

Galactic Lettering as a Reading Environment

Hand drawing, character encodings, and procedural transformation
in Flyxion’s font collection

Flyxion

Independent Researcher

October 2026

Abstract

The font collection in the `standardgalactic/fonts` repository contains hand-drawn faces for Standard Galactic and several related scripts, together with a procedural laboratory that turns those faces into thousands of transformed variants. This essay inventories the faces, separates what can be measured from the files from what rests on the maker’s account, and argues that the collection is best read through one distinction: the intensity of a process does not determine the state it leaves behind. A reader’s familiarity with a script accumulates along a path of exposure, and a degraded font is the product of either an independent sweep or a cumulative sequence. Treating both as paths, and keeping geometric, recognition, and provenance measures apart, yields a set of reader-side experiments that the repository is unusually well placed to support.

1 Introduction

Standard Galactic is a substitution alphabet that originated in the *Commander Keen* games of id Software and has since gathered a community of users and a block of positions in the UnderConScript Unicode Registry [1, 2]. Flyxion’s Galactic fonts have become ordinary instruments of reading in that community’s sense of the word. SGA Regular has served as a system font on two computers, and Cursive Galactic is used on a newer machine and for subtitles in VLC. These uses place the designs in conditions quite different from those of an alphabet chart or a decorative specimen. A system interface demands repeated recognition of familiar labels, filenames, numbers, and punctuation, and subtitles demand that recognition take place while speech, images, and narrative continue. The fonts therefore deserve analysis as parts of a working reading environment. Their unusual appearance is only the beginning of the subject. What the files supply, and what can be checked in them, are the conditions under which that appearance can become familiar.

The collection supports a more specific claim, and the essay is organized around it. Familiarity with a script is cumulative. A reader who has met a letterform ten thousand times in menus and subtitles occupies a position reached by a particular path of exposure, and two readers at the same nominal level of practice need not have followed the same path. The

procedural tools in the repository make the same distinction available on the side of the stimulus. Its sequence code separates an independent sweep, in which every state is generated from the original source with its own parameter value, from a cumulative sequence, in which each state is generated from its predecessor. Its order-effect experiments ask whether applying one operation and then another gives the same outline as the reverse. In each case the intensity parameter fails to determine the resulting state, because history enters. The argument of the essay is that a font state, like a reader state, is a point on a path, and that the repository’s habit of keeping geometric measurements, recognition scores, and provenance records apart is the right response to that fact.

The essay examines the font files in the repository, with particular attention to SGA Regular and Cursive Galactic, and draws on Flyxion’s account of making glyphs manually with a Sharpie, using Calligraphr online, and assembling or editing fonts with FontForge. That account supplies production history which the binaries alone cannot recover. The analysis distinguishes direct measurements, observations from rendered specimens, documentary statements in the repository, and interpretations of what these practices make possible. It does not presume that every included glyph was drawn by Flyxion or that an inherited metadata field supplies the complete history of a modified face. The repository’s own licensing notes take the same position: they list the provenance of SGA Regular and Cursive Galactic, among others, as not yet recorded [5]. The snapshot examined is commit `0d57718b6b69058886de31fe788cbdbe821e77a9`.

2 The inventory

The inventory immediately complicates the idea that these are simply fonts for a twenty-six-letter substitution alphabet. SGA Regular contains 14,208 glyphs and 14,205 mappings in its preferred Unicode character map. Cursive Galactic contains 1,134 glyphs and 1,132 mappings. These figures count different things. A glyph is an entry in the font’s repertoire of visual forms, while a character mapping associates a code point with such an entry. Several characters can share a glyph, and a glyph can exist without a direct Unicode mapping. Neither number is equivalent to a count of independently invented signs. The difference matters throughout this collection, whose faces range from compact handwritten alphabets to extensive multilingual compilations.

The principal inventory measured directly from the files is given in Table 1. Private-use counts describe mapped positions, including any positions containing placeholders, rather than certified usable characters.

