From Framework to Practice

Youth Privacy with Smart Voice Assistants (SVA)
From systematic review to mixed-methods evidence

Privacy Aware AI Research

Dr. Ajay Shrestha

Office of the Privacy Commissioner of Canada (OPC) – Contributions Program

Agenda

Roadmap

- ❑ Evidence base: youth-centered privacy-by-design in smart devices (systematic review)
- ❑ Youth voices: convenience vs. control (focus groups)
- ❑ Why SVAs create distinctive privacy tensions for youth
- ❑ PEA-AI lens and the five measured constructs
- ❑ Key findings: overall patterns + subgroup differences
- ❑ From perceptions to action: modeling privacy-protective behavior (PLS-SEM)
- ❑ Design and governance implications, limitations, and next steps

Evidence base: youth-centered privacy-by-design in smart devices

PRISMA-guided review (n = 122) across home, school, and health contexts

Key messages

1) Exposure is ambient and multi-context

Smart devices expand beyond clicks to audio/video, location, and biometrics—often in shared spaces.

2) Evidence is technically rich, but uneven

Most studies emphasize technical mechanisms; policy and education are comparatively thin.

3) Scaling fails without alignment

Privacy-by-design must combine enforceable baselines, usable controls, and sustained privacy literacy.

4) Implication for SVAs

Always-listening assistants are a stress-test where usability and governance determine real protection.

Thematic distribution (n = 122)

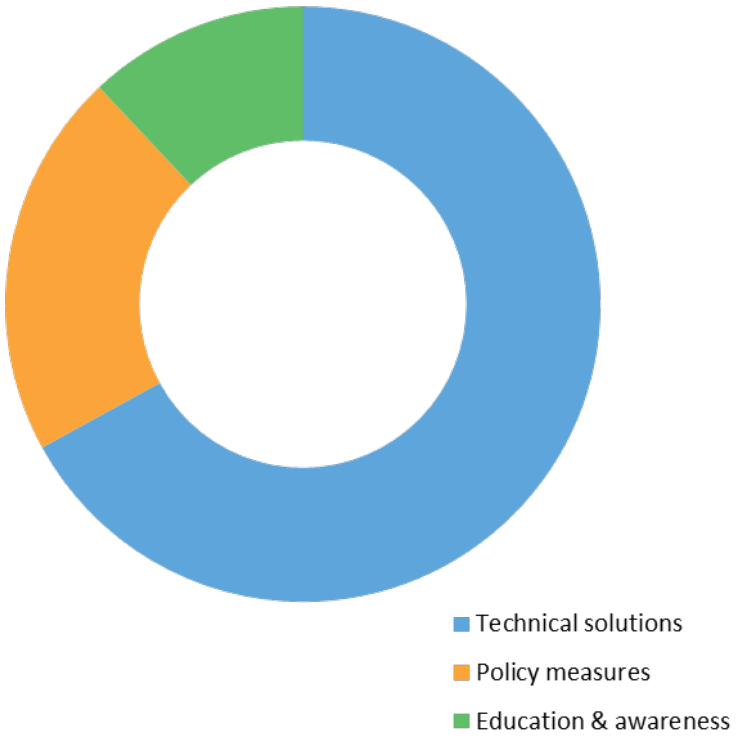

Technical: 67% • Policy: 21% • Education: 12%

Youth voices: convenience vs. control in SVAs

Focus groups (n = 26; ages 16–24) show where privacy control breaks down

Recurring themes

1) Benefits keep SVAs in daily routines

Micro-tasks, hands-busy contexts, and entertainment make utility immediate.

2) Always-listening concerns are persistent

Ambient listening, unclear retention, and perceived cross-app inferences undermine comfort.

3) Transparency friction drives a ‘wayfinding gap’

Policies are unreadable; settings are scattered; deletion feels unverifiable.

4) Protective behaviors are workaround-based

Mute/unplug, permission refusal, uninstall, and non-adoption often replace nuanced settings.

