AI-ASSISTED INVESTIGATIVE INTERVIEW TRAINING

A low-cost proof of concept for Calgary Police Service

Training first · measurable results · operational research later



Prepared for preliminary review · August 2026

EXECUTIVE PROPOSITION

Start with training. Earn the right to discuss operations.

The lowest-risk path is to test whether AI actually makes interviewers better.

PHASE I

Use AI as a training and coaching tool in simulated interviews.

No suspects

Volunteer actors and synthetic case files.

No live evidence

No confidential operational data.

Low cost

≈ C$2,000 incremental.

The immediate question is not whether AI can detect deception.

Can AI-assisted analysis measurably improve the training and performance of investigative interviewers?

WHY TRAINING FIRST

De-risk the idea before attaching it to a real investigation.

Legal exposure

No accused person; no live investigative jeopardy; no need to decide covert-analysis authority at the outset.

Privacy

Synthetic data and consenting participants dramatically reduce privacy and data-handling risk.

Scientific control

Known scenario facts allow objective scoring of what the interviewer and AI did or did not detect.

Cost

Existing rooms and equipment can be reused; the experiment remains intentionally inexpensive.

Institutional fit

Training is already a core CPS function and a 2026 priority.

Training creates evidence about the technology before CPS is asked to trust the technology.

That reverses the burden of persuasion: results first, ambition second.

CPS FIT

The proposal aligns with work CPS already says it is doing in 2026.

Training & development

The Annual Policing Plan calls for appropriate training, education, development and equipment, and for stronger investigative capacity.

Training feedback

CPS plans to implement a feedback process for training at the Chief Crowfoot Learning Centre to drive ongoing improvement.

AI governance

The plan identifies AI governance and standards, enhanced digital evidence management, and technology to improve investigative efficiency.

Human verification

A 2026 KPI specifically measures the accuracy of AI-assisted analysis compared with human verification.

Evidence processing

Another KPI tracks eligible video files processed using AI tools.

This project would be a narrow, controlled training experiment within an existing modernization direction.

Source: Calgary Police Service 2026 Annual Policing Plan (Calgary Police Commission, Nov. 2025).

THE TRAINING PROBLEM

A good interviewer is already doing several jobs at once.

Listen

Track meaning, wording and changes in account.

Observe

Watch behaviour without over-interpreting it.

Recall

Remember case facts, prior answers and chronology.

Compare

Notice contradictions and evidentiary conflicts.

Decide

Choose the next productive question.

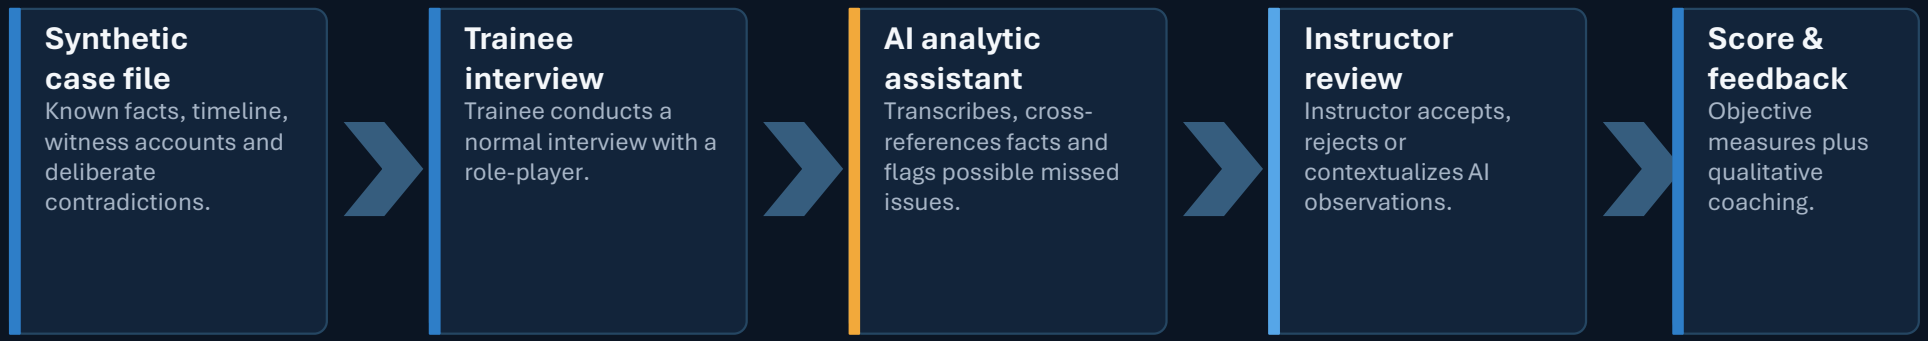
Document

Preserve a reliable record for later review.

