*Applied Human Factors and Ergonomics (AHFE2025),* Vol. 199, 2025, 1203–1213
https://doi.org/10.54941/ahfe1006931

# Using Generative AI Personas to Study Consent in Educational Data Use: Views of Parents and Elementary School Students

**Mitsuhiro Takasaki and Tetsuro Kakeshita**

Faculty of Science and Engineering, Saga University, Honjo, Saga 840-8502, Japan

## ABSTRACT

The digital transformation of education requires new consent models that enable the use of diverse learning data while safeguarding privacy. This study employed generative AI to create 50 pairs of parent–child personas and simulated their consent decisions for 40 types of learning data under systematically varied conditions. Descriptive statistics showed that Conditional Allow dominated across categories, while refusal clustered around surveillance-like data (e.g., IP addresses, logs, geolocation). A k-means clustering of purposes and justifications identified six distinctive clusters: educational improvement, direct learning support, regional policy and safety, special needs education, privacy distrust toward surveillance-related data, and privacy distrust toward third-party sharing. In addition, consolidating data types into six categories revealed intergenerational differences: children were more permissive toward evaluation and psychological data, while parents were cautious, particularly regarding family background and administrative information. These findings highlight that one-size-fits-all consent approaches are inadequate. Persona-based simulations offer an exploratory method for identifying patterns and testing safeguards, supporting the design of multi-layered, category-sensitive consent frameworks for educational data governance.

**Keywords:** Educational data, Generative AI, Persona simulation, Consent formation, Privacy, Parent–child differences

## INTRODUCTION

The digital transformation of education (often termed Educational DX) has created high expectations for improving learning processes, enhancing instructional quality, and supporting evidence-based decision-making. Central to this transformation is the effective use of diverse forms of learning data – ranging from attendance and academic performance records to health information, behavioral logs, and digital activity traces. In Japan, for example, the Digital Agency emphasizes that leveraging such data can maximize each child's potential through personalized learning and enable evidence-based improvements in teaching and administration (Digital Agency, 2025). International frameworks similarly underscore data-driven insights as a pillar of educational innovation; UNESCO's global guidelines

© 2025. Published by AHFE Open Access. All rights reserved. 

note that using data and evidence is central to informing and improving educational experiences in a digital age (UNESCO, 2023). These approaches highlight how robust data analytics can turn intuition-driven decisions into objective, actionable feedback for educators, thereby enriching learning outcomes and organizational effectiveness (Digital Agency, 2025).

However, the use of learning data inevitably raises significant challenges concerning personal information protection and privacy. In Japan and many other contexts, data use must comply with strict personal data protection laws, which require explicit consent from data subjects for collection or sharing of identifiable information. When minors are involved, parental consent becomes essential, as reflected in regulations worldwide—for example, the EU's General Data Protection Regulation (GDPR) mandates that processing a child's data is lawful only with consent from a parent or guardian if the child is under age 16 (European Union, 2016). In practice, this means schools and educational authorities cannot simply tap into student data without obtaining clear permission under defined conditions. While the potential benefits of learning data utilization are widely recognized, both parents and students often hesitate to give consent when data about themselves or their family is requested. Surveys indicate that parents worry student information could be misused—for example, repurposed for marketing or shared without proper safeguards (Future of Privacy Forum, 2020). Likewise, many students are cautious; studies have found that a lack of trust in how institutions will handle their data reduces their willingness to share, whereas greater confidence that data will be used appropriately correlates with higher consent rates (Liu & Khalil, 2023).

