

Teaching Audio Deepfake Detection in Undergraduate Computing Courses

An Exploration with Scenario-Based Caselets Combining Computer Science and Linguistics

A. Komlodi¹, C. Kulkarni², L. Neelakandan¹, V.P. Janeja¹, C. Mallinson¹, L. K. Chen¹

¹University of Maryland, Baltimore County (UMBC) | ²Towson University |



Acknowledgment: This work is supported by NSF Award #2346473.
QR code: Community Infrastructure to Strengthen AI for Audio Deepfake Analysis (CISAAD)

Context & Problem Statement

Audio deepfakes: AI-generated speech that mimics real human voices and can enable fraud, impersonation, and misinformation [3].

Curricular gap: Deepfake detection is a key skill for computing students, but it is hard to teach because it spans computer science, signal processing, and linguistics.

Research questions

1. Do caselets improve students' self-reported knowledge?
2. Do results differ by teaching approach, major, or class standing?
3. How can caselets be improved based on student input?

Interdisciplinary Framework

Computer Science
Detection algorithms: features, training data, model choice, classification, and evaluation.

Signal Processing
Analyzing the audio aspects of recordings, including preprocessing.

Linguistics
Natural human language: pitch, pauses, consonant sounds, breath, and audio quality [4].

Caselet: "Stopping Scam Calls"

A short case study solved in one class session [1]: a banking-fraud scenario in which students develop an automated system to detect deepfake scam calls, in three phases [2, 4]:

Problem formulation and model selection
Listen to audio clips, judge authentic vs. deepfake, and choose a detection approach.

Feature engineering
Linguistic features (pitch, pauses, breath, audio quality) and preprocessing.

Model validation
Confusion matrix interpretation, overfitting analysis, and false positives vs. false negatives.

Methodology

Comparative study: one caselet, two teaching approaches, two undergraduate computing courses (N = 53).


Course A: Computational Thinking (n = 29)
For students without a programming background. Guided format: a 30-minute instructor-led activity with short lectures and structured debriefing.

Course B: Data Science (n = 24)
Analytics with real-world datasets. Collaborative format: a brief ML overview, then peer work on the caselet with instructor facilitation.

Measures: pre and post-caselet self-reported knowledge (5-point Likert); three open-ended questions on new learning, lesson improvements, and desired topics.

Analysis: data were not normally distributed (Shapiro-Wilk), so non-parametric tests were used: Wilcoxon signed-rank (gains), Cohen's d (effect size), Mann-Whitney U and Kruskal-Wallis H (group differences). Two researchers thematically coded open-ended responses.

Results: Self-Reported Knowledge Gains



No significant differences in gain between groups: course p = .87, major p = .87, class standing p = .91.

Study Strengths

Large effect: d = 1.04 overall from one short classroom activity.

Both teaching approaches worked: guided (d = 1.17) and collaborative (d = 0.94).

Broad reach: large effects for non-tech (d = 1.42) and tech (d = 0.98) majors; no group differences.

Mixed methods: survey gains plus feedback to guide redesign.

Student Feedback Themes

More audio practice: more, and more diverse, deepfake examples.

Real-world applications: how organizations implement prevention.

Gamification: requested even in the collaborative course.

Linguistic focus: valued learning to identify linguistic features.

Current research: context on recent developments in the field.

Recommendations for Educators

Anchor in scenarios: ground abstract ML in scam-call threats.

Integrate disciplines: link ML, signal processing, and linguistics.

Flexible delivery: guided and collaborative formats both worked.

Discuss uncertainty: examine detection errors and their costs.

Stay classroom-ready: fit short caselets into existing courses.

Limitations and Next Steps

Limitation	Planned next step
Knowledge was self-reported, right after the activity.	Add formal learning assessments.
Long-term retention and skill transfer were not measured.	Measure retention and skill transfer over time.
Actual detection skill was not formally tested.	Measure detection skill, not just perceived knowledge.
Small subgroups: seniors (n = 6) improved but not significantly.	Expand caselet adoption across courses and universities.
	More practice examples, real-world applications, interactive elements, linguistic instruction, and research connections.

References

- [1] Cherukuri, J. (2025). Beyond lectures: Innovative approaches to promote student engagement through scenario-based learning (SBL). *Journal of Engineering Education Transformations*, 570–575.
- [2] Janeja, V. P., & Mallinson, C. (2024). Transdisciplinarity in audio deepfake discernment with expert-in-the-loop AI models. *Proc. UMBC Computing Research Symposium*, 12–18.
- [3] Khanjani, Z., Watson, G., & Janeja, V. P. (2023). Audio deepfakes: A survey. *Frontiers in Big Data*, 5, 1001063.
- [4] Nwosu, K., Evered, C., Khanjani, Z., Bhalli, N., Davis, L., Mallinson, C., & Janeja, V. P. (2023). Auto annotation of linguistic features for audio deepfake discernment. *Proc. AAAI Symposium Series*, 2(1), 242–244.