

CampusSecure: AI That Protects Access

Sahil Shashikant Thorve Department of Computer Engineering Alard College of Engineering and Management	Rushikesh Kailas Dube Department of Computer Engineering Alard College of Engineering and Management
Manmath Vishvanath Somure Department of Computer Engineering Alard College of Engineering and Management	Aditya Dileep Gayakwad Department of Computer Engineering Alard College of Engineering and Management
Prof. Vrushali More Department of Computer Engineering Alard College of Engineering and Management	

Abstract

Access control plays a central role in protecting digital systems, especially in environments like academic institutions where multiple users interact with sensitive data on a daily basis. While models such as Role-Based Access Control (RBAC) have been widely adopted due to their structured approach, they are not without limitations. As systems grow more complex, static access control methods often struggle to handle dynamic user behavior and evolving security threats.

This review examines the progression of access control models, starting from traditional approaches like Discretionary Access Control (DAC) and Mandatory Access Control (MAC), and moving toward more flexible frameworks such as RBAC and Attribute-Based Access Control (ABAC). In addition, recent developments involving artificial intelligence are explored, particularly in the context of behavioral analysis and anomaly detection.

Rather than focusing only on theoretical strengths, this paper also highlights practical challenges observed in real-world implementations, especially in academic systems. Based on the analysis, it becomes evident that no single model is sufficient on its own. A combination of structured access control and intelligent monitoring appears to be a more effective direction for modern systems.

Keywords

Access Control, RBAC, ABAC, Artificial Intelligence, Cybersecurity, Behavioral Analysis, Academic Systems

1. Introduction

Over the last decade, academic institutions have increasingly shifted toward digital platforms for managing their operations. Tasks that were once handled manually — such as maintaining student records, managing grades, and coordinating administrative workflows — are now performed through centralized systems. While this transition has improved efficiency, it has also introduced new security concerns.

One of the key challenges in such systems is controlling access to information. Not every user should have the same level of access, and even authorized users can sometimes misuse their privileges. This is where access control mechanisms come into play.

Earlier models like Discretionary Access Control (DAC) allowed users to decide who could access their resources.

While flexible, this approach often lacked proper oversight. On the other hand, Mandatory Access Control (MAC) enforced strict rules, but at the cost of flexibility.

These limitations eventually led to the adoption of Role-Based Access Control (RBAC), which offered a more organized way of managing permissions [1], [2].

RBAC works well in structured environments, but problems start to appear when systems become more dynamic. For example, a user with valid access rights may still behave in a suspicious way, such as accessing data at unusual times or making repeated requests. Traditional models are not designed to detect such behavior.

This is where newer approaches, particularly those involving artificial intelligence, are gaining attention. Instead of relying only on predefined roles, these systems attempt to understand how users behave and identify patterns that may indicate potential threats.

The purpose of this review is not just to describe different access control models, but to examine how they perform in practical scenarios, especially within academic systems. By analyzing existing approaches and their limitations, the paper aims to identify gaps and explore possible directions for improvement.

2. Evolution of Access Control Models

2.1 Early Models: DAC and MAC

Discretionary Access Control (DAC) was one of the earliest approaches, where control over access was left to the resource owner. While this gave flexibility, it also created inconsistency. Different users could define permissions in different ways, which made enforcement difficult at scale.

Mandatory Access Control (MAC), in contrast, attempted to solve this by enforcing strict policies based on predefined classifications. This worked well in highly secure environments, but it came with its own drawback — lack of flexibility. In systems where roles and responsibilities frequently change, rigid policies become a limitation rather than a strength.

DAC and MAC were among the earliest access control approaches proposed for securing information systems [3].

In practice, neither of these models scaled effectively when applied to large systems with diverse user groups, such as universities.

Problem:

Neither model scales well in environments with thousands of users and complex relationships.




Figure 2.1 Evolution of Access Control Models

2.2 Role-Based Access Control (RBAC)

RBAC introduced a more structured way of managing permissions by linking them to roles instead of individual users [1], [2]. This simplified administration and made it easier to enforce consistent policies.

However, while RBAC solves the problem of scalability, it introduces a different issue — rigidity. Roles are typically predefined, and once assigned, they do not change dynamically. This becomes problematic in situations where user behavior does not align with their assigned role.