Why smart voice assistants are a privacy stress-test

SVA context amplifies privacy negotiation

- Always-on listening (wake-word architecture) increases ambient surveillance concerns
- Shared household devices concentrate control in a primary account holder, reducing youth agency
- Data flows and retention are often opaque, weakening transparency and trust
- Convenience remains high—supporting a persistent “privacy paradox” in practice

Research questions (RQ1–RQ3): patterns • negotiation • implications

PEA-AI lens: privacy as negotiation (not a one-time choice)

Five dimensions; this study operationalizes key tensions for youth

PRB — **Perceived Risks & Benefits**
Trade-offs shaping the privacy calculus

TT — **Transparency & Trust**
Openness that supports confidence and consent

DOC — **Data Ownership & Control**
Ability to act on privacy preferences

EA — **Education & Awareness**
Knowledge and skills to navigate settings

PDS — **Parental Data Sharing**
Household mediation vs youth autonomy

Study focus: tensions made measurable

We measure five constructs aligned to the youth perspective:

- PPR — perceived privacy risk
- PPBf — perceived privacy benefits
- ATT — algorithmic transparency & trust
- PSE — privacy self-efficacy
- PPB — privacy-protective behavior

Core takeaway: privacy is negotiated through risk–benefit, transparency–trust, and control–action trade-offs.

Study overview

Phase 3 survey study (a different study) within a multi-study program

Three-phase program

Phase 1

PEA-AI stakeholder study

Phase 2

Youth SVA focus groups

Phase 3 (study)

Youth SVA survey + mixed-methods integration

Sample and procedure

Survey (Microsoft Forms), 5-point Likert items.

- N = 469 Canadian participants, ages 16–24
- Eligibility: ≥1 SVA use in prior month
- 20 items across 5 constructs; subgroup comparisons
- Ethics approval: VIU-REB #103597

Key demographics

Gender (n):

Male 241 • Female 174 • Non-binary/Other 15 • Prefer not 35

SVA use frequency (n):

Rarely 190 • Daily 126 • Weekly 113 • Monthly 38

Mean age 18.65 (SD 2.30)

Measurement instrument

20 items across five constructs aligned to PEA-AI dimensions

Constructs and what they capture

Construct	Definition (SVA context)	PEA-AI alignment
PPR	Concerns about recording, retention, and unauthorized access	PRB
PPBf	Convenience / personalization benefits that offset concern	PRB
ATT	Perceived transparency of data practices and trust in companies/algorithms	TT
PSE	Confidence and capability to manage settings and reduce exposure	DOC + EA
PPB	Reported privacy-protective behaviors (permissions, deletion, refusals)	DOC + EA (+TT as antecedent)

Note: PPBf denotes benefits (to distinguish from PPB behaviors).

Measurement quality

Reliability and convergent validity checks (PLS-SEM)

Reliability (all ≥ 0.70)

- Cronbach’s alpha: 0.714–0.880 across constructs
- Composite reliability (rho_c): 0.818–0.917
- Dillon-Goldstein’s rho (rho_a): 0.712–0.907

Convergent validity (AVE)

All AVE values exceed 0.50:

Construct	AVE
ATT	0.529
PPB	0.539
PPBf	0.711
PPR	0.734
PSE	0.626

Interpretation

The instrument is suitable for descriptive profiling and comparison across youth subgroups.

Key results (construct level)

Mean scores on a 1–5 Likert scale (N = 469)

