# PHARMA-AI Ethical Guardrail Builder

## Documentation & User Guide

Release Date: October 2026 | Application Version: 1.5.0 | Methodology: PHARMA Rules 1.2
Platform: Universal HTML5 / JavaScript - Client-Side Decision Support

**Central principle**

The learning outcome should determine the role of AI. The application helps educators make that reasoning visible, consistent and documentable; it does not replace institutional authority or professional judgment.

## 1. Introduction

The PHARMA-AI Ethical Guardrail Builder is a browser-based decision-support application developed for pharmacy educators, preceptors, assessment leaders and academic governance teams. It converts ethical principles into a structured review of a proposed educational use of artificial intelligence and generates practical, editable policy language.

Unlike a static checklist, the application recalculates its findings whenever the user changes the educational setting, learning outcome, stakes, AI function, data classification or safeguard. This allows users to test alternative designs and understand why a recommendation changes.

### Key features

1. **Dynamic PHARMA risk engine:** Assesses Privacy, Human judgment, Academic integrity, Reliability, Maintenance of competence, and Access and equity.
2. **Quick and Full review modes:** Supports an 11-minute workshop activity as well as detailed institutional review.
3. **Critical override rules:** Prevents serious incompatibilities from being hidden by a reassuring average score.
4. **Live decision trace:** Explains which inputs raised or reduced each ethical-risk dimension, with change markers on the live bars.
5. **Scenario exploration:** Includes 12 preloaded pharmacy-education scenarios and a visual side-by-side comparison.
6. **Policy Studio:** Produces four editable outputs for faculty, students and experiential settings, with a formatted document preview.
7. **Local profiles and outcomes:** Saves work in the browser and supports portable JSON backup and restoration.
8. **Professional reporting:** Provides copyable language, A4 print layouts, browser-based PDF saving and a downloadable standalone report.
9. **Visual results:** A Guardrail gauge, PHARMA hexagon and level ladder show the decision at a glance.
10. **Presenter mode:** Larger text and stronger contrast for projectors; press P to switch it on or off.

**Professional review required**

The Guardrail Index is a structured decision-support indicator, not a psychometrically validated risk instrument. The application does not provide legal, regulatory, accreditation or institutional approval.

## What's new in versions 1.2 to 1.5

Versions 1.2 to 1.4 changed how results are shown and how your work is protected. Version 1.5 introduces PHARMA Rules 1.2: six new override rules and policy language that follows the decision. Every score is unchanged; the 12 workshop scenarios keep their scores, and two of them (s4 and s12) now receive a stricter decision.

1. **Safer policy language:** Permitted uses now follow the decision level, never the proposal. A Do not proceed decision produces a Not for release notice and a No-AI declaration, and clinical reasoning, care plans, dose calculations and decision recommendations are never permitted on real patients.
2. **New override rules:** Unsupervised clinical decision support, AI marking without human review, health data outside an approved environment, outsourced reflection and unverified AI references now set a minimum decision.
3. **Guidance sources:** Every override shows its basis, and the Methodology view, reports and section 16 list the WHO, ASHP, FIP, UNESCO, EU, UAE, DOH, ICMJE and COPE sources in full.
4. **Clearer results:** A Guardrail gauge with the five decision bands, a PHARMA hexagon comparing risk before and after safeguards, and a Level 0–4 ladder.
5. **Severity colours:** Bars and decision panels run from green (lower risk) to red (critical); each dimension keeps its own letter colour.
6. **Guided Builder:** A clickable step bar that ticks off completed steps, icons on key choices, and +/− markers showing which bars moved after each answer.
7. **Visual Compare:** Overlaid hexagons, paired dimension bars and a table listing only the answers that differ between two reviews.
8. **Policy preview:** Policy Studio shows each output as a formatted document, with an Edit text view for changes.
9. **Protected work:** The app asks before replacing edited policy wording or an unsaved draft.
10. **Reports:** A Download report button saves a standalone HTML report; printed and downloaded reports include the gauge and hexagon.
11. **Presenter mode:** Press P for larger, higher-contrast text on projectors.
12. **Tablets and projectors:** Navigation now works at every screen width, including 1024-pixel projectors and tablets in landscape.

## 2. The PHARMA Ethical AI Lens

The application organizes ethical review around six pharmacy-specific domains. Each domain is scored separately so that a low concern in one area cannot obscure a major concern in another.

| Letter | Domain | Guiding question |
|---|---|---|
| P | Privacy and confidentiality | Are patient, student, research or institutional data being exposed? |
| H | Human judgment and accountability | Who remains responsible when the AI output is wrong, unsafe or misleading? |
| A | Academic integrity and attribution | Is AI use permitted, transparent and appropriately acknowledged? |
| R | Reliability, bias and verification | Could the output be inaccurate, fabricated, outdated, incomplete or biased? |
| M | Maintenance of competence | Is AI supporting development, or performing the learner's thinking? |
| A | Access and equity | Can all learners participate fairly, safely and accessibly? |

### The three cognitive risks

1. **Deskilling:** A previously developed ability weakens through repeated dependence on AI.
2. **Never-skilling:** Learners outsource cognitive work before the underlying competence is developed.
3. **Mis-skilling:** Learners internalize and repeat inaccurate or biased AI-generated reasoning.

**Recommended sequence**

Human first -> AI second -> Human verification

## 3. Decision Methodology

### The Guardrail Index

Each response contributes to one or more PHARMA dimensions. The engine combines baseline risk, contextual modifiers, selected AI functions and safeguards, then normalizes each dimension and the overall Guardrail Index to a 0-100 scale.

The overall score is intentionally influenced by both the average risk and the highest individual PHARMA risk. Critical override rules can set a minimum decision level when an incompatibility should not be averaged away.

| Index | Decision category | Operational meaning |
|---|---|---|
| 0-19 | Suitable for purposeful integration | The proposed use aligns with the learning purpose under standard safeguards. |
| 20-34 | Permitted with documented safeguards | Use is reasonable if the specified controls are implemented and monitored. |
| 35-54 | Restricted use only | Only selected AI functions should be permitted under documented limitations. |
| 55-74 | Redesign and approval required | Major risks remain and require redesign plus appropriate institutional review. |
| 75-100 or critical override | Do not proceed in the proposed form | A critical incompatibility must be removed before proceeding. |

In the app, these five bands are the colours of the Guardrail gauge and of every severity bar: green, teal, amber, coral and red.

