

Context-Aware Authorization Framework for Agentic AI: Adaptive Access Decisions, Risk Assessment, and Trust Management

Minh Anh Nguyen

Department of Artificial
Intelligence Research Specialist, Vietnam

Quang Huy Tran

Machine Learning
Research Analyst, Vietnam

VOLUME 03 ISSUE 09 (2026)

Published Date: 10 September 2026 // Page no.: - 08-14

ABSTRACT

Agentic artificial intelligence systems differ from conventional software applications because they can interpret objectives, invoke tools, access heterogeneous resources, and execute multi-step actions with limited human intervention. Consequently, authorization mechanisms based exclusively on static identities, predefined roles, or binary permission checks are insufficient for environments in which access requirements change according to task context, resource sensitivity, behavioral evidence, and evolving risk. This paper proposes a Context-Aware Authorization Framework for Agentic AI that combines contextual representation, adaptive risk assessment, attribute-aware decision mechanisms, and dynamic trust management. The methodological foundation is derived from the provided literature on attention modeling, attribute-specific representation learning, similarity prediction, iterative learning, and knowledge distillation. These techniques are conceptually transferred to authorization by treating an agent's request as a multidimensional representation rather than as an isolated identity-permission pair. The framework introduces a context aggregation layer, risk evaluation engine, policy decision mechanism, and continuous trust monitor. A conceptual evaluation demonstrates how increasingly fine-grained contextual information can improve authorization discrimination while introducing computational and governance trade-offs. The resulting framework provides a structured approach for risk-adaptive authorization in agentic environments, particularly where agents operate across heterogeneous tools and dynamically changing operational contexts.

Keywords: Agentic AI, Context-Aware Authorization, Adaptive Access Control, Risk Assessment, Trust Management, Attribute-Based Authorization, AI Security, Dynamic Policy Enforcement, Contextual Decision-Making

INTRODUCTION

The emergence of agentic AI changes the security assumptions underlying conventional access-control systems. A conventional application generally performs actions within a relatively predictable permission boundary. An AI agent, in contrast, can interpret a high-level objective, decompose it into subtasks, select tools, retrieve information, communicate with external systems, and execute actions whose consequences may vary substantially according to context. An authorization decision therefore cannot always be reduced to the question of whether a particular identity possesses a particular permission.

The central problem addressed in this research is the development of an authorization mechanism capable of determining whether an agent should be permitted to perform a particular action under a particular context and risk condition. Context may include the agent's current task, requested operation, resource sensitivity,

environmental state, previous behavior, confidence level, provenance of information, and trust history. The authorization outcome should consequently be adaptive rather than permanently determined by static privileges.

Recent work on risk-adaptive authorization for agentic AI provides a direct motivation for treating authorization as a dynamic security decision rather than a static access-control lookup (Bhushan, Pappu, Jadav, and Jaiswal, 2026). However, the broader methodological literature supplied for this study comes from visual retrieval and similarity-learning research. Studies on multi-attribute similarity learning, attribute-specific embeddings, attention mechanisms, and iterative similarity learning provide useful conceptual mechanisms for representing and comparing multidimensional contextual information (Ak et al., 2018; Dong, 2021; Ma, 2020; Yan et al., 2022).

The objective of this paper is therefore to construct a transferable framework in which contextual authorization can be modeled as a structured similarity and risk-assessment

problem. Rather than evaluating identity independently from environmental attributes, the proposed approach creates a contextual representation of each authorization request and compares it against policy, behavioral, and trust requirements.

The primary contribution is a four-layer architecture consisting of contextual representation, adaptive risk assessment, policy decision, and continuous trust management. The framework also establishes a conceptual relationship between attention-based feature selection and authorization, attribute-aware representations and contextual policies, and iterative similarity learning and continuous trust evaluation.

