

DGIN 5401: Operationalized Machine Learning in Healthcare

Course Syllabus – Fall 2026

Faculty of Computer Science, Dalhousie University

Course identification	DGIN 5401; 3 credit hours
Term	8 September–9 December 2026; first meeting is Thursday, 10 September
Learning modality	In person
Class meetings	Thursdays, 17h30–20h30, Studley Campus, McCain Arts and Social Sciences Building, room 2102
Course website	https://web.cs.dal.ca/~rudzicz/Teaching/DGIN5401/2026/
Instructor	Frank Rudzicz; frank@dal.ca ; 1577 Barrington; office hours by appointment
Course systems	Course website, Microsoft Teams, the private class GitHub repository, and Dalhousie email. Brightspace is not used for routine course activity.
Syllabus version	1 September 2026

1 Key dates

Unless an instruction states otherwise, out-of-class deadlines are at 23h59 Atlantic Time on the date shown. Detailed requirements appear in the Course Project Handbook and on the course website.

Date	Milestone or assessment
16 September	Ranked research and stewardship interests due
21 September	Individual question card due
24 September	Project proposal and initial manuscript due
8 October	Literature-synthesis checkpoint due
22 October	Dataset-analysis checkpoint due
29 October	In-person quiz/essay 1; project design review and scope lock
5 November	Methods checkpoint due
19 November	In-person quiz/essay 2; project evidence review
26 November	Experiments-and-analysis checkpoint due
3 December	Individual project presentations and integrated demonstration
9 December	Final individual paper, reproducibility package, and shared system dossier due

2 Recognition of place and history

Land acknowledgement. Dalhousie University operates in the unceded territories of the Mi'kmaw, Wolastoqey, and Peskotomuhkati Peoples. These sovereign nations hold inherent rights as the original peoples of these lands, and we each carry collective obligations under the Peace and Friendship Treaties. Section 35 of the Constitution Act, 1982 recognizes and affirms Aboriginal and Treaty rights in Canada.

African Nova Scotian acknowledgement. We recognize that African Nova Scotians are a distinct people whose histories, legacies and contributions have enriched that part of Mi'kma'ki known as Nova Scotia for over 400 years.

3 Course purpose, content, and outcomes

3.1 Description

Operationalizing machine learning in healthcare requires more than training an accurate model. This course examines the chain from clinical purpose and care processes to data, study design, modelling, evaluation, workflow, deployment, monitoring, and accountable use. Topics include health-data provenance and missingness; leakage-resistant study design; prediction, calibration, uncertainty, thresholds, and decision consequences; temporal and dynamic prediction; clinical natural-language processing and foundation models; explainability, subgroup analysis, fairness, privacy, safety, and human factors; reproducible systems, containerization, MLOps, monitoring, governance, and incident response.

The course combines instructor-led material, critical discussion of current research, student-led seminars, and an applied research project. The class will use the fully synthetic AIRO-Sim cardiac-surgery dataset to build and study a shared, transparent risk-system prototype. Each student will pursue a distinct research question, steward at least one shared artefact, and submit an independently written research paper. The class will also produce a concise shared system dossier and an integrated demonstration.

3.2 Rationale

Health and machine-learning systems are sociotechnical systems. A model can fail because the clinical question is poorly specified, the data encode care processes or inequities, the evaluation is optimistic, the interface induces misuse, the system cannot be reproduced, or performance changes after deployment. This course prepares students to reason across those boundaries and to make claims that are useful, testable, and deployable. Breadth comes from the lectures and seminars; depth comes from a cumulative, evidence-driven project.

3.3 Prerequisites and expected preparation

Formal prerequisites: None.

Expected preparation: Students should be comfortable programming in Python and should have prior exposure to probability, linear algebra, statistics, and machine learning. Students without this background should consult the instructor at the beginning of the term and complete the supplied readiness material. The course will revisit essential ideas when they are needed, but it is not an introductory programming or introductory machine-learning course.

