Indo American Journal of Multidisciplinary Research and Review (IAJMRR)
ISSN: 2581 - 6292 / Impact Factor: 6.885 / Website: www.iajmrr.com

BENCHMARKING AUTHORIZATION DECISIONS FOR AUTONOMOUS SOFTWARE AGENTS USING HISTORICAL EMPLOYEE ACCESS REQUESTS

Divya Vishwanath* & Paul Johnson Asamoah**
* Janardhan Rai Nagar Rajasthan Vidyapeeth, Udaipur, Rajasthan, India
** Power Life University College, Accra, Ghana, West Africa

Cite This Article: Divya Vishwanath & Paul Johnson Asamoah, "Benchmarking Authorization Decisions for Autonomous Software Agents Using Historical Employee Access Requests", Indo American Journal of Multidisciplinary Research and Review, Volume 7, Issue 2, July - December, Page Number 140-150, 2023.
Copy Right: © IAJMRR Publication, 2023 (All Rights Reserved). This is an Open Access Article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Abstract:

This analysis examines whether apparent consistency in authorization benchmarks persists when category counts are compared with request volume. Using 32,769 employee access requests collected during 2010 and 2011 from the Amazon benchmark distributed through OpenML, we calculate category frequencies, approval proportions, outcome composition and request shares across resource, organizational and role classifications. Resource categories containing both approvals and denials constitute 11.20 percent of resources but contain 50.52 percent of requests and 82.76 percent of denials. Role title and role code generate identical category partitions, demonstrating informational redundancy. An exploratory exclusion of singleton resources increases the mixed category proportion to 22.44 percent while changing aggregate approval by only 0.17 percentage points. These patterns reflect unequal category sizes and the concentration of varied outcomes beneath aggregate approval dominance. The analysis refines authorization benchmark interpretation by separating classification consistency, observational support and request concentration. Governance and operational reviews should report category denominators alongside approval proportions and examine populous mixed categories before interpreting historical decisions as permission patterns.

Key Words: Access Control; Authorization Benchmarking; Category Sparsity; Descriptive Analysis; Request Concentration

1. Introduction:

Authorization records can suggest consistent access decisions while concealing substantial variation in the categories containing most requests. This distinction matters when organizations use historical approvals and denials to understand permission arrangements, review access practices or prepare data for automation. Counting how many categories contain one outcome answers a different question from counting how many requests occur within those categories. When category sizes differ, the two summaries can support different interpretations of the same authorization system. This analysis examines that distinction across resource identity, organizational context and role identity.

Attribute based access control provides the conceptual foundation for organizing these relationships. Authorization depends on the correspondence between requester characteristics, requested resources and applicable policies. Hu, Kuhn, and Ferraiolo, 2015, explain how discrete subject and object attributes support decisions based on combinations of relevant information. Resource identity identifies the object requested, organizational context locates the requester within company structures, and role identity describes functional classification. These dimensions provide complementary descriptions of a request, although their presence does not establish that each contributes distinct information or independently determines the outcome.

The closest literature establishes both the potential and the interpretive difficulty of learning authorization arrangements from existing records. Xu and Stoller, 2015, develop an algorithm that derives attribute-based policies from permission relationships and accompanying attributes. Their approach constructs, generalizes and simplifies candidate rules, demonstrating how observed information can support policy reconstruction. Servos and Osborn, 2017, situate such developments within a broader research agenda concerning authorization models, administration, auditability and attribute representation. Together, these contributions make the structure and meaning of recorded attributes central to access control research. They also motivate examining whether apparently different fields represent distinct classifications and whether observed decision consistency rests on substantial evidence. Sparse observations complicate that interpretation. Cotrini et al., 2018, show that authorization logs containing a limited subset of possible requests can support overly permissive mined rules and misleading validation results. Their evaluation uses organizational logs, including Amazon data, establishing that historical access records already constitute an empirical setting for authorization policy research. The unresolved descriptive question addressed here concerns how category coverage and request concentration jointly shape the interpretation of such records. Algorithmic policy quality and the distribution of observed decisions are related concerns, but they require different empirical summaries. A

category containing one approval provides a different evidential basis from a category containing repeated approvals and denials.

This analysis addresses that measurement question by separating category uniformity from the distribution of requests and denials. Category uniformity refers to whether an observed category contains approvals only, denials only or both outcomes. Request concentration concerns how much of the sample lies within each type of category. The contribution is a transparent comparison of these perspectives within an established benchmark. It also examines redundancy across role classifications and assesses how excluding singleton resources changes the apparent composition of authorization outcomes. This positioning extends the interpretation of benchmark evidence without claiming a new dataset or a newly identified causal mechanism.

The empirical setting consists of employee requests for access to corporate computer resources. The records originate from the Amazon Employee Access Challenge and are distributed through OpenML dataset 4135, version 1. The source identifies the collection period as 2010 and 2011. Each observation records an approval or denial alongside resource, managerial, departmental and role attributes. The verified analytical sample contains 32,769 complete labelled requests. Its outcome is the recorded administrative decision, allowing the analysis to examine how approvals and denials are distributed across the supplied classifications.

This study aims to establish how recorded authorization decisions vary across resource identity, organizational context and role identity, and how that interpretation changes when category counts are considered alongside request volume. The central research question asks whether categories displaying uniform outcomes also account for most observed requests. Three specific objectives operationalize this question: characterize approval and denial distributions across requested resources, compare these distributions across managerial and organizational classifications, and examine their variation across role titles, role families and role codes. The analysis uses categorical frequency tables, approval and denial proportions, checks of category correspondence and an exploratory sensitivity comparison excluding singleton resources.

