Annals of "Dunarea de Jos" University of Galati
Fascicle I. Economics and Applied Informatics
Years XXXII – n°2/2026
ISSN-L 1584-0409 ISSN-Online 2344-441X
www.eia.feaa.ugal.ro

DOI https://doi.org/10.35219/eai15840409594

Formalising Bidder-Experience Verification in Public Procurement: A Rule-Constrained Mathematical Model and Anonymised Romanian Proof of Concept

Mihnea-Tudor Orjan[*]

ARTICLE INFO

Article history:
Received July 19, 2026
Accepted July 30, 2026
Available online September 1, 2026
JEL Classification
H57, C61, C63, K23

Keywords:
supplier qualification; legal informatics; robust optimisation; decision support; auditability; document intelligence; institutional sustainability; Romania

ABSTRACT

This paper formalises bidder-experience verification as an auditable mathematical decision problem. It uses two document streams: an authority corpus, ordinarily collected in Romania from the Electronic System for Public Procurement (SEAP), and a bidder corpus containing contracts, activity schedules, acceptance records and role-specific evidence. The proposed architecture includes AI-assisted extraction, while the proof of concept tests a deterministic rule kernel on pre-structured anonymised data. Material experience, documentary demonstration and robust safety are separated. Binary activity classification is combined with reference-period, acceptance, attribution and maximum-contract constraints; exact subset optimisation excludes inadmissible or redundant references. A finite uncertainty set distinguishes plausible adverse interpretations from concentration tests. Counterfactual tests cover exact thresholds, missing evidence, expired references, zero qualifying activity and loss of the dominant contract. The model remains human-supervised, privacy-aware and intended for decision support, as a reusable module for future multi-procedure validation and econometric research.

1. Introduction

Public procurement combines substantial economic stakes with document-intensive legal reasoning. A contracting authority may define technical and professional ability through monetary thresholds, reference periods, maximum numbers of references, activity taxonomies, acceptance conditions, attribution rules and documentary requirements. The bidder responds with a different documentary universe: past contracts, addenda, activity schedules, acceptance records, certificates, association shares, subcontracting documents and confirmations issued by final beneficiaries. Even when the legal result is binary, the path to that result is multidimensional.

This paper asks whether bidder-experience verification can be formalised sufficiently to support automated calculation and structured human decision-making without suppressing the legally bounded judgment of the evaluator. The proposed answer is affirmative if automation is designed as an auditable decision-support system rather than as a black-box prediction. The authority corpus defines the target and constraints; the bidder corpus supplies candidate observations. A line-level classification and aggregation engine then determines whether the admissible references satisfy the lot-specific threshold together with period, contract-count, acceptance, attribution and documentary-chain rules.

The implemented core is a deterministic mathematical decision model. It is not an inferential econometric model: the proof of concept contains one transformed case and estimates neither coefficients nor causal effects. A future multi-procedure dataset could support econometric analysis of processing time, clarification success or human override, but such predictions would remain analytically subordinate to the legally traceable rule kernel.

The article makes six contributions. First, it translates public-procurement experience requirements into explicit variables and constraints. Second, it separates material possession of experience from its documentary demonstration. Third, it implements exact admissible-subset selection and a finite-scenario robust optimisation rather than a simple ranking of nominal contract values. Fourth, it demonstrates the internally reproducible proof of concept through a privacy-preserving transformation of a real Romanian professional exercise; it does not claim external validation from a single case. Fifth, it adds counterfactual boundary and concentration tests that exercise PASS, CONDITIONAL, FAIL and UNRESOLVED pathways

[*]Bucharest University of Economic Studies, Economics I Doctoral School, Bucharest, Romania. E-mail address: orjanmihnea22@stud.ase.ro (M. T. Orjan).

without misrepresenting them as independent empirical procedures. Sixth, it frames the experience engine as one reusable module of a broader eligibility platform and links modular reuse to institutional sustainability. The design originated in a specific professional need, but its abstraction enables cumulative research and implementation beyond the initial procedure.

2. Literature review

2.1. Normative and institutional framework

2.1.1. Selection criteria, reference periods and evidence

Article 58 of Directive 2014/24/EU permits selection criteria concerning suitability, economic and financial standing, and technical and professional ability. Contracting authorities may require the human and technical resources and experience necessary to perform a contract to an appropriate quality standard (European Parliament & Council, 2014). Romanian Law no. 98/2016 transposes this framework and applies the principles of non-discrimination, equal treatment, mutual recognition, transparency, proportionality and assumption of responsibility (accountability) (Romanian Parliament, 2016). Government Decision no. 395/2016 regulates the evaluation workflow and the verification of the European Single Procurement Document, known nationally as DUAE (Government of Romania, 2016). ANAP Instruction no. 2/2017 provides operational rules concerning similar experience, reference periods, partial acceptance, associations, subcontractors and acceptable evidence (National Agency for Public Procurement, 2017).

The legal test is not reducible to one contract value. In a mixed contract, only qualifying activities may be counted. Partially accepted services may be relevant when the accepted deliverable is independently usable. The reference period is a procedure-specific variable: for services it is commonly three years, but the Instruction permits a longer period in duly justified cases where required to ensure adequate competition (National Agency for Public Procurement, 2017). The maximum number of references is also variable and must be extracted from the active procurement documents rather than assumed by the software.

Experience obtained as an association member or subcontractor must be attributed to the operator according to the proven share or assigned activities. For subcontractor experience, the general contractor confirms the activities and the value attributable to the subcontractor, while the final beneficiary confirms that the operator acted as subcontractor; the complete evidence chain must also establish the relevant period and accepted value (National Agency for Public Procurement, 2017). Article 209 of Law no. 98/2016 permits requests for clarifications or supporting documents where information is incomplete, erroneous or missing, but any remedy must respect equal treatment and transparency and may not create an evident competitive advantage (Romanian Parliament, 2016).

2.1.2. The role and precedence of procurement documents

The contract notice identifies the procedure, lots, deadlines and high-level criteria. The instructions to tenderers or procurement data sheet ordinarily contain the operative thresholds, reference window, maximum number of contracts and proof requirements. The technical specifications define the functional content of the activities and are therefore central to semantic similarity. Clarification answers interpret and resolve ambiguities in published requirements, whereas formal corrigenda may amend or supersede earlier provisions and may change temporal anchors. The DUAE request defines preliminary declarations, while the draft contract may clarify performance, acceptance and subcontracting. A reliable system must therefore preserve document provenance, publication time and legal precedence; otherwise, a correctly extracted sentence can still generate the wrong active rule.

Table 1. Two-stream data architecture and principal variables.