3.4 Learning outcomes

By the end of the course, students should be able to:

1. trace how clinical work, measurement, documentation, coding, and extraction produce health data, and identify threats involving provenance, missingness, labels, timing, and linkage;
2. formulate a defensible machine-learning or analytical study by specifying the intended use, population, unit of analysis, prediction time, outcome or estimand, comparator, decision context, and claim boundary;
3. design reproducible, leakage-resistant data and modelling pipelines that respect patients, time, sites, and other grouping structures;
4. choose and compare transparent baselines and more complex approaches for structured, temporal, textual, or multimodal health data without treating complexity as a contribution by itself;

5. evaluate predictive and analytical claims using appropriate discrimination, calibration, measurement, workload, and decision-oriented analyses;
6. assess explainability, fairness, privacy, safety, human factors, workflow, governance, and accountability as properties of an entire system rather than of an algorithm alone;
7. build and document a locally runnable research prototype with versioned data and outcome contracts, a tested inference interface, a case viewer, a reproducible environment or container, and a monitoring demonstration;
8. detect and interpret data-quality, population, workflow, and model-behaviour shifts, and propose proportionate responses;
9. critically read, present, and discuss research papers for an interdisciplinary clinical–technical audience, distinguishing what the evidence shows from what it merely suggests; and
10. conduct an independent research study, communicate its methods and evidence clearly, attribute shared and automated contributions accurately, and explain what fully synthetic data can and cannot establish.

4 How the course operates

4.1 Weekly format and preparation

Each three-hour meeting may combine a lecture or applied activity, one or more student paper presentations, discussion, a project stand-up, peer review, and research work. Students should complete assigned readings before class and arrive prepared to discuss the paper’s question, design, decisive evidence, limitations, and relevance to the shared project. Students should bring a laptop when possible for project activity.

The course is intentionally adaptive. Student-selected papers, project findings, guest availability, and class needs may change the balance of activities within a week. The central learning outcomes, assessment weights, required project backbone, and final evidence standards remain in force.

4.2 Course website, Teams, GitHub, and email

The course website is the public home for this syllabus, lecture slides, project handouts, the seminar-paper list, and the AIRO-Sim tutorial. Microsoft Teams is the operational source for announcements, paper approvals, scheduling changes, weekly-reflection submission, project registries, and other submission instructions. The private class GitHub repository is the working home for shared code, issues, reviewed changes, versioned interfaces, and integration. Dalhousie email may be used for announcements and private communication.

Students are responsible for monitoring Teams and Dalhousie email. **Questions about grades, accommodations, absences, interpersonal or project difficulties, or personal circumstances must be sent privately through Teams or email.** Questions sent within 24 hours of a deadline may not receive an answer before that deadline. Do not include private or identifying information in a public Teams post or repository issue.

The syllabus states official course-level requirements. The Course Project Handbook governs the research process and deliverables. The AIRO-Sim Technical Brief governs release-specific data, prediction-time, outcome, and system-interface details. Current links to both documents will be maintained on the course website.

4.3 Attendance, participation, and devices

Attendance is not assigned a separate mark, but the course depends on in-person seminar discussion, quizzes, design and evidence reviews, peer feedback, and integration work. These activities normally cannot be reproduced through asynchronous or remote work. Participation credit is based on prepared and constructive contribution rather than physical presence alone or frequency of speaking.

During quiz/essays, computers, phones, smart watches, generative-AI tools, and other electronic devices are not permitted. Any permitted paper materials or calculators will be stated in the assessment instructions. Device use during other class activities should support the activity and should not distract other participants.

4.4 Recordings and redistribution

No recording of the full class is planned. Slides and other public materials will be linked from the course website. Students may not independently record or redistribute class activity involving other students, student presentations, discussion, or project material except where authorized through a formal accommodation or by the instructor.