### Suggested AI involvement levels

| Level | Label | Interpretation |
|---|---|---|
| Level 0 | No AI | AI use is incompatible with the proposed activity in its current form. |
| Level 1 | Mechanical support only | Limit use to low-cognition administrative or language functions. |
| Level 2 | Limited assistive use | Permit defined functions while protecting substantive independent work. |
| Level 3 | Declared and supervised collaboration | Use AI transparently with human-first work and verification. |
| Level 4 | Purposefully AI-integrated learning | AI use is itself part of the learning outcome and is intentionally taught. |

### Examples of critical override rules

1. Identifiable patient or sensitive information paired with a public or unevaluated AI environment.
2. Patient-facing or progression decisions without mandatory qualified human review.
3. AI performing the same substantive competence a high-stakes assessment is intended to judge, without independent evidence of performance.
4. Required paid AI use without an equivalent alternative, triggering an equity redesign requirement.
5. New in Rules 1.2: AI decision recommendations or dose calculations in patient-facing or experiential work without mandatory review, independent verification and independent demonstration (basis: ASHP 2025; FIP 2025; WHO 2024). With all three safeguards, the same uses still require redesign and approval.
6. New in Rules 1.2: AI marking or grading without mandatory human review (basis: EU AI Act 2024).
7. New in Rules 1.2: identifiable or sensitive health data in a personal paid AI subscription, or in a general enterprise tool without explicit approval (basis: UAE Federal Law No. 2 of 2019; DOH Abu Dhabi 2018).
8. New in Rules 1.2: AI drafting of reflective writing without independent demonstration (basis: UNESCO 2023).
9. New in Rules 1.2: AI-generated references without verification in graded, high-stakes or published work (basis: ICMJE; COPE 2023).

Every override sets a minimum decision without changing the numerical score, and every finding shows its guidance basis. Section 16 lists the sources in full.

### How permitted AI uses are decided

Policy language permits a proposed AI function only when the decision level allows it. Anything above that level is listed by name under Not permitted.

1. **Level 0:** no AI use for any part of the activity.
2. **Level 1:** grammar refinement, and translation of non-clinical text that is checked.
3. **Level 2:** adds brainstorming, explanation or tutoring, summaries of sources the learner has read, evidence retrieval with verification, and feedback on the learner's own draft.
4. **Levels 3 and 4:** add drafting after a human-first draft, synthetic patient simulation, faculty-validated question generation and de-identified data analysis. Clinical reasoning, care plans, dose calculations and decision recommendations are allowed only to check the learner's own completed work on synthetic cases.
5. **Never for learners:** AI-generated citations, AI marking, and any clinical function on real patients or real patient data.

## 4. Getting Started

### Access and browser requirements

1. Open the live application in a current version of Chrome, Edge, Safari or Firefox.

2. Use a laptop or tablet for the most detailed review. Mobile phones are fully supported for Quick Mode and scenario exploration.
3. Do not enter patient-identifiable, confidential, unpublished or otherwise sensitive information.
4. Choose Quick Mode for workshop use or Full Review for institutional documentation.

### Home screen

1. **Start a quick review:** Opens the five-step Builder with only high-impact decisions visible.
2. **Open full institutional review:** Adds activity identity, purpose, review date and extended safeguards.
3. **Explore scenarios:** Opens preconfigured cases that demonstrate contextual differences.
4. **Saved work:** Displays locally stored assessments and institutional profiles.

**Local data notice**

Saved work belongs to the current browser and website origin. Export a JSON backup before changing devices, clearing browser data or moving the application to another domain.

## 5. Quick Workshop Mode

Quick Mode is designed for rapid group use during the workshop. It retains the complete decision engine while reducing the visible inputs to the factors most likely to change the ethical recommendation.

| Step | Focus | What the user decides |
|---|---|---|
| 1 | Purpose and setting | Select the educational setting, stakes and intended learning outcomes. |
| 2 | Proposed role of AI | Identify every function AI will perform. |
| 3 | Data and environment | Classify the information involved and the type of AI environment. |
| 4 | Human judgment and competence | Specify human review, human-first reasoning, independent demonstration and verification. |
| 5 | Integrity, disclosure and equity | Clarify AI-use rules, disclosure and equivalent access. |

The live panel recalculates after every selection. Bars that moved show the size and direction of the change (for example −18 ▼ in green), and the panel's colour follows the decision category. The step bar above the questions ticks off completed steps; click any step to jump to it. On phones and tablets, tap the floating score button to open the same information in a touch-friendly drawer.

