



Independent AI Tool Ratings and Evaluations



Copyleaks

Product Research & Evaluation

Produced on October 27, 2026

Table of Contents

1.0 Quality	7
1.1 Benchmarking and Validation	7
1.2 Detection Accuracy and Performance	10
1.3 Detection Explainability and Confidence	14
1.4 Model Adaptability and Coverage	17
1.5 Robustness Against Evasion	21
1.6 Usability and Workflow Integration	23
2.0 Privacy	25
2.1 Artificial Intelligence Data Usage	25
2.2 Data Governance and Ownership	27
2.3 Privacy Compliance and Regulatory Alignment	28
2.4 User Privacy Controls	29
3.0 Security	30
3.1 Application and API Security	31
3.2 Detection Engine Integrity	34
3.3 Identity and Access Management	36
3.4 Platform Security	37
3.5 Security Assurance and Incident Response	39
4.0 Impact, Ethics, and Safety	41
4.1 Abuse Prevention and Public Trust	42
4.2 Detection Fairness and Transparency	43
4.3 Responsible Artificial Intelligence Governance	45

4.4 Responsible Detection Practices	46
5.0 Sustainability and Reliability	47
5.1 Detection Model Maintenance	47
5.2 Detection Service Reliability	48
5.3 Environmental and Social Responsibility	51
5.4 Operational Excellence	53
5.5 Vendor Sustainability and Business Viability	54
6.0 Use Cases & Pricing	57
6.1 Commercial Packaging and Pricing	57
6.2 Detection Analysis and Reporting	60
6.3 Enterprise Deployment and Administration	61
6.4 Market Adoption and Customer Success	62
6.5 Workflow Integrations and Extensibility	65
Disclosures	68

How to Read This Report

This public report is part of huby's commitment to product transparency. It evaluates Copyleaks across six evaluation categories and 29 subcategories. At every level the product is scored on a scale of 1–5 and accompanied by a written assessment. Each score is backed by verifiable public references to ensure an objective assessment for our community.

Product Classification

This report evaluates **Copyleaks** as an **AI-detection** product, subtype **Text-detection**. Classification determines which subcategories and atomic factors are activated within the evaluation framework, and which peer products form the basis of comparison.

Scoring Framework

All scores are reported on a **1.0–5.0 scale** at category and subcategory levels (and at atomic factor level in the owner report). Scores between integers reflect graduated evidence quality within a band (e.g., 3.7 indicates the upper range of 'Adequate' but below the 'Strong' threshold).

Score	Band	Definition
5.0	Exemplary	Strong, independently verified evidence of best-in-class practice. Formal audits, third-party certifications, or peer-reviewed assessments confirm performance.
4.0 – 4.9	Strong	Multiple credible, corroborated sources substantiate the claim. Minor gaps exist but do not materially undermine the finding.
3.0 – 3.9	Adequate	Evidence is present but limited, mixed, or partially corroborated. The product meets a basic standard but with notable gaps or inconsistencies.
2.0 – 2.9	Weak	Evidence suggests meaningful shortfalls. Claims made by the product owner are not substantiated by independent sources, or documented failures are present.
1.0 – 1.9	Deficient	Material failures or absence of the capability/control. Significant risk to users or enterprise buyers exists.

The overall product score is a weighted aggregate of the six category scores. The weights below reflect the relative importance of each evaluation dimension for the **Text-detection** product subtype.

Category	Weight	Subcategories
Quality	30%	Detection Accuracy and Performance (30%) Robustness Against Evasion (22%) Model Adaptability and Coverage (16%) Detection Explainability and Confidence (14%) Benchmarking and Validation (10%) Usability and Workflow Integration (8%)
Privacy	16%	Data Governance and Ownership (30%) Artificial Intelligence Data Usage (28%) Privacy Compliance and Regulatory Alignment (24%) User Privacy Controls (18%)
Security	17%	Detection Engine Integrity (24%) Platform Security (22%) Application and API Security (20%) Security Assurance and Incident Response (20%) Identity and Access Management (14%)
Use Cases & Pricing	13%	Commercial Packaging and Pricing (24%) Enterprise Deployment and Administration (20%) Workflow Integrations and Extensibility (20%) Market Adoption and Customer Success (20%) Detection Analysis and Reporting (16%)
Sustainability and Reliability	10%	Detection Service Reliability (28%) Vendor Sustainability and Business Viability (22%) Detection Model Maintenance (20%) Operational Excellence (18%) Environmental and Social Responsibility (12%)
Impact, Ethics, and Safety	14%	Detection Fairness and Transparency (30%) Responsible Detection Practices (26%) Responsible Artificial Intelligence Governance (24%) Abuse Prevention and Public Trust (20%)

Overall Product Score: 3.7 · Adequate

Competitive Context

Where sufficient public data exists, scores are contextualised against the assessed product's competitive set. This does not change the score itself; it gives you relative positioning. The cohort below is every **Text-detection** product huby has evaluated under the same framework, with the same factors and the same tests applied to each.

	Copyleaks	Cohort median	Range	Rank
Overall	3.7	3.2	2.7 – 3.7	1 of 9
Quality	3.8	3.5	2.8 – 3.8	1 of 9
Privacy	3.5	3.3	2.3 – 3.5	1 of 9
Security	3.5	2.8	2.3 – 3.5	1 of 9
Use Cases & Pricing	4.2	3.5	3.0 – 4.2	1 of 9
Sustainability And Reliability	3.8	3.3	2.6 – 3.9	2 of 9
Impact, Ethics, And Safety	3.6	2.9	1.8 – 4.1	2 of 9

Peer products are reported as a median, a range and a rank rather than named individually: the cohort's individual scores belong to those products' own reports.

Independence and Conflict Management

huby is independent and does not charge product owners for evaluation, for listing, or for either report. There is no paid tier, no expedited review, and no way for a product owner to pay for placement, for a score, or for a score to be reconsidered.

Our analysis is not influenced by a product owner, and we fully disclose any work that we do with product owners. Owners can control certain visibility settings, but these do not govern the evaluation, the scores, or the written assessment, none of which are shared with the owner for approval before publication.

Disclosure for this report: huby has no commercial relationship with Copyleaks. Should any commercial relationship with a product owner exist in future, it will be disclosed in the relevant report. The full methodology, including the independence policy in sections 5.3 and 6.1, is published at huby.ai/methodology.

Citations and References

Throughout the report, inline citations appear as numbers in brackets, for example [1] or [3]. Each number corresponds to a source listed in the **References** section at the bottom of an area of analysis. Click a citation to open its source. Screenshots and other visual evidence are numbered separately as **Figure 1**, **Figure 2** and so on, running in a single sequence across the whole report. Each figure is shown once, at its first mention; later sections that rely on the same evidence point back to that figure number. Only sources with a verifiable URL are shown — citations without a retrievable source are omitted to keep the report clean.

1.0 Quality

Score: **3.8** · *Adequate*

The product's performance is characterized by a pattern of specialized strength that is undermined by fundamental weaknesses. It performs well against complex, meaning-based evasion methods but fails when faced with simple technical alterations or edited content. This variability is echoed in conflicting accuracy reports from different benchmarks and user-reported unreliability, which detracts from the system's otherwise strong workflow integrations. A significant gap exists between the vendor's claims and independently verifiable performance. The company's high-level accuracy statements are unsubstantiated by public data, and advertised capabilities like broad language support lack third-party validation. This evidence gap extends to the product's own output, as its confidence metrics are poorly calibrated and its detailed explanations cannot overcome the inconsistent detection results reported across multiple sources.

1.1 Benchmarking and Validation

Score: **4.4** · *Strong*

Copyleaks has achieved significant real-world validation, establishing itself as a standard tool in the AI detection field. Its widespread adoption is confirmed by integrations into major educational platforms like Canvas and Moodle [3] [8]. The product's market relevance is further highlighted by its inclusion in numerous independent comparisons [1] [4] [7] and by the fact that it is a specific target for "humanizer" tools designed to bypass AI detection [5]. This standing is supported by some strong independent performance metrics, including a reported 98.3% accuracy on one benchmark [1] and a finding in a peer-reviewed paper that it was resistant to certain adversarial attacks [1]. Despite its widespread use, the tool's performance is inconsistent across different evaluations, and the company's accuracy claims lack transparency. While one benchmark showed high accuracy, a separate academic study found its performance dropped to just 58.7% when tested against human-edited AI outputs [2] [4]. This stands in stark contrast to the company's marketing, which claims over 99% accuracy based on "independent third-party studies" [5] [6]. However, Copyleaks does not publish or provide links to these studies, making its central performance claims unsubstantiated [3] [5] [6]. Furthermore, while the company publishes high-level methodology documents, they lack the dataset and procedural details required for the public to independently verify its results [1] [2] [3]. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 1 — Screenshot of Copyleaks AI Detector product page claims




Figure 2 — Curated Evidence: Copyleaks marketing claim of 99% accuracy




Figure 3 — Image: GPTZero vs. Copyleaks AI Detector Comparison

References:

- [1] GitHub It-s-AI/llm-detection: It's AI — github.com/It-s-AI/llm-detection
- [2] Pangram 4 Technical Report arXiv — arxiv.org/html/2607.27183v1
- [3] AI for scientific integrity: detecting ethical breaches, errors... PMC — pmc.ncbi.nlm.nih.gov/articles/PMC12436494
- [4] Copyleaks AI Detector Most Accurate by 3rd-Party Studies — copyleaks.com/blog/ai-detector-continues-top-accuracy-third-party
- [5] Copyleaks Text Moderation Testing Methodology — copyleaks.com/text-moderation/testing-methodology
- [6] Copyleaks AI Detector Testing Methodology — copyleaks.com/ai-detector/testing-methodology
- [7] Copyleaks Case Study: AI for Detecting Plagiarism Info-Tech — www.infotech.com/research/copyleaks-case-study-ai-for-detecting-plagiarism
- [8] Protecting Your Work: Tracking Content Use for Plagiarism and... Silverchair — www.silverchair.com/news/protecting-your-work

1.2 Detection Accuracy and Performance

Score: **3.8** · *Adequate*

In a preliminary huby test, Copyleaks demonstrated an adequate but not exceptional ability to distinguish between AI and human writing, placing 6th out of 9 products tested and below the median for the group [1]. The tool's performance was strongest on straightforward content; in the limited test, it perfectly identified all purely AI-generated and purely human-written documents, achieving a 0% false positive rate on that sample (Figure 4) [1]. However, this strength did not extend to more complex cases. The detector failed on the single sample designed to evade detection and the single sample of AI text edited by a human, indicating a significant weakness against modified content [1]. This vulnerability is corroborated by an independent study that found accuracy dropped to 71% against 'humanized' AI text [2]. The product's real-world performance appears inconsistent, creating a discrepancy between controlled test results and user experience. While huby's test found no false positives, credible contradictory evidence includes a documented video showing a human article flagged as 100% AI and multiple user reports of the tool flagging human writing as AI-generated (Figure 5) [2]. This inconsistency, along with the measured below-median ranking, stands in contrast to the vendor's claim of over 99% accuracy [1] Figure 6. While the tool is notably fast, with one test measuring a 510ms response time (Figure 4), there is a lack of independent performance data for key features like long document analysis and scanned text detection (Figures 4 and 6). Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 4 — huby Test Capture: copyleaks result page for corpus sample ai_pure-336cee5261b12d92

