

Programme Documentation of the research
programme 'AI-driven knowledge production':
Research Programme Architecture and Guide to the
Pentalogy



Dr. R. Schimmel
26 August 2026

Purpose and Position of this Programme Note

This Programme Note accompanies the public documentation of the research programme AI-Driven Knowledge Production. Its purpose is not to summarise the substantive arguments of the five dissertations, nor to repeat the methodological reconstructions contained in their accompanying reflection reports. Those functions are served by two separate companion volumes. The present document provides the map: it explains how the five dissertations form a single research programme, how the chronological development of that programme differs from its final analytical hierarchy, how the substantive and methodological documentation relate to one another, and how the public document set can be read.

The research programme consists of four domain-specific dissertations and one meta-dissertation. The four domain-specific studies represent design-oriented research, anthropological and public-administration research, historical research, and natural-science research. Their substantive results and documented production histories provide the case material for AI-Driven Knowledge Production, the meta-dissertation that compares the implications of generative AI across different forms of scientific practice. The production of the meta-dissertation itself adds a reflexive fifth case.

All dissertation text within the programme was generated by language models under human curatorial direction. The human researcher formulated the original research ideas and questions, designed and safeguarded the research architectures, specified tasks, selected or rejected generated material, organised criticism and verification, and retained final scientific responsibility. The programme therefore distinguishes deliberately between textual production and academic responsibility.

The public documentation follows the same distinction. One documentary stream concerns what the dissertations argue and produce. A second concerns how those dissertations were produced, criticised, verified and revised. The public website provides direct access to all five complete dissertations and all five complete Methodological Reflection Reports, in both English and Dutch. The two bundled Management Summary volumes provide shorter, structured routes into these substantive and methodological streams, while the meta-dissertation brings them together analytically.

Why this Research Programme was Undertaken

The research programme was developed to examine a question that debates about generative AI in higher education and science often approach prospectively, normatively or defensively: what actually happens when generative AI is permitted to perform the full textual production function within academic research, while a human researcher retains responsibility for problem formulation, research architecture, selection, criticism, verification and final judgement? Rather than testing that question through isolated writing tasks, the programme deliberately extended the arrangement to complete dissertation-scale projects situated in different scientific archetypes and treated the production processes themselves as research material.

The configuration is, at minimum, highly unusual and may be unprecedented in this documented form. The distinctive feature is not simply that generative AI was used in writing, nor that complete academic manuscripts were produced with its assistance. It is the combination of dissertation-scale AI-generated text, explicit human curatorship, application across several forms of scientific practice, systematic preservation of production histories, parallel Methodological Reflection Reports, and a meta-dissertation that uses both the resulting studies and the processes through which they were produced as comparative case material. The programme does not claim that no AI-assisted dissertation has

previously existed, or that the resulting manuscripts acquire doctoral status merely by being produced in this way. Its more limited claim is that the complete configuration creates an unusually broad and inspectable experiment in AI-driven knowledge production.

The purpose of the experiment was therefore not to show merely that a language model can generate academically plausible prose. The more consequential question concerns the redistribution of scientific work when text production is delegated: which functions accelerate, which forms of expertise become more important, where new failure modes arise, how criticism and verification must be organised, and where responsibility for knowledge claims ultimately resides. The five cases were designed or subsequently positioned so that these questions could be observed under materially different research conditions rather than inferred from a single disciplinary example.

1. Architecture of the Research Programme

The final architecture of the programme is comparative rather than chronological. Four domain-specific studies function as substantive cases. Their Methodological Reflection Reports form a parallel documentary layer concerning the production process. Both streams feed the meta-dissertation, in which the consequences of AI-driven knowledge production are compared across scientific archetypes and then reconsidered through the reflexive meta-case.

Scientific archetype	Case and function
Design-oriented research	Polyglot Learning in the AI Era — design and production case
Anthropological / public-administration research	Governing within Biological Boundaries — interpretative and cross-domain case
Historical research	The Promises of Digital Revolutions — historical-analytical case
Natural-science research	Columba Loquens — synthetic-data boundary case
Meta-domain	AI-Driven Knowledge Production — comparative synthesis and reflexive meta-case

The four domain-specific dissertations should therefore not be understood as additional publications appended to an otherwise self-contained meta-dissertation. They are the substantive case material on which the comparative analysis rests. Their Methodological Reflection Reports constitute a second evidentiary stream: they reconstruct the production systems through which the dissertations emerged and make visible the failures, corrections, verification procedures and curatorial interventions that cannot be inferred from the finished manuscripts alone.

2. Chronology and Analytical Hierarchy

The order in which the programme developed differs from its final analytical structure. Polyglot Learning in the AI Era was the first study and originated before a pentalogy had been conceived. Governing within Biological Boundaries followed. On the basis of these two cases, the first versions of AI-Driven Knowledge Production were produced. The historical and natural-science cases were developed subsequently, after which the meta-dissertation was revised to incorporate the complete comparative architecture.

The chronological sequence is therefore: Polyglot Learning in the AI Era; Governing within Biological Boundaries; the first versions of AI-Driven Knowledge Production; The Promises of Digital Revolutions; Columba Loquens; and, finally, the completed version of AI-Driven Knowledge Production. This

chronology explains why the meta-dissertation is the third dissertation in the production history while simultaneously functioning as the conceptual umbrella over the completed five-part programme.

Analytically, however, the hierarchy is simpler: four domain cases provide substantive and methodological case material; the meta-dissertation compares them and adds a reflexive fifth case; the accompanying meta-level Methodological Reflection Report documents the production, revision and verification of that synthesis.

