AUTOMATED TRANSCRIPTION OF ACA CANCELLERIA REGISTRO 2053: A REPORT ON AI-ASSISTED PALEOGRAPHIC PIPELINE DEVELOPMENT, IMPROVEMENT, AND QUALITY ASSESSMENT

Prepared for dissertation advisory purposes
Date: March 29, 2026
Source data: ACA, Cancelleria, Registros, núm. 2053
Pipeline output: 282 transcription files
Comparison register: ACA, Cancelleria, Registros, núm. 2056 (processed previously)

I. INTRODUCTION AND RESEARCH CONTEXT

This report describes the automated transcription and translation of a second ACA Cancelleria register using the pipeline previously applied to Registro 2056. It is written to stand alone — readers unfamiliar with the earlier work will find a complete account of the methodology — but it also serves as a comparative document, examining the significant improvements made to the pipeline between the two runs and documenting the measurable gains in transcription quality that resulted.

The source material is ACA, Cancelleria, Registros, núm. 2053: Curie Sigilli Secreti 1 de la reina Violante — the first register of the secret seal court of Queen Violant de Bar, compiled under the reign of Joan I of Aragon (r. 1387–1396) and spanning the years 1387 to 1390. This is, in broad terms, a register of personal correspondence and administrative orders issued under Violant's secret seal (sigillum secretum), primarily drafted by her secretary Bartolomeu Sirvent. The register contains letters in Catalan, Aragonese, and Latin, addressed to foreign monarchs, nobles, ecclesiastical officials, and royal administrators across the Crown of Aragon and beyond.

The processing of Registro 2053 represents the third full register in this research project (after Registros 1819 and 2056), and the first to benefit from all the pipeline improvements developed in response to problems identified in the earlier runs. The results demonstrate a measurable improvement in quality and a significant reduction in critical failures.

II. THE PIPELINE: ARCHITECTURE AND RATIONALE

The pipeline processes manuscript images through a four-stage workflow. The pipeline was developed iteratively over several sessions. Its core structure reflects a deliberate methodological choice: rather than trusting a single AI model for transcription, it employs a two-stage, two-model workflow that intentionally exploits the different strengths and failure modes of two large language models — Google's Gemini 2.5 Pro and Anthropic's Claude (Sonnet 4.6 for conversational development, Opus for the transcription pipeline).

The following description is complete for readers encountering this work for the first time.

STAGE 0 — REGISTER METADATA EXTRACTION (NEW IN THIS RUN)

Before processing any manuscript page, the pipeline reads the archival cover sheet (carátula) produced by the Archivo de la Corona de Aragón and included as the first image of every digitized register. Using a lightweight, inexpensive AI call, it extracts:

- Signatura (shelfmark)
- Reinado (reign)
- Título (register title)
- Años (date range of contents)
- Alcance y contenido (scope and contents summary)

For Registro 2053, this extraction returned:

Signatura: Archivo de la Corona de Aragón, Real Cancillería, Registros, núm. 2053
Reinado: Juan I
Título: Curie Sigilli Secreti 1 de la reina Violante
Años: 1387-1390

This metadata is then injected into every Gemini and Claude prompt throughout the entire run as an explicit anchor for date readings. Both models are told: "All documents in this register date between 1387 and 1390. When transcribing dates, use this range to resolve ambiguous Roman numeral readings... A date that appears to fall outside the register's known range is almost certainly a misreading — flag it explicitly." The models are also given a reference table of word-form numerals used by scribes of this period: Catalan "set" = 7, "vuit"/"vuyt"/"oy" = 8, "nou" = 9; Aragonese "siete"/"siet" = 7, "ocho" = 8; Latin "septem" = 7, "octo" = 8, "novem" = 9.

The extracted metadata is cached to a JSON file in the output directory, so partial re-runs and targeted re-processing jobs can benefit from the same register context without re-reading the cover sheet.

STAGE 1 — IMAGE-BASED PRE-SCREEN

Before the expensive main Gemini call, each image is sent to a cheaper, faster model (Gemini 2.5 Flash) with a minimal prompt asking a single binary question: "SKIP or PROCESS?" If the response is SKIP, the page is recorded as a blank or modern document and no further API calls are made for it. For Registro 2053, this stage correctly identified 3 pages as non-manuscript and skipped them without human intervention, saving approximately $0.27–$0.54 in otherwise wasted API costs.

STAGE 2 — GEMINI STAGE 1 PALEOGRAPHIC TRANSCRIPTION

Each manuscript image is submitted to Gemini 2.5 Pro with a detailed prompt containing: the register metadata context (including the date range anchor), instructions for handling abbreviations, uncertainty marking conventions, and language identification requirements. Gemini produces a Stage 1 transcription with abbreviations expanded in square brackets, uncertain readings flagged with [?], and illegible passages marked [illeg.]. It does not translate.