Overall patterns

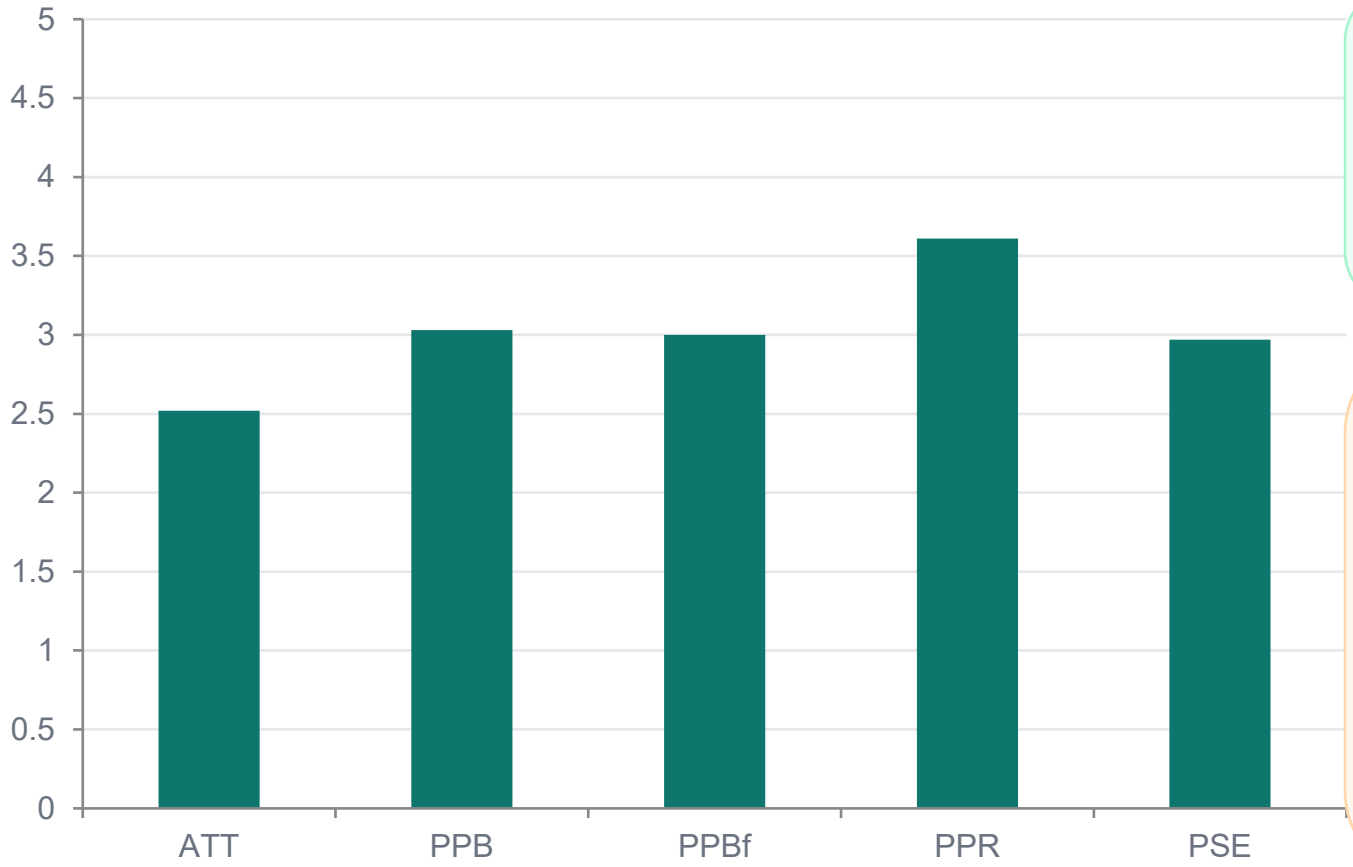

Highest: Risk perception (PPR)

PPR M = 3.61 (moderate–high concern)

Lowest: Transparency & trust (ATT)

ATT M = 2.52 (below midpoint)

PPBf, PPB, and PSE sit near the midpoint → a persistent privacy paradox and an efficacy gap.

