



The Human Factor
HUMAN-CENTRED BY DESIGN

2026 INDIGENOUS PUBLIC SAFETY CONFERENCE · SASKATOON · MONDAY, OCTOBER 26

AI-enabled emergency preparedness and decision support for Indigenous public safety

Human judgment is necessary, and not automatic.

Scott Ramey

I help people build safer systems for humans.

MSCHQ · PhD Candidate, Systems Design Engineering, University of Waterloo
Founder, The Human Factor (thehumanfactor.ca)
Former Assistant Fire Chief/Deputy Paramedic Chief, CCP



Slides, the questions card
and the guide

thehumanfactor.ca/events/ipsc-2026

ACKNOWLEDGEMENT

Treaty 6 Territory and the Homeland of the Métis

As we gather here today, we acknowledge we are on Treaty 6 Territory and the Homeland of the Métis. We pay our respect to the First Nations and Métis ancestors of this place and reaffirm our relationship with one another.

Thank you to the National Indigenous Fire Safety Council

and to the people in this room who have fostered the land and community for centuries.

THE HOUR

One idea, six connections

AI can help communities prepare, respond, evacuate and recover, but only if communities own the data, the tools work in their conditions, and the human judgment stays theirs.

You'll leave with six questions to ask any vendor, three to ask yourselves, and a way to tell a safeguard from a sentence in a policy.

1

What AI is

9 min

2

The reality here

4 min

3

People and AI

13 min

4

**What the
research
shows**

4 min

5

**Building your
system**

12 min

6

**What I've
learned, and an
offer to you.**

6 min

Warm cards are about this room. Blue cards are the core I bring everywhere.



The organization owns what its AI says

AIRLINE CHAT ASSISTANT · NOVEMBER 2022

My grandmother just died. Is there a bereavement fare?

Yes. Book now, and apply for the bereavement rate within 90 days of travel.

The real policy allowed no refund after travel. The chatbot made that up.

In effect, the tribunal said, the airline was arguing its chatbot was "a separate legal entity that is responsible for its own actions."

"A remarkable submission."

The tribunal held the airline responsible and ordered it to pay.

Whoever deploys the AI owns what it says: an agency, a vendor, a band office.



A 70-year-old idea, a 3-year-old tool

1955

The Dartmouth proposal names "artificial intelligence"

1966

ELIZA: a few hundred lines of pattern matching, and people confided in it

2017

The transformer: the architecture behind today's chatbots


2022

ChatGPT, November 30. An estimated 100 million monthly users two months later

What changed in 2022 wasn't the idea. Everyone got it at once, free, in plain English.



The machines were already learning



Lane departure warning: a camera reads the lines. Today's systems use trained recognizers.

1958

Neural networks. The perceptron: a few hundred weights.

1995

Support vector machines. Spam filters, fraud alerts.

2001

Random forests. Risk scores, weather, your credit card.

2020

175 billion weights. Same idea. Internet-scale text, GPUs, words as vectors.

A chatbot is the perceptron's grandchild, scaled a billion times.



"AI" is four different things, and each fails differently



Recognize

Smoke and fire detection cameras. Faces. Thermal imagery.

Fails by misidentifying



Predict

Wildfire spread. Flood forecasting. Who is at risk.

Fails by learning yesterday's patterns



Generate

Drafting an emergency plan. A public notice, and its translation.

Fails by making things up, fluently



Triage

Sorting calls and requests for help during an evacuation.

Fails by deciding who waits

First question for any vendor: which kind is this?

Ramey, Masoumi-Ravandi, Rehman and Cao, systematic review (in preparation): systems ranged from triage to forecasting to drone imagery; few were fully operational.



What a chatbot actually does

The river is expected to crest on _____



Illustration. It picks a likely next word. It has not looked at the river.

Generative AI models “generate content that is statistically likely in response to a prompt, for example, by predicting the next word in a sentence.”

Government of Canada, Guide on the use of generative AI

Very sophisticated autocomplete.

No built-in sense of true or false. Not a search engine. It doesn’t look anything up unless it’s connected to something that does.



Why it sounds so sure



Fluent is not accurate

It writes the wrong answer in the same confident voice as the right one. The industry word is "hallucination."