LITERATURE REVIEW

2.1 Attribute-Aware Representation as a Foundation for Contextual Authorization

Authorization requests contain multiple attributes that may interact rather than operate independently. For example, an agent requesting access to a financial database may be considered differently when performing an approved auditing task than when attempting an unrelated exploratory task. Similarly, the same operation may be acceptable during normal operation but inappropriate when the agent exhibits anomalous behavior.

The literature supplied provides extensive evidence that multidimensional attributes can be represented explicitly for similarity and classification purposes. Ak et al. developed multi-attribute similarity learning for attribute-based fashion search, demonstrating the value of jointly modeling multiple characteristics rather than treating them as isolated features (Ak et al., 2018). Ma similarly investigated attribute-specific embedding for fine-grained similarity learning (Ma, 2020). These ideas can be transferred conceptually to authorization by representing a request through dimensions such as identity, task, resource, operation, temporal context, behavioral history, and environmental risk.

Dong's work on attribute-specific embedding learning further supports the principle that individual attributes may contribute differently to a final similarity decision (Dong, 2021). In an authorization environment, this implies that the sensitivity of a resource, for instance, should not necessarily receive the same decision weight as the agent's task alignment or trust history.

This provides the theoretical basis for the proposed contextual representation layer.

2.2 Attention and Context Selection

Not every contextual attribute is equally relevant to every

authorization decision. Attention-based approaches provide a mechanism for emphasizing informative features while reducing the influence of less relevant information.

Su et al. proposed an attentional heterogeneous bilinear network for fashion image retrieval, illustrating how heterogeneous information can be integrated through attention mechanisms (Su et al., 2021). Zhang et al. examined task-aware attention for attribute prediction, emphasizing the importance of task-specific information in determining relevant features (Zhang et al., 2020). The relevance to agentic authorization is direct at the methodological level: authorization should prioritize context that is relevant to the requested operation rather than assigning uniform importance to every available signal.

For example, when an agent requests read-only access to a low-sensitivity document, task alignment and identity history may dominate the decision. For a request to modify critical infrastructure, resource sensitivity, action impact, behavioral anomalies, and recent trust changes should receive substantially greater attention.

Thus, task-aware contextual weighting becomes an important component of the proposed framework.

2.3 Similarity, Compatibility, and Authorization Decisions

Authorization can also be conceptualized as a compatibility problem. A request is compatible with a policy when its contextual characteristics sufficiently satisfy the policy's conditions.

Veit et al. examined conditional similarity networks, demonstrating the importance of learning similarity under different conditions (Veit et al., 2017). Xiao and Yamasaki investigated fashion compatibility through self-adaptive training and later relational knowledge distillation (Xiao and Yamasaki, 2022; Xiao and Yamasaki, 2024). Xiao, Zhang, and Yamasaki additionally considered compatibility under color distortion (Xiao, Zhang, and Yamasaki, 2023). Collectively, these works suggest that similarity should not always be interpreted as a fixed relationship; it can depend on contextual conditions and transformations.

For authorization, this means that a request should not be judged against an immutable notion of "allowed" or "denied." Its compatibility with a policy may depend on current risk, operational purpose, resource sensitivity, and trust state.

2.4 Retrieval, Generalization, and Cross-Domain Context

Context-aware authorization frequently operates across heterogeneous systems. An agent may interact with enterprise databases, cloud services, APIs, document repositories, and external tools. Cross-domain representation is therefore important.

Ji et al. investigated cross-domain image retrieval with attention modeling, providing a methodological analogy for combining information originating from different domains (Ji et al., 2017). Goenka's FashionVLP incorporated vision-language representations for fashion retrieval with feedback, while Liu et al. introduced a large-scale annotated dataset for robust clothes recognition and retrieval (Goenka, 2022; Liu et al., 2016). These studies reinforce the broader principle that robust decision systems require representations capable of handling heterogeneous information.

The authorization equivalent is a unified contextual representation capable of incorporating identity, task, resource, action, environment, and behavioral information.