Moreover, as analytical techniques become increasingly advanced, additional conditions must be considered—such as data anonymization, controlled third-party sharing, and limited data retention—which further complicate the consent process. Simply stripping names from records is often insufficient to fully protect privacy, especially when datasets can be cross-referenced or re-identified (Ohm, 2010). Therefore, ensuring meaningful protection requires robust measures (e.g., rigorous de-identification standards, encryption, clear rules on data retention and access) (UNESCO, 2023). Each of these safeguards needs to be communicated to and agreed by the data subjects (or their parents), making the consent agreements more complex. In short, educational institutions face a delicate task: they must balance the drive to utilize data for improvement with the obligation to uphold privacy and comply with regulations (Reidenberg & Schaub, 2018). This tension has been widely acknowledged as a major challenge that can hinder the expansion of learning analytics and data-driven education if not managed properly (Liu & Khalil, 2023).

These difficulties highlight the urgent need for reliable consent-formation models that enable the benefits of educational data use while preserving privacy. Researchers argue that demonstrating the efficacy of data-driven learning tools and embedding strong privacy protections and accountability are both essential to align with societal values (Reidenberg & Schaub, 2018). Recent advances in generative artificial intelligence (AI) offer a novel opportunity to explore this issue from a new angle. In particular, large language models now make it possible to simulate diverse personas—in other words, have an AI role-play as students or parents with particular backgrounds, beliefs, and concerns, and then predict how they might respond to data-use requests. Early work suggests this method is feasible: for instance, researchers have shown that AI agents can be created to mimic real individuals' decision-making and answer patterns with impressive fidelity. In one study, over a thousand people's personalities were simulated by feeding interview data to an LLM-based agent, which then answered questions in ways that closely mirrored the real respondents (Park et al., 2025). In the educational domain, AI-driven persona bots have even been used to provide realistic interview practice for graduate students, underscoring their ability to emulate human-like responses in a learning context (Sabbaghan & Brown, 2024). By harnessing such technology, we can explore how "virtual" students or parents might react when asked to consent to different data uses—for example, sharing academic records or behavioral logs for analytics. This study leverages that approach to investigate what conditions or communication strategies lead to smoother consent formation in educational data use. Ultimately, the goal is to inform the design of consent models that earn stakeholder trust and cooperation, allowing education systems to tap into data-driven innovations without compromising personal privacy. A virtual experiment platform for constructing consensus models that balance personal data protection with learning data analysis has also been proposed (Kakeshita et al., 2025).

## STUDY DESIGN

We conducted a virtual experiment using generative AI personas to model consent decisions for educational data use. The study combined three elements: (1) a structured set of learning data types, (2) multiple requester categories, and (3) systematically varied request conditions.

- **Learning Data**: Forty types of learning data were identified through iterative role-play interviews with GPT-4, instructed to simulate an experienced elementary school teacher.
- **Requesters**: Five major categories were defined, covering 12 specific entities:
  1. **School administrators and boards of education** (Board of Education, School, Health Office)
  2. **Teachers** (Homeroom teacher, Other teacher)
  3. **Researchers and research institutions** (University/Company, Individual researchers)
  4. **Parents** (Individual parent, Parent association)
  5. **Others** (Educational company, School bus operator)
- **Request Conditions:**
  - **Usage period:** three levels (unspecified, explicit end date, immediate disposal after use).

- **Third-party provision**: four levels (not mentioned, not allowed, re-consent required, disclosure of recipients and reasons).
- **Additional safeguards**: six options (e.g., anonymization, aggregated data, minimal fields only, feedback of insights, opt-out, encryption), with up to two conditions applied. To reduce the experimental scale, an orthogonal array design was employed.

### Persona Generation

Fifty pairs of personas, each consisting of an elementary school student and one parent, were generated using GPT-4.

- Students: defined by personality traits, hobbies, subject strengths, age, and gender.
- Parents: defined by occupation, hobbies, educational interest, attitudes toward data use, age, and gender.

### Dataset Construction and Consent Simulation

To evaluate requests for the use of learning data, we first constructed a comprehensive dataset of data types and request conditions. This was achieved through iterative discussions with GPT-4 instructed to role-play as an experienced elementary school teacher, in dialogue with the researcher. Through this process, we identified 40 categories of learning data (e.g., academic records, attendance logs, health information, digital traces) and systematically combined them with requesters, usage purposes, and conditional factors to create structured request scenarios.