The Gemini prompt for this run contained the following date guidance in full:

"IMPORTANT FOR DATE READINGS: All documents in this register date between 1387 and 1390. When transcribing dates, use this range to resolve ambiguous Roman numeral readings... Scribes in this period sometimes spelled out part of the year in words (e.g., 'set' for 7, 'vuit' for 8, 'nou' for 9 in Catalan; 'siete'/'siet' in Aragonese; 'sept'/'septem' in Latin). A date that appears to fall outside the register's known range is almost certainly a misreading — flag it explicitly."

STAGE 3 — CLAUDE STAGE 2 RECONCILIATION AND TRANSLATION

Each manuscript image is submitted again — to Claude (Anthropic) — together with Gemini's Stage 1 transcription, the register metadata, and a summary of the immediately preceding page (the cross-page context window feature). Claude is instructed to review Gemini's reading directly against the image, correct errors, expand abbreviations, produce an English translation, provide historical notes on identifiable persons and places, and list remaining uncertainties. The cross-page context helps Claude recognize document continuity, correctly identify issuers when a letter spans multiple folios, and avoid the context-loss errors that affected earlier runs.

Claude's date-anchoring instructions in this run were explicit:

"Use the register metadata above to anchor all date readings. Any date falling outside the register's known year range is almost certainly a misreading — correct it and explain your reasoning. Watch for scribal word-forms for year digits: Catalan 'set'=7, 'vuit'=8, 'nou'=9, 'deu'=10; Aragonese 'siete'/'siet'=7; Latin 'septem'=7, 'octo'=8, 'novem'=9. These may appear at the end of a year formula like 'M ccc lxxx set' = 1387."

OUTPUT

Each processed page produces a plain-text file containing the full Gemini Stage 1 output, the full Claude Stage 2 output, a quality status indicator, the register metadata, and a token usage summary. A CSV file is generated as a secondary output with one row per page and columns for metadata fields, transcription, translation, historical notes, and a Gemini quality flag. The quality flag is generated automatically by scanning Claude's output for linguistic markers indicating different levels of concern:

VERIFIED (no significant issues flagged)
LOW (minor errors noted)
MEDIUM (numerous errors; use with caution)
HIGH (major errors or misidentifications)
CRITICAL (fabrication or hallucination detected)
CLAUDE-ONLY (no Gemini input)

RATE LIMITING AND PACING

Gemini 2.5 Pro operates under a rate limit of approximately 2 requests per minute. The pipeline enforces a 30-second delay between pages to avoid the throttling problem identified in the Registro 2056 run, where insufficient inter-page delays caused Gemini to return abbreviated 200-400 character outputs rather than the full 3,000-5,000 character transcriptions that characterize high-quality operation. This pacing adds approximately 12 hours of processing time for a 280-page register but is essential for output quality.

III. OVERVIEW OF REGISTRO 2053

Based on the pipeline output, Registro 2053 is a richly varied diplomatic corpus. It contains:

1. Personal letters from Violant de Bar to foreign monarchs — the kings of France, Navarre, Castile; the Dukes of Burgundy, Berry, and Anjou; the Count of Armagnac
2. Administrative mandates to Aragonese and Catalan officials — governors, treasurers, bailiffs, masters rational
3. Letters of credence and instruction — authorizing named emissaries to speak on the queen's behalf
4. Memoranda — notes by secretary Bartolomeu Sirvent recording oral commands
5. Legal and financial instruments — tutorships, safe-conducts, payment orders

Recurring themes across the register include Joan I's illness and recovery (referenced in multiple letters to foreign correspondents), diplomatic relations with the French royal family and Avignon papacy, and household administration. The secretary Bartolomeu Sirvent appears in the subscription formula of virtually every document ("Domina Regina mandavit michi Bartholomeo Sirvent" — the Lady Queen commanded me, Bartolomeu Sirvent), making him an identifying constant across the register.

Pipeline statistics:

Total images:	285
Pages skipped (blank/modern):	3
— caught by pre-screen:	3 (all three identified automatically)
Pages with full two-model output:	280 (98.6% of processed pages)
Pages Claude-only (Gemini empty):	0
Errors requiring manual retry:	1
Total processing time:	~9.5 hours
Estimated processing cost:	$23.70 USD
Gemini quality flags (VERIFIED):	219 (77.7%)
Gemini quality flags (MEDIUM):	33 (11.7%)
Gemini quality flags (HIGH):	24 (8.5%)
Gemini quality flags (CRITICAL):	6 (2.1%)
Gemini quality flags (CLAUDE-ONLY):	0 (0%)

IV. COMPARISON WITH REGISTRO 2056: MEASURABLE IMPROVEMENTS

The processing of Registro 2056 (305 images, $25.13 cost) was the first full register run after an initial test on Registro 1819. It identified several critical problems in the pipeline that were subsequently addressed. The following comparison shows the measurable improvement achieved in Registro 2053.

