

AI Erasmus — A Platform for a New Era in Digital Humanities

The philosophy, the operations of an AI famulus for scholars

Gerd Graßhoff 

Humboldt-Universität zu Berlin, Deutsches Museum
gerd.grasshoff@hu-berlin.de

ABSTRACT

Critical editions make sources available; they do not sit beside the scholar as working instruments. AI Erasmus closes that gap. It is a research platform that treats editions as the substrate of inquiry and organises the scholar's work around operations: reading with an edition, querying a corpus with a visible method, consulting a dictionary whose senses are drawn from one declared scope, tracing persons and correspondence, and comparing concepts across authors and fields. The platform rests on a philosophy stated here in full — the amanuensis stance, the page as the unit of evidence, meaning as use, index-first answering, and provenance on every result — and the guide shows each commitment at work in the instrument, with screenshots, method traces, and measured numbers. Editions of Kepler, Copernicus, Euler, Pirckheimer, and von Kármán appear only as material inside examples, never as catalogues. The aim is to show how digital humanities can move from publishing sources to supporting the everyday craft of historical and conceptual research.

KEYWORDS digital humanities · critical editions · PDAI · research operations · historical dictionaries · lemma search · correspondence networks · concept history

Table of contents

0.1	Introduction: The New Erasmus	2
0.1.1	The Famulus and What It Inherits	2
0.2	Basel, 1521–1529: a household of famuli	2
0.3	Nuremberg, 1471–1547: instruments, presses, and scholarly households	3
0.4	The modern scholar's inverse situation	4
1	The Philosophy of AI Erasmus	5
1.1	The amanuensis, not the oracle	5
1.2	The page as the unit of evidence	6
1.3	Meaning as use: the dictionary's philosophy	7
1.4	Index first, generation second	7
1.5	Provenance, or the condition of trust	7
1.6	The speed argument	8
2	The Substrate in One Chapter	8
3	Operation I — Working with Editions	9
4	Operation II — Querying: What to Ask, and the Fascination of Speed	10
4.1	The families of askable questions	10
4.2	An inventory answered	10
4.3	Speed as experience	11
4.4	The score behind the answer	12
4.5	What falls off the fast path	13
5	Operation III — The Dictionary as a Focused Instrument	13
5.1	The anatomy of an entry	13
5.2	Focused on one body of writing	15
5.3	One person against another	15
5.4	Meaning in time	16

5.5 Meaning in an academic field	16
6 Operation IV – Prosopography and Correspondence Networks	16
7 Operation V – Concept History and Comparison	17
8 The Full Scope of Application	17
8.1 Four kinds of material, one instrument	17
8.2 The corpus at the time of writing	18
8.3 Teaching and collaboration	19
8.4 What AI Erasmus does not do	19
9 Implications for Scholarly Practice	19
10 Goal Expansion and the Famulus Handlungsmodell	20
10.1 Goals and Partiturs	20
10.2 The Action Repertoire	20
10.3 Iterative Goal Expansion	21
10.4 Repertoire Growth – the Generation-Partitur Loop	21
10.5 Failure Diagnosis as Repertoire Refinement	21
10.6 The Philosophy of the Handlungsmodell	21
11 Conclusion and Availability	22
12 Acknowledgements	22
Bibliography	22

■ 0.1 Introduction: The New Erasmus

0.1.1 The Famulus and What It Inherits

Famulus Erasmus is a working tool, and its name states what kind of tool it is. In the humanist workshop the famulus copied, collated, fetched and checked: he held the second copy against the first, carried the letter to the press, and kept the scholar's desk supplied with what the argument required. The system described here occupies that position. It does not settle the argument; it prepares the material on which the scholar settles it.

Three tasks defined the famulus, and the same three define this tool. The first is care for the reliable text: someone compares the witnesses and records where they differ. The second is the scholarly dialogue: the letter, the query, the answered objection, by which a claim is tested before it is printed. The third is dissemination: the press turns the checked text into copies that reach readers who will never see the manuscript. In each task a person does something to a document, and the tool takes over parts of each without taking over any of them whole.

The excursus that follows traces these three tasks where they were first organised at scale, in the printing houses of Basel and Nürnberg. It is not an ornament. The division of labour worked out in those workshops – who corrects, who compares, who signs the result – is the division of labour this tool proposes to redistribute, and the sections after the excursus measure the proposal against it.

■ 0.2 Basel, 1521–1529: a household of famuli

For the decade of his first Basel residence (1521–1529), Erasmus of Rotterdam kept a workshop in the printing house of Johann Froben: *amanuenses*¹ – in the sources usually called *famuli* – who wrote letters from dictation, collated manuscripts against variant witnesses, prepared the printer's fair copies (*exemplaria*),

¹The word goes back to the Latin phrase *servus ā manū* – literally “slave from the hand”, that is, one whose function was the handwriting. The substantive *āmanuēnsis* is formed from the second element, *ā manū*, with the belonging-suffix *-ēnsis*. The construction *ā* + ablative is the regular Roman formula for household service-roles: *a rationibus* (accountant), *a libellis* (petition-secretary), *ab epistulis* (correspondence-secretary). The term constitutively implies executing, not judging – the hand, not the head. The English borrowing dates only from about 1610. Scandinavian readers may note that *amanuensis* in the present text always means the person who holds the pen, never the one who holds the chair: the Norwegian academic title *førsteamanuensis* (roughly, associate professor) carries the exact inverse meaning.

and read proof. Among the identifiable names from this Basel decade: Lieven Algoet, Karl Harst, Quirinus Talesius, Nicolaas Kan, and Felix Rex [1].

The editions that came out of that work — the Greek New Testament, the church fathers, the *Adagia* — were the product of a division of labour that requires a little precision. Three groups contributed. The *famuli* in Erasmus's own household performed the repeatable clerical moves: dictation, copying, collation, preparation of the *exemplaria*. Froben's *castigatores* — the press correctors, employed by the printing house and not by Erasmus — were of different standing: Beatus Rhenanus, Sigismundus Gelenius, and for the *Novum Instrumentum* of 1516 also Oekolampadius, each a philologist who made textual decisions of his own. The compositors, pressmen, and factor ran the workshop proper.

The formula “one judgment supported by many hands” holds for the *famuli*: their function was the hand, the judgment stayed at the desk. It is already too strong for Froben's *castigatores*, whose hands also thought. Rhenanus and Gelenius made philological choices; the division of labour was blurred at precisely the point where the analogy needs it sharp. That concession belongs in the open, not buried in qualifications.

ERASMUS OF ROTTERDAM, A WORKSHOP

Erasmus (1466–1536) is remembered as the prince of the humanists; his working method is the better lesson. He edited texts against manuscripts, annotated them line by line, cross-referenced a classical corpus he carried largely in memory, and conducted a correspondence that functioned as the learned network of his age [2]. None of this was solitary genius at a bare desk: it was a staffed operation, embedded in a printing house, with trained assistants performing the repeatable moves so that editorial judgment could be spent where only judgment serves. The platform is named for the workshop, not for the genius.

■ 0.3 Nuremberg, 1471–1547: instruments, presses, and scholarly households

[Scope note. The preceding Erasmus paragraph is not in the segment supplied to this lector; it cannot be revised here without inventing its wording. Return it with the segment for a second pass. Slots marked [CITATION NEEDED] mark the load-bearing assertions; supplying the references falls to the bibliographer, not to this lector, and no reference has been composed here.]

Nuremberg issued the instrument alongside the edition; Basel did not. The Basel household worked a pattern found across Europe, and Nuremberg shows what varied within it. There the same tasks fell to the same household members. But the household worked in two media: it cut, graduated and sold brass instruments, and it set, corrected and printed books. The chronology rules out a generational gap. The high period of new scientific printing opens in the 1520s in the shops of [PRINTERS — NAMES NEEDED], some fifty years after Regiomontanus died [CITATION NEEDED: dating of the Nuremberg scientific press; Regiomontanus's death]. The Nuremberg and Basel households were therefore largely contemporaries. Nor did Nuremberg forgo editorial work: its printers issued Ptolemy's *Geography* and the *Almagest* as recensions of Greek texts prepared by [EDITORS — NAMES NEEDED] [CITATION NEEDED: editions, editors, places and dates]. The difference lies in the second medium, not in a division between science and philology. V3★

[Inserted paragraph, drafted to the user's instruction. It is offered as a shape for the argument, not as an assertion of fact: every proposition below carries a citation slot, and none should go to press unfilled.]

