



AI Skills Programme Guide

Practical Intelligence for the Third Epoch

Institute for Applied Intelligence — Skills Division
Fitzherbert University · 2025

ABSTRACT

The definitive guide to the Fitzherbert AI Skills Programme, open to all students regardless of college affiliation. Level I (Foundation) covers prompt engineering and output evaluation, on the basis that being unable to construct a coherent prompt is the Epoch III equivalent of being unable to read. Level II (Practitioner) covers RAG architecture and agent orchestration. Level III (Specialist) covers multi-agent systems and deployment pipelines. Level IV (Sovereign) was added in 2025 after it became apparent that Level III graduates were being hired directly by AI governance bodies without completing the institutional ethics component of their education. Level IV cannot be completed without demonstrating, in a live examination before three human examiners, that the student can articulate a coherent argument for why human governance of AI systems is preferable to the alternative. The examiners are not required to find the argument convincing. They are required to find it human.

Classification: Public — Academic Resource

Edition: 6 page institutional edition (approx.)

Published: 2025

Archive: Fitzherbert University Institutional Repository

This document is published by Fitzherbert University in accordance with the Transparency Mandate of 2003.

Table of Contents

I. Programme Rationale and Educational Doctrine	3
The AI Skills Programme was established on the premise that literacy in the Third Epoch includes the disciplined use, evaluation, and governance of artificial intelligence systems.	
II. Level Structure and Progressive Competency	4
Level I develops baseline competence in prompt construction, output evaluation, factual checking, and hallucination recognition.	
III. Assessment Philosophy and Live Demonstration	5
Assessment in the Programme privileges demonstrable judgment over memorised enthusiasm.	
IV. Ethics, Governance, and Human Responsibility	6
Ethics within the Programme is taught as an operating discipline rather than as a final lecture on unintended consequences.	
V. Credentialing, Progression, and External Recognition	7
Completion of each level is recorded through the University's credential infrastructure and linked to a clear statement of assessed competencies.	
VI. Review, Revision, and Institutional Ambition	8
The Programme is reviewed each year against labour market conditions, research developments, student outcomes, and internal evidence about where the curriculum is becoming stale, inflated, or insufficiently demanding.	

I. Programme Rationale and Educational Doctrine

— 3 —

The AI Skills Programme was established on the premise that literacy in the Third Epoch includes the disciplined use, evaluation, and governance of artificial intelligence systems. The University does not consider this proposition trendy. It considers it comparable to earlier recognitions that educated persons ought to be able to read critically, write coherently, reason quantitatively, and detect when a polished answer has outrun its evidence.

The Programme therefore treats AI competency as neither a niche technical add-on nor an excuse for boosterish inevitabilism. Its purpose is to develop graduates who can work with advanced systems without surrendering judgment to them. This balance is harder to teach than mere tool usage and materially more valuable. Fitzherbert has made it the organising principle of the curriculum rather than a concluding slogan in the brochure.

All levels of the Programme are designed to be practically assessable, ethically supervised, and legible to non-specialist employers, academic departments, and governance bodies. The University is not interested in issuing credentials that sound futuristic but fail to convey what the holder can responsibly do.

II. Level Structure and Progressive Competency

— 4 —

Level I develops baseline competence in prompt construction, output evaluation, factual checking, and hallucination recognition. The aim is not to create prompt maximalists but to produce users who understand that the quality of an answer depends on both the system and the human framing the task. Students learn early that vague instructions tend to return confident vagueness, which is one of the Programme's least romantic but most transferable lessons.

Level II moves into retrieval-augmented generation, fine-tuning concepts, workflow design, and domain-specific validation. Students are expected to build systems that can retrieve evidence, cite it accurately, and decline to improvise when the evidence is insufficient. The capacity to say "I do not know" is treated as a practical competency rather than an emotional failure, a pedagogical decision that some students find unexpectedly difficult.

Levels III and IV extend the work into multi-agent systems, deployment discipline, observability, governance, and alignment reasoning. By the end of the Programme, students must be able not only to use AI systems effectively but to situate them institutionally: who supervises them, how they fail, what their outputs can justify, and why human accountability remains non-transferable.

