

LORACRAFT PROJECT

From a photograph of an inscription to an evidence-based translation



To Bernard Mathieu,
*a world-renowned Egyptologist and an
exceptional teacher, who passed on far
more than the meaning of the signs,
and guided me so well when the path
became uncertain. Thank you.*



Montpellier, 18 August 2026.

01 A book to explain the project

The first document¹ published on the **Loracraftt** project is a reference technical architecture. It describes in detail the data objects, interfaces, formats and rules that allow several software applications to work together. This precision is useful to developers, but does little to answer a simpler question: what are we actually trying to do with a photograph covered in hieroglyphs?



¹ Scientific Architecture of **Loracraftt**, Didier DTL Morandi, 17 August 2026.

The relief reproduced on the previous page belongs to the White Chapel of Pharaoh Senusret I (12th Dynasty), now reconstructed in the open-air museum at Karnak, Egypt. On the left, Senusret I, wearing the Red Crown of Lower Egypt, faces the god Amun², the principal deity of the Karnak sanctuary. Amun holds a group made up of the ankh and djed signs to the king's nose: S34-R11, "life and stability". In their lowered hands, the god and the pharaoh each hold a single ankh. The scene depicts a divine gift: Amun grants or renews life and stability for the king. The figures, royal cartouches and inscriptions are meticulously carved in raised relief, using the particularly refined technique characteristic of the White Chapel. © Google AI.

"What are we actually trying to do with a photograph covered in hieroglyphs?" This document answers that question. It follows an inscription from its arrival in the software through to its translation, explaining the difficulties as they arise. Whenever a specialist term is essential, it is defined in the sentence. Code details remain in the technical volume.

IN PLAIN TERMS **Loracrafft** is not a program that looks at a stela and instantly announces what it says. It is a digital workshop: it prepares the image, segments the composition, proposes identifications, preserves disagreements and helps the researcher build a verifiable reading.

Five test images provide the running examples (see the appendix). Some are straightforward; others are damaged, vertical or photographed at an angle. They do not yet form a representative scholarly corpus. They serve to reveal errors in the prototype and ensure that one improvement does not undo another.

The cover image and the text of this document were produced by © OpenAI ChatGPT 5.6 and reviewed and corrected by the author (in particular, all the footnotes).

² Amun is the tutelary god of Thebes and the principal deity of the temple of Karnak. This is not the god Min, contrary to the identification supplied by Google AI.

02 Two volumes, two readerships

This volume

It is intended for an educated reader who does not spend their time programming in Python or practise Egyptian epigraphy every day. Here we discuss photographs, signs, words, evidence and decisions. The aim is to understand the journey of an inscription and the prototype's actual state of development.

The technical volume

The Scientific Architecture of *Loracraft*, by contrast, remains the developers' reference. It contains the data schemas, unique identifiers, programming interfaces, JSON formats, versioning rules and precise contracts between modules. That volume is not being discarded: it becomes the operating manual for machines and technical contributors.

WHY DOES IT MATTER? Mixing the two levels produces either a vague text that cannot guide software development or an accurate text that is unreadable to the person commissioning it. Separating them preserves rigour without forcing the reader to decipher every paragraph or read three hundred footnotes.

The two books complement each other. When a design decision needs to be implemented, consult the technical architecture. When the project needs explaining, a priority needs setting or a result needs presenting, return to this volume.

03 The complete pipeline in ten steps

1. Receive the photograph and keep the original intact.
2. Create a working copy, then adjust its orientation and contrast if necessary.
3. Identify the inscribed surface, registers, lines or columns, and then the quadrats.
4. Determine the direction in which the signs should be read, horizontally or vertically.
5. Isolate the visual objects and, where possible, propose codes from Gardiner's sign list.
6. Submit uncertain proposals for human review.
7. Represent the signs and their groupings using the Manuel de Codage.
8. Transliterate, segment the words and propose grammatical analyses.
9. Look for attestations and parallel passages in corpora, where access permits.
10. Develop and publish a translation linked to its evidence and variants.