AI is most credible where it reduces cognitive load, not where it claims to read minds.

PHASE I CONCEPT

A synthetic scenario creates a closed-loop training environment.



The AI sees only what the exercise designer authorizes it to see.

No hidden database. No real case material. No operational decision-making.

LEVEL 1 CAPABILITY

What the Phase I AI assistant actually does.

Live transcript

Create a searchable, time-linked record.

Timeline

Maintain who, what, where and when.

Evidence match

Compare statements with scenario facts.

Contradictions

Flag possible internal or evidentiary inconsistencies.

Unanswered questions

Identify where the response did not resolve the question.

Follow-up ideas

Suggest neutral clarification questions for instructor or trainee review.

HARD BOUNDARY

What the system must not claim.

The training-first model is deliberately conservative.

- “This person is lying.”
- “This answer proves guilt.”
- “This facial movement means deception.”
- “The AI score should determine credibility.”

AI OUTPUT

“Possible inconsistency. The subject’s answer conflicts with the scenario timeline. Consider clarifying the sequence.”

Observable behaviour may be discussed as a training cue only. Anxiety, trauma, fatigue, disability, culture and ordinary interview stress can all alter behaviour.

EXAMPLE

A training scenario with known ground truth.

SCENARIO FACTS

- Vehicle recorded downtown at 9:17 p.m.
- Parking receipt issued at 9:21 p.m.
- Subject claims to have remained home until 10:00 p.m.

AI TRAINING OBSERVATION

Possible missed issue

At 14:22 the subject said he remained home until 10:00 p.m. The scenario file contains a vehicle-location event at 9:17 p.m. No clarifying question has yet been asked.

Suggested neutral follow-up

“Help me understand where your vehicle was between 9:00 and 9:30.”

TRAINING FEEDBACK

The strongest early use may be after the interview, not during it.



1 · Replay

Link every observation to the transcript and video timestamp.

2 · Explain

Show the evidence conflict or missed line of inquiry.

3 · Coach

Instructor decides whether the observation is meaningful and discusses alternatives.

4 · Measure

Record performance over repeated scenarios.

HUMAN IN COMMAND

The comparison is instructor vs. instructor + AI.

Not instructor vs. machine.



Value exists only if the AI adds reliable observations that improve coaching without creating excessive false alerts, distraction or misplaced confidence.

EXPERIMENT

A simple comparative design can test the hypothesis.

GROUP A · STANDARD TRAINING

Trainee interviews a role-player. Instructor observes and scores using the existing rubric or a study-specific rubric.

No AI assistance.

GROUP B · AI-ASSISTED

Same or equivalent scenario. AI generates transcript-linked observations. Instructor decides what, if anything, to use.

AI never grades autonomously.

COMPARE RESULTS

Contradictions detected, follow-up quality, critical facts used, unanswered issues, false alerts, instructor time and trainee improvement.

Repeat across several scenarios and experience levels before drawing conclusions.

MEASUREMENT

Define success before the experiment begins.

Contradictions recognized	% of planted material contradictions detected
Relevant follow-ups	Count / instructor quality rating
Critical facts integrated	% of designated scenario facts used appropriately
Unanswered issues	Number of material issues left unresolved
AI false alerts	Rate and severity of unsupported flags
Instructor workload	Minutes required for review and debrief
Skill transfer	Performance change on later unaided scenarios

COST

Phase I can remain deliberately small.

