

Practice-Based Article

Building AI Literacy Through Library-Led Copilot Training: A Practice-Based Case Study

Mairéad Mc Keown, SALAI, Library, Knowledge & Market Intelligence Manager, Bord Bia

David Lombard, EU Projects Accounts Administrator, Bord Bia

ABSTRACT

Libraries across all sectors are increasingly expected to support their organisations in navigating artificial intelligence (AI) tools while also building confidence, critical judgement and ethical awareness among staff. This practice-based article summarises a conference presentation delivered by Mairéad McKeown and David Lombard at the Health Sciences Libraries Group (HSLG) Annual Conference 2026. It describes the design and delivery of a Microsoft Copilot graduate upskilling programme at Bord Bia and positions librarians and knowledge professionals as capability builders and trusted guides in organisational AI adoption. It outlines a structured, time-bound learning programme that combined foundations in AI literacy with hands-on experimentation embedded in everyday work, supported by knowledge management (KM) tools and reflective practice. While developed in a specific organisational context, the programme design, learning principles and librarian-led interventions are transferable across library sectors, including health, academic, public and special libraries.

KEYWORDS

AI literacy, capability building, knowledge management, libraries, practice-based learning

INTRODUCTION

Artificial intelligence (AI), and generative AI in particular, is rapidly reshaping how information is created, accessed and used. For librarians, this shift presents both opportunity and responsibility. Library professionals are increasingly asked to support colleagues who are curious about AI tools, uncertain about their limitations and concerned about issues such as accuracy, bias, transparency and professional judgement. Building organisational AI capability therefore requires more than technical training; it demands a focus on literacy, critical thinking and ethical use.

This article summarises a presentation delivered at the Health Sciences Libraries Group (HSLG) Annual Conference 2026. It shares a practical case study from Bord Bia, Ireland's food board, on a Microsoft Copilot graduate upskilling programme. The programme was designed and delivered by the Library, Knowledge and Market Intelligence team in partnership with a colleague seconded from Finance, with the aim of building confident, responsible and value-driven use of generative AI among early-career professionals. The purpose of this article is to offer librarians a concrete, adaptable model for leading AI capability building within their own organisations.

CONTEXT AND RATIONALE

Bord Bia operates in a complex, fast-moving environment where digital demands on graduates are accelerating year on year. Graduates are expected to synthesise large volumes of data and information to support evidence-based decisions/recommendations, while also balancing rotation learning with the needs of the organisations' functions within which they are placed.

The introduction of Microsoft Copilot generated both enthusiasm and uncertainty. Although graduates were comfortable with digital tools, many were not yet confident Copilot users and needed support to collaborate with it appropriately and responsibly in professional contexts, while applying human judgement.

Recognising this early-career inflection point, Bord Bia's Industry Talent team identified an opportunity to intervene by embedding AI capability and literacy into the existing graduate rotation learning framework. The Library, Knowledge and Market Intelligence team was asked to design and deliver a programme that would build confidence, critical thinking and responsible practice from the outset. Rather than positioning Copilot as a simple productivity shortcut, the programme framed it as a collaborative assistant that could support synthesis, creativity and efficiency when used with appropriate human judgement, organisational context and domain knowledge.

This approach aligns closely with the traditional strengths of librarianship, including the critical evaluation of information, reflective practice and ethical stewardship. It was also deliberately future-focused. The programme was designed to align closely with all seven of the IFLA's Trends on *Facing the Future of Information with Confidence* (IFLA, 2024) ensuring that graduates were not only learning how to use an emerging tool, but were also developing the skills, mindsets and critical awareness required to navigate an increasingly complex information environment. This alignment helped shift the focus away from simply providing tool-specific training towards transferable information literacy capabilities, such as critical evaluation, synthesis, ethical awareness and professional judgement, that remain relevant regardless of how technologies evolve. By grounding generative AI capability building within these internationally recognised trends, the programme helped ensure relevance, resilience and transferability beyond any single technology or organisational context.