5 Best AI Content Detection Tools

	Originality.ai	GPTZero	Turnitin	Copyleaks	ZeroGPT
Strengths	Very strong detection especially for paraphrased/edited AI content; good multilingual support; includes plagiarism checks.	Good balance of affordability and detection; performs well in academic contexts; offers free/freemium tiers; decent performance with raw AI-text.	Widely trusted in academia; massive database; good for institutional use; high detection of full AI content.	Strong multilingual support; good in third-party benchmarks; works well with paraphrased/rewritten content in certain cases.	Decent speed; free or lower-cost entry; good for quickly checking obvious AI content; decent for users who need lightweight detection.
Weaknesses	No free tier; costs can add up for high volume; even this tool can sometimes false-flag human text or struggle with highly hybrid content.	Accuracy drops with heavily edited/paraphrased AI text; some false positives for polished human writing; might lack as many advanced features (reporting, multilingual) compared to more expensive tools.	Struggles more with hybrid content or paraphrased AI; sometimes high false positive concern when content is very polished or formulaic; pricing/licensing tends to be institution-oriented.	According to users, sometimes flags human writing as AI; results may be inconsistent across content types; can lag; interface/polish may be less refined compared to premium tools.	Lower performance on edited/paraphrased AI content; higher false positives for formal or polished human writing; sometimes inconsistent on long/mixed content.
Best For	Publishers, SEO/content agencies, teams that need high confidence; or academic settings that want both plagiarism + AI detection.	Educators, students, people checking drafts or lighter content; anyone who wants a good detection tool without big cost.	Universities, schools, formal academic assessment settings; places needing solid policies and institutional oversight.	Enterprise / higher education needing both plagiarism + AI detection; multilingual content; larger document sets.	Quick checks; casual users; people wanting a lightweight tool just to flag obvious machine-written bits, not for high-stakes adjudication.

dNOVO Group | dnovogroup.com

Figure 5 — Curated Evidence: Copyleaks_image-7.png

ILLUSTRATIVE — shows the interface, not the measured result

The screenshot displays the huby interface for document scanning. At the top, it shows 'My Scans > New Scans 2:38 PM' and a 'Credits Left' indicator. The main document area shows a file named 'your-file.html (300 words)' with a 'Matched Text' progress bar at 27.7% and an 'AI Content' progress bar at 0%. The document text is visible, with some parts highlighted in red. The right sidebar, titled 'Matching Text Results (5)', shows a 'Categorized Summary' with 'Internet 100%' and '5 Results'. It lists several search results, including 'your-approach-hiring-all-wrong-ashish-singh' and 'Your Samurai's Guide To Achieving Worklife Balance - SamuraiNOW'. The bottom of the interface shows 'Powered by Copyleaks'.

Figure 6 — huby Test Capture [illustrative, different run]: copyleaks reported 0.0% AI on a sample whose true AI content is 100% (




Figure 7 — Curated Evidence: Copyleaks_05-copyleaks.jpg

References:

- [1] huby Test Data: results_copyleaks_20260921-manual.jsonl — storage.googleapis.com/.../results-copyleaks-20260921-manual-4b0635eb.jsn
- [2] The Do's and Don'ts You Need To Know Before Utilizing AI Detectors — copyleaks.com/.../the-dos-and-donts-you-need-to-...-before-utilizing-ai-detectors
- [3] [2507.17944] Evaluating the Performance of AI Text Detectors, Few... — arxiv.org/abs/2507.17944
- [4] Curated Evidence: Copyleaks_2507.17944v1.pdf — storage.googleapis.com/huby-media/copyleaks/evidence/2507-17944v1-0fb05bc0.pdf
- [5] AI Detector Free AI Checker for ChatGPT, GPT-5, Gemini & More — copyleaks.com/ai-detector
- [6] Copyleaks: Plagiarism Detector SMU — www.smu.edu/oit/services/copyleaks
- [7] ChatGPT and Academic Integrity Concerns: Detecting Artificial... — www.langedutech.com/letjournal/index.php/let/article/view/49
- [8] OCR Supported Languages Copyleaks API Docs — docs.copyleaks.com/reference/actions/miscellaneous/ocr-supported-languages
- [9] Pangram 4 Technical Report arXiv — arxiv.org/html/2607.27183v1
- [10] Curated Evidence: Copyleaks_Winston_AI_vs_Copyleaks_Best_AI_Detector_Accuracy_Test_False_Positives_Real_Results.mp4 — storage.googleapis.com/.../winston-ai-vs-copyleaks-b...se-positives-4e97c203.mp4
- [11] AI Detector Comparison 2026 — GPTZero vs Originality.ai vs... — aidetector.ac/compare

1.3 Detection Explainability and Confidence

Score: **3.9** · Adequate

Copyleaks provides strong, multi-faceted explanations for its findings, particularly for plagiarism detection. The platform offers a clear, side-by-side comparison that highlights matching text and provides direct links to the original source materials, creating a verifiable evidence trail for users [4] [6]. For AI detection, the tool also operates at a granular level, highlighting specific sentences and assigning each a distinct probability score, which provides a strong

basis for explainability [5]. This is enhanced by a feature called "AI Logic," which aims to explain the reasoning behind a detection to improve transparency [1] [3]. However, the explanation for AI detection is inherently probabilistic and cannot provide a citable source like a plagiarism report, a fundamental limitation of the underlying technology. While the explanations are detailed, the reliability of the confidence scores themselves presents a more complex picture. In huby's controlled evaluation, Copyleaks' confidence scores had an average error of 17.6 percentage points, meaning a score of 80% confidence does not reflect a true 80% probability [2]. Although this absolute error is significant, the tool's performance was strong relative to peers, ranking 2nd out of 9 products tested [2]. This measurement, however, was based on a preliminary sample size of 17 documents, which is below the preferred minimum [2]. The finding of overconfidence is corroborated by user reports and third-party analysis, which note issues with false positives and inconsistency where human writing is flagged as AI-generated [3]. This can create ambiguity, as can potential user confusion between the percentage of text flagged as AI and the separate confidence score assigned to those flags [2] [5]. The tool's core algorithms are designed to be deterministic, providing a foundation for reproducible results against a fixed set of data (Figure 4). However, its real-world consistency is challenged by its reliance on scanning the live internet, where source material is constantly changing (Figure 4). More significantly, user reports suggest that results can be inconsistent across different types of content, indicating a potential gap between the tool's theoretical design and its practical performance [2]. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 8 — User interface of Semihuman.ai showing Copyleaks detection categories




Figure 9 — Curated Evidence: Copyleaks_Copyleaks-Screenshot-1-1600x973.png




Figure 10 — Curated Evidence: Copyleaks_Copyleaks-1.png

References:

- [1] Does Copyleaks Detect AI? Accuracy and Workflow Guide | Leap — www.tryleap.ai/blog/does-copyleaks-detect-ai
- [2] Understanding Copyleaks AI Detection Results LTSA — ltsa.sheridancollege.ca/digital-learning-support-hub/copyleaks-ai
- [3] huby Test Data: results_copyleaks_20260921-manual.jsonl — storage.googleapis.com/.../results-copyleaks-20260921-manual-4b0635eb.jsonl
- [4] Copyleaks AI Logic, AI Source Match Faculty & Staff Support — staffsupport.spcollege.edu/.../Copyleaks-AI-Logic-AI-Source-Match
- [5] Copyleaks Plagiarism and AI Content Detector plugin Moodle plugins — marketplace.moodle.com/plugins/plagiarism_copyleaks
- [6] AI Content Detection via Open LMS and Copyleaks Partnership — www.openlms.net/blog/news/open-lms-partners-with-copyleaks
- [7] Moodle Forms Solution Partnership With Copyleaks for AI Content... — www.newswire.com/.../moodle-forms-solution-partne...-for-ai-content-and-22019907

1.4 Model Adaptability and Coverage

Score: **3.8** · Adequate

Copyleaks presents a comprehensive and forward-looking approach to content authenticity, bundling AI and plagiarism detection into a single, integrated tool [1] [3]. In preliminary huby testing, the platform performed perfectly, showing zero variation in its ability to detect content from different AI models and ranking first among the six products tested on the same material [1]. This result, however, is based on a small sample size and should be considered with caution [1]. The company's strategy is built on a model-agnostic framework, supported by academic research, which is designed to identify text from new and unseen language models without requiring specific training [4]. This allows Copyleaks to claim immediate coverage for emerging models like Claude and Grok, a key feature in the rapidly evolving AI landscape (Figure 4) [2]. While the company's strategy is well-defined, its performance claims are not consistently supported by independent, third-party evidence. For instance, Copyleaks advertises extensive multilingual capabilities, with its AI detector supporting over 30 languages, but no independent data was available to validate the accuracy of this coverage [7]. Furthermore, independent benchmark testing on the tool's ability to generalize across different content domains shows mixed results, with one test yielding a score of 58.7%, indicating its performance can be inconsistent when faced with novel or edited AI-generated text [3]. This lack of transparency

also applies to its model update frequency, as Copyleaks provides no public documentation or changelog detailing how often its detection engine is improved to counter new AI technologies. Supporting charts and figures illustrating these findings are included below.

Supporting Evidence:




Figure 11 — Curated Evidence:
 Copyleaks_Best-AI-Content-Detectors-2026-Originality.ai-vs-Copyleaks-vs-Winston-AI-Compared-1200x670.png




Figure 12 — Collected Image: Pangram 4 Technical Report

References:

- [1] huby Test Data: results_copyleaks_20260921-manual.jsonl — storage.googleapis.com/.../results-copyleaks-20260921-manual-4b0635eb.js onl
- [2] Copyleaks APIs | Real-Time Content Integrity Solutions — copyleaks.com/api
- [3] AI Detector Free AI Checker for ChatGPT, GPT-5, Gemini & More — copyleaks.com/ai-detector
- [4] AI Detection Supported Languages Copyleaks API Docs — docs.copyleaks.com/.../ai-detection-supported-languages
- [5] Cross-Language Detection Copyleaks API Docs — docs.copyleaks.com/concepts/features/cross-language-detection
- [6] Maintaining Academic Integrity with the Copyleaks API — docs.copyleaks.com/concepts/use-cases/academic-integrity
- [7] Enterprise Content Governance Copyleaks API Docs — docs.copyleaks.com/concepts/use-cases/enterprise-content-governance
- [8] Copyleaks Plagiarism and AI Content Detector plugin Moodle plugins — marketplace.moodle.com/plugins/plagiarism_copyleaks
- [9] Copyleaks: AI Content & Text Authenticity Detection — copyleaks.com
- [10] What's New Copyleaks API Docs — docs.copyleaks.com/resources/updates/release-notes
- [11] Updates, News & Announcements Archives Copyleaks — copyleaks.com/blog/category/news-and-announcements

1.5 Robustness Against Evasion

Score: **3.2** · *Adequate*

Copyleaks demonstrates a mixed but often strong ability to resist common evasion techniques, positioning it as a recognized benchmark in the AI detection market. Its defenses are particularly effective against translation-based attacks; a 2024 peer-reviewed academic paper verified that the tool is "robust" against attempts to disguise AI text by translating it to another language and back [3]. This is further supported by its prominence as a target for AI "humanizer" tools, whose marketing often highlights the ability to bypass Copyleaks, paradoxically underscoring its general effectiveness [5] [9] [10]. Independent evidence confirms this strength, with one demonstration showing Copyleaks successfully identifying two separate humanized texts as 100% AI-generated. Despite these strengths, the detector has significant and exploitable vulnerabilities. In huby's direct testing, the tool's detection rate dropped from 100% to 0% after a single, simple character-level perturbation was applied to an AI-generated text, a preliminary result that was the worst performance observed among the eight products tested by huby on the identical corpus [2]. This specific weakness is corroborated by an independent academic study that found the tool susceptible to "multiple character-level mutations" [3]. The tool's performance against paraphrasing is also inconsistent. While one study found it resistant to paraphrasing attacks [3], huby's preliminary test showed it failed to detect AI in the single paraphrased document tested [1], a weakness supported by other research showing high failure rates for machine-paraphrased texts (Figure 6). This conflicting evidence suggests that while Copyleaks is a formidable tool, its reliability varies, and it can be completely bypassed by certain technical evasion methods.