Stream	Typical documents	Variables extracted	Legal function
Authority	Notice; instructions; specifications; clarification answers; corrigenda; DUAE; draft contract	Thresholds, lots, reference window, contract cap, activity family, proof rule, deadlines	Defines the target and active constraints
Bidder	Contracts; addenda; schedules; acceptance records; certificates; association/subcontract records; final-beneficiary confirmations	Parties, role, activities, dates, quantities, prices, accepted values, share and provenance	Supplies candidate evidence

2.1.3. SEAP and the state of digital procurement

In Romanian practice, the authority-side series is ordinarily assembled from the Electronic System for Public Procurement (SEAP), the national electronic procurement platform accessible through www.e-licitatie.ro (Electronic System for Public Procurement, 2026). A procedure page may contain the contract notice, instructions, technical specifications, forms, draft contracts, corrigenda, ex-officio notices and consolidated clarification answers. The model therefore treats the SEAP record as a versioned documentary container, not as a single notice. Each acquisition event should retain its timestamp, document category, cryptographic hash and supersession relationship.

European procurement data are becoming more structured. eForms are the EU legislative open standard for publishing procurement notices and became mandatory in October 2023 (European Commission, 2023a). The eProcurement Ontology formally represents procurement concepts in human- and machine-readable formats (Publications Office of the European Union, 2026). The Public Procurement Data Space is intended to connect procurement data and advanced analytics (European Commission, 2023b). OECD reports identify repetitive-task automation, comparison, anomaly detection and supplier analysis as important opportunities, while also emphasising privacy, accountability and data-governance risks (OECD, 2025a; OECD, 2025b).

Romania already uses significant digital components, but not the specific line-level verification engine proposed here. SEAP supports electronic publication and interaction, and the National Integrity Agency's PREVENT system performs ex-ante conflict-of-interest checks (Electronic System for Public Procurement, 2026; National Integrity Agency of Romania, 2025). The reviewed official sources do not document a generally deployed Romanian system that automatically determines whether the activities, dates, values and evidence from past contracts satisfy a procedure-specific experience criterion. This is a statement about publicly documented large-scale adoption, not proof that no internal or private prototype exists.

2.2. Artificial intelligence, procurement systems and research gap

The interdisciplinary literature on artificial intelligence and public procurement has expanded rapidly. Sava (2023) reviews legal, technical and policy perspectives across a broad corpus, while Balkan and Akyuz (2025) provide a process-oriented taxonomy of artificial-intelligence and machine-learning applications in procurement and purchasing decision support more generally. Taken together, these reviews identify substantial potential, fragmented disciplinary treatment, implementation and governance challenges, and a continuing gap between prototype-heavy research and mature organisational adoption.

A first technical strand concerns information retrieval and document intelligence. Siciliani et al. (2023) integrate structured and unstructured tender information through semantic search, open information extraction and analytical dashboards. Zhao and Li (2024) propose a retrieval-augmented generation framework specifically for the production of semi-structured tender documents. Arroyo et al. (2022) combine optical character recognition, entity tagging and rule-based corrections to recover descriptions, quantities and prices from purchase documents. Although these studies address partly different document types and operational tasks, they motivate several components of the proposed extraction architecture. The replication prototype accompanying the present article nevertheless begins with pre-structured data and does not claim to implement the complete OCR and natural-language-processing pipeline.

A second strand concerns predictive classification. Silaschi, Pop and Coroiu (2024) apply machine-learning models to a large Romanian procurement dataset to classify bidder suitability and assess the consistency of previous selection decisions. Moiraghi et al. (2024) examine hierarchical zero-shot classification within the Common Procurement Vocabulary taxonomy, relying on class descriptions and the hierarchical structure of the CPV rather than merely reproducing categories observed in historical data. These systems generate statistical classifications rather than deterministic legal findings. The present system instead evaluates explicit legal inequalities and documentary evidence rules. Predictive methods may subsequently assist in prioritising cases for human review, but they should not silently replace the selection criteria published by the contracting authority.

A third strand addresses integrity, audit and risk detection. Schneider dos Santos et al. (2025) map 93 data-driven fraud-detection studies and identify substantial methodological diversity, limited availability of public datasets and very restricted reporting of deployment in operational investigative systems. INACIA integrates large language models into Brazilian public-audit workflows and evaluates generated outputs against human judgment (Pereira et al., 2025). OECD materials describe operational or developing systems such as Brazil's ALICE and data-driven audit selection and prioritisation in Portugal (OECD, 2025b; OECD, 2026). These applications support oversight, anomaly identification and risk-based control. The reviewed materials do not, however, directly formalise bidder-experience matching against the qualification requirements and evidentiary rules of an individual procurement procedure.

A fourth strand supplies the methodological foundations for uncertainty management and human oversight. Robust optimisation evaluates decisions against explicitly defined uncertainty sets and makes the cost of conservatism observable through the trade-off between nominal performance and protection against adverse scenarios (Ben-Tal et al., 2009; Bertsimas & Sim, 2004). Human–AI interaction research recommends communicating uncertainty, enabling effective correction and designing systems that support appropriate rather than automatic reliance on machine-generated outputs (Amershi et al., 2019). System-card approaches structure accountability across data, model, code and system dimensions and address development, assessment, mitigation and assurance (Gursoy & Kakadiaris, 2022). These principles support the article's finite-scenario safety test, the traceability of intermediate calculations and the separation between machine-assisted analysis and the official responsibility of human decision-makers.

Finally, sustainability-oriented research argues that artificial intelligence may assist public buyers in interpreting complex information across the procurement lifecycle, while also warning about opacity, privacy and data-security risks, market concentration, energy and water consumption, hardware production and disposal, and unequal access to digital infrastructure and capabilities (Andhov et al., 2025). The present article adopts a narrower and deliberately testable sustainability claim. Modular and structured reuse of procurement evidence may reduce duplicated documentary work, avoidable requests for clarification and inconsistent preliminary screening. These potential benefits are treated as hypotheses for subsequent empirical assessment rather than as outcomes demonstrated by the current prototype.

Commercial source-to-pay and supplier-management suites also provide relevant practical context. Public materials from JAGGAER, Coupa and Ivalua describe supplier onboarding, configurable qualification workflows, document and evidence collection, compliance gates, risk monitoring, supplier scoring and scorecards (Coupa, 2026; Ivalua, 2026; JAGGAER, 2026). These capabilities demonstrate that supplier qualification is already an established software market. However, within the publicly available product documentation reviewed for this study, no evidence was identified of a jurisdiction-specific engine that reconstructs the active rules of an individual public procurement procedure, classifies activity lines from mixed contracts, applies Romanian rules concerning contractual roles and acceptance evidence, and publishes a reproducible robust lower bound. This observation does not exclude proprietary or customer-specific configurations that are not publicly documented. The novelty claim is therefore limited to the transparent legal–mathematical bridge proposed in this article and does not extend to supplier-qualification software as a general category.