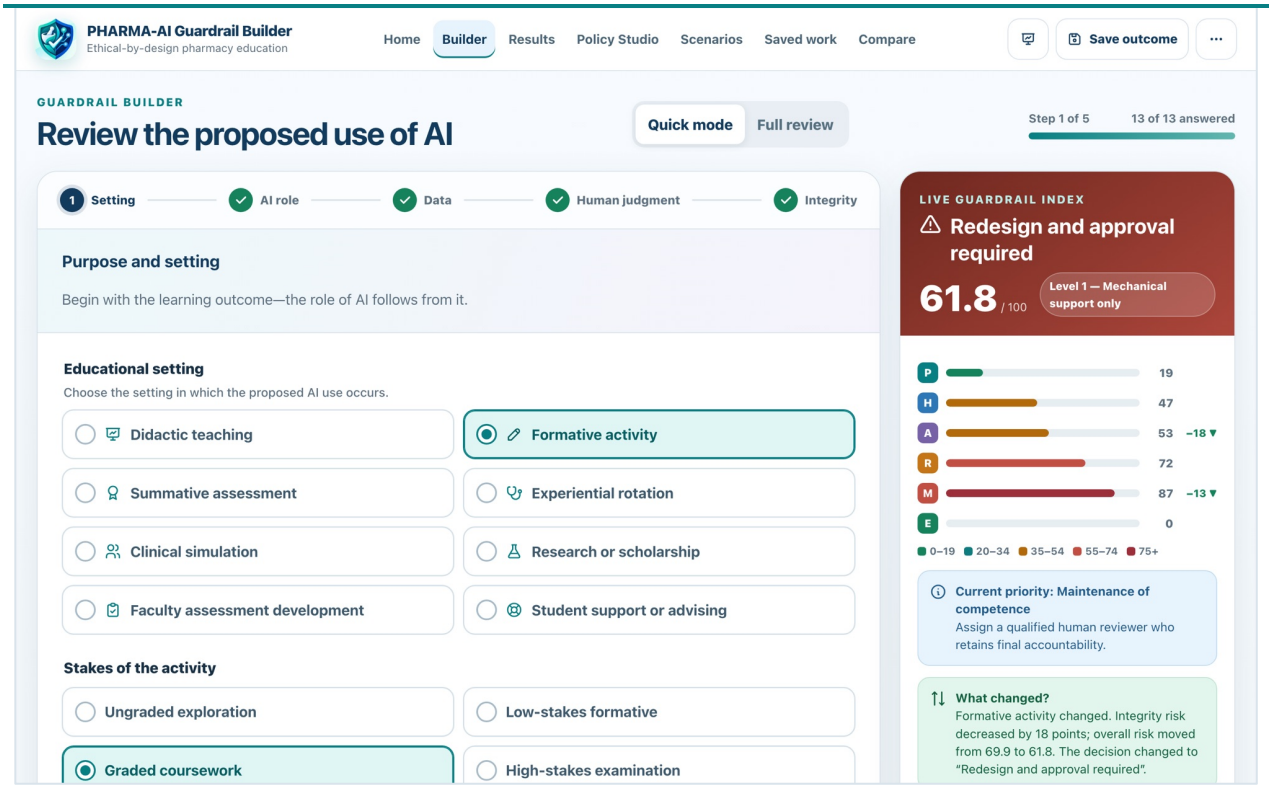



**Figure 1. The Builder in Quick Mode: step bar, choice icons and the live Guardrail Index with change markers (AI-assisted pharmaceutical care plan, setting changed to Formative activity).**

### Presenter mode

For workshops, press P (or use the presenter button in the header, or Presenter mode in the More menu) to enlarge the text and strengthen contrast for projectors. Press P again to return to the normal view. The setting is remembered in that browser.

## 6. Full Institutional Review

Full Review preserves the same five-stage structure but adds fields needed for governance, record keeping and later reassessment.

1. **Activity identity:** Assessment title, institution, course or rotation, faculty owner and review date.
2. **Educational purpose:** A brief activity description that must exclude names or sensitive information.
3. **Required use:** Whether the AI tool is mandatory, optional or not yet decided.
4. **Oral defence:** Whether learners must explain and defend the submitted work.
5. **Supporting controls:** Prompt retention, user training, accessibility arrangements and expert review.

### Good review practice

1. Assess the activity design, not an individual student's conduct.
2. Select the most sensitive information that could realistically be involved.
3. Describe what AI will do, not merely the name of the tool.
4. Treat 'not confirmed' as an unresolved control requiring follow-up.
5. Save a completed assessment only after reviewing the generated decision trace.

## 7. Understanding Live Results and the Decision Trace

### Live results panel

1. **Decision category:** The current operational recommendation.

2. **Guardrail Index:** A 0-100 comparison indicator based on the current responses.
3. **Suggested involvement:** The recommended Level 0-4 role for AI.
4. **PHARMA bars:** Separate scores for all six ethical domains. Bar colour shows severity (0–19 green, 20–34 teal, 35–54 amber, 55–74 coral, 75+ red); the letter badge keeps each domain's own colour.
5. **Priority alert:** The most important hard stop or unresolved safeguard.
6. **What changed?:** Explains how the most recent answer changed the risk profile. The box turns green when risk falls and coral when it rises, and the bars that moved show +/− markers.

## Full Results view

The Results view expands the live summary into these components:

1. Executive decision, with any hard-stop findings listed underneath.
2. Guardrail gauge: the needle shows the score and the highlighted band shows the decision. When a critical or approval override applies, a note explains why the decision is stricter than the score.
3. PHARMA hexagon: the dashed shape is risk before safeguards and the filled shape is risk after; the numbers read after / before.
4. Suggested AI involvement ladder from Level 0 (no AI) to Level 4 (AI-integrated learning).
5. Dimension-by-dimension profile: hatched bars show risk before safeguards and solid bars show risk after, so the gap is what the safeguards removed.
6. Critical findings, each with its guidance basis, and required safeguards.
7. Transparent decision trace: the strongest rules, each with a bar pointing right (adds risk) or left (a safeguard reduces risk).