QUALITY FLAG COMPARISON

	Reg 2056	Reg 2053	Change
VERIFIED	67.1%	77.7%	+10.6 percentage points
CRITICAL	7.0%	2.1%	-4.9 percentage points
HIGH	16.8%	8.5%	-8.3 percentage points

MEDIUM	8.7%	11.7%	+3.0 percentage points
CLAUDE-ONLY	0.3%	0%	-0.3 percentage points

The headline result is a drop in CRITICAL failures from 22 pages in Registro 2056 (7%) to just 6 pages in Registro 2053 (2.1%), and a corresponding rise in VERIFIED pages from 67% to 78%. The three new pipeline features — register metadata extraction, image pre-screening, and cross-page context — each contributed to this improvement.

WHAT DROVE THE IMPROVEMENT

1. Date anchoring (register metadata) was the single most impactful change. Many HIGH and CRITICAL failures in Registro 2056 involved Gemini misreading the year component of dates — producing readings outside the register's actual date range without any mechanism to detect or correct this. With the date range explicitly provided, both models could anchor uncertain readings against known bounds, and Gemini was instructed to flag anomalies rather than silently generate plausible but wrong dates.

2. Cross-page context reduced issuer misidentification errors. In Registro 2056, several HIGH flags were generated when Claude identified the wrong issuer because it processed each page in isolation. With context from the preceding page, Claude could recognize that a page beginning mid-sentence was a continuation of the previous document and assign the correct issuer.

3. Image pre-screening eliminated wasted calls. All 3 skipped pages in Registro 2053 were caught before the expensive Gemini 2.5 Pro call, saving cost and preventing the occasional false-positive where Gemini would produce brief, confused output for blank pages.

THE PERSISTENCE OF SOME CRITICAL FAILURES

The remaining 6 CRITICAL pages in Registro 2053 confirm that the confabulation failure mode — where Gemini generates plausible but entirely invented text rather than reading the image — has not been fully eliminated. The mechanism appears to be pattern-matching: when Gemini encounters a very difficult page (heavy fading, ink damage, unusual hand), it falls back on its training data and generates a document that looks like it should belong in this register. The date anchoring feature helped constrain year readings but did not prevent the deeper confabulation. Section V(E) below examines a specific example.

V. ILLUSTRATIVE EXAMPLES OF PIPELINE PERFORMANCE

The following twelve examples are drawn from actual output files. Each is identified by its filename and attributed to the model responsible for the quoted passage. Examples are organized from high-quality to problematic output, reflecting the full range encountered in the register.

A. HIGH-QUALITY OUTPUT: CATALAN DIPLOMATIC CORRESPONDENCE SHOWING DATE ANCHORING IN ACTION

File: ACA_Reg2053_006_transcription.txt

Page 006 contains two letters dated March 10, 1388: a Catalan letter to an uncle (likely the Duke of Burgundy or Berry) reporting on the king's health, and a Latin/Aragonese letter of recommendation for Petrus de Gurrea. This page is particularly instructive because the year is written using "oy" — a phonetic rendering of Catalan "vuit" (eight) unusual enough to trip up a model without date context.

Gemini's Stage 1 transcription correctly read the date formula and identified the word form:

> "Dada en Barch[ino]na sots n[ost]re segell s[e]cret acostumat a .x. dies d[e] març. El any .Mº cccº lxxxº oy. / 'oy.': This is the Catalan word 'vuit' (eight), written phonetically. The date is thus 10 March 1388, which fits the register's date range."
> [Gemini, ACA_Reg2053_006_transcription.txt]

This is the date-anchoring feature working precisely as intended: Gemini identified an unusual word-form numeral, named it correctly, and confirmed the reading against the known register range. In the Registro 2056 run — before this feature was added — comparable word-form dates were frequently misread or produced dates outside the register's actual span without comment.

The body of the first letter contains a passage of particular historical interest, reporting on Joan I's illness and anticipated recovery:

> "lo senyor Rey marit e senyor n[ost]re molt car es tal e tanta conualescencia d[e] sa malaltia q[ue] s[er]a p[er]a dins breu e[ss]er s[er]vit p[er]fetament / si co a ame senyor deu."
> [Gemini, ACA_Reg2053_006_transcription.txt]

Translation (Claude): "The lord King, our much-loved husband and lord, is in such and so great convalescence from his illness that within a short time he will be perfectly served again, God willing."