The completed analysis identifies a substantial difference between the prevalence of mixed categories and their share of requests. Resource categories containing both outcomes represent 11.20 percent of resources but account for 50.52 percent of requests and 82.76 percent of denials. Role title and role code generate equivalent category partitions, while broader role family categories containing both outcomes encompass almost the entire request sample. Excluding singleton resources nearly doubles the proportion of mixed resource categories while changing the aggregate approval proportion by only 0.17 percentage points. These findings show that aggregate stability can coexist with substantial differences in the apparent structure of category outcomes.

The theoretical contribution concerns the empirical interpretation of authorization attributes: classification meaning, observational support and request concentration must be considered together. Practically, the findings support reporting category denominators alongside approval proportions, recognizing equivalent fields and identifying where recorded denials concentrate before drawing conclusions about permission patterns. This approach is relevant to organizations evaluating authorization records across different systems and institutional settings. Its transferable contribution is the comparison procedure and its interpretation; the observed frequencies remain specific to the benchmark examined.

2. Theory, Conceptual Framework and Research Questions Development:

This analysis draws on attribute-based access control to explain authorization decisions through relationships among requesters, organizational classifications and requested resources. Authorization depends on how these elements correspond to applicable permission rules. Hu et al., 2014, establish this relational logic by defining access decisions through the evaluation of subject attributes, object attributes and requested operations against policy. The theoretical implication is that authorization patterns may become clearer when request characteristics are considered together. This expectation guides the analysis without presuming that any individual attribute determines approval.

The framework comprises authorization request attributes and the recorded access authorization decision. Request attributes contain three observed dimensions: resource identity, organizational context and role identity. The outcome distinguishes approval, coded one, from denial, coded zero. The available records support examining how decision frequencies and proportions vary across these categorical dimensions. We formulate exploratory research questions because the identifiers distinguish categories without ranking their authority, relevance or entitlement.

Resource identity defines the object for which access is requested. Its theoretical mechanism arises from the correspondence between a requester's characteristics and the permissions attached to that object. Identical requester characteristics can therefore accompany different decisions when the requested resources differ. Resource identity provides the object specific context within which organizational and functional classifications acquire authorization relevance.

Hu, Kuhn, and Ferraiolo, 2015, explain that attribute-based access control accommodates combinations of subject and object information when evaluating permissions. This establishes a conceptual basis for expecting variation across resources. It does not establish that resource categories

necessarily display different approval proportions in the available sample. Similar proportions remain possible when requests involve shared permission arrangements or comparable requester compositions.

The RESOURCE field provides the observed measure of resource identity. The descriptive comparison examines the number of requests, approvals and denials associated with each resource, together with the corresponding proportions. Counts remain essential because identical proportions can arise from substantially different numbers of observations. The comparison concerns recorded resource categories rather than an assumed ordering of their identifiers.

RQ1: How do approval and denial frequencies and proportions vary across requested resources?

Organizational context locates the requester within managerial, departmental and company grouping structures. Its theoretical role differs from resource identity because it describes the organizational position associated with a request rather than the requested object. Shared placement may correspond to common work requirements and permission arrangements. However, employees within the same organizational category can occupy different roles and request different resources. Organizational membership therefore provides context without implying uniform authorization outcomes.

Hu, Ferraiolo, and Kuhn, 2019, emphasize that access control attributes require defined meanings, consistent interpretation and management by an appropriate authority. Their framework supports examining organizational classifications as information used to describe the conditions surrounding authorization. The expected pattern depends on whether those classifications distinguish relevant permission arrangements. Greater similarity within a category is plausible when its members share access requirements, while variation remains plausible when their functions and requested resources differ. Each field remains a separate categorical measure. The analysis compares decision counts and proportions across categories within these fields. This preserves their distinct meanings and avoids combining managerial, departmental and company classifications into an arbitrary numerical score.

RQ2: How do approval and denial frequencies and proportions vary across managers, departments and organizational role groupings?

Role identity describes the functional classification associated with a request. Its theoretical mechanism concerns the alignment between work responsibilities and resource access. Similar functions may correspond to comparable permission requirements across organizational units. Nevertheless, a shared role does not imply identical access to every resource. Role related convergence is therefore conditional on the composition of requests represented within each category.

Xu and Stoller, 2015, develop an approach for deriving access control policies from permission relationships and accompanying attributes. Their work provides a basis for investigating whether functional classifications correspond to recurring authorization patterns. The relevant expectation is patterned variation across categories, rather than a positive or negative relationship with numerical role codes. Similar decision proportions across roles remain a possible outcome when the recorded classifications provide limited separation of permission arrangements. These fields represent related descriptions of functional position. The analysis examines their category specific request counts and approval and denial proportions without treating overlapping classifications as independent confirmation of a mechanism. This approach retains the distinction between the theoretical relevance of functional responsibility and the decision patterns observed in the records.

RQ3: How do approval and denial frequencies and proportions vary across role titles, role families and role codes?

The conceptual framework connects the three observed attribute groups with the recorded authorization decision through descriptive association. Resource identity identifies the requested object, organizational context identifies company placement, and role identity identifies functional classification. Their theoretical relationships motivate comparisons of decision distributions without assigning causal meaning to differences between categories. The Joint Task Force, 2020, distinguishes access control functionality from assurance, reinforcing the importance of defining the empirical outcome precisely. Here, the outcome remains the recorded approval or denial decision.

The three research questions are resolved through frequency tables, cross tabulations and category specific proportions. Counts accompany proportions so that the comparison retains information about the volume of supporting observations. The framework accommodates convergence, variation and sparse category coverage without imposing directional effects on nominal identifiers. Its contribution is a structured account of how observed authorization decisions are distributed across the resource, organizational and functional dimensions of access requests.