GROUND RULE Any stage may answer “I don't know”. That answer is better than an invented sign, a gap filled without warning or a translation that conceals its uncertainties.

04 Why hieroglyphs resist automatic recognition

The problem

An inscription is not a printed page with neatly aligned characters. Its surface may be broken, worn, reused or lit from an oblique angle. A carved feature may be mistaken for a crack. An inscription may be carved in sunk or raised relief. Two signs may touch; a single sign may be broken into several pieces. Nearby decoration may resemble a hieroglyph without belonging to the text.

The difficulty continues beyond the image. Signs may be grouped in a quadrat, the small square or rectangular space in which several signs are arranged together. The text may run from right to left or from left to right, in lines or in columns read from top to bottom. Honorific transposition sometimes places a divine name first without changing the expected grammatical order.

EXAMPLE A bird generally faces the beginning of the line. This clue helps establish the reading direction, but a symmetrical scene may contain figures facing both ways. The clue then becomes a proposal, not conclusive evidence.

Finally, recognising a shape is not enough to understand its function. The same sign may represent a sound, clarify a word's meaning or form part of a particular spelling. That is why **Loracrafft** separates what is seen, what is recognised and what is interpreted.

05 Seeing is not yet reading

Our approach

Loracrafft assigns each piece of information to the right level. A dark area or a contour belongs to observation. “This contour could be sign G17” is a hypothesis. “Here, this sign has the value *m*” is linguistic analysis. The word chosen in French belongs to the translation.

This separation may seem meticulous, but it prevents an early error from becoming invisible. If an identification changes, the program can show which analyses and French sentences depend on it. The image and contours, meanwhile, are not rewritten to support the new reading.

IN PLAIN TERMS The program does not replace an old proposal with a new one as though the first had never existed. It adds the new proposal and links the two. The researcher can therefore reconstruct the history of the analysis.

Expected outcome

The final result is more than a translation. It is a record in which each translated segment can be linked to the signs, image regions, validations and parallels that support it.

Current status

This principle is already built into *LoraDocument 1.0*: layout objects, sign hypotheses and human validations are recorded separately. The complete linguistic pipeline has yet to be built.

06 What happens to a photograph when it enters Loracraftt

The problem

A photograph can be renamed, cropped or compressed, leaving no way to know which version was used for an analysis. Yet a tiny difference in lighting or framing can change the detected contours. The original file must therefore remain identifiable.

Our approach

On import, **Loracraftt** preserves the original, records its dimensions and available information, documents its rights and calculates a digital fingerprint. This fingerprint acts as a signature for the file: if a single byte changes, the fingerprint changes too. Several views of the same monument can be linked without being confused with one another.

EXAMPLE An overall view locates the inscription on the monument; a photograph taken in raking light reveals incised lines; a high-resolution detail helps examine a sign. These are three different files attached to the same scholarly record.

Expected outcome

Another person must be able to retrieve the exact source used, understand its access conditions and identify the working images derived from it.

Current status

The core already preserves image provenance in the local workflow. The full acquisition module, with catalogues, licences and systematic checks, remains part of the target architecture.

07 The original never changes



test_5: a real photograph of a stela, used to test automatic straightening (Karnak).

Loracraftt works on copies. One version may be straightened, another contrast-enhanced and a third prepared for a particular engine. Each retains a link to the original and the parameters used to produce it.

WHY DOES IT MATTER? A more striking image is not necessarily more faithful. Excessive contrast can erase a fine line; noise reduction can remove an ancient trace. Preserving each stage makes comparison and reversal possible.

08 Straightening the image before segmentation

The problem

If the program looks for lines and columns in a tilted photograph, it creates incorrect geometry. Register boundaries become diagonal, quadrats overlap and sign crops are poorly framed. Recognition cannot repair inconsistent segmentation.

Our approach

The prototype first detects straight line segments visible in the image. It then estimates the dominant orientation while allowing for the fact that an inscription may be arranged vertically. In everyday terms, it looks for the overall tilt without assuming that all text must be horizontal.

