CARLETON UNIVERSITY
SPROTT SCHOOL OF BUSINESS
BUSI 6306

ARTIFICIAL INTELLIGENCE METHODS IN BUSINESS RESEARCH

Instructor: Dr. Ahmed Doha
Instructor Email: ahmed.doha@carleton.ca (only means of online communication)
Phone Number: (613) 520-2600 ext. 1909
Office: NI 5036

Schedule: One weekly three-hour class. Official meeting dates, time, room, and University-scheduled breaks are published in Banner and Brightspace for the current offering.

Delivery mode: In-class doctoral research seminar. Students' remote attendance is not permitted unless formally authorized.

Lecture Room: To be confirmed by the University; check Banner and Brightspace before classes begin.

Pre-requisites & precluded Courses: Experience in Python programming is recommended. Precludes additional credit for BUSI 6906 (no longer offered).

Course Calendar description from the current University calendar:

Foundations of AI methods, including machine learning, deep learning, and language processing, with emphasis on their applications in decision-making, hypothesis testing, and innovative, cross-disciplinary business research. Experience in Python programming is recommended. Includes: Experiential Learning Activity. Precludes BUSI 6906.

Course Description: This doctoral research studio examines how AI-enabled business research turns raw traces into representations, measures, generated outputs, and relational structures—and how researchers determine whether the resulting claims are reproducible, valid, and managerially meaningful. Four cycles—representation, measurement, LLMs and agents, and networks—each move from detailed methods and a focal paper, to instructor-led reconstruction, to a student reconstruction and research extension. Class 13 is an individual capstone proposal conference.

Learning Objectives:

- Translate a consequential business-research problem into an explicit analytical object, estimand or predictive target, validation plan, and decision implication.
- Construct and critically evaluate text representations, AI-augmented measurements, LLM and agent systems, and business networks.
- Explain, implement, and challenge consequential design choices in contemporary AI methods.
- Build provenance manifests, target registers, executable pipelines, discrepancy ledgers, and reproducibility archives for published empirical research.
- Distinguish computational reproduction, robustness analysis, empirical replication, and theoretical or methodological extension.
- Assess construct validity, leakage, temporal validity, evaluator validity, graph-construction validity, uncertainty, generalizability, and managerial usefulness.
- Communicate methodological reasoning across marketing, finance, strategy, organization studies, operations and supply chain, and information systems.
- Develop and defend a feasible research extension grounded in reconstructed evidence and publicly obtainable data.

Required Materials

The journal articles provide the business-research applications. The following books provide the systematic methods instruction. Assigned chapters and any reading adjustments will be posted in Brightspace.

- **Required student purchase:** Text as Data. Grimmer, Roberts, and Stewart (Princeton University Press, 2022). This is the research-design spine and the principal methods text for representation and measurement. Students who already have lawful full-text library access do not need to buy another copy.
- **Required student purchase:** Network Analysis. Rawlings, Smith, Moody, and McFarland (Cambridge University Press, 2023). This supplies network measurement, boundary specification, missing-data analysis, centrality, communities, two-mode networks, relational inference, diffusion, and influence. Students who already have lawful full-text library access do not need to buy another copy.
- **Carleton Library · O'Reilly Learning (no separate purchase):** Hands-On Large Language Models. Alammar and Grootendorst (O'Reilly, 2024). Selected chapters support Cycle 1 representation and Cycle 3 prompting, retrieval, retrieval-augmented generation, and agents.
- **Carleton Library · O'Reilly Learning (no separate purchase):** AI Engineering. Huyen (O'Reilly, 2025). Evaluation chapters support Cycle 2; system, retrieval, agent, data, monitoring, and reliability chapters anchor Cycle 3.
- **Carleton Library · O'Reilly Learning (no separate purchase):** Graph Neural Networks in Action. Stillman and Broadwater (Manning, 2025). This is Cycle 4's computational text for graph embeddings, message passing, graph convolutional networks, GraphSAGE, graph attention networks, link prediction, and project design.
- **Carleton Library · O'Reilly Learning (no separate purchase):** Network Science. Pinheiro (Wiley, 2022). Selected Cycle 4 chapters cover communities, projections, node similarity, centrality, graph optimization, and business applications.
- **Free author-hosted access (no purchase):** Speech and Language Processing, 3rd-edition online manuscript. Jurafsky and Martin (third-edition online manuscript). Selected chapters support tokens, classification, embeddings, large language models, transformers, information retrieval, and retrieval-augmented generation. Assigned chapters will identify the version used for the current offering.
- **Free official companion (no purchase):** Network Analysis with R companion. Rawlings et al. The companion is written in R and supplies runnable tutorials and data. The linked centrality and hierarchy chapter is assigned in Cycle 4; Python implementations used in class are supplied in the course repository.

