Empowering Young Canadians in the Smart Device Era:

A Privacy-by-Design Research and Public Engagement Initiative

This report is submitted in fulfillment of
the research project funded by
The Office of the Privacy Commissioner of Canada's (OPC)
Contributions Program

By

Dr. Ajay Shrestha
Principal Investigator
Professor, Computer Science Department
Privacy-Aware AI Research Team
Vancouver Island University
Nanaimo, BC, Canada

March 15, 2026

Table of Contents

Acronyms and abbreviations

Term	Meaning
AI	Artificial Intelligence
ATT	Algorithmic transparency and trust
COPPA	Children's Online Privacy Protection Act (United States)
GDPR	General Data Protection Regulation (European Union)
HTMT	Heterotrait-Monotrait Ratio of Correlations
IoT	Internet of Things
MGA	Multi Group Analysis
OPC	Office of the Privacy Commissioner of Canada
PEA-AI	Privacy-Ethics Alignment in AI
PIPEDA	Personal Information Protection and Electronic Documents Act (Canada)
PLS-SEM	Partial least squares structural equation modeling
PPB	Privacy-protective behavior
PPBf	Perceived privacy benefit
PPR	Perceived privacy risk
PSE	Privacy self-efficacy
PRISMA	Preferred Reporting Items for Systematic reviews and Meta-Analyses
SCOUR	Surveillance mechanisms; Consent and awareness; Operational data flow; Usage and exploitation; Regulatory/technical safeguards
SEM	Structural equation modeling
SVA	Smart voice assistant
UX	User experience

1. Executive summary

This final report documents activities and outcomes for the OPC-funded project "Empowering Young Canadians in the Smart Device Era: A Privacy-by-Design Research and Public Engagement Initiative." The project examined youth privacy in smart voice assistant (SVA) ecosystems and translated research evidence into actionable, youth-centred privacy-by-design guidance for developers, policymakers, educators, parents/guardians, and youth.

Across the reporting year, the work program progressed through five interlocking streams:

- Evidence synthesis: two structured literature reviews (SCOUR-guided review of covert surveillance; PRISMA-guided systematic review of youth-centred privacy-by-design solutions).
- Platform audit: privacy-by-design audit of Google Home Mini, Amazon Alexa (Echo Dot), and Apple Siri, combining heuristic evaluation, PIPEDA checklist assessment, and youth-centred usability testing.
- Youth qualitative study: five semi-structured focus groups (N = 26; ages 16–24) to derive construct-aligned themes and actionable design implications.
- Youth quantitative modelling: survey instrument (five constructs; four items each) analyzed via PLS-SEM with N = 469 valid survey responses; additional subgroup analyses by age (16–18 vs 19–24) and gender identity.
- Translation to practice: Privacy-by-Design Toolkit v1.0 (roundtable-refined; dated March 6, 2026) with tiered guidelines, verification steps, role-based checklists, and practical workflows.

These streams produced eleven primary deliverables used as the evidence base for this report: two literature-review papers [1], [2], one audit paper [3], one focus-group paper [4], five survey-based modelling papers (baseline [5], age-differentiated [6], gender-differentiated [7], privacy profiles [8], tension indices [9]), one integrative journal-style preprint (PEA-AI lens) [10], and the finalized toolkit document [11].

Key results (evidence-convergent across methods):

- Covert/ambient capture and retention uncertainty remain central youth-facing risks in smart-device contexts, with consent and awareness gaps recurring across the SCOUR lenses (covert surveillance review: 1,930 records identified; 171 included).
- The systematic review found that the research corpus strongly favours technical solutions (67%) over policy (21%) and education/awareness approaches (12%), indicating an implementation gap outside the technical domain (2,216 records identified; 122 included).
- The audit identified measurable usability–compliance trade-offs. PIPEDA totals differed across platforms (Google Home 16/20; Alexa 15/20; Siri 18/20), and a youth-centred UX task

set showed that disabling voice history was consistently the most time-consuming task across platforms.

- Qualitative analysis shows that privacy-protective behavior in smart voice assistants is shaped by perceived risks and benefits, transparency and trust, and privacy self-efficacy, with policy overload, fragmented settings, and unclear data retention reducing confidence and limiting protective action.
- PLS-SEM modelling indicates that privacy self-efficacy (PSE) is a primary lever for privacy-protective behavior (PPB). Algorithmic transparency and trust (ATT) influences PPB primarily through PSE (indirect effect), while perceived benefits show a tension-consistent role (positive via PSE, negative direct effect on PPB).
- Age and gender analyses suggest that core mechanisms are stable while certain pathways differ in strength. For age groups, ATT → PSE is stronger among older youth (ages 19-24; $\beta = 0.567$) than younger youth (ages 16-18; $\beta = 0.356$; $p = 0.024$). For gender groups, PPR → PPB is stronger for males than females (Male: $\beta = 0.424$; Female: $\beta = 0.233$; $p < 0.1$) and the indirect ATT → PSE → PPB pathway is stronger for females than males (Female: $\beta = 0.229$; Male: $\beta = 0.132$; $p < 0.1$).
- The integrative manuscript frames youth privacy as ongoing negotiation (not a one-time consent event) and articulates four actionable design principles, including making data flows visible/controllable, building skills through guided interaction, reducing friction for protective actions, and adaptable autonomy across developmental stages.
- The toolkit consolidates these results into tiered, verifiable guidance with role-based checklists and workflows, supporting evidence-to-implementation traceability.

Across research and engagement streams, the project advances a forward-looking, youth-centred model of privacy governance that moves beyond disclosure alone toward actionable transparency, measurable usability, and skills-based supports that enable young people to exercise meaningful control.

2. Project overview and objectives

Smart voice assistants and voice-enabled smart devices are increasingly used by youth in private bedrooms, shared living spaces, and mobile contexts [12], [13]. These ecosystems rely on always-on microphones, wake-word detection, cloud processing, and cross-service integration [1], [14], [15]. As a result, youth privacy risks are shaped not only by what data are collected, but also by how legible and verifiable privacy controls are in day-to-day use [16], [17], [18]. This project treats youth (ages 16–24, as implemented across the empirical studies) as young digital citizens and examines how perceptions and trust translate into privacy-protective behavior in SVA ecosystems.

Project goal and objectives (from the approved proposal, operationalized through the work program):

- Comprehensive assessment: investigate real-world usage and privacy implications of smart devices, focusing on voice-activated assistants and the interplay of perceived privacy risk and benefit, algorithmic transparency and trust, privacy self-efficacy, and privacy-protective behavior.
- Data collection and analysis: conduct focus groups and surveys with youth (ages 16–24) and analyze results using qualitative methods and structural equation modeling to identify drivers of protective action.
- Platform audit: evaluate leading commercial ecosystems (Google Home, Alexa, Siri) to assess both privacy compliance and usability of privacy controls using a combined framework.
- Translation to practice: develop and validate a Privacy-by-Design Toolkit that supports actionable transparency, verifiable controls, and skill-building to strengthen youth privacy self-efficacy.
- Public engagement and dissemination: translate findings through conference papers, a project website/repository hub, outreach materials, and public-facing engagement activities (webinar and workshop).

Core conceptual model used in the empirical program:

Construct	**Definition (operational focus)**
Perceived privacy risk (PPR) [19], [20]	Concerns about collection, access, and retention of voice interactions and associated data.
Perceived privacy benefit (PPBf) [21], [22]	Perceived value from convenience, personalization, and time-saving utility that may motivate data sharing.

Construct	Definition (operational focus)
Algorithmic transparency and trust (ATT) [23], [24]	Perceived openness, fairness, and responsible data handling by SVA providers; perceived clarity about data processing.
Privacy self-efficacy (PSE) [25], [26]	Confidence and capability to manage settings, prevent unwanted recording, update permissions, and manage privacy risks.
Privacy-protective behavior (PPB) [27], [28]	Actual protective actions such as reviewing permissions, deleting voice history, refusing features, and additional protective measures.

3. Work completed and deliverables

The project followed the staged work plan described in the proposal. Early phases focused on ethics approvals, recruitment, literature synthesis, and the platform audit. Later phases focused on empirical analysis, integration across methods, toolkit refinement and validation, and knowledge mobilization.

3.1 Deliverables mapping (proposal → completed outputs)

Deliverable category	Completion evidence
Two peer-reviewed full regular conference papers (literature reviews / analytical works)	Completed: PRISMA-based SCOUR covert surveillance review (IEEE UEMCON 2025) and PRISMA-based systematic review (IEEE CCWC 2026).
Audit-based peer-reviewed full regular conference paper on privacy-by-design in commercial smart devices	Completed: Privacy-by-design audit paper (IEEE CCWC 2026).
Survey and focus group findings and dissemination	Completed: focus-group qualitative analysis paper (IEEE CCWC 2026); survey PLS-SEM quantitative analysis paper (IEEE ICAIIC 2026); age-differentiated MGA paper (IEEE ICAIC 2026); gender heterogeneity MGA paper (submitted to IEEE CCECE 2026); privacy profiles paper (in draft version); tension indices paper (submitted to IEEE CSP 2026).
Privacy-by-Design Toolkit	Completed: Toolkit v0.1 (draft; dated February 4, 2026); Final Toolkit v1.0 (roundtable-refined; dated March 6, 2026) with tiered guidelines, verification steps, checklists, and workflows.
Engagement and educational outreach materials	Implemented: project website/repository hub and updates; social media outreach (Medium, LinkedIn, X); planned final webinar (Feb 24, 2026) and youth-centred workshop (Mar 3, 2026) in collaboration with VIU Library and Global Citizens Forum; explainer-style communication referenced via project Medium post link provided in project communications.