The Nuremberg workshops also stood inside the humanist correspondence network, and not merely beside it. Willibald Pirckheimer, the city's foremost Renaissance scholar, was a close friend of Albrecht Dürer [CITATION NEEDED: the friendship; Dürer's letters to Pirckheimer from Venice]. Dürer's household ran a workshop for printed art, in which his wife and family had a stated part in production and sale [CITATION NEEDED: the composition of the Dürer workshop; Agnes Dürer's role in selling prints at fairs — the extent of that role is contested and should be given as the sources give it, not as settled]. Pirckheimer in turn corresponded with Erasmus [CITATION NEEDED: Allen, *Opus epistolarum Des. Erasmi Roterodami* — specific letter numbers, not the edition at large]. What that correspondence shows about a shared outlook, and about any exchange of works between the two men, must be shown from the letters themselves: the

claim of a close spiritual attitude and a joint production cannot rest on the fact of an exchange of letters alone [CITATION NEEDED: the specific letters, and any secondary study that reads them this way]. Set beside the Basel household, the Nuremberg case therefore adds a third thing to the instrument and the edition — the letter — and it is through the letter that the two cities' workshops are connected at all.

[EDITORIAL NOTE, outside this lector's remit and forwarded to the chief: the instruction 'include much more references' cannot be executed by a style lector. Fabricated citations are worse than absent ones. The slots above are the deliverable this pass can honestly produce.]

In 1471 Johannes Regiomontanus (Johannes Müller von Königsberg, 1436–1476) settled in Nuremberg. His stated reasons for the choice are the paper's own thesis in fifteenth-century form: the city manufactured the finest astronomical instruments, and its commercial reach made it the communications hub of Europe — the best location from which to reform astronomy in correspondence with learned men anywhere. Together with his patron Bernhard Walther (1430–1504) he built an observatory, a mechanical workshop for instruments, and the first printing house in Europe dedicated to scientific books. Its 1474 broadside publishing programme announced editions of ancient, medieval, and contemporary mathematical and astronomical texts: the operations of copying, correcting, and diagram-printing delegated to other hands so that editorial judgment could be placed on the text.

Regiomontanus died in Rome in 1476; Walther continued the observations until his own death in 1504 and kept the manuscript remains. The transmission link to the next generation was Johannes Schöner (1477–1547), who acquired Walther's papers, edited and published Regiomontanus's posthumous works through the Nuremberg presses of the 1530s and 1540s, and made his own reputation as globe-maker and instrument maker. On the recommendation of Philipp Melanchthon, Schöner became professor of mathematics and astronomy at the newly founded Ägidien-Gymnasium on 25 July 1526 and held the post until his death, 16 January 1547; Melanchthon's epitaph placed him on the same level as Regiomontanus himself [3], [4].

Instrument making was the city's recognised specialty. Georg Hartmann (1489–1564), vicar at the Sebalduskirche, was among the most prolific instrument makers of his generation: quadrants, compasses, dials, celestial and terrestrial globes, astrolabes, gunners' equipment — and his *Collectio figurarum*, a rich gathering of engravings of sundials and astronomical instruments, witnesses the scale of the trade. He was personally acquainted with Pirkheimer, Dürer, Schöner, Peter Apian, and Melanchthon — the Nuremberg scholarly household in its full breadth [3]. Johannes Stöfler (1452–1531), a generation earlier, had already built celestial globes, wheel clocks, sundials, and equatoria that served the same conjunction of learned demand and artisanal precision.

Willibald Pirkheimer (1470–1530) — patrician, imperial councillor, translator of Greek texts, and Dürer's closest friend — represents the household in its patrician form. He was the centre of Nuremberg's humanist circle; his letter volumes are one of the platform's own collections and a first-person record of how the city's learned world organised itself.

What unified Basel and Nuremberg was not geography but an economic logic that Melanchthon stated plainly: without the export trade, the artists' workshops would lack the means. Sebastian Münster described Nuremberg from the outside as a city that compensated poor soil with *sinnreiche Werckmeister* — cunning craftsmen — and their work reached the learned world through Johann Petreius (c. 1496/7–1550), the Nuremberg publisher who made his name by bringing out high-grade scientific books [3]. The workshops of Basel and Nuremberg differ in medium — editions of texts on one side, instruments and diagram-printing on the other — but share the same structure: repeatable moves delegated to trained hands so that judgment can be spent where only judgment serves.

■ 0.4 The modern scholar's inverse situation

The modern scholar's situation is the inverse. The editions available today would have astonished the workshop — critical apparatus, complete correspondences, facsimiles of every page — and the assistance has vanished. One can open a letter volume of Kepler's *Gesammelte Werke* and still not know, without an

afternoon’s labour, how many letters it contains under the editor’s own numbering [5]. One can read Euler and still not see, at a glance, how *tempus* is used in a single treatise. One can hold a correspondence edition and still lack a fast path from a name to the pages where that person appears. The gap is not digitisation; most of these editions are scanned. The gap is the absence of *operations* — repeatable moves between a research question and page-grounded evidence, the moves the workshop staffed with people.

AI Erasmus is a research platform built to close that gap, and this guide is its full statement. The design decision from which everything else follows is organisational: the platform is organised by *operations*, not by collections. There is no module called “the Kepler corpus” and no section of this guide that surveys one. There are five operations — working with an edition, querying, consulting a focused dictionary, tracing persons and correspondence, comparing concepts — and the editions appear inside them as material, the way scores appear inside a conservatory’s curriculum.

The guide makes five contributions:

1. **A philosophy of the instrument.** Five commitments — the amanuensis stance, the page as evidence, meaning as use, index-first answering, and provenance — stated as design decisions with consequences (Chapter 2).
2. **An operations thesis.** Digital-humanities platforms should be designed around what scholars *do* with editions; the five operations are developed one chapter each (Chapters 4–8).
3. **Speed as scholarly experience.** Interactive latency — a boolean lemma search over a five-million-token collection returns in 67.1 ms in the run shown in Chapter 5 — changes which questions get asked.
4. **The dictionary as a focused instrument.** Senses are always scoped: one body of writing, one person against another, a span of time, or an academic field (Chapter 6).
5. **The full scope on one substrate.** Four kinds of material — early modern Latin science, humanist correspondence, a twentieth-century archive, a catalogue — served by the same five operations, measured in the July 2026 build (Chapter 9).

Three companion papers carry what this guide assumes. The PDAI substrate — the page format and the processing pipeline — is the subject of AI Scholar [6]. The orchestration behind every answer is written in Partitur scores, and the score language has its own guide [7], [8]; Chapter 5 opens the door between the two. And the question of whether an answer can be trusted — the golden-query testbed, its thirteen canonical goldens across seven forms of inquiry, the monitor log that makes a run inspectable, and the learning loop that turns a diagnosed failure into a promoted patch — is the subject of a third companion, *Golden Queries in AI Erasmus* [9]. This guide stays on the scholar’s side of that door and asks one question throughout: what can a scholar *do*, today, that the workshop in Basel would have recognised as its own work, done faster and without the staff?

1 The Philosophy of AI Erasmus

This chapter states the five commitments on which the platform is built. Each is a design decision, not a slogan: each has consequences that the later chapters exhibit in the instrument, and each rules something out. A reader who holds the five commitments can predict what AI Erasmus will refuse to do — which is the test of a philosophy that has actually been built.

1.1 The amanuensis, not the oracle

The first commitment fixes the division of labour. In Erasmus’s workshop the amanuensis copied, collated, fetched, and counted; the judgment about what a variant meant, which reading to print, what a passage implied, never left the scholar. AI Erasmus keeps that line.