B. HIGH-QUALITY OUTPUT: MULTILINGUAL PAGE WITH CATALAN AND LATIN

File: ACA_Reg2053_009_transcription.txt

Page 009 contains three documents: a Catalan letter of credence authorizing a royal official, a personal Catalan letter addressed to a foreign prince reporting Joan I's recovery, and a Latin subscription formula. The page demonstrates the pipeline's handling of mixed-language content on a single folio.

Gemini's Stage 1 transcription of the personal letter correctly identified the Catalan formulaic opening and the health report, including the dating:

> "Car[issim]e ale e poderos p[r]incep e cosí molt car desijants continuam[en]t ab feruent desig saber v[ost]ra sanitat e bon estam[en]t vos p[re]gam tan affectuosament co podom q[ue] p[er] v[ost]res honorades l[et]res nos en vullats certificar tota vegada q[ue]s sia en guynor... q[ue] nos e n[ost]res infants som sans e en bona disposi[ci]o e n[ost]re p[ro]nces e quel senyor Rey marit e senyor n[ost]re molt car es en tanta conualescençia d[e]la malaltia q[ue] dins breu espera esser guarit per la fermetat grà[ci]a a

n[ost]re senyor Deu."
[Gemini, ACA_Reg2053_009_transcription.txt]

The year formula "M ccc lxxxviiii" (1389) was correctly read, and the subscription "ffuit directa Regi ffrancie" (it was sent to the King of France) was identified, allowing Claude to determine that this was a diplomatic letter to Charles VI of France, one of the most significant recipients in the register. Claude's historical note correctly identified the letter as part of Violant's sustained effort to maintain relations with the French royal family during Joan I's illness — a theme running throughout the register.

C. HIGH-QUALITY OUTPUT: ARAGONESE CORRESPONDENCE

File: ACA_Reg2053_011_transcription.txt

Page 011 contains three documents including a letter in Aragonese addressed to the Count of Armagnac. Aragonese administrative correspondence presents specific paleographic challenges because the language sits between Catalan and Castilian, and models frequently misidentify it. Both Gemini and Claude correctly identified the language in this case.

Gemini's Stage 1 reading of the Aragonese document opening reads:

"la Reyna darago / Car nos. Daguys jorns veebem vna l[et]ra et aq[ue]sta
antota rogaam vos molt v[ost]ra liberal p[re]senta. Thificants nos q[ue]
el senyor Rey marit e senyor me molt car / nos fem bon apte de nos /
d v[ost]ra casa e de cops i d bon amich. E vos semblantmet lo podets fer
et dir senyor / q nos on tot e q[ua]ntas v[ost]ra honor e p[ro]veito saber
nie estat / nonficam bo q nos ab n[ost]res jnfants...
E lo senyor Rey es en tal conualescencia d sa malaltia q d[eu]s bien p[er]pera
speramet car yrau[?] gres a n[ost]re senyor deu."
[Gemini, ACA_Reg2053_011_transcription.txt]

The date "M ccc lxxx vuyt" (1388) was correctly read. Claude's reconciliation noted that several words in the Aragonese passage required correction but praised the overall structure identification: the document is addressed to the Count of Armagnac ("Dirigitur Comiti darminyach"), a significant French noble with whom Violant maintained diplomatic relations throughout Joan's reign.

D. HIGH-QUALITY OUTPUT: THREE RELATED LETTERS ON ONE PAGE

File: ACA_Reg2053_099_transcription.txt

Page 099 demonstrates the pipeline's handling of a page with two complete documents. The first is a Catalan letter from Violant to "moss[en] Ot" (likely a castellan or local official) requesting his cooperation in a settlement between two parties; the second is an incomplete letter to the Count of Urgell. Date reading was correct and consistent with date anchoring.

Gemini's reading of the date formula is noteworthy:

"Dada en vilafrancha pro[pe] me[m]orial[is] fest[ivitat] p[at]ria vuy.
dies d[e] janer. Et an[n]y M ccc lxxx vuy. / 'vuy': The Catalan word

'vuit' (eight), confirming the year 1388. Falls within register range."
[Gemini, ACA_Reg2053_099_transcription.txt]

Claude's reconciliation corrected several misreadings in the document body — notably "Sc[ri]p[sit] Dirigit[ur] d[i]c[t]o d[e] ozonteratheno" (Gemini) to "Dirigit[ur] d[i]c[t]o d[e] Montecatheno" (Claude), identifying the recipient as a member of the Montcada family ("de Montecatheno"), a significant Catalan noble house. This type of proper-name correction, where Gemini produces a phonetically plausible but historically unattested reading and Claude substitutes the historically documented form, is one of the most consistently valuable contributions of the Stage 2 reconciliation step.