3.2 Primary outputs (papers and toolkit)

- Covert Surveillance in Smart Devices: A SCOUR Framework Analysis of Youth Privacy Implications — IEEE UEMCON 2025 (published) [1].
- Toward Youth-Centered Privacy-by-Design in Smart Devices: A Systematic Review — IEEE CCWC 2026 (published) [2].

- Balancing Usability and Compliance in AI Smart Devices: A Privacy-by-Design Audit of Google Home, Alexa and Siri — IEEE CCWC 2026 (published) [3].
- Convenience vs. Control: A Qualitative Study of Youth Privacy with Smart Voice Assistants — IEEE CCWC 2026 (published) [4].
- Privacy by Voice: Modeling Youth Privacy-Protective Behavior in Smart Voice Assistants — IEEE ICAIIC 2026 (in press / hybrid conference / survey baseline model) [5].
- Age-Differentiated Pathways to Privacy Protection in Smart Voice Assistants: A multigroup PLS-SEM Study of Youth — IEEE ICAIC 2026 (published / hybrid conference) [6].
- Gender-Based Heterogeneity in Youth Privacy-Protective Behavior for Smart Voice Assistants: Evidence from Multigroup PLS-SEM — Preprint / submitted for peer review (IEEE CCECE 2026) [7].
- Privacy Profiles of Youth Smart Voice Assistant Users: A Cluster Analysis of Risk, Benefits Trust, and Behavior— Preprint (draft version only) [8].
- Negotiating Privacy with Smart Voice Assistants: Risk–Benefit and Control–Acceptance Tensions — Preprint / submitted for peer review (IEEE CSP 2026) [9].
- From Framework to Practice: Youth Negotiations of Privacy with Smart Voice Assistants Through the PEA-AI Lens — Integrative journal-style preprint (submitted for peer review to ACM Journal of Responsible Computing) [10].
- Privacy-by-Design Toolkit v1.0: Youth-Centred Guidelines for Smart Voice Assistants and Voice-Enabled Smart Devices — Toolkit (v1.0, March 6, 2026) [11].

3.3 Recruitment and dataset status (empirical stream)

This study received ethics approval from VIU-REB. The approval reference number #103597 was given for behavioral/amendment forms, consent form, and questionnaire. Data collection was completed on November 20, 2025 (extension approved through relevant governance processes). A total of 494 survey responses were collected; the survey modelling papers report 469 valid survey responses retained for PLS-SEM analysis. Five semi-structured focus groups (N = 26) were conducted with youth participants aged 16–24. Outside of Vancouver Island University, recruitment was enabled through approvals from Greater Victoria (SD61), Sooke (SD62), Nanaimo & Ladysmith (SD68), Campbell River (SD72), and Cowichan (SD79), as well as institutional recruitment through Carleton University and the University of Saskatchewan.

Dataset component	Collected	Used in analysis
Survey responses	494 total (unfiltered)	469 valid for PLS-SEM modelling
Focus groups	5 sessions	N = 26 participants (transcribed and thematically coded)

Data stewardship: study data were stored on institutional Microsoft OneDrive and removed from the Microsoft Forms collection environment on January 20, 2026.

4. Methods and implementation

4.1 Ethics, recruitment, and data stewardship

All youth-facing research activities proceeded under ethics oversight. Recruitment targeted Canadian youth aged 16–24 with experience using smart voice assistants. The empirical stream used two complementary data sources: (i) a cross-sectional survey analyzed using PLS-SEM, and (ii) semi-structured focus groups analyzed using thematic coding mapped to the five-construct model.

Data stewardship followed secure storage practices consistent with institutional requirements. Data were stored on Microsoft OneDrive and removed from Microsoft Forms once collection concluded.

4.2 Literature synthesis methods

Two literature reviews were conducted to ground the project in the broader youth smart-device privacy landscape and to identify intervention levers beyond product-specific settings. Both reviews used PRISMA-guided screening [29] but differed in analytic lens and scope.

4.2.1 SCOUR-guided covert surveillance review

This review addressed the question: "Can smart/AI-enabled devices and applications intentionally and surreptitiously record private conversations?" It used PRISMA screening and organized findings using the SCOUR framework: Surveillance mechanisms; Consent and awareness; Operational data flow; Usage and exploitation; Regulatory and technical safeguards.

Search and screening summary: the process resulted in 1,930 records (academic papers, conference papers, and reputable grey literature), with 171 papers included after screening and full-text review. Literature collection used Python scripts to scrape Google Scholar and Crossref via APIs, with a manual check to validate coverage. The review used a multi-term Boolean keyword strategy that combined smart-device terms, covert recording/surveillance terms, youth terms, consent/privacy terms, and policy/regulation terms.

SCOUR framework categories:

SCOUR lens	**Description**
S – Surveillance mechanisms	How does the device record data? Does it rely on constant listening, wake-word triggers, or background surveillance? Are there vulnerabilities (e.g., misfires, hidden features)?
C – Consent and awareness	How is consent sought (or not sought)? How transparent are the terms about audio/data capture? Is age-appropriate language used?

SCOUR lens	Description
O – Operational data flow	Where does recorded data go (cloud, third party, local)? Who has access? How long is it retained? How is it shared or sold?
U – Usage and exploitation	Are there documented cases of exploitation for marketing, advertising, or profiling? Do stakeholders use recorded data in unexpected ways?
R – Regulatory and technical safeguards	What technical safeguards and policy measures are proposed (including PIPEDA-related or similar regulatory measures)?

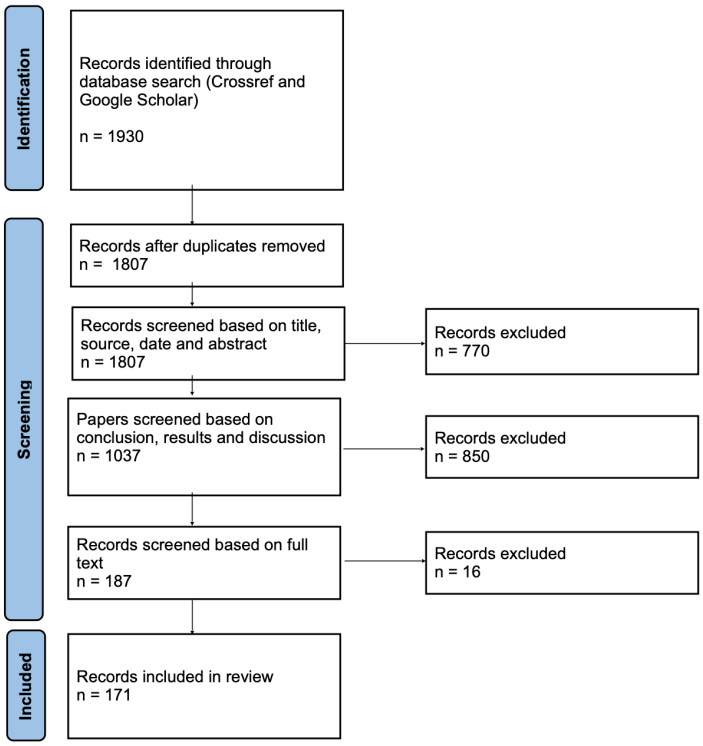

Figure 1. PRISMA flow diagram for the SCOUR-guided covert surveillance review (records identified n = 1,930; included n = 171).

4.2.2 Youth-centred privacy-by-design systematic review

This PRISMA-guided systematic review synthesized privacy-by-design solutions for youth in smart devices across technical mechanisms, policy/regulatory measures, and education/awareness strategies. The review adopted a scoping emphasis and synthesized findings qualitatively (not as effect-size estimates). It posed five research questions spanning technical approaches, policy gaps, education roles, implementation challenges, and multi-stakeholder collaboration.

Search and screening summary: the search identified 2,216 records across four databases (Google Scholar, Scopus, Springer, Crossref). After screening, 645 articles underwent eligibility

assessment, and 122 were included for analysis. The included corpus was categorized into three thematic categories: technical solutions, policy measures, and education/awareness strategies.

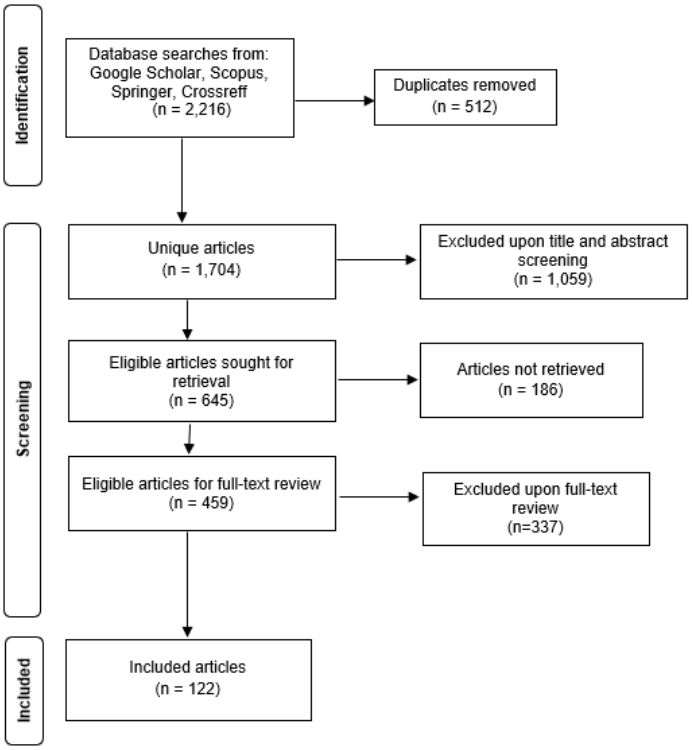

Figure 2. PRISMA flow diagram for the youth-centred privacy-by-design systematic review (records identified n = 2,216; included n = 122).