CONCEPT The amanuensis stance

The platform drafts, retrieves, counts, and assembles; it does not judge. Every output is offered as **material for** a scholarly act — a count to check, a citation set to read, a sense inventory to criticise — never as

the act itself. A platform that answered “what does this passage mean?” with a settled verdict would have changed sides.

AMANUENSIS, FAMULUS, HAND: A TERMINOLOGICAL NOTE

Two dimensions that are routinely conflated — whether a contributor **owns judgment** and whether a contributor **receives credit** — sort four roles into a matrix:

	credited	uncredited
own judgment	colleague · co-author	ghostwriter
no judgment	acknowledgement	amanuensis

The ghostwriter and the amanuensis are exactly complementary: one exercises judgment without credit, the other forgoes credit because judgment was never the function. This paper argues for the lower-right cell — “the judgment remained at the desk” — so the word is a thesis, not a description. That makes the choice of word consequential, especially when the analogy targets AI systems.

Why the archaic term is kept. Amanuensis is semantically uncontaminated by modern employment, authorship, and credit arrangements. No institutional expectation adheres to it; a reader brings no assumed hierarchy of co-authorship or qualification. The liability is narrow: the word strictly covers writing to dictation, and not the collating, indexing, and counting that Erasmus’s household also performed. That is why **hands** continues as the load-bearing metonymy in this text: the hand–head distinction is performed by the word itself, without the narrowness of the Latin term.

Two adjacent formulas carry a caution worth naming. Shapin’s **invisible technicians** [10] is science-historically established, but its argument is a critique of the invisibilisation itself — adopting it imports the charge that assigning all judgment to the desk was an ideological construction, which is not the claim here. **Human computers** [11] is the strongest candidate for the counting aspect and fits the AI analogy tightly, since there the name historically migrated from the person to the machine; the formula appears where counting is the explicit topic.

The stance has a practical edge: it decides what counts as a good answer. An oracle is judged by fluency; an amanuensis² is judged by accountability. A count with a method attached, a hit list with page addresses, a dictionary draft marked as draft — these are good answers even when they are inconvenient, because the scholar can check them. A fluent paragraph with no path back to evidence is a bad answer even when it happens to be right, because nobody can tell.

1.2 The page as the unit of evidence

Every scholarly claim eventually meets the question: *where does it say that?* The second commitment ensures the question always has an answer. The unit of evidence in AI Erasmus is the *page* — one raster image plus its structured transcription, sharing one address [6]. Every operation in this guide resolves to pages: a query returns pages, a dictionary quotation links to its page, a correspondence edge opens on the page of the letter.

The consequence is that disagreement has an address. When a scholar doubts a count, a sense, or an attribution, the dispute moves in one click to a facsimile and its transcription — the same ground a printed

²The sources — Erasmus’s letters and prefaces — call the household group *famuli* (house-companion, pupil, and servant at once) rather than *amanuenses*; *amanuensis* is the narrower term restricted to the writing function. How frequently Erasmus himself used *amanuensis* is a question for the index of Allen’s *Opus epistolarum*. In German, *Famulus* carries two live modern occupations that make it unreliable as a descriptive term: the *Famulatur* is still the standard term for the clinical clerkship, and Goethe’s *Famulus Wagner* — the learned, judgment-free imitator who in *Faust II* produces the Homunculus — makes the literary echo either a gift or a trap for an analogy targeting artificial assistance. The exact modern institutional analogue is the *wissenschaftliche Hilfskraft* (student research assistant). Terms deliberately avoided: *research assistant* implies employment, a qualification goal, and increasingly co-authorship — the opposite of the thesis; *secretary*, *scribe*, and *clerk* each carry modern institutional expectations the archaic Latin term does not [1].

apparatus would offer, reached faster. A platform whose answers ended in paraphrase would leave disagreement nowhere to stand; a platform whose answers end on pages turns every dispute into a reading.

The further consequence is that the pages on which this platform rests are *Grounding Evidence* in a precise sense: direct textual or observational bases for historical claims, distinct from secondary literature that interprets or summarises them [6]. The page is certified — structurally audited, provenance-tracked, and regression-anchored — before it enters the corpus. This is what allows the platform to print numbers and mean them. A secondary account of what Kepler wrote may paraphrase; the Kepler page in this corpus transcribes. The two kinds of evidence carry different epistemic weight, and the platform is built to serve only the former.

1.3 Meaning as use: the dictionary's philosophy

The third commitment governs everything the platform says about words. Wittgenstein's reminder — “the meaning of a word is its use in the language” [12] — is taken here as an engineering directive, not an epigraph. If meaning is use, then a claim about meaning is a claim about a *body of evidence*, and it can be right or wrong only relative to that evidence. The platform therefore never states what a word means; it surveys how a word is used, in a scope the scholar declares.

CONCEPT Meaning as use, operationalised

A dictionary entry in AI Erasmus is a **grammar of use**: a descriptive survey of how a term is actually used within one declared scope of evidence. Its form follows the Oxford English Dictionary — headword, forms, numbered senses, dated quotations [13] — because the OED entry is a **perspicuous representation**: an arrangement of examples that makes the pattern of use visible at a glance (Wittgenstein's *übersichtliche Darstellung*). Senses emerge from quotations; no sense is asserted without a page behind it.

The commitment rules out a familiar temptation: pasting a general lexicon onto a historical corpus. Latin *tempus* in an Euler treatise is not “time” in the abstract; it is whatever the treatise's own sentences make it. Chapter 6 builds the whole third operation on this point.

1.4 Index first, generation second

The fourth commitment fixes where answers come from. AI Erasmus is *index-first*: answers are produced from indices over pages — lemma tables, letter registers, gazetteers — and a language model assists at the edges, phrasing a question into a query or an answer into prose. The generation never supplies the facts.

The reason is an asymmetry of failure. An index-first answer can be empty, and an empty answer is honest: “no letters from this sender in this range” is scholarship's ordinary weather. A generated answer is never empty; pressed for something it does not know, a model produces fluent material with no page behind it, and the failure is invisible precisely where it is worst. The platform prefers the failure one can see. Wherever this guide shows a fast answer, the speed comes from an index; wherever it shows an empty answer, the emptiness is reported, not papered over.

1.5 Provenance, or the condition of trust

The fifth commitment binds the other four together. Every answer carries its *provenance*: the page addresses it rests on, the method that produced it, and the scope it was computed over. In the viewer this appears as a method trace attached to the answer; in this guide it appears as the trace inside every operation box.

CONCEPT Provenance

The recorded origin of a result: page address, method, scope, and the specialist that produced it. Provenance is what turns a fast answer into a checkable one — and it is the platform's only defence against its own speed, because an answer that arrives in milliseconds will be believed in milliseconds.

Provenance is deliberately boring machinery with a philosophical job. The amanuensis stance says the scholar judges; judging requires knowing what one is judging. A count of 122 letters is not judgeable; “122 letters, counted from the edition’s *Nr.* headers, over KGW 13” is. The trace does not make the answer right — it makes the answer *criticisable*, which is the most any method can promise.

1.6 The speed argument

The last commitment is the one that looks least philosophical and is not. Latency decides which questions are askable *in practice*. A query that costs a weekend is asked once, cautiously, after committee-grade deliberation; a query that costs fifty milliseconds is asked, refined, inverted, and asked again — inquiry becomes conversational. The difference is not comfort. Whole classes of questions — the hunch, the sanity check, the negative control — are only ever asked when asking is nearly free.

The workshop knew this. Erasmus could lean over and ask an assistant to check a reference *now*, and the texture of his annotation practice depended on it. Batch-mode digital humanities — submit the query, wait for the run — reproduced the archive’s reading-room rhythms in silicon. AI Erasmus is built for the workshop’s rhythm instead, and Chapter 5 shows what it feels like, with the measured numbers that back the claim.