Tuned to please

Assistants "consistently exhibit sycophancy." In 2025 one update was withdrawn for being "overly flattering or agreeable."



Professionals got caught

Lawyers who cited invented cases: sanctioned in New York, 2023; ordered to pay costs in Vancouver, 2024.

The tone that feels most helpful is the one to be most careful with.



2 · THE REALITY HERE

Any tool worth building has to be built for this

1,300+

emergencies in First Nations communities over 13 years

580+

evacuations · 130,000+ people affected

90

First Nations wildfire evacuations in 2023–24, the most on record

27,647

people evacuated from Saskatchewan First Nations since 2016, of 143,724 nationally

Auditor General of Canada 2022 · Indigenous Services Canada evacuation statistics to March 2026 (not final) · Christianson et al. 2024 · Photo: Jason Woodhead, Fort McMurray, May 1 2016, CC BY 2.0



Canada pays for response, not preparedness

Response and recovery



Preparedness and mitigation



Indigenous Services Canada spending, 2018–19 to 2021–22

3.5×

more on responding than on preventing

112

eligible mitigation projects unfunded; 74 had waited more than five years

38 of 70

Saskatchewan First Nations not covered by the province's wildfire agreement. The 2025 follow-up found fewer agreements than in 2022.



Design for the conditions, not the demo

Households with 50/10 Mbps unlimited internet, 2023

Canada overall



Rural Canada



First Nations Communities



A tool that needs a strong connection fails exactly where and when it's needed.

REQUIREMENTS, NOT FEATURES

Offline and local first

Low bandwidth

Works on a phone

Still standing when the tower isn't

One person, many hats: saves time, never adds work



Human factors in one slide

Design the system around the people who use it, under the conditions they really work in.

The more we automate, the more we depend on people for the rare, hard cases.

Bainbridge, "Ironies of automation," 1983 · Endsley, "Ironies of artificial intelligence," 2023

THE LENS, ONE IDEA PER SLIDE



The interface shapes the decision



Trust: calibrated, not maximized



Why we're built to believe it, and what stress does



Automation bias



A human in the loop is not enough



Built elsewhere doesn't fit automatically



The interface shapes the decision

ALERT MENU · HAWAII, JANUARY 13 2018

Amber Alert (CAE) – Kauai County Only

Tsunami Warning (CEM) – STATE ONLY

PACOM (CDW) – STATE ONLY

DRILL – PACOM (CDW) – STATE ONLY


Drill and live alert side by side, behind the same prompt. No plan to correct it: 38 minutes.

- More data on screen “may easily result in harming” awareness and performance
- Show the few things that matter
- Too many alerts become noise

Situational awareness is seeing what matters, understanding it, and projecting what happens next. The screen can build that, or bury it.



Trust: calibrated, not maximized



Rely on a tool exactly as much as it deserves. No more, no less.

Trust follows performance, and people drop an algorithm faster than a person after the same mistake. In emergency services, I think over-trust is the bigger danger.

In a survey of 99 firefighters, rescue pressure changed how capable they judged the AI, and so how much they trusted it.

Pressure moves trust, exactly when checking matters most.



Why we're built to believe it



Fluency

What is easy to read feels true. AI text is the most fluent text you will ever read.



The machine heuristic

"The computer says" feels objective. People don't.



Confirmation

We accept what agrees with us. The tool is tuned to agree.



Anchoring

The first plan on the table pulls every plan after it. The AI draft is always first.



And the reverse: after seeing it make mistakes, people abandon a tool that still beats them.

Algorithm aversion. Both errors come from the same place: we judge the tool by how it feels, not by its record.



Under stress, the shortcuts get stronger

CALM



Eight cues in view, including the one that says "how sure".

UNDER STRESS



Two cues. The loudest. The "how sure" line is the first to go.

- **Attention narrows.** Peripheral cues drop out first.
- **Experts take the first workable option.** A confident AI suggestion becomes the first option.
- **Stress closes the search early.** Fewer options considered.
- **Automation bias peaks under workload.** Exactly when the tool is most used.

Under pressure you will look at fewer options than you think. The tool will sound more certain than it is. Design the check for the moment you are weakest.



Automation bias

People accept automated advice when it's wrong, and miss what it misses, most of all under workload and time pressure.