4.3 Privacy-by-design audit methods

The audit evaluated three widely used ecosystems—Google Home Mini (1st generation), Amazon Echo Dot (3rd generation), and Siri on a MacBook Air (M4 2025). Device configuration and evaluation were conducted during September–October 2025 with region and language settings appropriate for Canadian users. Each device was set up using newly created accounts to assess default privacy experiences.

Evaluation framework: a structured three-part approach combining (i) a heuristic evaluation tailored to usability of privacy controls (seven criteria) [30], [31], [32], (ii) a PIPEDA compliance checklist assessing the ten fair information principles (scored 0–2 per principle) [33], [34], and (iii) youth-centred UX tasks focused on real-world privacy-management actions [35].

4.4 Focus group methods

The qualitative stream used five semi-structured focus groups (N = 26) with young Canadians aged 16–24. Analysis used a construct-aligned thematic approach across the five core constructs (PPR, PPBf, ATT, PSE, PPB). A codebook was developed with code families aligned to constructs

and subcodes capturing recurring themes (e.g., ambient listening anxiety, policy overload, low navigation efficacy, permission refusal, and physical mitigations).

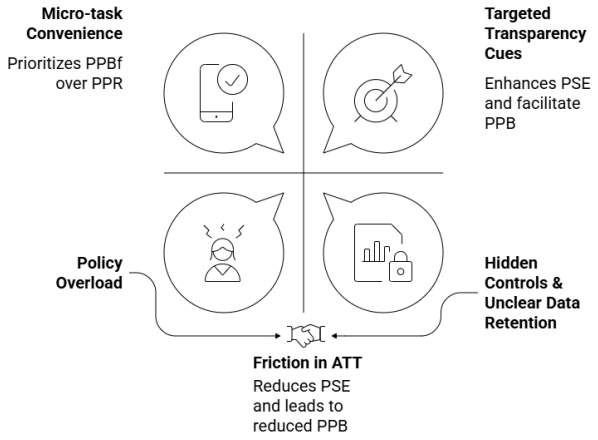

Figure 3. Construct-linked qualitative pathway summary from the focus-group study (ATT friction → reduced PSE → reduced PPB).

4.5 Survey modelling methods (PLS-SEM)

The quantitative stream used a cross-sectional survey analyzed using partial least squares structural equation modeling (PLS-SEM). The instrument operationalized the five constructs with four items per construct. The baseline model tested direct effects on privacy-protective behavior (PPB) and mediation through privacy self-efficacy (PSE).

The baseline modelling paper reports N = 469 valid survey responses. Subgroup analyses were implemented through multigroup analysis (MGA), including age-group comparison (younger youth 16–18 vs older youth 19–24) and gender-group comparisons.

5. Key findings

This section synthesizes findings across the project's core outputs. The findings are presented in a way that preserves the structure of each paper while enabling cross-method triangulation and evidence-to-toolkit traceability.

5.1 Covert surveillance review (SCOUR-guided)

The covert surveillance review identified 1930 records and included 171 papers after screening. It synthesizes evidence that always-on and wake-word architectures can create inadvertent or covert capture risks, and that user mental models often do not match actual data-flow realities. Across the SCOUR lenses, recurring issues include inadvertent activation, limited consent legibility, unclear retention and access controls, and gaps between regulatory guidance and user-verifiable controls.

Synthesis by SCOUR lens (condensed) [1]:

- Surveillance mechanisms: continuous listening and wake-word approaches create inadvertent activation risk and raise concerns about background capture; hidden features and vulnerabilities can amplify risk.
- Consent and awareness: terms and disclosures are often long and not youth-legible; consent mechanisms may not support meaningful understanding of audio/data capture.
- Operational data flow: uncertainty about where recordings are stored, who can access them, retention duration, and third-party sharing is a central driver of perceived risk.
- Usage and exploitation: documented concerns include profiling, targeted advertising, and unexpected secondary uses, including misuse by unauthorized actors.
- Regulatory and technical safeguards: recommendations emphasize stronger transparency, better default protections, and education to support informed control; the review also highlights the limits of current regulatory fragmentation.

5.2 Systematic review of youth-centred privacy-by-design solutions

The systematic review identified 2216 records; 645 articles underwent eligibility assessment, and 122 were included for synthesis. Across the included corpus, technical solutions dominate (67%) relative to policy measures (21%) and education/awareness strategies (12%).

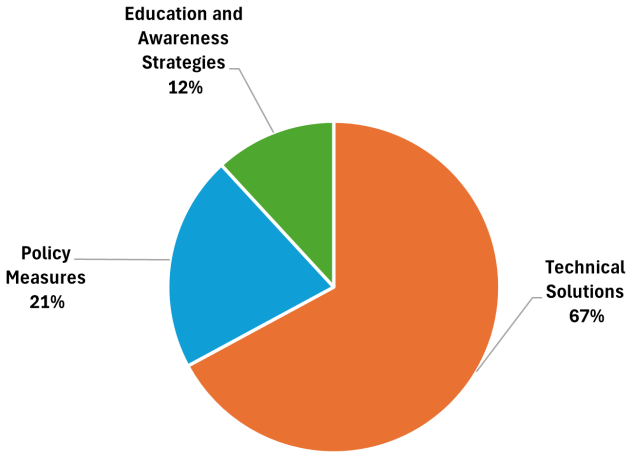

Figure 4. Distribution of reviewed studies by category: Technical (67%), Policy (21%), Education & awareness (12%).

This systematic review [2] further summarizes key themes across five research questions as shown below.

Research question focus	**Condensed key findings**
RQ1 (technical approaches)	Federated learning (including blockchain-enhanced variants), lightweight encryption, trusted execution environments, edge computing, and privacy-preserving authentication are repeatedly cited as mechanisms to enhance data security for youth-used smart devices.
RQ2 (policy frameworks)	Frameworks vary: GDPR/COPPA offer stricter protections while PIPEDA is described as lacking explicit youth protections; enforcement gaps and over-reliance on parental controls are recurring concerns.
RQ3 (education/awareness)	Curriculum integration, gamification, and stakeholder collaboration can improve youth privacy awareness; empowerment-based strategies can help bridge privacy paradox dynamics.
RQ4 (implementation challenges)	Barriers include resource constraints and compatibility issues (technical), weak enforcement and vague consent models (policy), and low digital literacy and privacy paradox barriers (education).
RQ5 (collaboration)	The review emphasizes multi-stakeholder collaboration: policymakers, manufacturers, and educators should coordinate to implement age-appropriate design principles, secure architectures, and scalable educational supports.

5.3 Privacy-by-design audit of Google Home, Alexa, and Siri

The audit combined heuristic evaluation, PIPEDA compliance scoring, and youth-centred UX tasks. Results highlight that compliance, usability, and verification are not perfectly aligned across platforms.

5.3.1 Heuristic evaluation (usability of privacy controls)

Table reproduced from the audit paper [3] (scores 0–2 per criterion; maximum total 14):

Criterion	Google Home	Alexa	Siri
1. Discoverability	2	2	2
2. Comprehensibility	2	2	2
3. Control	2	2	2
4. Granularity	2	2	2
5. Feedback	2	1	1
6. Reversibility	2	2	2
7. Consistency	2	2	2
Total	14	13	13

Interpretation: All three systems achieve near-ceiling scores on basic usability measures. Google performs the strongest overall, largely due to clearer feedback mechanisms such as confirmation pop-ups. By contrast, both Siri and Google place some voice-history controls outside the main assistant menu, increasing navigational uncertainty.

5.3.2 PIPEDA compliance checklist

Table reproduced from the audit paper [3] (scores 0–2 per principle; maximum total 20):

PIPEDA principle	Google Home	Alexa	Siri
1. Accountability	1	1	1
2. Identifying purposes	2	2	2
3. Consent	1	1	2
4. Limiting collection	2	1	2
5. Limiting use/disclosure	2	1	2
6. Accuracy	1	1	2
7. Safeguards	2	2	2
8. Openness	1	2	2
9. Individual access	2	2	1
10. Challenging compliance	2	2	2
Total	16	15	18

Interpretation: Siri achieves the highest overall privacy compliance score (18/20), supported by stronger consent transparency and more robust data accuracy controls. By contrast, Google and Alexa rely more heavily on opt-out defaults for certain optional uses, such as personalization. Performance on individual access and openness also varies depending on whether controls are available directly within the app or are distributed across external account dashboards.

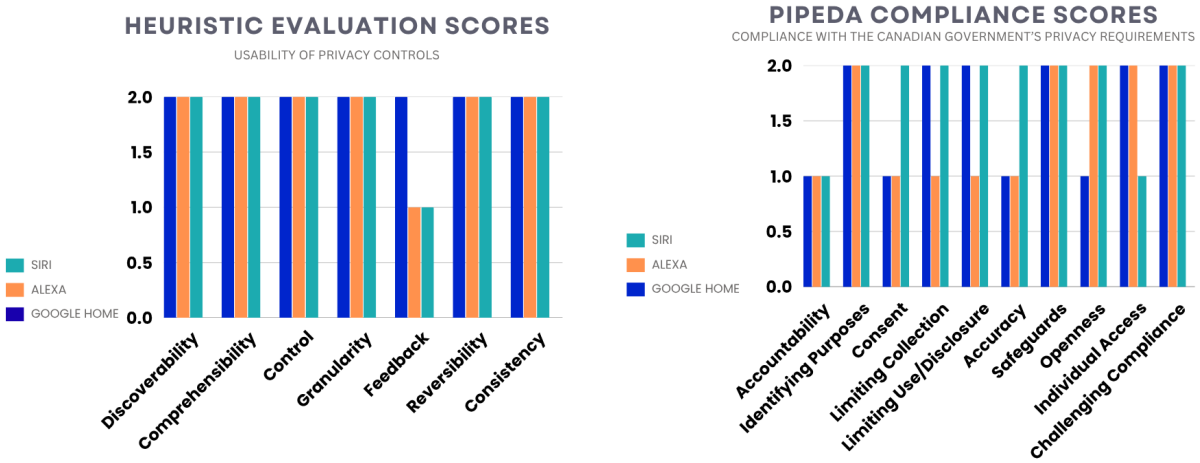

Figure 5. Audit heuristic evaluation and PIPEDA scores visualization.