For example, a user with legitimate access may still perform actions that are unusual or potentially harmful. RBAC does not account for this because it focuses only on *who* the user is, not *how* they are behaving.




Figure 2.2 Basic RBAC Model

Limitations of RBAC

- Static role definitions
- Role explosion in large systems
- No behavioral awareness
- Cannot adapt to dynamic threats

➤ This is where most systems fail.

2.3 Attribute-Based Access Control (ABAC)

To address some of the limitations of RBAC, Attribute-Based Access Control (ABAC) was introduced. Instead of relying solely on roles, ABAC evaluates multiple attributes such as user identity, resource type, and environmental conditions. ABAC extends traditional access control by incorporating attributes into authorization decisions [4].

This makes ABAC significantly more flexible. Access decisions can change depending on context, which is useful in dynamic environments. However, this flexibility comes at a cost. Defining and managing policies in ABAC systems can quickly become complex, especially as the number of attributes increases.

In many cases, organizations struggle to implement ABAC effectively because the policies become difficult to maintain and debug.




Figure 2.3 ABAC Framework

Limitations:

- Complex policy management
- Difficult to implement at scale

2.4 Hybrid Approaches

Because no single model is sufficient, hybrid approaches have started to emerge. These typically combine RBAC with ABAC, using roles for basic structure and attributes for dynamic adjustments.

While this approach improves flexibility, it still does not fully address behavioral aspects. The system can decide *when* and *where* access is allowed, but it still does not evaluate *how* the user interacts with the system over time.

Relationship-Based Access Control (ReBAC)

- Uses relationships between users and resources
- More granular control

Risk-Based Access Control

- Evaluates risk dynamically before granting access
- Used in banking and cloud systems

2.5 Shift Toward Intelligent Access Control

More recent approaches attempt to move beyond static rules by incorporating elements of intelligence. Instead of relying only on predefined policies, these systems monitor user activity and look for patterns.

This marks a shift from rule-based access control to behavior-aware systems. The idea is simple: even if a user has permission, their actions should still be evaluated for consistency and risk.

This shift is particularly relevant in environments where insider threats or credential misuse are common concerns.

3. Literature Review

The domain of access control has evolved significantly over the past decades, with researchers proposing various models to address limitations in scalability, flexibility, and security. This section presents a critical analysis of key contributions in the field, focusing on RBAC, AI integration, and secure academic systems.

3.1 Foundational Work in RBAC

Early work of Role-Based Access Control was formalized by **Ferraiolo and Kuhn [1]**, providing a structured approach to access management by linking users to roles and roles to permissions. This model simplified administration and reduced the complexity associated with user-level permissions.

Later, **Sandhu et al.** extended the RBAC model by introducing role hierarchies and constraints such as separation of duties [2]. This enhancement improved scalability and made RBAC suitable for large organizations.

➤ Critical Insight:

While RBAC simplifies permission management, it assumes that roles are static and predefined, which is unrealistic in dynamic environments.




Figure 3.1 AI-Based Behavioral Analysis

3.2 RBAC in Academic Systems

Several studies have explored the use of RBAC in educational platforms. These implementations generally focus on organizing access based on institutional roles such as students, faculty, and administrators.

While this approach works for basic access control, it often stops at the surface level. Most systems enforce permissions at the application layer without considering deeper issues such as API-level vulnerabilities or data-level protection.

This creates a situation where the system appears secure, but can still be bypassed under certain conditions.

- **Liu et al. (2017)** proposed an RBAC-based system tailored for academic administration. Their model aligned roles with institutional hierarchies, improving clarity and reducing conflicts.
- **Ratna et al. (2025)** focused on collaborative academic environments, demonstrating how RBAC can regulate access to course materials and shared resources.

Limitation Across Studies:

- Lack of real-time monitoring
- No database-level enforcement
- Vulnerable to API-level bypass

Translation:

They solved structure, not security.

3.3 Research on ABAC and Hybrid Models

To overcome RBAC limitations, researchers introduced **Attribute-Based Access Control (ABAC)**.

ABAC systems evaluate multiple attributes, including:

- User identity
- Resource sensitivity

- Environmental conditions

Research on ABAC highlights its ability to provide fine-grained control. Studies show that incorporating attributes allows systems to adapt to different scenarios more effectively than RBAC alone.

However, many of these studies also acknowledge the practical challenges of implementing ABAC. In real-world systems, managing a large number of attributes and policies can become overwhelming.