Within the academic literature and publicly available commercial documentation reviewed, the specific research gap is a transparent bridge between document intelligence and the deterministic application of legally binding, procedure-specific experience criteria. The proposed model addresses this gap by combining versioned rule extraction, lot-specific activity matching, exact contract selection, documentary admissibility rules and an operational finite scenario set. Its purpose is to support professional decision-making through reproducible calculations and an auditable evidentiary trail, not to automate or replace legally attributable adjudication.

3. Methods

3.1. Data architecture and analytical-series construction

3.1.1. Authority-side series

The first analytical series is generated from the authority corpus. Operationally, the collector queries SEAP by procedure or notice identifier, downloads every publicly available attachment and records the publication event (Electronic System for Public Procurement, 2026). The minimum acquisition set comprises the notice, instructions or data sheet, technical specifications, forms, draft contract, formal corrigenda, ex-officio communications and consolidated clarification answers. Each file is stored with a document identifier, source location, type, publication date, version, page range, cryptographic hash and precedence score.

The active requirement object contains the lot identifier l, minimum threshold T_l, reference period W_l, maximum number of references K_l, accepted activity set S_l, acceptance rule, attribution rule, documentary rule and submission deadline. Clarification answers are recorded as interpretative rule events, while corrigenda are recorded as formal amendment events. The earlier text remains in the audit log, but only the final coherent rule set is applied to the decision.

3.1.2. Bidder-side series

The second series is constructed at activity-line level. For every contract j and line i, the system records the activity description, family, quantity q_{ij}, unit price p_{ij}, line value v_{ij}, performance and acceptance dates, beneficiary, bidder role, association share or assigned activity, and supporting page references. Contract totals alone are unsafe for mixed contracts because non-qualifying activities may inflate the nominal value. The activity line is therefore the smallest economic observation used by the model.

Evidence may include contracts or relevant extracts, addenda, orders, service schedules, acceptance certificates, performance certificates, association or subcontract agreements, general-contractor confirmations, final-beneficiary confirmations and corroborative invoices. Native documents can be parsed directly; scans require OCR and table reconstruction. Every extracted field should retain a confidence score and page-level provenance. Low-confidence monetary values, inconsistent dates and unmatched references are sent to the human exception queue before the mathematical kernel runs.

3.1.3. Privacy-preserving transformation and confidentiality

The proof-of-concept data originate from a real Romanian professional review, but the replication package contains transformed observations. Monetary values are divided by the lowest relevant lot threshold. Direct names, the procedure identifier and exact dates are removed; acceptance dates are expressed as months before the deadline. This preserves inequalities, ratios, temporal conclusions and contract-count constraints. Residual re-identification risk cannot be excluded because distinctive combinations may still be recognisable to persons familiar with the underlying procedure.

Authority documents obtained from SEAP are public, whereas bidder evidence may contain prices, customer relationships, technical methods, personal data and trade secrets. Before general-purpose AI processing, the operator should apply data minimisation, pseudonymisation or anonymisation, role-based access, retention limits and deletion controls, and should verify training, subcontractor and international-transfer terms. An EDPB Support Pool of Experts report identifies practical privacy risks and mitigations for large language models (European Data Protection Board Support Pool of Experts, 2025); the GDPR and Directive (EU) 2016/943 provide the broader legal context (European Parliament & Council, 2016a; European Parliament & Council, 2016b). Local or contractually controlled enterprise processing is preferable for highly sensitive evidence. GPT@EC illustrates a secure institutional approach to general-purpose generative AI (European Commission, 2024).

Table 2. Practical confidentiality gate for AI-assisted processing.

Data class	Examples	Permitted processing approach	Minimum control
Public	SEAP notices and public attachments	General retrieval and extraction tools	Source log and integrity hash
Restricted	Anonymised schedules and acceptance records	Approved enterprise environment	Minimisation, access control and retention limit
Highly confidential	Unredacted contracts, prices and trade secrets	Local or contractually isolated processing	No external reuse; audit log; deletion verification

3.2. Rule-constrained mathematical decision model
3.2.1. Lot-specific material value

Let L be the set of lots, J the set of reference contracts and I_j the activity lines in contract j. Equation (1) defines each observed line value from accepted quantity and unit price. The active-criterion domain is T_l, $W_l > 0$ and $K_l \in N_0$; non-positive thresholds or periods are configuration conflicts.

$$v_{ij} = q_{ij} \cdot p_{ij} \tag{1}$$

Material eligibility is calculated without multiplying by documentary sufficiency. This preserves the distinction between possessing experience and proving it. The activity rule is deliberately non-compensatory: α_{ijl} is binary and equals one only when the activity line belongs to the accepted family for lot l under the relevant legal-technical scenario. A mixed line should be split into eligible and non-eligible value components rather than assigned an opaque fractional similarity score. The temporal coefficient τ_{ijl} and the acceptance coefficient r_{ij} are also binary. The attribution coefficient ρ_{ij} may be continuous because an association share or assigned subcontracted value can legitimately represent a fraction of the line.

$$e_{ijl}{}^{M} = v_{ij} \cdot \alpha_{ijl} \cdot \tau_{ijl} \cdot r_{ij} \cdot \rho_{ij} \tag{2}$$

$$E_{jl}{}^{M} = \Sigma_i \in I_j \; e_{ijl}{}^{M} \tag{3}$$

Equation (2) filters the observed line value by lot-specific similarity, temporal validity, acceptance and attribution, while equation (3) aggregates the surviving activity-line values at contract and lot level.

3.2.2. Documentary sufficiency and exact admissible-subset selection

For contract j, let R(role_j) denote the evidence types required for the operator's proven historical role, u_jm indicate the presence of required evidence item m, and z_j indicate an unresolved evidentiary conflict. Equation (4) defines role-specific documentary completeness. For a prime contractor, the minimum chain contains the contract or relevant extract, activity/value evidence and acceptance evidence. A consortium member additionally requires the association agreement and proof of the allocated activities or share. A subcontractor requires the subcontract or equivalent evidence, the general contractor's confirmation of the activities and attributable value, and the final beneficiary's confirmation of subcontractor status. The rule matrix is configurable because the legal relationship, not the nominal document title, determines the required chain.

$$d_j = (1 - z_j) \cdot \Pi_m \in R(role_j) \; u_{jm} \tag{4}$$

Equation (4) is not multiplied into material line value; it is applied only when constructing the demonstrated admissible set. The material maximum ignores documentary gaps and answers whether the experience exists on the available facts. The demonstrated maximum selects only documentary-complete contracts. The binary variable x_jl indicates whether contract j is used for lot l. The same contract may support multiple lots where the active documents allow this, but each lot is subject to its own cap K_l.

$$E_l{}^{M} = max_x \; \Sigma_j \; x_{jl} E_{jl}{}^{M}, \quad \Sigma_j x_{jl} \leq K_l \tag{5}$$

$$E_l{}^{D} = max_x \; \Sigma_j \; x_{jl} E_{jl}{}^{M}, \quad \Sigma_j x_{jl} \leq K_l, \quad x_{jl} \leq d_j \tag{6}$$

Equations (5) and (6) require optimisation over admissible subsets, not selection of the nominally largest contracts followed by a later validity check. This order prevents a false failure when a documentary-incomplete or out-of-period contract is redundant and another admissible contract already reaches the threshold. The replication code performs exact subset enumeration and includes dedicated regression tests for this case.

