# Adaptive Consent in Crisis AI: A Privacy-Ethics Framework for Continual Data Use in Resilient Communities

**Haseeb Javed**[1], **Shynar Mussiraliyeva**[2], **Hayoung Oh**[3*], **Farman Ali**[3*]

[1]Department of Computer Science and Engineering, Sungkyunkwan University, Seoul, South Korea
[2]Faculty of Information Technology, Al-Farabi Kazakh National University, Almaty, Kazakhstan
[3]Department of Applied AI, Sungkyunkwan University, Seoul, South Korea
haseebjaved@g.skku.edu, musiralievash@kaznu.kz, hyoh79@skku.edu, farman0977@skku.edu

## Abstract

With the growing use of adaptive AI in emergency settings and crisis response, healthcare organizations increasingly rely on highly sensitive personal data gathered under time-pressure conditions, where consent is fragile, incomplete, or difficult to maintain. In these contexts, consent often fails because people may say yes while under stress or dependent on essential services, as organizations gathered the data rapidly with limited opportunity for meaningful choice, and the same records are later reused for retraining the systems, shared with teams, or used for new goals in later stages of the crisis. To prevent misuse, we introduce Adaptive Consent, a framework that treats consent as a changeable state that must be enforced each time data is reused. By enabling real-time consent management, Adaptive Consent can reduce data-misuse risks by up to 40% and improve traceability of compliance across AI models. It provides guidance for deciding how organizations can determine whether crisis and healthcare data may be reused for AI, for what purpose, and within which time window during the response, recovery, and preparedness phases. We also outline an actionable workflow that lets patients, callers, and community members give consent, and we describe how it can be operationalized through machine-readable policies and reuse "gates" for training and retrieval (e.g., Retrieval-Augmented Generation (RAG)/search) to update/retrain the AI. Finally, we discuss key limitations and trade-offs that include traceability across collaborating partners, handling artifacts generated before consent changes take effect, and the risk of unintended reuse, where data collected for care is later used beyond its initial purpose.

## Introduction

Adaptive AI is being used more in crisis response and healthcare to support urgent decisions, such as patient capacity sorting in overcrowded hospitals, call center assistance, and decisions about how to distribute limited supplies. These systems are designed to improve over time by learning from new information and adapting as conditions evolve, where needs, risks, and resources change quickly. However, the same conditions that make the AI effective in certain emergencies also create significant ethical and privacy risks, such as in the COVID-19 pandemic, where large volumes of patient mobility data, contact-tracing records, and hospital capacity information were collected and shared across public health agencies and organizations and reused as systems are updated (Barthwal, Campbell, and Shrestha 2025), (Abou Ali, Dornaika, and Charafeddine 2025). These privacy and ethics concerns here are not driven solely by data scale or model capability, as they also arise from how consent works during emergencies. In crisis situations, with unclear or unstable consent, individuals may share sensitive information when they are upset to rely on care, without understanding how their data could be reused later (Sudhi et al. 2025), (Abou Ali, Dornaika, and Charafeddine 2025). As a result, consent is often captured as a single, point-in-time decision, while data may continue to circulate and be reused for retraining, retrieval-based access, and cross-organizational partnerships for long periods. Over time, this prolonged reuse increases the chances of accidental disclosure, re-identification, and downstream use that surpass the original intent of concern or emergency response (Tarantini et al. 2025). Existing privacy approaches are helpful, but they are not fully sufficient for this evolving data lifecycle. Practices such as encrypting data storage, de-identifying records, and limiting access do not fully cover situations for adaptive AI systems in which models are repeatedly retrained, policies evolve, and additional partners gain access in the data-sharing networks over time. In this context, the core problem shifts from "How do we protect the data?" to "How do we handle consent as it shifts across phases and over time?" To meet this need, we introduce "Adaptive Consent", a framework that ties the crisis phase (emergency, recovery, preparedness) to define rules about data usage, intended purposes, and time limits, while enabling individuals to update or withdraw consent. Keeping AI adapted to changing consent can help minimize harm, maintain trust, and promote responsible AI in sensitive environments.