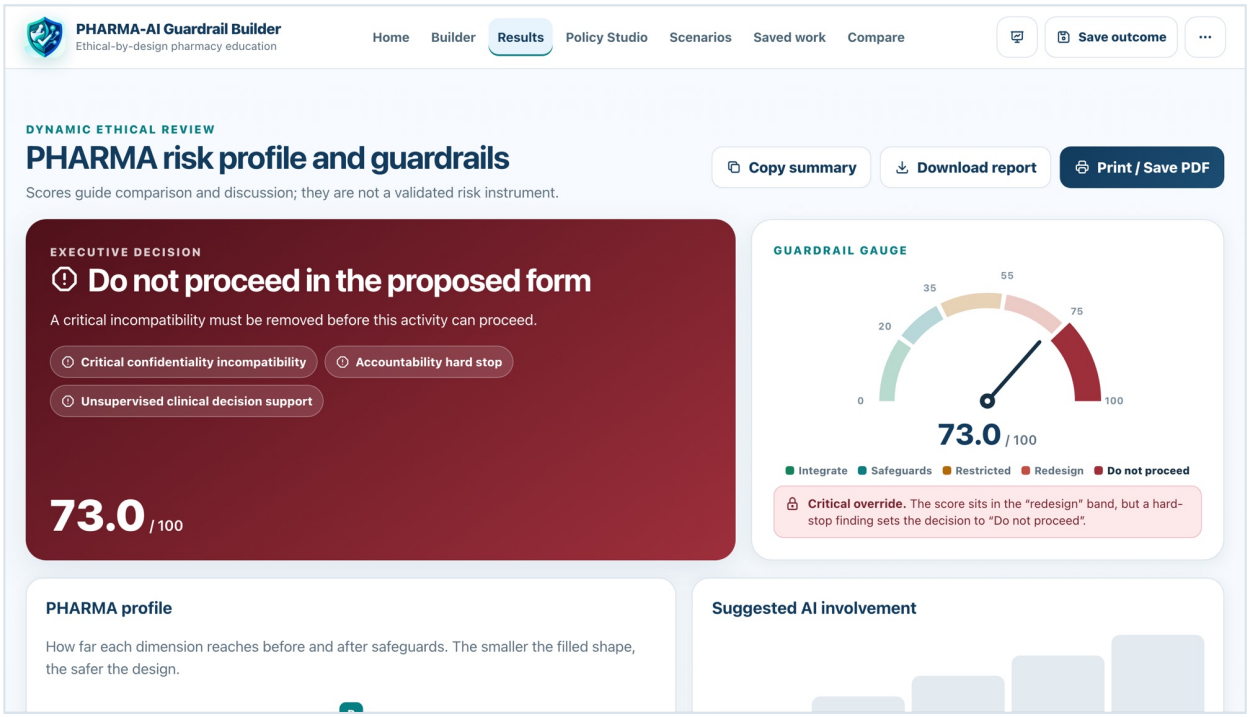



**Figure 2. Executive decision and Guardrail gauge for "Real patient case in a public chatbot". The score (73.0) sits in the Redesign band, but a critical override sets the decision to Do not proceed.**

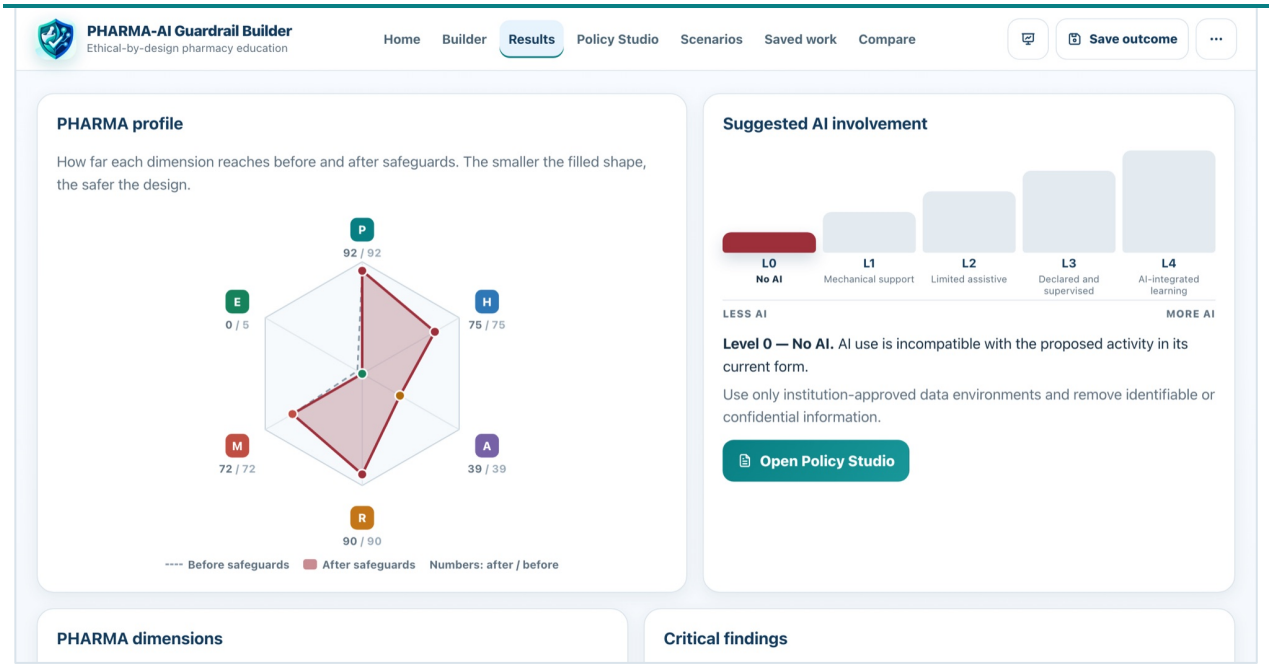



**Figure 3. PHARMA hexagon (before and after safeguards) and the suggested AI involvement ladder for the same scenario.**

### Interpretation rule

A lower overall score does not authorize use when a critical incompatibility remains. Read the critical findings and required safeguards before relying on the category label.

## 8. Scenario Library and Comparison

The Scenario Library contains 12 examples spanning didactic, assessment, experiential, research and faculty-development settings. Each scenario can be loaded into the Builder and modified.

| Scenario | Setting | Primary issue demonstrated |
|---|---|---|
| Real patient case in a public chatbot | Experiential | Confidentiality and patient-data exposure |
| Synthetic case with supervised AI critique | Formative | Purposeful supervised integration |
| AI-assisted pharmaceutical care plan | Summative | Integrity and competence replacement |
| Forty overnight AI-generated MCQs | Faculty | Unvalidated keys, references and bias |
| Biased patient-case generation | Didactic | Representational bias |
| Confidential manuscript in public AI | Research | Confidentiality and intellectual property |
| Required premium AI platform | Graded | Access and equity |
| Grammar refinement after a first draft | Summative | Limited assistive use |
| Formative drug-information exercise | Formative | Evidence verification |
| AI patient-counselling simulation | Simulation | Supervised communication practice |
| Faculty-reviewed OSCE case generation | Faculty | Expert validation |
| Reflective writing outsourced to AI | Experiential | Authenticity and reflective competence |

### Using Compare

1. Open Compare from the navigation, the More menu or a scenario card.
2. Choose two preloaded or saved assessments as A (blue) and B (purple).
3. Compare the summary cards and the overlaid PHARMA hexagons: A is drawn with a solid line and B with a dashed line.

4. Use Dimension by dimension to see how far B sits above or below A, then read What is different between A and B, which lists only the answers that differ.
5. Identify the input or safeguard responsible for the difference, then load either scenario and change that one answer to test it live.