Hybrid models attempt to address this by combining RBAC and ABAC, but they still face complexity issues and often lack clear implementation strategies.

➤ **Critical Insight:**

Flexibility increases → complexity explodes.

3.4 AI Integration in Access Control

Recent studies have started exploring the integration of artificial intelligence into access control systems. These approaches focus on analyzing user behavior rather than relying solely on predefined rules. Recent studies have even demonstrated that AI-driven access control systems can improve anomaly detection and reduce unauthorized access attempts [11], [12].

Instead of asking “Does the user have permission?”, AI-based systems also ask “Is the user behaving normally?”

This additional layer of analysis allows systems to detect:

- Unusual access patterns
- Rapid or repetitive actions
- Attempts to access restricted areas

However, many of these approaches remain at a conceptual level. Real-world implementations are limited due to challenges such as data availability and model accuracy.

- Research and Studies on **behavior-based access control** demonstrate how machine learning algorithms can identify abnormal user activities.
- **Akuthota (2025)** highlighted the effectiveness of AI-enhanced RBAC in reducing unauthorized access attempts in cloud environments.

Key Advantage:

AI shifts access control from:

- Static decision → Dynamic decision

Major Limitation:

- Requires training data
- Risk of false positives

3.5 Database-Level Security Approaches

Another important area of research focuses on securing data at the database level. Techniques such as **Row-Level Security (RLS)** allow systems to enforce access control directly within the database. PostgreSQL Row-Level Security provides fine-grained authorization directly at the database layer [7].

This approach reduces dependency on the application layer and provides an additional layer of protection. However, it

is often not integrated with higher-level access control mechanisms, which limits its effectiveness.

Advantage: Prevents data leakage even if API fails

Gap in Existing Work:

- Rarely integrated with AI or RBAC in academic systems

3.6 Security Standards and Frameworks

Based on the reviewed studies, several patterns become clear:

- RBAC is widely used but lacks adaptability
- ABAC improves flexibility but increases complexity
- AI introduces behavioral awareness but is not widely implemented
- Database-level security is effective but underutilized
- Most systems do not combine these approaches of **multi-layer integration**

effectively

4. Comparative Analysis

When comparing different access control models, it becomes clear that each approach solves a specific problem, but none of them provides a complete solution on its own.

For instance, DAC is simple and easy to implement, but it lacks strong control mechanisms. MAC, on the other hand, provides strict enforcement, yet it becomes impractical in environments where flexibility is required. RBAC strikes a balance between these two, which explains its widespread adoption.

However, once systems start dealing with dynamic user behavior, RBAC alone begins to fall short. ABAC attempts to solve this by introducing contextual decision-making, but in doing so, it also introduces complexity that is not always manageable in real-world systems.

AI-based approaches represent a different direction altogether. Instead of defining rules in advance, they attempt to observe patterns and make decisions based on behavior. This makes them powerful, but also dependent on proper tuning and data.

One important observation here is that these models are not necessarily competing with each other. In many cases, they are more effective when used together rather than in isolation.

4.1 Discussion

A closer look at these models suggests that the problem is not with any single approach, but with how they are applied. Systems often rely too heavily on one method, which creates blind spots.

For example, a system using only RBAC may successfully restrict access based on roles, but it cannot detect misuse within those roles. Similarly, a system relying only on AI may detect anomalies, but without a structured permission model, it may lack control.

This indicates that access control should not be treated as a single-layer mechanism. Instead, it should be designed as a combination of complementary approaches.

Table: Research Gap Analysis

Existing Research	Contribution	Limitation	Research Gap
DAC	Flexible ownership control	Weak centralized security	Poor scalability
MAC	Strong enforcement policy	Rigid implementation	Lack of flexibility
RBAC	Structured permissions	Static roles	No behavioral awareness
ABAC	Context-aware decisions	Complex management policy	Difficult large-scale deployment
AI-Based Systems	Dynamic monitoring	False positives	Limited real-world adoption
RLS-Based Systems	Strong data security	Database-specific	Limited integration with AI

5. Applications in Academic Systems

Academic platforms present a unique set of challenges when it comes to access control. Unlike many other systems, they involve multiple types of users interacting with shared resources on a regular basis. Multiple user roles (students, faculty, admins)

Students, faculty members, and administrators all require different levels of access. At the same time, the data being handled — such as grades, personal information, and institutional records — is highly sensitive.

