AI Governance Sprint for Publishers

A 90-Minute Workshop for Mapping AI Use, Risk, and Guardrails

A practical self-run workshop for small-to-mid-sized trade publishers

Who is this for? Marketing and publicity staff using general-purpose AI tools such as ChatGPT, Claude, Gemini, Copilot, Canva AI, Adobe Firefly, Grammarly, or similar tools to support campaign planning and content creation.

What's the issue to resolve? AI tools can help publishing teams brainstorm campaign angles, draft first-pass copy, summarize public information, and explore new workflows. The critical risk is not simply that AI might produce weak or inaccurate copy. The higher-risk issue is staff, reps, or freelancers pasting unpublished manuscripts, author information, internal title materials, sales data, unreleased artwork, or confidential campaign strategy into general-purpose tools without clear consent, governance, review, or understanding of tool-retention and training terms.

This workshop helps your team sort those uses into practical categories that align with 1) your firm's values and 2) with your governance decisions.

The goal is not to write a perfect AI policy. The goal is to create a working map of how AI is being used, what the risks are, and what guardrails are needed next.

Use this with the **Publisher AI Use-Case Mapping Tool**.

This is not a one-time policy exercise. It is a repeatable process for identifying risk, approving safe uses, and protecting authors, readers, staff, sales reps, and the publisher.

Workshop at a Glance

Who this is for

Small-to-mid-sized publishers who want a practical starting point for AI governance, especially across marketing, publicity, and sales. The workshop can be adapted for Editorial, Design, and Operations.

Who should attend

Include 3–6 people who understand day-to-day marketing and publicity workflows.

Suggested participants:

- Publisher or Associate Publisher
- Marketing Director or Marketing Manager
- Publicists and marketing assistants
- Sales reps or an in-house sales manager
- Optional: Design, Operations or admin lead

What you will produce

By the end of the session, you should have:

- A short set of AI principles
- A list of current or likely AI use cases
- A risk ranking for each use case
- A list of red-flag inputs
- A first version of practical AI guardrails
- A 30-day action plan

The framework: This workshop uses the NIST AI Risk Management Framework

NIST Function	Plain-English Meaning	Workshop Output
Govern	Who is accountable, and what are our boundaries?	Principles and red lines
Map	Where is AI being used, and what information is involved?	Use-case inventory
Measure	How risky is this use case?	Risk rating
Manage	What should we allow, review, escalate, or prohibit?	Guardrails and next steps

The agenda

Time	Step	What happens
0–10 min	Set the frame	Agree on purpose and tone
10–25 min	Govern	Draft principles and red lines
25–45 min	Map	Inventory real AI use cases
45–65 min	Measure	Assess risk using the mapping tool
65–80 min	Manage	Decide what needs guardrails
80–90 min	Commit	Assign owner and next steps

Minimum viable version

Only have 30 minutes? Do this:

1. List current AI use.
2. Identify red-flag inputs.
3. Name an AI owner.
4. Create an approved-tool list.
5. Schedule a follow-up review.

That is enough to get started.

Before You Begin

Set the tone

This is not an “AI is good” or “AI is bad” meeting.

This is a visibility meeting.

The goal is to understand what is happening, what is useful, what feels risky, and what needs clearer guidance.

Materials

- Publisher AI Use-Case Mapping Tool
- Shared screen
- Timer
- One facilitator
- One note taker

Facilitator note

Encourage honesty. Many teams discover that staff are already using AI informally. Treat that as useful information, not a failure.

1. Govern: Set the guardrails

Minimum viable time: 15–20 minutes.

Decide who owns AI use and what is clearly allowed, restricted, or prohibited. *See examples.

- **Allowed by default:** AI use based on public, non-sensitive material with human review.
- **Grey zone:** AI use involving mixed or uncertain sensitivity, rights, consent, or public-facing impact.
- **Prohibited unless approved:** AI use involving clearly confidential, restricted, personal, contractual, sales, rights, or substantial unpublished creative material in public/general-purpose tools.

Ask:

- Which tools and use cases are approved for staff use?
- Who can approve grey-zone uses?
- What information must not be entered into public AI tools?
- Who is accountable for final public-facing content?
- What cases require author consent?

Start with:

- Create a short list of prohibited inputs for public AI tools, i.e., unpublished manuscripts, substantial manuscript excerpts, author personal information, sales data, contracts, rights information, optional: confidential title sheets, internal marketing briefs, unreleased artwork, and interior spreads unless the use case and tool have been approved.
- Require a named human reviewer for AI-assisted public-facing content.
- Name an AI owner or small review group for grey-zone cases.

2. Map: Identify actual use cases

Minimum viable time: 20–30 minutes

Do not start with abstract principles. Start with what staff are actually doing or are likely to do.