The table describes the files rather than ranking them. Cheiro can possess fewer glyphs than mapped characters because several quotation-mark characters point to shared glyphs. A compact face can accomplish its intended purpose with dozens of forms, whereas a much larger face can still leave practical gaps. The repository also includes four Amiri faces, whose glyph counts range from 6,472 to 6,624 and whose preferred Unicode maps cover between 1,548 and 1,699 characters. Those are explicitly identified as third-party fonts in the repository documentation. Their presence is useful comparatively: substantial inventories can result

Font file	Glyphs	Unicode mappings	Private-use mappings
Sga-Regular.ttf	14,208	14,205	2,307
CursiveGalactic-Regular.ttf	1,134	1,132	244
NovaMonoStandardGalactic.ttf	6,024	6,021	2,304
Systada-Regular.ttf	669	667	380
Systada.ttf	669	667	380
dactyl.ttf	669	667	380
Logico_philosophicus-Regular.ttf	283	287	0
shapeform.ttf	81	78	0
Cheiro-Regular.ttf	66	70	0
Clypto-Regular.ttf	66	70	0
Lingojam_cipher-Regular.ttf	35	33	0

Table 1: Glyph and character-map counts measured from the font files at the examined snapshot.

from different design requirements, and the relation between mapped characters and available glyphs varies sharply between fonts.

SGA Regular’s extent is chiefly an extent of coexistence. In rendered samples, Latin letters become Galactic forms, accented Latin letters can become Galactic forms with distinguishing marks, and many other writing systems retain their conventional appearance. Greek remains recognizable Greek, Cyrillic remains Cyrillic, and sampled Arabic characters retain Arabic forms. Mathematical symbols, Braille, arrows, and other signs occupy further parts of the inventory. The file therefore combines a particular reading transformation with a broad supporting repertoire. This is especially relevant to its use in a desktop environment, where a font encounters more than continuous English prose. The practical surface includes punctuation, mathematical expressions, unexpected symbols, and occasional text in other scripts.

Nevertheless, a character map is not a guarantee of complete language support. A font can contain Arabic character outlines without supplying all the shaping machinery necessary for polished Arabic typography. SGA Regular has a GDEF table but neither a GSUB nor a GPOS table, the OpenType tables commonly used for contextual glyph substitutions and positioning. This does not make every rendering of complex text impossible, since application behavior and other mechanisms matter, but it limits the conclusions that can be drawn from coverage figures. The same caution applies to combining marks, rare letters, and script-specific conventions. It is more accurate to say that the file carries many kinds of characters than to claim that it fully supports every language whose characters appear in its map.

The absence of a shaping layer has a further consequence for the later sections. The inspected generators implement degradation primarily through changes to outline geometry. The absence of GSUB and GPOS does not require this approach: character remapping, metric changes, raster processing, and application behavior provide other ways to alter rendered text. Of SGA Regular’s 14,208 glyph entries, 14,175 contain outlines and 33 are empty, and the outlines together contain 638,671 points, roughly forty-five per glyph on average. Moreover, 11,710 entries have advance widths of either 512 or 1,024 units within a 1,000-unit em, so the full-width cells slightly exceed the em itself. This concentration is consistent with material

organized around half-width and full-width cells. Together with the sharp stylistic differences between character groups, it suggests a history of compilation and modification: the principal Galactic forms have rounded, somewhat irregular strokes, whereas numerous other characters have conspicuously pixel-like outlines. The evidence does not identify the exact source font for every group. Resemblance to a known broad-coverage font is a lead for further comparison, rather than proof of descent.