In such environments, even small weaknesses in access control can lead to serious consequences. For example, unauthorized modification of grades or exposure of student records can directly impact both individuals and the institution.

5.1 Practical Limitations Observed

In practice, many academic systems rely heavily on frontend controls. While this may work for basic scenarios, it creates a false sense of security. If backend APIs are not properly protected, these controls can be bypassed.

Another issue is the lack of behavioral monitoring. Most systems verify whether a user has permission, but they do not evaluate how that permission is being used. This makes it difficult to detect insider threats or compromised accounts.

Additionally, data protection is often handled at the application level, leaving the database layer relatively unprotected. This becomes a problem if the application layer is compromised.

➤ Basically:

Most systems are secure only “on the surface.”

6. Challenges and Research Gaps

Even though access control mechanisms have evolved significantly, several practical challenges continue to limit their effectiveness in real-world systems. These challenges are not always visible in theoretical models but become evident during implementation and deployment.

6.1 Lack of Dynamic Adaptability

Most traditional access control systems operate on predefined rules and static configurations. While this works in controlled environments, it becomes ineffective in situations where user behavior changes over time.

For example, a user who normally accesses the system during working hours may suddenly start making requests late at night or from unusual locations. Traditional models do not interpret this as suspicious because the user still holds valid permissions.

This highlights a critical gap — access control systems often verify identity, but they fail to evaluate context and behavior dynamically.

6.2 Insider Threats and Credential Misuse

One of the most underestimated risks in access control is insider threat. Unlike external attackers, insiders already have legitimate access, which makes detection significantly more difficult.

In many systems, once a user is authenticated, their actions are rarely questioned unless they explicitly violate predefined rules. However, misuse of valid credentials — whether intentional or due to compromise — can lead to serious security incidents.

This exposes a fundamental limitation:

Access control models focus on authorization, but not on **continuous trust evaluation**.

6.3 Complexity of ABAC Systems

While ABAC introduces flexibility, it also brings complexity that grows rapidly with scale.

As more attributes are added to improve decision-making, the number of policies increases, often leading to:

- Difficult policy management
- Increased chances of misconfiguration
- Performance degradation

In large systems, maintaining consistency across these policies becomes a challenge, and debugging access issues can become time-consuming.

This makes ABAC powerful in theory, but difficult to manage in practice.

6.4 Limited Integration of AI in Practical Systems

AI-driven access control is often presented as a solution to many existing problems. However, its practical implementation introduces its own set of challenges.

Firstly, AI models require sufficient and high-quality data to learn user behavior accurately. In many systems, such datasets are either unavailable or incomplete.

Secondly, there is a risk of false positives, where legitimate user actions are flagged as suspicious. This can

negatively impact usability and user trust.

Finally, AI models may lack transparency, making it difficult to explain why a particular decision was made. This becomes a concern in systems where accountability is important.

6.5 Weak Data-Level Enforcement

A major issue observed in many implementations is the lack of coordination between different layers of the system.

For example:

- Frontend may restrict access visually
- Backend may enforce RBAC
- Database may have minimal protection

However, these layers often operate independently rather than as a unified system.

As a result, if one layer is bypassed (such as through direct API access), the entire system becomes vulnerable.

6.6 Insufficient Data-Level Protection

Most access control implementations focus on application-level security, leaving the database layer relatively unprotected.

This creates a critical vulnerability. If an attacker manages to bypass the application logic, they may gain direct access to sensitive data.

Mechanisms such as Row-Level Security (RLS) address this issue, but they are still underutilized in many systems.

6.7 Absence of Continuous Monitoring Mechanisms

Another gap is the lack of continuous monitoring. Traditional systems perform checks only at the time of login or access request.

Once access is granted, there is little to no evaluation of ongoing user activity. This creates a window where malicious actions can occur without detection.

Continuous monitoring, combined with behavioral analysis, is necessary to address this issue effectively.

7. Proposed Direction: AI-Driven Multi-Layer Access Control

Based on the identified challenges, it is evident that future access control systems must move beyond static and isolated approaches. A more effective solution lies in combining multiple techniques into a cohesive framework.

7.1 Hybrid Access Control Model




Figure 7.1 Proposed Multi-Layer Security Architecture

A hybrid approach that combines RBAC, AI-based analysis, and database-level security can address the limitations of individual models.