Item-level patterns

Where concerns, trust, capability, and behaviors diverge

Figure 1: item means

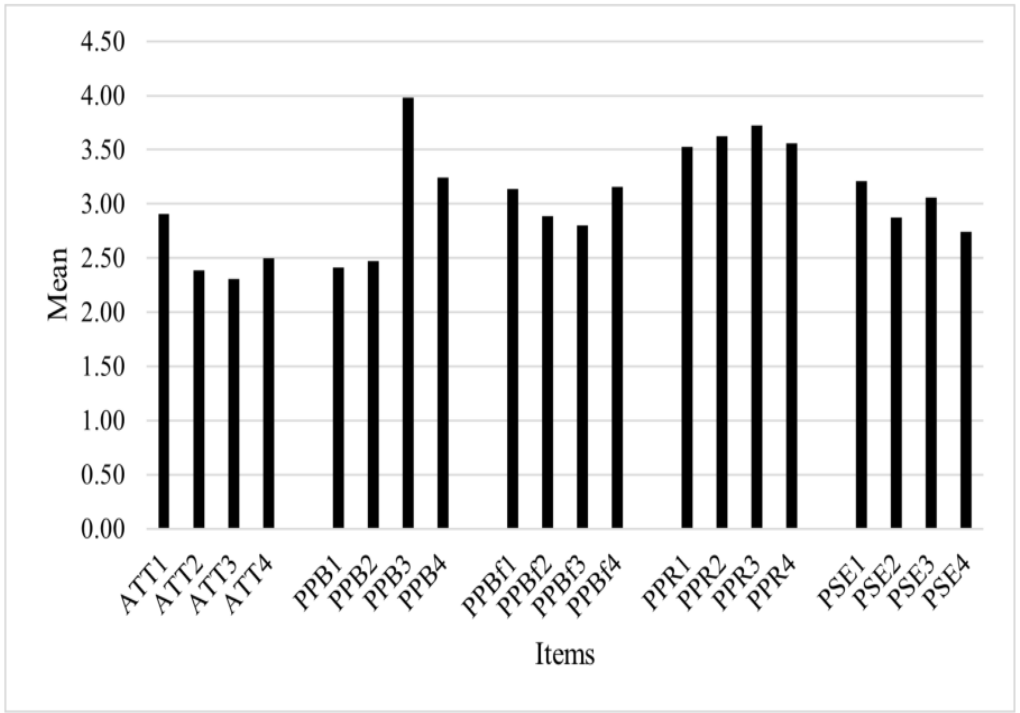

Figure 1: Distribution of means by survey item.

Highlights (selected items)

- Highest PPB: refusing features (PPB3 M = 3.98)
- Lowest PPB: reviewing permissions (PPB1 M = 2.41) and deleting history (PPB2 M = 2.47)
- Top risk: unauthorized access concern (PPR3 M = 3.72)
- Lowest transparency: “upfront about processing” (ATT3 M = 2.30)
- Lowest efficacy: managing overall risks (PSE4 M = 2.74)

Pattern: avoidance is more common than active, ongoing privacy management.

Subgroup finding 1: protective behavior profiles

High vs low PPB (top vs bottom quartile)

Item-level heatmap (means)

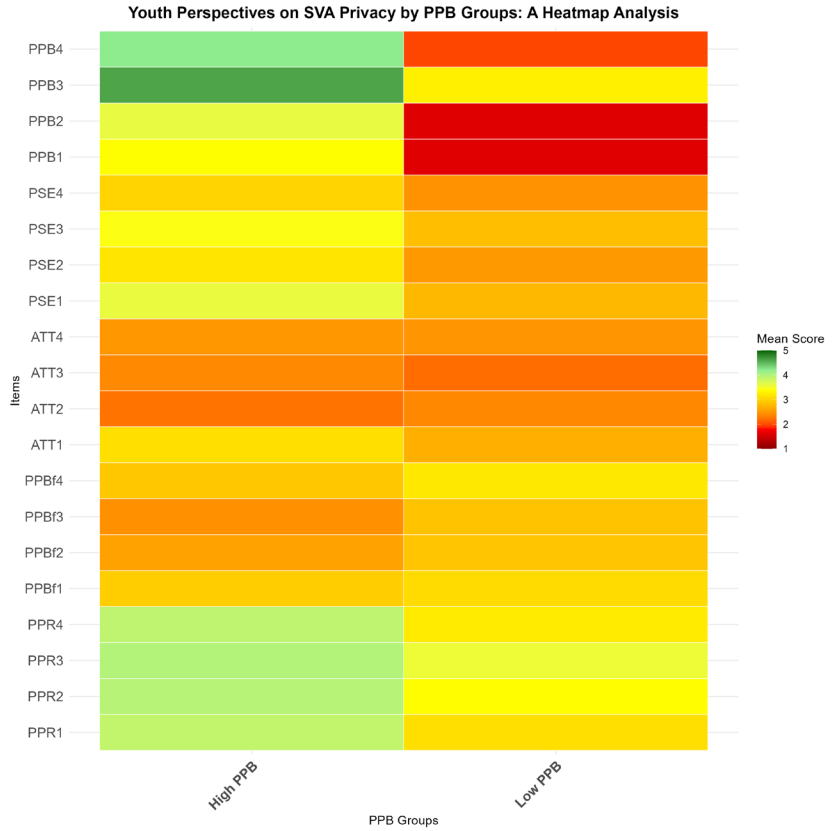

Figure 2: Heatmap of item-level means for high and low PPB.

