly, any output produced by an AI system.

Attribution and Usage Rights The ethical and legal responsibility to credit original creators of content – especially images, music, and writing – and to confirm you are permitted to use their work.

Bias An unfair skew in an AI system's outputs, often caused by unrepresentative or historically skewed training data. Bias can affect who a model works well for and who it underserves.

BRIEF Framework A structured approach to writing effective AI prompts. Stands for Background, Role, Input, End Goal, and Format. Introduced in Lesson 1 and applied throughout the curriculum.

Criteria The specific dimensions used to evaluate and compare items fairly. In Lesson 4, students define five criteria before directing an AI agent to conduct a structured comparison.

Data Storytelling Communicating complex information in a way that is easy to understand and memorable, with the goal of helping audiences make better decisions or take action. Requires three elements: data, story, and a call to action.

Data Visualization The graphical organization of information – using charts, graphs, and maps – to make patterns, trends, and outliers easier to see and interpret.

Ethical Data Use Responsible practices when collecting and presenting data, including crediting original sources, anonymizing survey responses, obtaining informed consent, using a variety of sources, and labeling all visualizations clearly.

Explainability The ability to clearly communicate how an AI system arrived at a result – even when that result is unexpected or unfavorable. An explainable system allows users to understand the logic behind an outcome.

Iteration The process of refining AI output through follow-up prompts and targeted feedback, rather than accepting the first response as final.

Large Language Model (LLM) AI software trained on billions of pages of text that predicts and generates language based on learned patterns. Examples include Claude, ChatGPT, Gemini, and Llama. LLMs do not look up answers – they predict the most likely next word based on what they have learned.

Lateral Reading Evaluating a source by opening new browser tabs and looking outside the source – at what other credible sources say about it – rather than reading deeply within it alone. Used by professional journalists and fact-checkers.

Market Research Researching background information to understand the context around a product, service, or target audience, including the size and characteristics of the population being served.

Model Card A document describing an AI model's intended use, performance metrics, limitations, training data, and potential biases. Required when a project involves training a model on a dataset. Analogous to a nutrition label for an AI system.

Model Comparison The practice of evaluating different AI tools against the same task to identify their relative strengths and weaknesses, rather than assuming

one model is universally superior.

Mockup A static visual representation of what an app or website screen will look like, used to plan and communicate design ideas before a product is built.

NotebookLM An AI-powered research tool developed by Google that only answers questions based on sources the user has uploaded. Reduces hallucination by grounding responses in user-provided material rather than general training data.

Prompt The text a human provides to an AI system in order to direct its behavior and shape its output. Prompt quality directly affects output quality.

Prompt Engineering The practice of crafting clear, specific, and well-structured prompts to improve the quality and usefulness of AI outputs.

Prototype An interactive or semi-functional model of a product, used to test and refine the user experience before full development begins.

Responsible Use Card An ethics documentation tool used when a project uses AI tools but does not involve training a model on a dataset. Based on IBM's 5 Pillars of Responsible AI, it documents how AI was used fairly, transparently, and responsibly.

Source Credibility The trustworthiness of a source, determined by factors including the author's reputation, publication date, editorial standards of the publisher, and corroboration by independent sources.

Source Curation The process of finding, evaluating, and selecting sources that are relevant, credible, and useful for a specific research purpose – as opposed to accepting whatever an AI tool returns by default.

Transparency Openness about how an AI system works – including what data it was trained on, who built it, and whether its process is consistent and fair across different users and contexts.

UI (User Interface) How a digital product looks – the visual design, layout, buttons, colors, and typography that users see and interact with.

UX (User Experience) How a digital product feels to use – the overall journey, ease of navigation, and whether users can accomplish their goals smoothly and intuitively.

Verification The process of independently confirming that AI-generated claims, citations, or data are accurate before using them as evidence. A non-negotiable habit when working with any AI output.

Teacher Background Reading

These resources are selected for educators, not students. The goal is not comprehensive expertise but enough working knowledge to answer unexpected student questions with confidence, facilitate nuanced discussions, and recognize when a student has raised something genuinely important. Estimated reading time is noted for each.

How Claude Works – Anthropic's Research on Tracing Thoughts in a Language Model
anthropic.com/research/tracing-thoughts-language-model ~20 minutes

The most directly relevant technical resource for this curriculum. Anthropic's own explanation of how large language models process and generate language – written accessibly enough for non-engineers. Particularly useful before teaching Lessons 1 and 2, where students ask why AI hallucinates and whether AI "understands" anything.
Having a clear mental model of what is actually happening inside the model will help you field those questions without oversimplifying.