Supporting Evidence:




Figure 13 — Curated Evidence:
Copyleaks_2e2efdf4-aa36-4974-a2dc-9133057ba33d_800x359.jpg




Figure 14 — Curated Evidence: Methodology for testing AI humanizer tools

References:

- [1] huby Test Data: results_copyleaks_20260921-manual.jsonl — storage.googleapis.com/.../results-copyleaks-20260921-manual-4b0635eb.js onl
- [2] Plagiarism Detection Levels Copyleaks API Docs — docs.copyleaks.com/concepts/features/detection-levels
- [3] Bypass Copyleaks | 100% Human Score EssayDone — www.essaydone.ai/bypass-copyleaks
- [4] Reduce Copyleaks AI False Positives — Humanizer Tool — humaniser.com/bypass-copyleaks
- [5] Bypassing copyleaks ? Need help in develop new prompt — community.openai.com/.../316541
- [6] Can AI Detectors Be Fooled? | AI Evasion & Detection Accuracy — copyleaks.com/blog/can-ai-detectors-be-fooled
- [7] Detecting AI-Generated Text: Factors Influencing Detectability... arXiv — arxiv.org/html/2406.15583v1
- [8] Cross-Language Detection Copyleaks API Docs — docs.copyleaks.com/concepts/features/cross-language-detection
- [9] Robustness of generative AI detection: adversarial attacks on black-box models (International Journal of Speech Technology, 2024) — storage.googleapis.com/.../Copyleaks_s10772-024-1014...%27s10772-024-10144-2.pdf
- [10] AI Video Detector for Enterprise Copyleaks — copyleaks.com/de/ai-detector/ai-video-detector
- [11] AI Video Detector for Enterprise Copyleaks — copyleaks.com/fr/ai-detector/ai-video-detector

1.6 Usability and Workflow Integration

Score: 4.5 · Strong

Copyleaks excels in its straightforward user experience and deep integration into existing workflows. The platform is designed to be intuitive, with a clean interface for submitting content and reports that use clear, visual formats for quick assessment [6] [9]. A significant strength is its seamless connection with major Learning Management Systems (LMS) like Canvas, Moodle, and Brightspace, allowing educators to check for plagiarism and AI-generated content directly within their grading environments [1] [2] [7]. For developers and businesses, the platform offers a robust and well-documented Application Programming Interface (API) to support custom integrations and large-scale batch processing, embedding detection capabilities directly into proprietary systems [2] [5] [8]. This focus on usability has been independently recognized, with one study awarding the platform top marks for its ergonomics and reporting interface [4]. Despite its strong workflow capabilities, there are notable areas for improvement in performance and report reliability. A competitive analysis points out that the application "can lag" and its interface may be "less refined compared to premium tools," suggesting a lack of polish compared to some competitors [3] [7]. More critically, evidence from the same analysis and user reviews indicates that report results can be inconsistent, particularly in the detection of AI-generated content, with the tool sometimes flagging human writing as AI [3]. While the platform's API is performant and designed for reliability, independent benchmarks show that at least one competitor offers significantly faster response times [1] [2]. The product also shows a commitment to accessibility with features like high-contrast design and public compliance documents, but the detailed report view can be challenging for users of screen readers [1] [2] [3] [4] [5]. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 15 — Screenshot of Copyleaks document upload interface




Figure 16 — Screenshot of Copyleaks dashboard with accessibility icon

References:

- [1] How Oakland University Used Copyleaks to Transform Plagiarism... — copyleaks.com/success-story/oakland-university
- [2] Canvas Copyleaks Teacher Scan Tool Analytic Dashboard — staffsupport.spcollege.edu/.../-Canvas-Copyleaks-Teach...Tool-Analytic-Dash-board
- [3] Export PDF Report Copyleaks API Docs — docs.copyleaks.com/concepts/features/export-pdf-report
- [4] How do I see the results of the Copyleaks plagiarism/AI scan for my... — bucknell.teamdynamix.com/.../How-do-I-see-the-results...for-my-Moodle-Assignment
- [5] Best AI Detector APIs in 2026: Complete Comparison — aidetector.ac/blog/best-ai-detector-apis
- [6] Copyleaks AI Detector Review 2026 — Accuracy, Pricing & Verdict — aidetector.ac/reviews/copyleaks
- [7] Copyleaks overview Brightspace Community D2L — community.d2l.com/brightspace/kb/articles/23680-copyleaks-overview
- [8] Submit File Copyleaks API Docs — docs.copyleaks.com/reference/actions/authenticity/submit-file

[9] Copyleaks Commitment to Accessibility — copyleaks.com/accessibility

[10] Compliance and Certifications Copyleaks — copyleaks.com/compliance-certifications

2.0 Privacy

Score: **3.5** · *Adequate*

A pattern of contradiction defines the product's approach to privacy, creating a gap between stated policy and operational practice. The company's strict prohibition on using customer data for AI training is undermined by conflicting product marketing, while its affirmation of user ownership is offset by indefinite data retention by default. This inconsistency extends to user controls, which are direct for submitted documents but require manual intervention for exercising personal data rights. The company's privacy and compliance posture relies almost entirely on unverified, self-reported claims. Foundational assertions about regulatory alignment, data lifecycle management, and AI data segregation are not substantiated by independent, third-party audits. This systemic lack of external validation requires users to trust vendor policies without objective proof of their implementation.

2.1 Artificial Intelligence Data Usage

Score: **3.5** · *Adequate*

Copyleaks presents a strong, privacy-first policy, stating that it does not use customer-submitted content to train its public AI models [1] [2]. This "never-in" approach is a superior privacy design, as it makes reactive opt-out controls unnecessary and ensures customer data is, by default, segregated from model development pipelines [3]. This core claim is supported by independent, peer-reviewed research describing a technical architecture capable of preserving privacy through methods like on-device analysis and the use of differential privacy [4]. The company also provides users with control to permanently delete their submitted files and associated reports. However, the company's own documentation and marketing introduce a significant contradiction to this otherwise clear policy. Copyleaks advertises a "Shared Data Hub" feature, which is described as enhancing the platform's detection accuracy [5] [6]. This implies that content from at least some users is incorporated into a shared dataset to improve the product, a practice that seems to conflict with the "never-in" promise. The available evidence lacks transparency regarding this feature, failing to clarify how data is used, whether participation is optional, or what anonymization techniques are applied. While Copyleaks offers some transparency to users, such as explaining the reasoning behind its AI detection results, its privacy posture is substantiated almost entirely by its own documentation. The company's privacy-centric claims about its commercial product, particularly regarding the segregation of customer data and the precise function of the Shared Data Hub, have not been verified by a formal, third-party audit. Supporting images illustrating these findings are included below.

Supporting Evidence:




Figure 17 — Curated Evidence: Copyleaks AI Detector user interface




Figure 18 — Curated Evidence:
Copyleaks_AI_Logic_Feature-1-1024x511.webp




Figure 19 — Curated Evidence: Copyleaks_featured-453.webp

References:

- [1] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [2] Terms Of Use Copyleaks — copyleaks.com/termsfuse
- [3] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [4] A privacy-preserving and robust AI-generated text detection system, Scientific Reports — storage.googleapis.com/.../s41598-026-64104-8-reference-fdfffba4.pdf
- [5] Data Hubs Copyleaks API Docs — docs.copyleaks.com/concepts/features/data-hubs
- [6] Self-Check Folders: Updated Data Hub Behaviour LTSA — itsa.sheridancollege.ca/.../copyleaks-self-check-folder
- [7] AI Video Detection API Copyleaks API Docs — docs.copyleaks.com/concepts/products/ai-video-detection-api
- [8] AI Image Detection API Copyleaks API Docs — docs.copyleaks.com/concepts/products/ai-image-detection-api

2.2 Data Governance and Ownership

Score: **3.7** · *Adequate*

Copyleaks' policies are clear that users retain full ownership and intellectual property rights over the content they submit for analysis (Figure 20) [2]. While users are provided with controls to delete their data, the platform's default approach is to store submitted content indefinitely unless a user actively requests its removal [5]. For institutional accounts, deletion is typically managed by an administrator, a standard practice for enterprise software. This model places the responsibility on the user to manage their data's lifecycle, as the service does not automatically delete content after a set period of time. The company discloses its data handling practices within its legal documents, which are linked directly from the content submission interface [4] Figure 2. These policies state that data may be shared with third-party partners for essential operations and, for most users, transferred to the US for processing [1]. Importantly, Copyleaks states that user content is not sold or used to train external AI models, limiting its use to providing the service and for internal research to improve its own systems [1] [4]. While enterprise customers are given greater control, such as options for data residency and opting out of research usage, transparency for individual users is more passive, relying on them to review the provided legal policies.

Supporting Evidence:



Figure 20 — Curated Evidence: Copyleaks_Custom-1024x667.webp

References:

- [1] Excluding and Preventing Indexing of Content Copyleaks API Docs — docs.copyleaks.com/concepts/features/exclude-content
- [2] Copyleaks: AI Content & Text Authenticity Detection — copyleaks.com
- [3] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [4] Delete Scans Copyleaks API Docs — docs.copyleaks.com/reference/actions/authenticity/delete
- [5] Data Privacy and AI Detection: What Happens to Your Papers After... — hub.paper-checker.com/.../data-privacy-ai-detection-...to-papers-after-submission
- [6] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [7] Copyleaks Research Reveals AI Has Unique Stylistic Fingerprints — copyleaks.com/blog/copyleaks-research-reveals-which-ai-wrote-what
- [8] AI in Action: 2025 Student AI Usage Report Copyleaks — copyleaks.com/blog/ai-in-action-2025-student-ai-usage-report

2.3 Privacy Compliance and Regulatory Alignment

Score: 3.5 · Adequate

Copyleaks demonstrates a commitment to privacy compliance through extensive public documentation outlining its alignment with regulations like the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) [1] [3]. The company provides key compliance features, such as a Data Processing Agreement (DPA) for customers and regional data storage options for institutions needing to keep data within specific geographic boundaries [2]. Its privacy policy is comprehensive and transparently discloses that user content may be used for AI training, while also providing a specific mechanism for users to opt out, showing a strong commitment to user control. This transparency is reinforced within the product, which makes the privacy policy and terms of use easily accessible before content is submitted for scanning. To substantiate its privacy and security posture, Copyleaks claims to hold multiple certifications, including SOC 2 and ISO 27001, which provide evidence of strong security controls foundational to data protection [1] [4]. However, the evidence for its regulatory compliance relies primarily on these related security certifications and the company's own statements. The assessment is limited because independent, third-party verification of its GDPR or CCPA compliance, such as a specific audit report, was not available for review.

Similarly, the security certifications themselves were based on vendor claims rather than viewable audit reports or certificates.