Methods Texts by Cycle

- Cycle 1 · Representation: Text as Data; Hands-On Large Language Models; and Jurafsky and Martin.
- Cycle 2 · Measurement: Text as Data plus Huyen's evaluation chapters.
- Cycle 3 · LLMs & Agents: AI Engineering plus selected Hands-On Large Language Models chapters and Jurafsky and Martin's retrieval material.
- Cycle 4 · Networks: Graph Neural Networks in Action and selected Network Science chapters, with Rawlings et al. supplying network measurement, research design, and inference.
- The four focal articles, online appendices, and faculty-curated readings remain required business-research applications of these methods.
- A laptop capable of running Python, Jupyter, and Git is required in every class. A precomputed path will be supplied for any exercise that would otherwise require paid API access.
- There is no conventional final examination.

Course Evaluation

Grade Distribution

Cycle 1 · Representation research dossier	*15%*
Cycle 2 · Measurement research dossier	*15%*
Cycle 3 · LLMs & Agents research dossier	*15%*
Cycle 4 · Networks research dossier	*15%*
Research workshop contributions (4 × 4%)	*16%*
Final written research proposal	*20%*
Final proposal presentation and oral defense	*4%*

Cycle Research Dossiers · 15% Each

Each student submits one dossier for the published paper reconstructed in each cycle's third class. All dossiers include a claim map, data and code provenance, registered targets, an executable focal result, a published-versus-reconstructed comparison, a discrepancy ledger, a research extension, a draft methods section, and one proposed exhibit. The cycle-specific analytical work is:

- Representation: define the corpus and unit of analysis; reproduce how text becomes tokens, embeddings, axes, topics, or concepts; recover the focal result; test construct validity, geometry, stability, and transportability; and specify an improved representation.
- Measurement: define the construct and estimand; reproduce how human, synthetic, or behavioural traces become a measure or decision quantity; recover the focal estimator or prediction; test bias, uncertainty, leakage, timing, and decision value; and specify an improved measurement design.
- LLMs & Agents: freeze the model and version, prompts, retrieval, tools, routing and stopping rules, evaluator, and aggregation; preserve outputs and logs; recover the focal result; test prompt, judge, drift, grounding, and agent-policy sensitivity; and specify the extension.
- Networks: define nodes, edges, direction, sign, and time; rebuild the graph, structural measures, and focal Graph Neural Network model or forecast against non-graph baselines; test communities, centrality, link prediction, null models, and topology sensitivity; and specify the business-research extension.

Research Workshop Contributions · 16%

Four workshop contributions are assessed at 4% each in Classes 3, 6, 9, and 12. Students contribute through assigned roles such as author, computational auditor, theory discussant, methods reviewer, and editor. Credit reflects documented preparation, command of the paper and evidence, specific and constructive criticism, and the quality of the resulting revision guidance—not frequency of speaking alone.

Final Written Research Proposal · 20%

Each student develops the strongest extension from the four cycle dossiers into one coherent and feasible business-research proposal. It must state the consequential business problem, theory gap, publicly obtainable data, proposed AI method, estimand or predictive target, preliminary feasibility evidence, validity threats, expected theoretical and managerial contributions, and reproducibility plan. Detailed length and submission requirements will be posted in Brightspace.

Final Proposal Presentation and Oral Defense · 4%

In Class 13, each student presents the final research proposal and responds to methodological, theoretical, feasibility, and managerial challenges. Assessment focuses on clarity, command of reconstructed evidence, direct responses to questions, and willingness to revise the design when the evidence warrants it.

