

FIELD PRACTICE

What 527 sealed field review reports reveal about Ontario fieldwork

We read 527 sealed field review reports from two Ontario structural and building-science practices and classified every observation. Concrete findings about how sealed fieldwork actually gets written, what photo evidence looks like across the corpus, and where the revision chains are longest.

April 16, 2026

[DOWNLOAD THE INSIGHTS REPORT \(PDF\) ↓](#)

What this is, and why it matters

A field review report sealed by a Professional Engineer is a legal document. It records what an engineer saw on site, how it compared to the approved drawings and the Ontario Building Code, and what the contractor is expected to do next. When a sealed report is later read by a code official, a condo board, a tribunal, or another engineer reviewing the same work, the words in that report carry the force of the engineer's professional opinion.

Writing these reports is not a fringe activity. The two Ontario structural and building-science practices whose corpus we analysed produced more than 550 sealed reports between them, across the fourteen building-science and structural categories that dominate Ontario restoration practice - from balcony repairs and garage decks to curtain wall, roofing, and hydro vault rehabilitation.

What this article reports is not opinion. Every number below traces back to the classified corpus. The article exists because sealed fieldwork is undermetered: the industry produces tens of thousands of these documents a year in Ontario alone, but almost no structured data has ever been published about how they are written, what they contain, or where the drafting patterns break down.

We are publishing this because Fermito is building a drafting tool for sealed engineering work, and we want our own product decisions, and the category's broader conversation, to start from data instead of anecdote.

The corpus

527 sealed field review reports. 523 of them contain at least one numbered observation; the remaining 4 are schedules, interim progress memos, or template files that were included in the source folders but do not follow the observation-report format. 4,569 classified observations in total.

The reports come from two Ontario structural and building-science practices. The firm names are withheld. What matters for this analysis is the combined corpus: a full cross-section of Ontario restoration fieldwork written by licensed engineers for real projects, spanning multiple years and the fourteen categories listed above. No single firm's style dominates; the observations hold up as an industry sample.

What sealed fieldwork actually documents

Every observation in the corpus was classified by a large language model into one of 26 category labels covering concrete and structure, waterproofing and roofing, cladding and openings, fire and safety, and process work. We checked the model's

output on a sample for accuracy. The labels are specific enough to be useful and general enough to hold up across different firms' drafting voices.

The top ten categories, and what share of the corpus each one accounts for, are:

1. **Progress observation** - 869 observations (19.0%)
2. **Windows and doors** - 474 observations (10.4%)
3. **Waterproofing membrane installation** - 468 observations (10.2%)
4. **Balcony guardrail** - 416 observations (9.1%)
5. **Concrete placement** - 380 observations (8.3%)
6. **Roofing assembly** - 189 observations (4.1%)
7. **Drainage and water management** - 183 observations (4.0%)
8. **Sealant degradation** - 164 observations (3.6%)
9. **Painting and coating** - 151 observations (3.3%)
10. **Concrete deterioration** - 151 observations (3.3%)

The top five categories account for 57% of all observations; the top ten account for 75%. Restoration fieldwork in the GTA is dominated by a relatively narrow set of recurring conditions. Concrete deterioration, waterproofing failures, and the envelope repairs that follow from them make up the bulk of what engineers are writing about week after week.

Pillar-group rollup

Rolling the 26 labels up into five engineering pillars makes the pattern clearer:

- **Envelope and water** - 1,272 observations (27.8%)
- **Cladding and openings** - 1,185 observations (25.9%)
- **Process and reference** - 1,049 observations (23.0%)
- **Concrete and structure** - 897 observations (19.6%)
- **Fire and safety** - 166 observations (3.6%)

The 2,169 observations in the concrete-and-structure and envelope-and-water pillars combined are not a coincidence. Ontario restoration work is predominantly water management and the consequences of water management failure. When water gets into a reinforced concrete assembly, the resulting corrosion, spalling, and membrane reinstatement work shows up in field review reports for years afterward.

Report length and observation density

Reports in the corpus are shorter than many engineers expect. The median report is 553 words - shorter than the article you are reading - and the middle 50% of reports fall between 440 and 685 words. The longest 10% run past 872 words and tend to be final-review summaries covering multiple phases. The shortest 10% are progress notes of 371 words or fewer, usually produced during the middle of a multi-month repair when the engineer is visiting weekly.