4.5 Research-paper seminars

Each student will give two regular research-paper presentations during the term, separate from the final project presentation. A regular presentation occupies approximately 30 minutes: ordinarily about 20 minutes of prepared presentation followed by facilitated questions and discussion. Papers are selected by the student and require written instructor approval more than two weeks before the presentation date. Students should normally send two to four candidate papers, with a short rationale for each, at least three weeks before the proposed date.

Papers should make a substantive contribution involving AI or machine learning in health or biomedicine and should support critical discussion of the clinical or operational problem, data, methods, evaluation, robustness, fairness, reproducibility, limitations, and path to use or non-use. Recent work is preferred unless an older paper is clearly foundational or directly relevant. The presentation schedule and exact paper choices will be developed with the class and posted in Teams. Full procedures, paper-finding guidance, and the presentation rubric are provided in the *Presentation Procedures, Expectations, and Paper-Finding Guide*.

4.6 Weekly reflections

Unless otherwise announced, each student submits an individual reflection through the designated Microsoft Teams assignment by 23h59 Atlantic Time on the Sunday following class. A normal reflection should be approximately 200–300 words and should identify:

1. one important claim, result, or idea from that week’s seminar papers or discussion;
2. one consequential strength in the evidence or study design;
3. one limitation, unresolved assumption, or alternative interpretation; and
4. one question or connection to the AIRO-Sim research project.

A student who presented that week submits a presenter reflection instead: identify one useful challenge or question raised during discussion, explain what you would revise in the presentation or interpretation, and state whether the discussion changed your view of the paper. Reflections are assessed for attentive, specific, constructive, and independent engagement, not for agreement with the presenter or instructor. No reflection is required during Fall Break; other exceptions will be announced in Teams.

5 AIRO-Sim research project

5.1 Project in one sentence

The class will build and study a shared, transparent cardiac-surgery risk system using fully synthetic AIRO-Sim data. Each student will lead a distinct research question, steward a shared artefact, and submit an individual research paper; the class will also produce a short shared system dossier and an integrated demonstration.

5.2 Note on clinical evidence

AIRO-Sim is fully synthetic and not for clinical use. It is related to the AIRO cardiac-surgery research program, but it contains no patient-level AIRO or Maritime Heart Centre data. Its field structure is schema-grounded and its temporal relationships are designed to be clinically plausible, while its frequencies, associations, notes, shifts, and benchmark performance are deliberate teaching choices.

Course findings may support claims about analyses conducted on AIRO-Sim. They do not establish clinical validity, safety, effectiveness, fairness, regulatory approval, institutional performance, documentation style, or patient benefit. Reports, presentations, repositories, interfaces, screenshots, logs, and exported cases must preserve this fact.

The frozen public release is [AIRO-Sim v1.0.0](#). The class may select postoperative delirium or a feasible sternal surgical-site-infection definition as its anchor outcome. A change in outcome normally creates a versioned derived outcome contract rather than a new source-data release.

5.3 Individual and shared responsibilities

Every student has two connected responsibilities:

- **Research lead:** define and pursue a claimable individual question; conduct decisive analyses or experiments; interpret limitations and alternative explanations; maintain reproducible code; and submit and defend an independently written paper.
- **Shared-artefact steward:** maintain at least one shared contract, dataset, script, model, schema, interface, test, figure, or decision record on which the class can rely, with acceptance checks, version history, known risks, and a backup reviewer.

The required shared backbone comprises a versioned outcome, cohort, snapshot, and split contract; a reproducible T0 feature and modelling pipeline with a transparent baseline; saved preprocessing and model artefacts with a documented inference interface; a synthetic-case viewer; a locally runnable package or container; a monitoring demonstration; a data statement and model/system card; and a concise shared dossier. T1, T2, NLP, multimodal, or richer service work must be earned after the T0 backbone and its acceptance checks are stable.

Leadership means accountability rather than exclusivity. Shared code, definitions, infrastructure, figures, and scientific ideas remain available to the class, but reuse must be attributed and each paper must make the student's independent contribution clear.

