Abstract: Trojan Keywords and Contextual Drift: Identifying AI Dependency in Case Study Analysis

The growing use of generative AI tools by students has complicated the evaluation of case study–based assessments, particularly in applied disciplines such as business, marketing, and hospitality. Existing AI-detection tools offer limited reliability and raise ethical and equity concerns, especially for ESL learners. This presentation introduces *Trojan keywords and contextual Trojan signals* as a practical, classroom-embedded method for identifying uncritical AI dependency during assessment evaluation.

Drawing on classroom implementation using the *Tiffin Box by Silk* case study, the presentation highlights a recurring phenomenon termed *contextual drift*, where AI-generated outputs silently override core case facts. In this instance, although the company in question is based in British Columbia, a majority of the student submissions developed customer personas that were geographically segmented in Toronto. This drift was traced to an AI system that overweighted a comparative reference provided for contextual understanding, resulting in misaligned geographic assumptions that were reproduced across submissions.

Rather than functioning as evidence of misconduct, Trojan signals are presented as diagnostic indicators that prompt instructional intervention. Findings from two course implementations suggest that this approach enhances student engagement with case facts, strengthens analytical justification, and promotes a more responsible use of AI as an assistive tool rather than a substitute for reasoning. The session concludes with an exploration of ethical considerations, implementation guidelines, and examples of how this low-tech, scalable method can be applied across disciplines without relying on automated detection software.