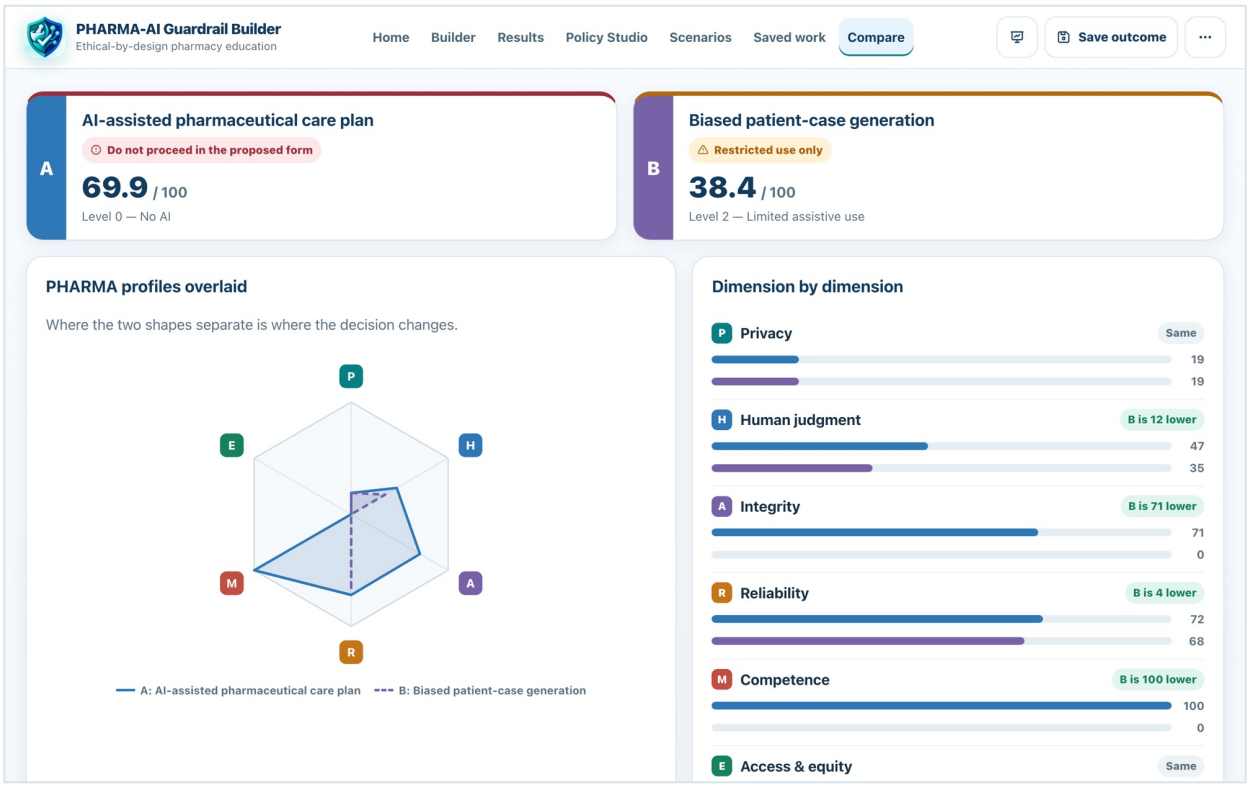



Figure 4. Compare view: AI-assisted pharmaceutical care plan (A) against Biased patient-case generation (B).

## 9. Policy Studio

Policy Studio converts the current ethical decision into editable language using deterministic templates. It does not send prompts or information to an external AI service.

| Generated output | Purpose |
|---|---|
| Faculty decision statement | Documents the decision, suggested AI involvement, principal concerns, safeguards and retained accountability. |
| Student-facing syllabus clause | States permitted and prohibited functions, required process, disclosure, privacy and consequences. |
| Experiential-rotation clause | Appears when the setting is experiential, simulated or patient-facing and emphasizes confidentiality and preceptor review. |
| Student AI-use disclosure statement | Provides a structured record of the tool, purpose, verification and learner responsibility. |

The language always follows the decision:

1. **Do not proceed:** the faculty statement lists the findings and redesign actions; the student clause states that AI is not permitted and is marked Not for release; the disclosure becomes a Student no-AI declaration.
2. **Redesign and approval required:** the student clause is marked Draft pending institutional approval and permits only Level 1 uses until approval.
3. **Faculty development and research:** the student clause becomes a Faculty and researcher AI-use statement, and the disclosure becomes a staff disclosure in line with ICMJE and COPE.

4. **Clinical activities:** every clause adds clinical safety rules: never act on AI output for a real patient, calculate every dose independently, and keep patient data out of tools not approved for health data in the UAE.

## Editing and copying

1. Select an output from the Policy Studio menu. It opens in Preview, laid out as a document with the logo, activity title and date.
2. Choose Edit text to change the wording, then Preview to check the result. Changes are kept in the current draft.
3. Use Copy this section for one output or Copy all outputs for a complete policy pack.
4. Use Regenerate from current answers to discard manual edits and rebuild the language.
5. If you change an answer in the Builder after editing, the app asks first. Choose Keep my wording to undo the answer change and keep your edits.