EXAMPLE In an early version of test_5, the measured tilt was approximately 44.7 degrees, which was incorrect. An initial version had begun segmentation too early. Once straightening was moved upstream, the stela regained a usable vertical orientation.

Expected outcome

The program provides a straightened working image and records the detected angle, the angle actually applied and the transformation needed to return to the original photograph's coordinates.

09 Straightening, perspective and contrast: three different operations

Straightening

This corrects rotation: the monument has been photographed at an angle, but its plane remains approximately parallel to the sensor.

Correcting perspective

This compensates for a trapezoidal distortion: one edge of the monument is farther away than the other. The operation requires the four corners of the relevant surface to be known or proposed.


Improving legibility

Contrast is adjusted locally, or certain kinds of noise are reduced, to make faint relief more visible. The result is an aid to analysis, never a new source.

CURRENT LIMITATION Automatic straightening is operational on the case tested with the currently available prototype. Fully automatic detection of the four corners and perspective correction are not yet operational. A broken edge, decoration or modern feature can mislead the estimate; visual checking remains essential during development.

Several prepared versions can coexist. A contour-detection engine does not necessarily need the same image as a human reader. The record preserves these variants without confusing them.

10 One inscription, two appearances



test_2: a modified³ version of the following image, bringing out the relief and structural boundaries.



*test_3: a natural-looking version of the same type of inscription, arranged in two registers
(Karnak, White Chapel of Senusret I).*

KEY POINT A prepared image may help an algorithm without being better for every purpose. **Loracraft** preserves the source image, its derivatives and the transformation recipe so that we always know what we are looking at.

³ All image modifications were made using OpenAI's ChatGPT 5.6 software, following the author's instructions.

11 Identifying the structure before the signs

The problem

An inscription has an organisation of its own. It may contain several registers arranged one above another, columns, lines, cartouches and damaged or missing areas. If every small mark is segmented immediately, that organisation is lost and the result is a set of fragments without context.

Our approach

Loracrafft works from large to small: the inscribed surface, registers or broad bands, lines or columns, quadrats, and finally signs. This is known as a “registers first” strategy. It was adopted after a global segmentation approach failed on a two-line inscription.

IN PLAIN TERMS Before looking for the letters on a page, we identify its columns, paragraphs and lines. For hieroglyphs, the exact sequence differs, but the idea is the same: understanding the layout prevents neighbours that should remain separate from being mixed together.

Expected outcome

The program produces a geometric map: each region has a shape, a position, a proposed type and relationships with neighbouring regions. No Gardiner code is assigned at this stage.

12 Why a mark is not a sign

The experiment that changed the plan

An early method isolated connected components: groups of pixels that touch one another. It implicitly assumed that one group corresponded to one sign. Real images disproved that assumption.

A carved sign may be interrupted by wear or a break and form several components. Conversely, two neighbouring signs may touch and form just one. An experimental merging process reduced 173 components to 56 groups, but it still merged some neighbouring signs and left some fragments separate.

EXAMPLE An animal that belongs to the decoration in the image may have a perfectly clear shape without being a hieroglyph. Forcing it into Gardiner's list would produce a technically confident but scientifically absurd answer.

Our approach

The system now distinguishes probable signs, groups, fragments, noise and unknown objects. Valid motifs that do not belong to the text must also be flagged. These categories guide the review instead of forcing every shape through a recogniser.

Current status

The prototype correctly detects the two registers in test_2, but it does not yet delimit all the actual hieroglyphic units. Segmentation remains the main obstacle in phase 1.

13 The vertical case: retaining context



test_4: a damaged vertical inscription, useful for testing columns, cartouches, fragments and reading order (Pyramid of Unas — © Morris Franken).

A vertical image reminds us why the conventions of a modern page should not be imposed. Column dividers, figures with a discernible orientation and cartouches provide clues that must remain linked to their locations.

WHY DOES IT MATTER? An isolated crop may resemble several signs. Its column, neighbours and position within a group sometimes reduce the ambiguity. Context must therefore not be discarded during segmentation.