5.3.3 Youth-centred UX task performance

Table reproduced from the audit paper [3] (average ease 1–5, average steps, average time in seconds):

Task	Ease (GH)	Ease (Alexa)	Ease (Siri)	Steps (GH)	Steps (Alexa)	Steps (Siri)	Time (GH)	Time (Alexa)	Time (Siri)
1. Disable Voice History	3.5	4.3	2.5	5	5	5	95	103	140
2. Find Data Sharing Info	4.3	4.5	4.3	3	3	3	80	41	100
3. Delete Commands	4.3	4.8	4.7	5	4	4	95	38	100
4. Verify Change	4.3	4.7	4.5	5	5	3	50	29	60

Interpretation: Across devices, disabling voice history required multiple steps and the longest time. The audit notes that Siri's relevant control was nested outside Siri-specific menus, increasing navigational complexity. The audit also reports user confusion when privacy changes take time to apply (notably on Alexa), which can undermine verification confidence.

5.4 Focus groups: convenience vs control

The focus-group study analyzed five sessions (N = 26; ages 16–24) and structured findings under the five core constructs. Across themes, youth described SVA use as a continuous negotiation between micro-task convenience and privacy concerns, with friction in transparency and control reducing self-efficacy and, in turn, protective action.

5.4.1 Codebook snapshot (construct-aligned)

The table below reproduces the codebook snapshot reported in the focus-group paper [4] (condensed).

Construct family	Subcode	Definition	Typical inclusion cues
PPR	Ambient "always-listening"	Anxiety that SVAs passively capture to await wake words	Mentions of microphones "always on"
PPR	Retention unknowns	Uncertainty about storage duration and secondary use	Deletion timelines; doubts about account deletion
PPBf	Micro-task convenience	Fast reminders/timers/ simple queries	Hands-busy routines
PPBf	Entertainment	Music/playback convenience	Always-on entertainment use
ATT	Policy overload	Long, unreadable notices	Calls for simpler text
ATT	Hidden controls	Hard-to-find privacy settings	Discoverability/jargon issues
PSE	Low navigation efficacy	Inability to find/use controls	"Don't know where to push"
PSE	Device-conditional efficacy	Confidence varies by device	"Phone yes; speaker no"
PPB	Permission refusal	Denying features to maintain privacy	Keeping prompts on "No"; uninstall
PPB	Physical mitigations	Hardware mute/unplug/ separate strategies	Sensor off-switches

5.4.2 Thematic results by construct (summary)

Themes reported in the focus-group study:

- A1: Ambient listening and uncertain retention raise baseline risk
- A2: Suspected cross-app inferences amplify surveillance concerns
- B1: Micro-task convenience sustains everyday use

- B2: Situational utility in hands-busy contexts and entertainment sustain use
- C1: Policy overload and hidden controls undermine transparency
- C2: Retention/deletion opacity depresses trust
- D1: Low navigation efficacy blocks protective action
- D2: Efficacy is device-conditional; youth ask for brief scaffolds
- E1: Permission and scope management are primary mitigations
- E2: Physical and situational strategies supplement software controls

Implication: The focus-group pathway model suggests that targeted, in-context transparency cues can strengthen PSE and facilitate PPB, while policy overload and hidden controls increase friction in ATT and reduce PSE and PPB.

5.5 Survey modelling: baseline PLS-SEM results

The baseline survey paper analyzed N = 469 youth using PLS-SEM. The model demonstrates moderate explanatory power, explaining 24.2% of the variance in privacy-protective behavior (PPB) and 24.1% of the variance in privacy self-efficacy (PSE).

5.5.1 Survey constructs and item operationalization

The table below represents the constructs and item wording categories from the baseline survey study [5].

Construct	Items (abbreviated)
PPR	PPR1: concern about the amount collected; PPR2: worry about recording without awareness/consent; PPR3: belief in unauthorized access risk; PPR4: unease about storage duration
PPBf	PPBf1: time/effort saving; PPBf2: personalization worth data sharing; PPBf3: benefits outweigh worries; PPBf4: value of preference learning
ATT	ATT1: understanding of information collected/stored; ATT2: trust responsible handling; ATT3: upfront explanation of processing; ATT4: belief in fair/unbiased recommendations
PSE	PSE1: know how to access/adjust settings; PSE2: can prevent recording when undesired; PSE3: confidence to update permissions; PSE4: belief in the ability to manage privacy risks
PPB	PPB1: review/update permissions; PPB2: delete voice activity history; PPB3: refuse features to maintain privacy; PPB4: additional protective measures

5.5.2 Descriptive results: Perceptions vs. capability

The bar chart shows a clear descriptive pattern across the five constructs. Perceived privacy risk has the highest mean score at approximately 3.61 out of 5, while algorithmic transparency and trust has the lowest at approximately 2.52; privacy self-efficacy remains near the midpoint at approximately 2.97.

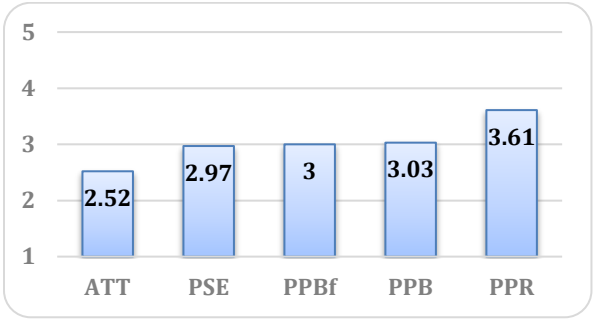

Figure 6. Mean construct scores (1-5 Likert)

Overall, these results suggest that youth express meaningful privacy concern, yet report relatively low trust, limited clarity about how smart voice systems handle data, and only moderate confidence in managing privacy settings. This pattern raises a central question for the model: whether the key barrier to privacy-protective behavior lies not in risk awareness itself, but in an efficacy gap between concern and the ability to act effectively.

5.5.3 Measurement model quality

Before interpreting the structural relationships, the measurement model was assessed to confirm reliability and validity. The results indicate that the model satisfies standard thresholds for reliability and convergent validity, with most indicator loadings at or above recommended levels, average variance extracted exceeding 0.50, and internal consistency measures above 0.70. Discriminant validity was also supported, as all HTMT ratios remained below 0.85; the highest value, observed between privacy self-efficacy and algorithmic transparency and trust, was 0.583 and remained well within the acceptable range. In addition, multicollinearity was not a concern, with a maximum variance inflation factor (VIF) of 1.47. Collectively, these results support the interpretation of the structural paths as relationships among well-measured and conceptually distinct constructs rather than artifacts of measurement error or item overlap.

5.5.4 Structural paths (direct effects)

The structural model explains a meaningful share of variance in the two key outcomes, with R^2 values of 0.242 for privacy-protective behavior and 0.241 for privacy self-efficacy, indicating that approximately one quarter of the variance in each construct is explained. The standardized path coefficients show that privacy self-efficacy is the strongest direct predictor of privacy-protective behavior ($\beta = 0.373$, $p < 0.001$), followed by perceived privacy risk ($\beta = 0.343$, $p < 0.001$). Perceived

privacy benefits show a weaker but notable trade-off effect, with a negative association with privacy-protective behavior (β = −0.130, p < 0.1). Among the antecedents of privacy self-efficacy, algorithmic transparency and trust has the strongest positive effect (β = 0.434, p < 0.001), while perceived privacy benefits also contribute a small positive effect (β = 0.121, p < 0.1). Notably, the direct effect of algorithmic transparency and trust on privacy-protective behavior is not significant, nor is the effect of perceived privacy risk on privacy self-efficacy. Overall, these results suggest that privacy behavior is shaped by both capability and risk-benefit evaluation, with confidence emerging as the strongest driver of action and transparency and trust influencing behavior primarily through their effect on self-efficacy.

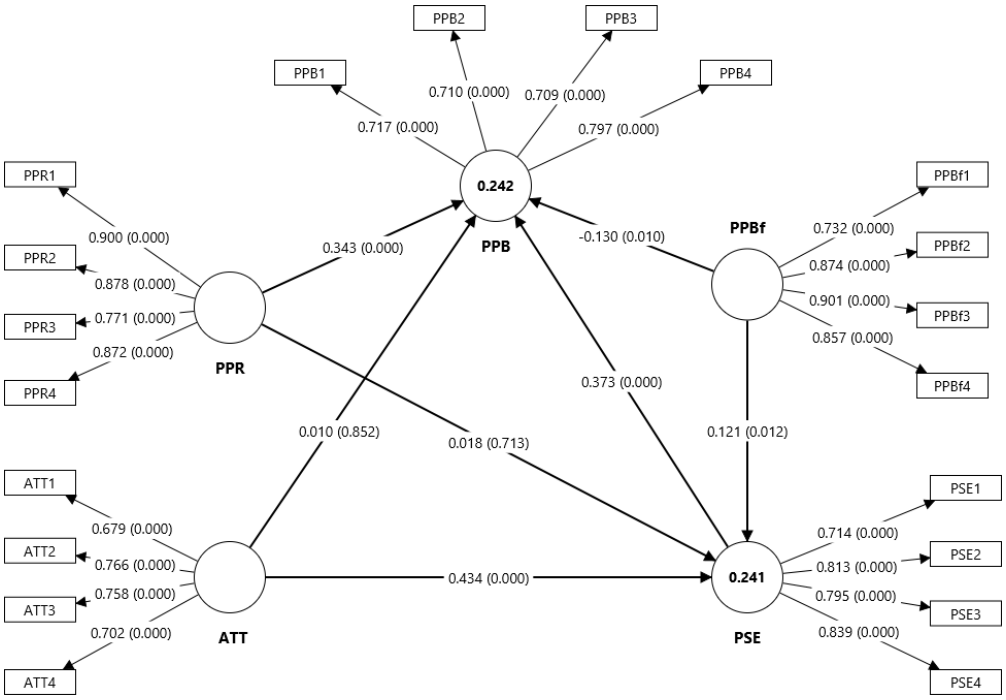

Figure 7. Baseline PLS-SEM structural model and measurement model (N = 469).