References:

- [1] Compliance and Certifications Copyleaks — copyleaks.com/compliance-certifications
- [2] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [3] Top Use Cases for AI Detection Across Industries Copyleaks — copyleaks.com/blog/ai-usage-best-practices-guide
- [4] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [5] Terms Of Use Copyleaks — copyleaks.com/termsfuse
- [6] Copyleaks AI Detector Review: Accuracy and Features — humanizethisai.com/blog/copyleaks-ai-detector-review

2.4 User Privacy Controls

Score: **3.2** · *Adequate*

Copyleaks provides users with a mix of direct, self-service controls and manual, request-based processes for managing privacy. The platform lacks a centralized privacy dashboard, meaning controls are distributed rather than consolidated in a single hub (Figure 2). Users have robust, direct control over the primary content they submit, with features to manage, delete, and organize documents and scan reports from their account history. Similarly, exporting scan results is a user-friendly process available directly through the interface and via API [1] Figure 2 [3] [4]. However, for actions concerning personal account data, such as requesting a comprehensive data export or account deletion, users must submit a formal request by emailing the company. While this manual process meets legal requirements, it places the burden on the user and lacks the convenience of an automated, in-app feature (Figure 2). The platform's approach to consent and communication follows a similar pattern of providing necessary but not always centralized controls. User consent is obtained by requiring agreement to the Terms of Use and Privacy Policy before using the service (Figure 2) [3]. For website tracking, Copyleaks implements a granular cookie consent banner that allows users to accept, reject, or customize their preferences, a practice that aligns with modern privacy standards [1] Figure 2. Users can opt out of marketing communications through standard methods, such as unsubscribe links in emails or by contacting support directly (Figure 2). Instead of an interactive control center, the user interface primarily offers links to static policy documents, requiring users to read through them to understand how to exercise their rights (Figure 2).

Supporting Evidence:




Figure 21 — Copyleaks Scan Results UI

References:

- [1] Terms Of Use Copyleaks — copyleaks.com/terms-of-use
- [2] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [3] Export Copyleaks API Docs — docs.copyleaks.com/reference/actions/downloads/export
- [4] Export PDF Report Copyleaks API Docs — docs.copyleaks.com/concepts/features/export-pdf-report

3.0 Security

Score: 3.5 · Adequate

A pattern of unverified claims and practical vulnerabilities undermines the product's security design. While the architecture incorporates formal controls and certifications, independent testing reveals significant performance failures in its core detection capabilities. This gap between design and demonstrated resilience is further reflected in observable application flaws and a lack of third-party validation for infrastructure defenses. The overall security program appears inconsistently applied and lacks public accountability. Protections are stronger for programmatic interfaces and integrated enterprise environments than they are for the primary web application and individual user accounts. This uneven posture is compounded by a systemic lack of public documentation for key processes, from security testing to incident management, preventing independent validation of company claims.

3.1 Application and API Security

Score: **3.3** · *Adequate*

Copyleaks provides foundational application and API security controls, but it is not a dedicated security product and lacks the comprehensive features of one. The platform's security measures, such as its SOC 2 certification and data encryption, are designed to protect the Copyleaks service itself rather than its customers' applications [1] [2]. For its API, the service implements a modern and secure authentication system where a permanent key is used to generate a temporary access token for making requests, a practice that reduces the risk of key exposure [9] [3]. Access to the API is further managed through rate limits that are tied to commercial subscription plans, which helps prevent system abuse and ensures service availability [6]. The platform demonstrates an awareness of common web vulnerabilities in its handling of user-submitted content, but evidence of a robust, in-depth strategy is limited. For example, Copyleaks correctly sanitizes uploaded HTML files to prevent them from executing malicious code in a user's browser (Figure 22), and its API offers a "sandbox" mode that suggests processing occurs in an isolated environment [8]. However, the assessment also found evidence of poor error handling, where an internal server error was exposed directly to a user after a file upload, creating an information leakage vulnerability. While these basic controls are present, there is no publicly available documentation or independent auditing that details protections against other critical threats, such as SQL injection, or more advanced file security measures like server-side malware scanning. Supporting images illustrating these findings are included below.

Supporting Evidence:




Figure 22 — Copyleaks UI demonstrating HTML sanitization

Tool	Implementation Complexity	Resource Requirements	Expected Outcomes	Ideal Use Cases	Key Advantages
[Grammarly](https://www.grammarly.com)	Low - easy multi-platform setup	Moderate - subscription-based	Accurate grammar, style, and plag	Everyday writers, students, profess	Integrated tools, quick editor feedback, citation help
[Scribbr](https://www.scribbr.com)	Low - web-based, pay-per-documen	Low to moderate - pay per check	Detailed Turnitin-level similarity rep	Students needing thorough pre-sub	Trusted Turnitin database, no institutional subscription needed
[Copyleaks](https://copyleaks.com)	Moderate - credit system, integrati	Moderate to high - plans with credit	Combined AI and plagiarism detect	Freelancers, teams, organizations	Enterprise integrations, broad language support
[Quetext](https://www.quetext.com)	Low - web-based with extensions	Moderate - subscription word limits	Clear plagiarism with AI detection b	Students, teachers, copywriters wit	Word-based pricing, bundles including AI and paraphrasing tools
[Copyscape](https://www.copyscape.com)	Low - simple web interface	Low - pay-per-search	Fast web-focused plagiarism results	Website owners, SEO professionals	Web-specific focus, monitoring services, pay-as-you-go
[ProWritingAid](https://prowritingaid.com)	Moderate - multiple platform integr	Moderate - subscription plus credits	Comprehensive writing analysis with	Authors, editors, serious writers	In-depth writing reports, flexible credit system
[Originality.ai](https://originality.ai)	Moderate to high - team setup, inte	High - professional pricing	Professional-grade AI and plagiarism	Content agencies, publishers, mark	Advanced AI detection, team features, website scanning

Figure 23 — Image of Copyleaks competitive analysis table




Figure 24 — Curated Evidence: Copyleaks_mceclip1.png

References:

- [1] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [2] Security Practices Copyleaks — copyleaks.com/security-practices

- [3] Account Actions Copyleaks API Docs — docs.copyleaks.com/reference/actions/account/overview
- [4] Authentication Issue Incident details Copyleaks Status — status.copyleaks.com/cmn2q33nt0y89onezzk0976cj
- [5] Engineering, Technology & Applied Science Research (ETASR) Security threats such as prompt injection attacks — storage.googleapis.com/.../Copyleaks_ETASR_13756.pdf...TF-8%27%27ETASR_13756.pdf
- [6] Rate Limits Copyleaks API Docs — docs.copyleaks.com/using-the-apis/rate-limits
- [7] Copyleaks API Overview — docs.copyleaks.com/using-the-apis/overview
- [8] Technical Specifications Copyleaks API Docs — docs.copyleaks.com/reference/data-types/authenticity/technical-specifications
- [9] Authentication Copyleaks API Docs — docs.copyleaks.com/using-the-apis/authentication

3.2 Detection Engine Integrity

Score: **3.5** · *Adequate*

Copyleaks demonstrates a proactive approach to securing its AI detection engine against manipulation. The system incorporates specific, documented features to counter evasion, such as detecting hidden characters and analyzing shifts in writing style to spot authorship anomalies (Figure 25) [2]. Its defenses are not just theoretical; a peer-reviewed study confirmed that Copyleaks successfully resisted certain adversarial attacks, providing third-party validation of its robustness [1]. Further evidence from a NIST publication details the company's use of sophisticated internal testing to harden its models against new threats [1]. The very existence of commercial tools designed to bypass Copyleaks underscores its prominence in the market and the real-world relevance of its security measures [3]. Despite these defenses, significant independent evidence shows that the engine's integrity can be compromised. A commercial market exists for "humanizer" tools designed specifically to fool Copyleaks, indicating known and exploitable vulnerabilities [1] [2]. Independent tests confirm their effectiveness, with one analysis showing a bypass tool succeeding 44% of the time [4] and a technical report finding Copyleaks detected only 58.7% of AI-generated text on a specific dataset [3]. These findings stand in contrast to the vendor's claim of over 99% accuracy, which is not supported by publicly available third-party studies [5]. This gap is compounded by user reports of significant false positives—where human writing is incorrectly flagged as AI-generated—and a lack of public information on how the company protects its own internal evaluation data from being compromised [6].

Supporting Evidence:




Figure 25 — Curated Evidence: Diagram of Integrity Mechanisms from Research Paper

Test Scores:

AI Generated Text Pass Rate – 4/9 (44%)

Human Written Text Pass Rate – 3/3 (100%)

One of a couple in our list of different style AI detector tools as Undetectable AI say that their tool integrates and checks through various AI detection tools – GPTZero, OpenAI, Writer, QuillBot, Copyleaks, Sapling, Grammarly, ZeroGPT – to help give its score even more validity.

Figure 26 — Curated Evidence: Undetectable AI Review Scores




Figure 27 — Curated Evidence: Copyleaks_ai-content-detection-tools-pangram.webp

References:

- [1] Academic Paper on Plagiarism Detection System Testing (2021) — storage.googleapis.com/huby-media/copyleaks/evidence/2201-03423v1-0311c918.pdf
- [2] Character Replacement & Plagiarism Detection: A Student's Guide — copyleaks.com/blog/character-replacement-plagiarism-detection-a-students-guide
- [3] AI for scientific integrity: detecting ethical breaches, errors... PMC — [pmc.ncbi.nlm.nih.gov/articles/PMC12436494](https://pubmed.ncbi.nlm.nih.gov/articles/PMC12436494)
- [4] Bypass Copyleaks | 100% Human Score EssayDone — www.essaydone.ai/bypass-copyleaks
- [5] Bypass Copyleaks AI Detector #1 Free Humanizer UnAI MyText — unaimytext.com/bypass-ai-detectors/bypass-copyleaks
- [6] Curated Evidence: Copyleaks_NIST.AI.100-4.pdf — storage.googleapis.com/.../nist-ai-100-4-18a10a05.pdf
- [7] Pros and cons of AI detection University at Albany — www.albany.edu/.../pros-and-cons-ai-detection
- [8] AI Detector Free AI Checker for ChatGPT, GPT-5, Gemini & More — copyleaks.com/ai-detector

3.3 Identity and Access Management

Score: **3.9** · *Adequate*

Copyleaks provides identity and access management features clearly tailored for its enterprise and institutional customers. The platform supports Single Sign-On (SSO), allowing organizations to integrate Copyleaks with their existing security systems and enforce their own authentication policies, including multi-factor authentication [1] [3]. Direct evidence confirms a dedicated "Login with SSO" button on the user interface, alongside integrations with major Learning Management Systems (LMS) for educational users (Figure 9). Furthermore, the platform includes role-based access controls, with visual evidence from the product showing dedicated sections for managing "Members" and "Repositories" within an organization. This structure allows administrators to control user permissions and segregate duties, which is essential for managing teams in a business or educational setting. While the platform's enterprise features are well-supported by evidence, there are significant gaps in documentation for individual user security and technical specifics. There is no primary evidence to confirm the availability of native multi-factor authentication (MFA), such as using an authenticator app, for individual accounts that do not use SSO [1]. This means users who are not part of a larger organization may have limited options for securing their accounts beyond a password. Similarly, there is a complete lack of public documentation regarding security controls for user web sessions, such as automatic logouts after a period of inactivity [1]. Although the platform's API uses a secure access token model, the security posture for the main web application remains undocumented, and further details on the granularity of user roles and supported SSO protocols are not publicly available. Visual evidence from the product's user interface is included for reference.