The table-level statistics sharpen several of these points. The 14,208 glyphs exceed the 14,205 mapped code points by exactly three, and the three unmapped glyphs are `.notdef`, `.null`, and `nonmarkingreturn`, the housekeeping entries of an ordinary TrueType font. Every other glyph entry is therefore reached by exactly one code point in the preferred character map. Unlike Cheiro, the file does not map multiple characters to the same glyph entry in that map. Distinct entries can nevertheless contain identical outline data, as the earlier outline inspection confirmed, so the count does not establish an equal number of visually distinct designs. The mapped code points divide into 13,821 in the Basic Multilingual Plane and 384 in Plane 1. By Unicode general category the largest groups are other letters (3,630), private use (2,307), other symbols (1,726), lowercase letters (1,223), and uppercase letters (1,076), followed by 969 code points that the Unicode data used for the tally does not assign. That last group should be read with care, since it may reflect characters added to Unicode after the data was compiled or positions filled with placeholders, and the category alone does not decide between them. The 816 nonspacing and 209 spacing combining marks are consistent with the zero-advance overlay mechanism described below for Cistercian numerals, given that no GPOS table exists to position marks, but this has been checked only for the sampled numerals. All thirty-two positions from U+EB40 to U+EB5F are mapped even though the registry proposal assigns twenty-nine, and rendered specimens confirm that the three unused positions display hexadecimal placeholders. The GDEF table occupies thirty bytes; its contribution must be determined from its contents rather than its size, since glyph classifications can inform layout even though GDEF does not itself specify substitutions or positioning adjustments, and the glyph outlines account for about 1.77 million of the file’s roughly 2.03 million bytes, close to 87 percent.

3 From marker to character address

Flyxion’s account of using a Sharpie gives the central Galactic forms an intelligible material origin. A felt-tip marker produces width through the contact between a relatively broad nib and paper, and curved turns, terminal shapes, and small differences in pressure or angle can enter the resulting silhouette. Digitization preserves selected features of that trace while removing its paper, scale, and immediate circumstances. The font can then repeat a hand-drawn sign thousands of times. Its individuality survives as a reproducible outline. What was once an event of writing becomes a resource for subsequent writing by anyone using the file.

It would be excessive to identify the tool solely from the outlines, since similar effects can be drawn digitally or produced through transformation. Here the material account comes from the maker, while the specimens show features consistent with it. SGA Regular’s strokes have

rounded ends and local irregularities that distinguish them from a strict grid construction. Several metadata records explicitly identify Calligraphr as the creation tool. SGA Regular’s unique identifier contains the date string 19 December 2020, Cursive Galactic’s contains 18 May 2024, and Cheiro and Clypto both contain 3 January 2025. These are traces of export or naming history. Later editing and merging can retain them, so they should not be mistaken for exhaustive dates of authorship or final revision.

The movement from marker to Calligraphr and FontForge is consequently more than a change of medium. It separates outline production from character assignment and deployment. The drawing supplies a form, a character map determines when that form appears, and spacing and other metrics determine how it behaves beside neighboring forms. A convincing isolated glyph can still fail in a line if its spacing is inappropriate, and an accomplished alphabet can still be inconvenient if the expected punctuation or whitespace is missing. Much of the labor of making a usable font occurs in these less spectacular relations. A person manually assembling a font is organizing the conditions under which drawn marks can function repeatedly as text.

SGA Regular’s two Galactic repertoires make this especially clear. Its ordinary Latin A–Z positions render Galactic forms, while the private-use range U+EB40–U+EB59 contains another rendering of the Galactic letters. The sampled private-use forms are angular and pixel-like, unlike the rounded Latin-position forms. The file thus distinguishes, in its construction, between changing how ordinary Latin-coded text looks and displaying separately encoded Galactic characters. A user can encounter similar alphabetic content through two different arrangements of character identity and glyph selection.

The private-use arrangement follows the Under-ConScript Unicode Registry proposal, which allocates U+EB40–U+EB5F to Standard Galactic. It assigns twenty-six letters, two quotation marks, and a full stop, leaving three positions unused. The registry coordinates private agreements about character positions, and it does not confer official Unicode assignments. Its practical value is nevertheless substantial. Two independently produced fonts can agree that the same private-use code point means a Galactic letter, allowing text to retain its intended script identity when moved between those fonts. The convention provides a shared address system for a community’s writing practices [1, 2].

Latin-position substitution serves a different practical aim. An existing English subtitle file can remain unchanged while its appearance becomes Galactic. Search, copying, and other text operations generally continue to encounter the original Latin characters, subject to the behavior of the application and document format. Dedicated private-use Galactic text changes the underlying character sequence and therefore requires appropriate conversion or input methods as well as a compatible font. These approaches have different conveniences. The former integrates easily with an existing computer interface, and the latter makes the script distinction explicit in the stored text. A font that accommodates both can participate in both practices without forcing every use into one model.