Table: Direct effect analysis

Path	**Std β**	**T**	**P**	**f^2**	**VIF**
PPR → PPB	0.343	7.096	0.000	0.134	1.158
PPBf → PPB	-0.130	2.574	0.010	0.017	1.310
ATT → PPB	0.010	0.187	0.852	0.000	1.466
PSE → PPB	0.373	7.869	0.000	0.139	1.317
PPR → PSE	0.018	0.367	0.713	0.000	1.157
PPBf → PSE	0.121	2.503	0.012	0.015	1.290
ATT → PSE	0.434	10.160	0.000	0.204	1.218

5.5.5 Mediation (indirect effects via PSE)

The mediation analysis clarifies the mechanism underlying the structural results. The most important indirect effect is the path from algorithmic transparency and trust to privacy-protective behavior through privacy self-efficacy ($\beta = 0.162$, $p < 0.001$), indicating full mediation. This suggests that transparency and trust influence protective behavior only to the extent that they strengthen users' confidence in their ability to manage privacy controls effectively. A smaller indirect effect is also observed for perceived privacy benefits through privacy self-efficacy ($\beta = 0.045$, $p < 0.1$), indicating partial mediation; although benefits slightly enhance self-efficacy, their overall direct effect still tends to reduce privacy-protective behavior. By contrast, perceived privacy risk does not significantly influence behavior through privacy self-efficacy. Overall, these findings point to an efficacy gap in which clarity and trust alone are insufficient to change behavior unless they also increase users' sense that they can act successfully and verify the outcome.

Table: Indirect effect analysis

Indirect path	**Std β**	**P**	**95% CI**
PPR → PSE → PPB	0.007	0.719	[-0.028, 0.044]
PPBf → PSE → PPB	0.045	0.019	[0.009, 0.085]
ATT → PSE → PPB	0.162	0.000	[0.114, 0.221]

Key Interpretation: PSE shows a strong positive association with PPB, and ATT influences PPB primarily through PSE (significant indirect effect). Perceived benefits have a positive association with PSE and a negative direct association with PPB, reflecting a risk–utility tension.

5.6 Age-differentiated pathways (MGA)

The age-differentiated analysis used survey data from 412 Canadian participants aged 16–24, comparing younger youth (16–18; N = 245) and older youth (19–24; N = 167). It reports mean differences across constructs and multigroup pathway comparisons.

5.6.1 Mean differences by age group

Construct	**16–18 mean (SD)**	**19–24 mean (SD)**	**p**
ATT	2.60 (0.69)	2.38 (0.77)	0.002
PPB	2.97 (0.75)	3.16 (0.78)	0.012
PPBf	3.09 (0.93)	2.84 (1.01)	0.012
PPR	3.44 (0.93)	3.85 (0.84)	< 0.001
PSE	3.04 (0.78)	2.92 (0.90)	0.153

5.6.2 Multigroup analysis (path differences)

Path	β (16–18)	β (19–24)	Difference	p
ATT → PPB	0.042	0.125	0.083	0.485
PPBf → PPB	-0.163	-0.152	0.012	0.912
PPR → PPB	0.307	0.337	0.030	0.791
PSE → PPB	0.330	0.376	0.046	0.655
ATT → PSE	0.356	0.567	0.212	0.024
PPBf → PSE	0.162	0.012	-0.150	0.159
PPR → PSE	0.004	0.045	0.040	0.696

Interpretation: The pathway from ATT to PSE is significantly stronger among older youth (β = 0.567) than younger youth (β = 0.356; p = 0.024), suggesting that transparency and trust translate into confidence more strongly as autonomy increases.

5.7 Gender-based heterogeneity (MGA)

The gender-focused analysis compared male (N = 241) and female (N = 174) participants (total N = 415) and also reported descriptive patterns for non-binary (N = 15) and prefer-not-to-say (N = 39) respondents.

5.7.1 Mean differences (male vs female)

Construct	Male mean (SD)	Female mean (SD)	p	SMD
ATT	2.57 (0.75)	2.54 (0.67)	0.702	0.038
PPB	3.03 (0.78)	2.96 (0.75)	0.307	0.102
PPBf	3.08 (0.96)	2.92 (0.85)	0.075	0.179
PPR	3.59 (0.96)	3.56 (0.87)	0.762	0.030
PSE	3.08 (0.87)	2.84 (0.74)	0.004	0.296

5.7.2 Multigroup analysis (selected direct and indirect differences)

Path	β (Female)	β (Male)	Difference	p
ATT → PPB	-0.071	0.069	-0.141	0.192
PPBf → PPB	-0.125	-0.111	-0.015	0.893
PPR → PPB	0.233	0.424	-0.191	0.062
PSE → PPB	0.469	0.330	0.140	0.170
ATT → PSE	0.489	0.400	0.089	0.308
PPBf → PSE	0.087	0.099	-0.012	0.909
PPR → PSE	0.106	-0.049	0.155	0.138
PPR → PSE → PPB	0.050	-0.016	0.066	0.088
PPBf → PSE → PPB	0.041	0.033	0.008	0.843
ATT → PSE → PPB	0.229	0.132	0.098	0.091

5.7.3 Descriptive statistics (non-binary and prefer-not-to-say)

Construct	Non-binary mean (SD)	Prefer not to say mean (SD)
ATT	1.83 (0.54)	2.39 (0.70)
PPB	3.32 (0.80)	3.24 (0.77)
PPBf	2.10 (1.07)	3.16 (1.02)
PPR	4.22 (0.45)	3.73 (0.83)
PSE	2.50 (0.80)	3.05 (0.84)

Interpretation: The analysis reports a significant mean difference for PSE (males higher than females, $p = 0.004$). Marginal pathway differences include PPR → PPB stronger for males than females ($p = 0.062$) and ATT → PSE → PPB indirect effect stronger for females than males ($p = 0.091$).

5.8 Integrative framing: youth privacy negotiations through the PEA-AI lens

The integrative preprint [10] brings together the survey findings and earlier qualitative insights through the PEA-AI lens, framing youth privacy in smart voice assistants as an ongoing negotiation among perceived risk, perceived benefit, transparency and trust, and the capacity to act. As shown in Figure a1, the overall item-level pattern is marked by consistently elevated privacy concern, weak transparency and trust, and only moderate privacy self-efficacy and privacy-protective behavior. The strongest reported protective action is refusal to use certain features, whereas more active and sustained practices, such as reviewing permissions and deleting voice history, remain comparatively infrequent. This pattern suggests that concern is present, but it does not consistently translate into ongoing privacy management.

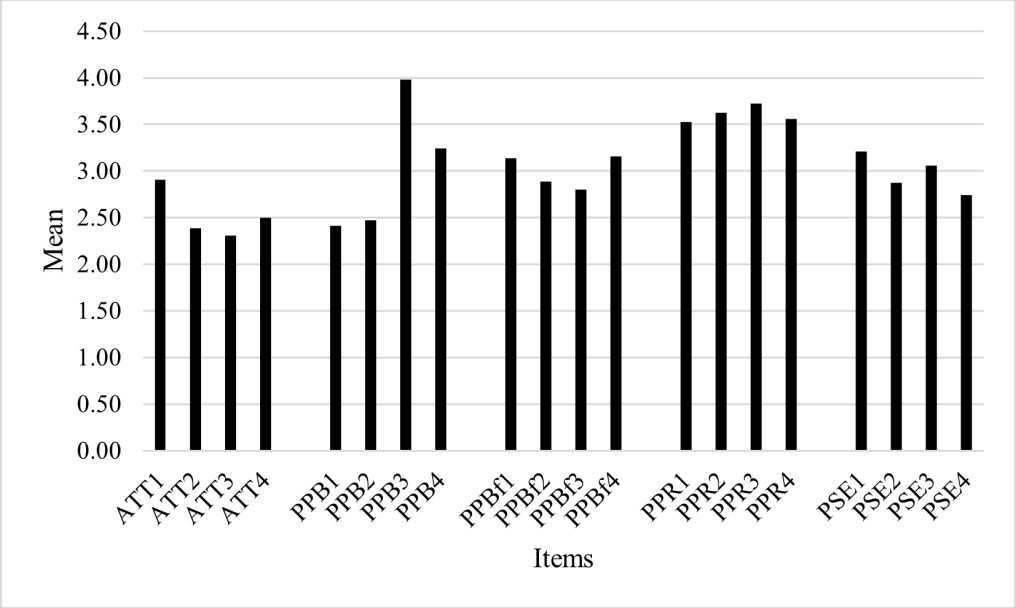

Figure 8. Distribution of means by survey item.

5.8.1 Distribution of means by survey item

This interpretation is reinforced by the subgroup heatmaps in Figures 9–12. Figure 9 shows that higher privacy-protective behavior is most clearly associated with stronger privacy self-efficacy and higher perceived risk, while transparency and trust differ little between the high- and low-protection groups. Figure 10 shows that heavy users report substantially greater perceived benefits, slightly higher transparency and trust, and somewhat lower perceived risk than light users, indicating that repeated use may normalize privacy trade-offs around convenience. Figure 11 indicates that participants aged 19–24 report higher perceived risk and somewhat stronger protective behavior, whereas those aged 16–18 report somewhat higher perceived benefits and transparency and trust. Figure 12 suggests additional variation by gender identity, with descriptively lower transparency and trust and lower self-efficacy among non-binary participants, although these results should be interpreted cautiously because of the small subgroup size. Taken together, these results show that youth privacy behavior is shaped not only by what risks they perceive, but by whether platforms make privacy understandable, actionable, and verifiable in everyday use.