3.2.3. Operational finite-scenario robustness

The human evaluator retains a legally bounded margin of appreciation where documents use open-textured concepts such as "similar", "ancillary", "independently usable" or "attributable". The model represents this through a finite scenario set Ω_l. Define $E_{jl}^s = \Sigma_{i}\in I_j\ v_{ij}\alpha_{ijl}^s\tau_{ijl}r_{ij}\rho_{ij}^s$, where α_{ijl}^s retains only unambiguously admissible lines and ρ_{ij}^s is the minimum defensible attributable share. Scenario h may exclude a disputed contract through $\delta_{jh} \in \{0,1\}$ and apply $m_{jh} \in [0,1]$ to its value.

$$\underline{E}_l = max_x\ min_h \in\Omega_l\ \Sigma_j\ x_{jl}(1 - \delta_{jh})m_{jh}E_{jl}^{s},\quad \Sigma_j x_{jl} \le K_l,\quad x_{jl} \le d_j \tag{7}$$

Equation (7) is implemented for the transformed data in the uniform-multiplier special case $m_{jh} = m_h$. The uncertainty set covers semantic, numeric, documentary and attribution uncertainty. Each scenario records its trigger, affected variables and classification as core or concentration stress. The six-percent correction is case-specific, not a universal parameter or legal probability; production bounds should come from OCR confidence, source reconciliation, documented ranges and explicit disputed activity classes.

3.2.4. Decision-theoretic status of the robust rule

The max-min rule is a normative internal certification rule, not an autonomous legal requirement. It reserves the label SAFE for a subset that reaches the threshold in every adverse interpretation included ex ante in the legally plausible uncertainty set. This expresses controlled risk aversion suitable for pre-submission screening and audit support. It does not require the contracting authority to adopt the most adverse interpretation, and it does not permit arbitrary deletion of undisputed evidence. Minimax regret or α-maximin rules could provide less conservative decision support, but they would change the meaning of SAFE. Catastrophic loss of an undisputed dominant reference is therefore reported as a concentration test rather than inserted into the core robust bound.

3.2.5. Decision classes, conflicts and robustness margins

The output distinguishes four classes. SAFE means $\underline{E}_l$ reaches T_l and no outcome-relevant conflict remains. CONDITIONAL means the material threshold is met but an apparently remediable documentary gap remains, or the central demonstrated value passes while the robust lower bound does not. FAIL means material value is below threshold or a non-curable admissibility or evidentiary defect prevents demonstration. UNRESOLVED is reserved for conflicts or missing variables that prevent a reliable bound.

$$SAFE_l = 1[\underline{E}_l \ge T_l] \cdot 1[Z_l = 0] \tag{8}$$

Equation (8) embeds documentary sufficiency as an admissibility condition. The five flags in equation (9) equal one, respectively, for indeterminate rule/version precedence, an absent outcome-relevant input, an unresolved historical role or evidence matrix, possible double counting or inconsistent duplicate values, and contradictory acceptance or attribution evidence. $Z_l = 0$ only if none prevents a reliable bound. Classification is hierarchical: $Z_l = 1$ gives UNRESOLVED; otherwise equation (8) gives SAFE; a material pass with a potentially curable gap gives CONDITIONAL; remaining non-qualifying cases give FAIL. Thus $d_j = 0$ for every contract does not mechanically imply FAIL while lawful curability remains unresolved.

$$Z_l = \max\{Z_l^{rule}, Z_l^{missing}, Z_l^{role}, Z_l^{duplicate}, Z_l^{conflict}\} \tag{9}$$

The principal diagnostics are robust margin M_l, coverage ratio CR_l, maximum absorbable loss AL_l, contract-count slack CS_l and temporal slack TS_l. AL_l is zero for a non-passing result and otherwise measures the share of robust demonstrated value that may be lost while T_l remains met.

$$M_l = \underline{E}_l - T_l,\quad CR_l = \underline{E}_l/T_l,\quad AL_l = 1 - T_l/\underline{E}_l \text{ if } \underline{E}_l \ge T_l > 0;\text{ otherwise } AL_l = 0 \tag{10}$$

$$CS_l = K_l - \Sigma_j x_{jl},\quad TS_l = W_l - max(age_{ij}) \tag{11}$$

Equations (10) and (11) convert the binary decision into interpretable value, contract-count and temporal safety margins.

3.2.6. Future empirical and econometric research

Econometrics is outside the present decision rule. A future multi-procedure expert-labelled study could model processing time, clarification success or override from robust margins and data-quality covariates, subject to endogeneity, selection, measurement error and evaluator heterogeneity; estimates would only prioritise review.

Table 3. Core variables and their legal-economic interpretation.

Symbol	Range	Meaning	Primary evidence
T_l	R+	Lot-specific minimum threshold	Instructions and active clarifications
W_l	R+	Reference period in months	Instructions, deadline and corrigenda
K_l	N_0	Maximum number of references	Instructions and active clarifications
α_{ijl}	{0,1}	Binary lot-specific activity-admissibility rule	Specifications, clarification answers and activity line
τ_{ijl}	{0,1}	Line-level temporal validity	Acceptance date and W_l
r_{ij}	{0,1}	Completion or independently usable accepted deliverable	Acceptance certificate or equivalent
ρ_{ij}	[0,1]	Attributable share or assigned activity value	Association/subcontract documents
d_j	{0,1}	Role-specific documentary completeness	Contractor and final-beneficiary evidence
Ω_l	finite set	Legally plausible adverse scenarios	Disputes, confidence bounds and attribution ranges
Z_l	{0,1}	Outcome-relevant unresolved conflict flag	Rule, data, role, duplicate and substantive checks
$\underline{E}_l$	R+	Robust demonstrated lower bound	Exact optimisation output

3.2.7. Computational tractability and auditability

Equations (5) and (6) reduce to selecting the highest-valued admissible references because values are non-negative and the only constraint is cardinality. The prototype retains exact enumeration to share a transparent implementation with equation (7) and support deterministic tie-breaking. For n references and cap K_l, it evaluates $\Sigma_{k=0}^{Kl}$ C(n,k); for n = 20 and K_l = 5, this is 21,700 subsets. Ties are resolved by fewer contracts and then lexical order; larger sets can use MILP or branch-and-bound. Pre-filtering may remove zero-compatible, expired or documentary-impossible references, but never a contract merely because another has a larger nominal value.

