

Measuring Readability at Scale

How iAccessible uses temporal stratified sampling to produce defensible, cost-efficient website readability metrics

Core approach. iAccessible begins with the complete crawler inventory, groups pages by when they were created or last modified, randomly samples within each time-based group, and applies statistical weights so results represent the full website, not just the pages selected for analysis.

WHY SAMPLING IS NEEDED

A full readability assessment of every page can be unnecessarily expensive and slow for large digital estates. A site with tens or hundreds of thousands of pages may contain substantial duplication, legacy content, and pages that have not changed in years. At the same time, recently created or modified pages deserve greater attention because they are more likely to reflect current publishing activity and emerging content-quality risk.

The iAccessible approach preserves statistical validity while concentrating analytical effort where change is occurring. It also supports configurable API call budgets, recurring monitoring, and time-based comparisons. iAccessible already provides large-scale readability and content-quality monitoring as part of its broader accessibility and web-governance capabilities.

METHODOLOGY: FROM CRAWL UNIVERSE TO REPRESENTATIVE METRICS

1	2	3	4	5
Crawl universe	Temporal strata	Sample size & allocation	Random selection	Weighted results & trends

Establish the sampling frame. The crawler identifies the eligible webpage universe for each domain or selected crawl depth. Duplicate, redirected, non-HTML, excluded, and failed URLs are handled under documented eligibility rules before sampling.

Assign temporal strata. Each eligible page is classified using creation date, last-modified date, and first-seen information from the crawler. Pages with missing or unreliable dates are retained in a separate “unknown metadata” stratum so they are not silently excluded.

Set the statistical target. The standard design uses a 95% confidence level. For estimating the percentage of pages meeting a readability target, a ± 5 percentage-point margin of error and a conservative 50% expected rate produce a baseline sample of about 384 pages for a 200,000-page site, before domain minimums or design adjustments. A finite population

correction is applied for smaller sites, and a census may be used when the eligible universe is sufficiently small.

Allocate and draw the sample. The baseline allocation is proportional to stratum size, with minimum sample floors and purposeful oversampling of newly created and recently modified pages. Selection is random within each stratum. Pages that fail analysis are replaced only from the same stratum using a pre-generated reserve list.

Weight and report results. Each analyzed page receives a weight based on the number of eligible pages it represents. Weighted estimates are used for website-wide rates and average scores; unweighted results are also retained for operational drilldown and remediation.

ILLUSTRATIVE TEMPORAL STRATA

Stratum	Illustrative rule	Sampling emphasis
New	Created or first seen within 30 days	Highest selection rate
Recently modified	Modified within 90 days	Elevated selection rate
Established	Last changed 91–365 days ago	Proportional coverage
Stable / legacy	Last changed more than 365 days ago	Representative coverage
Unknown metadata	Creation or modification date unavailable	Minimum floor; separately reported

LONGITUDINAL MONITORING AND SAFEGUARDS

For recurring scans, iAccessible can retain a stable panel of previously sampled pages while refreshing the remaining sample. This supports apples-to-apples trend measurement without losing sensitivity to new content. Newly created and recently modified pages automatically enter the eligible frame and receive the configured higher probability of selection.

The platform records the sampling frame, stratum definitions, random seed or selection record, page weights, API outcomes, and substitutions. Results can therefore be reproduced, audited, and explained. Estimates are accompanied by sample counts and confidence information, and findings can be viewed by domain, time stratum, content group, or individual page.

HOW CUSTOMERS USE THE RESULTS

Enterprise view. Weighted measures summarize the estimated percentage of pages meeting a selected readability target, the distribution of reading-grade levels or scores, and changes from prior measurement periods.

Change-focused view. Results for new and recently modified pages are displayed separately, helping communications and web-governance teams determine whether current publishing practices are improving or introducing new readability concerns.

Action view. Page-level findings remain available for drilldown, prioritization, and remediation. Customers can focus first on frequently visited, mission-critical, newly published, or repeatedly underperforming content while preserving an unbiased website-wide benchmark.

Customer value. Temporal stratification gives customers a practical balance: representative website-wide readability metrics, stronger visibility into current publishing activity, controlled analysis costs, and defensible evidence for content-governance decisions.

Important interpretation note: Automated readability scores are decision-support indicators. They should be interpreted with audience, purpose, required terminology, accessibility testing, and human editorial review. Readability analysis does not by itself establish WCAG or Section 508 conformance.