2 The Substrate in One Chapter

The philosophy needs machinery that can keep its promises. This chapter states the substrate in the minimum detail the rest of the guide requires; the full format and certified production pipeline are the subject of AI Scholar [6].

The ground layer is the PDAI page: one raster image and one structured transcription that preserves layout and reading order, under a single address. A page enters the corpus only after passing a structural audit gate: every defect is attributed to one of four root-cause layers (prompt, normaliser, CSS, or schema grammar), registered as a versioned learning case, and resolved through an A/B-gated instruction improvement loop [6]. This is what makes the page a trustworthy unit of evidence rather than an aspirationally transcribed scan: the audit gate is binary, not a quality aspiration. An unresolved structural blocker means the page is not yet in the corpus.

Above the certified pages, derived layers accumulate without ever overwriting them:

- translations
- linguistic trees and lemma tables built with Stanza under Universal Dependencies [14], [15] and stored in DuckDB for interactive query [16]
- gazetteers of persons, places, and domain concepts, fed in part by named-entity recognition [17]
- letter registers where an edition’s metadata resolves senders, recipients, and dates
- a LWgrammar lexicon of OED-form usage entries built from the certified pages, one entry per lemma with dated quotations linked to their sentence coordinates

Every operation in this guide reads one or more of these layers, and each chapter names which.

Three contracts govern the stack, and all three exist to keep the philosophy of Chapter 2 honest.

METHOD The audit gate

A page enters the corpus only when its structural transcription passes zero-blocker certification: every structural, schema, and visual defect resolved, every residual fault registered as a versioned learning case in the ai-query annotation system. A page that failed certification and was silently promoted anyway is not Grounding Evidence; it is an unresolved liability. The gate is enforced by the *certify-and-learn* stage of the *pilot-gated-ocr-pipeline* Partitur.

METHOD The integrity rule

A derived layer must never lag its source. An index older than the transcription it was built from is not a research instrument; it is a liability wearing one's clothes. When a page is re-transcribed, every layer above it is rebuilt before it may answer again. Existence of a layer is not freshness — the rule is checked, not assumed.

METHOD The provenance contract

Every page records which model and prompt produced its transcription; every derived layer records what it was built from and when; every answer records which specialists ran, which layers they read, and over what scope. The chain from answer back to page is never broken, which is what allows this guide to print numbers and mean them.

One further piece of machinery deserves a sentence here, because Chapter 5 walks through its door: the routing of a scholar's question to the layers is itself written notation. Every query is played from a Partitur score by a Conductor [8] — the orchestration is a document one can read, not code one must trust. This guide stays on the scholar's side; the score side has its own guide [7].

That is all the substrate this guide needs. What follows asks what a scholar can do once these layers are in place — which is where the workshop reopens.

3 Operation I — Working with Editions

The first operation is the oldest, and the platform's treatment of it sets the tone for the other four: work with an edition, digitally, without losing the edition. The scene is concrete. A scholar reading Kepler's correspondence opens volume 13 of the *Gesammelte Werke* [5] in the viewer and sees the *split view*: the facsimile on one side, the structured transcription on the other, both under one page address.




Figure 1 The split view on a Kepler letter page (KGW 13): facsimile and structured transcription share one page address. The edition's own segmentation — the editor's *Nr.* headers — drives navigation.

The split view
every other op
dictionary quo
correspondence

Three properties make this a workplace rather than a display. First, the reading surface is *edition-faithful*: navigation follows the editor’s own segments — letters under their *Nr.* headers, chapters, apparatus — so the scholar’s orientation in the digital volume matches their orientation in the printed one, and every citation resolves to edition coordinates a reader can take to the shelf. Second, layers are optional company: a translation can sit beside the Latin, linguistic annotation can be summoned and dismissed; the page is never replaced, only accompanied. Third, every location is a *deep link* — an address that can be put in a footnote, mailed to a colleague, or returned to from a query result.

The third property matters most, and quietly. Every operation in the next four chapters ends with a page link, and the link lands here, in the split view, with the facsimile in sight. The platform’s promise that disagreement has an address (Chapter 2) is kept by this surface: whatever an index claims, the page is one click away, and the page outranks the index.

What the operation deliberately does not do belongs in the same breath. The workplace does not re-edit: apparatus, stemmatics, emendation, and the authority of the printed volume remain editorial territory. The same split view serves a Kepler letter, an Euler treatise, and a typescript carbon from a twentieth-century archive; what changes across editions is the editorial convention, and the viewer follows the convention rather than flattening it. The machine sits beside the page. It does not sit on it.

■ 4 Operation II — Querying: What to Ask, and the Fascination of Speed

Editions make texts readable; indices make questions askable. The second operation is the practice of asking, and this chapter stages it in three beats: an inventory question answered as research, a lemma search fast enough to change the asker, and a look through the door at the notation behind both.

■ 4.1 The families of askable questions

The useful first move is classification. A scholarly question, phrased in ordinary language, is routed by *decomposition* — the deterministic classification of the question into a family with a declared evidence path. Six families cover the practice:

Table 1 The families of askable questions and the layer each one reads.

Family	Example question	Evidence path
Inventory	How many letters in this volume?	Edition headers, letter register
Lemma / boolean	Where do <i>circulus</i> and <i>centrum</i> co-occur?	Lemma index
Person / mention	Where is this name used?	Person gazetteer
Correspondence	Who wrote to whom, when?	Resolved letter register
Dictionary / sense	How is <i>tempus</i> used here?	Headword entry, dated quotations
Comparative	Same lemma under two scopes?	Scoped indices side by side

Table 1 is a working map, not a closed ontology; its point is the third column. Every family names the layer it reads, so every answer arrives knowing where it came from — the provenance contract of Chapter 3 applied to question-asking.

■ 4.2 An inventory answered

Inventory questions look humble and are not. “How many letters does KGW 13 contain?” is the edition checked against itself: the editor’s numbering, actually counted, on the scholar’s demand. Before the platform, the honest answer cost an afternoon with the volume and a tally; the question was therefore rarely asked, and editions went unchecked.

The screenshot shows a web browser window with the URL `http://ai-viewer/doc/kgw-13/viewer`. The main content area displays a Latin text from Kepler's 'De Revolutionibus' (Book 1, Chapter 10). The text is in a serif font and includes line numbers (Nr. 23, 1595, 35). The text discusses the nature of the world and the number of celestial bodies.

On the right side, there is a 'Query' panel. It shows the query: 'wie viele Briefe enthält der Band 13'. The results show: '122 letters** in 'kepler-prim/kgw-13' (Nr. 1-127). Counted via 'ocr_header_nr'. Below this, there is a table titled 'HOW THIS SEARCH RAN' with two rows: 'specialist' and 'assemble', both showing '122 pages'.

Figure 2 An inventory question answered in the viewer: the letter count for KGW 13 arrives with its method trace attached. The trace names the specialist and the evidence path — edition headers, not page totals.

OPERATION Inventory: letters in a volume 0.4 s

Q How many letters does KGW 13 contain?

1	fasttext-classify	classifies the question into the inventory family within its 30 ms budget
2	collection-select	fixes the scope to the Kepler collection — the score's iron rule keeps scope terms out of any content filter
3	volume-letter-count	reads the edition's <i>Nr.</i> headers in the structured transcription and counts unique letter numbers

→ 122 letters in KGW 13 in the measured build — counted from the editor's own numbering, with the trace naming every specialist that ran

The trace is the point, more than the number. The answer names its method — headers counted, not pages estimated, not a model's recollection — so a scholar who doubts the count knows exactly what to check and where. An inventory answer without a trace would be a statistic; with one, it is a verifiable act of edition-criticism that took under half a second.

4.3 Speed as experience

The second beat is the one that changes the asker. The platform's lemma search answers boolean questions over the linguistic layer — and it answers at conversational speed. The search shown below asks for pages where *circulus* and *centrum* co-occur, over the Copernicus print collection.