4 Registry material and composition

The registry-aligned material extends beyond Standard Galactic. SGA Regular has mappings throughout the ranges allocated to Tengwar, Cirth, Verdurian, Visible Speech, D’ni, Aurebesh, Orokin, and Deini, and selected rendered samples show recognizable forms from those systems. These observations support the claim that the file contains a collection of constructed and specialized scripts. They do not certify every character of each range. Some blocks contain assigned positions, unused positions, and placeholders together. A full conformance study would compare each assigned character with its chart and test its behavior in appropriate text sequences. The present analysis establishes a substantial pattern without turning sampled inspection into an exhaustive audit [2, 3].

Cistercian numerals offer a particularly revealing example because the convention governs composition as well as individual appearance. The registry proposal assigns spacing units digits and combining tens, hundreds, and thousands digits within U+EBA0–U+EBDF. In both SGA Regular and Cursive Galactic, the sampled units digit advances the cursor, while the sampled higher-place digits have zero advance and outlines extending back into the preceding character’s area. This is consistent with building a compound numeral by overlaying its components. It is a small typographic mechanism: the stored sequence has several characters, but their spatial arrangement can yield one assembled sign. Exhaustive testing of all combinations and applications remains outside the present inspection [4].

This example also explains why placeholders must be interpreted locally. The Cistercian proposal explicitly leaves positions such as U+EBAA–U+EBAF unused. Boxes bearing those hexadecimal addresses do not demonstrate a failure to draw the digits, because those positions do not designate digits in the first place. A placeholder at an assigned position would be a different matter. Broad fonts often preserve these diagnostic forms so that missing or undefined material remains visible. Their presence contributes to the raw glyph count without contributing the same kind of writing capacity as a usable character. The inventory is consequently both a collection of signs and a record of its own boundaries.

5 Cursive Galactic and reading in use

Cursive Galactic reorganizes this environment around a more continuous visual gesture. In the rendered specimens, its main letters lean, bend, and extend in ways that make their strokes feel more directional than those of SGA Regular. The word “cursive” should be understood carefully here. The font contains neither GSUB nor GPOS tables, and the inspected lettering does not depend on an OpenType system of contextual joins. Its cursiveness is visible chiefly in the individual forms and their arrangement, rather than demonstrated through a machinery of neighboring-letter substitutions. The distinction does not diminish the style. A font can evoke writing in motion without connecting every letter to the next.

The contrast is most evident at the scale of a line. SGA Regular’s mixture of dots, short bars, and discrete strokes gives sampled words a segmented texture, while Cursive Galactic draws more attention to the direction and extension of strokes. Whether one is easier to read cannot be decided from those descriptions alone, because size, spacing, display resolution,

familiarity, and the particular reader all matter. Flyxion’s preference for reading subtitles in the cursive face is evidence about an established personal practice. It is not a controlled comparison of population reading speeds. Its importance is that a design which may look difficult to an unfamiliar observer has become an effective reading surface for its user.

Subtitles make that practice especially instructive. Their disappearance imposes a temporal constraint, since the viewer cannot assume that a line will remain available indefinitely. Repeated use supplies opportunities to recognize common words and recurring patterns under that constraint. It is reasonable to interpret sustained comfortable viewing as evidence of practical familiarity, while leaving open how that familiarity is implemented psychologically. No particular neural mechanism need be posited, and no claim is needed that every letter is processed in a particular sequence. The observable fact is sufficient: an alternative graphic vocabulary has become compatible with following an ordinary film.

Using a system font expands that familiarity across tasks. Menus, error messages, filenames, and settings each furnish recurring encounters with the same repertoire. Their wording often has strong contextual support, but they also differ in length, importance, and frequency. Together they create a distributed practice environment without requiring a separate lesson. This gives the fonts an unusual relation to learning. They occupy activities that already matter to the reader, allowing familiarity to develop through use. The interface becomes both the object of attention and, intermittently, the background against which other attention takes place. It is also a clear case of a cumulative path: what the reader can do now depends on the order and spacing of encounters that cannot be recovered from the present level of skill.