## Related Work

### Consent During Crises and Changing Data Use

Research shows that privacy and ethics are highly sensitive to context and relational dynamics (Bhat 2025), (Ekmekci, Zhang, and Crawley 2025). In emergencies and other high-stress moments, people may share sensitive details simply to obtain help without understanding long-term reuse or shar-

*Corresponding Authors

| Aspect | Traditional Consent Model | Crisis-Specific Consent Model |
|---|---|---|
| Timing | Consent taken before data is used | Consent taken later or by another person |
| Scope | Narrow, purpose-specific permissions | Broader scope to support rapid response |
| User Control | High individual control right | Limited control during crisis |
| Data Use Risks | Lower due to clearly defined permissions | Higher due to evolving use cases |

Table 1: Traditional Consent Comparison with Crisis-Specific Consent Models

ing (Wei and Liu 2025), (Seidenberger and Maiti 2025), (Gilga et al. 2025). The concept of contextual integrity explains why this matters, as privacy expectations depend on the context of data sharing and why later reuse can violate norms when the original disclosure was reasonable in the moment (Ali et al. 2011), (Shah et al. 2025), (Li et al. 2025), (Goniewicz, Burkle, and Khorram-Manesh 2025). In addition, critiques of notice-and-consent argue modern data practices make it impractical to ask individuals to predict future data uses or manage permissions across complex, multi-actor systems (Ali et al. 2020), (Sloan and Warner 2014), (Khan, Levine, and Nguyen 2025), (Chomanski 2025). In crisis settings, these issues worsen because needs and objectives change quickly, and the same information can shift from immediate response to recovery and long-term planning, creating a gap between one-time consent and evolving systems. In general terms, traditional consent focuses on prospective, limited, fully informed consent, while crisis-specific models incorporate deferred or proxy choices, have wider, more flexible scopes that favor speed, and increase the risk of problematic data reuse. This is summarized in the Table 1, which is a comparison of these traditional and crisis-specific consent models.

## Continual and Adaptive Privacy Model Updating

Although many studies propose privacy reduction methods in machine learning, such as differential privacy and decentralized training (Wu et al. 2022), (Javed et al. 2025). However, these adaptive AI systems differ from static models because they are updated repeatedly with new data and changing goals, as they raise questions about withdrawing consent after data is influenced and reused in multiple models (Ali, Anwar, and Solehria 2013), (Hatherley 2025), (Luo and Ji 2025). Machine unlearning works on removing data influence (Nguyen et al., 2025), and methods such as SISA (Sharded, Isolated, Sliced, and Aggregated) make deletion more practical (Ferdous et al. 2025). Even so, adapting these approaches proves difficult at scale; they still do not answer policy-related questions about time bounds and which sections of pipelines could be considered "use" (training sets, logs, embeddings, and retrieval indexes). This needs to highlight the importance of consent control across continual updates, not just at data collection.

## Key Technical Challenges

Adaptive consent introduces several technical challenges: (1) crisis AI systems are modified over time; data is dispersed across multiple components such as databases, training sets, logs, and retrieval indexes; (2) consent revisions must be applied across the entire life cycle instead of just at the collection phase; and (3) enforcing such updates may require costly operations, including model retraining.

**Purpose and phase limits.** Using access controls and policy checks pre-sharing or pre-training can restrict systems from reusing by purpose (response vs. preparedness) and by phase (emergency, recovery, preparedness).

**Time-based retention.** Datastores have retention windows that expire records and subsequently delete them or lock them so they cannot be queried, shared, or used for future training of AI models.

**Consent-aware retrieval.** In search of RAG systems to protect individual records without permission and to obscure sensitive fields, teams can enforce consent at the time of the search, which often minimizes the risk more quickly than having to retrain the model.

**Basic audit trails.** Systems can document consent status with each record and log when that record is used in training or retrieval, allowing organizations to verify compliance

**Algorithm 1**: Consent-Aware RAG Pseudocode

```
Adaptive Consent Context (Crisis-Specific)
user_consent = {
    "user_id": "caller-456",
    "crisis_phase": "recovery",
    "granted_purposes": ["medical-triage"]
}
def adaptive_consent_retrieval(query, user_consent, vector_store):
Vectorize Crisis Query
    query_embedding = embedding_model.embed(query)
Multi-Gate Metadata Filter
    Enforces Purpose and Phase-based Reuse Gates
    metadata_filter = {
        "purpose": {
            "$in": user_consent["granted_purposes"] },
        "phase_allowed": user_consent["crisis_phase"]
    }
Filtered Retrieval (The "Reuse Gate")
    return vector_store.similarity_search(
        query_embedding,
        k=5,
        filter=metadata_filter
    )
Execute Protected RAG Pipeline
context = adaptive_consent_retrieval(
    "Where is the nearest oxygen supply?",
    user_consent,
    crisis_db
)
response = llm.generate(
    f"Context: {context}\nQuestion: {query}" )
```