PROGRAMME DESIGN PRINCIPLES

Our programme incorporated several design principles that may be of relevance across library sectors:

- **Library-led capability building:** The programme was designed and facilitated by the library and knowledge team in conjunction with a colleague seconded from the Finance team (this colleague participated on a previous Copilot up-skilling programme), reinforcing the role of librarians as AI educators and critical guides rather than passive tool adopters.
- **Plain English AI literacy:** Concepts such as large language models, hallucinations and prompt quality were explained in accessible, nontechnical language.
- **Timebound, realistic learning:** All peer-to-peer knowledge exchange sessions were limited to 45 minutes, respecting workload pressures while modelling sustainable learning behaviours.
- **Learning by doing:** Participants applied Copilot to real organisational tasks and day-to-day activities rather than hypothetical exercises.

- **Community connections and collective learning:** The programme was intentionally designed as a community of practice, where graduates learned with and from one another while experimenting with Copilot adoption and AI literacy together. Peer discussion, shared problem-solving and visible learning helped normalise uncertainty and supported collective sensemaking.
- **Psychological safety by design:** Experimentation was encouraged in a psychologically safe environment, where curiosity, questioning and iteration were valued over performance. This enabled participants to practise, learn and adapt together without fear of failure or judgement.
- **Ethics and judgement first:** Responsible use, in line with Bord Bia's generative AI policy, alongside verification and transparency, was embedded throughout rather than treated as an add-on.

PROGRAMME STRUCTURE

The Copilot graduate upskilling programme was delivered over three distinct phases.

Self-Directed Foundations:

In the foundations phase, graduates focused on AI literacy and critical thinking and completed a self-directed LinkedIn Learning pathway on Microsoft Copilot. This provided a shared baseline of learning covering responsible and appropriate use of generative AI, effective prompting, and a range of practical Copilot use cases across the Microsoft 365 suite of applications.

Completion of the learning pathway was reinforced through assessment. Graduates successfully completed an associated exam and were awarded industry recognised professional certificates, helping to build both confidence and credibility in their foundational Copilot knowledge before progressing to applied experimentation.

Peer to Peer Knowledge Exchange on Applications and Tools:

A core feature of the programme was encouraging graduates to use Copilot in the flow of their everyday work rather than treating it as a standalone tool. Participants were supported to test Copilot while drafting emails, summarising documents, preparing briefings, structuring presentations, planning meetings and reflecting on learning. This helped demystify AI use and positioned it as an assistive layer within existing workflows.

The facilitators emphasised that value emerged not from the tool itself, but from how thoughtfully it was applied. Graduates were encouraged to iterate on prompts, critique outputs and combine Copilot suggestions with their own subject knowledge and professional judgement. This approach fostered more creative ways of working, such as exploring alternative framings of problems, generating options rather than answers, and using Copilot to challenge assumptions rather than replace thinking.

Alongside Copilot, practical KM tools such as conversational leadership techniques, peer assist approaches and reflection were introduced. These tools helped participants situate AI outputs within a wider knowledge ecosystem and recognise when human expertise, collaboration or escalation were required, reinforcing that Copilot outputs were starting points for dialogue rather than finished answers. Every two weeks graduates met as a group with the facilitators, where they shared what was working well and what wasn't, the group were inspired by different use cases and application of the best practice prompt

framework and for many this was their favourite part of the programme.

Creative Finale and Impact

In the final phase, graduates completed a reflective lesson learned exercise and then fed their reflections into the Copilot Creator app using a structured prompt. They were encouraged to apply the RGCSE prompt framework (role, goal, context, source, expectations) to create a multimedia asset of their choice that captured the story of their learning journey. (RGCSE is a simple prompt checklist used to make expectations explicit and improve output quality.)

Graduates were not expected to produce polished or perfect outputs. Instead, they were encouraged to apply the skills and knowledge they had developed throughout the programme in a psychologically safe environment. This approach allowed participants to test creative ideas, experiment with different formats and reflect on impact without fear of failure. Graduates reviewed what worked, what did not and how their approach might change, reinforcing learning through reflection rather than performance. Graduates produced a range of multimedia assets, including videos, storyboards, posters, infographics, poems and presentations.