6 The smaller faces and layered authorship

The smaller faces help distinguish this environment from mere inventory size. Cheiro’s 66 glyphs produce recognizable Latin handwriting with a flowing, rounded character in the rendered sample. Clypto, also with 66 glyphs, gives the same Latin sentence a more angular and economical texture. Both support letters and a selection of punctuation, while ordinary decimal digits are absent from their preferred Unicode maps. A renderer using these files alone can therefore show missing-character boxes for a numerical specimen, whereas a browser or desktop with fallback may conceal that limitation by supplying digits from another face. Apparent completeness in an application can be a property of the font environment rather than of one file.

Cheiro and Clypto are useful comparisons because they preserve recognizable Latin reading while varying the drawn vocabulary. Their identity is located substantially in gesture, proportions, and the differentiation of familiar characters. SGA Regular and Cursive Galactic modify the familiar graphic vocabulary much more decisively, even when the underlying characters remain Latin. The collection therefore explores several degrees of departure from conventional typography. This observation concerns what the files do, rather than an assumption that every face was conceived as part of one systematic research program from the outset.

Logico_philosophicus supplies a further case. Its Latin-position specimen renders as an arrangement of mathematical-looking and logical-looking signs, rather than recognizable Latin

handwriting or the usual Galactic forms. Its 283 glyphs support 287 mapped characters. The effect invites a comparison between reading prose and inspecting formal notation, but the name and visual vocabulary do not establish a new logical language. Without an additional grammar and interpretation, substitution changes the appearance of the sentence rather than its inferential content. Its interest lies partly in that displacement: prose can acquire the visual atmosphere of formal argument while continuing to be ordinary prose at the character level.

The Lingojam cipher font demonstrates why a standard pangram is not always an appropriate test. Its 35 glyphs and 33 mappings do not include the ordinary Latin alphabet. Instead, the map contains selected Canadian syllabic characters, mathematical signs, a katakana character, and other Unicode symbols. An ordinary English sentence therefore produces missing characters when rendered in this font alone. That result should not be described as a failed alphabet without considering its intended input. The file appears designed to render an already converted symbol sequence. This is a third arrangement alongside Latin-position substitution and dedicated private-use encoding: existing Unicode characters serve as the carriers of a visual cipher.

Such reuse can be convenient because those characters can already be copied and exchanged through Unicode text systems. It also introduces a distinction between their conventional identities and their intended use in the cipher. A mathematical sign employed as a substitute letter remains that mathematical character to software which does not know the convention. The same visible result can therefore conceal quite different underlying structures. Studying these fonts requires attention to what was typed, what was stored, and what the font actually drew. A screenshot alone cannot reliably answer all three questions.

Shapeform pushes the comparison toward compact silhouettes. Its 81 glyphs and 78 mappings produce dense, irregular black forms in the inspected Latin-letter sample. Its preferred map does not include ordinary space or decimal digits, which again makes a conventional sentence specimen partly a test of unsuitable input. Its internal family name remains “Untitled1,” while a copyright field names Mechachleopteryx and the unique identifier records a FontForge date string in December 2023. These traces show why filenames, internal names, and actual visual behavior should all be examined. They need not have been brought into complete alignment during a manually evolving project.

Systada and Dactyl reveal another kind of continuity. Systada-Regular, Systada, and Dactyl each contain 669 glyphs, 667 mappings, and 380 private-use mappings. Their specimens share a Galactic character, while details such as the treatment of numerals differ. Equal counts do not prove that the files are identical or establish their exact genealogy. Their metadata provides further evidence: the Systada files and Dactyl preserve an attribution to Andrew O’Rourke, with acknowledgments to Fixedsys Excelsior and galax.xyz, while Dactyl has its own internal naming and a later FontForge identifier. The repository’s licensing notes separately describe Systada as a mixture of hand-drawn material and glyphs from another typeface whose source is not yet identified [5]. The repository’s `FONTLOG.txt` consists of a single line, “Terminal Friendly Standard Galactic Font,” and does not name a file. On Flyxion’s account it describes Systada, which gives that family a purpose the other faces lack: it was made to be read in a terminal. This is consistent with the Fixedsys Excelsior acknowledgment in the metadata, since that typeface is a fixed-width system face, but the link between the log line and the font

rests on the maker’s statement and not on anything in the file. These are traces of layered construction, not reasons to assign all the drawing to any one source.