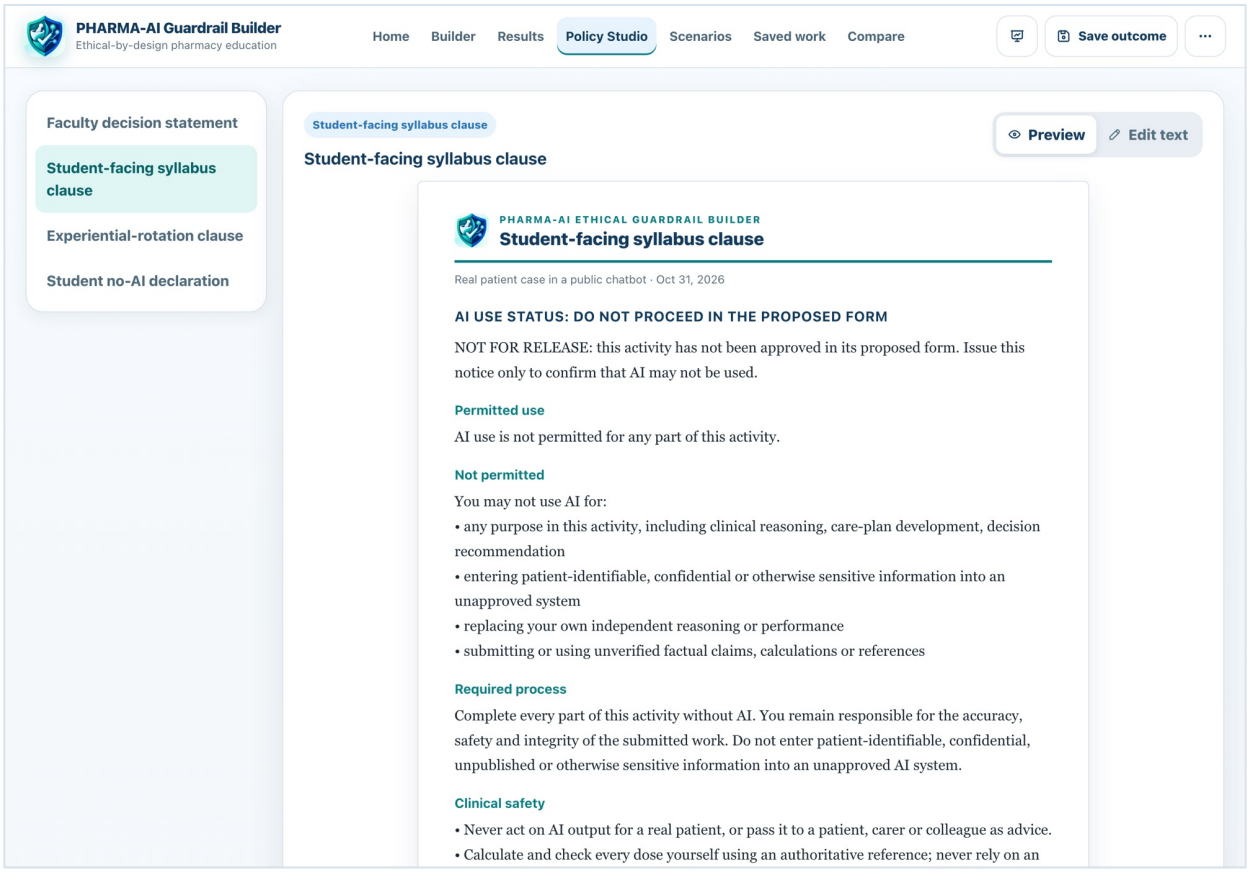



**Figure 5. Policy Studio preview for "Real patient case in a public chatbot": AI use is not permitted, the notice is marked Not for release, and clinical safety rules are added.**

**Important**

Generated language remains a draft. Review terminology, institutional policy references, misconduct procedures and local data-protection requirements before adoption.

# 10. Saving Profiles, Outcomes and Backups

## Saved assessments

Save outcome stores the current responses, PHARMA scores, decision, generated language, application version and methodology version in browser storage. Saved assessments can be opened, duplicated or deleted.

When the current draft has unsaved changes, loading a scenario, starting a new review, or opening or duplicating a saved assessment asks for confirmation first. Choose Keep my draft to stay where you are.

## Institutional profiles

An institutional profile stores reusable defaults such as institution, course or program, faculty owner, approved AI environment, disclosure expectations, equivalent-access arrangements and supporting controls.

## Backup and restoration

1. Open Saved work.
2. Choose Export backup to download a JSON file containing profiles, assessments and the current draft.
3. Store the file in an institutionally appropriate location.
4. Choose Import backup on another browser or device. Importing replaces the saved work in that browser after confirmation.

**Do not use browser storage as a clinical repository**

Browser storage is convenient local persistence, not an approved location for patient-identifiable, confidential or sensitive information.

# 11. Copying, Printing and PDF Reports

## Available report formats

1. Executive decision and PHARMA risk profile.
2. Detailed ethical audit with critical findings, safeguards and methodology.
3. Policy and syllabus pack containing the applicable generated statements.
4. Standalone HTML report, including the logo, Guardrail gauge, PHARMA hexagon and any edited policy language.

## Print or save as PDF

1. Open Results or Policy Studio.
2. Select Print / Save PDF or Print policy pack.
3. Choose A4 or Letter in the browser print dialog if required.
4. Select Save as PDF or a physical printer.
5. Review the PDF before distribution, particularly after extensive manual edits.

The print layout automatically hides navigation, touch controls and accessibility skip links. Reports remain readable in colour or grayscale and preserve genuine paragraph breaks in all generated statements. Printed and downloaded reports also include the Guardrail gauge and PHARMA hexagon.

### Standalone report export

Choose Download report on the Results view, or in the More menu, to save a self-contained HTML report of the current assessment. It includes the logo, gauge, hexagon and any edited policy language, and can be opened or printed without the app.

## 12. Mobile and Tablet Operation

The application is designed for touch use and adapts its information architecture rather than simply shrinking the desktop interface.

| Device or orientation | Interface behaviour |
|---|---|
| Mobile portrait | Single-column Builder, bottom navigation and full-height live-results drawer. |
| Mobile landscape | Compact single-column or split view where screen width permits. |
| Tablet portrait | Questionnaire with a floating score button that opens the live results, plus bottom navigation. |
| Tablet landscape | Side-by-side questionnaire and live PHARMA results, with bottom navigation up to 1120 pixels wide (including 1024-pixel projectors). |
| Desktop | Full two-panel workspace with persistent navigation and insights. |

### Touch and accessibility features

1. Large selection cards and touch targets of approximately 44-48 pixels.
2. Sticky Back, Next and Results actions during an assessment.
3. Readable body text and responsive headings without horizontal page scrolling.
4. Keyboard navigation, visible focus indicators and screen-reader labels.
5. Text, icons and labels used alongside colour for risk status.
6. Reduced-motion support and no essential information conveyed only through animation.
7. Confirmation dialogs keep keyboard focus inside until closed; Escape cancels.
8. Presenter mode (press P) gives larger, higher-contrast text on projectors.

## 13. Privacy, Security and Data Handling

### Privacy-by-design architecture

1. **No external processing:** All calculations and policy generation occur within the browser.
2. **No external AI API:** The application does not send assessment content to a language model.
3. **No analytics by default:** The self-contained application includes no tracking service.
4. **Local persistence only:** Drafts, profiles and outcomes use browser storage.
5. **Portable backup:** Users control transfer through explicit JSON export and import.
6. **Safe text rendering:** User-entered content is treated as text rather than executable HTML.

### Data-entry rule

Describe the educational activity without entering names, medical-record numbers, dates of birth, case identifiers, unpublished results or any other information that could identify a patient, student, colleague or confidential project.