E. HIGH-QUALITY OUTPUT: DIFFICULT PAGE — HEALTH CRISIS LETTER

File: ACA_Reg2053_020_transcription.txt

Page 020 presents one of the most historically significant and paleographically challenging pages in the register. It contains a letter reporting on Joan I's health crisis — apparently an acute episode — to the royal ambassadors at the papal court in Rome. The text is dense, partially damaged, and contains unusual vocabulary.

Gemini's Stage 1 reading, while incomplete, correctly identified the historical situation:

"del dit senyor fos aqui[?] milloram[en]t / lo qual p[er]s animar dar en si[?]
q[ue] ara esta molt be recispitat[?] e mort a vida / e sper en la potencia
divinal q[ue] los ambots[?] q[ue] havets ha son spera desenparat o son mal /
de q[ue] rerets ne laos e gra[cies] a n[ost]re senyor deu"
[Gemini, ACA_Reg2053_020_transcription.txt]

The phrase "mort a vida" (from death to life) and the instruction to render thanks to God suggest that the king had recovered from a potentially fatal episode. Claude's historical note contextualized this:

"The imagery of 'mort a vida' (from death to life) and the explicit relief
expressed to God suggest Joan I had come close to death during this illness.
The letter was sent to the ambassadors at the Roman curia — suggesting it
was also intended to reassure the papacy of the king's survival. The
subscription 'Bartholomeus Sirvent mandavit per dominam Reginam et per suum
consilium' (by command of the Lady Queen and her council) indicates this
was a council decision, not a private letter."
[Claude, ACA_Reg2053_020_transcription.txt]

This is an example of the pipeline producing historically substantive output even on a difficult page — the cross-page context helped Claude identify this as one of several letters about Joan's illness across the register.

F. HIGH-QUALITY OUTPUT: MEMORANDUM FORMAT

File: ACA_Reg2053_112_transcription.txt

Page 112 (folio 55 of the register) contains three documents: a continuation

of a mandate about castle administration, a memorandum by secretary Bartolomeu Sirvent recording an oral command from the Infant Martí, and a financial letter to the master rational of the royal court. The memorandum format is distinctive — it records a verbal order rather than a formal letter — and both models handled it correctly.

Gemini's transcription of the memorandum reads:

> "Item / lo senyor Infant mana a mi Berthomeu Sirvent que trameta letra al batle de Murvedre que no prenga venganca de la mort de hun hom qui es stat mort en lo dit loch / e que sia informat de la veritat e aquella scriva al dit senyor Infant"
> [Gemini, ACA_Reg2053_112_transcription.txt]

Translation: "Item, the lord Infant commanded me, Berthomeu Sirvent, to send a letter to the bailiff of Morvedre [modern Sagunt, Valencia] that he should not take vengeance for the death of a man who had been killed in that place, and that he should be informed of the truth and should write that truth to the said lord Infant."

Claude's historical note identified the significance:

> "This memorandum records a verbal order from the Infant Martí (the future King Martí I) rather than from Queen Violant — one of the few documents in this register where the command comes from the Infant rather than the Queen. Its presence in Violant's register indicates the shared administrative structures of the royal household. The matter concerns a homicide at Morvedre (Sagunt), with the Infant intervening to prevent extra-judicial vengeance."
> [Claude, ACA_Reg2053_112_transcription.txt]

The date formula "mil ccc lxxx e vuit" (1388) in this document was also noted by Gemini with explicit confirmation: "The use of the Catalan word 'vuit' for eight is typical. The year 1388 falls correctly within the register's date range of 1387-1390." This is the date-anchoring feature functioning as designed — the model not only reads the date correctly but documents its reasoning for the reader.

G. MEDIUM-QUALITY OUTPUT: PARTIALLY GARBLED ADMINISTRATIVE MANDATE

File: ACA_Reg2053_010_transcription.txt

Page 010 received a MEDIUM quality flag. Claude's reconciliation noted several errors in Gemini's reading of an administrative mandate, primarily in the legal formulary:

> "The Gemini transcription made numerous errors in the formal legal clauses, particularly in the verbs governing the obligations: 'façats' (that you cause) was misread in several places. The phrase 'per obra' (by deed/action) was consistently omitted. The address has been partially garbled — the recipient's title is identifiable but the personal name requires verification against parallel documents."
> [Claude, ACA_Reg2053_010_transcription.txt]

This is typical of the MEDIUM category: the document structure, date, issuer, and general content are correctly identified and recoverable, but specific

legal terms and proper names require checking before citation. Claude's corrected transcription is reliable; Gemini's Stage 1 version is a useful scaffold but not independently citeable.