NovaMonoStandardGalactic is particularly instructive because its filename could lead an analyst astray. An ordinary Latin pangram renders as recognizable Latin text in the inspected file. Its internal family name is “NovaMono for Powerline,” and its metadata records the NovaMono attribution. It also contains 2,304 private-use mappings. The responsible conclusion is that this file’s Galactic role cannot be inferred simply from its filename or from assuming that it performs the same Latin substitution as SGA Regular. The repository’s notes identify it as a Modified Version of NovaMono under the SIL Open Font License, and they draw the consequence that follows from the licence: a Modified Version may not keep the Reserved Font Name, so the family name inside the font as well as the filename should change [5]. The metadata observation and the legal rule therefore point at the same field. Further glyph-level comparison would be needed to delimit all of the file’s modifications.

These histories make authorship a layered matter. A font may include original drawings, borrowed supporting characters, altered outlines, retained metadata, and newly introduced mappings. Manual labor can be concentrated in selecting and arranging those relations as well as in drawing every contour. Flyxion’s recollection of working through several tools is compatible with this layered structure. The files do not preserve the whole sequence of decisions, but neither does the absence of a written design history imply that those decisions were absent. The licensing notes add a rule for derived material: a font generated from an OFL source inherits that licence, and one generated from a public-domain source inherits that status. Provenance in this collection is therefore recorded most completely for the transformations and least completely for the sources, an asymmetry that matters for any claim about what the experiments transform.

7 Procedural transformation

The newer repository extends that craft into procedural transformation. At the inspected snapshot, its complete recursive tree lists 2,418 TTF or OTF paths, including 15 root-level font files and 2,184 font paths under `experiments/`. These are file counts, not counts of unique designs or independent inventions. Additional groups hold earlier experimental versions, styled variants, debugging outputs, and alterations. The scale is nevertheless meaningful because it shows how a small set of source faces can become material for a large space of controlled variations. The handwritten outline becomes a starting condition for later geometric processes.

The inspected generator implements operations including jitter, shear, quantization, waves, rotation, contour displacement, and several nonlinear deformation fields. Some affect individual points, while others move larger parts of a contour coherently or impose spatial divisions. These differences matter aesthetically and experimentally. Independent noise can roughen a stroke without preserving the relation between its parts, whereas coherent movement can retain a recognizable local structure while relocating it. Quantization subjects coordinates to a grid, producing a different kind of simplification from simple scaling. The experiment names offer evocative readings of these operations, but their mathematical implementation determines

what they actually change.

The code also distinguishes an independent parameter sweep from a cumulative sequence. In an independent sweep, each output is generated from the original source with its own parameter value. In a cumulative sequence, each state is generated from its predecessor. The difference changes what a progression means. A sequence of increasingly strong deformations need not contain any accumulated history if every frame starts again from the source. Conversely, cumulative modification can retain earlier disturbances and alter the conditions under which later ones act. The repository records which kind of sequence produced a font in a sidecar file beside it, which makes the distinction a property of the data and keeps it out of the analyst’s memory.

Transformation order can matter for similar reasons. Quantizing a curve before deforming it need not produce the same outline as deforming it before quantization, because the second operation acts on a different intermediate state. The repository advertises paired order-effect experiments, including a wave-then-grid against grid-then-wave comparison at fixed intensity, and contains viewers for comparing such paths. This is a productive extension of the font work into questions of persistence and reconstruction. It does not require claiming that a glyph literally remembers. The final geometry can preserve consequences of the sequence through which it was produced, and history becomes observable through differences between resulting forms.

The repository also treats the schedule of intensity along a document as a separate variable. A small module of curve functions maps a position in the document to a distortion intensity, with linear, quadratic, cubic, smoothstep, and cliff shapes. A cliff schedule stays near clean until a chosen position and then ramps rapidly, and a quadratic schedule delays most of the damage to the end. The same endpoint fonts under different schedules give a reader different experiences of the same text, which is one more way in which the nominal intensity at a point underdetermines what is encountered.