Auditability requires more than a final label. Each execution should preserve the active rule version, input hashes, activity-line classifications, selected and rejected subsets, documentary exclusions, scenario-specific values and the reason for every human override. A rerun on the same checksum-verified inputs should reproduce the same output. When the evaluator changes a classification or lower bound, the system should create a new decision version rather than overwrite the previous state. This event-sourced design makes legal discretion visible and supports internal control, challenge review and scientific replication.

3.3. Software architecture and operational use

Figure 1 distinguishes the proposed architecture from the implemented prototype.

Figure 1. Auditable architecture of the proposed bidder-experience verification system

The full production architecture contains six modules. The ingestion module fingerprints files, stores originals and renders pages. The proposed extraction module applies native-text parsing, OCR and table reconstruction. The semantic module maps activities to a controlled taxonomy derived from the technical specifications and clarification answers. The legal-rule engine converts the active documentary package into versioned constraints. The decision kernel performs exact subset selection and finite-scenario robust optimisation. The explanation module produces a lot-by-lot report with margins, excluded lines and source-page links.

The accompanying prototype implements and tests the deterministic decision kernel on pre-structured anonymised data and instantiates the subcontractor-specific documentary chain used in the transformed case. It does not claim to implement autonomous OCR, rule extraction or semantic classification. Those components are architectural specifications supported by the literature and remain future engineering work. This distinction is important for reproducibility: readers can inspect exactly which functions are operational and which are proposed.

Generally accessible AI can assist clause retrieval, OCR correction, table normalisation, extraction scripts, scenario execution and narrative explanation. Monetary aggregation, date arithmetic, contract-cap optimisation and final rule checks should remain reproducible in deterministic code. Every extracted value must retain provenance, and every machine-suggested legal classification must be editable by a procurement expert. The system should distinguish missing information from zero and uncertainty from negative evidence. The operational workflow has seven stages: collection of the authority corpus from SEAP; hashed versioning; confidentiality classification and minimisation of bidder evidence; AI-assisted extraction in an approved environment; deterministic execution; human resolution of exceptions; and generation of a reasoned memorandum. The human reviewer works from exception queues rather than repeating all arithmetic. Typical exceptions include low-confidence table rows, unreconciled totals, inconsistent party names, ambiguous activity families and a clarification that changes the active rule.

The replication package contains the anonymised CSV series, explicit robust scenarios, Python code, 21 automated tests, deterministic outputs and figures. Tests cover exact-threshold passage, zero contract caps, fractional attribution, lot-specific activity compatibility, duplicate lines, missing values, invalid coefficients, documentary redundancy, temporal redundancy, evidentiary conflicts and the maximum-adversity lower bound.

3.4. Proof-of-concept case design and assumptions

The transformed exercise concerns three lots in a Romanian services procedure. The authority required similar activities accepted within a 36-month reference window and permitted no more than five reference contracts. The bidder relied on two contracts containing periodic acceptance records. The bidder had acted as subcontractor. The model therefore assumes that the contracts or relevant extracts, the general contractor's confirmation and the final beneficiary's confirmation of subcontractor status are filed. All ambiguous cleaning, sanitation and generic ancillary activities are excluded from strict value; horticultural maintenance, tree and shrub operations and irrigation-system activities remain included.

The lowest lot threshold is normalised to one. The other thresholds become 1.492 and 1.656. Strict material and demonstrated experience equal 4.707 units. The larger contract contributes 3.988 units. The oldest relevant acceptance is 25 months before the deadline, giving 11 months of temporal slack. Two of five permitted contracts are used, giving three unused slots.

4. Results

4.1. Central and robust results

Figure 2 compares the three thresholds with the strict demonstrated value and the robust lower bound.

Table 4. Normalised proof-of-concept results.

Indicator	Lot A	Lot B	Lot C
Threshold T_l	1.000	1.492	1.656
Material value E_l^M	4.707	4.707	4.707
Demonstrated value E_l^D	4.707	4.707	4.707
Robust lower bound $\underline{E}_l$	3.749	3.749	3.749
Robust margin M_l	2.749	2.256	2.093
Robust coverage CR_l	3.749×	2.512×	2.264×
Maximum absorbable loss AL_l	73.33%	60.19%	55.83%
Contracts used / permitted	2 / 5	2 / 5	2 / 5
Decision	SAFE	SAFE	SAFE

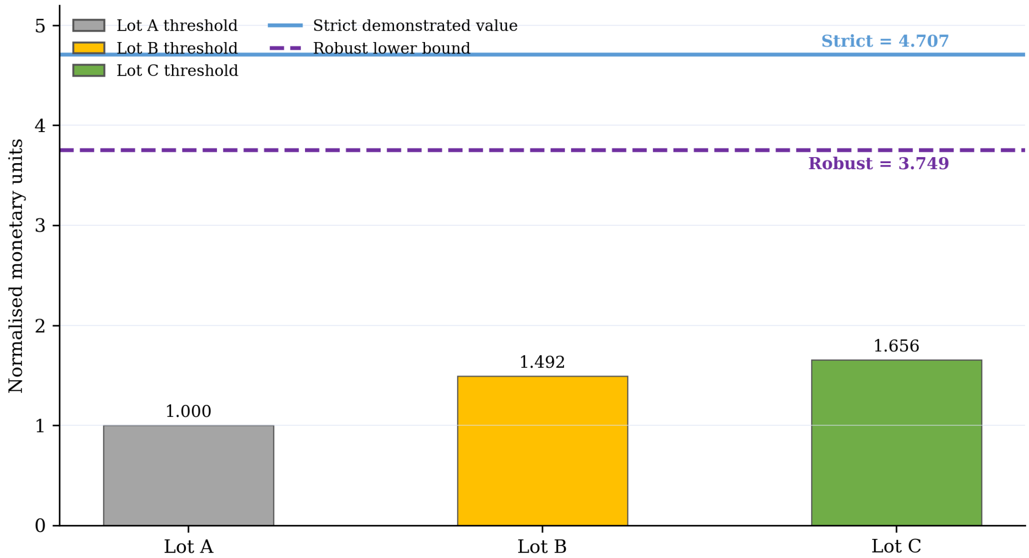

Figure 2. Strict demonstrated experience and the robust lower bound compared with the three anonymised lot thresholds

The plausible uncertainty set and the separate concentration tests are shown in table 5. The strict baseline already excludes ambiguous activity families and applies lower-bound attribution. Scenario S1 removes the specifically disputed smaller contract. S2 applies a declared six-percent lower numeric correction. S3 combines both adverse changes and produces the core robust value of 3.749 units. S4 and S5 remove the undisputed dominant contract and are reported separately as concentration tests; they do not enter the legally plausible max-min bound.

Table 5. Plausible adverse scenarios and separate concentration tests.

