CAMBRIDGE
UNIVERSITY PRESS & ASSESSMENT

Record this session

AI Lunch & Learn

Your Journey from AI Curious to AI Confident

Welcome to Lesson 1.3: AI Capabilities and Limitations

Today we'll cover:

- Identify tasks where AI excels vs. struggles
- Recognise when human oversight is critical
- Position AI as an augmentation tool
- Build effective human-AI collaboration

30th October, 2025

© Cambridge University Press & Assessment 2022

CAMBRIDGE
UNIVERSITY PRESS & ASSESSMENT

How We Learn: The 30-Minute Formula

Every session follows the same structure

Instruction

Core concepts explained clearly

10 minutes

Demonstration

Watch it work in real-time

10 minutes

Practice

Try it yourself with guidance

10 minutes

CAMBRIDGE
UNIVERSITY PRESS & ASSESSMENT

What We've Learned So Far

Lesson 1.1:

- AI recognises patterns, not rules
- It's a tool for augmentation, not replacement
- AI makes mistakes - always verify outputs

Lesson 1.2:

- LLMs predict patterns, not facts
- Tokens and context windows limit what AI "remembers"
- Hallucinations happen - verify specific claims

CAMBRIDGE
UNIVERSITY PRESS & ASSESSMENT

Where AI Shines vs. Where AI Struggles

AI EXCELS at:

- ✓ Pattern recognition
- ✓ Repetitive tasks
- ✓ Data processing at scale
- ✓ Generating multiple options
- ✓ Consistency across large volumes

AI STRUGGLES with:

- ✗ True creativity and originality
- ✗ Understanding context and nuance
- ✗ Common sense reasoning
- ✗ Emotional intelligence
- ✗ Tasks requiring real-world experience

What AI Does Well

1. Pattern-Based Work

Summarising documents, Drafting emails from templates, Categorising and tagging content, Translating text between languages

2. Data-Heavy Tasks

Analysing spreadsheets, Finding information across documents, Generating reports from data, Spotting trends in large datasets

3. Repetitive Processes

Formatting content consistently, Generating variations of text, Answering common questions, Creating first drafts quickly

AI is excellent at tasks with clear patterns and defined inputs/outputs

What AI Struggles With

1. Judgment Calls:

Strategic decisions, ethical considerations, understanding politics, reading between lines

2. Specialised Expertise:

Company processes, internal info, current events, real-time awareness

3. Human Touch:

Building relationships, sensitive conversations, unspoken context, creative problem-solving

4. Accuracy on Specifics:

Precise dates/numbers/citations, calculations, complex instructions, real-time data

The Augmentation Mindset

Think of AI like:

- A research assistant who does the initial legwork
- A drafting tool that gets you to 70%
- A brainstorming partner with infinite patience
- A tireless helper for tedious tasks

YOU bring: Domain expertise, critical thinking, creativity, relationship building, final decisions

AI brings: Speed, scale, pattern recognition, consistent execution, 24/7 availability

When Human Oversight is ESSENTIAL

High-Stakes Decisions

- Financial commitments
- Legal documents
- Medical advice
- HR decisions

Always have a human in the loop

Accuracy-Critical Content

- Published reports
- Customer communications
- Academic work
- Compliance documents

Verify facts, numbers, citations

Sensitive Situations

- Customer complaints
- Performance issues
- Confidential matters
- Emotional contexts

Human judgment required

Creative Strategy

- Brand voice and tone
- Strategic direction
- Original concepts
- Stakeholder alignment