Observation density matches this pattern. The median report contains **7 observations**, the 75th-percentile report contains 10, and the 90th-percentile report contains 16. Five to ten substantive observations per visit is the central tendency; the corpus rarely shows the kind of 30-item walkthrough list that some firm templates encourage.

This matters for drafting tooling. A generator that outputs 20 observations by default will feel wrong to most engineers. A generator that supports five to ten detailed observations per report - with the option to expand on the rare deeper review - matches what the corpus shows.

Photo evidence is near-universal, and photo density is high

90.3% of reports cite at least one photograph. The median report contains 5 photo references; the 75th-percentile report contains 7; the 90th-percentile report contains 9 or more.

51.6% of observations directly reference a specific photograph by number, using language like "Refer to Photos #3 and #4" or "(Photo 1)". Photo evidence is not

optional garnish in Ontario sealed fieldwork; it is woven into the observation itself. An observation without a photo reference is the exception, not the rule.

This has an uncomfortable implication for any drafting workflow. If the engineer cannot capture, caption, and attach photos to observations in the same session that produces the report text, the report is harder to write, and the photo-text linkage weakens. The corpus is effectively a proof that site-visit-to-sealed-draft workflows need photo-first tooling, not text-first tooling.

Regulatory citations are sparse at the observation level

2.6% of observations in the corpus cite a specific code section, CSA standard number, PEO regulation, or named drawing reference. The remaining 97.4% of observations carry the engineer's professional opinion about conformance with the drawings and specifications without anchoring that opinion to a named clause inside the observation text itself.

This is not a gap in the engineering; it is a pattern in the drafting. Individual observations in the corpus are short. The regulatory context typically lives in the report's opening boilerplate, the referenced specifications, and the shared project drawings. What the corpus does not show is engineers repeating the citation inside every single observation, the way a compliance-checklist template would.

Within the observations that do name a regulatory reference, the citations concentrate on a handful of families:

- Observation-level citations are rare enough that no family crosses the 10-citation threshold.

Interpretation. 2.6% is not a quality judgement. It is a structural fact about Ontario sealed fieldwork that matters for AI-assisted drafting. A tool that tries to force a code citation into every observation will produce reports that look wrong to experienced engineers, because they do not match the corpus baseline. A tool that surfaces the right citation only when the observation type calls for it - CSA material standards for

concrete placement, OBC Part 9 clauses for low-rise residential envelope, PEO Regulation 941 for sealing practice - is closer to what the corpus shows.

Revisions and multi-visit projects

1.3% of reports carry an explicit revision indicator in the filename or document header - "Rev 1", "Revision 2", "Re-issued", or an R-number suffix. This is a floor on the true revision rate because not every revised report is tagged that way; firms often replace the prior document without marking it as a revision.

A more reliable signal of multi-visit drafting is the project chain. Grouping reports by their project-root filename, 10.2% of projects in the corpus produced more than one sealed report. Most projects show up in the corpus once, usually because they were one-visit reviews or because only the final report was archived. The longest chain in the corpus reaches 52 reports; the 90th-percentile chain length is 1.

The most revision-heavy categories (among those with at least 10 reports):

- **Curtain wall** - 4.5% of reports carry a revision marker (1 of 22)
- **Windows and doors** - 3.0% of reports carry a revision marker (1 of 33)
- **Garage repair** - 1.9% of reports carry a revision marker (3 of 158)

The practical implication is that sealed drafting is not a one-shot activity. A firm producing 50 FRRs a month is re-issuing 10 to 15 of them as construction progresses. Drafting tooling that treats every report as independent misses the single biggest source of productivity leak in the workflow: the time spent re-reading, re-referencing, and re-constructing the context from the prior report when the next one is drafted.

Observation-to-recommendation ratio

A complete field review finding has three parts: what was observed, how it compares to the standard, and what should happen next. The corpus tells us the first two are nearly always present. The third is not.

34.5% of classified observations in the corpus contain an actionable recommendation. The remaining 65.5% report a condition without explicitly telling the contractor what to do about it. This is not a failure of the engineers; many observations are pure status notes ("progress is on schedule") where no recommendation is appropriate.

