

TECHNICAL BRIEF

How WelWrit Reviews a Work

Method, evidence, and the line between what is measured and what is projected.

WelWrit is a literary review engine. It reads an entire work and reviews it against seven literary standards, returning a marked-up copy and a report in which every score is tied to the exact words — or the public record — that produced it. This brief states plainly how that is done, what the engine measures directly, what it projects, and what it deliberately does not claim. We publish our method because a review instrument is only as trustworthy as the reasoning a reader can inspect.

1. One principle: every score is evidence

WelWrit does not report how a work made it feel. Every statement it produces rests on something checkable: a cited passage, a structural count, or an external record. A score is never shown alone — it is an index into the evidence beneath it, and a reader can always trace a number back to the passages or records that produced it.

This is a deliberate constraint. It rules out flattery and impression, and it means the only way to change a score is to change the work. The analysis is rigid by design.

2. The seven standards, and how each is assessed

Most standards are measured directly from the text. One — Permanence — is measured from public records when a work is published, and projected from comparable features only when it is not. The engine always states which basis produced a given assessment.

<p>Style</p> <p>The author's distinct way of expressing ideas and using literary devices.</p>	<p>Direct stylometric measurement</p>
<p>Suggestiveness</p> <p>Symbolism, nuance, and implied meaning carried by the text.</p>	<p>Motif & symbol structure mapping</p>
<p>Artistry</p> <p>Aesthetic craft — sound, parallelism, figurative density.</p>	<p>Craft analysis vs. a literary corpus</p>
<p>Intellectual Value</p> <p>Ideas that stimulate thought and reveal something true.</p>	<p>Idea-development mapping, cited</p>
<p>Spiritual Value</p> <p>Moral and human weight — values placed in tension.</p>	<p>Moral-landscape mapping, grounded</p>
<p>Universality</p> <p>Reach of theme and experience across readers.</p>	<p>Thematic-scope analysis</p>

Permanence

Records if published; projection if not

Endurance and continued relevance over time.

Why state the basis of each score? Because a literary work's reach and endurance unfold over time. For a published book those are facts in the record — library holdings, editions, reprints, citations, translations. For an unpublished manuscript they cannot yet exist, so the engine projects from measurable comparable features and says so. Conflating the two would be the easiest way to mislead, so we don't.

3. How the engine works

WelWrit is a pipeline, not a single opaque model. A deterministic backbone performs the measurable structural analysis; a language layer writes the interpretive context on top of it. The split is intentional: the measurable parts stay explainable, and every interpretive claim is required to cite the span that grounds it.

- Ingestion: the work is parsed and segmented into addressable spans, with exact offsets preserved so findings can be anchored to the page.
- Deterministic analysis: stylometry, motif and symbol mapping, craft features, citation auditing, and plagiarism screening — computation over the text, fast and explainable.
- Records layer: for published works, library and bibliographic databases supply real endurance data for Permanence.
- Interpretive layer: writes the supporting context for each standard, under a hard rule — no claim without a cited span.
- Assembly: evidence is merged into a marked-up PDF and a report, each carrying an intellectual-property statement.

4. What WelWrit will not do

- It does not write for the author. The engine reviews; it never drafts, rewrites, or completes prose. The words under review remain wholly the author's.
- It does not certify literary greatness. WelWrit reports an evidence-based review, not a verdict on merit — a claim no instrument can defend.
- It does not blend originality with quality. Plagiarism screening is a separate engine, reported as fact with sources, never folded into a quality score.
- It does not present a projection as a measurement. Where Permanence is projected rather than drawn from records, the engine says so.

5. Where a human comes in

The engine performs the review; a human expert is available by escalation. An author can dispute a finding or request expert human review as a premium add-on, and for cases the engine flags as low-confidence a person adjudicates. The reviewer sees the full evidence trail and may confirm, revise, or override any finding, with their rationale recorded — but does not rewrite the author's work either.

6. Performance and benchmarks

WelWrit intends to set the standard for automated literary review, which raises rather than lowers the bar for honesty: published performance figures must be measured on instrumented runs, not asserted. A benchmark that cannot be reproduced is not a benchmark.

Performance is anchored to fixed reference works, every pipeline layer is timed independently, and each result is published with its conditions — hardware, model version, and concurrency — so others can reproduce it. Speed is reported only at a fixed quality bar, so the engine is never quietly traded down for a faster number.

Status: the figures behind our performance targets are pending validation by benchmark runs. Until those runs exist, WelWrit states no specific average processing time as fact. Setting the benchmark means earning the number, then publishing how it was earned.

This brief reflects WelWrit's methodology as currently designed. It will be updated as the engine is built and benchmarked. Authorship of every reviewed work remains entirely with the author.