

From **Caversham House** this month: **3 strategic insights** on the trends reshaping how organisations think about AI, **3 tools** we've been testing at the frontline and **3 training courses** to help your teams build real capability and move from curiosity to confident action.

## 1. The Month in View: 3 Strategic Trends

Curated insights to help you read the room and lead the shift.

### Trend 1: AI Is Reshaping Who Gets Hired and Who Does the Hiring

AI has moved into the recruitment process itself. Algorithmic screening now filters candidates before a human reads a CV, hiring criteria have shifted toward AI fluency as a baseline expectation, and the entry-level pipeline is narrowing. There are two risks: the first is that the wrong people get filtered out, and the second is that organisations are hiring for AI skills they don't reinforce once people are through the door.

- LLM-based screening tools may already favour CVs written by the same model family doing the evaluating, creating a self-reinforcing loop that rewards AI-polished applications over real capability.
- Employers increasingly list AI literacy as a requirement, yet most still offer thin formal training, which leaves a gap between what they ask for and what they invest in.

**Call to action:** The organisations that win on talent will audit hiring models to assess candidate qualifications rather than AI-tool stylistic patterns, and close the gap between the AI skills they advertise for and those they actually develop.

1. Audit what your job postings ask for versus what your onboarding delivers
2. Test screening tools for bias toward AI-generated text over human substance
3. Train recruiters to assess judgment, not just tool familiarity

#### Relevant course:

[GenAI Foundation Course](#)

#### Read more:

[Nvidia exec: AI likes to use AI and algorithmic hiring may favour AI-written CVs | Employers want AI skills. What is the best way to learn them?](#) | [Airbnb CEO: two kinds of people who risk falling behind in the AI era](#) | [A reality check on the AI jobs hysteria](#)

### Trend 2: When Fluency Masks a Lack of Understanding

Last month we covered hallucinations: AI confidently making things up. This month the risk is subtler. AI assistants now produce work that looks rigorous because it's well-structured, plausible, professionally worded, but it can bypass the human reasoning it was supposed to support. The finished page looks fine while the thinking behind it may never have happened.

- AI writing tools smooth over gaps in reasoning rather than exposing them, producing persuasive structure even when the underlying question is unclear or flawed.
- Research shows that sustained AI-assisted writing weakens brain connectivity, reduces ownership of the work produced, and leaves users unable to recall or defend what they wrote.
- Business schools are now moving past tool training to teach when to trust, question or override AI outputs.

**Call to action:** Fluency is not the same as understanding. If your review processes can't tell the difference, the gap will show up in work, decisions and the credibility of your people.

1. Design reviews that test reasoning, not just the finished product
2. Alternate AI-assisted and unassisted practice to keep judgment sharp
3. Train people to challenge confident-sounding output, not just check for errors

**Relevant course:**

[AI Fluency: Pushing Back on Polished Output](#)

**Read more:**

[Eager to please AI assistants smooth over the gaps in our thinking](#) | [Your Brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing](#) | [Business schools move beyond the basics to teach collaboration with AI](#)

### **Trend 3: AI Is Disrupting the Business Models That Businesses Run On**



The disruption conversation has focused mostly on jobs and workflows. The bigger structural question is now surfacing: which businesses survive as AI absorbs the value proposition they were built on?

- AI is now being packaged directly into vertical tools that compete with design, legal and small-business software products.
- The most defensible competitive positions belong not to whoever controls the AI layer, but to organisations controlling the data AI needs to function. This makes data ownership a business model question, not just a technology one.
- First-mover advantage is real but closing: in real estate, only 25% of firms qualify as AI leaders against 40% across industries, and those treating AI as an add-on to legacy models are already falling behind on margins and timelines.

**Call to action:** The question is no longer whether AI will affect your business model but whether you are building the moats that survive what AI commoditises next.

1. Stress-test whether your core product is a scaffold or a business
2. Identify where proprietary data creates durable advantage and invest in it
3. Stay model-agnostic by treating AI providers as infrastructure, not strategy

**Relevant course:**

[AI Steering Committee and Governance](#)

**Read more:**

[What did Claude just kill? Six moats to protect yourself](#) | [The private capital opportunity in AI-enabled climate and sustainability sectors](#) | [Four ways to accelerate growth with AI and analytics](#) | [The AI-first real estate company: an opportunity for structural advantage](#)

---

## 2. The Toolbelt: 3 Updates from the Frontline

What we've been trying out at Caversham House and why it matters for your stack.