METHODOLOGY

3.1 Proposed Context-Aware Authorization Architecture

The proposed framework consists of four principal layers:

1. Contextual Representation Layer
2. Adaptive Risk Assessment Layer
3. Policy Decision and Enforcement Layer
4. Continuous Trust Management Layer

The conceptual architecture treats an authorization request as a structured contextual object rather than a simple identity-permission pair.

Let an authorization request be represented as:

$$R=\{I,T,A,O,S,E,B,P\}R=\{I,T,A,O,S,E,B,P\}$$

where II represents agent identity, TT task context, AA requested action, OO target object or resource, SS resource sensitivity, EE environmental context, BB behavioral evidence, and PP provenance-related information.

The objective is to compute a decision:

$$D=f(R,\Pi,\tau)D=f(R,\Pi,\tau)$$

where Π/Π represents applicable authorization policies and τ/τ represents the current trust state.

The decision need not be binary. A practical system can generate allow, allow-with-constraints, require-human-approval, or deny outcomes.

3.2 Contextual Representation Layer

The first stage transforms heterogeneous request attributes into a normalized contextual representation.

A static authorization model might represent an agent as:

$$\text{Agent} \rightarrow \text{PermissionAgent} \rightarrow \text{Permission}$$

The proposed model instead represents:

$$\text{Agent} + \text{Task} + \text{Action} + \text{Resource} + \text{Context} + \text{Risk} + \text{Trust} \rightarrow \text{Decision}$$

$$\text{Agent} + \text{Task} + \text{Action} + \text{Resource} + \text{Context} + \text{Risk} + \text{Trust} \rightarrow \text{Decision}$$

This distinction is fundamental because the same agent can legitimately require different privileges during different tasks.

Attribute-specific learning provides a conceptual foundation for this approach. Instead of compressing all information into an undifferentiated representation, individual attributes can maintain their semantic identity. This allows policy rules to distinguish, for example, between a high-trust agent performing a low-impact read operation and the same agent requesting a high-impact administrative action.

3.3 Contextual Attention Mechanism

The framework introduces task-sensitive contextual weighting.

For each contextual feature $x_{i,j}$, an attention weight $w_{i,j}$ is assigned according to its relevance to the current authorization request:

$$C=\sum_{i=1}^n w_i x_i C=\sum_{i=1}^n w_i x_i$$

where CC represents the contextual authorization representation.

The weights should change according to the requested task. This is conceptually consistent with task-aware attention mechanisms described in the supplied literature (Zhang et al., 2020).

For example, suppose an agent requests access to a customer database to generate a routine report. Task alignment, resource sensitivity, identity assurance, and historical behavior may be emphasized. If the agent instead requests permission to delete customer records, action impact and resource sensitivity should receive significantly higher weights.

The major advantage is contextual selectivity. The principal limitation is that dynamic weighting increases system complexity and creates additional requirements for explainability.

3.4 Adaptive Risk Assessment

Risk assessment converts contextual evidence into an operational risk value.

A conceptual risk function can be defined as:

$$Risk = \alpha R_s + \beta R_a + \gamma R_b + \delta R_e + \epsilon R_t$$

$Risk = \alpha R_s + \beta R_a + \gamma R_b + \delta R_e + \epsilon R_t$

where R_s is resource sensitivity, R_a is action criticality, R_b is behavioral anomaly, R_e is environmental risk, and R_t represents trust-related risk.

The coefficients determine the importance of each factor.

The key principle is that risk should be adaptive. An authorization request that was acceptable earlier may become unacceptable after a change in behavior or environmental conditions. This risk-adaptive orientation is central to the agentic authorization problem (Bhushan et al., 2026).

A hypothetical example illustrates this behavior. Assume an agent normally retrieves internal reports. Its historical trust level is high, and the request is consistent with its assigned task. The system may permit access. If the same agent suddenly attempts to access a highly sensitive database outside its normal task context, the contextual representation changes and risk increases. The resulting decision could transition from allow to human approval without permanently revoking the agent's identity.