Key differences

Substantial (≥ 0.50):

- PSE: +0.67
- PPR: +0.63

Modest:

- PPBf: −0.28
- ATT: n.s.

Interpretation: protective action aligns most strongly with
(1) higher risk perception and
(2) higher self-efficacy.

Subgroup finding 2: heavy vs light SVA users

Daily/weekly vs monthly/rare use

Item-level heatmap (means)

Figure 3: Heatmap of item-level means for SVA usage.

Key differences

Large (practically meaningful):

- PPBf: +1.11 (benefits)

Smaller (incremental):

- ATT: +0.34
- PPR: −0.22
- PPB, PSE: n.s.

Interpretation: frequent use is associated with greater perceived utility and slightly lower perceived risk.

Subgroup finding 3: age and gender patterns

Differences are informative, but some subgroups are small

Age (16–18 vs 19–24)

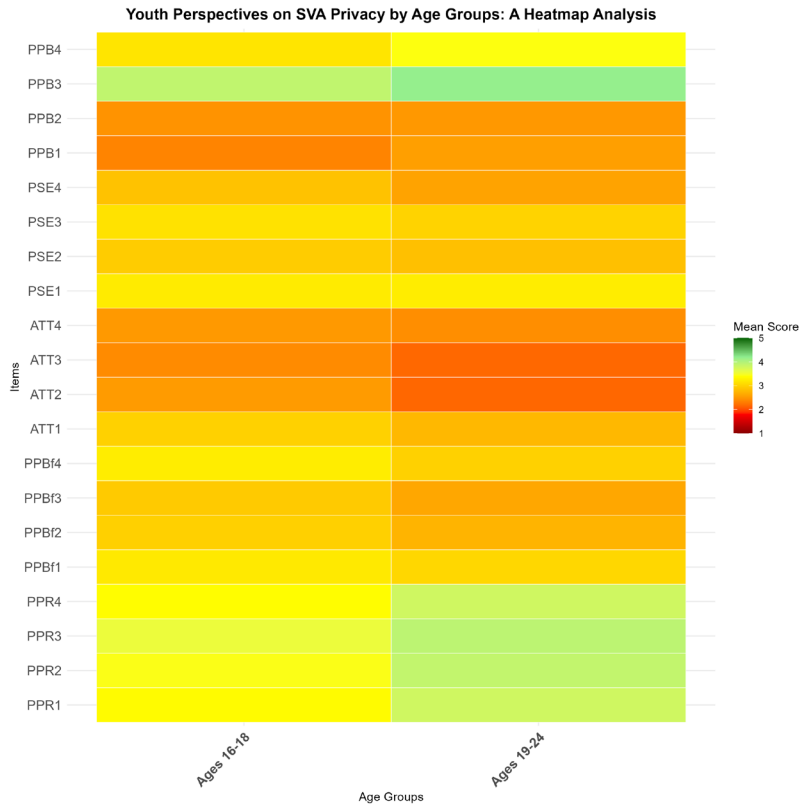

Figure 4: Heatmap of item-level means by age groups.

Largest gap: PPR (+0.41 for ages 19–24)

Gender identity (4 groups)