III. Assessment Philosophy and Live Demonstration

— 5 —

Assessment in the Programme privileges demonstrable judgment over memorised enthusiasm. Students are asked to critique outputs, design workflows, justify architecture decisions, respond to failure scenarios, and defend the boundaries of systems they have built. The University has found that AI education becomes quickly unserious when it rewards fluent description of tools more than the capacity to deploy them under accountable conditions.

Live exercises are central because the Programme wants to observe what students do when a system behaves strangely, incompletely, or seductively. A perfectly successful demonstration is not always the strongest performance. In many cases, the higher mark goes to the student who recognises risk, constrains the system appropriately, documents uncertainty, and resists the temptation to improvise a triumphant post-hoc explanation for a poor result.

Assessment records may include notebooks, build logs, retrieval evidence, red-team reports, oral defences, and governance memos. This breadth is intentional. Fitzherbert does not believe serious AI competence can be captured by a single glossy artefact unaccompanied by the process that produced it.

IV. Ethics, Governance, and Human Responsibility

— 6 —

Ethics within the Programme is taught as an operating discipline rather than as a final lecture on unintended consequences. Students are expected to evaluate bias, data provenance, deployment context, institutional incentives, misuse pathways, and the governance obligations that arise when systems influence real decisions. The University is impatient with educational models that teach technical capability first and ethical seriousness only after the capability has already been normalised.

Level IV in particular requires students to articulate a defensible account of why human governance remains necessary even when automated systems outperform humans in narrow and measurable tasks. The purpose of this requirement is not nostalgic human exceptionalism. It is to test whether the student can distinguish capacity from legitimacy and optimisation from authority.

The Programme committee is aware that some students initially experience these requirements as philosophically inconvenient. It considers that reaction pedagogically promising. A governance principle worth holding should be able to survive contact with technical sophistication rather than depending on its absence.

V. Credentialing, Progression, and External Recognition

— 7 —

Completion of each level is recorded through the University's credential infrastructure and linked to a clear statement of assessed competencies. The University has resisted pressure to make the credentials more expansive in language than the assessments justify. A credential is useful precisely to the extent that it can be believed by someone who was not present for the assessment event.

Progression between levels depends on demonstrated mastery rather than elapsed time. Students may move quickly if they can substantiate their readiness, but acceleration is not treated as a virtue in itself. Fitzherbert has seen enough technically gifted students underestimate the governance and interpretive demands of advanced work to regard pacing as part of educational seriousness rather than a drag on it.

External recognition of the Programme is supported by publication of competence statements, assessment principles, and verification pathways. Employers and partner institutions are encouraged to read the credential as evidence of disciplined AI capability under institutional oversight, which the University considers a more meaningful claim than generic familiarity with the latest tooling vocabulary.

VI. Review, Revision, and Institutional Ambition

— 8 —

The Programme is reviewed each year against labour market conditions, research developments, student outcomes, and internal evidence about where the curriculum is becoming stale, inflated, or insufficiently demanding. Review is especially important in AI education because a stable syllabus can become outdated while the validation language around it remains deceptively current.

Revision is undertaken with caution. Fitzherbert does not wish to become a programme that chases every technical trend only to rediscover, a year later, that students still need judgment, verification, architecture discipline, and ethical restraint. The University aims instead for durable competence with adaptive examples, an approach less flamboyant than rapid curricular churn and more educationally defensible.

This guide is therefore both curriculum statement and institutional wager. It assumes that the University can educate people not merely to use intelligent systems, but to govern their use in ways worthy of academic trust. Fitzherbert considers that wager ambitious, necessary, and very much the sort of thing a university ought to attempt.



INSTITUTIONAL NOTICE

This document is published by Fitzherbert University and archived in the Institutional Repository in accordance with the Transparency Mandate of 2003 (Charter Amendment IV).

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form without the prior written permission of the Office of the Chancellor, except for brief quotations in academic reviews and scholarly articles.

A cryptographic hash of this document is registered on the Fitzherbert Canonical Registry. The SHA-256 hash and associated metadata are available at the University's canonical verification endpoint.

Fitzherbert University · Established 1783 · Veritas per Verificationem