AI can assist, not decide

The Human-AI Partnership Framework

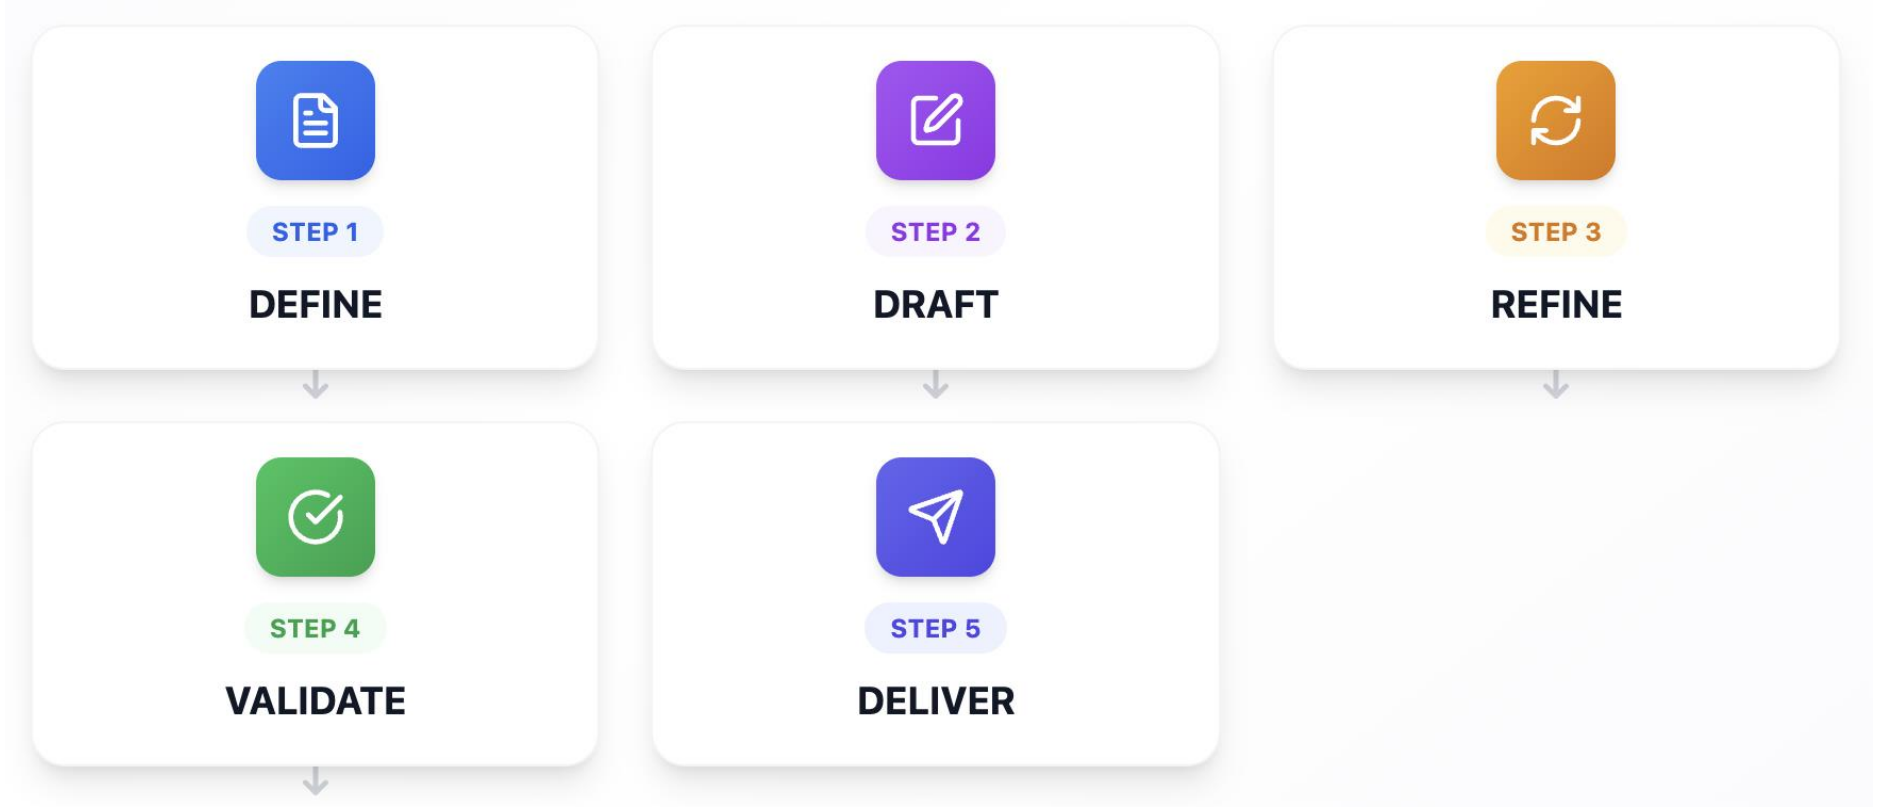

1. Set objectives, criteria, constraints (human leads)
2. Generate options, process data, create first drafts (AI does heavy lifting)
3. Apply domain knowledge, adjust for context, ensure accuracy (human leads)
4. Verify facts, check for hallucinations, ensure requirements met (critical human check)
5. Take ownership, make final adjustments, deliver with confidence (you're accountable)

Good Uses vs. Risky Uses

GOOD: AI for First Draft *Scenario:* Need to write a project update email

- **AI does:** Generate draft based on your bullet points
- **You do:** Adjust tone, add specific context, verify names/dates
- **Result:** Email done in 5 minutes instead of 20