The measurement code makes a useful epistemic distinction. It records representation-level quantities such as numbers of points and contours, bounds, and a contour-length measure. Such measurements characterize the stored geometry. They cannot, without reader evidence, establish how legible a word is or whether its meaning remains recoverable. Small geometric changes can remove a diagnostic dot, while large movements elsewhere can leave recognition relatively undisturbed. A useful investigation therefore relates geometric state to performance on a specified task instead of substituting one for the other. The repository documentation explicitly acknowledges this limit.

Its OCR experiments introduce another task with another kind of observer. The inspected core renders words against a plain high-contrast background, submits the image to Tesseract, and compares the output with the expected string. This arrangement helps separate font-related changes from additional background noise. However, an English OCR system failing on Galactic substitutions is not by itself evidence that the text has become unreadable to a practiced Galactic reader. Nor does such failure establish security against other recognition systems. It establishes a result for a particular recognition setup. Human reading, generic OCR, and an OCR model trained on the relevant script can have substantially different success conditions.

Even the scoring method deserves specificity. The core uses Python’s SequenceMatcher to derive an edit-like score by summing the larger side of each non-matching alignment block. Each block can be realized by substituting the shared length and inserting or deleting the remainder, so these operations define a valid edit sequence. Its cost is therefore an upper bound on the minimum Levenshtein distance, although it need not equal that minimum. With the same normalization and clipping, the resulting character-accuracy score cannot exceed the corresponding score calculated from the true minimum distance. This bound follows from the construction rather than from random sampling. It identifies what the reported quantity measures, and the same care that distinguishes character mappings from glyph counts should distinguish alignment-derived similarity from minimum edit distance.

8 Reader-side experiments

The principal opportunity opened by this collection is to study the relation between learned recognition and changing inscription. The practices described above supply an unusual combination: a script that one reader has used daily as a working interface, and a generator that can vary the inscription while holding the underlying character sequence fixed. The experiments sketched here are proposals. None has been run, and none of the observations in the preceding sections is evidence for how they would turn out.

The first design varies the history of exposure while holding its amount constant. Initially unfamiliar readers are randomly assigned to schedules containing the same material and the same total exposure, with order, spacing, or progression manipulated between groups. All groups are then assessed on a common test containing held-out items. Both groups accumulate experience: shuffling or resetting the displayed stimuli does not reset a reader’s learned state. Differences on the common test would provide evidence that the manipulated exposure history matters beyond its total amount, subject to controls for baseline ability and other differences between groups.

The second design varies the stimulus along a path while comparing readers with different levels of familiarity. Practiced and initially unfamiliar readers encounter matched text through a clean source face, an independent sweep, and a cumulative sequence ending at the same nominal intensity, with text assignment and presentation order counterbalanced. A related design compares wave-then-grid with grid-then-wave at matched parameter settings. A reliable difference in recognition accuracy or response time would support a perceptual consequence of transformation order under those conditions. Failure to detect a difference would leave the geometric distinction intact and could reflect task insensitivity, limited statistical power, or a difference without practical relevance to the tested task.

The third design asks how well the geometric measurements predict recognition. The sidecar measurements, the OCR scores, and reader accuracy and response time can be collected on the same set of rendered words and compared. The repository’s own caution that its measurements are geometric and do not establish legibility sets up the question, and the practiced and naive readers serve as two different observers alongside generic OCR. A script-trained OCR model would add a fourth. Cheiro and Clypto, which preserve recognizable Latin reading while

changing the drawn vocabulary, act as a control for effects of handwriting-like texture apart from substitution.

The fourth design uses the curve schedules. Because the position-to-intensity mapping is a separate function, a document can be rendered under a linear schedule and under a cliff schedule with identical endpoints. Recovery after added context, isolated letters against continuous text, and the point at which a reader stops following the narrative are all measurable against position, and the curve shape is the manipulated variable.

These designs share several constraints that should be stated at the outset. Display size, spacing, and fallback glyph substitution can change results, and a rendering pipeline should fix them. A single practiced reader is a case study and cannot support a population claim, so the practiced condition would need more than one reader. The recognition scores should be reported with their method, for the reasons given in the preceding section. Finally, the sources of the transformed faces should be recorded as carefully as the transformations, since an experimental result about a face of unrecorded provenance inherits that gap.