Figure 3 The lemma search for *circulus* AND *centrum* over the Copernicus print collection: 138 pages in 6 documents, returned in 67.1 ms. Each document row carries its per-lemma counts and a deep link that opens the top page in the split view.

The measured run returns 138 pages across 6 documents in 67.1 ms, ranked by lemma density, each hit one click from its page. At that latency the cost of a follow-up question is effectively zero, and the practice changes shape: add a conjunct, drop a distractor, switch the collection scope, chase the hunch — six questions in the time a catalogue search used to take to load. This is *interactive latency* in the precise sense of Chapter 2: not a benchmark, but the condition under which asking becomes continuous with thinking.

The engineering is compact to state and belongs to the companions: lemma-sorted columnar storage, DuckDB pruning, a tiered retrieval funnel, gazetteer expansion where names vary [16], [18], [19]. The scholarly point is what those choices buy: the negative control one actually runs, the sanity check one actually performs, the exploratory question one actually asks.

4.4 The score behind the answer

Both answers above were played from the same written notation. Every question in AI Erasmus is routed by a Partitur score — a single YAML document that a Conductor plays [8] — and the score can be read. The opening of `ai-query-phase1`, the score played on every scholar’s question before any language model is consulted, declares its contract:

```
PARTITUR ai-query-phase1 / top level

partitur:
  id: ai-query-phase1
  goal: >
    Decompose a raw scholarly corpus query into elementary evidence
    (persons, places, dates, collections, intent, attested lemmata)
    BEFORE any LLM is consulted.
  runtime: pdai:duckdb 1
  metadata:
```

67.1 ms in the
by the search s
with the result
reader sees is ti
saw.

iron_rules: 2

- "Collection names are SCOPE, never lemmas"
- "No costly searches of no value – audit every run/skip"

- 1 Deterministic instrumentation: index lookups by default – the index-first commitment of Chapter 2, written as a field.
- 2 Iron rules: editorial principles recorded in the score itself. The first one kept “Kepler” out of the lemma filter in the inventory answer above.

PLAYED BY **Conductor** / **pdai:duckdb**

What the scholar experiences as an answer with a method trace *is* the performance record of this score: the trace names the movements played and the movements skipped, in the score’s own vocabulary. The two sides – the scholar’s operation and the score’s performance – are one event in two notations, and the score side has its own guide [7]. A scholar who wants to hold the platform to its iron rules can read them; that is what it means for method to be citable.

■ 4.5 What falls off the fast path

Honesty about the boundary belongs in the same chapter as the speed. Questions that need an unresolved letter register, a stale linguistic layer, or senses not yet curated fall off the fast path – and the platform’s obligation is to say so. A specialist that finds nothing returns an empty answer with its trace; the integrity rule (Chapter 3) refuses layers that lag their sources rather than serving them quietly. Empty and partial answers, honestly labelled, are part of the practice this chapter describes – the amanuensis reporting “nothing found where you told me to look” is doing its job, not failing at it.

■ 5 Operation III – The Dictionary as a Focused Instrument

The third operation is where the philosophy of Chapter 2 pays its way. A dictionary in AI Erasmus is not a general lexicon pasted onto a corpus; it is a *focused instrument* for surveying use, and the focus – the declared scope of evidence – is a control the scholar sets, the way one focuses optics. LWgrammar, the platform’s dictionary layer, writes OED-form entries from one chosen body of evidence [13], [19].

■ 5.1 The anatomy of an entry

The entry below surveys Latin *tempus* under an Euler treatise scope – one body of writing named by its Eneström number [20].

The screenshot displays the AI Viewer interface. The main window shows a manuscript page titled "DE LINEA BREVISSIMA IN SUPERFICIE QUACUNQUE DUO QUAE LIBET PUNCTA IUNGENTE" by Johann Bernoulli, dated November 1728. The manuscript text discusses the shortest path between two points on a surface. On the right, a dictionary panel for "tempus" is open, showing its Latin form, etymology, and various senses with numbered quotations from the manuscript.

**DE LINEA BREVISSIMA
IN SUPERFICIE QUACUNQUE DUO QUAE LIBET
PUNCTA IUNGENTE**

Manuscriptum Novembris 1728

Commentatio 9 indicis ENESTROEMIANI
Commentarii academiae scientiarum Petropolitanae 3 (1728), 1732, p. 110–124

1. Cuique notum est et a multis tanquam axioma ponitur lineam seu viam brevissimam a dato puncto ad aliud quodcunque esse lineam rectam. Ex hoc facile intelligitur in superficie plana lineam brevissimam duo quaelibet puncta iungentem esse rectam, quae ab altero ad alterum ducitur. In superficie sphaerica, in qua linea recta duci non potest, statuitur a Geometris viam brevissimam esse circulum maximum, qui data duo puncta coniungit.

2. Quae autem in superficie quacunque sive convexa sive concava sive ex his mixta sit via brevissima, quae ex dato puncto ad aliud quodcunque ducitur, nondum est generaliter determinatum. Proposuit mihi hanc quaestionem Celeberrimus IOHANNES BERNOULLI, significans se universalem invenisse aequationem, quae ad lineam brevissimam determinandam cuique superficiei accommodari possit¹). Solvi ego etiam hoc problema solutionemque hac disser-

Dictionary — LWgrammar

tempus la

Whole collection

Look up Build / update + compare

enestrom: 866/866 docs dated, 1725–2021

WHOLE COLLECTION

tempus la

corpus freq 5548 · 252 doc(s) · rev 1 · 9 citation(s)

FORMS

tempus (la) · tempore (la) · temporis (la) · tempora (la) · temporibus (la)

ETYMOLOGY

From Proto-Indo-European *temh₂- ('to cut'), via Latin tempus ('time, portion of time, season').

I. Physical and mathematical time as a measurable quantity, variable, or interval of duration.

I.1. The duration or interval of physical motion, especially descent along a curve under the influence of gravity. [mechanics]

Rule of use: Used as a variable to be minimized (e.g., tempus descensus, tempore brevissimo) or equated (tempore aequali) in variational calculus and mechanics, typically governed by the integral of ds/v.

1744 Euler, *E1*, p. 1, s. o SMS descensus adorsum fuerit, semper aequali tempore ad punctum SAS perveniet 1.)

1744 Euler, *E101*, p. 28, s. 3 Ut si quaerenda sit curva brachystochrona, debeat tempus descensus, quod per $\int \sqrt{\frac{ds}{v}}$ exprimitur, esse minimum.

1744 Euler, *E101*, p. 249, s. 2 Invenire curvas SBrachystochronas SAS, super quibus corpus brevissimo tempore ex SAS in SMS perveniat, dum scilicet a gravitate naturali animatum descensum in puncto SAS inchoat.

I.2. An infinitesimal increment or element of time (elementum temporis). [mathematics]

Rule of use: Collocated with 'elementum' or 'differenziale', representing the differential dt in equations of motion, often expressed as ds/v.

1744 Euler, *E101*, p. 228, s. o functioni cuicunque binarum variabilium SAS et SYS, tum $\delta V \sqrt{\frac{ds}{v}}$, tum $\sqrt{\frac{ds}{v}}$ exprimit elementum temporis ideoque formula

1744 Euler, *E101*, p. 243, s. o Quare cum celeritas in SYS sit $\delta = \sqrt{\frac{ds}{v}}$, erit temporis, quo arcus SAS conficitur, elementum

II. Temporal reference, epoch, or historical period.

II.3. A specific point in time, epoch, or moment, often in

Figure 4 The LWgrammar entry for Latin *tempus* under an Euler treatise scope: headword, senses, and dated quotations, every quotation anchored to its page. The senses survey this treatise's usage — not Latin's.

The anatomy is the OED's, kept deliberately: headword, attested forms, numbered senses, and dated quotations that link to pages. The form is not nostalgia. An entry of this shape is a perspicuous representation — the arrangement itself makes the pattern of use visible, sense by sense, before any definition is read. And the entry observes the platform's grammar: senses emerge from quotations, every quotation carries its page, and where curation is incomplete the entry says *slip*, not certainty. The scholar remains the editor of last resort; the platform drafts and retrieves — the amanuensis stance, applied to lexicography.