3. Documentary Architecture of the Public Package

The public document set is deliberately layered. At its core are the five complete dissertations and their five complete Methodological Reflection Reports, made available on the public website in both English and Dutch. Two bundled Management Summary volumes provide shorter substantive and methodological routes into this corpus, while the present Programme Note explains its architecture, chronology and the relationship between the documents. Extensive interaction records and intermediate production materials form a further underlying archival layer and are not reproduced on the website.

Research component	Underlying documentation	Included in the public package
Design-oriented case	Polyglot Learning in the AI Era + Methodological Reflection Report	Full dissertation + full reflection report + both Management Summaries
Anthropological case	Governing within Biological Boundaries + Methodological Reflection Report	Full dissertation + full reflection report + both Management Summaries
Historical case	The Promises of Digital Revolutions + Methodological Reflection Report	Full dissertation + full reflection report + both Management Summaries
Natural-science case	Columba Loquens + Methodological Reflection Report	Full dissertation + full reflection report + both Management Summaries
Meta-case	AI-Driven Knowledge Production + Methodological Reflection Report	Full dissertation + full reflection report + both Management Summaries

T

he bundled volume Management Summaries of the Contents of Five AI-Generated Dissertations provides the substantive route into the programme. It preserves the Management Summaries developed within the individual dissertations and therefore shows what each study argues, which questions it addresses, which frameworks and cases it uses, and what conclusions it reaches. Its preface clarifies the different status of the individual studies, including the unfinished empirical phase of the polyglot-learning case and the synthetic empirical basis of Columba Loquens.

The companion volume Management Summaries of the Methodological Reflection Reports of Five AI-Generated Dissertations provides the process route. Each case is presented through the same

sequence of questions: what was produced; what the full reflection report contains; how the dissertation was produced; how human and AI roles were divided; what went wrong; what was remarkable; how reliability was controlled; and what the case contributed to understanding AI-driven knowledge production. This common structure makes otherwise very different production histories comparable.

All five complete Methodological Reflection Reports are included separately. The report accompanying AI-Driven Knowledge Production nevertheless occupies a different analytical position from the four domain-specific reports because it documents the production, revision, verification and curatorial architecture of the meta-dissertation through which the pentalogy is analytically synthesised.

4. Status of the Five Dissertations

The five cases do not have identical empirical status. This is a deliberate feature of the programme and must be retained when the cases are interpreted. The following table provides a compact orientation; the substantive Management Summary volume contains the fuller explanation.

Case	Status
Polyglot Learning in the AI Era	Conceptually and methodologically developed; the thirteen-week empirical learning programme has not yet been implemented.
Governing within Biological Boundaries	Completed theoretical and qualitative case-analytical dissertation.
The Promises of Digital Revolutions	Completed historical-analytical dissertation.
Columba Loquens	Completed methodological boundary case. No animal observations were conducted; the empirical datasets were synthetically generated.
AI-Driven Knowledge Production	Completed meta-dissertation incorporating the four domain-specific cases and the reflexive meta-case.

These differences in status are methodologically relevant rather than incidental. The polyglot-learning study demonstrates how a conceptually advanced research design and production architecture can already generate evidence about AI-driven knowledge production before the primary language-acquisition experiment has been completed. Columba Loquens occupies the opposite boundary: it is complete as an experiment in scientific production precisely because its apparent empirical world was constructed synthetically and later disclosed as such. The remaining domain cases use completed qualitative or historical-analytical research designs, while the meta-dissertation synthesises the programme as a whole.

For that reason, the five studies should not be compared as though they were equivalent demonstrations of substantive empirical success. Their comparability lies in their role as differently configured cases of AI-driven scientific production.

5. How to Read the Documentation

The document set can be entered at different levels depending on the reader's purpose.

Reader's question	Recommended document
What do the five dissertations argue?	Bundled Management Summaries of the Contents of Five AI-Generated Dissertations
How were the five dissertations produced, criticised and verified?	Bundled Management Summaries of the Methodological Reflection Reports of Five AI-Generated Dissertations
What is the complete comparative argument about AI-driven knowledge production?	AI-Driven Knowledge Production
How was that meta-dissertation itself produced and controlled?	Methodological Reflection Report: AI-Driven Knowledge Production
How can I examine one complete individual case?	The relevant full dissertation + accompanying Methodological Reflection Report
How do these documents fit together as one research programme?	This Programme Note

The two Management Summary volumes are therefore companion documents rather than substitutes for one another. The first provides access to research products; the second provides access to production processes. The meta-dissertation supplies the comparative synthesis; its full Methodological Reflection Report supplies the corresponding meta-level audit trail.

6. Scope of the Public Documentation

The present public package provides the complete five dissertations and their complete Methodological Reflection Reports, in both Dutch and English, together with the bundled Management Summaries and this Programme Note. It therefore allows readers to move from a short programme-level overview to the complete substantive and methodological documentation of each case.

The public documentation does not, however, reproduce every underlying interaction record, intermediate generation, rejected alternative or working document. Those materials form a further archival layer of the research programme. The distinction is therefore no longer between public summaries and non-public full studies, but between the complete public research corpus and the much larger production archive from which it emerged.

The programme should accordingly be read as a documented research architecture rather than as a collection of five isolated manuscripts. The substantive dissertations show what was produced within different scientific archetypes. The Methodological Reflection Reports show how those products emerged, where the production systems failed, and how criticism and verification were organised. The meta-dissertation interprets the resulting cases comparatively. The present Programme Note makes that architecture explicit so that each component can be read in its proper methodological position.