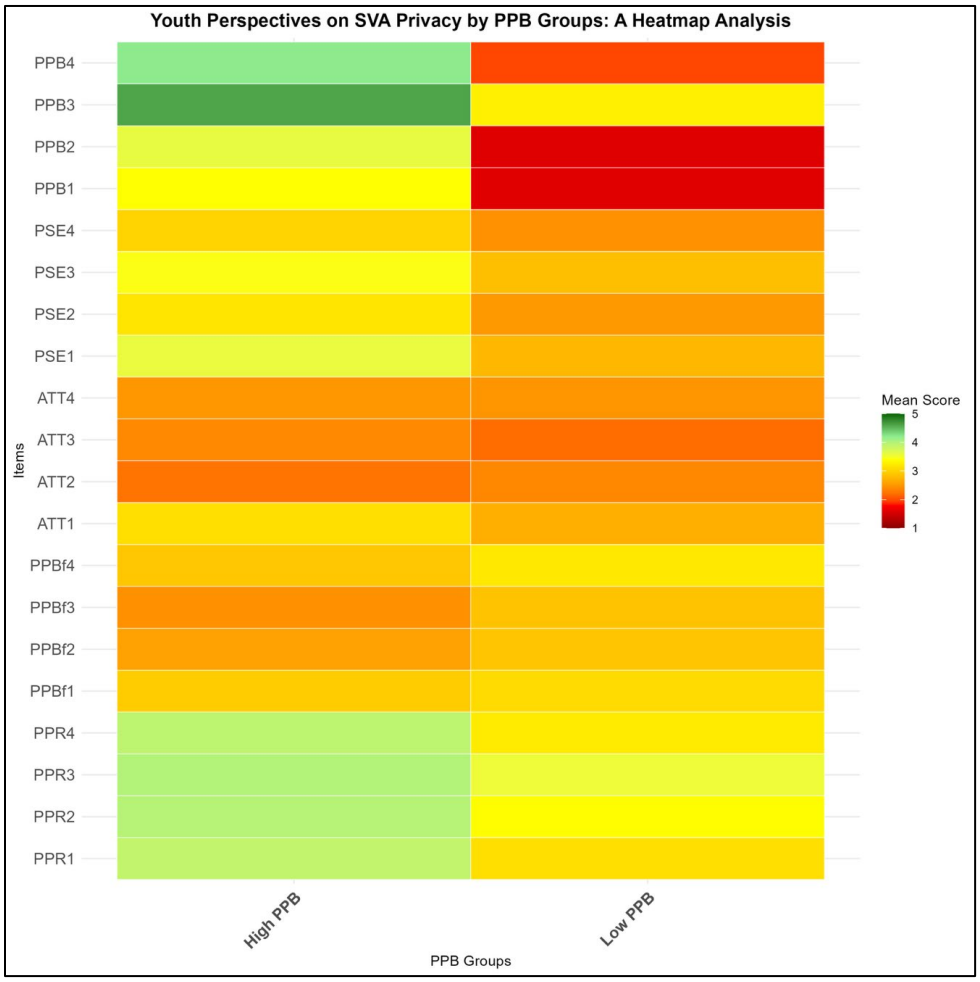

Figure 9. Heatmap of item-level means for high and low PPB.

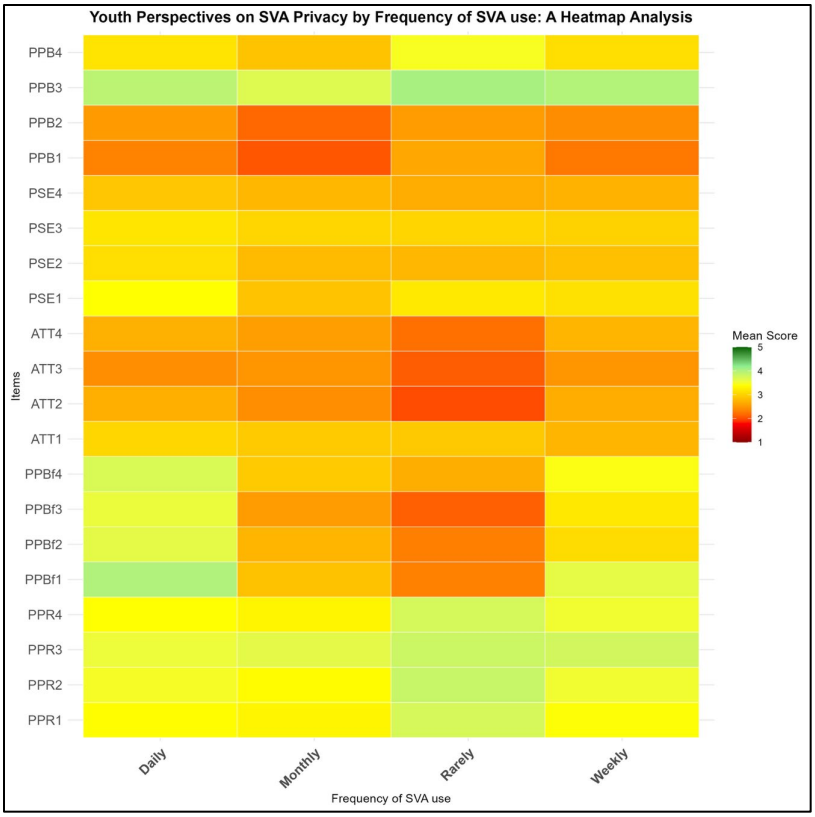

Figure 10. Heatmap of item-level means for SVA usage.

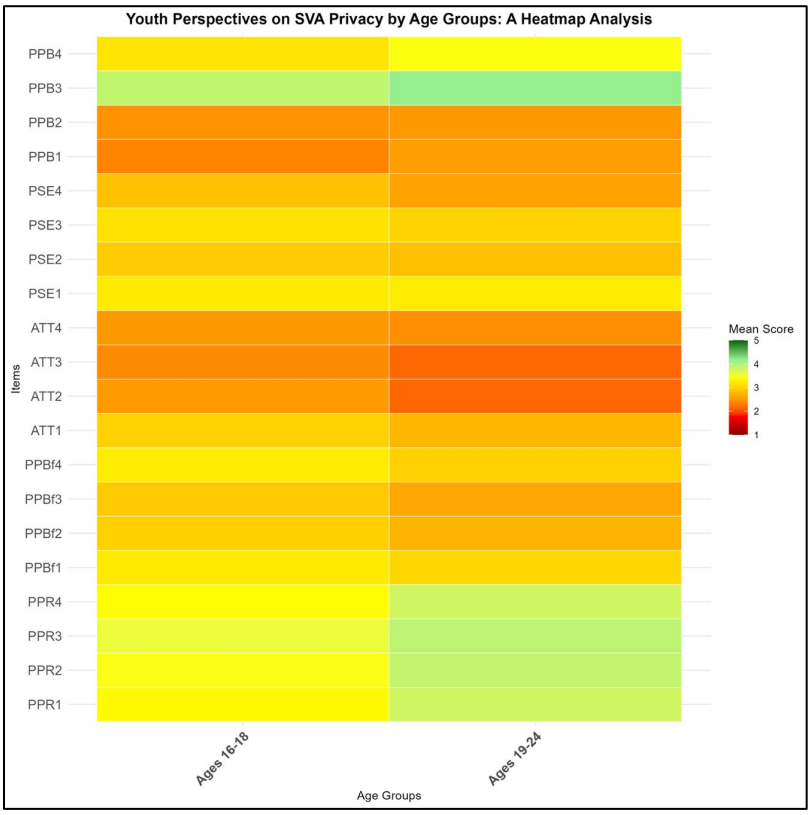

Figure 11. Heatmap of item-level means by age groups.

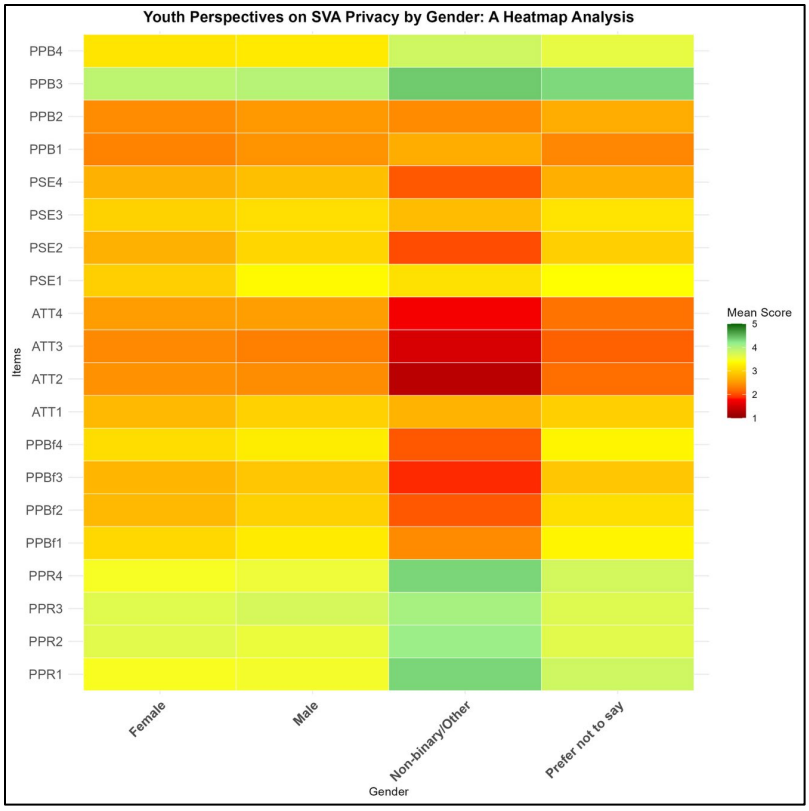

Figure 12. Heatmap of item-level means by gender identity.