≈ C$2,000

working proof-of-concept figure

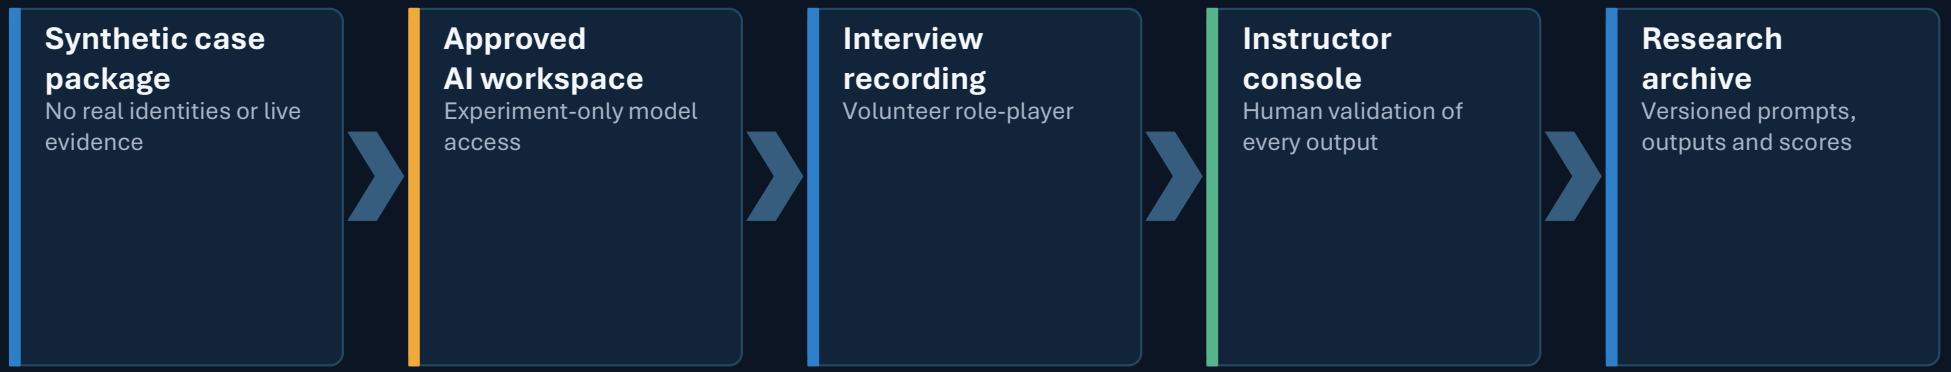
Existing CPS resources should be reused wherever practical. Labour and internal evaluation time are not included in the equipment estimate.

Existing room / computer / recording	$0
Approved experimental AI/model access	$0–300
4K webcam if needed	$150–300
Secondary camera / repurposed device	$0–200
USB microphone if needed	$100–200
LED lighting / ambient sensor	$150–350
Mounts, cables, calibration	$175–350
External SSD / display if needed	$325–650

Estimated incremental range: C$900–$2,350

TECHNICAL BOUNDARY

Keep Phase I isolated from operational systems.



Production-grade security architecture is a later problem only if the research justifies creating one.

SAFEGUARDS

Phase I privacy rules should be simpler than the technology.

- ✓ Synthetic scenario facts only
- ✓ Consenting participants / role-players
- ✓ No live police database connections
- ✓ No reuse of recordings for model training
- ✓ Access limited to the project team
- ✓ Defined retention and deletion schedule
- ✓ Complete logs of model, prompts and outputs
- ✓ Privacy/legal review before any expansion

FAIRNESS

Training data can reveal bias before operational stakes exist.

Scenario diversity

Vary language style, age, cultural background, disability, trauma presentation and communication patterns where ethically appropriate.

Measure disparity

Track whether the system produces materially different false-alert rates across participant groups or scenario types.

Human override

No score or flag should compel an instructor conclusion. Every output remains contestable and reviewable.

Behavioural “deception” claims should remain outside Phase I.

The first experiment does not need them to test the strongest part of the idea.

GO / NO-GO

Predefine reasons to stop.

STOP if...

- False alerts are frequent or distracting
- Instructor workload increases without benefit
- Trainee performance does not improve
- Outputs cannot be reproduced or audited
- Bias or privacy risks cannot be controlled

MODIFY if...

- Benefits appear only in certain scenarios
- Real-time alerts distract but post-interview review helps
- The model performs well only with narrower tasks
- Better prompts or data structure materially improve results

PROCEED if...

- Material improvement is repeatable
- False-alert burden is acceptable
- Instructor feedback is positive
- Controls are auditable
- CPS determines the next phase has a legitimate purpose

ROADMAP

Four stages, with a decision gate between each one.

1

Training proof of concept

Synthetic scenarios; post-interview AI coaching; ≈ C$2,000 equipment/services.

2

Instructor decision support

Test instructor + AI against instructor-only review; refine metrics and interface.

3

Retrospective research

If legally approved, analyze selected completed interviews after the fact. No live influence.

4

Prospective operational research

Only if justified: tightly governed Level 1 assistance in selected real interviews.

Each phase requires a fresh decision. No phase implies approval of the next.