**The first ethical guardrail**

A governance tool does not need sensitive data to evaluate whether sensitive data would be involved.

## 14. Troubleshooting

| Issue | Likely cause | Recommended action |
|---|---|---|
| The live app looks unchanged after an update | The browser may be using a cached copy. | Refresh the page. If needed, use a hard refresh and reopen the live link. |
| Saved work is missing on another device | Browser storage is local to a device, browser and domain. | Export a JSON backup from the original browser and import it on the new device. |
| Copy button does not work | Clipboard access may be restricted outside a secure connection. | Use the live HTTPS site, or select and copy the editable text manually. |
| Print dialog does not open | Pop-ups or printing may be blocked by the browser. | Allow printing, then retry from Results or Policy Studio. |
| PDF contains interactive controls | The page was printed without using the app's report action or from an old cached version. | Refresh the app and print through Print / Save PDF. |
| Literal backslash-n text appears | An older application version or old manually edited statement is open. | Refresh to the current version. Previously saved text is cleaned automatically when reopened. |
| Risk remains high after adding disclosure | Disclosure does not resolve privacy, safety or competence incompatibilities. | Review Critical findings and address the specific hard stop or required safeguard. |
| Imported backup is rejected | The file is not a valid PHARMA-AI backup or exceeds the size limit. | Use an unmodified JSON backup exported by the application and keep it below 5 MB. |
| Text is hard to read on the projector | The standard view is designed for laptops. | Press P to switch on Presenter mode, and press P again to switch it off. |
| Edited policy wording was replaced | An answer was changed and the replacement was confirmed. | When asked, choose Keep my wording. Copy edited text before changing answers you are unsure about. |
| The decision is stricter than the score suggests | A critical or approval override applies; the finding shows its guidance basis. | Read the override note under the Guardrail gauge and the hard-stop findings on the decision card. |

## 15. Technical Specifications and Deployment

| Specification | Implementation |
|---|---|
| Architecture | Single-page, self-contained browser application. |
| Core technologies | HTML5, embedded CSS3 and vanilla JavaScript. |
| Application size | Approximately 210 KB including the embedded logo and icons; no external runtime dependencies. |
| Processing | Deterministic client-side rules engine; no server calculation. |
| Persistence | Browser localStorage with JSON backup and restore. |
| Reporting | Clipboard API with fallback, Blob downloads (backup and standalone HTML report) and browser printing. |
| Responsive design | Mobile, tablet and desktop layouts with touch and keyboard support. |
| Hosting | Compatible with Cloudflare Pages or any static HTTPS hosting service. |
| Network behaviour | No application data transmission after page load. |
| Current version | 1.5.0, October 2026 (PHARMA Rules 1.2). |

### Cloudflare Pages deployment

1. Save the self-contained application as index.html, or retain the supplied filename and route to it.
2. Upload the file through Cloudflare Pages Direct Upload or connect a Git repository.
3. No build command, server function, database or environment variable is required for the standalone edition.
4. Use HTTPS so clipboard operations and browser security controls function consistently.
5. After changing domains, remind users that saved browser data does not automatically move to the new origin.

### Release notes

1. **1.5.0 (PHARMA Rules 1.2):** Six new override rules with guidance basis; policy language that follows the decision level; Not for release notices, No-AI declarations and staff statements; clinical safety rules; guidance sources in the app, reports and guide. Scores unchanged; s4 and s12 now Do not proceed.
2. **1.4.0:** Clickable step bar, icons on setting, information and environment choices, change markers on the live bars, a visual Compare view and a formatted Policy Studio preview.
3. **1.3.0:** Guardrail gauge with decision bands and override note, PHARMA hexagon, level ladder, before/after dimension bars, direction bars in the decision trace, and visuals in printed and downloaded reports.
4. **1.2.0:** Navigation fixed for tablets and 1024-pixel projectors; warnings before edited policy wording or unsaved drafts are replaced; Download report button; logo, icons and severity colours; in-app confirmation dialogs; Presenter mode.
5. **1.0.1:** Removed the skip link from printed reports, converted statement separators to genuine paragraph breaks, and added automatic cleanup of literal backslash-n characters in saved statements.

## 16. Methodological Limitations and Governance

The application structures educational judgment but cannot determine whether an activity complies with every institutional policy, law, regulator, accreditation standard or contractual tool term. Those requirements vary by jurisdiction, institution and time.

The scores are designed for consistent comparison and explanation, not for psychometric interpretation. Institutions adopting the tool should review the rule set, document any local modifications, identify accountable decision makers and schedule periodic review.

### Recommended governance use

1. Use the application before approving an AI-enabled activity, not only after an incident.
2. Attach the report to course, assessment or experiential-education documentation where appropriate.
3. Record the responsible reviewer and the institutional profile used.
4. Reassess when the AI tool, data environment, assessment stakes or learning outcome changes.
5. Investigate critical findings through the relevant privacy, assessment, clinical or research governance route.

### Guidance sources

The override rules, the permitted-use levels and the generated language follow the guidance below. The same list appears in the Methodology view and in every printed or downloaded report. The application remains decision support; check outputs against institutional policy and current regulations.

1. **WHO 2024.** World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2024. https://www.who.int/publications/i/item/9789240084759
2. **ASHP 2025.** American Society of Health-System Pharmacists. ASHP statement on artificial intelligence in pharmacy. Am J Health-Syst Pharm. 2025;82(19):e853. https://academic.oup.com/ajhp/article/82/19/e853/8240129
3. **ASHP Policy 2413.** American Society of Health-System Pharmacists. Policy position 2413: Role of artificial intelligence in pharmacy practice. 2024. https://www.ashp.org/Pharmacy-Practice/policy-positions-database/2024/2413
4. **FIP 2025.** International Pharmaceutical Federation (FIP). FIP statement of policy: Artificial intelligence in pharmacy practice. The Hague: FIP; 2025. https://www.fip.org/file/6354
5. **UNESCO 2023.** Miao F, Holmes W. Guidance for generative AI in education and research. Paris: UNESCO; 2023. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