- RBAC provides a structured foundation for managing permissions
- AI introduces behavioral awareness and anomaly detection
- Database-level controls ensure that data access is restricted even if other layers fail This combination creates a more balanced system that is both manageable and adaptable

7.2 Role of Artificial Intelligence

One of the key directions for improvement is the integration of behavioral intelligence into access control decisions. Instead of relying solely on predefined roles or attributes, systems should analyze how users interact with resources over time.

This includes:

- Monitoring access frequency
- Identifying unusual patterns
- Detecting deviations from normal behavior

By incorporating these factors, access control becomes more dynamic and responsive.



Figure 7.2 Continuous Verification in Zero Trust Architecture

7.3 Multi-Layer Security Architecture

A critical requirement for modern systems is the implementation of multi-layer security.

Each layer should independently enforce security while also contributing to a unified framework:

- Presentation Layer → Limits exposure at the interface level
- Application Layer → Enforces RBAC and API security
- AI Layer → Detects anomalies and suspicious behavior
- Data Layer → Ensures fine-grained access control

This approach minimizes the risk of single-point failure and improves overall system resilience.

7.4 Continuous Verification and Zero Trust Approach

Future systems should adopt a continuous verification model rather than relying on one-time authentication.

This aligns with the principles of Zero Trust Architecture, where:

- Every request is verified
- Trust is not assumed based on previous actions
- Access decisions are continuously evaluated

This approach is particularly effective in environments with high security requirements.

7.5 Practical Implementation Considerations

While the proposed approach is effective, it must also be practical. Key considerations include:

- Balancing security with system performance
- Minimizing false positives in AI detection
- Ensuring ease of policy management
- Maintaining system scalability

Without addressing these factors, even the most advanced models may fail in real-world deployment.

7.6 Relevance to Academic Systems

Academic systems benefit significantly from this approach due to their unique characteristics.

- Multiple user roles
- Frequent data access

- High sensitivity of information

By combining structured access control with behavioral monitoring and data-level protection, institutions can achieve a more secure and reliable system.

7.7 Future Research Directions

Future research can focus on developing adaptive access control systems capable of continuously learning from user behavior. The integration of machine learning techniques can improve anomaly detection accuracy while reducing false positives. Emerging concepts such as Zero Trust Architecture, federated identity management, blockchain-based audit logging, and risk-adaptive access control offer promising opportunities for enhancing system security.

Further studies may also investigate the practical deployment of AI-driven access control systems within large-scale academic environments, where performance, privacy, and usability must be balanced carefully.

8. Conclusion

Access control has evolved significantly over time, but the challenges it aims to address have also become more complex. While traditional models like RBAC continue to play an important role, they are no longer sufficient on their own.

This review shows that modern systems require a combination of structured access control, contextual decision-making, and behavioral analysis. Each of these components addresses a different limitation, and together they provide a more comprehensive solution.

In academic systems, where multiple users interact with sensitive data, this becomes even more important. A layered approach not only improves security, but also makes the system more resilient to different types of threats.

Looking ahead, further improvements are likely to come from better integration of these approaches, along with advancements in machine learning and real-time analysis. However, the key idea remains the same — access control should not be static, but adaptive.

Overall, AI-driven multi-layer access control represents a critical step toward building secure, scalable, and resilient systems in an increasingly complex digital landscape.

9. REFERENCES

- [1] D. Ferraiolo and D. Kuhn, "Role-Based Access Controls," 1992.
- [2] R. Sandhu et al., "Role-Based Access Control Models," IEEE Computer, 1996.
- [3] M. Bishop, Computer Security: Art and Science, 2018.
- [4] V. C. Hu et al., Guide to Attribute Based Access Control, NIST SP 800-162.
- [5] OWASP Foundation, Application Security Verification Standard.
- [6] M. Jones et al., JSON Web Token (JWT), RFC 7519.
- [7] PostgreSQL Documentation, Row-Level Security Policies.
- [8] OpenID Foundation, OpenID Connect Core.
- [9] Google Cloud, Zero Trust Security Model.

- [10] Microsoft Security, Zero Trust Architecture.
- [11] N. Zhang et al., Behavior-Based Access Control Using Machine Learning.
- [12] S. Sharma and P. Gupta, Anomaly Detection in User Behavior Using AI.