The programme concluded with a creative finale where participants shared their reflections and voted on the work that most inspired them. This peer led recognition reinforced a sense of community and embodied a core KM principle: valuing and learning from the collective wisdom of the group.

OUTCOMES AND EVALUATION

As a practice-based initiative, the programme prioritised observable changes in confidence, judgement and day-to-day use over purely technical mastery. Evaluation was kept lightweight and embedded within delivery to reduce burden while still capturing learning and improvement opportunities.

- **Participation and completion:** attendance across the peer sessions, completion of the LinkedIn Learning pathway, and exam/certificate attainment by all graduates.
- **Confidence and capability lift:** brief pre- and post-programme self-assessment (for example, confidence using Copilot, confidence evaluating outputs and confidence knowing when not to use generative AI).
- **Quality of practice indicators:** evidence of iterative prompting, explicit verification behaviours, and appropriate use of sources (e.g., citing internal documents, checking figures, confirming policy constraints).
- **Applied outcomes in the flow of work:** participant-reported examples where Copilot supported drafting, synthesis or planning, alongside the human review steps taken before outputs were shared.
- **Reflective learning capture:** themes emerging from the “lessons learned” exercise and creative finale (what changed in how graduates approached tasks, what pitfalls they encountered, and what guardrails helped).
- **Continuous improvement:** a short retrospective with facilitators to identify what to keep, change or add for the next cohort (for example, more peer-to-peer exchange in smaller groups or shorter check-ins between self-directed learning modules).

GUARDRAILS FOR RESPONSIBLE USE

Responsible use was reinforced throughout in line with Bord Bia's generative AI policy, with a consistent emphasis that Copilot outputs require human oversight. The following guardrails summarise the types of controls and behaviours that supported safe experimentation and professional judgement.

- **Protect sensitive information:** only use Copilot in line with the do's and don'ts set out in Bord Bia's generative AI policy; responsibility for compliance sits with each user.
- **Assume errors are possible:** treat outputs as drafts; check facts, figures, citations and any claims that could affect decisions, stakeholders or reputation.
- **Verify against trusted sources:** confirm outputs against authoritative internal material (policies, reports, approved messaging) and external primary sources where relevant.
- **Keep human accountability:** the user remains responsible for what is sent, published or decided; Copilot does not replace professional judgement.
- **Be transparent when appropriate:** follow internal expectations for disclosing gen AI assistance, particularly for externally facing material or high-stakes work.
- **Use AI for options, not authority:** favour prompts that generate alternatives, outlines and questions, and avoid treating the tool as a single source of truth.
- **Know when to escalate:** seek advice from the library/knowledge team, IT, corporate governance or line management when a task involves uncertainty.

HOW TO ADAPT THIS MODEL

Although delivered in a specific organisational context, the approach can be adapted to suit different library sectors, workforce profiles and technology environments. The following options help translate the model without requiring extensive new resources.

- **Start with a minimum viable programme:** run a pilot with one cohort, evaluate and reflect. If the proof of concept is viable, begin to scale and iterate.
- **Use whatever tools are available:** if Microsoft Copilot is not licensed, focus on AI literacy, prompting, verification and workflow design using organisation-approved alternatives (or non-tool-based scenarios).
- **Embed into existing structures:** attach the programme to induction, continuing professional development (CPD) time, performance management programmes, quality improvement forums, or graduate/trainee rotation learning.
- **Localise use cases:** co-design examples with teams (e.g., briefing notes, evidence summaries, communication drafts, meeting preparation, policy comparison) so practice aligns with real work.
- **Scale support through community and partnerships:** if resourcing is a constraint, look for co-leads and champions outside of the library (e.g., in human resources or learning and development, IT, governance, business units) and set up a community of practice, peer champions or office hours, where dedicated support is available, so learning continues after the formal programme ends. In this practical

case, the programme co-lead was a member of the first colleague Copilot programme and was not a member of the library team; this brought valuable diversity of thinking and firsthand learner experience into design and delivery.