POSSIBLE PHASE III

Retrospective research: test against what investigators later learned.

Only after Phase I/II demonstrate value and CPS approves the legal/privacy framework.

Completed interview

Use a lawfully available historical recording selected under an approved research protocol.

Blind AI review

Ask the system to identify Level 1 contradictions, chronology gaps and missed follow-ups without influencing the original interview.

Compare outcomes

Did the AI notice issues later corroborated by the investigation? How many unsupported alerts did it create?

Retrospective research tests operational relevance without changing an investigation in real time.

POSSIBLE PHASE IV

Live operational assistance is a future research question, not the opening request.

Potential Level 1 use

Live transcript, chronology, evidence comparison, contradiction alerts and suggested neutral follow-up questions.

Human authority

Investigator remains responsible for every question and every investigative decision. AI output is advisory.

Separate behavioural research

Facial, voice or physiological inference should require independent scientific validation and separate legal approval.

A successful training tool does not automatically become an operational tool.

LEGAL POSTURE

Defer the hardest legal questions until there is evidence the technology is worth them.

Phase I–II

Synthetic training: consent, privacy, records, procurement, responsible AI governance.

Phase III

Historical interviews: lawful access, retention, disclosure, privacy and research authorization.

Phase IV

Live interviews: Charter, voluntariness, counsel, right to silence, disclosure, privacy, and whether any technique requires prior judicial authorization.

The proposal does not ask CPS to decide Criminal Code s. 487.01 or covert-analysis questions in order to approve a synthetic training study.

SUGGESTED REVIEW PATH

A cross-functional preliminary review would be more useful than a single-unit decision.

Training / CCLC

Does this solve a real training problem?

Investigative support

Are the Level 1 functions operationally meaningful?

IT / AI governance

Can the experiment be isolated, logged and controlled?

Legal / privacy

Are the proposed Phase I safeguards adequate?

Evaluation / strategy

Can CPS design a fair test and measure results?

Possible first home: Chief Crowfoot Learning Centre, with referral to the appropriate investigative, technology, legal/privacy and evaluation personnel.

The Chief Crowfoot Learning Centre is listed in the CPS public directory and 2026 Annual Policing Plan.

DECISION REQUESTED

The immediate ask is intentionally modest.

1. Refer the concept to the appropriate CPS training/research group for preliminary assessment.

2. If reviewers see merit, authorize design of a synthetic, low-cost Phase I proof of concept.

3. Use measured results to decide whether the project stops, changes, or advances.

No request is made for operational deployment.

APPENDIX A

Proof-of-concept cost estimate.

Existing interview-room computer	$0
Existing interview-room video/audio recording	$0
Approved AI/model access	$0–$300
Additional 4K webcam, if required	$150–$300
Secondary camera / repurposed smartphone	$0–$200
Quality USB microphone, if required	$100–$200
Adjustable LED lighting	$100–$200
Ambient-light sensor	$50–$150
Camera mounts, tripod, cables	$75–$150
1–2 TB external SSD	$125–$250
Interviewer/instructor display, if needed	$200–$400
Misc. setup/calibration materials	$100–$200

Estimated total incremental cost: C$900–$2,350 · Recommended working figure: approximately C$2,000

APPENDIX B

Illustrative trainee scorecard.

Measure	Scenario target	Trainee result	AI contribution	Instructor note
Material contradictions	4	3	1 additional	Useful
Critical facts used	8	6	2 prompted	Moderate
Unanswered issues	≤2	3	2 flagged	Useful
Unsupported AI alerts	0	—	1	Reject / document
Follow-up quality	Instructor 1–5	4	—	Strong
Debrief time	Baseline	32 min	+6 min	Acceptable?

Illustrative only. CPS instructors should define the actual rubric and thresholds.

APPENDIX C

Principal public sources and framing notes.

Calgary Police Service — Chief and Executive

Confirms Chief Katie McLellan and current executive information.

Calgary Police Commission — “Katie McLellan to continue as police chief until mid-2028” (25 Mar. 2026)

Confirms continuation as Chief through mid-2028.

Calgary Police Service 2026 Annual Policing Plan

Training, investigative capacity, AI governance/standards, AI-assisted analysis vs. human verification, and AI video processing.

Calgary Police Service public directory

Chief Crowfoot Learning Centre and current CPS contact information.

Proposal note

All costs, workflows, scorecards and phased research designs in this presentation are illustrative proposals, not CPS policy or procurement estimates.