6. **EU AI Act 2024.** Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III point 3(b) and Article 14. Used as an international benchmark; not binding in the UAE. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
7. **UAE Federal Law 2/2019.** United Arab Emirates. Federal Law No. 2 of 2019 concerning the use of information and communications technology in health fields. https://uaelegislation.gov.ae/en/legislations/1209
8. **DOH Abu Dhabi 2018.** Department of Health Abu Dhabi. Policy on use of artificial intelligence (AI) in the healthcare sector of the Emirate of Abu Dhabi (Policy/AI/0.9). 2018. https://www.doh.gov.ae/-/media/E9C1470A575146B18015DEBE57E47F8D.ashx
9. **ICMJE.** International Committee of Medical Journal Editors. Recommendations, section V: Use of artificial intelligence in publishing. https://www.icmje.org/recommendations/browse/artificial-intelligence/
10. **COPE 2023.** Committee on Publication Ethics. COPE position: Authorship and AI tools. 2023. doi:10.24318/cCVRZBms. https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools

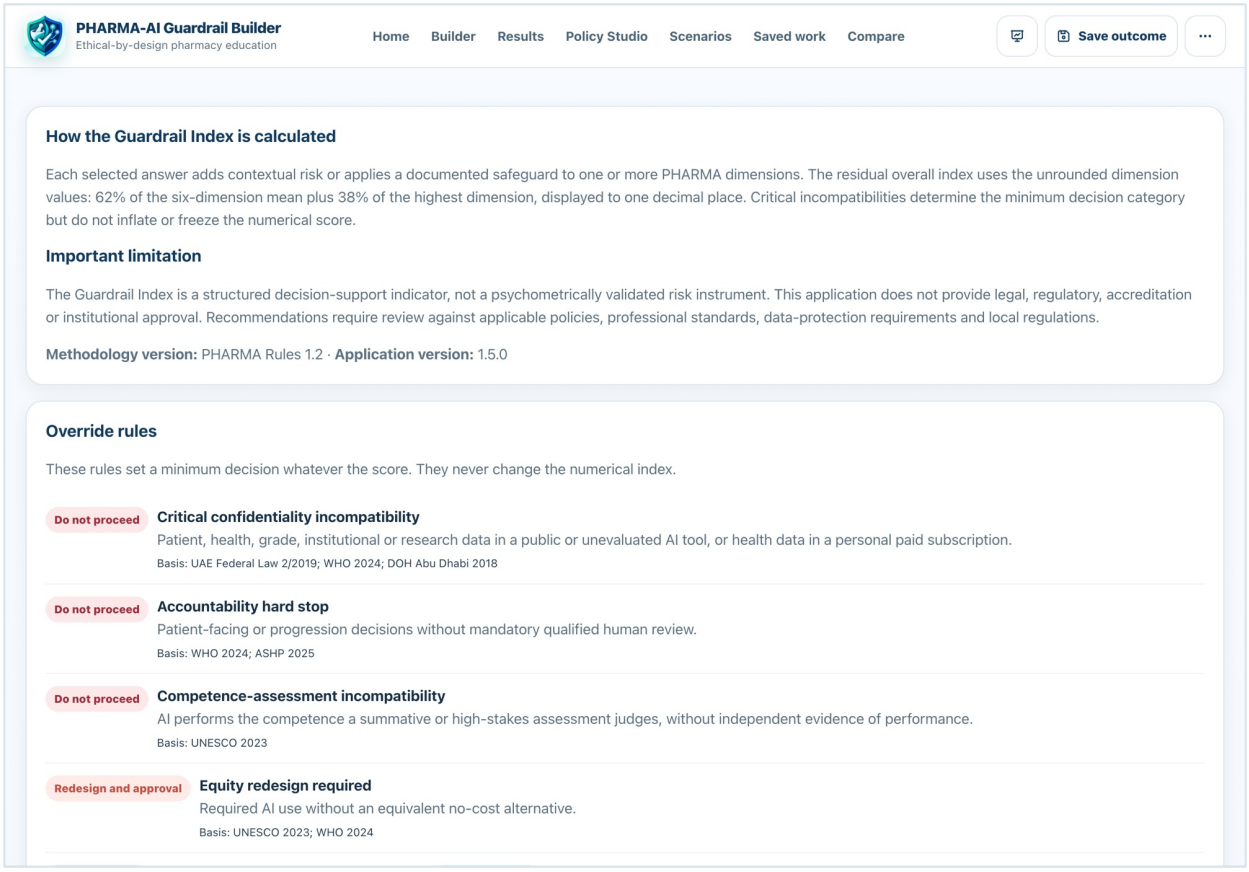



**Figure 6. Methodology view: versions, and override rules each showing its guidance basis. The permitted-use levels and the full list of guidance sources follow below.**

## 17. Glossary

1. **AI environment:** The technical and contractual setting in which an AI system processes information.
2. **Critical override:** A rule that prevents a serious incompatibility from being reduced by averaging or minor safeguards.
3. **Guardrail gauge:** The semicircular dial on the Results view showing the Guardrail Index against the five decision bands.
4. **Guardrail Index:** The 0-100 decision-support indicator generated from the current responses.
5. **Hard stop:** A condition requiring removal or redesign before the activity proceeds.
6. **Human-first workflow:** Learners attempt or document reasoning before consulting AI.
7. **Inherent risk:** Risk present before safeguards are considered.

8. **No-AI declaration:** The student form generated when a decision does not permit AI, confirming that no AI tool was used.
9. **Override rule:** A rule that sets a minimum decision whatever the score, shown with its guidance basis.
10. **PHARMA hexagon:** A six-axis chart of the PHARMA dimensions; the dashed shape is inherent risk and the filled shape is residual risk.
11. **Presenter mode:** A projector-friendly view with larger text and stronger contrast, switched on and off with P.
12. **Residual risk:** Risk remaining after selected safeguards are applied.
13. **Synthetic information:** Invented data that do not represent a real identifiable person or confidential record.

## 18. Development and Attribution

1. Concept and ethical logic: Prof. Muhammad Al-Shorbagy, Dean, College of Pharmacy, Gulf Medical University.
2. Development approach: AI-assisted VIBE coding, with faculty ownership of the educational and ethical decision logic.
3. Application family: Acabuild AI LabHub - applications by academicians for academicians.
4. Methodology: PHARMA Ethical AI Lens, Rules Version 1.2.

***Innovation is not measured by how much AI is inserted into a process, but by how appropriately it is used - or intentionally not used.***