Supporting Evidence:




Figure 28 — Comparison of AI Content Detectors featuring Copyleaks

References:

- [1] Security Practices Copyleaks — copyleaks.com/security-practices
- [2] Video of Copyleaks Login User Interface — storage.googleapis.com/.../ai-content-detection-tool...-improve-seo-9294ff61.mp4
- [3] Copyleaks AI Detector Review: Accuracy and Features — humanizethisai.com/blog/copyleaks-ai-detector-review
- [4] Leap AI vs Copyleaks — Humanizer vs Enterprise Plagiarism Platform — www.tryleap.ai/vs/copyleaks
- [5] Account Actions Copyleaks API Docs — docs.copyleaks.com/reference/actions/account/overview
- [6] Authentication Copyleaks API Docs — docs.copyleaks.com/using-the-apis/authentication

3.4 Platform Security

Score: **3.8** · Adequate

Copyleaks demonstrates a strong commitment to data protection through robust encryption and a resilient infrastructure. The company protects customer data using industry-standard TLS encryption when it is in transit and AES-256 encryption when it is stored, or at rest [1] [2]. This foundational security is built upon Amazon Web Services (AWS), which provides high availability and redundancy. The platform's disaster recovery plan is not just a claim; it has been independently validated by a third-party audit as part of its SOC 2 Type II compliance, confirming that Copyleaks has a tested and formalized strategy to ensure service availability and data integrity [3]. The company also

employs advanced architectural features, including the use of a blockchain to create immutable, tamper-proof records and an AI-based system to detect anomalous activity [3] Figure 29. The company's broader security posture is supported by its claims of holding multiple industry certifications, including SOC 2, SOC 3, and ISO 27001 [2] [3] [4]. These certifications imply that an independent auditor has reviewed and validated a wide range of security controls, from infrastructure monitoring and threat detection to the processes used for secure software development. While these certifications provide a significant level of assurance, the full third-party audit reports were not available for review. However, tangible evidence of secure development practices is visible within the platform itself, which includes a feature for analyzing third-party code to manage potential security risks (Figure 20). Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 29 — Copyleaks Architectural Diagram

References:

- [1] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [2] Integrations & Technical SchoolsPLP — schoolsplp.com/integrations
- [3] Compliance and Certifications Copyleaks — copyleaks.com/compliance-certifications
- [4] Security Practices Copyleaks — copyleaks.com/security-practices

3.5 Security Assurance and Incident Response

Score: **3.1** · Adequate

Copyleaks attests to maintaining a formal security program, evidenced by its claims of holding key industry certifications such as SOC 2, SOC 3, and ISO 27001 [1] [2]. These certifications, which are corroborated by an independent source [3], suggest that the company's security controls and management systems have been audited

by external experts. A SOC 2 attestation, in particular, provides a significant level of assurance that security controls are not only designed properly but also operate effectively over time, indicating a structured approach to security. Despite these certifications, the company offers limited public documentation to substantiate its day-to-day security practices. While Copyleaks advertises its compliance, the underlying audit reports and certificates are not publicly available for direct review [2]. Furthermore, there is no publicly accessible information detailing the company's process for responding to security incidents, nor is there mention of a formal program for testing its systems against cyberattacks, often known as penetration testing [1] [6]. This lack of transparency makes it difficult for external observers to verify the specific details and rigor of Copyleaks' security operations, even as evidence confirms the existence of threats that would require a formal response [2]. In terms of transparency, Copyleaks maintains a public status page that provides real-time and historical data on operational issues like service outages, demonstrating a commitment to communicating about system availability [8] [9]. However, this transparency does not appear to extend to security-specific events. The company does not provide a dedicated security advisory page or a public record of how it has handled past security breaches, focusing its public disclosures on operational performance rather than security incidents. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 30 — Screenshot of Semihuman.ai User Interface

References:

- [1] Trust & Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/overview
- [2] Compliance and Certifications Copyleaks — copyleaks.com/compliance-certifications
- [3] Best Free AI Content Detectors & Optimizers (2026) CopyCrest — copycrest.com/resources/best-free-ai-content-detectors-optimizers
- [4] Curated Evidence: Copyleaks_CCD-QMS-2026-v1.0-FINAL.pdf — storage.googleapis.com/.../ccd-qms-2026-v1-0-final-1eb10dfe.pdf
- [5] Is Copyleaks down? StatusGator — statusgator.com/services/copyleaks
- [6] Security Practices Copyleaks — copyleaks.com/security-practices
- [7] Curated Evidence: Copyleaks_NIST.AI.100-4.pdf — storage.googleapis.com/.../nist-ai-100-4-18a10a05.pdf
- [8] Copyleaks's Services status page Status — status.copyleaks.com
- [9] Authentication Issue Incident details Copyleaks Status — status.copyleaks.com/cmn2q33nt0y89onezzk0976cj

4.0 Impact, Ethics, and Safety

Score: **3.6** · Adequate

A systemic gap exists between the company's stated principles and the public evidence supporting them. The product's marketed performance and responsible AI posture are consistently challenged by a lack of independent verification and contradictory external research. This pattern of unsubstantiated claims creates a trust deficit and prevents independent validation of the technology's fairness and reliability. The product's core design philosophy shifts the final responsibility for interpretation to the user. However, this reliance on human oversight is not supported by sufficient vendor guidance or the model transparency needed for users to make informed judgments. This creates a risk that users may misapply the tool in high-stakes decisions, particularly given the opacity of the model's

decision-making and performance limitations.

4.1 Abuse Prevention and Public Trust

Score: **3.6** · *Adequate*

Copyleaks positions itself as a key tool for ensuring authenticity and preventing various forms of abuse. The platform offers specific products designed for high-stakes fraud detection, such as combating financial aid fraud by identifying forged documents and "ghost student" applications [3]. Its capabilities extend beyond text analysis to include AI-generated image detection, video analysis for identity verification, and scanning source code to prevent intellectual property theft [1] [5] [6]. To prevent misuse of its own system, Copyleaks implements technical controls like API rate limiting and provides detailed reports that encourage human oversight [2] [3] Figure 32. Central to its public trust strategy is a widely advertised claim of over 99% accuracy, which the company states is backed by independent third-party studies [7]. However, there are significant gaps between the company's claims and the available evidence, impacting its trustworthiness. The core claim of ">99% accuracy" is not substantiated with the independent studies themselves, and it is directly contradicted by an academic study that measured the tool's AI detection accuracy at 75%. This discrepancy is amplified by widespread user reports of false positives, which question the tool's real-world performance. This lack of independent verification extends to its platform controls, for which evidence of effectiveness comes almost exclusively from the vendor (Figure 23) [5]. Furthermore, security measures are inconsistently applied; while text-based reports are designed to be tamper-resistant, the AI image detector's on-screen results can be easily falsified [1] [2]. While Copyleaks is transparent about system uptime via a public status page, it does not have a public disclosure process for security vulnerabilities or data breaches, limiting its overall safety transparency [1] [2].

Supporting Evidence:



Figure 31 — Curated Evidence:
Copyleaks_video-detection-homepage-hero.webp




Figure 32 — Screenshot of Copyleaks AI Image Detector Interface

References:

- [1] AI Image Detection API Copyleaks API Docs — docs.copyleaks.com/concepts/products/ai-image-detection-api
- [2] Online Text Compare Diff Tool | Compare 2+ Documents | Copyleaks — app.copyleaks.com/text-compare
- [3] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [4] Rate Limits Copyleaks API Docs — docs.copyleaks.com/using-the-apis/rate-limits
- [5] Copyleaks vs. GPTZero: Which AI Detector Is More Accurate? — copyleaks.com/blog/copyleaks-vs-gptzero
- [6] AI for scientific integrity: detecting ethical breaches, errors... PMC — pmc.ncbi.nlm.nih.gov/articles/PMC12436494
- [7] Notice history Copyleaks's Services status page — status.copyleaks.com/history/1
- [8] Copyleaks's Services status page Status — status.copyleaks.com

4.2 Detection Fairness and Transparency

Score: **3.6** · Adequate

Copyleaks makes efforts to be transparent by acknowledging technical limitations and the potential for its AI detector to produce false positives, encouraging users to treat its output as a statistical probability rather than a final verdict [1] [3]. However, a significant gap exists between the company's performance claims and publicly available evidence, particularly concerning fairness. For instance, while Copyleaks reports a high accuracy rate for writing by non-native English speakers, independent academic research found that AI detectors like it are more likely to incorrectly flag text from this group [1] [2]. This discrepancy is compounded by the system's "black box" nature. The tool highlights text it deems likely to be AI-generated but does not explain the specific reasons for its judgment, making it difficult for writers or institutions to meaningfully audit or challenge a finding. The product's design does include features that support a fairer process. By localizing detections to specific sentences rather than issuing a single verdict for an entire document, it allows for a more nuanced and targeted review. Despite these positive design choices, the company's broader claims of accuracy and fairness are not substantiated by public data. Copyleaks states its accuracy is validated by "independent third-party studies," but these reports are not provided, and a separate comparative review omits independent accuracy data for the product while providing it for competitors [2] [3]. Similarly, while the tool

supports over 100 languages, there is no evidence demonstrating that its performance is consistent and unbiased across them [3]. This lack of verifiable, independent data on performance and bias remains a critical transparency issue for users. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 33 — Curated Evidence: Copyleaks_13-69c8efcf-b314-424f-9a43-cc5bf5f22569.png

References:

- [1] Rate Limits Copyleaks API Docs — docs.copyleaks.com/using-the-apis/rate-limits
- [2] Technical Specifications Copyleaks API Docs — docs.copyleaks.com/reference/data-types/authenticity/technical-specifications
- [3] AI Detector Accuracy for Non-Native English Content Copyleaks — copyleaks.com/.../accuracy-of-ai-detection-model...or-non-native-english-speakers
- [4] Auditing the Fairness of AI-Detection Tools: A Comparative Study of ESL Published and AI-Generated Texts and Their Misclassification Risks — ijtle.com/.../auditing-the-fairness-of-ai-dete...nd-their-misclassification-risks
- [5] Plagiarism Detection Levels Copyleaks API Docs — docs.copyleaks.com/concepts/features/detection-levels
- [6] AI Detector & Plagiarism Checker for Legal Services Copyleaks — copyleaks.com/use-case-legal-services
- [7] Copyleaks AI Logic, AI Source Match Faculty & Staff Support — staffsupport.spcollege.edu/.../-Copyleaks-AI-Logic-AI-Source-Match
- [8] Copyleaks vs ZeroGPT: Which AI Detector Is More Accurate? — copyleaks.com/blog/copyleaks-vs-zerogpt
- [9] Does Copyleaks Detect AI? Accuracy and Workflow Guide | Leap — www.tryleap.ai/blog/does-copyleaks-detect-ai
- [10] AI Text Detection API Copyleaks API Docs — docs.copyleaks.com/concepts/products/ai-text-detection-api

4.3 Responsible Artificial Intelligence Governance

Score: **3.3** · Adequate

Copyleaks presents its AI detection technology as a decision-support tool designed to keep a human in control. The company publishes a detailed set of AI principles centered on accuracy, fairness, and human accountability, and aligns its governance with external standards like ISO [3] [4] [5]. This philosophy is reflected in the product, which provides users with detailed reports and sentence-level confidence scores to show where the model detected AI-generated content [4]. The platform also gives users direct control to guide the analysis by calibrating sensitivity or instructing the tool to ignore specific sections like quotes and references [2]. These features reinforce the company's stated "human-in-the-loop" principle, positioning the AI as an investigative aid rather than a final arbiter. Despite providing tools for human review, the platform has significant gaps in guidance and transparency. The evidence does

not show that Copyleaks provides explicit warnings or educational materials to help users interpret results responsibly and mitigate "automation bias," where a high probability score might be accepted without critical investigation. This is a critical omission, as third-party tools designed to evade AI detection make it essential that users treat the tool's output as a starting point, not a final verdict (Figure 30). While Copyleaks claims its high accuracy is validated by independent studies, these reports are not publicly available, leaving users to rely on vendor claims [1, 10]. Furthermore, the detection model itself remains a "black box," with no public disclosure of its architecture or training data, which limits independent assessment of its potential flaws or biases [0, 2].