It "cannot be prevented by training or instructions, and can affect decision making in individuals as well as in teams."

Parasuraman and Manzey, 2010



What helps: making a named person clearly accountable for the decision.



What happens if the system provides an incorrect recommendation and the provider follows it? Who is accountable for any resulting medical errors?

The researchers, after interviewing paramedics (Bai et al., 2025)



A human in the loop is necessary. It is not sufficient.

41 OVERSIGHT POLICIES REVIEWED

"People are unable to perform the desired oversight functions."

Green, 2022

319 WORKERS, 936 WORK EXAMPLES

"Higher confidence in GenAI is associated with less critical thinking."

Lee et al., 2025

What works is a check built into the task, not a signature at the end.

Green 2022 · Lee et al. 2025, CHI · Spearing et al. 2026 · Tan et al. 2026 · both reviews

The screenshot displays a web-based 'Incident Action Plan' form. At the top, it says 'Incident Action Plan' and 'ICS-compliant IAP based on chemical, weather, and plume data'. Below this, there's a section for 'PROPERTY/ENVIRONMENT CONSERVATION' with a trigger condition: 'If liquid phosgene or contaminated water is present'. The owner is 'Environmental Unit Leader' and the decision point is 'Notify Ontario Ministry of the Environment if waterway contact is confirmed'. A 'DISPOSITION REQUIRED' section has two buttons: 'Acknowledge' (checked) and 'Override'. A 'High Confidence' status is shown. The '1. SITUATION SUMMARY' section describes a 'Large phosgene release at University of Waterloo campus (43.4683, -80.5429), Waterloo ON — dense heavier-than-air gas drifting SSE at 14 km/h; no safe occupancy downwind to 6.93 km without respiratory protection.' It lists the chemical as 'Phosgene, CAS 75-44-5, UN 1076, a colourless gas with faint hay-like odour' and includes a 'Public Use 00-0171' link. At the bottom, a progress bar shows '0 of 9 recommendations dispositioned', '0 acknowledged', and '0 pending review'. A 'Required before Finalize: 0 sections pending review' message is also present. Navigation buttons for 'Back' and 'Finalize' are at the bottom.

Real screen, test data: the draft cannot be saved without being opened and edited.



AI built elsewhere doesn't automatically fit

RECOGNIZE

On law enforcement images, the highest false positive rates were for American Indigenous faces.

PREDICT

Suicide-risk models tested with a tribe were poorly calibrated until they were recalibrated locally.

GENERATE

Whether a general chatbot reflected tribal-specific information affected how far tribal members trusted it.

Accurate at what, for whom, measured by whom?

A model trained on someone else's data has to be tested, and often recalibrated, with the community it will serve.



My research program

How do commanders calibrate trust in, and appropriately rely on, AI decision support under uncertainty?

- | | | | | | |
|----------|---|------------------|----------|---|------------------|
| 1 | Systematic review: 32 studies | written | 2 | Scoping review: governing AI in emergency services | drafted |
| 3 | Interview study: AI governance in Canadian fire services | planned | 4 | Study 1: how experienced incident commanders decide | ethics approved |
| 5 | Study 2: do commanders catch an AI error in a simulated incident? | ethics in review | 6 | Study 3: how trust changes over repeated use | ethics in review |



Everyone talks about trust. Almost nobody measures it.

1,007

records screened, 2000 to June 2025

32

studies included

7

used any structured
measure

0

validated
scales

- Largest group: paramedicine and emergency departments. Few in fire. None in policing.
- None measured workload, awareness or trust with a standard test.
- Few systems were fully operational. Most were prototypes and simulations.

So when a vendor says “proven in the field,” ask which field, and measured how.



What responders accept

FIREFIGHTERS ON DATA ABOUT THEMSELVES

Never punitive

De-identified

Restricted access

Voluntary

They can see their own data

Where US law said little, some unions bargained data limits into their contracts. Protections varied.

AI review of body camera footage felt less fair to officers when they had no control over recording and review.

Acceptance is about control and purpose, not technology.

For a community, that is the bridge to data sovereignty.



Data sovereignty is a design requirement

OCAP® · FIRST NATIONS

Ownership

Control

Access

Possession

"Possession is the mechanism by which ownership can be asserted and protected."