14 Finding the reading direction

The problem

As we know, hieroglyphs may be arranged horizontally and read from right to left or from left to right, or they may be arranged vertically. In a vertical text, signs are read from top to bottom, but the order in which the columns should be read still needs to be determined.

Reading direction is indicated primarily by hieroglyphs depicting living beings: people, birds and other animals almost always⁴ face the beginning of the text. The reader therefore reads towards them. The arrangement of lines, columns, quadrats and dividers provides further clues.

A scene may, however, contain several independent inscriptions. A ruler and a deity facing each other may, for example, be accompanied by texts oriented in opposite directions.

A single direction should therefore not be imposed on the whole composition: each inscribed area must first be delimited and studied separately.

Our approach

Loracrafft begins by identifying the different text regions in the image. For each one, it examines the orientation of signs depicting living beings, the arrangement in lines or columns and the organisation of signs within the quadrats.

The software then proposes a reading order for each inscription.

If a section is erased, damaged or difficult to distinguish, it can retain several hypotheses. Their consistency is checked against vocabulary, grammar and comparable texts from other monuments.

⁴ There are, however, rare inscriptions known as retrograde inscriptions, in which the signs face the end of the text. The arrangement of lines, columns and quadrats, together with the grammar, can then confirm or correct the initial hypothesis (see Henry George Fisher, <https://www.metmuseum.org/fr/met-publications/the-orientation-of-hieroglyphs-part-1-reversals>).

EXAMPLE In a scene where a pharaoh and a god face each other, the inscriptions in front of them may be read in opposite directions. The figures' orientation alone is not enough. **Loracraftt** focuses on the signs that make up each text, then handles the god's words, the ruler's titles and the offering formula separately.

15 Proposing a Gardiner code without inventing one

The problem

Egyptologists commonly use Gardiner codes such as G17 or N35 to identify signs in a reference list organised into broad categories. These codes are useful, but a worn or poorly segmented shape may match several entries.

Our approach

For each object, *LoraGlyph* can produce a shortlist of candidates. Each candidate specifies the reference catalogue consulted, the engine used and the type of score. If there is no credible candidate, the list remains empty and the object is sent for review.

IN PLAIN TERMS The program does not say, “This sign is G17.” It says, “For this crop and with this catalogue, G17 is a possibility ranked ahead of two others, with this score.” Scholarly judgement comes afterwards.

Expected outcome

The software sees the crop, the region in the full image, the candidates and the reasons for their ranking. It can accept, reject, correct, merge fragments or request new segmentation.

Current status

LoraGlyph 1.0.1 already produces crops and hypotheses that are separate from the geometric objects. Recognition quality has not yet been scientifically validated.

16 Comparing shapes: useful but fragile

The first experiments

The prototype compares silhouettes, contours, pixel distributions and seven numbers that summarise the overall shape. These are often called *Hu moments*⁵. They help identify a shape despite changes in size or position.

None of these measures is sufficient on its own. A poorly segmented contour gives a poor comparison; two different signs may have similar silhouettes; a change in stroke thickness or framing can lower the score of an otherwise obvious example.

EXAMPLE A test on sign G1 incorrectly returned a recognition score of 0.748% with a raw comparison, although the case looked visually straightforward. This result rules out claiming that a single threshold, such as 0.75, can distinguish correct from incorrect matches in every case.

Next steps

Several methods will be evaluated separately and may then be combined. Developing a neural network is not a priority until the reference images are properly segmented and the corpus is organised into reliable training and evaluation sets. Trained on poor crops, such a model would mainly learn to reproduce our errors.

PRINCIPLE A powerful engine cannot compensate for poor geometry or a poorly constructed corpus.

⁵ Ming-Kuei Hu (1918–2002) was a Chinese-American engineer and professor at Syracuse University. In 1962, he published a method for describing shapes using seven mathematical values, now known as Hu invariant moments. These values allow silhouettes to be compared independently of their position, size and, to a large extent, orientation.