References:

- [1] Export PDF Report Copyleaks API Docs — docs.copyleaks.com/concepts/features/export-pdf-report
- [2] Copyleaks AI Detector Testing Methodology — copyleaks.com/ai-detector/testing-methodology
- [3] How Do AI Detectors Work? Copyleaks — copyleaks.com/blog/how-does-ai-detection-work
- [4] Enterprise Content Governance Copyleaks API Docs — docs.copyleaks.com/concepts/use-cases/enterprise-content-governance
- [5] Generative AI Governance – Risk & Compliance Copyleaks — copyleaks.com/governance-risk-and-compliance
- [6] What is Responsible AI? Copyleaks — copyleaks.com/blog/what-is-responsible-ai
- [7] Updates, News & Announcements Archives Copyleaks — copyleaks.com/blog/category/news-and-announcements
- [8] AI Detector Free AI Checker for ChatGPT, GPT-5, Gemini & More — copyleaks.com/ai-detector

4.4 Responsible Detection Practices

Score: **3.9** · *Adequate*

Copyleaks consistently positions its AI detector as a tool to support, not replace, human decision-making. This "human-in-the-loop" philosophy is embedded in the product's design, which provides users with a detailed report rather than a simple pass/fail judgment. Instead of a single verdict, users in education, HR, and legal fields receive multiple data points, such as the percentage of AI-detected text, plagiarism details, and comparisons to an author's past work [4] [6]. This design encourages a more thoughtful investigation and is a responsible approach, given that AI detectors can make mistakes and third-party tools exist specifically to evade them (Figure 8) [3] [1]. The company also frames the tool as an evidence provider for internal compliance workflows, reinforcing that the final judgment should rest with a person [7]. While the tool's design implicitly guides users toward responsible interpretation, the company provides limited explicit, documented advice on how to use the tool ethically, particularly in high-stakes employment and academic scenarios [1] [4]. Much of the specific guidance on handling potential false positives or understanding the tool's limitations is developed by the educational institutions using the software, not by Copyleaks itself [1] [2] [3] [4]. This guidance gap is significant because the company markets its tool with claims of "over 99% accuracy," but the third-party studies cited to support this were not available for review [1]. This may create a risk of user over-reliance on the tool's scores, a concern echoed in some user reviews about perceived inaccuracies.

References:

- [1] Copyleaks AI Detection Tool: The Honest Read CASRAI — casrai.org/guides/copyleaks-ai-detection-tool
- [2] Academic Integrity Tools and Statements | CTI SUU — www.suu.edu/cti/academic-integrity-tools-and-statements.html
- [3] AI Ethics in Education | Copyleaks 2025 AI in Education Trends Report — copyleaks.com/blog/ai-ethics-in-education-2025-student-ai-usage-report
- [4] The Rise of AI in the Recruitment Process | Human Resource Services — hrs.wsu.edu/2025/08/15/the-rise-of-ai-in-the-recruitment-process
- [5] Top Use Cases for AI Detection Across Industries Copyleaks — copyleaks.com/blog/top-ai-detection-use-cases-across-industries
- [6] Privacy Policy Copyleaks — copyleaks.com/privacy-policy
- [7] What is a DMCA Takedown? Copyleaks — copyleaks.com/blog/what-is-dmca-takedown

5.0 Sustainability and Reliability

Score: **3.8** · *Adequate*

A systemic pattern emerges where strong, observable indicators of operational health and market acceptance are paired with a lack of publicly verifiable commitments. The company provides real-time operational transparency and has secured investor funding, yet it does not publish formal performance guarantees, financial statements, or a product roadmap. This creates a recurring gap between the service's observed stability, as seen in huby's testing, and its unverified performance against formal standards or long-term stress. This opacity extends to the core detection technology and social claims, where vendor-reported accuracy rates are not substantiated by accessible third-party studies. The absence of a disclosed methodology for core model maintenance, combined with independent data showing performance issues, casts uncertainty on the long-term sustainability of the product's primary function. Consequently, while the business appears viable and the service is operationally resilient, customers cannot independently assess the sustainability of its model accuracy or its performance against contractual obligations.

5.1 Detection Model Maintenance

Score: **2.7** · *Weak*

While Copyleaks claims a high degree of accuracy for its AI detector, including a very low 0.2% rate of false positives—instances where human writing is incorrectly flagged as AI-generated—independent evidence suggests performance issues in practice [1]. Multiple user reviews and a third-party analysis indicate that the tool incorrectly flagging human work is a noticeable problem for users (Figure 5). At the same time, the detector can also struggle with false negatives, or failing to identify AI-generated content. A technical report measured a high false negative rate of 58.7% on a benchmark testing against AI-imitated writing styles, a finding that is supported by anecdotal user reports of the tool missing AI-generated text [2]. These performance gaps are coupled with a significant lack of transparency regarding the company's model maintenance processes. To remain effective against constantly evolving AI writing tools, a detector must be regularly tested and updated. Although Copyleaks confirms that internal testing occurs, it provides no public schedule for how often it revalidates its model's performance or refreshes its benchmark datasets with content from new AI models [4]. Furthermore, there is no available documentation describing a process for monitoring, analyzing, and using error reports from the field to systematically retrain and improve the detection model. This absence of a publicly documented maintenance cadence makes it difficult to assess the long-term reliability and sustainability of the tool's accuracy.

References:

- [1] The Do's and Don'ts You Need To Know Before Utilizing AI Detectors — copyleaks.com/.../the-dos-and-donts-you-need-to-...-before-utilizing-ai-detectors
- [2] Pangram 4 Technical Report arXiv — arxiv.org/html/2607.27183v1
- [3] AI Detector Accuracy for Non-Native English Content Copyleaks — copyleaks.com/.../accuracy-of-ai-detection-model...or-non-native-english-speakers
- [4] Copyleaks AI Detector Testing Methodology — copyleaks.com/ai-detector/testing-methodology
- [5] CCD Quality Management System 2026 (CCD QMS 2026 v1.0) — storage.googleapis.com/.../ccd-qms-2026-v1-0-final-1eb10dfe.pdf
- [6] Copyleaks vs. GPTZero: Which AI Detector Is More Accurate? — copyleaks.com/blog/copyleaks-vs-gptzero
- [7] Top 20 AI-Generated Text Detectors Comparison AIMultiple — aimultiple.com/ai-generated-text-detector

5.2 Detection Service Reliability

Score: **4.3** · *Strong*

Copyleaks demonstrated exemplary processing reliability in huby's independent testing, successfully processing 100% of all 19 submissions without any failures [3]. This perfect success rate was the best performance observed among the nine AI detection products tested on the identical corpus [3]. The service's reliability is further corroborated by independent market signals, such as its use as a benchmark by other industry tools, which indicates it is perceived as a consistently available standard. The platform is built on a scalable, API-first architecture designed for high-volume processing, with features like batch processing and load balancing to manage demand and ensure performance consistency [5] [4] [7]. The company maintains strong operational transparency through a public status page that provides real-time system status and a detailed history of past incidents, with its availability reporting corroborated by third-party monitoring services [1] [2]. To manage service load and maintain stability, Copyleaks uses an asynchronous, queue-based architecture with API rate limits and a tiered priority system for paying customers [8]. While its adoption by enterprise clients and an industry award recognizing its "scalable infrastructure" suggest a robust platform, the available evidence does not include a public Service Level Agreement (SLA), independently audited uptime reports, or quantitative load-testing results to verify performance against a contractual guarantee. Supporting charts illustrating these findings are included below.

Supporting Evidence:



Figure 34 — Curated Evidence: Copyleaks_Features-Image-3-1.webp




Figure 35 — Copyleaks Plagiarism Detector Screenshot




Figure 36 — Curated Evidence: Copyleaks-GPTInfer detector-FxyUMtvVBL8BqRMPMv1ZYTzTDIqk87.webp

References:

- [1] Copyleaks's Services status page Status — status.copyleaks.com
- [2] Is Copyleaks down? StatusGator — statusgator.com/services/copyleaks
- [3] huby Test Data: results_copyleaks_20260921-manual.jsonl — storage.googleapis.com/.../results-copyleaks-20260921-manual-4b0635eb.jsonl
- [4] Best Practices for Working with Texts & Documents Copyleaks API — docs.copyleaks.com/concepts/performance/best-practices
- [5] Copyleaks APIs | Real-Time Content Integrity Solutions — copyleaks.com/api
- [6] Handling Failures Copyleaks API Docs — docs.copyleaks.com/concepts/performance/handling-failures
- [7] Copyleaks API Overview — docs.copyleaks.com/using-the-apis/overview
- [8] Rate Limits Copyleaks API Docs — docs.copyleaks.com/using-the-apis/rate-limits

5.3 Environmental and Social Responsibility

Score: **3.9** · *Adequate*

Copyleaks' social responsibility is centered on its core function of upholding academic and professional integrity. The company demonstrates significant collaboration within the educational community through partnerships with major Learning Management Systems like Canvas [2] and academic publishers such as Morressier [1]. Its widespread adoption is independently confirmed by universities that integrate the tool into their academic processes [4]. Beyond simply detecting plagiarism, the product also serves an educational role by offering tools that provide writing feedback and help with proper citations, directly contributing to student learning [3] Figure 38. In contrast, the company does not have a public, formal environmental sustainability strategy. There are no published commitments regarding renewable energy use, carbon emissions, or other environmental targets. While Copyleaks provides documentation on how to use its software more efficiently [5] and makes marketing claims about the speed of its tools (Figure 38), these efforts appear driven by performance and cost optimization rather than a stated environmental goal. The company relies on major cloud providers for its infrastructure, but it does not disclose information about its data center energy sources or report on its own operational footprint. Supporting charts illustrating these findings are included below.

Supporting Evidence:




Figure 37 — Curated Evidence: Copyleaks_best-ai-and-plagiarism-checkers-gptzero-scaled.png



Figure 38 — Curated Evidence: Copyleaks AI Grader efficiency claim

References:

- [1] Morressier and Copyleaks partner to bring automated detection of... — www.morressier.com/post/morressier-and-copyleaks-partner
- [2] Copyleaks Inc Partners with Canvas LMS to Offer Plagiarism... — copyleaks.com/.../copyleaks-inc-partners-with-ca...-using-ai-and-machine-learning
- [3] How to Fight plagiarism in Education Copyleaks — copyleaks.com/plagiarism-resources/fight-educational-plagiarism
- [4] Copyleaks Update & AI Workshop University of Illinois Springfield — www.uis.edu/.../copyleaks-update-ai-workshop
- [5] Best Practices for Working with Texts & Documents Copyleaks API — docs.copyleaks.com/concepts/performance/best-practices

5.4 Operational Excellence

Score: **3.7** · *Adequate*

Copyleaks demonstrates a strong commitment to operational transparency, primarily through a comprehensive public status page that is independently monitored by third-party services [2] [3]. This platform provides customers with real-time service health, a detailed history of past incidents, and specific reports on service disruptions, such as a notable authentication outage [4] [6]. Users can also subscribe to receive proactive email notifications about the system's status [5]. This practice of open communication extends to its product development, where the company publishes regular API release notes for developers and general product updates for all users, indicating a mature process for managing and announcing changes [7] [8]. While the company's transparency is a positive indicator of reliability, there is a gap between its operational practices and publicly verifiable commitments. Copyleaks offers formal Service Level Agreements (SLAs) to its enterprise customers, a claim corroborated by an independent industry source [1]. However, without a public SLA document, specific guarantees regarding uptime, performance, or financial remedies for service failures cannot be independently confirmed. A similar pattern is observed in customer support; while the service provides multiple support channels, including an in-app chat feature, there are no published metrics like average response times or hours of availability to substantiate the quality and speed of this support (Figure 39). This lack of public documentation and performance data limits a full assessment of the company's service commitments.