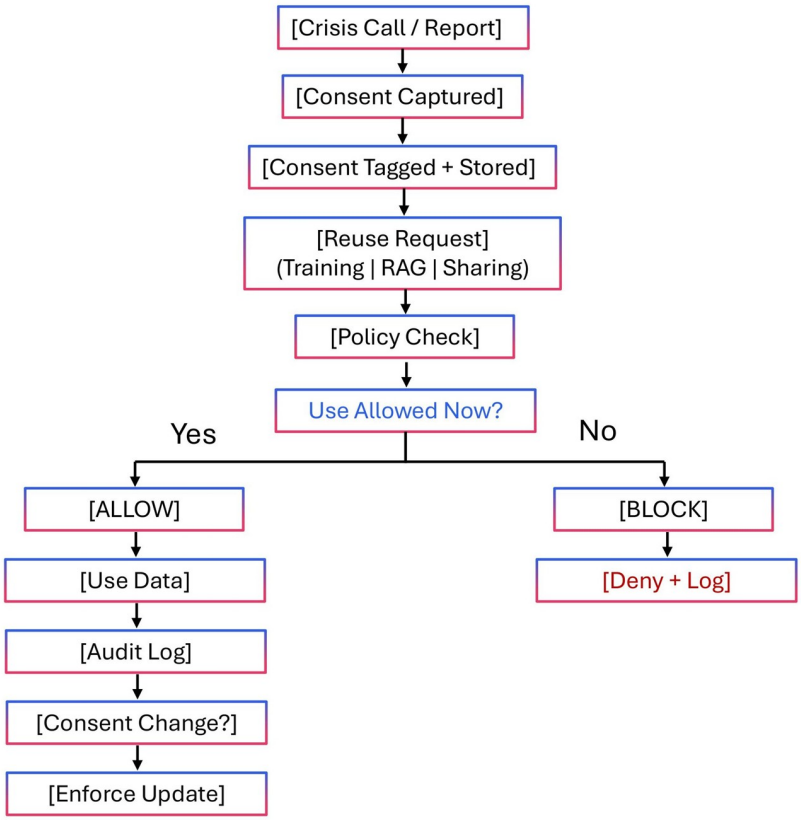


Figure 1: Adaptive Consent Workflow for Continual Crisis AI Data Use.

and examine potential misuse later. Algorithm 1 shows a detailed example that demonstrates how an AI retrieval system enforces adaptive consent by filtering data based on the user's allowed purpose and crisis phase before generating a response in a protected RAG pipeline.

## Limitations and Open Technical Gaps

**Consent withdrawal after model updates.** When data undergoes several model versions, its effects become embedded and it becomes challenging to withdraw. Because of ongoing retraining, a record may be assimilated into training, features, indexes, or reports. To get rid of it, one would have to retrain, rebuild indexes, and do some artifact clearing. Current machine unlearning methodologies provide partial assistance; however, they are not yet fully integrated for many real-world deployments.

**Data sharing across partners breaks enforcement.** Crisis response operates across many organizations and vendors. Consent conditions may be respected within one system but are lost elsewhere due to different tagging, policies, or retention. Without systems with compatible policies and shared audit criteria, adaptive consent will continue to be highly inconsistent throughout the ecosystem.

**Maintaining Consent Controls Under Emergency Constraints.** When deploying new adaptive features in emergencies, teams may skip detailed metadata design, consent tracking, and logging, with the assumption that these can be addressed later. However, the problem is that emergency shortcuts tend to harden into long-term practice as systems expand. Adaptive Consent, therefore, requires a privacy-by-default foundation that still holds up when decisions must be made quickly, not just when circumstances are perfect. In addition, organizations must embed scalable governance mechanisms so that they can ensure accountability under high-pressure and in rapidly changing environments.

Such challenges demonstrate the need for an adaptive consent framework that clearly defines where consent must be enforced, such as at the stages of data collection, model training, and system update cycles, while also clearly mentioning the distinctions between what is immediately actionable and what remains resource-intensive. Furthermore, it should provide structured guidance and reasonable counselling to prioritize the implementation efforts, while balancing organization operational needs with long-term ethical, legal, and technical sustainability.