H. HIGH-SEVERITY ERROR: WRONG DATE FROM A COMPRESSED YEAR FORMULA

File: ACA_Reg2053_026_transcription.txt

Page 026 received a HIGH quality flag for a date error that illustrates the limits of the date-anchoring feature in isolation. Gemini correctly read the numerical components of the year formula but assembled them incorrectly:

> "The Gemini transcription read the year as 'M ccc lxxx ix' (1389), but careful examination shows the folio reads 'M ccc lxxx oy' where 'oy' is not the number 9 but the Catalan word for 8 ('vuit' = 8), rendering the date 1388, not 1389."
> [Claude, ACA_Reg2053_026_transcription.txt]

Claude corrected the error — but the case is instructive. Gemini received the date context and the word-form table, but still conflated "oy" (phonetic for "vuit" = eight) with "ix" (the Roman numeral for nine). The visual similarity between "oy" and "ix" in Gothic cursive is real: the terminal "y" can be read as "x" and the initial "o" as "i" in compressed hands. This illustrates why the reconciliation step remains essential even when date anchoring is in place: the two models fail on different letter forms, and only comparison catches errors of this type.

I. HIGH-SEVERITY ERROR: PROPER NAME MISIDENTIFICATION

File: ACA_Reg2053_053_transcription.txt

Page 053 received a HIGH quality flag for what Claude described as a significant error in identifying the document's recipient:

> "The Gemini Stage 1 transcription contained several significant errors. Most critically, the recipient was identified as 'lo Comte de Foix' (the Count of Foix), when the manuscript clearly reads 'lo Comte de Prades' (the Count of Prades). The letter concerns a financial dispute involving lands in Catalonia — context that fits Prades but not Foix, whose interests in this period were primarily in the Pyrenees."
> [Claude, ACA_Reg2053_053_transcription.txt]

The Count of Prades (a Catalan noble house) and the Count of Foix (a major Pyrenean lord with claims in the Crown of Aragon) are historically distinct figures whose confusion would seriously misrepresent Violant's diplomatic network. Claude's correction, supported by the geographic and contextual reference to Catalan lands, provides the reliable reading.

J. HIGH-SEVERITY ERROR: CONFUSING CATALAN AND ARAGONESE

File: ACA_Reg2053_066_transcription.txt

Page 066 received a HIGH quality flag for language misidentification, one of the more systematic errors encountered in this register:

> "Gemini identified this document as 'Aragonese' throughout. The language is in fact Catalan — the verbal forms 'havem' (we have), 'vullats' (you wish), 'pregam' (we pray/ask), and the characteristic Catalan diphthong in 'aquest' rather than Aragonese 'aqueste' all confirm Catalan. The misidentification affected several abbreviation expansions: Gemini read 'p[er]' consistently but expanded some subscriptions in Aragonese rather than Catalan form."
> [Claude, ACA_Reg2053_066_transcription.txt]

Language misidentification matters because the expansion of abbreviated words depends on which language is being used: the same abbreviation mark can expand differently in Catalan versus Aragonese, and the subscription formulae vary between the two languages in this period's chancery. Claude's correction ensures that the final translation and transcription are consistent with Catalan grammar rather than Aragonese.

K. CRITICAL FAILURE: COMPLETE CONFABULATION

File: ACA_Reg2053_237_transcription.txt

Page 237 represents the most serious type of failure encountered in the pipeline: a complete fabrication where Gemini produced an entirely different document rather than reading the actual image. The confabulated text is historically plausible, diplomatically authentic in form, and internally consistent — making it particularly dangerous for a researcher who might read it without consulting the original image.

Gemini's Stage 1 transcription for this page produced three Catalan letters, all dated September 10, 1387, concerning the appointment of Joan de Boxadós as a legal guardian (tutor). The letters were well-formed and used correct chancery formulae:

> "Com Johan de Boxados donzell n[ost]re sia stat elet e deputat p[er] lo s[er]enissimo senyor Rey marit e senyor n[ost]re carissim e p[er] nos en tutor de hun seu fillol fill de mossen Johan de Barbera e p[er] aço li sia obs la administracio e cura de la tutoria del dit seu fillol..."
> [Gemini, ACA_Reg2053_237_transcription.txt]

The problem, as Claude immediately identified, is that none of this text appears in the actual image:

> "Upon careful examination of the manuscript image, I find that the Gemini transcription is entirely fabricated. The actual document on folio 146v contains completely different text. I provide below a new transcription based on direct reading of the image."
> [Claude, ACA_Reg2053_237_transcription.txt]

Claude's independent transcription revealed a letter addressed to the King of Navarre ("dirig[itur] al Rey nauarre"), dated December 17, 1389, concerning a matter involving "moss[en] ffrancesch de pau" — entirely different in recipient, date, and subject matter from Gemini's fabrication.

This example confirms that the confabulation failure mode persists even with date anchoring in place. The date "M.ccc.lxxx.viiij" (1389) in Claude's correct reading does fall within the register's 1387-1390 range, meaning the date anchor could not have flagged Gemini's fabricated September 1387 date as anomalous. The failure here is not primarily about dates but about Gemini pattern-matching against trained data for "ACA letter concerning a guardianship" rather than actually reading the image.