17 The 6,722 glyph files: a library, not 6,722 signs

The available resource

A local collection drawn from Raymond Monfort's resources⁶ contains 6,722 glyph files. This figure counts files and graphic variants. It does not mean that Gardiner's list contains 6,722 distinct signs.

The work required

Each file needs an inventory entry: associated code, variant, format, dimensions, fingerprint, source, licence and possible duplicates. Freely usable shapes, references available for consultation and resources that cannot be redistributed must be distinguished.

WHY DOES IT MATTER? A large catalogue may give an impression of broad coverage while containing many variants of the same sign, misnamed files or uncertain rights. Quantity is no substitute for proper documentation and assessment.

Current status

The files are available for experimentation, but the scholarly and legal manifest remains to be produced. Alongside them, a set of 160 real scans could support evaluation, provided that all views of the same inscription remain in the same dataset partition.

⁶ <https://papocle.com/RaymondMonfort/>

18 The computer proposes, the researcher decides

The problem

A human correction must neither erase the automatic hypothesis nor automatically become a universal truth. Two specialists may examine the same region and disagree. Their disagreement contains useful scholarly information.

Our approach

Each review action becomes a dated event: accept, reject, correct, mark as illegible, merge, split or defer. The action specifies its target, author, degree of certainty and, where possible, justification.

EXAMPLE An engine proposes G43. A first annotator accepts it with reservations; a second proposes G17 and justifies the choice through context. **Loracraftt** preserves both opinions. An editorial decision may select G17 for a published version without erasing G43 or the first opinion.

Expected outcome

Review is a contribution whose history is preserved. Agreement between annotators can be measured, but the number of votes does not replace the argument.

Current status

The multiple-validation model exists in *LoraDocument 1.0*. An earlier console-based review tool already writes a correction log. The visual interface for redrawing, merging and splitting regions remains a priority.

19 What the review screen should show

Full image. It preserves the monumental context and the sign's position.

Proposed region. Its crop remains visible and its outline can be corrected.

Working images. They assist examination without hiding the original photograph.

Gardiner candidates. The screen shows their ranking, scores and the engine that produced them.

Previous decisions. Disagreements and justifications remain available for consultation.

Review actions. Accept, reject, correct, merge, split, mark as illegible or defer.

IN PLAIN TERMS The priority is not a “translate” button. It is a screen where the segmentation of test_2, for example, can be properly corrected and that correction recorded without overwriting the past.

A validation does not automatically enter the training data. It must first be checked, carry compatible rights and be assigned to a partition that does not contaminate the tests.

20 From a validated sign to the Manuel de Codage

The problem

A linear list of codes is not always enough to represent an inscription. Within a quadrat, signs may be placed above, below or beside one another. This arrangement must be preserved and linked to the visible objects.

Our approach

The Manuel de Codage, abbreviated to MdC, provides conventions for representing signs and their groupings. **Loracrafft** builds a structured representation: each element of the encoding will remain aligned with the visual object or objects from which it derives.

EXAMPLE If two signs are arranged one above the other in a quadrat, the encoding must indicate their relationship. If a section is missing, the gap remains marked. If two orders are possible, two branches can coexist.

Expected outcome

The researcher can display the image, quadrats, MdC encoding and a hieroglyphic rendering side by side. Correcting a sign changes only the representations that depend on it.

Current status


LoraMdC remains part of the target architecture. Its development will follow the stabilisation of segmentation, reading order and sign validations.

21 Transliteration is not translation

Transliteration

Transliteration represents the selected sign values using specialised Latin characters. It facilitates linguistic analysis, but it already reflects choices: phonetic value, complement, determinative, group order and any proposed restoration.

Example:



English translation of the image text: "King N adores Re at dawn as he emerges, when he breaks through his clay and climbs into the sky as Khepri."

(Book of Gates, sarcophagus of Seti I, 19th Dynasty; transcription, transliteration and translation by Raymond Monfort, <http://raymond.monfort.free.fr>).