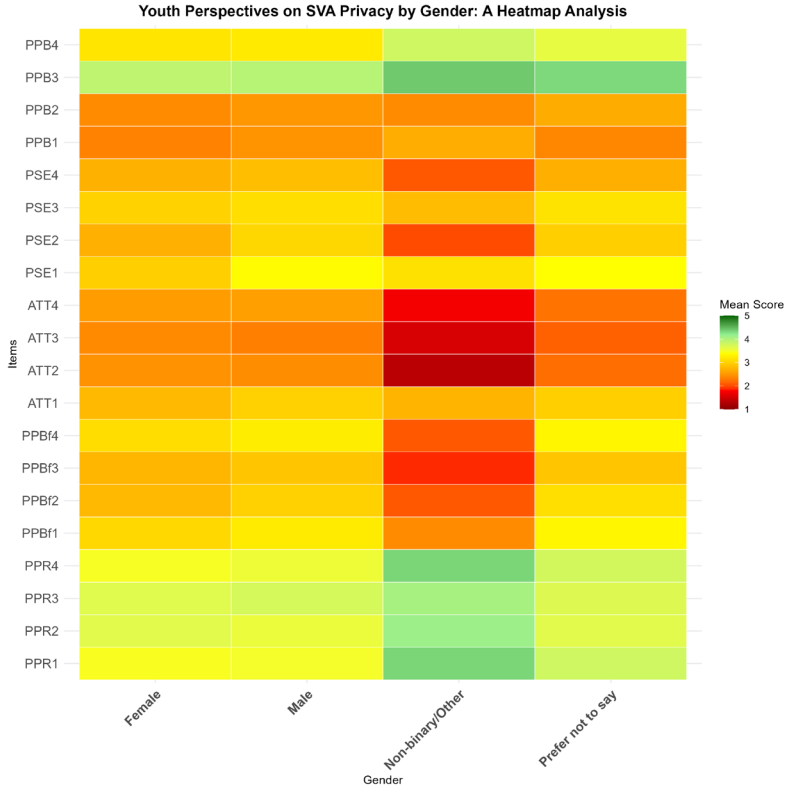

Figure 5: Heatmap of item-level means by gender identity.

Caution: non-binary/other subgroup n = 15; PPR test n.s. (p≈0.094).

From perceptions to action: modeling privacy-protective behavior

Survey (N = 469) PLS-SEM: risk/benefit calculus + transparency/trust + self-efficacy

Key results

1) Self-efficacy is the strongest driver

PSE → PPB $\beta = 0.373$ ($p < 0.001$): confidence translates into action.

2) Risk increases protection—but not through efficacy

PPR → PPB $\beta = 0.343$ ($p < 0.001$); PPR → PSE is not significant.

3) Benefits can dampen protection

PPBf → PPB $\beta = -0.130$ ($p < 0.1$), with a small indirect uplift via PSE.

4) Transparency/trust matters via efficacy

ATT → PSE $\beta = 0.434$ ($p < 0.001$); ATT → PPB not significant; $R^2 \approx 0.24$.

PLS-SEM structural model

Standardized paths and explained variance (N = 469)