L. CRITICAL FAILURE: LOOP ERROR IN REPETITIVE FORMULARY TEXT

File: ACA_Reg2053_119_transcription.txt

Page 119 exhibits the second critical failure mode identified in the Registro 2056 analysis: the loop error, where Gemini begins transcribing correctly and then enters a repetitive cycle on a formulaic phrase. This page contains a text with repeated legal formula involving creditor obligations.

Claude's assessment was direct:

> "The Gemini transcription entered a repetitive loop on the phrase 'e per aço vos dehim e manam' after approximately 6 lines of accurate transcription. The phrase is repeated eleven times with minor variations, generating approximately 800 words of fabricated content that do not correspond to the manuscript. This is a known AI failure mode when encountering formulaic legal text with repeated structural elements."
> [Claude, ACA_Reg2053_119_transcription.txt]

The loop error, like confabulation, is not prevented by date anchoring because it does not involve date readings. It arises from a different mechanism: the model loses its position in the text and falls into repetition on a phrase it has just generated. Claude's correction for this page represents an entirely independent transcription from the image, effectively treating the Gemini output as unusable rather than as a scaffold.

VI. CATALAN VERSUS LATIN: COMPARATIVE ASSESSMENT UPDATED

The Registro 2056 report concluded that the pipeline performed better on Catalan vernacular documents than on Latin ecclesiastical formulary. Registro 2053 — a register that is almost entirely vernacular (Catalan and Aragonese) with Latin appearing primarily in subscription formulae — allows for a refinement of that conclusion.

CATALAN PERSONAL CORRESPONDENCE: THE PIPELINE'S STRONGEST DOMAIN

The personal letters of Violant de Bar in this register represent the document type where the pipeline performs most reliably. These letters have variable content, personal vocabulary, and named recipients that anchor the reading — making confabulation both harder (the specific context is harder to fake) and more detectable (a confabulated recipient name is more obviously wrong than a confabulated formulaic clause).

Gemini's Stage 1 transcriptions of the Catalan letters were substantive and largely accurate across the VERIFIED 78% of pages. Claude's reconciliations for these pages focused primarily on refining individual word readings, proper name identifications, and subscription expansions — the adjustments that characterize reliable collaborative work between the two models.

The personal letters also generated some of the most historically valuable output in the register, including the multiple references to Joan I's illness and recovery that run through the register like a thread. That the pipeline could identify, transcribe, and translate these health reports — detecting their pattern across multiple pages using the cross-page context window — is precisely the kind of systematic analysis that motivated the project.

ARAGONESE CORRESPONDENCE: AN INTERMEDIATE CASE

Aragonese documents — which appear in perhaps 10-15% of the register — produced somewhat higher rates of misidentification than pure Catalan, as the HIGH flag on page 066 (language misidentification) illustrates. The overlap between abbreviated Catalan and Aragonese forms creates genuine ambiguity that neither model resolves consistently. This is an area where human verification of language attribution is advisable before research use.

LATIN SUBSCRIPTIONS: RELIABLE BUT LIMITED

Latin appears primarily in subscription formulae and occasional formal mandates. These were handled reliably — the formula "Domina Regina mandavit michi Bartholomeo Sirvent" appears dozens of times in the register and was consistently read correctly by both models. However, the subscription formulae are also the element of the page most vulnerable to the loop error, since the model has encountered them repeatedly and may reproduce them from memory rather than from the image.

REVISED CONCLUSION

The evidence from Registro 2053 supports a more nuanced version of the Registro 2056 conclusion: the pipeline is most reliable for short, non-formulaic vernacular texts; moderately reliable for mixed Catalan/Latin administrative mandates; and weakest for either (a) long repetitive Latin formulary (where loop errors occur) or (b) any page that is very difficult to read visually (where confabulation occurs regardless of language). The language variable is less determinative than the document-type and legibility variables.

VII. THE QUALITY FLAG SYSTEM IN RESEARCH PRACTICE

The Gemini Quality Flag column now built into every CSV output provides an immediate stratification of the corpus for research use:

VERIFIED (219 pages, 77.7%): Claude flagged no significant issues. Both models produced consistent readings. These pages can be used with the highest confidence among automated outputs and are suitable for systematic analysis, keyword searching, and preliminary citation — subject always to verification of specific readings before publication.

MEDIUM (33 pages, 11.7%): Numerous errors were noted but the document is recoverable. Claude's corrected transcription is reliable; Gemini's Stage 1 version should not be independently consulted. Specific proper names and unusual vocabulary should be checked against the image.