RISKY: AI for Final Decision *Scenario:* Deciding whether to approve a budget request

- **AI does:** Analyses data, suggests recommendation
- **You do:** Consider AI input + organisational context + relationships + strategic priorities
- **Result:** Informed decision with AI assistance, not AI decision

GOOD: AI for Research *Scenario:* Understanding a new topic quickly

- **AI does:** Summarises key concepts, explains in plain language
- **You do:** Cross-check important details, dive deeper on critical points
- **Result:** Faster learning curve, better foundation

RISKY: AI for Sensitive Communication *Scenario:* Responding to an upset customer

- **AI does:** Generates template response
- **You do:** Heavily customise with empathy, specific context, genuine care
- **Result:** Better to write yourself or use AI only for structure

Remember These Core Principles

1. **AI Amplifies, Not Replaces** Your expertise + AI capability = Enhanced output
2. **You're Always Accountable** Never publish AI output without your review and approval
3. **Garbage In, Garbage Out** The quality of AI output depends on the quality of your input
4. **Verify Don't Trust** Always fact-check important details, especially dates, numbers, citations
5. **Human for High-Stakes** Critical decisions, sensitive situations, accuracy-critical content need human judgement

Demo Time

Now let's see this in action...

Demo 1: Where AI Excels

Summarisation

Format Conversion

Brainstorming Variations

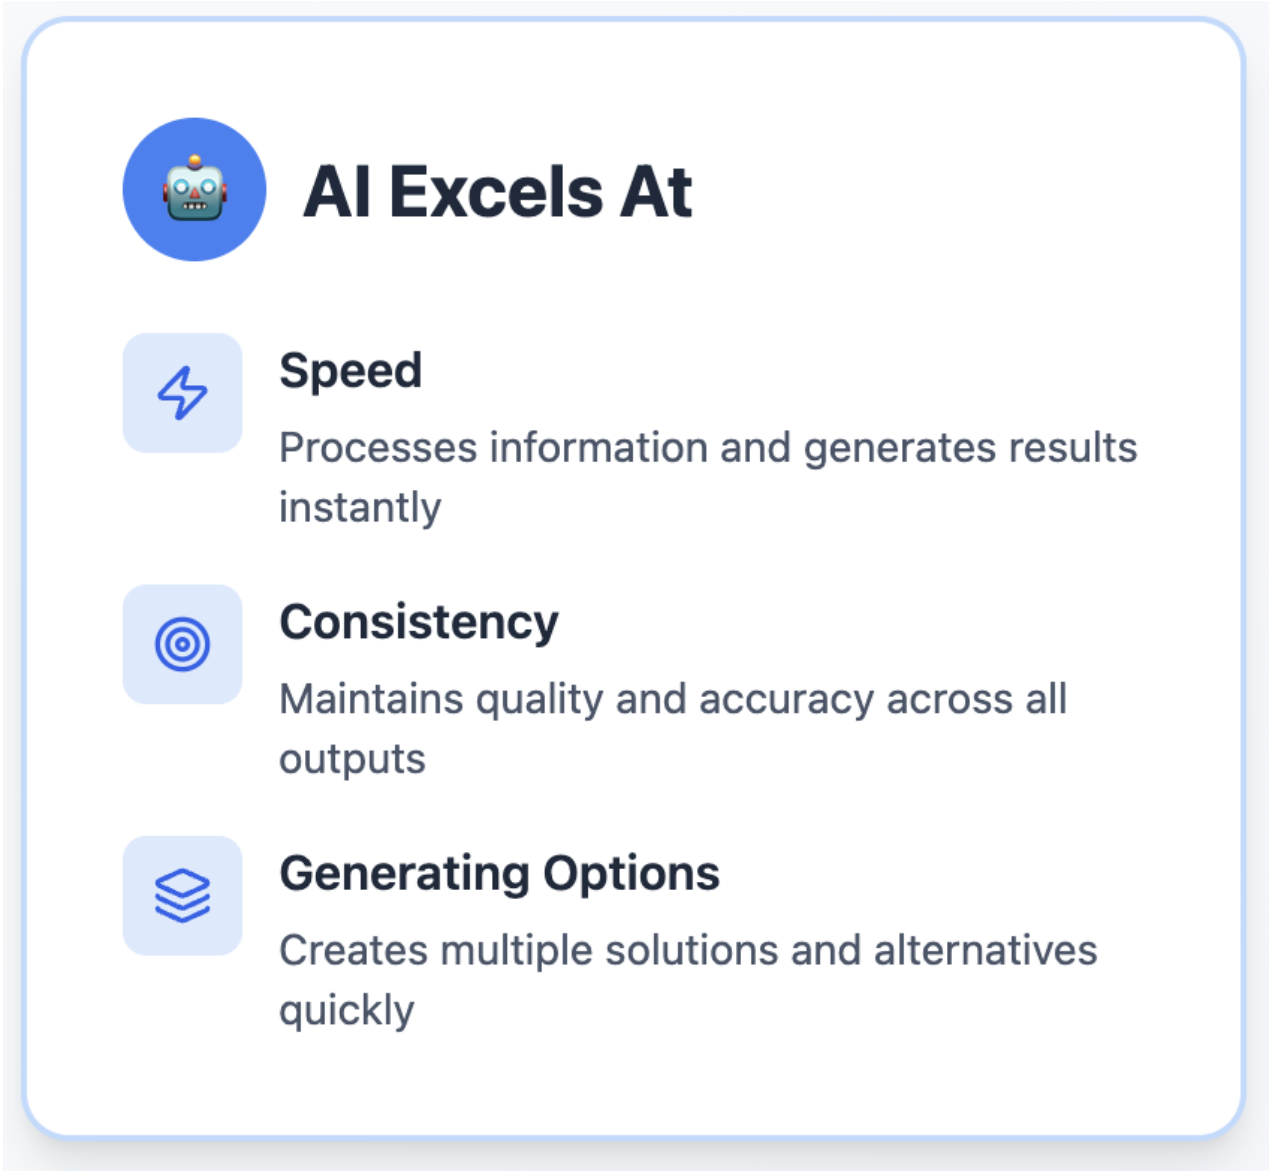

Demo 2: Where AI Needs Help

Common Sense Test

Judgement Call Test

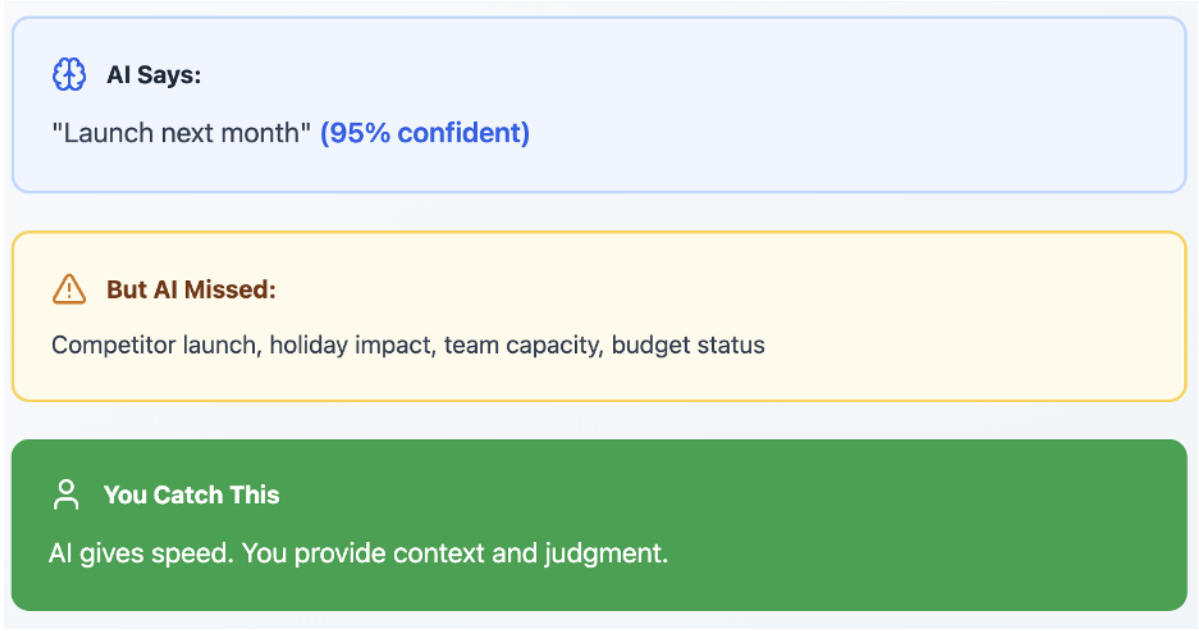

Exercise: AI Capability Assessment

Identify and test an AI-friendly task

Part 1 Identify 3 tasks

Part : Categorise them

Part 3: Test one with AI

Part 4: Quick sharing in chat

Common tasks could be:

Writing emails

Analysing data

Scheduling meetings

Creating presentations

Responding to routine questions

Categories:

AI-Friendly: AI could do most of this (with human review)

AI-Assisted: AI could help but I lead significantly

Human-Essential: AI would struggle, needs my expertise

Exercise: AI Capability Assessment

Part 1: Identify Your Use Cases (2 minutes)

Think of 3 tasks you do regularly at work. Write them down.