Scenario	Class	Adverse operation	Value	Highest-threshold result
S0	Core	Strict binary activity and lower-bound attribution baseline	4.707	Pass
S1	Core	Exclude specifically disputed smaller contract	3.988	Pass
S2	Core	Apply illustrative 6% lower numeric correction	4.425	Pass
S3	Core	Combine S1 and S2	3.749	Pass - robust bound
S4	Concentration	Exclude undisputed dominant contract	0.719	Fail
S5	Concentration	Combine S4 with 6% correction	0.676	Fail

4.2. Counterfactual boundary, concentration and sensitivity tests

The contract-count test yields 40% utilisation and a slack of three references. This is flexibility, not an additional legal entitlement: unused slots permit alternative admissible references but cannot cure insufficient value or defective evidence. The temporal test yields 11 months of slack, meaning that the oldest selected acceptance remains 11 months inside the active window; no special legal threshold attaches to that margin.
Counterfactual software tests exercise decision pathways not independently observed in the single source case. They are synthetic boundary tests, not additional procurement procedures. A redundant incomplete smaller contract is ignored and the larger complete contract remains SAFE. If the essential larger contract lacks the required confirmation, material value passes but the result becomes CONDITIONAL pending a separate Article 209 assessment. Exact equality passes because the criterion uses $E \geq T$. Zero qualifying activity, zero acceptance or zero attribution produces FAIL. An unknown historical role or outcome-relevant evidence conflict produces UNRESOLVED.

Table 6. Counterfactual validation tests using the transformed dataset.

Test	Controlled change	Expected class	Purpose
Exact threshold	Set T equal to demonstrated value	SAFE	Boundary and numeric precision
Redundant incomplete reference	Remove smaller-contract attestation	SAFE	Admissible-subset optimisation

Test	Controlled change	Expected class	Purpose
Essential documentary gap	Remove dominant-contract attestation	CONDITIONAL	Material versus demonstrated value
Expired references	Place all acceptance dates outside W	FAIL	Temporal non-compensation
No similar activity	Set all α to zero	FAIL	Activity non-compensation
Unknown role/conflict	Make decisive role unresolved	UNRESOLVED	Operational conflict flag
Dominant-reference loss	Exclude larger contract	FAIL	Concentration, not core legal bound

Figure 3 presents deterministic retained-value sensitivity. The exact critical retained shares of strict demonstrated value are 21.25% for Lot A, 31.71% for Lot B and 35.18% for Lot C. Thus, a committee could exclude up to 78.75%, 68.29% and 64.82%, respectively, before the central demonstrated value falls below the threshold. Whether losses of that magnitude are legally plausible depends on the identified disputed lines; the ratios diagnose fragility but do not decide similarity. The core adverse scenario retains 79.64% and remains above every threshold, whereas loss of the dominant contract retains only 15.27% and fails all lots.

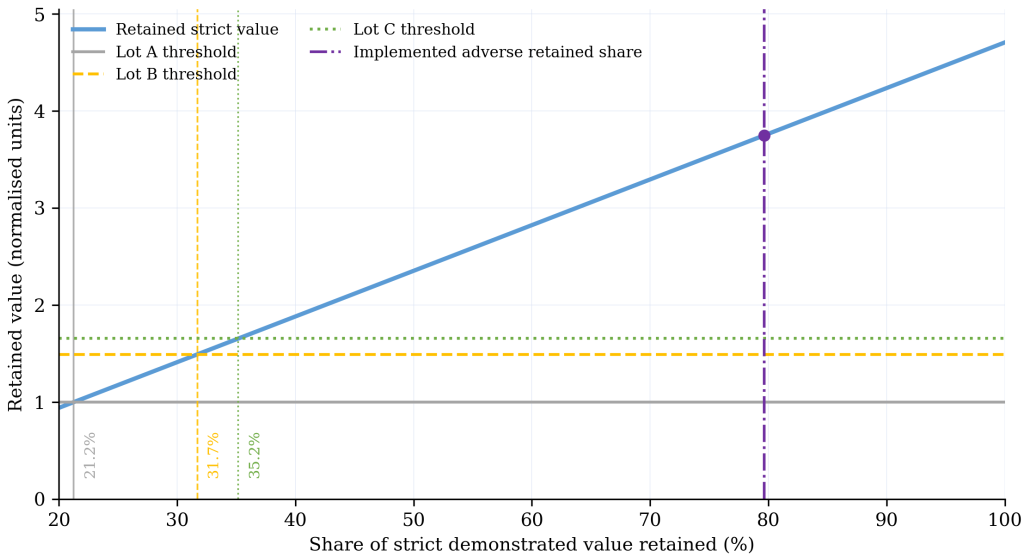

Figure 3. Deterministic retained-value sensitivity and exact critical shares for the three lot thresholds.

5. Discussion

5.1. Interpretation and limits of the proof of concept

The case demonstrates internal reproducibility and the practical usefulness of separating material, demonstrated and robust values. It does not validate the system across authorities, sectors or evaluators. The source case is transformed and the strict classifications were reviewed manually before coding. The proof of concept does not test empirical discriminative power across real decisions. The added counterfactual tests verify software logic and boundary behaviour only. A wider validation programme requires multiple procedures, independent expert labels, prospective shadow testing and reports of precision, recall, false acceptance, false rejection and inter-evaluator agreement.

5.2. Practical utility

The model reduces repetitive extraction and arithmetic while making every judgment visible. Each included line has a lot-specific rule, value and provenance link; each excluded line has a reason. Exact subset optimisation ensures that redundant defective contracts do not contaminate otherwise sufficient evidence. The report directs the expert toward decisive uncertainties instead of requiring a complete manual recalculation for every scenario.

The model can serve bidders before submission, contracting authorities during evaluation, internal controllers and legal teams reviewing challenges. Its strongest immediate outputs are not the labels alone, but the audit trail, missing-evidence list, critical retained share, contract-count slack and temporal slack. These variables explain why the conclusion is robust or fragile.

5.3. Comparison with alternative approaches and commercial systems

A nominal threshold rule is fast but overcounts mixed contracts and ignores period, role and evidence. Human-only review is legally flexible but consumes time and can produce inconsistent arithmetic or subset selection. A machine-learning classifier may scale across labelled decisions but predicts past outcomes and may obscure the active legal rule. Commercial procurement suites support supplier onboarding, qualification workflows, evidence collection and risk management (Coupa, 2026; Ivalua, 2026; JAGGAER, 2026), yet public product descriptions do not show the procedure-specific legal-mathematical audit trail proposed here. The present approach trades broad plug-and-play deployment for explicit rule configuration and reproducibility.

Table 7. Comparison with alternative qualification approaches.

