

DGIN 5401: Presentation and Reflection Guide

Operationalized Machine Learning in Healthcare

Fall 2026 – Faculty of Computer Science, Dalhousie University

1 Requirements and preparation

A seminar should help a mixed clinical–technical audience understand a paper’s contribution, methods, decisive evidence, and limitations. Explain what it establishes and what it does not; do more than summarize the authors’ conclusions.

Assessment at a glance. Each student gives **two regular paper presentations, worth 5% each**. Individual weekly reflections are worth **5% in total**. The final project presentation and defence is separate.

Each regular slot is approximately **30 minutes**: about **20 minutes presenting**, **8 minutes discussing**, and **2 minutes for transition**. The balance may vary with the paper and discussion.

Preparation timeline

When	What to do
At least 21 days before, where possible	Send two to four candidate papers, with two or three sentences each on course fit, what the class could learn, and any scope, quality, or access concern.
More than 14 days before	Obtain written agreement on the final paper through Teams or email. The paper is not scheduled until approved.
About 7 days before	Share a draft deck for brief feedback: optional, but strongly recommended, especially for the first talk.
About 3 days before	Post two proposed discussion questions in the designated Teams space: optional, but strongly recommended.
Presentation day	Arrive at least five minutes early. Test the audiovisual setup; bring editable slides and a PDF backup. Prepare an offline fallback for demonstrations.

Raise scheduling problems early. Missing the approval window may require a different paper or date and may reduce the mark when it prevents adequate preparation or coordination.

2 Choosing and appraising a paper

Choose work in which AI or machine learning is substantive to a method, system, or deployment question in health, biomedicine, public health, or healthcare operations. Clinical data or a clinical trial alone does not necessarily fit.

Prefer peer-reviewed work from a credible venue, usually from approximately the last three years. Older foundational papers and carefully vetted preprints may be suitable, with instructor approval. Check a preprint's subsequent publication status.

Look for enough detail to support independent analysis; negative, null, or mixed findings are welcome. Across your two talks, consider balancing an application paper with a methods, evaluation, implementation, or governance paper. We aim to achieve breadth and depth.

Finding candidates

Starting points include Machine Learning for Healthcare (MLHC), the Conference on Health, Inference, and Learning (CHIL), *The Lancet Digital Health*, and *npj Digital Medicine*. arXiv and medRxiv carry preprints. A venue is not a guarantee of study quality.

When searching (e.g., through Google Scholar), combine a domain, method, and evaluation concern: for example, “sepsis early warning calibration external validation.” Use year filters and follow references and later citing papers. Look for validation, critique, correction, retraction, implementation, or negative evidence. Inspect associated code, data, and protocols: availability alone does not establish reproducibility.

Questions to guide your reading

Use these questions to select and appraise the paper. Focus on what matters to its claims; critical analysis should run throughout the talk.

1. **Problem and intended use.** What need is being addressed, who would act on the result, and what could go wrong?
2. **Data and study design.** Who is represented or excluded? Identify sources, sample sizes, sites, dates, cohort rules, labels, missingness, and governance. For prediction studies, establish the unit of analysis, prediction time, target, horizon, and evaluation population.
3. **Methods and comparisons.** How is the model trained, and what assumptions matter? Are simple baselines, comparators, and ablations appropriate and treated fairly? What can actually be reproduced?
4. **Evaluation and interpretation.** Which result is decisive? Examine validation design, calibration, uncertainty, errors, and practical consequences where relevant. Could a good aggregate score conceal low precision, excessive workload, or harmful threshold choices?
5. **Validity and harm.** Could patient, site, temporal, documentation, template, or preprocessing leakage explain the result? Consider shortcuts, distribution shift, subgroup evidence and its statistical power, fairness, privacy, and safety.
6. **Conclusions and next evidence.** Does the evaluation match the intended use? What alternative explanation or missing evidence most limits the claim, and what analysis would change your interpretation?

3 Building and delivering the presentation

A suggested 20-minute structure

This is a planning aid only. Adapt the following to the paper and your own better judgment, and consult the instructor as needed. Naturally, a methods paper is more likely to have a focus on the Methods section. Practise with a timer, leaving room for discussion, for example.

Part	Minutes	Main task
Title and hook	Up to 1	Explain why the question matters to the relevant stakeholders.
Problem and workflow	2–3	Establish the unmet need, intended use, current workflow, and plausible harm.
Data	2–3	Explain the population, sources, labels, timing, missingness, and important selection processes.
Methods	4–5	Explain the approach, training, assumptions, and most important baselines or comparisons.
Results	3–4	Interpret the decisive figures or tables, including uncertainty and relevant validation evidence.
Limitations, fairness, and safety	2–3	Examine failure modes, subgroup effects, privacy, deployment risks, and monitoring needs.
Takeaways and discussion	1–2	Give two or three defensible conclusions and introduce the discussion.