5.8.2 Four actionable design principles (condensed)

Design principle	Condensed implementation intent
Principle 1: Make data flows visible and controllable	Move beyond subtle indicators; provide persistent, legible mic status and an explicit transparency channel. Integrate in-context controls (e.g., delete last interaction) to make protective actions part of normal use.
Principle 2: Build skills through guided interaction	Shift from static information to interactive learning: youth-friendly onboarding, step-by-step tutorials in settings, and scaffolds that convert procedural guidance into confidence and repeatable routines.
Principle 3: Reduce friction for protective actions	Lower the cost of key protective actions (review, delete, permission management) and provide clear feedback/receipts to support verification; minimize multi-hop navigation.
Principle 4: Adaptable autonomy across developmental stages	Support developmentally appropriate levels of autonomy and control, with age-appropriate language and graduated guidance that respects emerging independence while reducing hidden complexity.

Together, these principles translate the PEA-AI interpretation into practical design and governance guidance by emphasizing legibility, in-flow control, capability-building, and developmentally appropriate privacy support.

6. Privacy-by-Design Toolkit v1.0

The Privacy-by-Design Toolkit translates the multi-method evidence base into tiered, verifiable recommendations and practical implementation supports. Toolkit v0.1 (draft; dated February 4, 2026) was subsequently refined through the roundtable process into the final Toolkit v1.0 (dated March 6, 2026), which includes a traceability summary, evidence matrix, alignment mapping, gap register, tiered guidelines, role-based checklists, templates, and example workflows.

6.1 Toolkit structure and intended audiences

The toolkit specifies three primary audiences:

- Educators (classroom and workshop settings): practical activities and checklists that support youth privacy self-efficacy and informed device use.
- Developers/product teams (UX/privacy engineering): design and implementation requirements focusing on actionable transparency, verifiable controls, and protective defaults.
- Policymakers/regulators: minimum expectations and governance-oriented guidance grounded in youth evidence.

6.2 Evidence consolidation and prioritization

The toolkit consolidates evidence from the project papers and uses explicit prioritization criteria to assign tiered priorities (Tier 1–4). It also includes an alignment check against referenced frameworks and a gap register to document unresolved issues and areas needing further validation.

6.3 Tiered guidelines (summary)

Tier 1–2 guidelines are prioritized for immediate implementation. The table below summarizes the Tier 1–2 guideline intent (abbreviated).

Tier	Guideline focus (abbreviated)
Tier 1	G1 Listening/recording states explicit and controllable; G2 Retention/deletion time-bounded and provable; G3 Controls discoverable and youth-legible; G4 Consent meaningful and defaults protective.
Tier 2	G5 Standardized plain-language transparency summaries (“data labels”); G6 Bound third-party sharing and cross-service personalization with verifiable controls; G7 Granular access to recordings and derived data; G8 Youth-legible safeguards and actionable security cues;

Tier	Guideline focus (abbreviated)
	G9 Build privacy self-efficacy through repeatable routines.
Tier 3–4	Monitor and address as capacity allows; The toolkit includes additional concerns (Tier 3 and Tier 4) and explicit gap notes.

6.4 Verification and "prove it" checks

A distinctive feature of the toolkit is its verification orientation. Recommendations are paired with steps that help users confirm whether a privacy action took effect (e.g., deletion confirmation, state indicators, and observable receipts), reflecting audit and focus-group evidence that uncertainty about whether changes are applied reduces trust and self-efficacy.

6.5 Practical supports (checklists, templates, workflows)

The toolkit includes role-based checklists (educator, developer, policymaker), templates (e.g., a "Privacy Hub" task list; classroom worksheet; policymaker one-pager), and example workflows for common privacy tasks (disabling voice history, finding sharing/personalization settings, deleting stored commands, and verifying changes). These elements operationalize "reduce friction" and "build skills" principles in concrete, repeatable ways.

6.6 Roundtable refinement and version control

The toolkit document includes a change log and version-control appendix. It also documents that stakeholder roundtable feedback was treated as a validation source for implementation feasibility and item prioritization in v1.0.

7. Public engagement and knowledge mobilization

Knowledge mobilization combined academic dissemination, an online evidence hub (website and repository), and public-facing engagement activities. The project outputs were disseminated through in-person and hybrid/virtual conferences.

7.1 Academic dissemination (conferences and publications)

- In-person conferences: IEEE UEMCON 2025 (New York)[1] and IEEE CCWC 2026 (Las Vegas)[2].
- Hybrid/virtual conferences: IEEE ICAIC 2026 (Houston)[3] and IEEE ICAIIC 2026 (Tokyo)[4].
- Peer-reviewed and in-press outputs spanning literature synthesis, audit evaluation, qualitative themes, survey modelling, subgroup analyses, and integrative framing.

7.2 Project website and repository hub

The project website[5] and GitHub repository hub[6] were established to provide centralized access to publications, summaries, and outreach materials. The toolkit document notes that certain resources were uploaded on the project website and that v1.0 reflects triangulated evidence across sources.

The website and repository were created before October 31, 2025, and continue to be updated to reflect emerging findings and new outputs.

7.3 Outreach and public engagement activities

Outreach and public engagement activities were advanced through Medium, LinkedIn, X, and Meetup events, alongside the recorded webinar held on February 24, 2026, and the in-person workshop held on March 3, 2026, at Vancouver Island University in collaboration with the VIU Library. The Medium post[7] framed this dissemination strategy through three complementary video-based engagement pathways: a short explainer[8] introducing the privacy implications of smart voice assistants, a recorded workshop[9] emphasizing practical and participatory privacy learning, and a recorded webinar[10] presenting project evidence, tensions, and design directions. Together, these activities helped move the project from awareness to action by making youth privacy issues in smart voice technologies more understandable, participatory, and accessible to students, parents, educators, and broader community stakeholders.

[1] https://medium.com/@PrivaLab/covert-surveillance-in-smart-devices-reflections-from-ieee-uemcon-2025-3c19f7c7cef4
[2] https://medium.com/@PrivaLab/research-dissemination-at-ieee-ccwc-2026-7e82ab419451
[3] https://medium.com/@PrivaLab/conference-presentations-at-ieee-icaic-2026-evidence-informed-privacy-for-ai-and-smart-voice-4f37cdd87245
[4] https://medium.com/@PrivaLab/when-always-listening-meets-youth-agency-notes-from-icaiic-2026-912e71f20949
[5] https://csci-viu.github.io/privyouth-smart/
[6] https://github.com/csci-viu/privyouth-smart
[7] https://medium.com/@PrivaLab/from-awareness-to-action-three-ways-were-reframing-youth-privacy-in-smart-voice-tech-8f45a116bf4f
[8] https://youtu.be/DmX7-l73Iiw?si=MT3U7jTEv_WsLQTI
[9] https://www.youtube.com/watch?v=eAFxlqBjLNM
[10] https://youtu.be/ddOp-Cyx4_M?si=XrlVkzJ6rqnPdqnd

8. Challenges, risks, and mitigation

This section documents key risks and mitigation actions based on evidence emerging from the research outputs.

- Recruitment and multi-site ethics timing: delayed approvals from some external institutions created recruitment risk. Mitigation included an extension of data collection to November 20, 2025 and prioritization of approved school districts and institutional channels to reach participation targets.
- Stakeholder scheduling: scheduling conflicts affecting the stakeholder roundtable timeline. Mitigation included an established roundtable plan and structured feedback approach, with refinement incorporated into toolkit v1.0.
- Usability–compliance gap: audit evidence indicates that compliance scores do not guarantee youth-legible, verifiable controls. Mitigation is embedded in toolkit guidelines emphasizing feedback, reversibility, task-based discoverability, and explicit state indicators.
- Evidence limitations and future-facing gaps: the toolkit explicitly flags unresolved issues (e.g., third-party deletion limits, cross-border handling and the risk of introducing youth age-identification mechanisms that could create new harms) to avoid overstating what is currently supported by evidence.

9. Conclusion and sustainability

This project delivers a coherent evidence-to-practice pathway for youth privacy in smart voice assistant ecosystems. Across literature synthesis, platform audit evidence, focus-group themes, and survey modelling, results converge on a consistent implication: youth privacy protection is strengthened when transparency is actionable, controls are usable and verifiable, and privacy self-efficacy is actively supported.

Toolkit v1.0 operationalizes these insights through tiered, verifiable guidance and role-based implementation supports. In line with the proposal's engagement objectives, the project's website/repository hub and public engagement activities provide pathways for sustained uptake.

Forward-looking next steps grounded in the evidence base include expanding validation across additional smart-device contexts, increasing inclusive sampling for subgroup stability, and evaluating the toolkit's adoption and effectiveness through educator- and developer-facing implementations and iterative stakeholder feedback cycles.

Acknowledgment

This project has been funded by the Office of the Privacy Commissioner of Canada (OPC); the views expressed herein are those of the authors and do not necessarily reflect those of the OPC.

> *For any queries or additional information regarding this report, kindly contact the Principal Investigator. The research team appreciates the continued support of the OPC, VIU, and all participating school districts, universities, volunteers, and stakeholders who made this project possible.*

References (including project output)

[1] A. Shouli, Y. Bobkova, and A. K. Shrestha, “Covert surveillance in smart devices: A SCOUR framework analysis of youth privacy implications,” in *2025 IEEE 16th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, New York, NY, USA: IEEE, Oct. 2025, pp. 0124–0133. doi: 10.1109/UEMCON67449.2025.11267573.