Approach	Primary strength	Primary limitation	Role of proposed model
Nominal value summation	Speed	Ignores legal filters and mixed activities	Replaces gross value with admissible line value
Human-only review	Contextual legal judgment	Time, inconsistency and arithmetic risk	Automates calculation and exposes exceptions
Predictive ML	Scalability across labelled cases	Opacity and dependence on historical labels	Keeps the legal result deterministic
Commercial suites	Workflow, onboarding and supplier data integration	Publicly undocumented procedure-specific legal logic	Provides an auditable specialist module

5.4. Human discretion and lawful clarification

The human evaluator is not noise to be eliminated. Similarity and independent usability can remain contestable, and the official committee retains responsibility for applying the active procurement documents. The system therefore presents strict and central classifications, identifies value dependent on disputed lines and calculates the maximum-adversity scenario. A missing document may be remediable through a lawful request for supporting evidence or clarification, depending on what was initially declared and provided, but the system must not assume that every gap can be cured without testing equal treatment, transparency and the prohibition of evident advantage under Article 209 of Law no. 98/2016 (Romanian Parliament, 2016).

5.5. International and Romanian relevance

The proposal is consistent with eForms, the Public Procurement Data Space, the eProcurement Ontology and OECD work on data-driven procurement (European Commission, 2023a; European Commission, 2023b; OECD, 2025a; OECD, 2025b; Publications Office of the European Union, 2026). International examples show that automated analysis can support audit, integrity and knowledge management (OECD, 2025b; OECD, 2026; Pereira et al., 2025). Romania's SEAP supplies the public authority corpus and PREVENT demonstrates institutional acceptance of automated ex-ante checks (Electronic System for Public Procurement, 2026; National Integrity Agency of Romania, 2025). The contribution here is narrower: a transparent engine for a substantive qualification criterion.

5.6. Confidentiality, accountability and adoption limits

Public procurement transparency does not make all bidder evidence open data. Deployment must distinguish public authority documents from restricted commercial evidence and must retain access, retention and deletion controls. Human-AI guidelines and system-card approaches support visible limitations, editable outputs and auditability (Amershi et al., 2019; Gursoy & Kakadiaris, 2022). General-purpose AI should remain outside the deterministic legal kernel unless each output is independently validated.

Although the method is scientifically specifiable and technically feasible, it is not yet a publicly documented large-scale standard among Romanian procurement specialists or bidding firms (European Commission, 2020). Diffusion is constrained by scanned and versioned files, scarce labelled decisions, confidentiality, human legal responsibility and the cost of maintaining rule updates. The article therefore presents a sendable research prototype and roadmap, not a description of an already generalised practice.

5.7. Institutional sustainability and modular research agenda

The experience engine can be one module in a broader eligibility platform that separately evaluates exclusion grounds, financial standing, professional authorisations, personnel, quality and environmental management systems, conflicts of interest and procedure-specific technical conditions. Modular design is preferable to a monolithic model because legal changes and errors can be isolated, tested and updated without retraining the entire system.

Its sustainability contribution is institutional and operational rather than automatically environmental. Structured reuse can reduce duplicated documentary work, unnecessary model calls, repeated data transfers and avoidable clarification cycles. These claims must be measured through controlled studies of review time, error rates, rework, computing use and user confidence. AI itself has energy, hardware and market-concentration costs, so deterministic code should replace generative processing wherever it can perform the task adequately (Andhov et al., 2025).

Table 8. Sustainability-oriented modular research and development roadmap.

Module or stage	Next development	Evaluation metric	Expected contribution	Principal risk
Experience engine	Extend lot-specific taxonomies, evidence roles and jurisdictional rules	False acceptance/rejection; expert agreement	Reusable qualification analysis	Rule drift or overfitting
Additional eligibility modules	Add financial, authorisation, personnel and management-system checks	Cross-module consistency and exception rate	Broader platform with isolated legal updates	Error propagation
Cross-procedure validation	Test multiple authorities, sectors and contract types	Precision, recall, calibration and processing time	Evidence of transferability	Dataset shift
Privacy-preserving deployment	Use local, enterprise or federated processing	Retention, access and re-identification audits	Secure reuse of commercial evidence	Residual disclosure
Human-AI governance	Shadow mode and reviewer override analysis	Time saved, override reasons and confidence	Accountable adoption	Automation bias
Resource governance	Measure model calls, energy and storage per procedure	Energy and computing cost per correct decision	Computational restraint	Rebound effect

A staged implementation strategy should begin with retrospective reconstruction, continue in shadow mode during live procedures and move to assisted production only after accuracy, confidentiality and efficiency targets are met. Future controlled studies can test whether model assistance reduces median analyst time without increasing false acceptance, whether maximum-adversity reports reduce inter-evaluator dispersion, and whether modular reuse lowers repeated extraction and computing costs.

6. Conclusions

Bidder-experience verification can be represented as an auditable mathematical decision problem. The authority corpus defines lot-specific activity families, thresholds, reference periods, contract caps and proof rules; the bidder corpus supplies activity-level values, dates, roles and documentary links. The model separates material experience from documentary demonstration, applies binary non-compensatory legal filters, selects the best admissible subset of contracts and calculates a robust lower bound over an explicit legally plausible scenario set. Catastrophic concentration tests are reported separately rather than used to redefine the legal criterion.

The anonymised Romanian proof of concept demonstrates internal reproducibility and remains SAFE for all three transformed lots under the core scenario combining exclusion of the disputed smaller reference and a six-percent lower numeric correction. A separate concentration test shows that loss of the dominant reference would produce failure, making the model's dependence on that contract explicit. The result is also stable under the contract-count and temporal tests. The exercise does not establish external validity, and the proposed OCR/NLP layer is not yet implemented in the accompanying prototype.

General-purpose AI can support extraction, coding and explanation, but confidential evidence requires controlled processing and the final legal assessment remains human. The model's practical value lies in reducing repetitive work, preventing avoidable arithmetic and subset-selection errors, identifying missing evidence and explaining the role of each document. Its longer-term research value lies in modular expansion, cross-procedure human-controlled validation and, only after suitable data exist, separate econometric analysis of review behaviour and processing outcomes.

Acknowledgements

The proof-of-concept dataset is a privacy-preserving transformation of a professional analytical exercise. Direct identifiers have been removed and monetary and temporal data have been transformed; residual re-identification risk cannot be excluded.

Data and code availability

An anonymised replication package is supporting the manuscript and can be provided on request. It contains the two analytical series, the explicit adverse-scenario set, the deterministic eligibility engine, automated tests, generated outputs, figures, a variable dictionary and a checksum manifest. The transformed data reproduce the decision structure without disclosing direct identifiers or original monetary amounts.

Appendix A. Minimum implementation and validation protocol

Before deployment, a procedure-specific implementation should pass four gates. First, the authority corpus must be complete, versioned and legally frozen at the applicable submission deadline; every clarification answer and corrigendum must be linked to the rule it interprets or amends. Second, the bidder corpus must reconcile activity-line values with contract and acceptance totals, identify duplicated evidence and distinguish missing data from zero. Third, role-specific documentary requirements must be configured for the relevant legal relationship, rather than inferred from nominal party labels. Fourth, an independent procurement specialist should review the activity taxonomy, adverse scenarios and explanation report before any operational classification is released.