Make the evidence understandable

Do not assume deep technical or clinical knowledge. Define specialized terms and explain the measures needed to assess the claim. Distinguish development, validation, test, temporal, site, and prospective evaluations. Separate what the study shows from what its authors infer or propose.

Use readable, high-contrast slides, ordinarily with text of at least **24 points**. Give each slide one main purpose. Explain important visual patterns aloud; use labels or symbols rather than colour alone. Avoid dense page screenshots or walls of text. When cropping or redrawing a figure, preserve its meaning, relevant axes, and denominators, and identify the source and modifications. Speak to the audience rather than reading the slides.

Lead the discussion

Prepare **at least two questions** that invite analysis. They should draw on the paper and course concepts while allowing more than one defensible position. For example:

- Which design choice most limits the claim, and how would you test its effect?
- What evidence would be needed before moving from retrospective evaluation to workflow testing?
- Who bears the cost of false positives, false negatives, or added workload?

Invite alternative interpretations, ask follow-up questions, and clarify disagreements. You need not know every answer: identify unresolved questions and explain how they could be investigated.

4 Assessment and academic conduct

Presentation rubric

Each regular presentation is marked out of 100. The two presentation percentages are equally weighted, with each contributing **5% of the course grade**. The following is the default rubric only.

Criterion	Points	What is assessed
Problem framing and significance	15	A clear, consequential clinical, public-health, scientific, or operational need, connected to an intended use or analytical purpose.
Technical accuracy	20	Correct, concise explanation of methods and assumptions at an appropriate level for the audience.
Evaluation and interpretation	15	Appropriate interpretation of measures, calibration, validation, uncertainty, errors, and practical consequences.
Critical analysis	20	Independent examination of limitations, fairness, safety, external validity, trade-offs, alternative explanations, and claim boundaries.
Communication and visuals	15	Structure, pacing, slide design, accessibility, citation practice, and oral explanation that support understanding.
Discussion leadership	10	Substantive questions, respectful engagement, and development of alternative interpretations.
Time management	5	A prepared talk close to 20 minutes, with time allocated according to the importance of the evidence.
Total	100	

The rubric is applied with judgement appropriate to the paper: a methods paper and an implementation study may need different balances of clinical, mathematical, and operational detail.

Attribution, tools, and conduct

Provide lawful access. Confirm that the class can lawfully access the paper.

Credit sources. Cite borrowed figures, tables, text, ideas, code, and schematics. Caption figures and tables and cite their sources on the relevant slide. Include a DOI or stable source link in the deck or references, and identify modifications to source figures.

Own and verify your work. Your slides, explanation, and critical analysis must accurately represent your contribution. Never fabricate citations, experiments, results, quotations, or access to unavailable evidence. Follow the syllabus policy on automated tools: disclose material use, verify the output independently, and do not let it replace your reading, judgement, or ability to defend the work.

Respect others. Critique ideas, assumptions, methods, and evidence rather than people. Do not independently record or redistribute another student's presentation or class discussion unless authorized through an accommodation or by the instructor.

5 Weekly reflections

Reflections can connect seminar learning to the AIRO-Sim Research project or another course topic. Show your judgement and insight rather than reciting elements of the paper. This individual component is worth **5% of the course grade**.

Submit through the designated **Microsoft Teams assignment** by **23h59 Atlantic Time on the Sunday following class**, unless otherwise announced. A normal reflection should be approximately **200–300 words**.

Submit **one coherent reflection for the week**, not one per paper. No reflection is required during **Fall Break**; other exceptions will be announced in Teams. For an eligible absence or approved accommodation, follow the syllabus procedure and contact the instructor privately.

Normal weekly reflection

Address the following, using specific evidence from that week’s seminar papers or discussion:

1. **Claim or idea:** identify one important claim, result, or idea.
2. **Strength:** explain one consequential strength in the evidence, study design, analysis, or implementation.
3. **Concern:** examine one limitation, unresolved assumption, alternative explanation, or concern about the limits of the claim.
4. **Connection:** raise one question, connection, or implication for the AIRO-Sim Research project or another course topic.

Be concrete. “The paper had a large dataset but more work is needed” does not identify what matters. Explain which dataset property affects which claim, and what analysis or evidence would resolve the concern.

When you are the presenter

Submit a **presenter reflection instead of the normal reflection**. Rather than reviewing how well the talk went or restating its conclusions, address:

1. one useful challenge, question, or alternative interpretation raised during discussion;
2. what you would revise in your presentation, evidence selection, or interpretation; and
3. whether the discussion changed your view of the paper, and why.

How reflections are assessed

Reflections are assessed for attentive engagement, specificity and accuracy, constructive critical reasoning, a meaningful question or connection, and clear, concise, independently written prose. Agreement with the presenter or instructor is **not** required.

Acknowledge external sources used. The tool-use requirements on page 4 apply.