3. Data:

This analysis establishes source provenance, documents sample retention and translates authorization attributes into observed categorical measures. The data audit links each descriptive comparison to verified request records, explicit denominators and reproducible calculations.

3.1 Data Sources, Population, and Sample Construction:

The analysis uses Amazon Employee Access, OpenML dataset 4135, version 1, distributed as ARFF file 1681098. The underlying provider is Amazon through the Employee Access Challenge hosted by Kaggle. OpenML provides the repository copy and machine-readable metadata. The repository release dates to November 2015, while the underlying requests were collected during 2010 and 2011. The

extraction timestamp is preserved separately in the accompanying workbook. Repository identification follows the distinction between datasets, metadata and computational tasks described by Vanschoren et al., 2013.

OpenML labels the licence Public without identifying a standard reuse licence. Public availability therefore documents accessibility rather than establishing unrestricted redistribution rights. The repository description states that training and test records were merged, but the downloaded file contains 32,769 labelled observations. Sample construction follows the verified file contents rather than inferring additional observations from that description.

The unit of analysis is one recorded employee request for access to a computer resource. The empirical setting is corporate information access administration. The data form a cross section of historical requests with irregular occurrence, rather than an annual panel or time series. Individual request dates and geographical identifiers are absent. Consequently, the collection window identifies overall coverage without supporting country comparisons or annual aggregation. These boundaries reflect the distinction between operational records and a sampling design emphasized by Hand, 2018.

The target population comprises employee authorization requests within the corporate access system represented by the benchmark. The accessible sampling frame consists of the published labelled records. Its coverage of the underlying corporate population cannot be quantified. We retain the complete accessible frame rather than selecting a probability sample. As boyd and Crawford, 2012, explain, extensive digital coverage does not establish population representativeness.

Eligibility requires a recorded authorization outcome and the nine supplied request attributes. All 32,769 records satisfy these conditions. No eligibility revisions or exclusions were introduced during the audit. The original extraction process remains distinct from the retention rules applied here, consistent with the importance of understanding data generation mechanisms identified by Lazer et al., 2014. The resulting sample supports description of recorded authorization patterns, preserving the distinction among descriptive, explanatory and predictive objectives established by Shmueli, 2010.

Table 1: Sample construction

Stage	Units	Observations	Reason or Rule
Published labelled frame	32,769 requests	32,769	OpenML dataset 4135, version 1
Excluded for missing substantive fields	0 requests	0	All nine attributes and the outcome are populated
Excluded as exact duplicates	0 requests	0	No identical records across the ten source fields
Retained after quality checks	32,769 requests	32,769	Original records and outcome values preserved
Final descriptive sample	32,769 requests	32,769	Complete accessible frame used throughout

Counts refer to requests rather than distinct employees. The source does not provide an employee identifier from which a verified employee population could be constructed.

3.2 Data Integration, Cleaning, and Quality Assurance:

The analysis uses one substantive dataset. No external linkage, temporal merge or supplementary record matching was performed. The nine original attribute fields were preserved, and the source outcome named target was renamed ACTION without changing its values. A sequential source row identifier records the original observation order. This separation of original information from processing metadata follows the transparent data management practices discussed by Wilson et al., 2017.

Each request occupies one row, and each field occupies one column. This structure follows the observational organization described by Wickham, 2014, and permits category summaries without restructuring the source observations. The audit confirmed 32,769 populated values in every substantive field. No deletion, imputation, interpolation or replacement with zero was required.

Duplicate checks compared all nine attributes together with ACTION. A separate consistency check examined whether identical attribute profiles received conflicting outcomes. Both checks returned zero cases. Repeated values within individual fields were retained because multiple requests can legitimately share a resource, manager or role category. This distinction between repeated categories and erroneous records is consistent with the constraint-based approach to data quality described by Chu et al., 2016.

The audit verified that ACTION contains only zero and one. Attribute identifiers were retained as nominal categories and stored as text in the workbook. Their numerical appearance was not used to define magnitude, distance or an outlier threshold. Rare categories remained in the sample because low frequency is an observed feature rather than evidence of a recording error. This treatment preserves the distinction between completeness, consistency and substantive accuracy emphasized by Sadiq, 2013.

Quality checks also reconciled category totals with the complete sample and outcome totals. Every attribute summary reproduced 32,769 requests, 30,872 approvals and 1,897 denials. Recalculation of the stored profile hashes and partition assignments produced zero mismatches. The checks did not alter

records or category definitions. Their executable implementation preserves the connection between inputs, processing steps and reported quantities advocated by Sandve et al., 2013.

3.3 Variable Measurement, Construction, and Validation:

Authorization request attributes comprise resource identity, organizational context and role identity. These dimensions organize observed categorical information rather than indicators of an estimated latent construct. The distinction follows the attribute based authorization logic described by Hu et al., 2014, in which request characteristics acquire meaning through their relationship with access rules.