5.4 Repository practice and change control

The class will use one private shared GitHub repository during the term. Material shared tasks, decisions, and blockers should have a visible record. Substantive changes to shared interfaces should occur through

reviewed pull requests once dependencies begin to stabilize. A change that can alter another student's results must be reviewed, versioned, announced, and linked to the analyses that may need to be rerun.

Exploratory notebooks should remain separate from scripts or documented notebooks that generate final evidence. Environments, seeds, model versions, prompt or schema versions, and important sources of nondeterminism must be recorded. Never commit credentials, real patient information, copied clinical notes, instructor-only files, answer keys, hidden challenge data, or restricted material.

6 Assessment and grading

AIRO-Sim Research project (70%) A cumulative individual research paper and reproducibility package, shared-artefact stewardship, contribution to the integrated system and dossier, and an individual oral presentation and defence. The detailed project requirements are in the Course Project Handbook.

Two in-person quiz/essays (10%) Two individual assessments of approximately 60 minutes, worth 5% each. They emphasize explanation, evidence appraisal, interpretation of figures or results, and defensible study or system decisions rather than code execution.

Two research-paper presentations (10%) Two individual seminar presentations, worth 5% each. Evaluation considers problem framing, technical accuracy, evaluation and interpretation, independent critical analysis, communication and visual design, discussion leadership, and time management.

Weekly reflections (5%) Concise individual reflections submitted through Teams as described above. The component rewards specific engagement with evidence, limitations, alternative interpretations, discussion, and project connections.

Prepared participation, peer review, and interaction (5%) Constructive seminar participation, project stand-ups, review of shared work, response to feedback, and reliable contribution to the shared research. Quality, preparation, and usefulness matter more than speaking frequency.

6.1 Project assessment

The project constitutes 70% of the course grade. The percentages below are proportions of the *project mark*, not additional course-level percentages.

Project component	Weight	Standard
Literature review	10%	Extensive, current, accurate synthesis of relevant work and evidence.
Dataset analysis	15%	Judicious statistics and analysis of provenance, timing, outcome construction, missingness, event counts, leakage, and synthetic-evidence boundaries.
Methods	15%	Correct, reproducible technical description and well-justified experimental or analytical design.
Experiments and analysis	30%	Decisive evidence, valid procedures, uncertainty, error and robustness analysis, appropriate comparisons, and thoughtful interpretation.
Technical quality and reproducibility	15%	Implementation and testing; reproducible environments and commands; reliable stewardship, versioning, and integration of shared artefacts.
Written presentation	10%	Coherent research-paper argument, precise writing, effective figures and tables, and correct attribution.
Oral presentation and defence	5%	Clear, accurate, audience-aware presentation and demonstrated understanding of individual and shared work.

The literature review, dataset analysis, methods, and experiments-and-analysis components are evaluated at their checkpoint and again in the final paper; the higher mark for each component is retained. The remaining components are assessed from the final submission, stewardship record, presentation, and defence.

Each checkpoint revises and extends the same individual manuscript. There is no shared class paper. The shared dossier and integrated demonstration are required class artefacts, but there is no single class-wide mark that can hide unequal contributions. Grades are based on the student's own paper, analyses, code, decisions, stewardship, attribution, and demonstrated understanding.

Question registration, the question card, proposal, design review, evidence review, contract and risk-register updates, and other stages are required even when they do not carry separate percentage weights.

6.2 Quiz/essays

Quiz/essay 1 is planned for **Thursday, 29 October 2026**; quiz/essay 2 is planned for **Thursday, 19 November 2026**. Each will normally occupy approximately 60 minutes of class. The first emphasizes foundations, data and study design, and honest evaluation. Quiz/Essay 2 assesses material through the 5 November class, including temporal and textual systems, explainability, subgroup analysis, fairness, human factors, privacy, governance, reproducibility, inference interfaces, containers, and MLOps.