Supporting Evidence:




Figure 39 — Copyleaks User Interface Screenshot

References:

- [1] Copyleaks vs Originality.ai for Agencies (2026) Human Writes — humanwritesai.com/blog/copyleaks-vs-originality-compared
- [2] Copyleaks's Services status page Status — status.copyleaks.com
- [3] Is Copyleaks down? StatusGator — statusgator.com/services/copyleaks
- [4] Authentication Issue Incident details Copyleaks Status — status.copyleaks.com/cmn2q33nt0y89onezzk0976cj
- [5] Get updates by email Copyleaks's Services status page — status.copyleaks.com/subscribe/email
- [6] Notice history Copyleaks's Services status page — status.copyleaks.com/history/1
- [7] What's New Copyleaks API Docs — docs.copyleaks.com/resources/updates/release-notes
- [8] Updates, News & Announcements Archives Copyleaks — copyleaks.com/blog/category/news-and-announcements

5.5 Vendor Sustainability and Business Viability

Score: 4.4 · Strong

Copyleaks demonstrates strong business viability, underpinned by significant investor confidence and validated market growth. The company has secured multiple venture capital funding rounds, including a notable \$6 million round in 2022, which provides a substantial runway for development and expansion [1] [2] [4]. This financial backing is complemented by strong performance indicators, such as its placement at No. 153 on the Inc. 5000 list of fastest-growing private companies, which serves as third-party validation of its rapid growth [5]. The company's stability is further reinforced by a founder-led team that has shown strategic agility, successfully pivoting to capitalize

on the market demand for AI content detection and expanding beyond its traditional academic base. The company has built a sustainable market position by embedding its technology directly into the workflows of its target customers. This is most evident in its extensive ecosystem of strategic partnerships with major Learning Management Systems (LMS) like Canvas, Moodle, and D2L, which places its tools at the core of the education market [3] [5] [8] [9]. Its reach also extends into the publishing sector through collaborations with firms like Morressier [6]. However, as a private company, Copyleaks' overall sustainability assessment is limited by a lack of public transparency. It does not disclose audited financial statements or provide a forward-looking product roadmap, which prevents a direct assessment of profitability and limits customers' ability to plan for its long-term evolution [1] [2]. While specific user-base claims are not independently audited, the combined evidence of funding, strategic partnerships, and recognized growth points to a stable and expanding business.

Supporting Evidence:




Figure 40 — Curated Evidence: Copyleaks_best-ai-content-detection-tools-ai-detector.jpg




Figure 41 — Curated Evidence: Copyleaks_AI_vol2.webp

References:

- [1] Copyleaks 2026 Company Profile: Valuation, Funding & Investors — pitchbook.com/profiles/company/117837-19
- [2] Copyleaks raises \$6 million to uncover plagiarism | Ctech — www.calcalistech.com/ctechnews/article/sj6rhjoeq
- [3] Growing popularity of plagiarism and AI-detection technology... — hartfordbusiness.com/.../growing-popularity-of-plag...makes-copyleaks-one-of-cts

- [4] Copyleaks: AI Content & Text Authenticity Detection — copyleaks.com
- [5] About Us Copyleaks — copyleaks.com/about-us
- [6] Open LMS Partners With Copyleaks, Adding Advanced AI-Driven... — learningnews.com/.../open-lms-partners-with-copyleaks-adding-advanced-ai-driven-content-detection
- [7] What's New Copyleaks API Docs — docs.copyleaks.com/resources/updates/release-notes
- [8] AI Image Detection for Enterprises | Copyleaks — copyleaks.com/blog/introducing-copyleaks-ai-image-detection
- [9] Copyleaks Inc Partners with Canvas LMS to Offer Plagiarism... — copyleaks.com/.../copyleaks-inc-partners-with-canvas-lms-offering-plagiarism-detection-using-ai-and-machine-learning
- [10] Instructure Names Copyleaks Industry Recommended Partner AI — www.globenewswire.com/.../instructure-names-copyleaks-industry-recommended-partner-ai-and-higher-education.html

6.0 Use Cases & Pricing

Score: 4.2 · Strong

Copyleaks' API-first strategy is central to its commercial model, enabling the deep integrations and administrative functions that define its enterprise offerings. This integration-driven approach successfully fosters market adoption and high switching costs, yet the core API and its various plugins lack formal third-party performance validation. The product's extensibility is also leveraged by external tools that challenge its core value, while some customers report technical friction in the integrated user experience. A consistent pattern emerges where foundational capabilities are independently confirmed, but their depth and granularity for enterprise-scale operations are not. This gap between confirmed features and unverified depth applies to administrative reporting, bulk user management, and the specific terms of volume discounts. Consequently, the product's proven success in education is not matched by equivalent third-party validation of its cost-effectiveness or administrative robustness for large, non-academic organizations.


6.1 Commercial Packaging and Pricing

Score: 4.4 · Strong

Copyleaks centers its commercial strategy on a flexible, credit-based system that allows customers to align costs directly with consumption. Users purchase credits—where one credit typically scans 250 words—through either recurring subscriptions or as a one-time, pay-as-you-go option, avoiding the rigid tiers of some competitors [1] [3]. A perpetually free tier is available, but its limited allowance of roughly 10 credits per month makes it suitable only for occasional, low-volume use [2] [3]. The company further segments its offerings for different markets, with distinct packages for Business, Education, and API/Developer clients [1] [2]. This strategy reserves high-value features like deep Learning Management System (LMS) integrations for institutional customers and dedicated API access for enterprise clients, who are served through a flexible, custom-quoted licensing model [4] Figure 43 [5]. This pay-for-what-you-use model is positioned as a key competitive advantage, and evidence suggests it is effective in the market. Independent analysis confirms the pricing is competitive for smaller buyers, while public procurement data shows Copyleaks winning significant six-figure contracts with universities, validating its viability against established incumbents [2] [4]. The value proposition for these larger clients is augmented by features that position the tool as an integrated part of their content integrity infrastructure [5] [6]. However, the overall value for money is tempered by a significant challenge to its core function: the emergence of commercially available 'humanizer' tools designed specifically to bypass AI detectors like Copyleaks (Figure 5). This issue, coupled with some user reviews questioning if the tool's performance justifies its cost, raises questions about the ultimate return on investment. Supporting charts

illustrating these findings are included below.


Supporting Evidence:



The screenshot displays the Copyleaks pricing page with four main plans. The first plan, 'AI + Plagiarism Detection', is highlighted as the 'BEST VALUE' with a 'SAVE 20%' tag. It costs \$16.99/mo and includes up to 100 credits, detection across multiple languages, multiple forms of paraphrasing detection, detection of the most popular AI models (ChatGPT, Gemini, Claude) with over 99% accuracy, and two users included. The other three plans are 'Plagiarism Detector' at \$10.99/mo, 'AI Detector' at \$9.99/mo, and 'Enterprise & LMS' which requires contacting sales. Each plan includes a 'Get Instant Access' button and a 'Cancel anytime' note.

Plan	Price	Key Features	Action
AI + Plagiarism Detection	\$16.99/mo	Up to 100 Credits, Detection across multiple languages, Multiple forms of paraphrasing detection, Detection of the most popular AI models (ChatGPT, Gemini, Claude) with over 99% accuracy, Two users included	Get Instant Access
Plagiarism Detector	\$10.99/mo	Up to 100 Credits, Multiple forms of paraphrasing detection, Detection of source code plagiarism, Cross-language detection, Two users included	Get Instant Access
AI Detector	\$9.99/mo	Up to 100 Credits, Detection in over 30 languages, with more in the works, Detection of the most popular AI models (ChatGPT, Gemini, Claude) with over 99% accuracy, Two users included	Get Instant Access
Enterprise & LMS	Contact us	API integration (white label available), Turnkey LMS integrations, Enterprise grade security, Customer permissions + SSO, Dedicated Customer Success Team	Schedule a Demo

Figure 42 — Screenshot of Copyleaks Pricing Page



The screenshot shows the Copyleaks website navigation bar and hero section. The navigation bar includes links for Products, Solutions, Pricing, Resources, Login, Get a Demo, and Start for Free. The hero section features the headline 'Content Integrity & AI Detection For Marketing' and a sub-headline 'From Fortune 500 companies to the world's top universities, our AI-powered content analysis platform protects intellectual property, governs AI use, and sets the standard for responsible AI adoption.' A 'Get a Demo' button is prominently displayed. Below the hero section, there is a 'FEATURED IN' section listing various media outlets like Inc. 500, Business Insider, TechCrunch, Axios, Fast Company, Wired, and Forbes. A small inset image shows a sample report titled 'Blog Post - GA4' with a '100% AI Content Found' result.

Figure 43 — Copyleaks Website Navigation Bar

References:

- [1] Copyleaks Review: How Accurate Is It at Detecting AI Content? — walterwrites.ai/copyleaks-review
- [2] Pricing Copyleaks — copyleaks.com/pricing
- [3] Copyleaks Pricing (2026): Plans & Coupons CASRAI — casrai.org/guides/copyleaks-pricing
- [4] Manage Your Credits Copyleaks API Docs — docs.copyleaks.com/concepts/management/manage-your-credits
- [5] Enterprise Content Governance Copyleaks API Docs — docs.copyleaks.com/concepts/use-cases/enterprise-content-governance
- [6] AI Detector Free AI Checker for ChatGPT, GPT-5, Gemini & More — copyleaks.com/ai-detector
- [7] Curated Evidence: Copyleaks_Turnitin_vs_Copyleaks_Which_One_Is_Actually_Better_2026.mp4 — storage.googleapis.com/.../turnitin-vs-copyleaks-which-one-is-actually-better-2026-9ba5804d.mp4
- [8] 6 Best Copyleaks Alternatives for AI Detection in 2026 Leap AI — www.tryleap.ai/alternatives/copyleaks

6.2 Detection Analysis and Reporting

Score: **3.5** · *Adequate*

Copyleaks provides comprehensive analysis reports that allow users to document and share the results of a content scan. Users can download a formal PDF report or share findings through a secure, interactive online link [1] [2] [5]. These reports are designed for transparency, packaging the full analysis—including similarity scores, highlighted text, and source comparisons—into a single document. This functionality supports a variety of use cases, from academic record-keeping and legal due diligence to providing verifiable proof of originality to publishers or clients. In addition to formatted reports, the platform's developer tools allow for the export of raw scan data for more technical analysis [2]. Beyond single-scan reports, the platform offers features for managing and analyzing historical data to provide deeper context. Copyleaks retains past scans, which can be organized into folders for long-term auditing and content tracking [5]. The system uses this history to power analytical tools, such as graphing originality trends over time or comparing a new document against an author's previous work to identify unusual patterns [4] [5]. For institutional and enterprise clients, Copyleaks offers extensive customization, allowing organizations to generate reports that feature their own branding, logos, and color schemes, effectively integrating the detection service into their own platforms [1] [2]. This white-labeling capability is positioned as a premium feature for customers using the developer API [1]. Supporting visual evidence illustrates these customization and reporting features.

