

Responsible AI in Digital Learning

A qualitative pilot from Falmouth University

Julian Green & John Sumpter



"This isn't really a story about AI. It's a story about how a university can adopt AI safely."



The Challenge

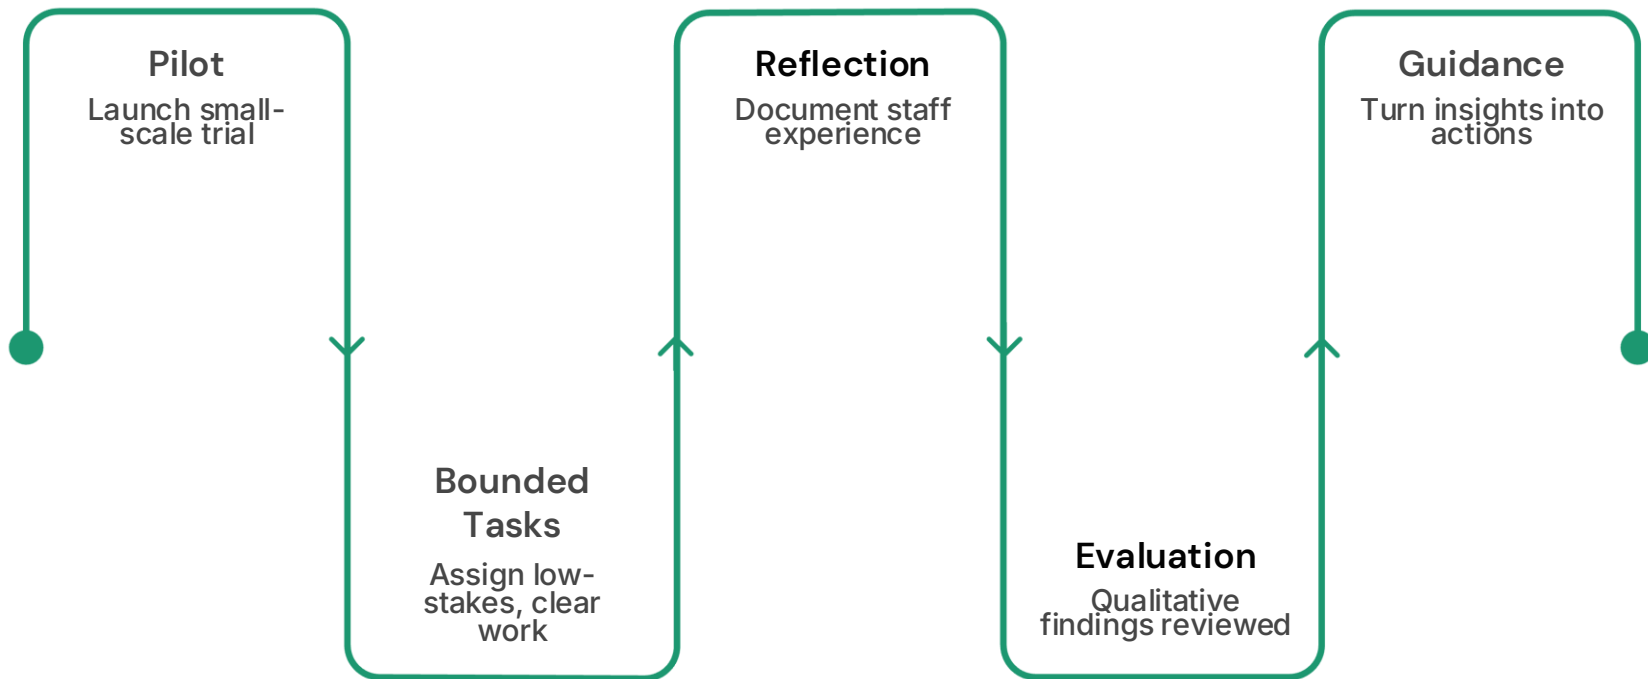
Should we use AI?

How do we use it safely?

How do we keep human expertise central?

The real challenge was governance, not technology.

What We Did



Bounded Tasks

Low-stakes, well-defined work only



Workshops

Collaborative sense-making across teams



Reflective Practice

Staff documented their experience throughout



Qualitative Evaluation

Findings shaped into actionable guidance

What We Tested

Four carefully chosen, low-risk task categories.



Email Summaries



Canvas Development



Accessibility



Meeting Notes

We deliberately avoided high-risk academic decision-making.

What We Found

Benefits

- Faster turnaround
- Less repetitive work
- Better first drafts
- Improved accessibility

Risks

- Hallucinations
- Oversimplification
- Over-trust in outputs

Human judgement remained essential.

Our Principles

Progressive

Embrace AI where it adds genuine value



Safe

Protect data and institutional integrity



Transparent

Always disclose use and review outputs



Human in the loop. Always.

Do's and Don'ts

DO ✓

Brainstorm and draft

Structure and organise

Improve accessibility

Use as a thinking partner

DON'T ✕

Upload confidential data

Upload student information

Delegate academic judgement

Assume AI is always better!

AI is an assistant—not the decision maker.



The Biggest Lesson

The innovation wasn't AI.

It was creating a repeatable method for responsible experimentation.



Discussion

Where are you already using AI?

Knowingly or not—it may already be in your workflows.