The assessments are completed individually and in person. No computers, phones, smart watches, generative-AI tools, or unapproved outside assistance are permitted. Any permitted paper materials will be stated in advance. A missed quiz without an approved arrangement receives zero. A student who misses a quiz because of a documented or otherwise eligible absence will be offered a make-up or alternate individual assessment; the instructor will determine its date and format.

6.3 Project deadlines and lateness

The checkpoint schedule coordinates shared dependencies as well as individual progress. Students should raise foreseeable problems privately and before the deadline whenever possible. A short extension that does not damage another student’s work or the shared critical path may be approved in writing.

Without an approved arrangement, a late checkpoint may lose the opportunity for timely formative feedback and may be assessed only when the corresponding component appears in the final paper. Missing a formative stage does not remove the underlying final requirement. The final paper, reproducibility package, presentation, and shared-system obligations remain due as stated unless an approved accommodation or absence arrangement applies.

6.4 Absences and alternate work

Students should notify the instructor privately as soon as possible and follow the applicable Dalhousie and Faculty of Computer Science absence process. Alternate arrangements depend on the missed requirement: a quiz may require a make-up assessment; a presentation may be rescheduled or replaced by an equivalent format; a reflection or participation activity may be excused or replaced; and a project deadline may be revised where the circumstances and shared dependencies permit.

An absence by one student does not automatically alter deadlines for shared artefacts on which others depend. Serious circumstances affecting individual work or the class’s ability to integrate the system should be raised early so that responsibilities, reviewers, and fallbacks can be adjusted fairly.

6.5 Graduate grade scale and minimum grade requirement

The Faculty of Graduate Studies grading scale is used:

A+	A	A-	B+	B	B-	F
90–100	85–89	80–84	77–79	73–76	70–72	0–69

A minimum grade of B- is required to pass this graduate course. A lower grade is recorded as F under Faculty of Graduate Studies regulations.

7 Use of automated tools and academic integrity

7.1 Automated and generative tools

Automated tools, including generative-AI systems, may be used for learning, coding assistance, interpretation of error messages, brainstorming tests, editing, and approved research methods, provided that:

- the use is consistent with the instructions for the particular assessment;
- the student independently verifies outputs, citations, code, calculations, and claims;
- borrowed or generated text, code, figures, prompts, procedures, and ideas are attributed or disclosed as required;
- the student preserves enough of the analytical record to explain how each reported result was obtained; and
- the student can explain and defend every submitted claim, result, design choice, and implementation decision.

Automated output may not substitute for the student’s independent reading, interpretation, synthesis, or scientific judgement. A tool-generated weekly reflection or paper critique does not satisfy the purpose of that assessment. Material automated contributions to the project must be disclosed in `contributions.txt` and, when they affect the scientific work, in the methods or acknowledgements.

Generative AI and other outside assistance are prohibited during the in-person quiz/essays. Students must not upload another student’s work, grades, private information, credentials, instructor-only material, answer keys, unreleased challenge data, or restricted course material to a third-party service. Fabricated citations, data, experiments, outputs, or statistical results are academic-integrity violations.

NLP and LLM work within an approved project is optional. Any such study must record the model or provider, version or access date, prompt or extraction schema, data sent, outputs retained, and reproducibility limitations. Permission to use fully synthetic AIRO-Sim notes does not imply that the same service or workflow would be acceptable for real clinical data. Proofreading and local suggestions for clarity are permitted. Generating or substantively rewriting passages that carry the student’s analysis, synthesis, interpretation, or argument may not replace the student’s independent authorship and must be disclosed when it constitutes a material contribution.

7.2 Collaboration and independent authorship

Collaboration is expected on the shared system and permitted where it genuinely serves more than one research question. Students may discuss concepts, review pull requests, share designated class artefacts, and reuse shared code, definitions, figures, or infrastructure with accurate attribution.

The two quiz/essays, two regular paper presentations, weekly reflections, and final individual paper are individual work. Final papers must be independently authored; students may not copy or lightly paraphrase another person’s or tool’s prose, even when projects share a setting or artefact. Each student must distinguish inherited shared work from their own contribution and must understand every part of the submitted analysis that supports their claims.