Late Work, Academic Consideration, and Reappraisal

Workshop documents structure the work of classmates, so deadlines are meaningful. Students should contact the instructor before a deadline when possible. Academic consideration and accommodation requests follow Carleton's current procedures. A request to reappraise a mark must identify the criterion and evidence at issue and should be submitted within seven calendar days of receiving feedback. Reappraisal may raise, maintain, or lower the grade.

Tentative Schedule

Class	Topics
1	**Representation · Methods and focal-paper introduction** Represent language as sparse vectors, contextual embeddings, semantic axes, neighbourhoods, and guided concepts. Interrogate when geometry becomes measurement. **Focal reading:** Lee, Cheng, Mao & Manzoor (2025), "Guided Diverse Concept Miner," Information Systems Research. **Evidence produced:** representation model card, validity portfolio, focal claim map, and reconstruction targets.
2	**Representation · Instructor reconstruction** Audit 23,720 WCAI journeys; execute clean and released paths; diagnose split and configuration discrepancies, perfect row-index leakage, AUC gaps, concept recall, and instability. **Evidence produced:** executable audit, reported-versus-rerun panel, leakage verdict, concept-validity ledger, and research gaps.
3	**Representation · Student research workshop** Defend a faculty-curated representation-paper reconstruction and propose an extension in construct recovery, stability, transportability, or managerial actionability. **Evidence produced:** reconstruction dossier, representation multiverse, extension, and draft methods section.
4	**Measurement · Methods and focal-paper introduction** Separate prediction from measurement; derive synthetic-data bias, naïve pooling, AI-augmented estimation, prediction-powered inference, uncertainty, and equivalent-human-sample value. **Focal reading:** Wang, Zhang & Zhang (2026), "Large Language Models for Market Research," Marketing Science. **Evidence produced:** estimand table, bias decomposition, AAE derivation, claim map, and targets.
5	**Measurement · Instructor reconstruction** Audit vaccine and sports-car conjoint objects; rebuild coefficient-error curves and data-saving calculations; test why naïve synthetic substitution fails. **Evidence produced:** trial-object audit, executable error curves, recovered 24.9%–79.8% range, and sensitivity ledger.
6	**Measurement · Student research workshop** Defend a curated measurement and decision-policy paper; connect predictive fit with construct validity, uncertainty, cost, and action. **Evidence produced:** measurement dossier, utility/validity verdict, individual extension, and methods section.
7	**LLMs & Agents · Methods and focal-paper introduction** Treat model, prompt, role, order, retrieval, tools, routing, stopping rules, judge, aggregation, and version as components of an empirical treatment. **Focal reading:** Doshi, Bell, Mirzayev & Vanneste (2025), "Generative Artificial Intelligence and Evaluating Strategic Decisions," Strategic Management Journal. **Evidence produced:** LLM-and-agent system manifest, evaluator design, drift register, claim map, and targets.

Class	Topics
8	**LLMs & Agents · Instructor reconstruction** Audit two 60-business-model studies and 28 targets; reconstruct diversity-versus-scaling arithmetic; expose the boundary between a published-table audit and raw-response reproduction. **Evidence produced:** evidence ladder, Tables 3–4 audit, aggregation decomposition, access verdict, and robust evaluator extension.
9	**LLMs & Agents · Student research workshop** Defend a faculty-curated 2024+ LLM or agent reconstruction; interrogate treatment bundles, agent policy, judge validity, drift, contamination, and inferential level. **Evidence produced:** frozen output corpus and logs, full system manifest, evaluator audit, proposed study, and methods section.
10	**Networks · Methods and focal-paper introduction** Define nodes and edges; derive adjacency and message passing; analyze signed and directed Graph Neural Networks, centrality, communities, null graphs, and topology sensitivity. **Focal reading:** Liu et al. (2025), "Beyond Complements and Substitutes," Information Systems Research. **Evidence produced:** edge contract, hand-worked Graph Neural Network, network validity plan, retail decision map, and targets.
11	**Networks · Instructor reconstruction** Audit 3,000,888 Favorita rows; build a 33-family signed directed graph; compare own-history and graph-aware forecasts; analyze communities, centrality, and topology sensitivity. **Evidence produced:** graph ledger, 161-edge network, forecast panel, community analysis, scoped verdict, and extensions.
12	**Networks · Student research workshop** Defend a curated network and Graph Neural Network reconstruction and stress-test node, edge, time, projection, and aggregation decisions. **Evidence produced:** near-independent dossier, topology multiverse, research proposal, and rigorous methods section.
13	**Individual capstone proposal conference** Defend the strongest extension developed during the four cycles as one coherent, feasible, and publicly auditable business-research study. **Evidence produced:** individual proposal defense, editorial decision, revision contract, and archival plan.