Word segmentation

Inscriptions do not mark word boundaries as modern French does. The program must propose where each word begins and ends. Two segmentations may lead to two different analyses.

Analysis

For each proposed word, **Loracraft** will seek a lemma (the dictionary form), a grammatical category, morphological features and syntactic relationships. Each proposal will refer to an identifiable rule or resource.

IN PLAIN TERMS Moving from Gardiner codes to French is not a single leap. Between them lie order, grouping, values, transliteration, words, translation, grammar and context.

Current status

The grammar module has not yet been implemented. The lexicon module is being validated: a SQL database of 7,830 entries compiled from the 2019 edition of Yvonne Bonnamy's dictionary. The technical document describes their contracts so that they can be built without confusing dictionary information, hypotheses and editorial decisions.

22 Carrying uncertainty all the way into French

The problem

An uncertain sign can change a word; that word can change the grammar; the grammar⁷ can change the translation. If uncertainty disappears halfway through the pipeline, the French sentence appears more certain than the evidence allows.

Our approach

Each result declares the elements on which it depends. When two readings remain plausible, the program can create two branches. A synthesis compares them and states exactly what changes.

EXAMPLE Reading A produces a well-attested construction. Reading B requires a rare spelling but remains compatible with the damaged area. The edition may prefer A, retain B as a variant and explain why.

Gaps and restorations

A destroyed area is a material observation. The text proposed to fill it is a restoration. The two must never be recorded as a single thing. In the export, restored text remains marked and linked to the arguments supporting it.

WHY DOES IT MATTER? Caution does not prevent conclusions. It simply requires us to show where observation ends and reconstruction begins.

⁷ The current prototype uses a blend of Champollion 1836, Gardiner 1957 and Malaise–Winand 1999, with a (generous) pinch of Grandet–Mathieu.

23 How the TLA can help

The TLA's place in our pipeline

The **Thesaurus Linguae Aegyptiae** is the world's largest electronic corpus of lemmatised Egyptian texts. It provides access to texts, lemmas, their attestations and their contexts. It is not a service to which a photograph can be sent for direct sign recognition. It comes into play after the initial visual and linguistic hypotheses.

Our approach

Loracraft prepares a query from possible transliterations, sign sequences, words, periods and object types. A connector will query only officially accessible services or data and will preserve the resource version, consultation date and external identifiers.

EXAMPLE An uncertain sequence might correspond to a formula found on a Middle Kingdom stela. A search retrieves several attestations. These passages become material for comparison; they do not automatically turn the hypothesis into a fact.

Expected outcome

Each proposed parallel states what was compared, why it occupies that place in the ranking and where the cited passage can be consulted, subject to rights restrictions.

24 A parallel is a clue, not a seal of truth

The risk

Finding a very similar formula can strengthen a reading, but the absence of a result does not prove that the reading is impossible. The corpus may be incomplete, the spelling different or the query too narrow.

There is also a risk of circular reasoning. A modern edition may already have used the inscription under study to reconstruct a passage. Presenting that edition as independent evidence would amount to asking the text to confirm itself.

Our approach

The system distinguishes the work, its physical witness, the passage and the edition. It preserves the genealogy of sources, index versions and normalisation rules. Earlier results are not rewritten when a corpus changes; they are flagged as having been calculated with an earlier version.

CURRENT LIMITATION *LoraCorpus* and the TLA connector have yet to be built. Before large-scale indexing, rights, normalisation methods and a policy for specialist evaluation must be established.

Expected outcome

The researcher receives a list of comparable passages, not a verdict. They can check the source, discard a proposed parallel and explain why another is relevant.

25 Laying out the arguments

The problem

At the end of the pipeline, several clues may converge or contradict one another: sign shape, reading order, grammar, lexicon, dating and parallel passages. A single overall score would conceal the nature of these arguments.

Our approach

LoraReasoner will construct scenarios. Each conclusion will state its premises, the rules applied, the counterarguments and the assumptions required. Geometric, linguistic, documentary and statistical rules will remain identifiable.