### Claude Design

Our Learning Director took to Claude Design to see how it compared with how Google's NotebookLM, which she currently uses to test UX options for our new pop-up courses. This is how she used Claude Design:

- Gave it our brand: fonts, colours, the full palette
- Shared one of our longer courses as a style reference
- Wrote prompts focused on UX (flow, navigation, readability for a shorter format)
- Reviewed what came back on the canvas, then refined through chat and inline comments

**Why we love it:** By the end, she had a couple of strong prototypes ready to take to the design team, which were way better than NotebookLM. Her favourite part was being able to leave inline comments on specific items that she wanted changed, which sped things up massively.

### First draft in the canvas for Copilot in Outlook

Launched in March

Our programme co-ordinator handles many emails from multiple addresses. They're enjoying using an Outlook update, where Copilot writes a first draft directly in the canvas and can then iterate with the user to continue improving it. Instead of generating a one-off draft, Copilot writes in place and asks clarifying questions about the goal, audience or tone, which is very useful for different inboxes that deal with distinctly different issues. The email is then updated as the user responds.

**Why we love it:** It lets users refine a message in a few quick turns until it's ready to send and keeps every change visible in Outlook with no copy/paste or formatting surprises.



# Google AI Studio

If you're working with AI models and haven't tried Google AI Studio, it's worth ten minutes of your time. It's a free browser-based tool that lets you test prompts against Google's latest Gemini models (currently 3.5 Pro, 3.5 Flash, and 3.1 Flash-lite) without writing a line of code.

**Why we love it:** The killer feature for us is the ability to re-run a single step in a conversation in isolation: rather than replaying an entire chat to test one change, you just tweak and re-fire that turn, which makes prompt refinement much faster.

We also use it to run the same prompt across different model sizes back-to-back, which quickly tells you whether you need the full power of 3.5 Pro or whether the lighter Flash variant does the job just as well.

**Value:** For more technical use cases, it's excellent for stress-testing agentic workflows: you can probe function calls and structured outputs with awkward edge-case inputs before anything goes near production.

---

### 3. Training Update: Learning Spotlight

Highlighting our newest course builds and strategy-focused updates designed to help your organisation lead the transition.

#### NEW COURSE

##### **AI Fluency: Pushing Back on Polished Output**

AI tools now produce work that looks rigorous because it is well-structured, plausible and professionally worded, but the underlying reasoning may never have happened. This module helps employees recognise the difference between fluency (slick output) and understanding (the thinking behind it), and build the habit of challenging confident-looking work, including their own.

**Takeaway:** A personal fluency self-check: a one-page guide for testing whether you can defend, recall and own AI-assisted work before it leaves your desk.

**We've been adding just-in-time pop-up modules to our training programme. If you're using AI Agents, these two are for your business:**

For Team Leaders

**When AI Agents Join Your Team: Governance for Leaders**

Agentic AI (tools that act autonomously rather than just responding) is arriving in team environments now. This module gives leaders the governance framework they need to set appropriate boundaries before something goes wrong, not after.

**Takeaway:** A one-page team agent governance template: a simple framework for defining what agents can and can't do in your team's name, who is accountable for agent actions, and what approval checkpoints must stay in human hands.

## For Teams

### **AI Agents: When AI Starts Doing Things for You**

AI has shifted from answering questions to taking autonomous actions. This module helps all staff understand what agentic AI is, what it means for their oversight responsibilities, and how to work with it safely.

**Takeaway:** A one-page oversight checklist: the questions to ask before letting an AI agent act on your behalf, with a simple framework for deciding when human approval is required and when it isn't.

---

## **Ready to take the next step?**

**Already a subscriber?** These modules are live in your training portal now.

**Thinking about subscribing?** Take a look at our [2026 Training Subscriptions](#) to see what's included and how it works for your team.

**Want to talk strategy?** [Book a briefing](#) or [contact us](#) and we can work through your 2026 AI priorities together.