For each use case, record:

Tool | Task | Department | Input type | Data sensitivity | Risk level | Human review | Owner | Repeatability | Workflow opportunity | Notes

Example use cases include drafting social copy from public-facing metadata, rewriting an author bio for variations in length and purpose (social media bio, pitch bio, about the author, official back cover bio), generating campaign angles from a publicly available excerpt, summarizing a marketing brief, brainstorming comparable titles, analyzing publicly available reader reviews, creating Amazon A+ content ideas, or summarizing trends.

3. Measure: Rate the risk

Minimum viable time: 15–20 minutes

Use a simple rating: **Low / Medium / High / Prohibited unless approved.** The point is shared judgment, not mathematical precision.

Factor	Lower Risk	Higher Risk
Input sensitivity	Non-copyrighted public metadata, published reviews	Unpublished manuscript, sales data, author personal information
Rights or consent	Public or already approved material	Author-owned, unpublished, licensed, or unclear rights
Tool risk	Approved/private tool with clear terms, training turned off	Public tool with unclear retention or training terms
Output impact	Internal brainstorm	Public-facing copy, claims, positioning, or retailer content
Human review	Named reviewer required	No defined reviewer

Practical rule: If the input is confidential, unpublished, personal, commercially sensitive, copyright protected, or rights-sensitive, treat the use case as high risk until reviewed.

Grey-zone AI use refers to a workflow where AI may be helpful, but the risks are context-dependent because the task involves unpublished, internal, rights-sensitive, personal, commercially sensitive, or reputationally important material. Grey-zone uses are not automatically prohibited, but they should be treated as medium to high risk until reviewed by tool, input type, consent/rights implications, human review process, and accountable owner.

For publishers, grey-zone use often includes:

Use case	Why it is grey-zone
Generating campaign angles from an unpublished manuscript excerpt	Helpful for marketing, but involves confidential creative work and possible author consent issues.
Rewriting an author bio	Low risk if based on public bio; higher risk if using private details or sensitive identity/context.
Summarizing a title information sheet	Useful internally, but title sheets may include confidential positioning, sales expectations, rights info, or strategy.
Brainstorming comp titles from internal positioning	Useful, but outputs may create inaccurate or misleading comparisons if not checked.
Creating Amazon A+ ideas from cover/interior spreads	May involve unreleased or licensed artwork with rights restrictions.
Using reader reviews of comp titles to infer audience sentiment	Reviews may be public but protected by copyright. AI may overgeneralize, misrepresent sentiment, or fabricate patterns.

4. Manage: Decide what to do next

Minimum viable time: 15–20 minutes

Assign a response based on risk level.

Risk Level	Recommended Response
Low	Allow with human review. Example: drafting social copy from public metadata.
Medium	Allow with documented review and approved inputs. Example: rewriting an approved author bio.
High	Require manager or publisher approval, approved tool, and possible author/rights review. Example: using manuscript excerpts or unreleased artwork.
Prohibited unless approved	Do not use public AI tools. Example: uploading sales data, contracts, rights information, confidential briefs, or substantial unpublished manuscript text.

5: Commit

Turn the workshop into a living process

Time: 10 minutes
Output: 30-day action plan

This is where the work becomes governance.

AI Principles & AI in Practice

Use two layers so the work does not become stale.

Layer	Purpose	Examples
Layer 1: Principles	Stable commitments	Protect author trust, preserve human accountability, safeguard confidential information. Layer 1 should change slowly. This is good public-facing information to have in a policy.
Layer 2: Practice	Living controls	Approved tools, prohibited inputs, review roles, escalation rules, use-case inventory. Layer 2 should be reviewed regularly. This is good for internal use.

Ask: Who owns this going forward?

Possible owners:

- Publisher
- Associate Publisher
- Marketing Director
- Operations Lead
- Small AI working group

Do: Within 30 days, decide who will:

- Finalize AI principles
- Create or update the approved-tool list
- Confirm prohibited inputs
- Review high-risk use cases
- Share guidance with staff, reps, vendors, and contractors
- Schedule the next review

Recommended review cadence

- Monthly for the first 3 months
- Quarterly after that
- Any time a new tool, workflow, contract issue, or author concern emerges

Recommended first implementation

Start with a 60–90 minute working session:

1. Choose one department: i.e., marketing/publicity.
2. List 6–8 real or likely AI use cases.
3. Complete the Use-Case Mapping Table.
4. Identify low-risk uses that can proceed with human review.
5. Identify high-risk or prohibited inputs.
6. Name the person responsible for approving grey-zone uses.
7. Revisit the map quarterly or when a new tool, workflow, or concern emerges.

This approach gives staff room to experiment responsibly while creating clear red lines around unpublished creative work, author information, confidential business data, and public-facing claims.