Subsequently, these request scenarios were presented to 50 AI-generated parent–child persona pairs. Each pair was asked whether they would consent to providing the specified learning data under the given conditions. Responses—including approval, conditional approval, or refusal, along with justifications and any additional conditions—were generated using the OpenAI GPT-4 model via the LangChain framework. This procedure allowed us to simulate realistic variations in consent behaviors across diverse stakeholder profiles, while maintaining experimental control through the predefined dataset and orthogonal array design.

### Analytical Approach

- Descriptive statistics:

1. Descriptive statistics were applied to identify data types with high frequencies of permission, conditional permission, and refusal. To improve interpretability, the 40 types of learning data were consolidated into six categories:
2. Health/Life (e.g., health records, fitness records, insurance, geolocation)
3. Learning Activities (e.g., attendance, online logs, portfolio, sensor data)
4. Grades/Evaluation (e.g., test answers, evaluation data, certificates)
5. Psychological/Survey/Emotion (e.g., psychological assessments, questionnaires, emotion data)

6. Family/Personal Background (e.g., family situation, parental occupation, scholarship forms)
7. Identification/Administrative (e.g., roster, student profile, IEP, IP address)

Consent outcomes (Allow, Conditional Allow, Deny) were aggregated by category, separately for parents and children. The results were visualized using 100% stacked horizontal bar charts and radar charts to highlight differences in acceptance tendencies.

- Text mining and clustering:

Free-text justifications provided by personas were combined with the stated utilization purposes from the requesters. Texts were vectorized using TF–IDF and clustered using k-means ($k = 6$).

- Interpretation:

Clusters were reviewed qualitatively and labeled with human-readable themes (Educational improvement and progress monitoring, Direct learning support, Regional policy and safety management, Special needs education and research use, Privacy distrust). These labels are used throughout the Results and Discussion sections, while numeric codes (Cluster A–F) are provided for reference.

## RESULTS

### Overall Consent Patterns

Across all scenarios, conditional permission dominated the responses. Both student and parent personas frequently recognized the potential value of learning data utilization but simultaneously expressed reservations about privacy and data management.

- Refusals were concentrated on data types that appeared intrusive or surveillance-like, such as IP addresses, online activity logs, and geolocation data.
- Permissions were more frequent for data types directly linked to educational or health benefits, such as test scores, interest surveys, and health condition data.
- Conditional permissions prevailed in cases where the stated purpose was legitimate but respondents required safeguards, including anonymization, time-limited storage, and restrictions on third-party sharing.

Figure 1 illustrates the overall distribution of consent decisions (Permission, Conditional Permission, Refusal) for all learning data requests, comparing students and parents. Both groups showed a strong preference for conditional permission, though students were somewhat more willing to grant full permission, while parents more frequently chose refusal. This contrast suggests that while students tend to take a more straightforward stance of either granting or denying permission, parents are more cautious and prefer to set conditions, reflecting their greater awareness of potential risks.

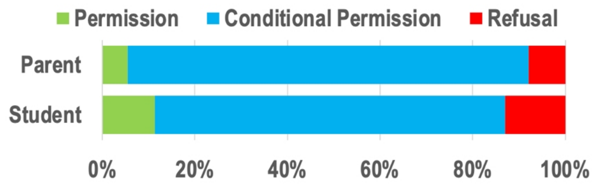


**Figure 1:** Overall distribution of consent decisions (permission, conditional permission, refusal) by students and parents across all learning data requests.

## Consent Patterns by Requester Type

Consent choices varied depending on who made the request. Figure 2 shows the distribution of consent decisions by requester type, separated for students (left panel) and parents (right panel).