Exercise: AI Capability Assessment

Part 2: Assess Each Task (3 minutes)

For each task, decide:

AI-Friendly: AI could do most of this (with human review)

AI-Assisted: AI could help but I lead significantly

Human-Essential: AI would struggle, needs my expertise

Exercise: AI Capability Assessment

Part 3: Test One (4 minutes)

Pick your most AI-friendly task and:

Ask AI to help with it right now

Evaluate the output

Note what you needed to add/change/verify

Exercise: AI Capability Assessment

Part 4: Share & Discuss (1 minute)

In the chat, share:

What task did you test?

What did AI do well?

What did you need to change?

What We've Learned

AI Excels at:

- Pattern-based tasks
- Data processing
- Repetitive work
- Generating first drafts

AI Struggles with:

- Judgement calls requiring context
- Specialised/recent/internal knowledge
- Emotional intelligence
- Accuracy on specifics

Effective Collaboration:

- Define → Draft → Refine → Validate → Deliver
- You lead and verify, AI assists and accelerates
- Human oversight essential for high-stakes work

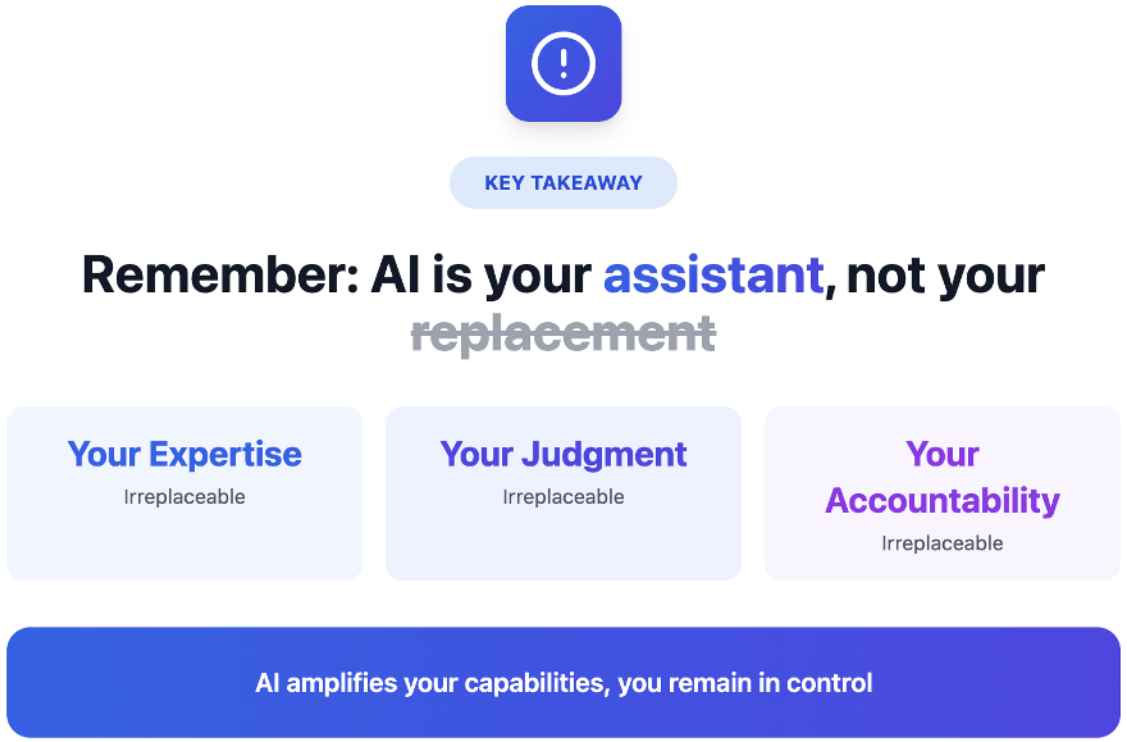

Homework

The AI Collaboration Experiment

Choose ONE work task to collaborate with AI on this week

Document:

What was the task?

How did you use AI?

What did AI do well?

What needed changing?

Time comparison?

Use again?

Share your experience in Teams Chat!

Next Week

Lesson 1.4: AI Ethics

We'll explore: AI bias, privacy, intellectual property, ethical guidelines

Get ready to think critically about ethical AI use!

www.cambridge.org

© Cambridge University Press & Assessment 2022