9 Preservation and conclusion

The collection invites a more precise account of preservation. Preserving a font file can preserve outlines without preserving knowledge of its expected input. Preserving a private-use character sequence can preserve code points without preserving the convention that gives them meaning. Preserving a screenshot can preserve appearance while discarding the underlying character identities. The repository's mixture of fonts, source projects, scripts, text, and explanatory material can help preserve several of these relations together. A future account of the manually constructed faces would add another layer by recording which characters were drawn, borrowed, transformed, or reassigned and why.

The scale of SGA Regular becomes intelligible in this setting. Its 14,208 entries record much more than the visible alphabet on an introductory chart. They include practical accommodations, other scripts, registry-aligned material, diagnostic placeholders, and the residues of a layered production history. Cursive Galactic selects a smaller repertoire while giving its everyday reading forms a different movement. Cheiro and Clypto investigate handwritten Latin texture, Logico_philosophicus shifts prose toward the appearance of symbolic notation, Lingojam_cipher depends on converted Unicode input, and the procedural experiments turn established outlines into sequences of possible states. Their differences are part of the collection's substance.

What joins these activities is the work of making marks available for continued interpretation, and the distinction announced at the start returns here at every joint. A Sharpie stroke becomes a vector outline, an outline acquires a character address, and a character becomes part of an interface or subtitle, where repeated encounters let a reader proceed through it with increasing familiarity. Each step is an operation on the previous state, and none of them is fixed by a single intensity parameter. The same holds for the transformed faces, whose states depend on whether they were generated from the source or from one another, and on the order in which the operations were applied. The fonts are consequently valuable both as artifacts

of manual making and as instruments for studying how a reading environment, and a reader's place in it, come to be established. The strongest evidence available now is an ordinary act, a viewer following a film in letters that were once drawn with a marker, and the proposals above are an account of how to turn that act into a measurement.

Sources and method

The font counts, mappings, metadata, table inventories, spacing observations, and outline comparisons were obtained directly from downloaded files using FontTools. Visual observations refer to specimens rendered with Pillow, and no claim of exhaustive multilingual shaping or perceptual validation is made. SGA Regular and Cursive Galactic in the fonts repository were byte-for-byte identical to the copies previously examined in the alphabet repository. The units-per-em value, the advance-width counts, and the table list of SGA Regular were rechecked against the file at the examined commit. The discussion of the scoring bound follows analytically from the edit sequence represented by the opcode calculation; no random-string experiment is claimed. Private-use block occupancy was checked against registry allocations, with selected outlines inspected rather than every assigned character certified. The counts of repository files came from an untruncated recursive Git tree. Font metadata dates are retained date strings rather than verified completion dates. Statements about Sharpie drawing, manual work, system-font use, and VLC subtitles come from Flyxion's account in the accompanying conversation. The essay offers interpretation of those practices without presenting that interpretation as experimental psychology.

References

- [1] Rebecca Bettencourt, *Standard Galactic registry proposal*, 14 September 2019. <https://www.kreativekorp.com/ucsur/charts/sga.html>
- [2] Kreative Korporation, *Under-ConScript Unicode Registry*, allocations and project description. <https://www.kreativekorp.com/ucsur/>
- [3] Doug Ewell, *Aurebesh registry proposal*, 22 June 2016. <https://www.kreativekorp.com/ucsur/charts/aurebesh.htm>
- [4] Paul Hardy and Rebecca Bettencourt, *Cistercian Numerals registry proposal*, 18 November 2020. <https://www.kreativekorp.com/ucsur/charts/cistercian.html>
- [5] Flyxion, *fonts repository snapshot*, especially the font binaries, README.md, LICENSES.md, FONTLOG.txt, make_experiments.py, advanced_experiments.py, experiment_support.py, degradation_curve.py, and adversarial_ocr_core.py. <https://github.com/standardgalactic/fonts/tree/0d57718b6b69058886de31fe788cbdbe821e77a9>. Repository documentation was used as evidence of the project's stated organization and component history, rather than as a substitute for inspecting the font files.