- For both groups, conditional permission was the dominant choice across all requester categories.
- Students showed slightly higher unconditional permission rates toward teachers and parents as requesters, suggesting stronger trust in familiar figures.
- Parents tended to be more cautious, showing higher refusal rates when the requester was researchers or external parties.
- Across both groups, requests from administrative boards and researchers were generally accepted only conditionally, highlighting concerns about broader or external data use.

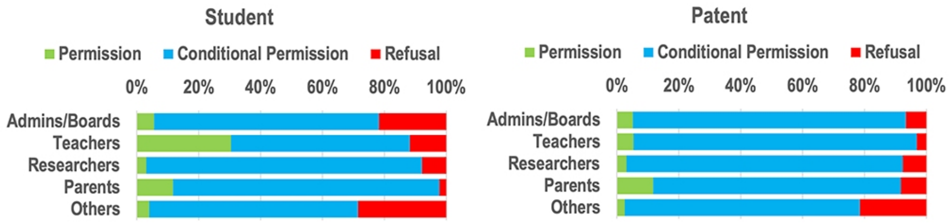


**Figure 2:** Distribution of consent decisions (permission, conditional permission, refusal) by requester type, comparing students and parents.

## Clustering of Purposes and Justifications

The integration of requesters' utilization purposes with respondents' justifications and subsequent clustering ($k = 6$) revealed six distinctive groups. These groups were labeled according to their thematic focus to aid interpretation. Table 1 summarizes the distribution of consent choices across the clusters.

- **Cluster A:** *Educational improvement and learning progress monitoring* (e.g., aggregated analyses of test scores, attendance logs, or progress data for policy and curriculum development)
  → The majority of responses indicated Conditional Allow (81%), often requiring anonymization and explicit data retention limits.
- **Cluster B:** *Direct learning support and classroom feedback* (e.g., remedial instruction, targeted feedback, qualification support)
  → This cluster showed the highest rate of unconditional Allow (15.8%), reflecting strong acceptance of purposes directly beneficial to individual learning.
- **Cluster C:** *Regional policy and safety management* (e.g., monitoring school bus operations, regional health statistics, or safety initiatives)
  → Responses were dominated by Conditional Allow (86.4%); refusals were relatively rare (4.8%).
- **Cluster D:** *Special needs education and research use* (e.g., individualized education plans, support for students requiring accommodations, research projects)
  → Responses overwhelmingly took the form of Conditional Allow (94.9%), emphasizing safeguards even when the purpose was well accepted.
- **Cluster E:** *Privacy distrust – surveillance-related data* (e.g., online learning records, IP addresses, geolocation, activity logs)
  → Responses were a mix of Conditional Allow and Deny. Unconditional permission was nearly absent, indicating strong concerns about monitoring and surveillance.
- **Cluster F:** *Privacy distrust – third-party sharing* (e.g., external research use, educational companies, re-provisioning of data)
  → This cluster was characterized almost exclusively by Conditional Allow, with refusals when re-consent or disclosure safeguards were missing.

**Table 1:** Distribution of consent decisions by cluster.

| Cluster | Refusal (%) | Conditional Permission (%) | Permission (%) |
|---|---|---|---|
| **A (Educational improvement)** | 10.8 | 81.0 | 8.1 |
| **B (Direct learning support)** | 23.0 | 61.3 | 15.8 |
| **C (Regional policy & safety)** | 4.8 | 86.4 | 8.8 |
| **D (Special needs education)** | 3.0 | 94.9 | 2.0 |

Continued

**Table 1:** Continued

| Cluster | Refusal (%) | Conditional Permission (%) | Permission (%) |
|---|---|---|---|
| E (Privacy distrust I) | 1.0 | 99.0 | 0.0 |
| F (Privacy distrust II) | 0.0 | 100.0 | 0.0 |

## Category-Based Comparison Between Children and Parents

To further explore differences in attitudes toward specific data types, the 40 learning data items were consolidated into six categories: Health/Life, Learning Activities, Grades/Evaluation, Psychological/Survey/Emotion, Family/Personal Background, and Identification/Administrative.

