

Defensible AI in Practice

A Hands-On Workshop for Biomedical Research

PRESENTED BY THE DIGITAL PRECLINICAL SOCIETY (DIPS) · HOSTED BY NESBR AT TAKADA

Tuesday, October 20, 2026 · 9:00 AM – 4:00 PM

About the workshop

Generative AI is already in the daily work of biomedical research (literature, protocols, manuscripts, posters, correspondence), usually without a shared standard for how it is done well. This workshop covers that standard by demonstration. This is the day you do it yourself, on your own material, with someone in the room when it matters. You bring real work from your own desk; you leave having produced output you can defend, along with the rules and records to keep producing it after you go home.

Who this is for

- Laboratory animal veterinarians
- IACUC and IBC coordinators
- Research compliance and administration staff
- Science communications teams
- Scientists at academic and industry member institutions

No prior experience with AI tools is assumed. Skepticism is welcome and expected.

What you leave with

- A working prompt rebuilt from one of your own tasks
- A confidentiality decision rule, filled in for the material you actually handle
- A verification record you can reuse on any AI-assisted output
- One recurring task turned into a workflow with guardrails written down

How the day works

Six hands-on labs, three short breaks, and lunch. You work on your own laptop throughout. Each lab is a short setup, then you do the task, then a live walk-through of real examples from the room, including your own, if you volunteer them. The last block reverses that order: two problems from the room get solved live first, then you work yours.

Labs use Claude; the techniques transfer to other assistants, and where they don't, the day says so and explains why. Where a lab touches a rule, the rule is named by its exact type. A policy notice, a draft guidance, a set of principles, a reflection paper, a position statement, and a binding Regulation are not interchangeable, and the day treats them accordingly.

Before you arrive

- Bring a laptop you can run an AI assistant on, and test it from that exact machine beforehand, not from a personal device you're leaving at home.
- Confirm your institution permits the tool. Many restrict them; find out before the day, not during Lab 1.
- If your institution blocks access, come anyway and say so at registration. You'll be paired with someone who has access: you drive the thinking and the writing, your partner drives the tool. Everything you produce is still yours.
- Bring one real task from your own desk. Not your best one. The one on top.

Agenda

Time	Segment	What you do
9:00 – 9:20	Opening: What "Defensible" Means Today	The day's throughline: technique, the lines you can't cross, verification. Tool check and pairing. Name the task you're bringing.
9:20 – 10:15	Lab 1: Build a Prompt That Holds	Rebuild a weak prompt from your own task using task, context, constraints, and output format. Work the four failure modes. Check for text carried invisibly when you paste from a document. A live A/B on two widely repeated techniques, so you stop using the ones the evidence no longer supports.
10:15 – 10:30	Break	
10:30 – 11:20	Lab 2: The Line You Cannot Cross	Sort your own working material against a confidentiality decision rule, by descriptor, never by opening the document. Learn to tell a binding obligation from a recommendation off the verb. The one federal prohibition that is absolute rather than advisory. Draft a disclosure sentence and find the placement rule for the journal or funder you're actually aiming at.
11:20 – 11:30	Break	
11:30 – 12:25	Lab 3: Verify It Like a Reviewer Would	Why models compose sources rather than retrieve them, and why a working link is not verification. Verify a set of AI-produced references containing planted fabrications. Build a verification record you can reuse.
12:25 – 1:10	Lunch	
1:10 – 2:00	Lab 4: Synthesis Without the Drift	Run a real literature- or evidence-synthesis task through the checklist live. Find where nuance gets flattened. Work out what a position statement actually asks of you, where the declare / don't-declare line sits, and where preclinical practice differs.
2:00 – 2:50	Lab 5: The Work Your Colleagues See	Bring one of your own outputs (a poster panel, a slide, a manuscript or nonclinical section) and apply the standard: message hierarchy, cutting to the finding, figure integrity, disclosure, and close human supervision with a quality review step.
2:50 – 3:05	Break	
3:05 – 3:45	Lab 6: Bring Your Own Problem	The core of the day. Two problems from the room get worked end to end on screen; then you work yours. Bring the one you didn't think AI could help with.
3:45 – 4:00	Close: Your Guardrails, Written Down	Prompt libraries, where agents break, and the guardrail rule. Walk out with your confidentiality decision rule, your verification record, and one guardrailed workflow.

The labs draw on the ten-session Defensible AI webinar series. Prior attendance is not assumed; the workshop stands alone.

Presenter

Szczepan Baran, MS, VMD

Chief Scientific Officer, Instem

Co-founder, Digital Preclinical Society