Important limit: This framework is not legal advice and does not replace review of contracts, author agreements, privacy obligations, funder requirements, tool terms, copyright considerations, or internal policies.

Measuring Success

You do not need a perfect AI policy to start.

You need:

- Visibility into current AI use
- Shared language for risk
- Clear red lines
- Named accountability
- A practical way to revisit decisions

That is enough to move from informal experimentation to responsible AI governance.

The goal is not unrestricted adoption or blanket prohibition.

The goal is responsible experimentation with author trust, staff judgment, and publisher accountability intact.

Example Use Cases

Allowed by Default: Public, non-sensitive information + human review

These are generally good candidates for experimentation.

Case	Why It's Usually Allowed
Draft social media posts from an approved catalogue description	Input is already public and easily verified
Generate newsletter subject line options from approved copy	Low risk and easy to review
Rewrite approved marketing copy for different audiences	Source material already approved
Summarize public reader reviews for internal use	Uses public information
Brainstorm blog post ideas for a publishing website	Creative support only
Generate interview questions for an author event	No sensitive input required
Suggest keywords for retailer metadata optimization	Public-facing but low consequence
Create alternative headlines for a press release	Easy human review
Turn an approved book description into retailer-specific formats	Based on approved source material
Summarize public industry articles for internal discussion	Public information, internal use

Rule of caution: If the information already appears publicly on your website, in your catalogue, on retailer sites, or in media coverage, it's often a reasonable starting point for AI-assisted drafting. Beware that third-party material may be protected by copyright and the implications of use in AI prompts should be reviewed.

Grey Zone

Useful, but context matters. These are often the most interesting publishing use cases. They are not automatically prohibited. They require review before becoming standard practice.

Use Case	Why It's Grey Zone
Generate campaign angles from an unpublished manuscript excerpt	Useful, but involves confidential creative.
Rewrite an author bio using public and internal information	May introduce inaccuracies or unwanted framing. Needs author review/approval.
Summarize a title information sheet	May contain internal positioning/strategy
Brainstorm comparable titles using unpublished book details	Helpful, but may expose confidential information, and may be inaccurate if the AI prompt lacks enough context.
Draft publicity pitches from an unpublished manuscript, excerpt, book description	Depends on what content is provided
Analyze reader reviews of comparable titles to identify themes	May overgeneralize or misrepresent sentiment
Create retailer copy from a marketing brief	Brief may contain confidential info
Generate event or speaking topics for an author	May require author review and alignment
Use cover images to brainstorm Amazon A+ content concepts	Depends on rights and intended use
Summarize internal meetings or campaign discussions	Internal information may not belong in public AI tools

Questions to ask

- Is the material unpublished?
- Does it contain internal strategy?
- Could an author be surprised that this was entered into an AI tool?
- Is there a human reviewer?
- Would you be comfortable emailing this information to an external vendor?

If the answer is "maybe," you're likely in grey-zone territory.

Prohibited Unless Approved

High-risk inputs and workflows

These should trigger review before use and may be prohibited in public AI tools altogether.

Use Case	Why It's High Risk
Uploading a full unpublished manuscript to AI tools	Confidential creative work
Uploading substantial manuscript excerpts	Copyright and consent concerns
Uploading author personal information	Privacy concerns
Uploading sales reports	Commercially sensitive
Uploading royalty information	Confidential financial data
Uploading contracts	Legal and rights implications
Uploading rights information	Contractual and business risk
Uploading media contact databases	Third-party personal information
Uploading customer or newsletter subscriber data	Privacy and compliance concerns
Uploading unreleased cover concepts	Copyright and pre-publication concerns
Uploading interior spreads before publication	Copyright and licensing concerns
Asking AI to evaluate manuscript submissions for acquisition decisions without clear governance	Significant fairness, accountability, and transparency concerns

Rule of caution. If the information is:

- unpublished
- personal
- contractual
- financial
- Rights-related or copyrighted
- commercially sensitive

then it should be treated as high risk until someone explicitly reviews the workflow.

A Publishing Example That Shows All Three Categories

Marketing a New Novel

Allowed by Default

- Use ChatGPT to generate Instagram captions from the approved book description.
- Generate newsletter subject lines from approved marketing copy.

Grey Zone

- Paste a chapter excerpt into Claude to brainstorm publicity angles.
- Summarize the title information sheet to identify positioning opportunities.

Prohibited Unless Approved

- Upload the entire manuscript.
- Upload author contract details.
- Upload sales projections and royalty information.

Thinking through the progression helps teams understand that the same project can contain workflows that belong in all three categories. The question is not "Are we using AI?" The question is "Which parts of this workflow carry which risks?" That's the mindset the workshop is trying to build.

Further Reading

The US Department of Commerce's National Institute of Standards and Technology AI Risk Management Framework (NIST) is a practical and simple framework for governing, mapping, measuring, and managing AI risk. It was developed in collaboration with the private and public sector.