The instrument's controls sit in the viewer beside the page. The dictionary panel takes a term, a language, a year range, and a scope; *Look up* reads an existing entry, *Build* drafts one from the evidence, *compare* opens a second scope beside the first.



Figure 5 The dictionary panel beside a KGW page: term, language, year range, and collection scope are the instrument's controls. Look up reads an entry; Build drafts one from the scoped evidence.

OPERATION Building an entry

Q *How is tempus used in this treatise?*

1	collection-select	resolves the scope — one Eneström-numbered treatise, no other evidence admitted
2	stanza-filter	attests the headword in the lemma index and retrieves every occurrence with its sentence and page
3	lwgrammar	drafts sense branches from the gathered quotations and dates each quotation from the edition's metadata

→ a draft entry — senses with dated, page-anchored quotations, marked as draft until a scholar has read it

The box narrates what the *Build* button performs. Note what is absent: no step consults a general lexicon, and no step invents a quotation. The entry can only say what the scoped pages support — which is precisely the design.

5.2 Focused on one body of writing

The first focus is the one just shown: one treatise, one volume, or one author's works under a stable identifier. The dictionary then answers how *this* writing uses the word — the local grammar of use, portrait-sharp because nothing else contaminates the evidence. For a scholar editing or translating the treatise, this is the reference work that has never existed for any but the most-studied texts: a concordance with judgment, compiled on demand.

5.3 One person against another

The second focus is differential. Scope the entry to author A, then to author B, and read the two sense inventories side by side — same headword, different habits of collocation and sense weight. The output is not a verdict about who uses the word correctly; it is two portraits, each accountable to its own pages, and the

differences are the finding. This is where lexicography meets prosopography without collapsing into either: the persons are scopes here, not network nodes.

5.4 Meaning in time

Where imprint or letter dates exist, the scope can be a slice of time: up to a year, from a year, between years — the year-range controls in the panel above. The entry then shows persistence, narrowing, or drift across decades, without claiming to be a historical dictionary of the language. The claim is smaller and sounder: *within this dated evidence*, here is how the word’s use moved. For concept history (Chapter 8), these time-scoped portraits are the raw material.

5.5 Meaning in an academic field

The fourth focus follows the domain gazetteer: scope the entry to writing labelled astronomy, or mechanics, or theology, and the same headword yields a disciplinary portrait. A term’s senses inside one field’s writing need not match its senses in another — *momentum* in mechanics is not *momentum* in rhetoric — and the field-scoped entry makes the divergence visible instead of averaging it away.

Across all four foci the instrument is one and the anatomy is one; only the scope changes. That is the point of calling it focused: the same optics, aimed.

6 Operation IV — Prosopography and Correspondence Networks

Editions of letters make networks possible; they do not compute them. The fourth operation moves from names on pages to persons, and from persons to correspondence edges that can be queried, dated, and opened as pages. It is the guide’s most sober chapter, because its material is the most uneven — and the instrument’s honesty about that unevenness is the lesson.

Two layers must be kept strictly apart. *Mention* records who appears in whose text: the person gazetteer, applied to pages, answers “where is this name used?” — useful for reading circles, reception, and the company an author keeps on the page. *Correspondence* records who writes to whom, and it requires more than a gazetteer: it requires a resolved *letter register*, the edition’s metadata of sender, recipient, and date per letter. Mention is cheap and broad; correspondence is expensive and exact. Conflating them produces the bad prosopography in which being named in a letter and receiving one look alike.

The registers in the July 2026 build show what resolution looks like at three different scales, across three centuries of material. The von Kármán papers resolve 15,484 letters, with 34,204 person links and 19,208 place links in the letter index; Euler’s *Opera Omnia* correspondence resolves 1,296 letters [21]; the Pirckheimer letter volumes resolve 354 [22], [23]. The numbers are read from the collection indexes, and their spread is instructive: a twentieth-century archive with typed letterheads resolves richly; a sixteenth-century humanist edition resolves what its editors encoded, and no more.

OPERATION A correspondence roster

Q Who corresponds with Kármán in 1938, and where are the letters?

1	fasttext-classify	routes the question to the correspondence family
2	person-lookup	resolves the name against the person gazetteer — 5,797 canonical persons with their variant forms in the Kármán index
3	mention-in-letters	reads the resolved letter register, filters by year, and returns the roster with a page link per letter

→ a dated roster of senders and recipients, each row one click from its letter’s page in the split view

What the operation yields is deliberately modest in form: shortlists for prosopographical follow-up, correspondence intensity over time, and the one-click return to Operation I’s split view where the letter itself waits. What it must not hide are its limits, and here the platform’s honesty is structural. Gazetteer coverage

The gazetteer of von Kármán, typist’s “von Kármán” person before t

is asymmetric across the collections — the Kármán index resolves 5,797 persons where the Pirckheimer index resolves 35, and the Pirckheimer collection has no domain gazetteer at all in the July 2026 build — so the same question is answered richly in one collection and thinly in another. Ambiguous persons, undated letters, and editions whose conventions do not encode sender and recipient cleanly all leave gaps in the register. The specialist that meets such a gap returns no edges and says so; the empty roster with its trace is the honest answer, and the platform prefers it to a plausible guess, every time.

7 Operation V — Concept History and Comparison

The fifth operation composes the third and fourth, and it is short because they did the work. Concept history needs more than keyword hits: it needs senses, scopes, and comparable samples. AI Erasmus treats comparison as its own operation with one form: *fix the concept, vary the scope, read the citation sets side by side*.

The scopes are the ones the dictionary already knows — person, period, field — now taken two at a time. Where the dictionary gives a portrait within one focus, comparison gives the differential across foci: the same lemma under Kepler’s letters and under Euler’s mechanics, under writing before 1600 and after 1700, under astronomy and under theology. Where dependency indices allow, the comparison can be grammatical as well as distributional — who is *subject* of what, which verbs govern the concept in each scope [15], [18].

OPERATION A scoped comparison

Q Does inertia behave the same in Kepler’s letters and in Euler’s mechanics?

1	stanza-filter	attests the lemma in both scopes and retrieves each occurrence with sentence and page
2	collection-select	keeps the two evidence sets strictly apart — one per scope, no pooling
3	lwgrammar	assembles two citation sets side by side, each with its dated, page-anchored quotations

→ two scoped citation sets, read together — the differences are the finding, and every difference has pages behind it

The relation to Operation IV is situational rather than lexical: persons and correspondence locate *who is speaking to whom* when a concept travels through letters — the network says where the concept moved, the scoped entries say what it became there.

The operation’s ambition is deliberately bounded. The platform does not deliver a finished *Begriffsgeschichte*; it shortens the path to the evidence such a history requires, from weeks of collation to an afternoon of reading. The judgment about what the differences *mean* never leaves the scholar — Chapter 2’s first commitment, holding at the last operation as at the first.

8 The Full Scope of Application

One instrument, the same five operations — across material that spans four centuries, three languages, and every genre from printed treatise to typescript carbon. This chapter shows the breadth without breaking the guide’s own rule: the collections appear as *kinds of material* that exercise the operations differently, never as catalogues.

8.1 Four kinds of material, one instrument

Early modern Latin science. Printed treatises — Kepler’s works, Copernicus editions, Euler under his Eneström numbers — with formulae, diagrams, and the long sentences of learned Latin. This material exercises the lemma index and the dictionary hardest: the searches of Chapter 5 and the *tempus* entry of Chapter 6 live here. The editorial conventions are stable and the print is good; the challenge is the language, and the linguistic layer carries it.

Humanist correspondence. The Pirckheimer letter volumes: names, networks, and rhetoric, in Latin and early modern German. The material is Operation IV's sixteenth-century case — a resolved register of 354 letters in the July 2026 build — and its thinner gazetteer is the honest-limits example of Chapter 7. What the operations return here is correspondingly precious: humanist networks are exactly where a fast path from name to page changes working practice.