## Proposed Framework

To address these gaps, we propose **Adaptive Consent**, a framework that treats consent as a dynamic state that can evolve over time and must be re-evaluated whenever crisis data is reused. Instead of claiming that permissions given during an emergency carry over to subsequent stages, new partners, or new versions of the model, Adaptive Consent

requires that each reuse event positively correlates with an explicit policy. Every decision has a purpose, a time frame (expiration/review), and an accountability document, so that the ongoing adaptive learning and iterative deployments stay in sync with the expectations of the person or the community. Consent is documented in ways that fit with crisis conditions: the system registers who is giving authorization (an individual or a delegated community channel), for which purpose the data may be used, the duration of the consent, and what forms of reuse may be prohibited.

The intention is to reduce legal warnings and present a limited number of concise and workable options that can still be exercised under time constraints. Consent is recorded and translated into machine-readable policies via tagging each record with a consent tag, as shown in Figure 1. This tag records what the user consented to, including the use of the record for what purpose, what phases it can be used/shared, for what duration, the constraints on use and sharing, and if the record can be used for training and/or retrieval (e.g., RAG/search). The framework facilitates the enforcement of consent at the record level to minimize "silent reuse," whereby data collected for immediate support is subsequently used for unrelated development and/or long-term modeling without obtaining renewed consent.

Adaptive consent applies where actual reuse takes place, not just where data is stored. When a training job is started, only records that are policy-compliant for training concerning the requested purpose, phase, and time window can be included in the training dataset. During a retrieval request, only those records with the appropriate permissions can be made available to a model, and obscuring sensitive or confidential information or filtering is performed where necessary. With external data sharing, downloads and exports are permitted only if the recipient can abide by the same consent restrictions, so permissions do remain when data crosses organizational boundaries.

Given that consent is subject to change, triggers for updates and withdrawals result in definitive actions in the system. At the very least, the system restrains future training and future retrieval regarding the records in question. Where possible, impacted records are also removed from retrieval indices, and the system is set to perform records retrieval updates and unlearning for the affected modules, so that future versions of the models will incorporate the updated consent. The system clearly separates actions that can be taken immediately and those that may need to be deferred, which ensures that 'consent' does not become a system operationally unenforceable promise.

Finally, Adaptive Consent relies on auditability and review processes. Every access for training, retrieval, or sharing creates a log entry that connects the activity to the relevant consent state, purpose, and version of the system. Additionally, each log captures contextual metadata such as user identity, access pathways, and strengthens traceability. This makes it possible to track compliance, aids in resolving disputes when misuse is unclear, and allows organizations to show proof of compliance without impeding emergency operations. In addition, version-controlled consent records enable retrospective analysis that allows multiple organizations to verify the data usage which is aligned with exact consent conditions. Those structured audit trails support regulatory reporting requirements and third-party audits without requiring disruption to ongoing system processes.

## Conclusion

As traditional one-time consent frameworks prove too rigid and limited for resilient communities, crisis AI systems need ongoing data collection and continuous model updates throughout emergency response, recovery, and preparedness. We presented Adaptive Consent, a framework that treats consent as a dynamic, enforceable state: permissions are recorded in a crisis-relevant manner, converted into a machine-readable policy, and enforced at every reuse point, including training, retrieval (RAG/search), and sharing, to ensure that data collected for emergency crises cannot be reused later without the individual consent in the future. Adaptive Consent presents a constructive way to maintain privacy-preserving, accountable crisis AI that can evolve without losing trust. This is particularly the case when enforcement is coupled with audit logging, and the processes for updating and withdrawing consent (immediate blocking of future use and cleanup when feasible) are handled clearly.

## Acknowledgments

This research was supported in part by the Sports and Tourism R&D Program through the Korea Creative Content Agency, funded by the Ministry of Culture, Sports and Tourism (grant number: RS-2024-00344893), and in part by an Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. RS-2025-25442569, AI Star Fellowship Support Program (Sungkyunkwan Univ.)).

## References

Abou Ali, M.; Dornaika, F.; and Charafeddine, J. 2025. Agentic AI: a comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*, 59(1): 11.

Ali, S.; Anwar, S.; and Solehria, S. 2013. User interaction based framework for protecting user privacy in online social networks. *Proceedings of the ICISO*.

Ali, T.; Alam, M.; Nauman, M.; Ali, T.; Ali, M.; and Anwar, S. 2011. A scalable and privacy preserving remote attestation mechanism. *Information-An International Interdisciplinary Journal*, 14(4): 1193–1203.

Ali, T.; Khan, Y.; Ali, T.; Faizullah, S.; Alghamdi, T.; and Anwar, S. 2020. An Automated Permission Selection Framework for Android Platform: T. Ali et al. *Journal of Grid Computing*, 18(3): 547–561.