Machine Bias – ProPublica's Investigation of the COMPAS Algorithm
propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing ~25 minutes

The original investigative journalism piece behind the criminal justice example in Lesson 6. Read this before teaching the ethics session so you can speak to the specifics with confidence if students push back or ask follow-up questions. Also a good model of what rigorous, evidence-based AI criticism looks like – which is the disposition you are trying to cultivate in students throughout the unit.

The Beginner's Guide to Model Cards – Google's Model Card Documentation
modelcards.withgoogle.com ~15 minutes browsing across 2–3 examples

Students examine these in Lesson 6, but educators should review them first. Spend time with the Object Detection and Gemma model cards in particular.

Understanding what professional documentation actually looks like will help you guide students toward something meaningful rather than a superficial fill-in-the-blank exercise. Pay attention to the limitations sections – those are the most instructive.

A People's Guide to AI – AlgorithmWatch / MIT Media Lab
https://alliedmedia.org/wp-content/uploads/2020/09/peoples-guide-ai.pdf
~30 minutes for the core sections

A plain-language primer on how AI systems are built, how they fail, and who bears the consequences. Written specifically for non-technical readers. Covers bias, data collection, automation, and accountability in accessible terms. Particularly useful background before Lessons 6 and 4, and valuable for building your own vocabulary around AI ethics before facilitating student discussions.

The SIFT Method – University of Chicago Library Guide
guides.lib.uchicago.edu/c.php?g=1241077&p=9082322 ~10 minutes

A companion framework to lateral reading, used by librarians and media literacy educators. SIFT stands for Stop, Investigate the source, Find better coverage, Trace claims. Reading this before Lesson 2 gives you a second vocabulary for discussing source verification – useful when students need a different framing than what the lesson provides, or when the lateral reading explanation isn't clicking for a particular student.

If Time Allows...

The Gender Shades Project from MIT Media Lab (gendershades.org) is about visually compelling research summary showing that facial recognition systems performed significantly worse on darker-skinned women. It takes under ten minutes to review and provides a concrete, student-accessible example of AI bias that complements the Lesson 6 discussion without requiring technical background to understand.

Cumulative Evidence of Understanding

This section consolidates the observable evidence of learning across all seven sessions. Because this curriculum is ungraded, these indicators are intended to help educators recognize whether students are developing genuine fluency – not to evaluate performance. Look for these behaviors cumulatively as the unit progresses.

Prompting & AI Fluency

- Student can explain in their own words what a large language model does – and what it cannot do
- Student writes prompts that include all five BRIEF elements without being prompted to do so
- Student can articulate why two different LLMs produced different outputs for the same prompt
- Student uses follow-up prompts to iterate on AI output rather than accepting the first response
- Student can identify at least one instance where an AI response was confidently wrong and explain why that happens

Critical Evaluation & Verification

- Student independently applies lateral reading strategies to check a source without being told to
- Student can correctly distinguish a real source from a fabricated one and explain the specific evidence that led to that conclusion
- Student can identify what perspective or information is missing from a source
- Student can name an alternative source they would use if their primary source could not be verified
- Student treats AI output as a starting point for investigation rather than a final answer

Research & Knowledge Construction

- Student selects and imports sources into NotebookLM that are credibly relevant to their chosen topic – not just the first results returned
- Student generates and uses at least one NotebookLM study aid to deepen their understanding rather than just skim the summary.

- Student produces a project goal statement that includes all three required components: an AI opportunity, a specific challenge, and a personal connection
- Student can explain how AI is being used or could be used in a real career field of their choosing

Data Communication

- Student can state the story their data visualization is telling in one clear sentence
- Student selects a chart type that is appropriate for their data – and can explain why
- Student's Datawrapper visualization includes a clear insight-driven title, labeled axes, and a cited source
- Student can distinguish between a UI problem and a UX problem when evaluating a design
- Student's app mockup reflects intentional design decisions made for a specific target audience, not just default AI-generated output

Ethics & Accountability

- Student can describe at least one real-world example of AI bias and explain how it happened
- Student correctly identifies whether their Samsung Solve for Tomorrow project requires a Model Card or a Responsible Use Card – and can explain the difference
- Student's completed card accurately reflects their project's data sources, limitations, and intended use
- Student can explain the difference between explainability and transparency in their own words
- Student contributes a substantive perspective to the group discussion on AI accountability – not just agreement with what others say

Overall Disposition

- Student approaches AI tools with calibrated skepticism – neither uncritical acceptance nor reflexive dismissal
- Student can articulate at least one decision they made that improved an AI output
- Student can name something they would do differently the next time they use an AI tool for real work

Created in partnership with the Mark Cuban Foundation, markcubanai.org

MARK CUBAN FOUNDATION