EXAMPLE Reading A is preferred because it agrees with the shape, the syntax and two parallels; reading B is retained because the area is damaged; reading C is rejected because it requires an impossible order.

Expected outcome

The system can answer a counterfactual question: if G43 were replaced by G17, which analyses and sentences would change?

Current status

This module is part of the long-term target. An initial version based on explicit rules and a dependency graph is preferable to a more impressive but less controllable black box.

26 Producing several useful translations

Three purposes

A literal translation shows structure and repetitions. A philological translation explains choices, gaps and technical terms. A fluent translation aims for natural French⁸. These meet different needs and can coexist.

Our approach

Each French segment remains aligned with the transliteration and analysis supporting it. Restorations and alternatives do not disappear. A generative model may help rephrase or compare solutions, but it will not be allowed to invent a source or conceal an uncertainty.

WHY DOES IT MATTER? An elegant sentence can be wrong with great confidence. **Loracraftt** seeks a verifiable translation first, then readable wording.

Human validation

The published version is reviewed by a specialist during the development phase. The system preserves automatic proposals and corrections so that later revisions can be understood.

Current status

LoraTranslate has not yet been implemented. Following sign recognition, ordering, transliteration and linguistic analysis, it will produce translations in French, English and German.

⁸ We chose to move from transliteration to French because more of the data required for the knowledge base are available in that language, particularly lexicon and grammar resources. The software will then offer an English or German translation according to the user's settings.

27 Making a late correction without starting again

The scenario

A working edition has selected G43. After examination, a researcher proposes G17. This correction affects the transliteration, two possible words, a grammatical analysis and a French sentence. It does not affect the photograph, the straightening or the other columns.

Our approach

The program creates a new validation and a new working branch. It recalculates only the results that depend on the corrected sign. Independent stages remain valid. A new revision is produced; the old one remains accessible and citable.

IN PLAIN TERMS We do not repaint the whole building site because one stone has been more accurately identified. We trace the chain of consequences and redo only what needs to be redone.

Expected outcome

A revision comparison identifies the changed sign, the affected transliteration, the analysis now preferred and the updated translation. The reader can check the justification for the change.

Current status

The prototype can already record each new result without overwriting earlier ones. When a user corrects a region, sign or word, **Loracrafft** will then need to recalculate only the results that depend on that correction, all the way through to the translation, without unnecessarily restarting the entire process.

28 The five test images as a development log

The test images do not all serve the same purpose. The clean line checks the basics. The inscription in two registers reveals segmentation errors. Its visual variant helps compare image preparation methods. The vertical document tests columns, cartouches and damage. The stela photographed at an angle tests straightening on a real object.

Case	Question	Precaution
test_1	Are already-correct detections preserved?	No improvement should introduce a regression.
test_2	Are the two registers and each unit correctly delimited?	Do not process test_3 until this stage has been validated.
test_3	Does the method hold up on a more natural-looking variant?	Deliberately deferred case.
test_4	Does the damaged vertical structure remain legible?	Preserve context and alternatives.
test_5	Does straightening take place before geometric processing?	Always check orientation visually.

29 Where the prototype really stands

Level	What it means today
Already tangible	<i>LoraDocument 1.0</i> ; image straightening and contrast enhancement; separation of the two registers in test_2; production of crops and several recognition hypotheses; support for validating multiple results; local, versioned recording of results without overwriting earlier ones.
Experimental	Quadrat and sign segmentation; cartouches; comparison using templates and shape descriptors; Gardiner classification.
Still to be built	Complete reading order; MdC; grammar; lexicon; corpus search and TLA connector; reasoning; translation.
Later option	Audio reconstruction, clearly presented as conventional or reconstructed pronunciation.

30 What works, what fails and what we have learnt

What works

The *LoraDocument 1.0* core, non-overwriting storage, stable identifiers, image provenance, multiple validations and local orchestration under Windows 11 and Python 3.12. Straightening test_5 validated the overall sequence of visual preparation.

What still fails

Accurate segmentation of quadrats and signs in test_2. Open cartouches or those cut off by the edge remain difficult. Shape comparisons are fragile when framing, noise or stroke thickness change. No current score should be presented as scientific certainty.