A prospective validation study should compare human-only and model-assisted review on the same independently labelled procedures. The release gate should require pre-specified ceilings for false acceptance and false rejection, stable results under exact-threshold and redundant-evidence tests, reproducible execution from a checksum-verified package, and documented resolution of every high-severity exception. Operational monitoring should record review time, clarification requests, overrides, rule changes, data-retention events and computing use. Material drift in legislation, procurement documents or activity taxonomies should suspend automatic classification until the affected module is reviewed and regression-tested. These controls convert the model from a one-off calculation into an auditable research and governance process while preserving the evaluation committee's legal responsibility.

References

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human–AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Article 3. https://doi.org/10.1145/3290605.3300233

Andhov, M., Darnall, N., & Andhov, A. (2025). Leveraging AI for sustainable public procurement: Opportunities and challenges. Frontiers in Sustainability, 6, 1603214. https://doi.org/10.3389/frsus.2025.1603214

Arroyo, R., Yebes, J., Martínez, E., Corrales, H., & Lorenzo, J. (2022). Key information extraction in purchase documents using deep learning and rule-based corrections. Proceedings of the First Workshop on Pattern-Based Approaches to NLP in the Age of Deep Learning, 11–20.

Balkan, D., & Akyuz, G. A. (2025). Artificial intelligence and machine learning in procurement and purchasing decision-support: A taxonomic literature review and research opportunities. Artificial Intelligence Review, 58, 341. https://doi.org/10.1007/s10462-025-11336-5

Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). Robust optimization. Princeton University Press. https://doi.org/10.1515/9781400831050

Bertsimas, D., & Sim, M. (2004). The price of robustness. Operations Research, 52(1), 35–53. https://doi.org/10.1287/opre.1030.0065

Coupa. (2026). Supplier qualification: Coupa glossary and supplier-management documentation [Official product information]. Retrieved July 17, 2026.

Electronic System for Public Procurement. (2026). SEAP public portal. Retrieved July 17, 2026, from https://www.e-licitatie.ro/

European Commission. (2020). Emerging technologies in public procurement. Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs.

European Commission. (2023a). eForms: EU legislative open standard for publishing public procurement notices. Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs.

European Commission. (2023b). Public procurement: A data space to improve public spending, boost data-driven policy-making and improve access to tenders for SMEs (2023/C 98 I/01).

European Commission. (2024). Commission launches a new general-purpose AI tool—GPT@EC. Retrieved July 17, 2026.

European Data Protection Board Support Pool of Experts. (2025). AI privacy risks & mitigations—Large language models.

European Parliament and Council of the European Union. (2014). Directive 2014/24/EU of 26 February 2014 on public procurement and repealing Directive 2004/18/EC. Official Journal of the European Union, L 94, 65–242.

European Parliament and Council of the European Union. (2016a). Directive (EU) 2016/943 of 8 June 2016 on the protection of undisclosed know-how and business information. Official Journal of the European Union, L 157, 1–18.

European Parliament and Council of the European Union. (2016b). Regulation (EU) 2016/679 of 27 April 2016 (General Data Protection Regulation). Official Journal of the European Union, L 119, 1–88.

Government of Romania. (2016). Government Decision no. 395/2016 approving the methodological norms for Law no. 98/2016, consolidated version. Romanian Legislative Portal.

Gursoy, F., & Kakadiaris, I. A. (2022). System cards for AI-based automated decision systems. arXiv. https://doi.org/10.48550/arXiv.2203.04754

Ivalua. (2026). Supplier management: Onboarding, qualification, risk, performance and compliance [Official product information]. Retrieved July 17, 2026.

JAGGAER. (2026). Supplier management and performance: Onboarding, compliance checks and configurable workflows [Official product information]. Retrieved July 17, 2026.

Moiraghi, F., Palmonari, M., Allavena, D., & Morando, F. (2024). Zero-shot hierarchical classification on the Common Procurement Vocabulary taxonomy. Proceedings of the 2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC), 273–282.

National Agency for Public Procurement. (2017). Instruction no. 2 of 19 April 2017 issued for the application of Articles 178 and 179(a)–(b) of Law no. 98/2016 and Articles 191 and 192(a)–(b) of Law no. 99/2016. Official Gazette of Romania, no. 300 of 27 April 2017.

National Integrity Agency of Romania. (2025). PREVENT integrated information system for ex-ante verification of potential conflicts of interest in public procurement. Retrieved July 17, 2026.

OECD. (2025a). Digital transformation of public procurement: Good practice report (OECD Public Governance Policy Papers No. 77). OECD Publishing. https://doi.org/10.1787/79651651-en

OECD. (2025b). Governing with artificial intelligence: The state of play and way forward in core government functions. OECD Publishing. https://doi.org/10.1787/795de142-en

OECD. (2026). Anti-corruption and integrity outlook 2026: Harnessing the integrity advantage. OECD Publishing. https://doi.org/10.1787/16708b78-en

Pereira, J., Assumpcao, A., Trecenti, J., Airosa, L., Lente, C., Cléto, J., Dobins, G., Nogueira, R., Mitchell, L., & Lotufo, R. (2025). INACIA: Integrating large language models in Brazilian audit courts—Opportunities and challenges. Digital Government: Research and Practice.

Publications Office of the European Union. (2026). The eProcurement Ontology. TED Developer Documentation. Retrieved July 17, 2026.

Romanian Parliament. (2016). Law no. 98/2016 on public procurement, consolidated version. Romanian Legislative Portal.

Sava, N.-A. (2023). Artificial intelligence and public procurement—Deciphering the interdisciplinary perspectives of the literature. European Review of Digital Administration & Law, 4(2), 79–88. https://doi.org/10.53136/97912218126648

Schneider dos Santos, E., Machado dos Santos, M., Castro, M., & Tyska Carvalho, J. (2025). Detection of fraud in public procurement using data-driven methods: A systematic mapping study. EPJ Data Science, 14, 52. https://doi.org/10.1140/epjds/s13688-025-00569-3

Siciliani, L., Taccardi, V., Basile, P., Di Ciano, M., & Lops, P. (2023). AI-based decision support system for public procurement. Information Systems, 119, 102284. https://doi.org/10.1016/j.is.2023.102284

Silaschi, I. O., Pop, I. D., & Coroiu, A. M. (2024). Improving transparency in Romanian public procurement: Machine learning to classify bidders and validate decisions. Procedia Computer Science, 246, 1210–1219. https://doi.org/10.1016/j.procs.2024.09.547

Zhao, Y., & Li, D. (2024). A large language model-based framework for semi-structured tender document retrieval-augmented generation. arXiv. https://doi.org/10.48550/arXiv.2410.09077