7.3 Academic integrity

Academic work at Dalhousie is governed by honesty, trust, fairness, responsibility, and respect. Sources, reused code, data, figures, prompts, and permitted tools must be acknowledged as appropriate. Shared artefacts must identify contributors and versions. Suspected academic-integrity violations are handled through Dalhousie’s academic-integrity procedures. See [Dalhousie Academic Integrity](#).

8 Responsible research and public release

8.1 Synthetic data, privacy, and safety

Do not add real patient information, real clinician names, copied clinical notes, credentials, or restricted material to the course project. Student-release and instructor-only materials must remain separate. Student work must not depend on latent generator truth, anomaly keys, hidden annotations, benchmark predictions, or an unreleased final challenge set.

No paid model, API, cloud account, or other commercial service is required. Every required deliverable must have a no-cost local or institutionally available path. The course does not reimburse optional service costs. Clinician input may be useful if available, but it is optional, not guaranteed, and must not be described as clinical validation.

8.2 Public release after the course

Selected work may be prepared for public release after grading, subject to instructor review, university policy, appropriate licences, accurate attribution, and removal of grades, peer-review comments, credentials, instructor-only materials, answer keys, and unreleased challenge data. A student's grade does not depend on agreeing to publish an individual paper or code contribution.

Any public release must identify the data and software versions, distinguish AIRO-Sim from AIRO, state that the work was produced in DGIN 5401 using fully synthetic data, identify contributors accurately, and preserve the non-clinical-use boundary.

9 Resources and tools

9.1 Required computing environment

Students require access to a free Python environment, Jupyter or an equivalent development environment, Git, and the open-source libraries used in the course and project, including NumPy, pandas, Matplotlib, and scikit-learn. Particular projects may use additional open-source packages. The shared system will also use a locally runnable package or container; Docker and Podman should both be considered where feasible.

Installation and AIRO-Sim verification instructions will be provided on the course website and in the release. No paid software, subscription, hosted API, or commercial account is required or recommended for required work. Students who cannot reliably access a suitable computer should contact the instructor privately as early as possible so that a workable arrangement can be discussed.

9.2 Course and project resources

The course website will link the current syllabus, Course Project Handbook, AIRO-Sim Technical Brief, lecture slides, seminar papers, presentation procedures, and dataset tutorial. The frozen public dataset release is available at:

<https://github.com/Dalhousie-A2I2/AIRO-Sim/releases/tag/airo-sim-v1.0.0>

The following freely available references may be useful for background; there is no required textbook purchase:

- Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor, *An Introduction to Statistical Learning*.
- Christoph Molnar, *Interpretable Machine Learning*.
- Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong, *Mathematics for Machine Learning*.
- Kevin P. Murphy, *Probabilistic Machine Learning*.

Required and recommended scientific papers will be selected throughout the term and linked from the course website or Teams.

10 Tentative schedule

The sequence below provides the course's planned intellectual progression. Student-selected seminar papers will be distributed across the term and may cause local adjustments in emphasis or order. Presentation dates and papers are confirmed individually with the instructor and posted in Teams. Changes will be announced through Teams and Dalhousie email and reflected on the course website.