Table 2: Variable dictionary

Variable or Indicator	Analytical Role	Definition and Construction	Source and Identifier	Unit and Frequency	Transformation and Direction	Missing Value Rule	Coverage and Limitation
ACTION	Outcome	Recorded access decision	OpenML 4135 v1; target	Binary value per request; irregular	Renamed only; one means approved, zero means denied	No missing values; all retained	32,769 requests
RESOURCE	Resource identity	Requested resource identifier	OpenML 4135 v1; RESOURCE	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	7,518 categories; resource descriptions unavailable
MGR_ID	Organizational context	Identifier of the requester's manager	OpenML 4135 v1; MGR_ID	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	4,243 categories; does not identify the requester
ROLE_ROLLUP_1	Organizational context	First company role grouping	OpenML 4135 v1; ROLE_ROLLUP_1	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	128 categories
ROLE_ROLLUP_2	Organizational context	Second company role grouping	OpenML 4135 v1; ROLE_ROLLUP_2	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	177 categories
ROLE_DEPTNAME	Organizational context	Department classification	OpenML 4135 v1; ROLE_DEPTNAME	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	449 categories
ROLE_TITLE	Role identity	Business title classification	OpenML 4135 v1; ROLE_TITLE	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	343 categories
ROLE_FAMILY_DESC	Role identity	Extended role family classification	OpenML 4135 v1; ROLE_FAMILY_DESC	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	2,358 categories
ROLE_FAMILY	Role identity	Role family classification	OpenML 4135 v1; ROLE_FAMILY	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	67 categories
ROLE_CODE	Role identity	Company role code	OpenML 4135 v1; ROLE_CODE	Category per request; irregular	Preserved as nominal text; no ordering	No missing values	343 categories; maps one to one with ROLE_TITLE

The measurement audit verified a one-to-one correspondence between ROLE_TITLE and ROLE_CODE. Their separate summaries therefore represent equivalent partitions of the observations rather than independent evidence. Other attributes retain their original category boundaries. This preserves attribute meanings rather than introducing undocumented recoding, consistent with the semantic considerations developed by Hu, Ferraiolo, and Kuhn, 2019.

Validation concerns field completeness, admissible outcome values, category correspondence and reconciliation of counts. The grouping of attributes does not create an additive authorization score. Xu and Stoller, 2015, distinguish attribute information from the permission relationships used to derive policies, supporting the retention of these separate empirical objects.

ACTION has a valid count of 32,769, a mean of 0.9421, a sample standard deviation of 0.2335, a median of one and a range from zero to one. Approvals account for 94.21 percent of requests and denials for 5.79 percent. For nominal attributes, category counts and frequencies provide meaningful summaries; arithmetic means of their codes do not. The outcome retains its recorded administrative meaning, consistent with the distinction between access control functionality and assurance made by the Joint Task Force, 2020.

Three auxiliary fields support traceability. SOURCE_ROW preserves observation order, PROFILE_SHA256 records a deterministic attribute profile hash, and PARTITION records the corresponding allocation flag. These fields do not enter the substantive descriptive comparisons.

3.4 Analytical Framework, Identification, and Reproducibility:

The implemented analysis is descriptive. It examines how approval and denial counts and proportions vary across resource, organizational and role categories. Each field is summarized separately using the same complete sample. The quantities describe requests in the published benchmark rather than effects of changing a requester's characteristics. This distinction between reproducible calculations and the interpretation of their results follows Goodman et al., 2016.

Equation 1 defines the category specific approval and denial proportions:

$$\hat{p}_{jc} = \frac{A_{jc}}{N_{jc}}, \qquad \hat{d}_{jc} = \frac{D_{jc}}{N_{jc}}$$

Here, j identifies an attribute field, cidentifies an observed category, $N_{jc}$is its request count, A_{jc} is its approval count, and $D_{jc}$is its denial count. The quantities $\hat{p}_{jc}$and $\hat{d}_{jc}$are the corresponding proportions. Percentages equal these proportions multiplied by 100. Every reported category has a positive request count.

Category frequencies accompany all proportions. This is necessary because 3,766 of the 7,518 resource categories contain one request. Singleton categories also occur in managerial and role classifications. An observed proportion of zero or one in such a category summarizes one decision and does not establish a stable authorization rule. Keeping the denominator visible supports the distinction between an exact computational result and its evidential scope discussed by Peng, 2011.

The executable audit reconstructs category counts directly from the supplied workbook. It checks substantive missingness, duplicate records, conflicting profiles, outcome totals and deterministic allocation fields. It also produces nine category summary files with request counts, approval counts, denial counts and both proportions. These explicit inputs and outputs implement the disclosure principle emphasized by Stodden et al., 2016.

The workbook preserves the source identifier, variable dictionary, transformation record and source file checksum. The accompanying reproduction script supplies the calculations needed to regenerate the audit and category summaries. Providing executable procedures alongside the records addresses the distinction between a general availability statement and usable computational materials identified by Stodden et al., 2018.

The descriptive questions remain exploratory. The completed audit introduced no outcome dependent exclusions, category pooling or changes to recorded decisions. All comparisons retain the same observation base and variable definitions. This explicit account of analytical choices supports the transparent reporting advocated by Munafò et al., 2017, while keeping the interpretation aligned with the observed request categories and administrative outcomes.

4. Methods:

4.1 Research Design and Analytical Scope:

The analysis uses a descriptive cross-sectional design to examine authorization outcomes across resource identity, organizational context and role identity. The unit of analysis is one employee access request. The quantities of interest are the observed approval and denial proportions within each category and the distribution of requests across categories containing different outcome combinations. This design addresses the three research questions through comparisons of empirical distributions. The distinction between descriptive, explanatory and predictive objectives follows Shmueli, 2010.

The analytical sample follows the construction documented in Table 1 and comprises 32,769 complete requests. Each request contributes once to each attribute specific summary. Comparisons use variation between categories within the same field. All substantive fields retain the definitions and coding documented in Table 2.

Interpretation remains confined to the published request records. Differences between category proportions can reflect the composition of requesters, requested resources and organizational classifications. The comparisons describe these differences without isolating the contribution of one characteristic from the others. Consequently, the results characterize observed authorization patterns rather than effects of changing category membership.