The observations that do carry a recommendation distribute across types as:

- **Document for record** - 1,678 observations
- **Contractor to clarify** - 495 observations
- **Rework / replace** - 451 observations
- **Repair per detail** - 251 observations
- **No action / acceptable** - 137 observations
- **Further investigation** - 110 observations
- **Monitor** - 30 observations
- **Retest after cure** - 18 observations

Categories with the highest recommendation density (observation types where engineers almost always close the loop):

- **Structural crack** - 72.7% carry an explicit recommendation
- **Waterproofing failure** - 71.3% carry an explicit recommendation
- **Sealant degradation** - 61.6% carry an explicit recommendation

Categories with the lowest recommendation density (observation types that tend to be status-only):

- **Photo reference only** - 1.3% carry an explicit recommendation
- **Progress observation** - 4.9% carry an explicit recommendation
- **Parking deck** - 22.2% carry an explicit recommendation

What the data does not say

It is worth being specific about what this analysis cannot tell you.

- **Nothing about quality.** The report analyses the shape of sealed fieldwork, not whether the underlying engineering was sound. A report that is short, lacks explicit code citations, and omits a recommendation may still be defensible in context.
- **Nothing about individual firms.** The corpus is blended across two practices. Per-firm breakdowns are deliberately excluded to protect the identity of the source firms.
- **Nothing about site outcomes.** We did not match reports against subsequent construction outcomes, warranty claims, or tribunal records. The report is a descriptive analysis of documentation patterns, not an outcome study.
- **Nothing about small or emergency reports.** Reports under 25 words (extracted text) were included in the count but did not produce observations. Those are likely one-line site memos, not failures of the corpus.

How FermiTo uses these findings

A short note, since the report is not a product pitch. FermiTo is building a drafting assistant for sealed engineering work. The corpus findings above directly shape what the product does:

- Photo-first, not text-first. 90.3% of reports cite photos; 51.6% of observations are photo-anchored. A drafting flow that starts with the photos and writes text around them matches the corpus shape.
- Five-to-ten observation default, not thirty. The median report has 7 observations, not 30.
- Recommendation suggestion as an option, not a default requirement. Only 34.5% of observations close the loop with a recommendation; forcing one on every observation creates reports that do not look like the corpus.

- Explicit revision chaining. 10.2% of projects produce multiple reports. Drafting each one from scratch leaves productivity on the floor.

Methodology

The corpus was loaded in full from two on-disk source directories. Every `.docx` file was extracted via the `mammoth` library; every `.pdf` file via `pdf-parse`. Where the same report existed in both formats, the `.docx` version was preferred. Duplicates across formats were deduplicated by normalized filename stem.

Firm names, engineer names, contractor names, project numbers, addresses, and municipal specifics were stripped from every observation before any downstream analysis. The anonymization ran on the raw extracted text using a combination of known-identifier regex and case-insensitive stem matching; the resulting text contains generic tokens like `[FIRM]`, `[PERSON]`, `[PROJECT]` in place of identifying content. A residual-leakage scan across the entire anonymized corpus returned zero hits against the sentinel list.

Observation classification used Anthropic's Haiku 4.5 model. Each observation was passed with the full anonymized context and assigned a primary category, an optional secondary category, a recommendation type, and a boolean recommendation flag. Classifications were cached on disk by content hash so the analysis is deterministic and re-running after a successful first pass is near-free.

Statistics were computed deterministically from the classified output. No claim in this article required the model to perform the aggregation step; the model's job was category assignment, not statistical reasoning.

[DOWNLOAD THE INSIGHTS REPORT \(PDF\) ↓](#)

Every number in this article traces back to a field in a machine-generated stats file. If you run sealed work in Ontario and want the full categorical breakdown for your own team's training, contact FermiTo.

THE FERMITO NOTES

One email when a new piece or template drops.

Practice notes on sealed engineering work, the regulation around it, and what AI drafting does and does not change. No cross-promotion, no re-selling the list.

Subscribe

ONE EMAIL PER RELEASE. UNSUBSCRIBE ANY TIME.

[← All articles](#)



The drafting assistant for engineers who sign. Built for Ontario structural and building-science practice.

ENGINEERS SIGN. FERMITO DRAFTS.

Capture

For principals

About

Generate

For firms

Security

Revise

Case study

Pricing

Export

Articles

Report templates

Book a demo