A twentieth-century archive. The von Kármán papers: 1,066 archival folders of typescript, carbon copies, and dated letterheads [23]. The material is the platform's densest network case — 15,484 resolved letters, 5,797 gazetteer persons — and it proves the operations are not tuned to early print. The same split view reads a 1938 carbon; the same roster query returns its senders.

A catalogue. Eneström's catalogue of Euler's writings [20] enters the platform not as a text to read but as a spine to hang operations from: its numbers are the stable identifiers by which treatises are scoped, cited, and compared. A catalogue, treated as an index layer, becomes infrastructure — the platform's quiet homage to the bibliographers.

■ 8.2 The corpus at the time of writing

The five collections measured in the July 2026 build, read directly from the collection indexes:

Collection	Documents	Pages	Tokens	Named entities
Kepler (KGW)	26	14,277	5,646,156	117,232
Euler Opera Omnia (Series I–III, works)	766	27,047	5,444,262	109,453
Euler Opera Omnia (Series IVA, correspondence)	8	6,259	2,183,255	—
Pirckheimer	2	662	145,864	4,829
von Kármán papers	1,066	38,024	8,691,066	423,016
Total	1,868	86,269	22,110,603	654,530

: The July 2026 corpus snapshot: five collections, one substrate. Token counts equal lemma-index rows (one row per token); the Copernicus print collection shown in Chapter 5 was in ingest at the time of the audit. {#tbl-corpus}

?@tbl-corpus is a snapshot with provenance, not a boast: roughly 22.1 million lemma-indexed tokens over 86,269 pages, every one addressable by the operations of this guide. The table's real argument is the *rows' variety* — a critical edition, a catalogue corpus, a correspondence edition, a letter book, an archive — all served by the same five operations without re-ingesting anything per tool.

The instrument's front door shows the same breadth in working form: the document browser lists every collection's documents with their page counts, one click from any document's split view.




Figure 6 The document browser: the platform's collections listed with page counts, each document one click from its viewer. The lemma search of Chapter 5 is the door at the top.

8.3 Teaching and collaboration

The operations transfer to the seminar room because they are inspectable. A dictionary scope is a file a student can be handed; a query with its method trace is an argument a class can criticise; a score's iron rules (Chapter 5) are a reading assignment about method. A seminar can hold identical evidence with the professor — same pages, same indices, same traces — which was never true of the archive-visit model of training. And because every answer ends in a page link, teaching with the platform is still teaching with sources, not with summaries of sources.

8.4 What AI Erasmus does not do

The negative scope is part of the scope. The platform does not produce editions: apparatus, stemmatics, and emendation remain editorial crafts, and the platform consumes their results. It does not replace the printed volume as scholarly object; it cites into it. It does not deliver finished interpretation — no operation ends in a verdict, only in evidence arranged for one. And it does not answer beyond its indices: where the register is unresolved or the layer is stale, the honest empty answer of Chapter 5 is the designed behaviour. The workshop drafts, fetches, and counts. The desk still judges.

9 Implications for Scholarly Practice

If the five operations become ordinary, the craft changes in specific places, and it is worth saying which.

Reading stays primary — and gains a shorter supply line. The split view does not replace slow reading; it removes the dead time between a question raised at the page and the evidence that bears on it. The scholar's attention returns to where the workshop always wanted it: on the page, with the fetching done by other hands.

Method becomes shareable. A dictionary scope, a query with its trace, a score with its iron rules — these are files, and files can be mailed, versioned, cited, and refereed. Scholarship has always demanded that conclu-

sions be criticisable; the platform extends the demand to the *instruments*, because an inspectable method is one a colleague can disagree with precisely.

The scholar acquires an editorial relation to the instrument itself. Scopes are theirs to declare, gazetteers theirs to correct, and — through the score language — budgets and iron rules theirs to argue [7]. The division of labour is Erasmian: the platform is staff, and staff are directed.

The risks are named rather than waved at. OCR ceilings bound every derived layer, and early modern German and abbreviated Latin still test transcription and lemmatisation. Gazetteers are uneven across collections — Chapter 7 printed the asymmetry. And fast answers invite fast belief: an answer that arrives in 67 milliseconds will be trusted in 67 milliseconds. Every one of these risks is met by the same design answer, provenance — the trace that makes speed accountable — which is why Chapter 2 called provenance the condition of trust rather than a feature.

10 Goal Expansion and the Famulus Handlungsmodell

This chapter steps behind the operations the previous chapters demonstrated and describes the action model — the *Handlungsmodell* — that animates them. Every workshop move Erasmus’s secretary performed was an intelligible step directed toward a goal: retrieving a letter, copying a passage, looking up a word. The same structure governs AI Erasmus. Understanding it matters both for scholars who want to know what the platform is actually doing and for agents that use the platform as a substrate for their own reasoning.

10.1 Goals and Partiturs

A **goal** is a natural-language statement of what is to be achieved: “*count the letters Erasmus exchanged with Pirckheimer*”, “*find the term Kepler uses for the force driving planetary motion*”. A **Partitur** (from the German *Partitur*, a musical score) is a YAML document that specifies how to pursue a goal: a sequence of stages, each with a runtime (a DuckDB query, a GLiNER extraction pass, an LLM call), an evaluator that grades the result, and a goal contract stating what a successful outcome must contain.

The iron rule governing the platform:

Every goal corresponds to an agent playing a Partitur.

Python functions are never orchestration steps. They appear only as leaf-stage runtimes embedded in a Partitur — callable tools that a stage invokes, never the stage itself. This separation is not a stylistic preference. It makes every action inspectable (the Partitur is the record), patchable (the Partitur is the patch target for the learning loop), and composable (a new goal is served by selecting and binding an existing Partitur, not by writing new glue code).

10.2 The Action Repertoire

The famulus holds a **repertoire** of registered Partiturs — scores that have been drafted, tested against the golden set, and promoted to active status. At the time of writing the repertoire contains thirty-seven entries, covering:

- **Filter construction** (`ai-query-phase1`): the stage-one Partitur that decodes a query, grounds its entities against the archive, and builds the minimal evidence set.
- **Correspondence register** (`ai-query-phase1` with `correspondence-roster` specialist): answers letter-count and correspondent-list questions from `letters_resolved` without retrieving a single sentence.
- **Terminology finder**: identifies technical terms for a concept across a corpus via lexical search and LLM verification.
- **Manuscript inventory, entity roster, concept statements**: further specialisations each matched to a declarative result contract.
- **Planning** (`plan-composer-agent`), **synthesis**, **Partitur generation**: meta-level scores that reason about the repertoire itself.

10.3 Iterative Goal Expansion

A query rarely maps to a single Partitur. The platform resolves it through **iterative goal expansion**: a loop, one round per open subgoal, that binds each subgoal to the best-fitting registered Partitur before running it.

Round zero expands the raw query into the filter-construction subgoal and binds it to `ai-query-phase1`. After that Partitur plays and the evidence set is assembled, the context accumulated so far — the grounding card, the filter set, the collection memory, any prior failure rails — is richer than at the start. The next expansion round reads all of it and selects the specialist Partitur from narrower candidates, guided by what the evidence has already established.

This is the crucial asymmetry with a static pipeline: each successive round benefits from the evidence the prior round produced. Expansion becomes progressively more **guarded and rail-directed**: hard rails from grounding (has this collection a letter register? does the person exist in `ent_canonical`?) prune infeasible Partiturs before they play, so the later rounds almost never face an open OR tree.

10.4 Repertoire Growth — the Generation-Partitur Loop

When no registered Partitur fits a subgoal, the famulus does not improvise Python code. Instead it invokes the **generation-Partitur loop** — itself a scored agent — which drafts a new minimal score, registers it as a candidate in the version registry, and subjects it to the same A/B gate every instruction patch must pass (the baseline is measured *before* the draft, a discipline enforced since the ordering bug that once compared a patched state against itself). Only after promotion does the new Partitur enter the repertoire and become available for future expansion.