First Nations Information Governance Centre. OCAP® is a registered trademark of FNIGC.

CARE · INDIGENOUS DATA

Collective benefit

Authority to control

Responsibility

Ethics

Free, prior and informed consent runs through UNDRIP. Canada's 2021 Act affirms its application in Canadian law.

Carroll et al. 2020; Global Indigenous Data Alliance

"All technical systems are cultural and social systems."

The community owns its data

Sharing is an explicit, logged act

Export or delete, any time

Knows where it lives



One principle per screen, around the whole cycle



Preparedness The data stays with the people it belongs to

Response One clear picture, not twelve layers

Response Data is not awareness: one picture, people make sense of it

Evacuation Built for real conditions: a phone, offline, gloves on

Decisions Shows its reasoning and how sure it is

Recovery The human decides, and the decision is recorded

Resilience Not every problem needs AI

The screens that follow are illustrations of each principle, not a product. The principle carries the evidence.



The data stays with the people it belongs to

Sharing and permissions		Data location: community server · backup in Canada	
	Hazard inventory and risk assessment	Community only	
	Emergency plan and evacuation routes	Shared with Tribal Council · logged	
	Registry: who needs help to leave	Emergency coordinator only	
	Live incident map	Province, during a declared event, expires	
	Sacred, burial and harvesting sites	Not held, by default	
		Export everything	Delete everything

Purpose limits, access control and local ownership, set before go-live.

Sharing with a province or agency is an explicit act, logged, and never the default. The community can take it all back.

A Canadian data centre run by a US company is still reachable under US law (CLOUD Act, 2018). Ask whose law reaches it, not only which country.

Who holds the registry of people who need help to leave, and under what consent, is a decision only the community can make.

Illustration. Fictional settings.



RESPONSE

One clear picture, not twelve layers



Show the few things that matter now, and project what comes next.

Where the fire is. Where it will be. Whether the road out is open. Everything else is a tap away, not on the screen.


The projection is the part AI can carry under high workload. The decision to go is still a person's.

Illustration. Fictional terrain and incident.



RESPONSE

Data is not awareness



One row, four things

Chemical, weather, isolation, protective action. Everything else is a tap away.

"Moderate confidence"

Up top, in colour, where narrowed attention still finds it.

"Wind: 7 min old, stale"

The tool says what it doesn't know. That is a degraded mode, by design.

People do the sensemaking. The AI's job is to say what changed.

Real screen, hazmat incident, test data. Same principle for a wildfire picture. · Weick 1993 · Klein et al. 2010 · Guyeux 2026 · both reviews



Three levels of awareness, and where AI fits

LEVEL 1

Perceive

What is out there?

AI CARRIES

Every feed in one place. Flags what changed. Says what is stale.

PEOPLE HOLD

Which feeds matter here.

LEVEL 2

Understand

What does it mean for us?

AI CARRIES

Fuses and summarises, with its confidence showing.

PEOPLE HOLD

The local knowledge no feed contains.

LEVEL 3

Project

What happens next?

AI CARRIES

The projection under workload, with its uncertainty.

PEOPLE HOLD

The check, and the decision.

What each level needs comes from the people who do the work. That is why human factors sits with the end users before anything is built.



Built with responders, for real conditions



Hands busy, gloves on, low bandwidth, and designed with the people who will use it.

- Three big targets, not a form. Works on the phone the coordinator already has.
- Keeps taking registrations with no signal, and syncs when the tower is back.
- No double work: registered once at the hall, counted everywhere.
- Co-design was our review's most consistent recommendation. Firefighters judge new routines against what already works, and their consent is central.

Illustration. Fictional counts.