Supporting Evidence:




Figure 44 — Screenshot of Copyleaks "Historical Author Data" feature

References:

- [1] Export PDF Report Copyleaks API Docs — docs.copyleaks.com/concepts/features/export-pdf-report
- [2] Export Copyleaks API Docs — docs.copyleaks.com/reference/actions/downloads/export
- [3] Manage Your Credits Copyleaks API Docs — docs.copyleaks.com/concepts/management/manage-your-credits
- [4] Prevent Self-Plagiarism and Author Conflicts Copyleaks API Docs — docs.copyleaks.com/concepts/features/self-plagiarism
- [5] Ways to Display Reports Copyleaks API Docs — docs.copyleaks.com/concepts/features/how-to-display
- [6] Curated Evidence: Copyleaks_AI_Content_Detection_Tool_Copyleaks_Avoid_Plagiarism_Improve_SEO.mp4 — storage.googleapis.com/.../ai-content-detection-tool...-improve-seo-afe2f64d.mp4

6.3 Enterprise Deployment and Administration

Score: 4.0 · Strong

Copyleaks is designed with features to support deployment and administration within large, complex organizations. The platform offers significant flexibility, providing a standard cloud-based service as well as an on-premise or private cloud option for enterprises with strict data security or regulatory requirements [6]. Its architecture is built to handle distinct divisions within a single company, allowing a central administrator to create and oversee separate, isolated accounts for different departments or business units. This structure enables organizations to set specific policies for content governance, such as detection thresholds and exclusion lists, that can be applied across the entire enterprise or tailored to individual teams [1] [2]. Independent evidence confirms this capability, noting the product's support for "tenant-specific management" when integrated into third-party platforms (Figure 5). To manage the platform at scale,

Copyleaks provides a suite of administrative tools. It supports automated user provisioning through the SCIM standard, which allows companies to connect Copyleaks to their central identity systems (like an HR platform) to automatically manage user accounts [1]. For monitoring, administrators have access to a centralized analytics dashboard to track platform usage, a feature confirmed by documentation from several universities [4] [5]. This interface provides high-level metrics, such as credit consumption, for at-a-glance monitoring (Figure 5). The platform's flexibility is further enhanced by a comprehensive API and software development kits (SDKs), which enable deep integration into custom workflows and applications [7]. In addition to custom integrations, Copyleaks offers pre-built connections for common enterprise systems, including code repositories and learning management platforms.

References:

- [1] Enterprise Content Governance Copyleaks API Docs — docs.copyleaks.com/concepts/use-cases/enterprise-content-governance
- [2] How University of Michigan-Dearborn Prioritized Pedagogical... — copyleaks.com/success-story/university-of-michigan-dearborn
- [3] ChatGPT Dominates US Campuses: 700k Licenses, 14M Interactions — windowsforum.com/.../chatgpt-dominates-us-campuse...-and-14m-interactions.394273
- [4] Canvas Copyleaks Teacher Scan Tool Analytic Dashboard — staffsupport.spcollege.edu/.../-Canvas-Copyleaks-Teach...Tool-Analytic-Dash-board
- [5] How to Use Copyleaks in D2L OIT Knowledge Base — oitkb.siu.edu/knowledge-base/using-copyleaks-in-d2l
- [6] AI Detector Market Share and Industry Analysis 2040 — www.rootsanalysis.com/ai-detector-market
- [7] Copyleaks API Docs — docs.copyleaks.com

6.4 Market Adoption and Customer Success

Score: **4.5** · *Strong*

Copyleaks has established a strong market presence across educational, enterprise, and individual user segments, driven by deep product integrations and widespread industry trust. For large organizations, its API-first strategy enables embedding into core enterprise workflows, a focus validated by partnerships with platforms like the D2L Learning Management System [2] [3]. This deep integration into essential systems creates high switching costs, which is a strong indicator of customer retention. Perhaps the most compelling evidence of its standing is its use as a validation benchmark by competing AI detection tools, which signals a high degree of market trust and perceived reliability (Figure 45) [9]. The company also successfully attracts individual users by lowering the barrier to entry with a prominent free tier and easy-to-use browser extensions. The product's adoption is substantiated by direct and indirect indicators of customer success and industry recognition. Evidence from independent reporting shows Copyleaks winning contract renewals with institutional customers like universities, a direct measure of retention [6]. The company also provides specific customer case studies, offering concrete examples of successful adoption in demanding environments like higher education [4]. Its position is further confirmed by its inclusion in multiple third-party competitive analyses alongside established tools (Figure 45) [9] and by media coverage highlighting its growing popularity [7]. This market traction is also reflected in its ability to secure funding from venture capital firms [10]. Despite these strengths, the evidence for Copyleaks' market success has limitations. While user feedback is positive, the evidence lacks large-scale quantitative data on customer churn rates or active user counts to fully verify the scale of its adoption. The company claims "Fortune 500" and top university clients, but the publicly available evidence does not include comprehensive, named client lists to substantiate these broad marketing claims [8]. Furthermore, some independent tests and user reviews raise concerns about the tool's reliability and effectiveness in all situations [5]. While recognized within its niche, the product has not yet appeared in analyses by major research firms like Gartner

or Forrester, indicating its recognition is not yet at the highest enterprise level. Supporting images illustrating these findings are included below.

Supporting Evidence:




Figure 45 — Competitive Analysis of AI Detection Tools

References:

- [1] AI Image Detector Price, Features, Reviews & Ratings Capterra India — www.capterra.in/software/185429/copyleaks
- [2] Enterprise Content Governance Copyleaks API Docs — docs.copyleaks.com/concepts/use-cases/enterprise-content-governance
- [3] D2L Announces New Partnership with Copyleaks to Better Detect... — www.d2l.com/.../d2l-announces-new-partnership-w...iarism-and-ai-generated-content
- [4] Success Stories Copyleaks — copyleaks.com/success-stories
- [5] Pros and cons of AI detection University at Albany — www.albany.edu/.../pros-and-cons-ai-detection
- [6] What Universities Spend on AI Detection — \$15M+ in Data GradPilot — gradpilot.com/news/university-ai-detection-procurement-2025-investigation
- [7] Growing popularity of plagiarism and AI-detection technology... — hartfordbusiness.com/.../growing-popularity-of-plag...makes-copyleaks-one-of-cts
- [8] About Us Copyleaks — copyleaks.com/about-us
- [9] Compare Plagiarism Checkers — Honest Side-by-Sides | Noplag — noplag.com/compare
- [10] News JAL Ventures — www.jalventures.com/news

6.5 Workflow Integrations and Extensibility

Score: **4.3** · *Strong*

Copyleaks is designed to integrate its detection capabilities directly into the existing workflows of its users, rather than requiring them to work on a separate platform. For the education sector, the service provides plugins for major Learning Management Systems (LMS) like Canvas, Moodle, and Blackboard, which is confirmed by independent university support documentation [1] [2]. This allows educators to check student submissions for plagiarism and AI-generated content within their familiar grading environment. Beyond academia, Copyleaks offers add-ins for widely used document platforms like Google Docs and Microsoft Word, as well as a browser extension for Chrome and Edge that has a significant user base [3] [4] [11]. This multi-pronged approach enables users to access detection features from within the tools they use for content creation and management. For organizations that require custom solutions, Copyleaks provides an Application Programming Interface (API) and webhook support, which are described by independent sources as central features of its product strategy [5] [6] [7]. This allows developers to build the detection engine into proprietary software, and evidence shows it has been adopted by other technology platforms for this purpose. This extensibility also supports high-volume enterprise and institutional users (Figure 5). However, the user experience is not always seamless, with at least one report of a cumbersome process in an LMS integration and visual evidence of a technical error. While the integrations are extensive, the evidence supporting them comes largely from the vendor, with a notable lack of formal, third-party technical audits or in-depth performance reviews for the API and various plugins.

Supporting Evidence:

Comparison Chart		
	Evalato	Award Force
Program templates	✓	✗
Custom email	✓	✗
Embeddable registration form	✓	✗
Points voting	✓	✗
Positional voting	✓	✗
Score normalization	✓	✗
Plagiarism detection	✗	✓ but requires a paid Copyleaks account
Intercom integration	✓	✗
Salesforce integration	✗ but available through Zapier	✓
Free trial	✓	✗

Figure 46 — Curated Evidence: Award Force Comparison Chart




Figure 47 — Curated Evidence: Copyleaks screenshot_7838a51f.webp

References:

- [1] Canvas Plagiarism Checker & AI Detection Plugin Copyleaks — copyleaks.com/learning-management-systems/canvas-plagiarism-checker
- [2] Using Copyleaks: A Guide for Students Student Support — studentsupport.spcollege.edu/.../-Using-Copyleaks-A-Guide-for-Students
- [3] Copyleaks Google Docs Add-on | AI & Plagiarism Detection — copyleaks.com/blog/copyleaks-launches-google-docs-add-on

- [4] AI and Plagiarism Checker Copyleaks Google Workspace — workspace.google.com/marketplace/app/ai_and_plagiarism_checker_copyleaks/
- [5] Copyleaks APIs | Real-Time Content Integrity Solutions — copyleaks.com/api
- [6] Copyleaks Review: Plagiarism Veteran, Newer AI Detector Leap AI — www.tryleap.ai/review/copyleaks
- [7] Webhooks Overview Copyleaks API Docs — docs.copyleaks.com/reference/data-types/authenticity/webhooks/overview
- [8] Webhooks Security Copyleaks API Docs — docs.copyleaks.com/concepts/security/webhooks
- [9] 6 Best Copyleaks Alternatives for AI Detection in 2026 Leap AI — www.tryleap.ai/alternatives/copyleaks
- [10] AI Detector Extension for Chrome & Edge: Verify AI Content in Browser — copyleaks.com/ai-detector/extension
- [11] AI Content Detector Copyleaks Chrome Web Store — chromewebstore.google.com/.../gplcmncpkldjicbknjjkoidpgkcakd

Disclosures

About this report. huby produced this evaluation report for Copyleaks, using its own tests, publicly available research and evaluation data, and data available from the product company. Scores follow huby's published methodology: six categories, a category → subcategory → factor hierarchy, and weighted scores from 1.0 to 5.0. The full methodology is available at huby.ai/methodology.

Point-in-time. Information was collected between September 8, 2026 and September 21, 2026. AI products change often, so features, pricing, policies and performance may have changed since then. Check anything important with the vendor.

Sources and accuracy. This report draws on the results from a large dataset including huby's test harness, independent research, documentation, third-party reviews and the vendor's public materials. AI-assisted analysis solely based on this data has also been used in producing this report. AI can misread or leave out information. huby does not guarantee that this report is accurate or complete.

Independence. The company behind Copyleaks did not pay for, commission or review this report. In general, huby does not have a financial relationship with a vendor. If it does, these relationships do not affect the analysis and scores.

Not advice. This report is for informational purposes only. It is not legal, financial, security, compliance or procurement advice. Base purchase decisions on your own evaluation.

Trademarks. Product names and logos belong to their owners. Mentioning them does not imply endorsement.

Corrections. If you see an error, contact support@huby.ai and cite the report name. huby reviews reported errors and updates reports when needed.