

Exploring AI

A Faculty Seminar on Practical Integration

Dr. Frazier A. Smith | Department Chair, Central Piedmont Community College

Join at menti.com

Code: [YOUR CODE]

THE SCENARIO

Your best student submits a final paper. It's polished, well-argued, and uses sources you've never seen her cite before. You're 80% sure she used AI to write at least part of it.

She has a 4.0 GPA and has never given you a reason to doubt her work.

What do you do?

What do you do?

- A. Grade the paper as submitted
- B. Ask the student to explain her process
- C. Run it through a detection tool
- D. Fail the assignment
- E. Something else

What the Next 80 Minutes Are About

"We don't yet have shared language for when AI use is acceptable and when it isn't."

0-8 min	The Opening Dilemma
8-12	Where We Are Right Now
12-28	How AI Actually Works
28-45	The Stoplight Framework
45-60	Prompt Design
60-75	Ethics and Edge Cases
75-85	Practical Applications
85-90	Close and Commitment

How Large Language Models Work

What Is a Large Language Model?

1

A prediction engine

Software trained on massive text to predict what words come next

2

"Large" = scale

Billions of adjustable parameters learned from trillions of words of text



3

Pattern recognition, not knowledge


It learned how language works, not facts about the world

Analogy: An LLM is like an extraordinarily well-read autocomplete — not a thinking mind, but a pattern-completion engine.

What Is a Large Language Model?



Where the "Knowledge" Comes From



**A large language model (by itself) doesn't browse the internet.
It produces text from patterns learned during training.**

AI, without tools, is the most confident liar you'll ever meet.

Tokens, Not Words

LLMs don't read words: they process tokens (subword pieces)

under

stand

ing

"understanding" = 3 tokens

Why this matters

- Models handle misspellings and new words because they work at the sub-word level
- Input limits are measured in tokens (~750 words ≈ 1,000 tokens)
- The model never "reads" your text the way you do — it processes numerical token sequences

Next-Token Prediction

Everything an LLM does is this one operation, repeated.



By default, there is no reasoning engine

No comprehension module. No understanding. Just extraordinarily sophisticated pattern completion.

The output looks intelligent

Because the patterns were learned from trillions of words of human-written text. The sophistication is inherited, not generated.

The Transformer: Why These Models Work

The breakthrough: the Attention Mechanism

When predicting the next word, the model weighs which previous words matter most

"The professor who taught the advanced seminar last fall is now writing a book about ____"

To predict the blank, the model attends most strongly to "professor," "seminar," and "book" — not just the word immediately before it.

This is why AI output can feel surprisingly coherent!
The model tracks relationships across the entire passage.

Three Concepts You'll Encounter

Temperature

Controls randomness in output

Low temperature = predictable,
conservative responses

High temperature = creative, varied,
riskier responses

Context Window

**How much text the model
can "see" at once**

Measured in tokens
(thousands to millions)

Larger window = handles longer
documents, but costs more

Fine-Tuning

Specializing a general model

Base model: general text predictor

Fine-tuned: specialized for
conversation, coding, etc.

ChatGPT = fine-tuned GPT

Where We Are Right Now

How often do you believe your students are currently using AI for coursework?

A. Rarely or never

B. A few students, occasionally

C. Many students, regularly

D. Most students, for most assignments

The Perception Gap

60%+

of undergraduates have used generative AI for coursework

Stanford University Student Survey, 2024

Faculty beliefs about AI usage vary widely,
and those beliefs shape policy decisions.

How AI Actually Works

Statistical Prediction, Not Comprehension

AI models predict the next most probable word based on patterns in training data.

It doesn't "know" anything. It doesn't "understand" content.

It assembles statistically likely sequences of text.

"Hallucination"

Assembling plausible components, not retrieving real sources

The output reads like a real citation
because the model has seen
thousands of them

AI detection tools work the same way
— statistical patterns, not reading for
meaning

Same Model. Same Knowledge. Different Instruction.

Bare Prompt

"Explain photosynthesis."

Generic, textbook-level output

CORE Prompt

"Explain photosynthesis to a community college biology student who struggles with abstract processes. Use a kitchen analogy. Under 200 words."

Targeted, tailored, useful output

What We Just Saw

1

AI predicts patterns — it doesn't understand content

The mechanism is statistical prediction, not comprehension

2

AI fabricates confidently — it doesn't retrieve from a database

"Hallucination" produces plausible but invented content

3

AI responds to instruction quality — the prompt is the lever

Specific, structured prompts produce dramatically better results

The Stoplight Framework

The Stoplight Model



GREEN — AI use is encouraged or expected

Brainstorming topics | Generating practice quiz questions | Self-study aids

YELLOW — AI use is permitted with conditions

Attribution required | Reflection on changes | Revision log | Instructor permission

RED — AI use is not permitted

Timed in-class essays | Unassisted writing assessments | Proctored exams

Green Zone: AI Encouraged or Expected

Brainstorming Essay Topics

Students use AI to generate potential angles on an assigned topic, then select and develop their own thesis

Practice Quiz Generation

Students prompt AI to create practice questions from their notes, then self-assess before the real exam

Research Exploration

Students use AI to survey a topic landscape and identify areas for deeper investigation

Yellow Zone: AI Permitted with Conditions

AI-Assisted Outline

Student uses AI to help structure a paper, then writes all content independently

Code Debugging Help

Student uses AI to identify bugs, but must explain the fix in their own words

Conditions that enable yellow assignments:

- ✓ Attribution of AI use
- ✓ Written reflection on what was changed and why
- ✓ Revision log showing the process
- ✓ Prior instructor permission

Red Zone: AI Not Permitted

Timed In-Class Essays

Assesses the student's unassisted writing ability under controlled conditions

Oral Examination Responses

Verbal demonstration of knowledge that can't be delegated to AI

Portfolio Reflection Pieces

Personal narrative and self-assessment that must represent authentic voice

"Most academic integrity conflicts come from ambiguity, not dishonesty."

