MARLUC

CASE STUDY · MARLUC SYNOPSIS

Vanderbilt University and Turnitin's AI detector

IMPLEMENTED 2023 · EDUCATION · US · BAND E · REVERSED

VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

In August 2023 Vanderbilt switched off Turnitin's AI writing detector for the foreseeable future, after several months of testing. It cited lack of insight into how the tool works, false positive risk and privacy.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

No measurable effect: withdrawn after a test.

AI used: AI writing detection feature inside Turnitin's plagiarism checking product.

Implemented: 2023. Grade: band E, No measurable effect, Reversed.

Headline result, as reported: Vanderbilt disabled Turnitin's AI detector in August 2023 and said it does not believe AI detection software is an effective tool

1 · The case

The facts, as reported by the 4 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Vanderbilt University. US private research university. The decision applied to its use of Turnitin through its Brightspace learning platform, as published on the Brightspace site. Canvas block: 8. Governance and risk. Also 6. Data and technology, 5. People and skills.

Business challenge

The challenge. In April 2023 Turnitin released an update that showed how much of a submitted paper it judged to have been written by AI. Vanderbilt says staff had concerns about how the product worked and how reliable the results would be.

Who implemented it

Turnitin (Software vendor): Supplied the plagiarism checking product and the AI writing detection feature..

Technology (SEAM pillar E)

SEAM types: T2 Embedded AI in existing software. Deployment: Vendor SaaS. Vendors and products: Turnitin, Turnitin AI writing detection.

Where it was built in: A feature added to Turnitin's existing paper-checking product; according to Vanderbilt it was enabled for customers with less than 24 hours' notice.

What changed

Before: the AI indicator was available to Vanderbilt through Turnitin. After: the university turned it off and pointed instructors to early conversations with students, clear expectations, citation of AI use and redesigned assignments. [1] [2] [3] [4]

What it cost

Not reported.

What it saved

Not reported. [1] [2] [3] [4]

Governance and control

Vanderbilt's stated reasons were short notice with no way to switch the feature off, no insight into how it works, false positive risk, evidence of higher flagging of non-native English writers, and privacy concerns. All of this is the university's own account.

S2 · Accountable owner · Not reported
Not reported. The post is published on the university's Brightspace site without a named author or owner. [1]

A1 · Safe use · Claimed
Vanderbilt's guidance tells instructors to set expectations early, require disclosure and citation where AI is allowed, and consider redesigning assignments. [1]

A2 · Human in the loop · Claimed
For suspected misuse the guidance suggests comparing the writing with the student's other work, checking sources and facts, and talking to the student. [1]

A3 · Monitoring and assurance · Claimed
The university says it used and tested the tool for several months, met Turnitin and other AI leaders, and spoke to other universities before deciding. [1]

A4 · Vendor AI · Claimed
The university cites the short notice, the lack of an option to disable the feature, no insight into how it works, and privacy concerns about sending student data to a third party. [1]

A5 · Agent controls · Not reported
No agents described.

A6 · Standards and regulation · Not reported
Not reported.

What worked, what did not, what is unproven

WORKED

– The university made a decision on its own review and published its reasons.

– The post gave instructors practical alternatives: early discussion, disclosure and citation, and assignment redesign.

DID NOT

– Not reported: the post does not report any measured results from the testing period, such as the number of papers flagged.

UNPROVEN

– The 750 figure is the university's estimate from the vendor's stated rate applied to 2022 submission volume, not a count of real cases.

– How many papers the detector flagged at Vanderbilt: Not reported.

– Whether the decision reduced or increased AI misuse: Not reported.

Timeline

Apr 2023 Turnitin releases an update that reports how much of a submitted paper was written by AI [1]

Jun 2023 Turnitin's Chief Product Officer publishes document-level and sentence-level false positive rates [2]

Aug 2023 Vanderbilt announces it has disabled Turnitin's AI detection tool for the foreseeable future [1]

Sept 2026 An Australian academic review of the sector reports that unis are moving away from AI-detection software [4]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: E, No measurable effect. Proof: Reversed. Level of change: Not graded.

No level assessed: band E. The tool was tested and then withdrawn, with no change to processes or roles reported.

Vanderbilt tested the detector for several months and then disabled it before using it to judge student work, as far as its post says. No harm, cost or sanction is reported, so this is band E, withdrawn after a test.

The SEAM view and lessons learnt

The decision rested on three points the university could check or ask about: how the vendor's tool works, what its error rate means at the university's volume, and where student data goes. The vendor's own later figures separate document-level and sentence-level rates, which matters when a flag is used in a misconduct process. For any larger organisation buying AI features inside existing software, the transferable point is to test, understand the error rate at your own scale and decide the rules for use before switching a feature on.

1. Ask a vendor how an AI feature works and what its error rate means at your volume.
2. Decide who may act on a flag, and on what evidence, before use.
3. Keep the option to switch a vendor feature off.

Evidence

Score 1 of 4. Own account: the university's published decision and reasons, with the vendor's published false positive figures. No independent evaluation of the decision.

Evaluation method. Organisation's own published statement, read against the vendor's published statistics; no independent check.

Sources disagree. Turnitin's June 2023 post puts its document-level false positive rate at less than 1% for documents with 20% or more AI writing and its sentence-level rate at around 4%. Vanderbilt applied a 1% rate to all its 2022 submissions to reach about 750 papers. Turnitin's rate is stated for documents with at least 20% AI writing, so the university's figure is an illustration.

Status. Decision announced 16 August 2023. Current status of Vanderbilt's use: Not reported.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/vanderbilt-turnitin-ai-detection-disabled

[1] Vanderbilt University: Guidance on AI Detection and Why We're Disabling Turnitin's AI Detector
https://www.vanderbilt.edu/brightspace/2023/08/16/guidance-on-ai-detection-and-why-were-disabling-turnitins-ai-detector/

[2] Turnitin: Understanding the false positive rate for sentences of our AI writing detection capability
https://www.turnitin.com/blog/understanding-the-false-positive-rate-for-sentences-of-our-ai-writing-detection-capability

[3] arXiv: GPT detectors are biased against non-native English writers (Liang et al.)
https://arxiv.org/abs/2304.02819

[4] The Conversation: Unis and schools are moving away from AI-detection software. They should stop using it altogether
https://theconversation.com/unis-and-schools-are-moving-away-from-ai-detection-software-they-should-stop-using-it-altogether-292593

Vanderbilt University website: https://www.vanderbilt.edu

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/vanderbilt-turnitin-ai-detection-disabled

VERSION HISTORY