4.2 Measurement and Descriptive Specification:

The recorded outcome distinguishes approval from denial. Resource identity contains one categorical field, organizational context contains four fields, and role identity contains four fields. Each field is summarized separately. Identifiers retain their nominal meaning, so numerical differences between codes carry no analytical interpretation.

Category specific approval and denial proportions follow Equation 1. Their denominators are the numbers of requests observed within the corresponding categories. Overall proportions are calculated from total approvals or denials divided by total requests. They therefore preserve the contribution of every request rather than assigning equal weight to categories of different sizes.

Each observed category is classified as containing approvals only, denials only or both outcomes. Equation 2 distinguishes the prevalence of mixed categories from their shares of requests and denials:

$$C_j = \frac{\sum_c M_{jc}}{K_j}, \qquad R_j = \frac{\sum_c N_{jc}\, M_{jc}}{N}, \qquad L_j = \frac{\sum_c D_{jc}\, M_{jc}}{D}$$

Here, j identifies an attribute field and c identifies an observed category within that field. M_{jc} equals one when the category contains at least one approval and one denial, and zero otherwise. K_j is the number of observed categories, N_{jc} is the category request count, and D_{jc} is its denial count. N and D are the total request and denial counts. C_j measures the proportion of categories containing both outcomes, R_j measures their share of requests, and L_j measures their share of denials. Reported percentages multiply these proportions by 100.

Equation 2 supplies complementary denominators for interpreting the same records. A field can contain few mixed categories while locating a substantial share of requests within them. Because every attribute partitions the same requests, shares calculated across different fields are not additive.

The analysis also compares the two largest categories within selected fields. Categories are ranked by request count, with identifier order breaking ties. Selection therefore depends on observation volume rather than approval proportion. Approval contrasts are expressed in percentage points. Table 3 presents the category classifications and request shares, while Table 4 presents the selected category comparisons.

4.3 Data Processing and Consistency Validation:

Processing begins with the Dataset worksheet in the supplied workbook. The original substantive fields are read without changing category labels or recorded outcomes. The working structure preserves one request per row and one field per column, consistent with the observational organization described by Wickham, 2014.

The audit checks completeness across the nine attributes and the outcome. It verifies that outcome values belong to the recorded binary coding. Exact duplicate checks compare all ten substantive fields, while profile consistency checks group the nine attributes and identify whether identical profiles contain conflicting outcomes. These procedures distinguish repeated category membership from duplication of an entire record.

The substantive summaries retain every request because the audit identifies no missing substantive values, exact duplicate records or conflicting identical profiles. Category frequency does not determine baseline eligibility. Singleton categories remain identifiable within the complete sample rather than being removed during cleaning.

For each attribute, requests are grouped by the unchanged category identifier. Approval and denial counts are accumulated separately, and their sum supplies the category denominator. Validation then reconciles the category totals with the complete sample and checks that aggregated approvals and denials reproduce the overall outcome counts. The separation of source records, processing steps and derived outputs follows the computational documentation practices described by Wilson et al., 2017.

A correspondence check examines mappings between role title and role code in both directions. A one-to-one mapping requires each title to correspond to one code and each code to correspond to one title. Identical partitions are interpreted as equivalent descriptions of the same observations.

The stored profile hashes and allocation flags are also recalculated and compared with their workbook values. These checks assess record traceability. The substantive descriptive summaries use the complete sample and exclude the auxiliary identifiers from category comparisons.

4.4 Sensitivity Analysis, Interpretation and Reproducibility:

An exploratory sensitivity analysis addresses the influence of singleton resources on apparent outcome uniformity. The restricted sample excludes requests belonging to resources observed once in the complete sample. Resource counts, mixed category prevalence and the overall approval proportion are then recalculated. The exclusion rule uses resource frequency rather than the recorded authorization outcome.

The comparison evaluates two distinct quantities: the composition of resource categories and the aggregate outcome distribution. Changes in mixed category prevalence indicate sensitivity to observational sparsity. Changes in the overall approval proportion indicate how the same restriction alters sample composition. Differences between the complete and restricted samples are interpreted jointly because the restriction changes both the observation count and the set of represented resources.

All reported proportions are empirical summaries of the accessible benchmark. Interpretation emphasizes denominators, category coverage and differences in observed magnitude. The sensitivity comparison does not establish stability across independently collected organizational records.

Calculations were implemented in Python 3.12.14 using standard library modules for workbook extraction, grouping, hashing, descriptive statistics and output generation. The procedures are deterministic. The supplied reproduction script regenerates the completeness audit, duplicate checks, profile checks and nine attribute specific frequency files. Supplementary outputs document category classifications, request shares and comparisons of the largest categories.

The replication materials preserve the workbook checksum, processing instructions and derived counts. This documentation connects each reported quantity to its input records and calculation procedure, following the reproducibility principles articulated by Sandve et al., 2013. The methodological contribution is the coordinated interpretation of established descriptive measures: category proportions, request concentration, classification redundancy and sensitivity to sparse observations.

5. Findings:

This analysis examines authorization decisions across resource identity, organizational context and role identity. The findings address the three research questions through category frequencies, approval proportions and the distribution of denials. Comparisons distinguish patterns across categories from patterns across the requests contained within them.

The analytical sample contains 32,769 employee access requests collected during 2010 and 2011. All requests have a recorded outcome and complete values for the nine substantive attributes. The audit identified no exact duplicate records or conflicting outcomes for identical attribute profiles. All descriptive comparisons therefore use the same sample. These counts describe the accessible benchmark, consistent with the distinction between operational records and population coverage emphasized by Hand, 2018.