What failure has taught us

Structure must be detected before small units; a pixel component is not a sign; objects outside the text must be allowed to remain outside it; execution reports and outputs must be archived by version and test case.

What has not been started

Complete reading order, MdC, grammar, lexicon, parallel search, reasoning and translation are planned modules. Describing them precisely in the ***Scientific Architecture*** does not mean that they have been implemented.

31 The four-phase roadmap

1. See and validate: stabilise the geometry of test_2, build the review interface and measure results on real inscriptions.
2. Encode and analyse: produce aligned MdC, transliteration, words and grammatical analyses linked to sources.
3. Compare and argue: search reproducibly for parallels, particularly through authorised TLA resources, then construct explainable scenarios.
4. Translate and publish: propose several levels of translation, submit them for human review and publish a fully traceable revision.

IMMEDIATE PRIORITY Complete the ground truth for test_2 and the geometry-correction interface before seeking a more spectacular recogniser, then move on to test_3.

The foundations—shared record, provenance, rights, tests, archiving and migrations—support every phase. They are not decorative extras: they make the results reusable.

32 How will we know that Loracraftt is useful?

Loracraftt will be useful when a researcher's tool can start from a photograph, correct the geometry, examine proposed sign identifications, develop a reading, find comparisons and publish a translation without losing track of the choices made.

Its ambition is not to eliminate the uncertainty inherent in ancient inscriptions. It is to make that uncertainty visible, localised and productive. An alternative preserved today may become tomorrow's best reading, thanks to a new photograph, reassembled fragments or a better-understood parallel.

THE CENTRAL IDEA An image, a region, a proposed sign, an analysis and a translation are not five successive versions of the same truth. They are five different objects connected by evidence.

The software may change its engine, catalogue or interface. What must endure is the scholarly record: what was observed, what was proposed, who decided, on what evidence and in which version.

The next success is therefore not an automatic-translation demonstration. It is a correct and revisable segmentation of test_2, examined by humans and preserved in a record that allows the process to be rerun. On that foundation, each additional layer will have a purpose.

33 A short glossary

Term	Plain-language definition
Quadrat	A small graphic space in which one or more signs are arranged together.
Gardiner code	A reference code used to identify a sign in a list organised by category.
Corpus	An organised collection of texts and witnesses used to find attestations and parallels.
Digital fingerprint	A value calculated from a file that can verify its identity and detect changes.
Hypothesis	A revisable proposal, such as a possible sign identification.
Lemma	The dictionary form to which a form encountered in the text is assigned.
Manuel de Codage (MdC)	A convention for representing hieroglyphic signs and their arrangement.
Parallel	A comparable passage that may clarify a reading without automatically validating it.
Provenance	The documented history of a file, transformation, proposal or decision.
Register	A broad horizontal band or structured area within a composition.
Transliteration	The representation of Egyptian linguistic values using a specialised alphabet.
Ground truth	Reference segmentation or annotation established and checked by humans to evaluate the program.

34 Sources, limits and links to the technical volume

Main source

Scientific Architecture of **Loracraftt**, Didier DTL Morandi, version 1.1, 17 August 2026. This volume restates its content for a non-specialist reader. For implementation purposes, the technical volume remains the reference for contracts, formats, interfaces and naming rules.

Resources mentioned


Gardiner's sign list for sign classification; the Manuel de Codage for their representation; the *Thesaurus Linguae Aegyptiae* for accessible lemmas, attestations and corpora; Raymond Monfort's glyph resources; and the standards for editions, images and provenance cited in the technical volume.

Limits of this document

The numerical results describe development experiments conducted to date. They are not yet measurements that can be generalised to a representative corpus. The five images are working cases, not evidence of overall performance.

A rule for future notes

Footnotes should be used for sources, historical details and scholarly discussion. A definition needed for understanding belongs in the main text. A data contract, JSON example or detailed interface belongs in the technical volume.





test_5.png