Figure 3 presents a radar chart comparing Allow and Deny rates by category for children and parents. Children were more permissive in Grades/Evaluation and Psychological/Survey/Emotion, while parents showed stronger reluctance, particularly regarding Family/Personal Background and Identification/Administrative data.

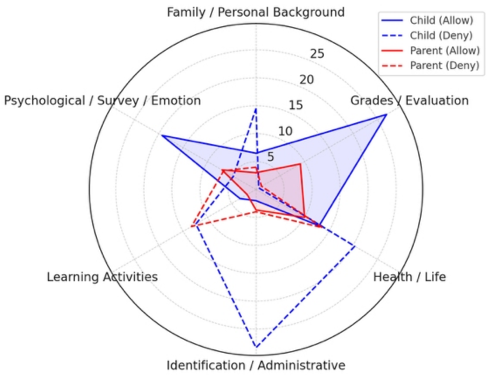


**Figure 3:** Comparison of parental and child consent and refusal across categories of learning data.

## DISCUSSION

This study provides three key contributions to understanding consent behaviors for educational data use.

1. **Conditionality as the dominant mode of consent.**
   Across both clustering and category analyses, Conditional Allow overwhelmingly outweighed unconditional permissions or refusals. This underscores the importance of offering conditional consent mechanisms (e.g., anonymization, limited retention, restricted sharing) in educational data governance frameworks.

2. **Purpose-driven and data-driven sensitivities.**
   The clustering analysis (A–F) revealed that acceptance was highest when data use served direct educational benefits (Cluster B) or well-accepted policy purposes (Cluster A, C, D). In contrast, Clusters E

and F captured fundamental distrust, differentiated into surveillance-related concerns and third-party sharing concerns. This suggests that stakeholders distinguish between "being monitored" and "losing control over data once shared externally."

3. **Category-specific variations and intergenerational differences.**
   The category-based comparison highlighted that children were more willing to share Grades/Evaluation and Psychological/Survey data, while parents were consistently more cautious, especially for Family/Personal Background and Identification/Administrative data. These findings imply that consent procedures should not only be data-sensitive but also account for different risk perceptions between parents and children.

4. **Implications for educational data governance.**
   Effective governance requires a dual-layered approach:
   By purpose (clusters A–F): enabling consent frameworks that differentiate between educational, policy, and privacy-distrust contexts.
   By data type (six categories): designing safeguards tailored to data sensitivity, particularly for surveillance-related and family-background information.
   Together, these results suggest that one-size-fits-all consent models are insufficient. Instead, multi-layered consent frameworks—integrating both purpose-based and data-type perspectives—are essential to balance educational innovation with privacy protection.

## Practical Implications

The findings suggest several concrete directions for practice:

- Dynamic consent models: Educational institutions should adopt consent processes that allow respondents to express conditional preferences (e.g., anonymization, opt-out, data minimization) rather than binary decisions.
- Category-specific safeguards: Policies should explicitly address the heightened sensitivity of Family/Personal Background and Identification/Administrative data, while promoting transparency in the handling of Grades/Evaluation and Learning Activities.
- Stakeholder engagement: Interventions should account for differences between children's and parents' perceptions. Mechanisms such as age-appropriate explanations, parental involvement, and co-consent procedures may improve trust and legitimacy.
- Risk communication: Distinguishing between surveillance risks and external-sharing risks can help policymakers and educators design clearer communication strategies that resonate with stakeholders' actual concerns.