Approvals account for 30,872 requests, or 94.21 percent, while denials account for 1,897 requests, or 5.79 percent. With approval coded one and denial coded zero, the outcome mean is 0.9421, the sample standard deviation is 0.2335, and the median is one. This approval dominance provides the reference against which category differences must be interpreted. A category can contain predominantly approved requests while still accounting for a substantial number of denials.

Table 3 distinguishes categories containing only approvals, only denials, or both outcomes. The final column reports the percentage of all requests located in categories containing both outcomes. This denominator distinguishes the prevalence of mixed categories from their importance within the request distribution.

Table 3: Authorization outcomes across observed attribute categories

Attribute	Categories	Singleton Categories	Approval Only Categories	Denial Only Categories	Categories With Both Outcomes	Requests in Categories With Both Outcomes, Percent
RESOURCE	7,518	3,766	6,384	292	842	50.52
MGR_ID	4,243	845	3,429	68	746	32.46
ROLE_ROLLUP_1	128	10	51	1	76	97.29
ROLE_ROLLUP_2	177	17	71	2	104	96.88
ROLE_DEPTNAME	449	24	182	4	263	89.24
ROLE_TITLE	343	32	189	3	151	91.92
ROLE_FAMILY_DESC	2,358	643	1,974	46	338	67.33
ROLE_FAMILY	67	0	20	0	47	99.24
ROLE_CODE	343	32	189	3	151	91.92

For RQ1, resource categories display substantial variation in observed authorization outcomes. Of the 7,518 resources, 6,384 contain only approvals, 292 contain only denials, and 842 contain both outcomes. However, the 842 mixed resources represent only 11.20 percent of resource categories while containing 16,554 requests, or 50.52 percent of the sample. They also contain 1,570 denials, equivalent to 82.76 percent of all denials. Counting resources alone therefore understates the concentration of requests and denials within resources exhibiting outcome variation.

The two most frequently requested resources illustrate differences without selecting categories according to their approval rates. Resource 4675 records 836 approvals among 839 requests, giving an approval proportion of 99.64 percent. Resource 79092 records 468 approvals among 484 requests, giving 96.69 percent. The difference is 2.95 percentage points. These results resolve RQ1 by demonstrating variation in recorded approval proportions across resources. Their interpretation is consistent with the subject and object combinations described by Hu, Kuhn, and Ferraiolo, 2015, without establishing that resource identity independently determines authorization.

For RQ2, organizational classifications differ in category concentration and within category outcome variation. Manager identifiers contain 746 mixed categories, representing 17.58 percent of manager categories and 32.46 percent of requests. By comparison, mixed categories contain 97.29 percent of requests under ROLE_ROLLUP_1, 96.88 percent under ROLE_ROLLUP_2 and 89.24 percent under ROLE_DEPTNAME. Thus, most requests within the broader organizational classifications belong to categories that contain both approvals and denials.

Organizational concentration also affects the interpretation of aggregate proportions. The largest ROLE_ROLLUP_1 category, 117961, contains 21,407 requests, equivalent to 65.33 percent of the sample. Its 1,087 denials represent 57.30 percent of all denials, despite an approval proportion of 94.92 percent. A high approval proportion therefore coexists with a large contribution to the total denial count. Among the two largest departments, approval proportions are 93.74 percent for category 117878 and 91.74 percent for category 117941, a difference of 2.00 percentage points. RQ2 is resolved through these observed differences in organizational composition and decision proportions.

For RQ3, role classifications combine apparent category uniformity with extensive request level variation. Among 343 role title categories, 189 contain only approvals, three contain only denials and 151

contain both outcomes. The mixed categories contain 30,122 requests, or 91.92 percent of the sample, and 1,893 denials, or 99.79 percent of all denials. Consequently, the majority of role title categories exhibit a single observed outcome, but almost all denials occur within categories exhibiting both outcomes.

The role title and role code fields map one to one and generate identical category partitions. Their matching results therefore represent the same empirical structure. This verified redundancy reinforces the importance of interpreting attribute meanings and correspondence, as emphasized by Hu, Ferraiolo, and Kuhn, 2019.

The two largest role title categories have approval proportions of 92.04 and 96.76 percent, a difference of 4.72 percentage points. At the broader role family level, 47 of 67 categories contain both outcomes. These categories account for 32,520 requests, or 99.24 percent of the sample, and all 1,897 denials. RQ3 is therefore resolved by identifying variation between role categories alongside variation within the categories containing most requests. This pattern is compatible with examining combinations of attributes and permission relationships, as developed by Xu and Stoller, 2015. It does not establish a complete authorization rule from role membership alone.

Table 4: Authorization proportions in the two largest categories within selected attributes

Attribute	Category Identifier	Requests	Approvals	Denials	Approval Percentage
RESOURCE	4675	839	836	3	99.64
RESOURCE	79092	484	468	16	96.69
ROLE_DEPTNAME	117878	1,135	1,064	71	93.74
ROLE_DEPTNAME	117941	763	700	63	91.74
ROLE_TITLE	118321	4,649	4,279	370	92.04
ROLE_TITLE	117905	3,583	3,467	116	96.76
ROLE_FAMILY	290919	10,980	10,347	633	94.23
ROLE_FAMILY	118424	2,690	2,616	74	97.25

An exploratory sensitivity check examined whether singleton resources account for the resource level pattern. Removing the 3,766 resources represented by one request retains 29,003 requests across 3,752 resources. The retained sample contains 27,374 approvals and 1,629 denials, producing an approval proportion of 94.38 percent. This differs from the complete sample by 0.17 percentage points. Overall approval dominance therefore persists after this exclusion.