Barthwal, A.; Campbell, M.; and Shrestha, A. K. 2025. Privacy ethics alignment in ai: A stakeholder-centric framework for ethical ai. *Systems*, 13(6): 455.

Bhat, M. 2025. Toward an Ethic of Synthetic Relationality: Identity, Intimacy, and Risk in AI-Mediated Roleplay Environments. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 416–429.

Chomanski, B. 2025. The challenge of regulating digital privacy. *Critical Review of International Social and Political Philosophy*, 1–25.

Ekmekci, P. E.; Zhang, L.; and Crawley, F. P. 2025. Ethics Underpinning Data Policy in Crisis Situations. *Data Science Journal*, 24.

Ferdous, S.; Omi, N. H.; Mahi, I. H.; Hemel, M. H. N.; and Badhon, M. S. S. 2025. *Machine unlearning for class removal through SISA-based deep neural network architectures*. Ph.D. thesis, Brac University.

Gilga, C.; Hochwarter, C.; Knoche, L.; Schmidt, S.; Ringler, G.; Wieland, M.; Resch, B.; and Wagner, B. 2025. Legal and ethical considerations for demand-driven data collection and AI-based analysis in flood response. *International Journal of Disaster Risk Reduction*, 122: 105441.

Goniewicz, K.; Burkle, F. M.; and Khorram-Manesh, A. 2025. Transforming global public health: climate collaboration, political challenges, and systemic change. *Journal of infection and public health*, 18(1): 102615.

Hatherley, J. 2025. A moving target in AI-assisted decision-making: dataset shift, model updating, and the problem of update opacity: J. Hatherley. *Ethics and Information Technology*, 27(2): 20.

Javed, H.; Ali, F.; Shah, B.; Dilshad, N.; and Kwak, D. 2025. MediGuard: Protecting Sensitive Healthcare Data with Privacy-Preserving Language Models. *IEEE Journal of Biomedical and Health Informatics*.

Khan, L. M.; Levine, S. A.; and Nguyen, S. T. 2025. After Notice and Choice: Reinvigorating" Unfairness" to Rein In Data Abuses. *Stan. L. Rev.*, 77: 1375.

Li, H.; Fan, W.; Chen, Y.; Jiayang, C.; Chu, T.; Zhou, X.; Hu, P.; and Song, Y. 2025. Privacy checklist: Privacy violation detection grounding on contextual integrity theory. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 1748–1766.

Luo, H.; and Ji, C. 2025. Cross-cloud data privacy protection: Optimizing collaborative mechanisms of ai systems by integrating federated learning and llms. In *2025 IEEE 7th International Conference on Communications, Information System and Computer Engineering (CISCE)*, 230–233. IEEE.

Seidenberger, S.; and Maiti, A. 2025. From Big Data to Valued Data: A Dataset Value Taxonomy for AI-Native Empirical Research. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 2294–2305.

Shah, B.; Junaid, M.; Rustam, H.; Habib, M.; and Anwar, S. 2025. AI-Driven Fog-Edge Computing for IoMT Systems: Architecture and Use Cases. In *Proceedings of the AAAI Symposium Series*, volume 6, 42–48.

Sloan, R. H.; and Warner, R. 2014. Beyond notice and choice: Privacy, norms, and consent. *J. High Tech. L.*, 14: 370.

Sudhi, M.; Aishwarya, T.; Shetty, D. K.; Balakrishnan, J. M.; Ahmad, S.; and Sankaran, P. P. 2025. Ai-driven innovations in emergency and disaster response: strategies for effective planning. *Proc Eng Sci*, 7: 1293–304.

Tarantini, G.; Fraccaro, C.; Porzionato, A.; Van Mieghem, N.; Treede, H.; Shammas, N.; Szerlip, M.; Thourani, V.; Gerosa, G.; Marchese, A.; et al. 2025. Informed consent and shared decision-making in modern medicine: case-based approach, current gaps and practical proposal. *The American Journal of Cardiology*, 241: 77–83.

Wei, W.; and Liu, L. 2025. Trustworthy distributed ai systems: Robustness, privacy, and governance. *ACM Computing Surveys*, 57(6): 1–42.

Wu, X.; Zhang, Y.; Shi, M.; Li, P.; Li, R.; and Xiong, N. N. 2022. An adaptive federated learning scheme with differential privacy preserving. *Future Generation Computer Systems*, 127: 362–372.