- **Be always guided by governance:** ensure all instruction is delivered in line with your organisation's generative AI policy.
- **Keep evaluation lightweight:** collect a small number of indicators (confidence lift, examples of applied use, verification behaviours, skills, knowledge) and use these to refine the next cohort.

ROLE OF THE LIBRARY

A key message here is that librarians and co-programme leads are well placed to lead AI capability initiatives. In this programme, the library team acted as facilitators, critical friends and ethical guides. They supported participants in validating outputs, identifying gaps and understanding when not to use AI. This role mirrors the advisory function librarians already perform across sectors, including evidence appraisal, research support, information governance and professional education.

Relevance Across Library Sectors

Although this case study was developed within a single organisation, the approach has relevance for librarians working across health sciences, academic, public and special libraries. The emphasis on embedding AI experimentation into everyday tasks, alongside critical appraisal, ethical awareness and reflective practice, aligns with core professional values common to all library contexts. The programme demonstrates how AI capability can be built incrementally within existing training, induction or continuing professional development structures, without requiring deep technical expertise or significant additional resources.

CONCLUSION

The programme aimed to move graduates from being Copilot-aware but not yet confident or consistent in extracting value, to becoming confident, responsible users who continuously apply the knowledge, skills and habits needed to collaborate with Copilot to enhance creativity and enable smarter ways of working. This practice-based case study demonstrates how libraries can move beyond tool demonstrations to lead meaningful AI capability building. By encouraging experimentation in the flow of daily work, supported by AI literacy, knowledge management (KM) tools and reflective practice, librarians can help colleagues develop more responsible and more creative ways of working.

Importantly, learning did not end when the formal programme concluded. Graduates are now transitioning into a Bord Bia Copilot Alumni community, designed to sustain momentum and embed learning as an ongoing, collective endeavour. This community brings together a wider, multigenerational and cross-functional cohort of colleagues who have also undertaken the Copilot programme, enabling continued sharing of practice, peer learning and reflection. By positioning AI capability as something that is continually practised, discussed and refined within a supportive professional community, the programme reinforces the idea that AI literacy is not a one-off intervention, but a sustained capability built over time.

The approach described offers a flexible, transferable model that can be adapted by librarians across sectors seeking to build organisational confidence and capability in AI.

Figure 1.

Image by Ronan Sweeney

**ACKNOWLEDGEMENTS**

The authors acknowledge the use of Microsoft Copilot as a brainstorming and editing support tool during the preparation of this article.

AUTHOR CONFLICT OF INTEREST STATEMENT

The authors declare no known conflicts of interest.

REFERENCES

International Federation of Library Associations and Institutions. (2024). *IFLA trend report 2024: Facing the future of information with confidence*. Digital Media Research Centre. <https://www.ifla.org/trend-report/>

FURTHER READING

Libfocus. (2024). *Webinar review: AI and workforce development in libraries: Overcoming barriers and embracing lifelong learning*. Libfocus Irish Library Blog. <https://tinyurl.com/yrfku4t>

Digital Leadership Forum. (2024, 19 September). *Digital transformation and AI: Collaborating in the new world of work* [Webinar]. Retrieved March 4, 2026, from <https://www.thedigitalleadershipforum.com/>

Bord Bia. (2025, January). *Copilot skills building programme induction* (Copilot Skills Building Pilot Programme) [Internal resource].

McKeown, M., & Lombard, D. (2026). *Building AI capability through a Copilot graduate upskilling programme* [Conference presentation]. Health Sciences Libraries Group Annual Conference, Dublin, Ireland. <https://hslg.ie/2026/03/13/conference-2026-presentations/>

Microsoft. (2024). *Responsible AI principles*. Retrieved March 4, 2026, from <https://www.microsoft.com/ai/responsible-ai>