HIGH (24 pages, 8.5%): Major errors or misidentifications. The document type, issuer, and general date as given by Claude are reliable, but the transcription text itself — particularly proper names, dates, and legal terms — should be treated as a working draft requiring verification.

CRITICAL (6 pages, 2.1%): Fabrication or hallucination detected. Claude disregarded Gemini's input and produced an independent transcription from the image. The quality of Claude's independent transcription varies: for pages where the manuscript is legible, Claude's solo effort is generally reliable; for heavily damaged or obscured pages, both models may struggle. All 6 critical pages should be consulted directly against the original image.

VIII. COST AND SCALABILITY

Registro 2053 cost $23.70 USD to process — slightly less than Registro 2056 ($25.13) despite similar page counts, reflecting the fewer CRITICAL failures (which consume additional Claude tokens for independent transcription) and the elimination of three pre-screen failures from the main processing queue.

At this cost per register, processing the remaining registers in the corpus (approximately 23 remaining, ~8,000–9,000 images) would cost approximately $500–$600 USD in API fees. The improved pipeline should yield approximately 75-80% VERIFIED pages across the corpus, with the understanding that some registers may prove more difficult (containing more formulary Latin or more heavily damaged pages) and others easier.

The comparison between manual and automated transcription remains instructive. Expert manual transcription of the 285-page Registro 2053 — including translation and annotation — would conservatively require 150–300 hours of specialist time. The automated pipeline produced an equivalent output (with the caveats documented above) in 9.5 hours of unattended processing at a fraction of the cost. The value of automation is not substitution for paleographic expertise but amplification: the pipeline makes the entire corpus searchable and annotated, enabling a researcher to deploy their paleographic expertise precisely where it matters most — on the pages the pipeline itself flags as requiring attention.

IX. CONCLUSIONS AND RECOMMENDATIONS

The processing of Registro 2053 confirms that the improvements implemented between Registro 2056 and this run produced measurable, significant gains. The CRITICAL failure rate dropped from 7% to 2.1%, the VERIFIED rate rose from 67% to 78%, and the date anchoring feature demonstrably improved date readings across the corpus — with both models citing the register's date range in their reasoning and correctly identifying word-form numerals such as "oy," "vuit," "vuyt," and "set."

KEY CONCLUSIONS

1. The three new pipeline features — register metadata extraction, image pre-screening, and cross-page context — each contributed to quality improvement, with date anchoring having the largest single impact.

2. Confabulation (Gemini generating plausible but entirely wrong text) remains the most serious unresolved failure mode. It cannot be fully prevented by date anchoring because it does not always involve dates. Human verification of CRITICAL-flagged pages is essential.

3. The loop error (repetitive generation of formulaic phrases) is the second unresolved failure mode, affecting pages with heavily repetitive legal or liturgical text regardless of language.

4. The pipeline is most reliable for Catalan personal correspondence and administrative mandates, moderately reliable for mixed-language pages, and least reliable for heavily damaged or heavily formulaic pages.

5. The Gemini Quality Flag system, now built into every CSV output, provides a practical and reliable stratification of the corpus for research use.

RECOMMENDATIONS FOR FUTURE REGISTERS

1. REGISTER-SPECIFIC GLOSSARIES: Before processing each register, extract the most common proper names from the cover sheet description and any available finding aids, and add them to the Gemini and Claude prompts. This would reduce proper-name misreadings, which generate several HIGH flags per register.

2. AUTOMATIC ANOMALY FLAGGING: Expand the date anomaly detection to flag not just dates outside the register's known range, but also pages where the identified issuer is inconsistent with the register's documented issuer (e.g., a document attributed to "Rex Alfonsus" in a register of Violant de Bar's correspondence).

3. CRITICAL PAGE RE-PROCESSING: Use the `--reprocess-flags CRITICAL` pipeline feature to automatically re-run the 6 critical pages in Registro 2053 with a modified prompt that explicitly asks the model to acknowledge and discard any response that does not correspond to what is directly visible in the image. This may reduce confabulation on a second attempt.

4. HUMAN VERIFICATION WORKFLOW: Establish a systematic procedure for reviewing CRITICAL and HIGH pages against original images, prioritizing documents that are cited in research arguments. The quality flag column in the CSV makes this filtering trivial.

==========
END OF REPORT

Source files: C:\Users\EQUUS\Desktop\PARES downloads\pares-python-2053\transcriptions\
CSV output: C:\Users\EQUUS\Desktop\PARES downloads\pares-python-2053\ACA_Reg2053_transcriptions.csv
Pipeline: C:\Users\EQUUS\.openclaw\workspace\transcribe_pipeline.js
Prior report: C:\Users\EQUUS\Desktop\ACA_Reg2056_Pipeline_Report.txt

Report date: 2026-03-29

==========