- NIST AI Risk Management Framework
  https://www.nist.gov/itl/ai-risk-management-framework
- NIST AI RMF Playbook
  https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

Sarah Downey is a BC-based consultant and former Executive Director who helps Canadian nonprofits adopt AI ethically. Her focus is on AI adoption that protects the organization's values, frees up capacity for higher-level tasks, and keeps people at the centre. Her approach to AI Policy recommends two layers: one that is a baseline of principles and values that are unlikely to change regardless of technology. The second is focused on internal communications about day-to-day use of approved tools.

- Sarah Downey — Board-Ready AI Policy Checklist
  https://www.sarahdowneyconsulting.com/post/board-ready-ai-policy-checklist
- Sarah Downey — Policy Design in the AI Era
  https://www.sarahdowneyconsulting.com/post/policy-design-in-the-ai-era

Here are two examples of AI Policy. Wiley's is comprehensive and focused on AI as part of content creation. It highlights risks and prohibited use that are worth noting. New Society's policy is the model used for the Mapping Tool mentioned in this sprint document.

- Wiley AI Guidelines for Book Authors
  https://www.wiley.com/en-us/publish/book/resources/ai-guidelines/
- New Society Publishers AI Policy
  https://newsociety.com/about-us/our-commitments/ai-policy/

Last, the EU's approach to AI literacy programs may be of interest to those looking to learn more about the opportunities to use AI in various workflows.

- https://artificialintelligenceact.eu/ai-literacy-programs/

Appendix: Transparency Statement on AI Use

This document was developed with AI assistance. AI was used as a drafting, structuring, and synthesis tool, not as a replacement for human judgment, subject-matter expertise, or final editorial responsibility.

Purpose of AI Use

AI was used to help translate concepts from the Responsible AI Professional certificate offered by BC+AI, along with publicly available AI risk management frameworks, into a practical guide for small-to-mid-sized Canadian book publishers. The goal was to explore how publishers can assess AI use in marketing and publicity workflows, especially where staff may be working with unpublished manuscripts, author information, internal title materials, sales data, or unreleased visual assets.

Human and AI Contributions

Area	Author Contribution (Monique Sherrett)	AI-Assisted Contribution
Project concept	I identified the publishing-sector use case, audience, practical constraints, and focus.	AI helped refine the project scope into a responsible AI risk assessment format.
Sector expertise	I contributed knowledge of publisher marketing/publicity workflows, including title information sheets, marketing briefs, author bios, comparative titles, metadata, Amazon A+ content, and campaign planning.	AI helped organize those workflows into a risk assessment structure.
Framework selection	I chose to incorporate NIST's AI Risk Management Framework as an agile governance approach.	AI helped map the project to Govern, Map, Measure, and Manage.
Risk identification	I identified confidentiality, copyright, author consent, and unpublished manuscript use as the central risks.	AI helped articulate the risk in formal language and connect it to practical publisher scenarios.
Drafting	I provided direction on tone, audience, scope, use case, and desired deliverables.	AI generated draft language, tables, headings, and implementation steps.

Review and editing	I reviewed the outputs for relevance, usefulness, and fit with publishing-sector realities.	AI suggested wording, structure, and concise framing.
Source use	I provided source materials and links, including NIST resources, public publisher AI policies, and Responsible AI Professional course-related materials.	AI summarized and incorporated those sources into the draft, with attribution where available.
Final accountability	I am responsible for the final selection, editing, interpretation, and use of this document.	AI did not independently verify legal obligations, contractual requirements, or current tool terms.

Review and Verification

All AI-generated content was reviewed by me, Monique Sherrett, before inclusion. I assessed whether the recommendations were appropriate for small-to-mid-sized trade publishers and revised the content to reflect the intended audience: busy publishing professionals who may range from cautiously enthusiastic to strongly skeptical about AI tools.

Where the document refers to external frameworks or public AI policies, reference links are provided. Readers should consult those sources directly and seek legal, privacy, or rights advice where appropriate. This document is intended as a practical governance framework, not legal advice or a substitute for a formal AI policy.

Limits of AI Assistance

The AI system did not have independent knowledge of a specific publisher's contracts, author agreements, privacy obligations, internal systems, insurance requirements, or tool licenses. It also did not conduct a legal review. Recommendations should therefore be treated as a starting point for internal discussion, not as definitive compliance guidance.

Accountability Statement

I used AI to support brainstorming, structuring, synthesis, and drafting. I retained responsibility for the project framing, publishing-sector judgment, final edits, risk prioritization, and recommendations. The final document reflects my assessment of how small-to-mid-sized Canadian trade book publishers can begin mapping AI use in a practical, transparent, and responsible way.