- Week 1: 10 Sep** Course orientation; what would have to be true for machine learning to improve healthcare; intended use; the AIRO-Sim research; project and seminar procedures. Project exploration begins.
- Week 2: 17 Sep** Clinical data science foundations: how care becomes data; units, identifiers, provenance, timing, labels, missingness, leakage, synthetic-data limitations, and a first data audit. Ranked interests are due 16 September; anchor outcome and provisional leads are posted by 18 September.
- Week 3: 24 Sep** Honest models and meaningful evaluation: study cards, transparent baselines, event budgets, splitting, discrimination, calibration, scoring rules, thresholds, workload, decision consequences, uncertainty, and claim boundaries. Question cards are due 21 September and confirmed by 22 September; proposals are due 24 September. The last day to add or change full-term fall courses is 22 September.
- Week 4: 1 Oct** Temporal and dynamic prediction: time origins, prediction landmarks, fixed-horizon and survival questions, censoring, competing events, updating, and time-dependent evaluation. Student paper seminars begin or continue as scheduled.
- Week 5: 8 Oct** Clinical text, NLP, multimodal methods, and foundation models: representation, extraction, temporal and template leakage, factuality, evidence attribution, and evaluation beyond fluency. Literature-synthesis checkpoint due. The last day to drop a full-term fall course without a W is 7 October.
- Week 6: 15 Oct** Explainability, uncertainty, subgroup analysis, and fairness: denominators, instability, correlated features, measurement and selection, incompatible criteria, and limits of synthetic subgroup evidence.
- Week 7: 22 Oct** Human factors, workflow, interfaces, and clinical decision support: users, tasks, risk communication, workload, automation bias, actionability, scenario evaluation, and the case-viewer contract. Dataset-analysis checkpoint due.
- Week 8: 29 Oct** In-person quiz/essay 1 during approximately the first 60 minutes; privacy, safety, governance, accountability, and design review. Rescope requests are due 26 October; accepted scopes and major contracts lock on 29 October.
- Week 9: 5 Nov** Reproducibility, inference interfaces, containers, testing, versioning, and MLOps for a deployment-oriented teaching prototype. Methods checkpoint due. The last day to drop a full-term fall course with a W is 5 November.
- Week 10: 12 Nov** Fall Break; no class.
- Week 11: 19 Nov** In-person quiz/essay 2 during approximately the first 60 minutes; monitoring, shift, delayed outcomes, alert interpretation, incident response, recalibration, and rollback. Project evidence review.
- Week 12: 26 Nov** Integration, error autopsy, robustness, final evidence clinic, system documentation, and demonstration rehearsal. Experiments-and-analysis checkpoint due; shared artefacts freeze for final integration.
- Week 13: 3 Dec** Individual project presentations and oral defence; integrated AIRO-Sim system demonstration; collective synthesis and next evidence stages.
- Final submission: 9 Dec** Final individual paper and reproducibility package; shared dossier, tagged repository state, runnable package or container, case viewer, monitoring demonstration, model/system card, data statement, and contribution record. There is no Registrar-scheduled final examination.

11 Accessibility, respect, and support

11.1 Accessibility and accommodation

Dalhousie is committed to an accessible learning environment. Formal accommodations are managed by the Student Accessibility Centre for Halifax courses. Students who encounter a barrier related to a protected characteristic should contact the Centre as early as possible and should communicate approved course arrangements to the instructor in time for implementation. See [Accessibility and Accommodations](#).

Because the two quiz/essays and many seminar and project activities occur in class, students who require testing or participation accommodations should arrange them early and verify the applicable procedures and deadlines.

11.2 Culture of respect and inclusivity

Substantive disagreement and critical discussion are welcome. Harassment, discriminatory conduct, personal attacks, and disclosure of another student's private information are not. Everyone in the classroom, Teams space, and shared repository is expected to contribute to a respectful, accessible, and inclusive learning environment. Critique claims, evidence, methods, and systems rather than people. Faculty-specific information and support are available through [CoReCS: Culture of Respect in the Faculty of Computer Science](#).

11.3 Faculty of Computer Science Learning Centre

The Faculty of Computer Science Learning Centre provides support with computer-science and mathematics course work. Current hours and access information are available from the [FCS Learning Centre website](#).

11.4 Responsible computing

Use of Faculty and University computing resources must comply with the [Dalhousie Acceptable Use Policy](#) and the [Faculty of Computer Science Responsible Computing Policy](#).

12 Course Syllabus Appendix

Current University statements, policies, guidelines, and student supports are maintained in the [Course Syllabus Appendix](#) (Dalhousie login required). Students must consult it together with this syllabus.