[2] M. Campbell, M. S. Al Jasem, and A. K. Shrestha, “Toward youth-centered privacy-by-design in smart devices: A systematic review,” in *2026 IEEE 16th Annual Computing and Communication Workshop and Conference (CCWC)*, IEEE, Jan. 2026, pp. 0444–0453. doi: 10.1109/CCWC67433.2026.11393846.

[3] T. De Clark, Y. Bobkova, and A. K. Shrestha, “Balancing usability and compliance in AI smart devices: A privacy-by-design audit of Google Home, Alexa, and Siri,” in *2026 IEEE 16th Annual Computing and Communication Workshop and Conference (CCWC)*, Las Vegas, USA: IEEE, Jan. 2026, pp. 0378–0387. doi: 10.1109/CCWC67433.2026.11393731.

[4] M. Campbell, T. De Clark, M. S. Al Jasem, S. Joshi, and A. K. Shrestha, “Convenience vs. control: A qualitative study of youth privacy with smart voice assistants,” in *2026 IEEE 16th Annual Computing and Communication Workshop and Conference (CCWC)*, Las Vegas, USA: IEEE, Jan. 2026, pp. 0369–0377. doi: 10.1109/CCWC67433.2026.11393734.

[5] M. Campbell and A. K. Shrestha, “Privacy by voice: Modeling youth privacy-protective behavior in smart voice assistants,” in *2026 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, Tokyo: IEEE, Feb. 2026, “in press”.

[6] Y. Bobkova, M. Campbell, T. De Clark, and A. K. Shrestha, “Age-differentiated pathways to privacy protection in smart voice assistants: A multigroup PLS-SEM study of youth,” in *2026 IEEE 5th International Conference on AI in Cybersecurity (ICAIC)*, Houston, Texas, USA: IEEE, Feb. 2026, pp. 1–8. doi: 10.1109/ICAIC67076.2026.11395739.

[7] M. Campbell, Y. Bobkova, and A. K. Shrestha, “Gender-based heterogeneity in youth privacy-protective behavior for smart voice assistants: Evidence from multigroup PLS-SEM,” in *39th Annual Canadian Conference on Electrical and Computer Engineering (CCECE 2026)*, Montreal, Canada: IEEE, 2026, “in review”.

[8] A. K. Shrestha, “Privacy profiles of youth smart voice assistant users: A cluster analysis of risk, benefits, trust, and behavior,” in *Preprint Version*, 2026, “draft”.

[9] M. Sheikho Al Jasem and A. K. Shrestha, “Negotiating privacy with smart voice assistants: Risk-benefit and control-acceptance tensions,” in *2026 10th International Conference on Cryptography, Security and Privacy (CSP 2026)*, Sapporo: IEEE Xplore, 2026, “in review”.

[10] M. Campbell, Y. Bobkova, and A. K. Shrestha, “From framework to practice: Youth negotiations of privacy with smart voice assistants through the PEA-AI lens,” *ACM Journal of Responsible Computing*, 2026, “in review”.

[11] A. K. Shrestha, *Privacy-by-design toolkit v1.0: Youth-centred guidelines for smart voice assistants and voice-enabled smart devices*. Nanaimo: Vancouver Island University, 2026. [Online]. Available: https://csci-viu.github.io/privyouth-smart/ToolKit/Final_Privacy-by-Design_Toolkit.pdf

[12] V. Steeves, “Privacy, sociality and the failure of regulation: Lessons learned from young Canadians’ online experiences,” in *Social Dimensions of Privacy: Interdisciplinary Perspectives*, Cambridge University Press, 2015, ch. 13, pp. 244–260. doi: 10.1017/CBO9781107280557.014.

[13] E. A. Vogels, R. Gelles-Watnick, and N. Massarat, “Teens, social media and technology 2022: TikTok has established itself as one of the top online platforms for U.S. teens, while the share of teens who use Facebook has fallen sharply,” 2022. Accessed: Oct. 01, 2025. [Online]. Available: https://www.jstor.org/stable/resrep63507

[14] M. Tabassum *et al.*, “Investigating users’ preferences and expectations for always-listening voice assistants,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 3, no. 4, pp. 1–23, 2019, doi: 10.1145/3369807.

[15] J. Lau, B. Zimmerman, and F. Schaub, “Alexa, are you listening?: Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers,” *Proc. ACM Hum. Comput. Interact.*, vol. 2, no. CSCW, pp. 1–31, Nov. 2018, doi: 10.1145/3274371.

[16] A. Shouli, A. Barthwal, M. Campbell, and A. K. Shrestha, “Unpacking youth privacy management in AI systems: A privacy calculus model analysis,” *IEEE Access*, vol. 13, pp. 115780–115803, 2025, doi: 10.1109/ACCESS.2025.3585635.

[17] S. Solera-Cotanilla, M. Vega-Barbas, J. Pérez, G. López, J. Matanza, and M. Álvarez-Campana, “Security and privacy analysis of youth-oriented connected devices,” *Sensors*, vol. 22, no. 11, May 2022, doi: 10.3390/s22113967.

[18] A. Barthwal, M. Campbell, and A. K. Shrestha, “Privacy ethics alignment in AI: A stakeholder-centric framework for ethical AI,” *Systems*, vol. 13, no. 6, p. 455, Jun. 2025, doi: 10.3390/SYSTEMS13060455.

[19] N. K. Malhotra, S. S. Kim, and J. Agarwal, “Internet users’ information privacy concerns (IUIPC): The construct, the scale, and a causal model,” *Information Systems Research*, vol. 15, no. 4, pp. 336–355, 2004, doi: 10.1287/ISRE.1040.0032.

[20] H. J. Smith, S. J. Milberg, and S. J. Burke, “Information privacy: Measuring individuals’ concerns about organizational practices,” *MIS Q.*, vol. 20, no. 2, pp. 167–195, 1996, doi: 10.2307/249477.

[21] T. Dinev and P. Hart, “An Extended Privacy Calculus Model for E-Commerce Transactions,” *Information Systems Research*, vol. 17, no. 1, pp. 61–80, Mar. 2006, doi: 10.1287/ISRE.1060.0080.

[22] A. Acquisti and J. Grossklags, “Privacy and rationality in individual decision making,” *IEEE Secur. Priv.*, vol. 3, no. 1, pp. 26–33, Jan. 2005, doi: 10.1109/MSP.2005.22.

[23] A. Beldad, M. De Jong, and M. Steehouder, “How shall I trust the faceless and the intangible? A literature review on the antecedents of online trust,” *Comput. Human Behav.*, vol. 26, no. 5, pp. 857–869, Sep. 2010, doi: 10.1016/J.CHB.2010.03.013.

[24] H. Jeff Smith, T. Dinev, and H. Xu, “Information privacy research: An interdisciplinary review,” *MIS Q.*, vol. 35, no. 4, pp. 989–1015, 2011, doi: 10.2307/41409970.

[25] A. Bandura, “Self-efficacy: Toward a unifying theory of behavioral change,” *Psychol. Rev.*, vol. 84, no. 2, pp. 191–215, Mar. 1977, doi: 10.1037/0033-295X.84.2.191.

[26] G. R. Milne and M. J. Culnan, “Strategies for reducing online privacy risks: Why consumers read (or don’t read) online privacy notices,” *Journal of Interactive Marketing*, vol. 18, no. 3, pp. 15–29, Jan. 2004, doi: 10.1002/DIR.20009.

[27] P. A. Norberg, D. R. Horne, and D. A. Horne, “The privacy paradox: Personal information disclosure intentions versus behaviors,” *Journal of Consumer Affairs*, vol. 41, no. 1, pp. 100–126, Jun. 2007, doi: 10.1111/j.1745-6606.2006.00070.x.

[28] S. Youn, “Determinants of online privacy concern and its influence on privacy protection behaviors among young adolescents,” *Journal of Consumer Affairs*, vol. 43, no. 3, pp. 389–418, Sep. 2009, doi: 10.1111/j.1745-6606.2009.01146.x.

[29] M. J. Page *et al.*, “The PRISMA 2020 statement: an updated guideline for reporting systematic reviews,” *BMJ*, vol. 372, Mar. 2021, doi: 10.1136/BMJ.N71.

[30] J. Nielsen and R. Molich, “Heuristic evaluation of user interfaces,” in *Proceedings of the SIGCHI conference on Human factors in computing systems Empowering people - CHI ’90*, New York, New York, USA: ACM Press, 1990, pp. 249–256. doi: 10.1145/97243.97281.

[31] V. Bellotti and A. Sellen, “Design for privacy in ubiquitous computing environments,” in *Proceedings of the Third European Conference on Computer-Supported Cooperative Work*

13–17 September 1993, Milan, Italy ECSCW '93, Dordrecht: Springer Netherlands, 1993, pp. 77–92. doi: 10.1007/978-94-011-2094-4_6.

[32] E. Luger and T. Rodden, "An informed view on consent for UbiComp," in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, New York, NY, USA: ACM, Sep. 2013, pp. 529–538. doi: 10.1145/2493432.2493446.

[33] Office of the Privacy Commissioner of Canada, "Guidelines for obtaining meaningful consent," Aug. 2025. Accessed: Oct. 11, 2025. [Online]. Available: https://www.priv.gc.ca/en/privacy-topics/collecting-personal-information/consent/gl_omc_201805/

[34] Office of the Privacy Commissioner of Canada, "PIPEDA fair information principles," May 2025. Accessed: Oct. 27, 2025. [Online]. Available: https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/p_principle/

[35] J. Nielsen, "Why you only need to test with 5 users," Nielsen Norman Group. Accessed: Oct. 11, 2025. [Online]. Available: https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/