DECISIONS

It shows its reasoning, and how sure it is

The screenshot displays a software interface for an Incident Action Plan (IAP). At the top, it says 'Incident Action Plan' and 'ICS-compliant IAP based on chemical, weather, and plume data'. Below this, a section titled 'PROPERTY/ENVIRONMENT CONSERVATION' describes a scenario: 'Prevent runoff and atmospheric spread toward Laurel Creek corridor'. It includes a 'Trigger' (liquid phosgene or contaminated water), an 'Owner' (Environmental Unit Leader), and a 'Decision point' (Notify Ontario Ministry of the Environment). A 'DISPOSITION REQUIRED' section has 'Acknowledge' and 'Override' buttons. A 'High Confidence' badge is visible. The '1. SITUATION SUMMARY' section details a phosgene release at the University of Waterloo campus, including coordinates, wind speed, and chemical properties. At the bottom, a status bar shows '0 of 9 recommendations dispositioned' and 'Required before Finalize: 9 sections pending review'. Navigation buttons for 'Back' and 'Finalize' are at the bottom.

Trigger · Owner · Decision point

Every recommendation says why, who, and when to act. With its source.

"High confidence. What does this mean?"

An honest "how sure," and a link that explains it.

Acknowledge, or override with a reason

0 of 9 dispositioned. You cannot finalize until a person has touched each one.

In a computer test, AI claims that passed an evidence check were right 71% of the time; raw AI output, 52%. No people tested yet.



The human decides, and the decision is recorded

Decision record · Wildfire, day 1		Becomes the after-action review	
16:22	Evacuate zones 2 and 3 by 17:30	Accepted	Emergency coordinator
16:40	Suggested route via North Road. Overridden: bridge washed out last spring, not on the map	Overridden	Fire chief
17:05	Shelter capacity alert. Modified: open the school gym	Modified	Reception centre lead
17:48	Share live map with Province for the duration of the event	Approved	Chief and Council
Day 1: 14 suggestions · 9 accepted · 3 modified · 2 overridden		What the map didn't know → add to next year's plan	

An override is one tap, and it is recorded, not punished.

The record of what was decided and why becomes the community's after-action review, and its memory for the next event.

You can't measure reliance you don't record. Two overrides out of fourteen is how you know the tool is being checked.

Illustration. Fictional incident; roles, not names.



Not every problem needs AI



SAFETY

Learn without blame

From events and from normal work: learning reviews, not investigations.



SYSTEMS DESIGN

Fix the work itself

See how the work is really done, then the plan, the forms, the radio protocol, the handover.



HEALTHY WORKPLACES

The people who ran it

Rest after the event, consent to change, and no one punished for an override.



HUMAN FACTORS

Design with, not for

Co-design with the people who do the work, and a measured pilot before any purchase.

**Start with the problem and the people.
Technology is one answer among several, and
often not the first.**

After an emergency: what went wrong that software could fix, and what went right that software must not disturb.



Four things I've learned

1 Communities already run these events, with less than anyone else has. The knowledge of what floods first and who needs help is already there. A tool's job is to hold it, not supply it.

2 No system should decide. Software can propose and record. People decide.

3 Don't arrive with a finished product for a community you haven't sat with.

4 The community is a partner in research, not its subject. The findings go back to the community first.



Three questions only a community can answer



Evacuees get registered by the province or the Red Cross. Who should hold the list of who needs help to leave, under what consent, and does it come back?



The province's wildfire map asks for "values at risk." Do sacred, burial and harvesting places belong on any map, visible to whom? My default is not to hold them.



2023, then 2025. What went wrong in your last emergency that software could actually have fixed, and what went right that software must not disturb?



Six questions to ask

1 Accurate at what, and tested with communities like ours?

2 Wrong in which direction, and what happens then?

3 Who owns the data, where does it live, and can we export or delete it?

4 Does it work offline, and on a phone?

5 Who in our community decides what it holds and shows?

6 Who is accountable when it's wrong?



The one-page card and
the short guide, free

thehumanfactor.ca/events/ipsc2026

What makes oversight real



The person can judge

They have the information, the time and the training to see when the tool is wrong.



The person can disagree

They have the authority to override it, and sometimes they do.



Someone measures it

How much people rely on the tool, and how often they override it, is counted and reviewed.

Take any one away and “a human makes the final call” is a sentence in a policy, not a safeguard.



One hour. Online. Free. Bring your last emergency.



Concept image

- 1 You bring the event, or the worry. I bring the six questions and the human factors lens.
- 2 You leave with a one-page note: what to fix with no software at all, what software could help with, and what to ask before you buy anything.
- 3 If you want to, help shape the software I am building: HazReady, for community emergency preparedness.