3.5 Policy Decision Engine

The policy engine evaluates the contextual representation against authorization requirements.

A policy may specify:

Policy = (Task, Resource, Action, Conditions, Threshold)
Policy = (Task, Resource, Action, Conditions, Threshold)

The decision engine calculates contextual compatibility and compares the resulting risk with a policy threshold.

For example:

- Low risk + high policy compatibility → Allow
- Moderate risk + acceptable compatibility → Allow with restrictions
- Elevated risk + uncertain compatibility → Human approval
- High risk + low compatibility → Deny

This approach avoids the limitations of purely binary access control.

Similarity-learning research provides a useful conceptual basis for this mechanism. Conditional similarity and

compatibility methods demonstrate that relationships can depend on contextual conditions rather than being universally fixed (Veit et al., 2017). In authorization, the “similarity” represents how closely a request matches the intended conditions of a legitimate policy.

3.6 Continuous Trust Management

Trust should not be treated as a permanent property of an agent.

The proposed trust state is updated using behavioral evidence:

$$\tau_{t+1} = g(\tau_t, B_t, D_t)$$

where τ_t is current trust, B_t represents observed behavior, and D_t represents previous authorization decisions and outcomes.

Successful behavior consistent with policy can preserve or increase trust, whereas repeated policy violations, unusual resource access, or suspicious action sequences can reduce trust.

This mechanism is particularly important for autonomous agents because a single authentication event cannot capture the complete security state of a long-running agent.

3.7 Constraint-Based Authorization

Adaptive authorization should not always respond to increased risk with complete denial. Instead, the framework supports graduated restrictions.

For example, an agent may be allowed to:

- read but not modify a resource;
- access aggregated rather than raw information;
- execute a tool with restricted parameters;
- perform an operation only after human approval;
- access a resource for a limited period.

Such restrictions transform authorization from a binary gate into a risk-sensitive control mechanism.

The approach also aligns with the broader literature's emphasis on fine-grained attribute and compatibility modeling (Ak et al., 2018; Ma, 2020).

RESULTS

The conceptual evaluation of the proposed framework produces four principal findings.

First, contextual representation provides greater decision

granularity than static authorization. When identity, task, resource, action, behavior, and environment are evaluated together, the framework can distinguish requests that would otherwise appear identical under role-based authorization. This is particularly important for autonomous agents whose permissions can vary substantially according to operational objectives.

Second, attention-based contextual weighting improves decision relevance. The supplied literature demonstrates the usefulness of task-aware and heterogeneous attention for identifying informative attributes (Su et al., 2021; Zhang et al., 2020). Transferred to authorization, this means that the decision engine can emphasize the characteristics most relevant to a specific operation. High-impact actions can therefore trigger stronger consideration of resource sensitivity and behavioral evidence than routine operations.

Third, adaptive risk assessment enables graduated responses. Rather than treating risk as an independent monitoring metric, the proposed framework directly integrates risk into the authorization outcome. The resulting decision space includes restricted access and human escalation in addition to allow and deny. This creates a more flexible security posture and corresponds closely with the risk-adaptive perspective proposed for agentic AI authorization (Bhushan et al., 2026).

Fourth, continuous trust provides temporal sensitivity. An agent's authorization status can evolve as new behavioral evidence becomes available. This prevents a high initial trust score from becoming a permanent security assumption. The framework therefore supports continuous reassessment rather than one-time authorization.

The primary trade-off is computational complexity. A static permission lookup can be extremely efficient, whereas contextual authorization requires feature extraction, risk calculation, policy evaluation, and trust updates. Furthermore, dynamic contextual models introduce explainability challenges. If the system denies a request because multiple contextual attributes collectively increased risk, administrators need sufficient evidence to understand the decision.

The findings therefore indicate that context-aware authorization is most beneficial in environments where the security consequences of autonomous actions justify additional decision complexity. For low-risk applications, a simpler authorization mechanism may remain preferable.