The same check changes the proportion of resource categories containing both outcomes from 11.20 to 22.44 percent. The number of mixed resources remains 842 because none is a singleton. Apparent category uniformity is therefore sensitive to the prevalence of sparsely observed resources, even though the aggregate approval proportion changes little. This sensitivity concerns category composition and does not establish predictive stability.

Across the three research questions, the central finding is the distinction between category uniformity and the distribution of requests. Numerous categories contain only approvals, yet categories containing both outcomes account for most requests under the broader organizational and role classifications. Denials are particularly concentrated within these mixed categories. The evidence describes authorization decision patterns while retaining the distinction between recorded functionality and assurance articulated by the Joint Task Force, 2020.

The findings answer all three research questions descriptively: authorization proportions vary across resources, organizational classifications and roles, while the interpretation of that variation depends on category frequency and request concentration. The contribution lies in showing why category counts and request weighted summaries provide different accounts of the same authorization records.

6. Discussion:

Apparent authorization consistency depends on whether attention falls on categories or on the requests they contain. Table 3 shows that relatively few resource categories contain both approvals and denials, yet these categories account for more than half of all requests. Equation 1 makes the category denominator explicit, while the comparison with request shares reveals how uneven category sizes change interpretation. The empirical contribution is this separation between category uniformity and the concentration of requests. A benchmark dominated by categories with one observed outcome can still contain substantial decision variation where requests concentrate. These finding redirects interpretation from the number of apparently uniform categories toward the volume and composition of the observations supporting them.

Role identity provides organizational structure without fully separating recorded authorization outcomes. Table 3 shows that mixed role title categories contain most requests, while the one-to-one correspondence between role title and role code demonstrates that additional fields need not supply additional information. The differences between frequently observed categories in Table 4 further show that functional classifications accompany distinct approval proportions. Servos and Osborn, 2017, describe attribute-based authorization through combinations of subject attributes, object attributes and policy conditions. The evidence adds a measurement implication to that framework: attribute diversity must be distinguished from informational diversity. Equivalent classifications should receive one

substantive interpretation, and a category containing both outcomes should not be treated as a complete representation of the conditions governing access.

The distribution of denials identifies a further distinction between category prevalence and operational concentration. Mixed resource categories contain 82.76 percent of denials, while mixed role family categories contain every recorded denial. These are alternative partitions of the same requests, so their shares describe overlapping concentrations rather than additive contributions. Table 3 therefore locates the observed denial burden without attributing it to independent causal channels. Cotrini et al., 2018, demonstrate that sparse authorization logs can support misleading conclusions about permission rules and evaluate their approach using organizational logs, including Amazon data. The present analysis complements that work by showing where decision variation resides within the observed benchmark. Its contribution concerns the distribution of supporting evidence, rather than the introduction of a new dataset or proof of a governing policy.

Sparse categories create an interpretive challenge that aggregate proportions alone conceal. Excluding singleton resources nearly doubles the proportion of resource categories containing both outcomes, although the overall approval proportion changes by only 0.17 percentage points. Aggregate stability therefore coexists with a substantial change in the apparent composition of category outcomes. Luque et al., 2019, show that class imbalance affects the interpretation of classification performance measures, while Krawczyk, 2016, emphasizes the importance of class structure beyond the overall imbalance ratio. Here, the descriptive evidence establishes a related reporting requirement: approval dominance and category coverage must be examined together. The sensitivity check changes what category uniformity means, even though it scarcely changes the prevailing outcome. This sharpens the distinction between a stable aggregate and a stable empirical structure.

The international relevance lies in a transferable approach to interpreting authorization records across systems with different category structures. Ferraiolo et al., 2016, compare access control standards that share authorization objectives while differing in policy representation and implementation. Against that architectural diversity, Tables 3 and 4 and Equation 1 provide a common descriptive basis: report category denominators, distinguish equivalent attributes, and examine where requests and denials concentrate. Administrators can use these summaries to prioritize review of populous mixed categories and avoid interpreting singleton outcomes as established rules. Benchmark reporting should likewise present category coverage beside approval proportions. The theoretical refinement is that the informational value of an authorization attribute depends on both its classification meaning and the observations supporting each category. This perspective broadens the contribution beyond one corporate setting without assuming identical authorization patterns across countries or institutions.

7. Conclusion and Implications:

Apparently uniform access records can conceal where authorization decisions vary most. This analysis shows that category consistency and request concentration provide different interpretations of the same system. Its contribution is a descriptive framework that connects decision patterns with the volume of supporting observations. The findings refine the empirical application of authorization theory by distinguishing classification structure from informational contribution. Equivalent classifications provide repeated representations, while sparsely observed categories offer limited evidence of consistent decision patterns. For managers, this distinction supports directing review toward categories that combine substantial request volumes with diverse outcomes. Governance policies should require approval proportions to be accompanied by their underlying counts and coverage. Operational teams can apply these checks before interpreting historical decisions or preparing authorization benchmarks. Such practices offer a transferable basis for accountable access administration across organizations.

The historical corporate sample provides a foundation for comparative research using independently collected organizational records. Longitudinal data could establish whether category patterns persist as responsibilities and permissions change. Cross institutional comparisons could assess transferability under different governance arrangements. Linking requests to explicit policy rules and subsequent access outcomes would distinguish administrative consistency from policy correctness and support direct evaluation of authorization mechanisms.