## Limitations and Future Work

Several limitations must be acknowledged:

- Use of generative AI personas: The study relied exclusively on AI-generated parents and children to simulate consent decisions. While

this approach offers scalability, diversity, and experimental control, it cannot fully replace responses from real participants. The validity of persona-based simulations depends on how well generative models capture social and cultural nuances of data privacy attitudes. Recent studies have further suggested that LLM-based agents may respond more rationally and consistently than humans, potentially overlooking the variability and irrationality often observed in real decision-making (Feng et al., 2025).
- Lack of real-world validation: Future studies should compare persona-based outcomes with surveys or interviews conducted with actual students and parents, to assess alignment and identify systematic biases.
- Static request scenarios: Although the orthogonal design allowed systematic variation of conditions, real-world consent processes are often iterative, negotiated, and context-dependent. Future research could explore dynamic consent simulations, including repeated decision-making over time.
- Cluster interpretation: While six clusters (A–F) were interpretable, the boundaries between privacy-related clusters (E and F) may vary across datasets. Larger-scale replication with different cultural and institutional contexts is necessary to confirm their stability.

### Usefulness of Persona-Only Simulations

Despite these limitations, persona-only simulations provide unique benefits in specific contexts:

- Early-stage exploration: When designing new policies, guidelines, or consent models, persona simulations allow rapid prototyping and hypothesis generation before committing resources to large-scale human surveys.
- Ethically sensitive topics: For issues involving children's privacy, health, or surveillance, persona simulations enable preliminary discussion without exposing real participants to potential risks.
- Scenario stress-testing: Personas can be exposed to systematically varied, even extreme, request conditions that may be difficult to test in real-world studies, helping policymakers anticipate boundary cases.
- Comparative analysis: Persona responses can serve as a structured baseline, against which real-world survey results can be compared to identify gaps, biases, or cultural effects.

Thus, while persona-only studies should not be seen as substitutes for empirical surveys, they are valuable as exploratory tools, particularly in policy design, ethics debates, and methodological development for educational data governance.

## CONCLUSION

In sum, the combination of cluster-based purpose analysis and category-based parent–child comparison provides a multi-layered perspective on educational data consent. While persona-based simulation cannot substitute for real-world consent studies, it offers a valuable exploratory method for identifying patterns, generating hypotheses, and designing more nuanced governance models. Future research should integrate both simulated and empirical approaches to achieve a balanced understanding of stakeholder attitudes.

## ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number 23K02633 (Grant-in-Aid for Scientific Research (C)).

## REFERENCES

Digital Agency, Japan. (2025) Guidelines for the utilization of learning data. Available at: https://www.digital.go.jp/en/news/511df327–5ba3–456e-a5cd-2ebeddd8c960 (Accessed 23 September 2025).

European Union. (2016) General Data Protection Regulation (GDPR). Official Journal of the European Union, L119.

Feng, Y., et al. (2025) "Noise, adaptation, and strategy: Assessing LLM fidelity in decision-making", Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP, 2025), to appear. Available at: https://arxiv.org/abs/2508.15926.

Future of Privacy Forum. (2020) Parents' attitudes toward student data privacy. Washington, DC: FPF.

Kakeshita, T., Akiba, Y., Harada, S. and Takasaki, M. (2025) "A virtual experiment platform for building a consensus model that balances personal data protection and learning data analysis", Journal of Information Processing, 66(12), in press. (in Japanese).

Liu, Y. and Khalil, M. (2023) "Trust and willingness to share learning data: A systematic review", Journal of Learning Analytics, 10(2), pp. 45–62.

Ohm, P. (2010) "Broken promises of privacy: Responding to the surprising failure of anonymization", UCLA Law Review, 57(6), pp. 1701–1777.

Park, J., Lee, H. and Kim, S. (2025) "Simulating human decision-making with large language model personas", Proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI 2025).

Reidenberg, J. R. and Schaub, F. (2018) "Achieving big data privacy through accountability", Stanford Technology Law Review, 21(2), pp. 1–54.

Sabbaghan, S. and Brown, M. (2024) "AI-driven persona bots for interview practice in higher education", Computers & Education, 205, 104853.

UNESCO. (2023) Guidelines for the use of generative AI in education and research. Paris: UNESCO.