DISCUSSION

The proposed framework changes the conceptual foundation of authorization from identity-centric control

to context-centric decision-making. This distinction is particularly significant for agentic AI because autonomous systems can generate action sequences that were not explicitly anticipated when their initial permissions were configured.

The literature supplied supports the methodological components of this transition. Attribute-specific approaches demonstrate the value of preserving semantic distinctions among features (Dong, 2021; Ma, 2020). Attention-based approaches demonstrate how task relevance can influence feature importance (Su et al., 2021; Zhang et al., 2020). Conditional similarity and compatibility approaches demonstrate that relationships can be evaluated under contextual conditions (Veit et al., 2017; Xiao and Yamasaki, 2022). The proposed authorization architecture synthesizes these principles into a security decision process.

The principal theoretical implication is that authorization can be understood as a contextual compatibility problem constrained by risk. A request is not simply "authorized" because an agent possesses an identity. Instead, it is evaluated according to whether the current combination of task, action, resource, behavior, and environmental conditions is compatible with the policy's intended security boundary.

This perspective also clarifies an important trade-off between security and autonomy. Excessively restrictive authorization can undermine the operational value of autonomous agents by forcing frequent human intervention. Conversely, excessive autonomy can create unacceptable risk when agents possess broad tool access. Risk-adaptive authorization offers an intermediate approach in which intervention intensity increases as contextual uncertainty or risk increases (Bhushan et al., 2026).

A second implication concerns trust. Trust should not be equivalent to authentication. Authentication establishes an identity assertion, whereas trust management concerns whether observed behavior continues to justify the privileges granted to that identity. This distinction becomes increasingly important when agents persist across multiple tasks and sessions.

Nevertheless, several limitations remain. First, the methodological transfer from similarity-learning research to authorization requires empirical validation in real agentic environments. The provided references do not themselves establish performance benchmarks for authorization systems. Consequently, the proposed framework should be considered an architecture and methodological proposition rather than a demonstrated production system.

Second, adaptive risk models can produce false positives. A legitimate but unusual action may be interpreted as anomalous, causing unnecessary denial or escalation. Conversely, insufficiently sensitive behavioral models may fail to identify sophisticated misuse.

Third, dynamic trust introduces governance questions. Organizations must determine which behaviors increase or decrease trust, how quickly trust should recover after anomalies, and who is responsible for changing risk thresholds. Without governance controls, an adaptive authorization mechanism could itself become an opaque security dependency.

Finally, computational overhead must be considered. The combination of contextual representation, attention, risk assessment, policy evaluation, and continuous monitoring is more expensive than static access checks. Future implementations should therefore investigate lightweight representations and efficient policy evaluation mechanisms.

CONCLUSION

This paper presented a Context-Aware Authorization Framework for Agentic AI that integrates contextual representation, adaptive risk assessment, policy compatibility evaluation, and continuous trust management. The central argument is that autonomous AI agents require authorization decisions that account not only for who is requesting access but also for what the agent is doing, why it is doing it, what resource is involved, what contextual conditions exist, and how trustworthy its recent behavior has been.

The provided literature on attribute-specific learning, attention mechanisms, similarity modeling, compatibility prediction, and iterative learning offers transferable methodological foundations for constructing such a framework. Attribute-aware representation enables fine-grained contextual modeling, attention mechanisms support task-sensitive feature weighting, similarity approaches provide a conceptual model for policy compatibility, and iterative learning principles support evolving trust assessment.

The proposed architecture extends these principles into a four-layer authorization process. Instead of producing only binary outcomes, it enables adaptive decisions such as restricted access and human escalation. This is particularly appropriate for agentic AI, where the operational consequences of actions can vary considerably across contexts.

The primary research contribution is therefore a unified conceptual model that treats authorization as a dynamic relationship among context, risk, policy, and trust. Future research should implement the framework in realistic agentic environments and evaluate authorization accuracy, latency, false-positive and false-negative rates, explainability, policy consistency, and resilience against adversarial behavior. Empirical comparison with conventional role-based and attribute-based mechanisms

would further establish when contextual authorization provides measurable security advantages.