The sequence above contains 13 instructional classes. Official meeting dates, University-scheduled breaks, academic deadlines, assignment deadlines, and reading-pool allocations are published through Banner and Brightspace for the current offering.

ADDITIONAL INFORMATION

Course Sharing Websites

Materials created for this course—including the browser sites, teaching materials, datasets prepared for instruction, exercises, assignments, and assessments—remain the intellectual property of their author(s). They are intended for enrolled students' personal academic use and may not be reproduced or redistributed without prior written consent. Third-party papers, software, and data remain subject to their own licences and terms.

Group Work

The Sprott School of Business encourages group work because it develops interpersonal, communication, leadership, followership, and integrative problem-solving skills. In this course, some reconstruction execution and workshop preparation may be collaborative. Every student remains individually responsible for the extension argument and methods writing. Group work never authorizes copying another member's prose, analysis, or interpretation without attribution.

Use of Generative AI and Computational Tools

Generative AI is an object of study and an authorized research tool when used transparently. Unless an assignment states otherwise, students may use AI to troubleshoot code, generate alternative specifications, critique prose, or explore designs. Every submitted portfolio must record the provider and model, date or version, prompts or interaction protocol, relevant outputs, human verification, and the student's substantive decisions. Students remain responsible for correctness, citations, confidentiality, licensing, and originality.

Prohibited practices include fabricating data, citations, outputs, logs, or execution claims; uploading confidential or restricted data to an external model; presenting model-generated analysis as personally executed; and using AI to complete another student's assigned role. Undisclosed use may constitute misrepresentation under Carleton's Academic Integrity Policy.

In accordance with the Carleton University Graduate Calendar, the letter grades assigned in this course have the following percentage equivalents:

A+ = 90–100	B+ = 77–79	C+ = 67–69	D+ = 57–59
A = 85–89	B = 73–76	C = 63–66	D = 53–56
A− = 80–84	B− = 70–72	C− = 60–62	D− = 50–52
F = Below 50			

Grades entered by Registrar:
WDN = Withdrawn from the course
DEF = Deferred

Academic Regulations

University rules regarding registration, withdrawal, academic standing, appeals, grades, and other graduate regulations can be found in the Carleton University Graduate Calendar.

Requests for Academic Accommodation

Carleton is committed to providing academic accessibility for all individuals. You may need special arrangements to meet your academic obligations during the term. The accommodation request processes—including academic consideration for medical and other extenuating circumstances, disability-related accommodation, pregnancy, religious obligations, survivors of sexual violence,

and student activities—are outlined at students.carleton.ca/course-outline. Students should initiate requests as early as possible and follow the University's stated timelines.

Academic Integrity

Violations of academic integrity—including plagiarism, unauthorized collaboration, misrepresentation, fabrication or misrepresentation of research data, and completing work for another student—are serious academic offences. Students are expected to understand and follow the Carleton University Academic Integrity Policy. In this course, reproducibility records are assessed evidence; fabricated runs, altered logs, invented citations, or false statements about data access are treated as research-data misrepresentation.

Centre for Student Academic Support

The Centre for Student Academic Support provides academic writing, study, and learning support. Visit carleton.ca/csas for current services.

Important Information:

- Students must retain copies of all submitted work, data, code, and execution records.
- All final grades are subject to the Dean's approval.
- Course correspondence must be sent from the student's valid Carleton email account and include the student's full name and student number.
- Official class details, dates, deadlines, and reading adjustments will be communicated through Brightspace.