Book the hour

thehumanfactor.ca/events/ipsc2026

No product on the table in the hour. No pitch. If it is useful, you will know what to do next.

Human judgment is necessary, and not automatic.

AI can help you think. It can't do your thinking for you, and it can't carry your accountability.

If your community wants a say in how these tools are designed, I'd like to hear from you.



The Human Factor

HUMAN-CENTRED BY DESIGN

I help people build safer systems for humans.

SAFETY · SYSTEMS DESIGN · HEALTHY WORKPLACES ·
HUMAN FACTORS

Scott Ramey · scott@thehumanfactor.ca

thehumanfactor.ca · University of Waterloo, Systems Design Engineering



Slides, the questions card
and the guide

thehumanfactor.ca/events/ipsc
2026

SOURCES

Where the evidence comes from

Auditor General of Canada (2022). Report 8: Emergency Management in First Nations Communities; (2025) Follow-up on Programs for First Nations.

Indigenous Services Canada. Wildland fire and flood evacuation statistics, 2016 to 2026 (not final).

Christianson et al. (2024). Wildland fire evacuations in Canada from 1980 to 2021. Int. J. Wildland Fire.

CRTC (2025). Canadian Telecommunications Market Report, Fig. 22.

FNIGC. The First Nations Principles of OCAP®. · Carroll et al. (2020). The CARE Principles. Data Science Journal.

Lewis, J. E. (Ed.). (2020). Indigenous Protocol and Artificial Intelligence Position Paper. · UNDRIP Act, S.C. 2021, c. 14.

Grother, Ngan and Hanaoka (2019). NIST IR 8280. · Haroz et al. (2024). JAMA Netw Open. · Gupta, Chen and Tsai (2024). IUI Companion.

Moffatt v. Air Canada, 2024 BCCRT 149. · Mata v. Avianca (S.D.N.Y. 2023). · Zhang v. Chen, 2024 BCSC 285.

Treasury Board Secretariat. Guide on the use of generative AI. · Sharma et al. (2024). ICLR. · OpenAI (2025). Sycophancy in GPT-4o.

Bainbridge (1983). Automatica. · Endsley (2023). Ergonomics. · FCC (2018). Report and recommendations: Hawaii Emergency Management Agency January 13, 2018 false alert.

Endsley (1995). Toward a theory of situation awareness. Human Factors. · Jones and Endsley (1996). Aviat. Space Environ. Med. · Endsley, Bolté and Jones (2003). Designing for Situation Awareness. · Endsley (2017). Human Factors.

Parasuraman and Manzey (2010). Human Factors. · Goddard, Roudsari and Wyatt (2012). JAMIA. · Green (2022). CLSR. · Lee et al. (2025). CHI.

Hancock et al. (2011). Human Factors. · Dietvorst, Simmons and Massey (2015). JEP: General. · Zou et al. (2026). Ergonomics. · Dell'Acqua et al. (2023). HBS WP 24-013.

Parush et al. (2020). CogSIMA. · Bai et al. (2025). BMC MIDM. · Park et al. (2025). Appl. Ergon. · Berndt and Herczeg (2019).

Spearing et al. (2026). JARMAC. · Tan et al. (2026). AAAI SS. · Adams et al. (2026). J Exp Criminol. · Adams (2025). J Crim Just.

Topazian et al. (2025). JOEM; (2026) New Solutions. · Saulnier and Abbatangelo (2024). Can. Public Policy. · Brooks, Grugulis and Cook (2022). Human Relations.

Weick (1993). ASQ. · Klein, Moon and Hoffman (2006). · Klein, Wiggins and Dominguez (2010). · Gyeux (2026).

Pupulidy and Vesel (2017). · Hollnagel et al. (2013). Safety-I to Safety-II. · ISO 9241-210:2019. · Conklin (2019).

CPSE Center for Innovation (2025). AI Strategic Scan, n = 156, via FireRescue1. · Tao et al. (2026). Policing 49(1).

Ramey, Masoumi-Ravandi, Rehman and Cao (in preparation). Humans, algorithms, and alarms: a systematic review (32 studies). University of Waterloo.

Ramey and Rehman (draft, 2026). Governing AI use in emergency services: a rapid scoping review (42 sources; provisional counts).