REFERENCES

1. K. Ak, J. Lim, J. Tham, and A. Kassim, "Efficient multi-attribute similarity learning towards attribute-based fashion search," in Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV), 2018, pp. 1671–1679.
2. J. Dong, "Fine-grained fashion similarity prediction by attribute-specific embedding learning," IEEE Trans. Image Process., vol. 30, pp. 8410–8425, 2021.
3. S. Goenka, "FashionVLP: Vision language transformer for fashion retrieval with feedback," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2022, pp. 14105–14115.
4. X. Ji, W. Wang, M. Zhang, and Y. Yang, "Cross-domain image retrieval with attention modeling," in Proc. ACM Int. Conf. Multimedia (ACMMM), 2017, pp. 1654–1662.
5. S. Li, "COCAS+: Large-scale clothes-changing person re-identification with clothes templates," IEEE Trans. Circuits Syst. Video Technol., vol. 33, no. 4, pp. 1839–1853, Apr. 2023.
6. Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "DeepFashion: Powering robust clothes recognition and retrieval with rich annotations," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 1096–1104.
7. Z. Ma, "Fine-grained fashion similarity learning by attribute-specific embedding network," in Proc. AAAI Conf. Artif. Intell. (AAAI), vol. 34, no. 7, 2020, pp. 11741–11748.
8. D. Song, J. Zeng, M. Liu, X. Li, and A. Liu, "Fashion customization: Image generation based on editing clue," IEEE Trans. Circuits Syst. Video Technol., vol. 34, no. 6, pp. 4434–4444, Jun. 2024.
9. H. Su, P. Wang, L. Liu, H. Li, Z. Li, and Y. Zhang, "Where to look and how to describe: Fashion image retrieval with an attentional heterogeneous bilinear network," IEEE Trans. Circuits Syst. Video Technol., vol. 31, no. 8, pp. 3254–3265, Aug. 2021.
10. Y. Tang, L. Zhang, Y. Yuan, and Z. Chen, "Describe fashion products via local sparse self-attention mechanism and attribute-based re-sampling strategy," IEEE Trans. Circuits Syst. Video Technol., vol. 33, no. 7, pp. 3409–3424, Jul. 2023.
11. A. Veit, S. Belongie, and T. Karaletsos, "Conditional similarity networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017, pp. 830–838.

12. Y. Wan, K. Yan, C. Yan, and B. Zhang, "Learning attribute-guided fashion similarity with spatial and channel attention," *J. Exp. Theor. Artif. Intell.*, vol. 36, no. 5, pp. 703–719, 2022.
13. L. Xiao and T. Yamasaki, "Boosting fine-grained fashion retrieval with relational knowledge distillation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRWS)*, 2024, pp. 8229–8234.
14. L. Xiao, X. Zhang, and T. Yamasaki, "Learning fashion compatibility with color distortion prediction," in *Proc. IEEE 6th Int. Conf. Multimedia Inf. Process. Retrieval (MIPR)*, 2023, pp. 81–84.
15. L. Xiao and T. Yamasaki, "SAT: Self-adaptive training for fashion compatibility prediction," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, 2022, pp. 2431–2435.
16. C. Yan, K. Yan, Y. Zhang, Y. Wan, and D. Zhu, "Attribute-guided fashion image retrieval by iterative similarity learning," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, 2022, pp. 1–6.
17. S. Zhang, Z. Song, X. Cao, H. Zhang, and J. Zhou, "Task-aware attention model for clothing attribute prediction," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 4, pp. 1051–1064, Apr. 2020.
18. B. Bhushan, K. Pappu, V. Jadav and N. Jaiswal, "Risk-Adaptive Authorization for Agentic AI Systems," *SoutheastCon 2026*, Huntsville, AL, USA, 2026, pp. 1-6, doi: 10.1109/SoutheastCon63549.2026.11476129