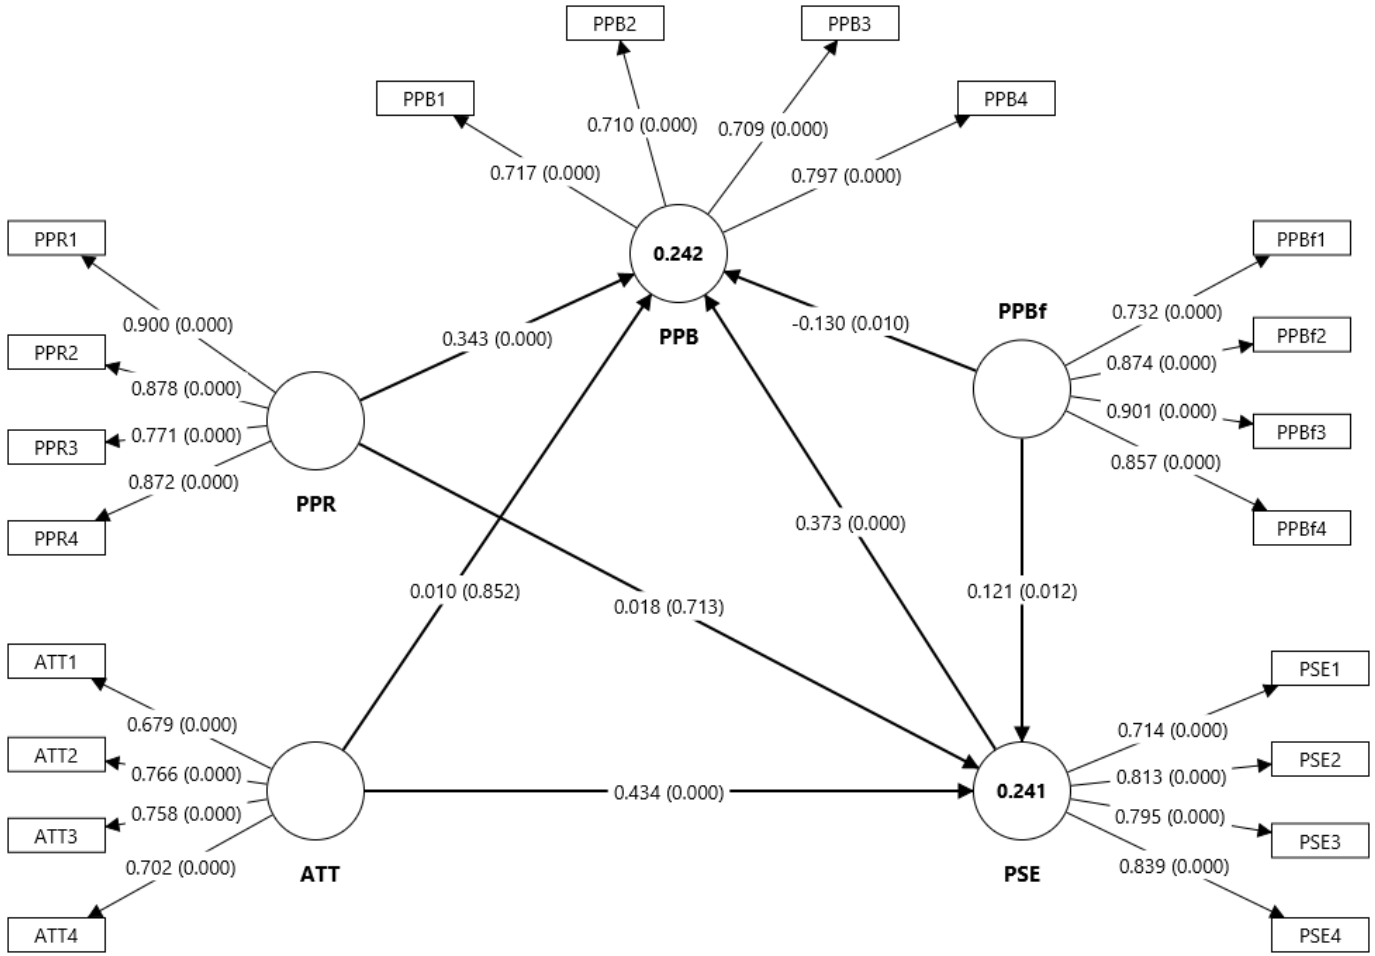

Note: β values are standardized path coefficients; significance as reported in the paper.

Design implications for SVA platforms

Four actionable principles grounded in item-level frictions

Principles

1) Make data flows visible and controllable

Persistent transparency channels (screen where available; companion app; standardized device cues) + in-flow deletion/review actions.

2) Build skills through guided interaction

Youth-friendly onboarding and just-in-time explanations that translate knowledge into action.

3) Reduce friction for protective actions

One-touch privacy shortcuts, granular consent, and individual profiles on shared devices.

4) Adaptable autonomy across developmental stages

Scaffolded controls and inclusive threat models; treat subgroup differences as design signals.

Governance and education implications

Multi-stakeholder levers mapped to PEA-AI dimensions

Implication matrix

Stakeholder lever	PRB (risk–benefit)	TT (transparency–trust)	DOC (control)	EA (agency/skills)
Platform design	Make benefits explicit (avoid “dark patterns”)	Layered, in-context explanations	In-flow controls (delete last interaction; one-touch privacy)	Guided walkthroughs; just-in-time prompts
Education	Privacy calculus exercises using SVA scenarios	Teach “policy fatigue” and how to verify claims	Practice logs/permissions/auto-delete settings	Skills checklist aligned to PSE items
Regulation / governance	Youth-protective defaults (e.g., auto-delete)	Standards for layered notices + auditability	Requirements for usable, accessible controls	Accountability for youth-facing UX evidence

Limitations, next steps, and closing takeaways

Limitations

- Canadian online sample may limit generalizability
- Small subgroup sizes (e.g., non-binary n = 15) affect stability of comparisons
- Multiple comparisons: interpret as exploratory alongside practical magnitude
- Self-reported PPB may diverge from observed settings behavior
- Instrument is SVA-focused; portability to other conversational AI requires validation

Future work

- Longitudinal tracking and intervention studies (improve PSE and PPB)
- Cross-platform and cross-jurisdiction comparisons
- Triangulate surveys with ethically designed observational/log-based measures

Closing takeaways

1) Risk remains salient; transparency/trust is weakest.
2) Protective action hinges on self-efficacy.
3) Design for negotiation: legibility + in-flow control + capability building.

Q&A