References:

1. Boyd, D., & Crawford, K. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. Information, Communication & Society, volume 15, issue 5, pages 662 to 679. https://doi.org/10.1080/1369118X.2012.678878
2. Chu, X., Ilyas, I. F., Krishnan, S., & Wang, J. 2016. Data cleaning: Overview and emerging challenges. Proceedings of the 2016 International Conference on Management of Data, pages 2201 to 2206. Association for Computing Machinery. https://doi.org/10.1145/2882903.2912574
3. Cotrini, C., Weghorn, T., & Basin, D. 2018. Mining ABAC rules from sparse logs. Proceedings of the 2018 IEEE European Symposium on Security and Privacy, pages 31 to 46. IEEE. https://doi.org/10.1109/EuroSP.2018.00011
4. Ferraiolo, D. F., Chandramouli, R., Hu, V. C., & Kuhn, D. R. 2016. A comparison of attribute based access control standards for data service applications: Extensible Access Control Markup Language

and Next Generation Access Control. NIST Special Publication 800-178. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-178
5. Goodman, S. N., Fanelli, D., & Ioannidis, J. P. A. 2016. What does research reproducibility mean? Science Translational Medicine, volume 8, issue 341, article 341ps12. https://doi.org/10.1126/scitranslmed.aaf5027
6. Hand, D. J. 2018. Statistical challenges of administrative and transaction data. Journal of the Royal Statistical Society: Series A, Statistics in Society, volume 181, issue 3, pages 555 to 605. https://doi.org/10.1111/rssa.12315
7. Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. 2014. Guide to attribute based access control definition and considerations. NIST Special Publication 800-162. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-162
8. Hu, V. C., Ferraiolo, D. F., & Kuhn, D. R. 2019. Attribute considerations for access control systems. NIST Special Publication 800-205. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-205
9. Hu, V. C., Kuhn, D. R., & Ferraiolo, D. F. 2015. Attribute based access control. Computer, volume 48, issue 2, pages 85 to 88. https://doi.org/10.1109/MC.2015.33
10. Joint Task Force. 2020. Security and privacy controls for information systems and organizations. NIST Special Publication 800-53, Revision 5. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53r5
11. Krawczyk, B. 2016. Learning from imbalanced data: Open challenges and future directions. Progress in Artificial Intelligence, volume 5, issue 4, pages 221 to 232. https://doi.org/10.1007/s13748-016-0094-0
12. Lazer, D., Kennedy, R., King, G., & Vespignani, A. 2014. The parable of Google Flu: Traps in big data analysis. Science, volume 343, issue 6176, pages 1203 to 1205. https://doi.org/10.1126/science.1248506
13. Luque, A., Carrasco, A., Martín, A., & de las Heras, A. 2019. The impact of class imbalance in classification performance metrics based on the binary confusion matrix. Pattern Recognition, volume 91, pages 216 to 231. https://doi.org/10.1016/j.patcog.2019.02.023
14. Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E. J., Ware, J. J., & Ioannidis, J. P. A. 2017. A manifesto for reproducible science. Nature Human Behaviour, volume 1, issue 1, article 0021. https://doi.org/10.1038/s41562-016-0021
15. OpenML. 2015. Amazon_employee_access, version 1, dataset 4135. Dataset repository record. https://www.openml.org/d/4135
16. Peng, R. D. 2011. Reproducible research in computational science. Science, volume 334, issue 6060, pages 1226 to 1227. https://doi.org/10.1126/science.1213847
17. Sadiq, S., editor. 2013. Handbook of data quality: Research and practice. Springer. https://doi.org/10.1007/978-3-642-36257-6
18. Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. 2013. Ten simple rules for reproducible computational research. PLOS Computational Biology, volume 9, issue 10, article e1003285. https://doi.org/10.1371/journal.pcbi.1003285
19. Servos, D., & Osborn, S. L. 2017. Current research and open problems in attribute based access control. ACM Computing Surveys, volume 49, issue 4, article 65, pages 1 to 45. https://doi.org/10.1145/3007204
20. Shmueli, G. 2010. To explain or to predict? Statistical Science, volume 25, issue 3, pages 289 to 310. https://doi.org/10.1214/10-STS330
21. Stodden, V., McNutt, M., Bailey, D. H., Deelman, E., Gil, Y., Hanson, B., Heroux, M. A., Ioannidis, J. P. A., & Taufer, M. 2016. Enhancing reproducibility for computational methods. Science, volume 354, issue 6317, pages 1240 to 1241. https://doi.org/10.1126/science.aah6168
22. Stodden, V., Seiler, J., & Ma, Z. 2018. An empirical analysis of journal policy effectiveness for computational reproducibility. Proceedings of the National Academy of Sciences, volume 115, issue 11, pages 2584 to 2589. https://doi.org/10.1073/pnas.1708290115
23. Vanschoren, J., Van Rijn, J. N., Bischl, B., & Torgo, L. 2013. OpenML: Networked science in machine learning. ACM SIGKDD Explorations Newsletter, volume 15, issue 2, pages 49 to 60. https://doi.org/10.1145/2641190.2641198
24. Wickham, H. 2014. Tidy data. Journal of Statistical Software, volume 59, issue 10, pages 1 to 23. https://doi.org/10.18637/jss.v059.i10
25. Wilson, G., Bryan, J., Cranston, K., Kitzes, J., Nederbragt, L., & Teal, T. K. 2017. Good enough practices in scientific computing. PLOS Computational Biology, volume 13, issue 6, article e1005510. https://doi.org/10.1371/journal.pcbi.1005510
26. Xu, Z., & Stoller, S. D. 2015. Mining attribute based access control policies. IEEE Transactions on Dependable and Secure Computing, volume 12, issue 5, pages 533 to 545. https://doi.org/10.1109/TDSC.2014.2369048