A student uses AI to generate a first draft of an essay, then substantially revises it. What color is this?

A. Red — not permitted

B. Yellow — permitted with conditions

C. Green — encouraged

Prompt Design

The Quality of the Output Depends on the Quality of the Instruction

The CORE Prompting Framework

C

Context

Who is the audience? What do they already know?

O

Objective

What should the output accomplish?

R

Requirements

Format, length, constraints, tone

E

Expectations

What does a good result look like? What should it avoid?

CORE Elements

- ✓ **C**ontext: [audience + knowledge]
- ✓ **O**bjective: [goal of output]
- ✓ **R**equirements: [format/length/tone]
- ✓ **E**xpectations: [quality criteria]

Building the Prompt Live

[Prompt constructed live during session]

Example Faculty Workflow Prompt:

"Write three multiple-choice questions assessing whether a student understands the difference between correlation and causation. Target introductory statistics students. Each question should include one distractor that sounds correct but confuses the two concepts. Include answer explanations."

C: introductory statistics students | O: assess understanding | R: MC format, 3 questions, explanations | E: distractors that confuse correlation/causation

Iteration is the process. Revise the prompt, not just the output.

Prompt Design: Key Takeaways

1

Specific prompts produce dramatically better output

Vague in, vague out — every CORE element earns its place

2

Removing any element degrades results

Context, Objective, Requirements, and Expectations each do measurable work

3

Iteration is normal and expected

Refine the prompt based on what you see — this is skill, not magic

Ethics and Edge Cases

Modeling How to Think Through Hard Situations

Scenario A: The Verbal Mismatch

"A student submits an excellent paper but can't explain any of its arguments in a follow-up conversation. They become defensive when asked."

It could be...

- Legitimate work by a student who struggles verbally
- Anxiety about being accused, not guilt
- Different skills in writing vs. speaking

Or it could be...

- AI-generated work the student can't defend
- Defensiveness masking dishonesty
- A pattern worth investigating further

You don't know. What's your process?

Scenario B: The Equity Gap

"You teach two sections of the same course. One has reliable laptop and internet access. The other does not. How do you set AI policy across both?"

Equal Treatment

Same policy for both sections. Consistent expectations. But a blanket ban penalizes students who could benefit.

Equitable Treatment

Tailored policies for each section. Meets students where they are. But blanket permission advantages those with access.

Equal and equitable point in different directions.

Scenario C: The Threshold Problem

"You allow AI for brainstorming. A student brainstorms with AI, then organizes with AI, then drafts a paragraph with AI. At what point did brainstorming become writing?"



There is no clean line. The student followed the logic of each step.

**A student submits an excellent paper but cannot explain it verbally.
Your first response:**

- A. Ask them to walk through their writing process
- B. Grade the paper on its merits regardless
- C. Require a revised submission with process documentation
- D. Report a suspected academic integrity violation

**Policies need to be specific about process,
not just classification.**

Clear expectations before the assignment
prevent detective work after it.

Practical Applications

Rubric Generation

Rubric Descriptors in Minutes, Not Hours

Sample Prompt:

"Create a 4-level rubric for a persuasive essay assignment in a first-year composition course. Levels: Exemplary, Proficient, Developing, Beginning. Criteria: thesis clarity, evidence use, organization, and mechanics. Use student-friendly language."

**AI produces serviceable first drafts of rubric language.
Edit for your standards — don't start from scratch.**

Discussion Board Seeding

Thought-Provoking Discussion Prompts on Demand

Sample Prompt:

"Generate three discussion prompts for an undergraduate sociology class reading about social media and identity. Target Bloom's 'Analyze' level. Each prompt should invite productive disagreement without a single correct answer. Audience: sophomores with mixed engagement levels."

Specifying Bloom's level and desired controversy changes the output dramatically.

Study Guide Creation

From Learning Objectives to Student Study Guide

Sample Prompt:

"Create a student study guide for a midterm covering [3 learning objectives]. Include key terms with definitions, practice questions (mix of recall and application), and a self-assessment checklist. Format for easy printing."

Immediately recognizable as a time sink. AI handles the structure; you refine the content.

Email Drafting

Routine Communications

Sample Prompt:

"Draft a professional but warm email to a student who has missed three consecutive classes. Express concern, reference institutional support resources, and ask them to meet during office hours. Tone: supportive, not punitive."

Tasks you do repeatedly with minor variations. AI handles them competently.

Assessment Item Refinement

Sharpen Your Questions

Sample Prompt:

"Take this recall-level question: 'What is the definition of photosynthesis?' Convert it to an application-level question suitable for an introductory biology exam. Include four answer options with one correct answer and three plausible distractors. Explain why each distractor is wrong."

Connects back to the CORE framework. AI improves distractors and elevates question rigor.

One Change

**What is one specific change you will make
in your course this semester?**

Open-ended response

THE SCENARIO — REVISITED

Your best student submits a final paper. It's polished, well-argued, and uses sources you've never seen her cite before. You're 80% sure she used AI to write at least part of it.

She has a 4.0 GPA and has never given you a reason to doubt her work.

Has your answer changed?

The answer depends on decisions you haven't made yet

what your AI policy says,
how you communicated it,
and what process you have for
conversations like this one.

Those decisions are yours.
Today was about giving you the framework
and the language to make them well.

Thank You!

Slides and resources will be emailed to all participants.

Email: f.smith@wingate.edu