The generation loop is the platform's growth mechanism. Each new kind of question the archive raises adds one registered score to the repertoire; each subsequent query of the same kind finds its score ready.

10.5 Failure Diagnosis as Repertoire Refinement

When a goal-check FAIL occurs, the **failure-cause analysis agent** — its own scored Partitur — diagnoses the cause and proposes a corrective rail. Causes are typed:

- *Wrong anchor* — the collection the grounding step selected was not the right one.
- *Over-narrow filter* — the AND logic in the filter was too strict and excluded relevant evidence.
- *Missing capability* — the bound Partitur assumed `letters_resolved` where the collection has none.
- *Unresolved variant* — a person span in the query was not found in `ent_canonical` or its surface-variant list.
- *Gap* — the archive genuinely does not contain this data; the honest answer is a capability-gap report, not a failed search.

The diagnosed cause serves two functions. First, it is injected into the current run's accumulated context so the *next expansion round* avoids the same branch. Second, it is written to the collection's portable memory (`agent_context.json`, `learned.failure_causes`) so that *future runs* on the same collection start with that rail already in place.

Together, successful bindings (written on PASS) and failure rails (written on FAIL) make the collection memory a living record of what the platform has learned about each archive — portable, inspectable, revisable by the scholar who knows what the archive actually contains.

10.6 The Philosophy of the Handlungsmodell

The Handlungsmodell is a specific answer to a specific philosophical question: what kind of action can an AI agent perform in a scholarly context without overstepping the amanuensis stance?

The answer the platform gives is: *an action whose specification is public, testable, and revisable*. A Partitur is the specification. The golden set makes it testable. The learning loop makes it revisable. A famulus that acts within these constraints cannot silently substitute one thing for another (the silent-substitution bug this platform was designed to eliminate), cannot invent bibliographic data, and cannot claim to have answered a question the evidence does not support.

The amanuensis stance from Chapter 2 and the Handlungsmodell of this chapter are the same commitment at two levels of description. The famulus acts only within its repertoire; the repertoire grows only through a gated, publicly traceable process; and every step of every action is on the record.

11 Conclusion and Availability

This guide opened in a Basel printing house and has argued, at book length, that the workshop is buildable now. AI Erasmus is that argument made operational: a platform that treats editions as the substrate of inquiry, keeps judgment at the desk, and staffs the repeatable moves — reading with an edition, querying with a visible method, consulting a focused dictionary, tracing persons and correspondence, comparing concepts across scopes — with specialists whose every answer ends on a page.

The philosophy is what makes the engineering scholarly. The amanuensis stance decides what counts as a good answer; the page as evidence gives disagreement an address; meaning as use makes the dictionary honest; index-first answering makes failure visible; provenance makes speed trustworthy. None of these commitments is decoration on the platform — each one rules something out, and the chapters have shown the refusals working.

AI Erasmus runs within the AI Platform deployment, with the viewer as the primary workplace; the companion papers document the PDAI substrate [6], the score language and the Famulus Handlungsmodell detailed in the preceding chapter [7], [8], and the golden-query testbed that holds every answer to a declared oracle — thirteen canonical questions, count and floor oracles, monitor log, learning loop [9]. Questions, their checks, and the development of the famulus are written down in notebooks whose blocks declare scope, plan, and oracle, and whose logs keep trace, verdict, and artefacts; a notebook can read from and run on other machines [24]. The corpus will grow — the Copernicus collection was in ingest as this guide went to press — and the operations vocabulary is meant to stay stable while it does. Planned work is concrete: richer sense curation in LWgrammar, tighter letter-register resolution for the early modern collections, and comparative views that keep the page one click away. The editions remain the ground. The workshop, at last, is staffed.

12 Acknowledgements

Thanks are due to the archives and editorial projects whose editions underpin the illustrations, and to collaborators in the AI Platform project for index, viewer, and query infrastructure.

Bibliography

- [1] F. Bierlaire, *La familia d'Érasme: contribution à l'histoire de l'humanisme*, no. 17. in *De Pétrarque à Descartes*. Paris: Vrin, 1968.
- [2] J. Huizinga, *Erasmus*. New York: Charles Scribner's Sons, 1924.
- [3] K. Reich, “Philipp Melanchthon im Dialog mit Astronomen und Mathematikern. Ausgewählte Beispiele,” in *Mathematik und Naturwissenschaften in der Zeit von Philipp Melanchthon*, no. 26, F. Fuchs, Ed., in *Pirckheimer-Jahrbuch für Renaissance- und Humanismusforschung*, Wiesbaden: Harrassowitz, 2012, pp. 27–57.
- [4] M. Maruska, “Johannes Schöner — ein angesehener Wissenschaftler und genialer Wissensvermittler,” in *Mathematik und Naturwissenschaften in der Zeit von Philipp Melanchthon*, no. 26, F. Fuchs, Ed., in *Pirckheimer-Jahrbuch für Renaissance- und Humanismusforschung*, Wiesbaden: Harrassowitz, 2012, pp. 121–161.
- [5] J. Kepler, *Gesammelte Werke*. München: C. H. Beck, 1938.
- [6] G. Graßhoff, “PDAI: A Certified Document Infrastructure for Computational Historical Scholarship,” 2026.
- [7] G. Graßhoff, “The AI Partitur — A Score Language for the Humanities,” 2026.

- [8] G. Graßhoff, “The AI Agent Orchestrator: Partitur, Conductor, and Harvest.” 2026.
- [9] G. Graßhoff, “Golden Queries in AI Erasmus: A Testbed for Scholarly Inquiry,” 2026.
- [10] S. Shapin, “The Invisible Technician,” *American Scientist*, vol. 77, no. 6, pp. 554–563, 1989.
- [11] D. A. Grier, *When Computers Were Human*. Princeton: Princeton University Press, 2005.
- [12] L. Wittgenstein, *Philosophical Investigations*. Oxford: Blackwell, 1953.
- [13] Oxford English Dictionary, *The Oxford English Dictionary*, 2nd ed. Oxford: Oxford University Press, 1989.
- [14] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, and C. D. Manning, “Stanza: A Python Natural Language Processing Toolkit for Many Human Languages,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, Association for Computational Linguistics, 2020, pp. 101–108. doi: 10.18653/v1/2020.acl-demos.14.
- [15] J. Nivre and others, “Universal Dependencies v2: An Evergrowing Multilingual Treebank Collection,” *Proceedings of LREC*, 2020.
- [16] M. Raasveldt and H. Mühleisen, “DuckDB: An Embeddable Analytical Database,” in *Proceedings of the 2019 International Conference on Management of Data (SIGMOD)*, ACM, 2019, pp. 1981–1984. doi: 10.1145/3299869.3320212.
- [17] U. Zaratiana, N. Tomeh, P. Holat, and T. Charnois, “GLiNER: Generalist Model for Named Entity Recognition using Bidirectional Transformer,” in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*, Association for Computational Linguistics, 2024.
- [18] G. Graßhoff, “Querying with AI Linguistic Trees: Stanza-Backed Dependency Parses, Lemma Normalisation, and DuckDB,” 2026.
- [19] G. Graßhoff, “A Multi-Layer Index and Retrieval Architecture for Historical Document Collections,” 2026.
- [20] G. Eneström, *Verzeichnis der Schriften Leonhard Eulers*. in Jahresbericht der Deutschen Mathematiker-Vereinigung, Ergänzungsband~IV. Leipzig: Teubner, 1910.
- [21] L. Euler, *Opera Omnia, Series~IV: commercium epistolicum*. Basel: Birkhäuser / Springer, 1975.
- [22] W. Pirckheimer, *Briefwechsel*. München: C. H. Beck, 1940.
- [23] California Institute of Technology Archives, “Theodore von Kármán Papers: Finding Aid.”
- [24] G. Graßhoff, “Famulus Notebook: The Notebook as the Principal Medium of Scholarly Querying, Checking, and Development, Federated Across Machines,” *arXiv preprint*, 2026.