

RESEARCH ARTICLE



Reassessing Consumer Contract Formation and Execution in Autonomous and AI-Driven Contexts

Abdelhakim Adnani^{1,*}

¹*Faculty of Legal and Political Sciences, University of Hassan First, Morocco*

Abstract: This article proposes a hybrid computational legal analysis of consumer contract formation and execution in AI-driven contexts, treating Moroccan Law 31-08 and Law 09-08 as doctrinal foundations and computational legal methods as the primary analytical engine. The central research question is: how can defeasible logic and algorithmic compliance modeling operationalize consumer protection obligations in AI-mediated contractual environments? The study addresses three interlocking challenges: the validity of AI-mediated consent under conditions of algorithmic opacity, the legal recognition of personal data as contractual consideration, and the operationalization of the professional duty of conformity for autonomous and evolving AI products. The article introduces defeasible logic rule-sets (IC-1, DC-1, CF-1), two computational compliance models (the Algorithmic Consent Validation Engine (ACVE) model and the Automated Conformity Assessment Model (ACAM)), figures depicting model architecture and knowledge structures, and a multi-scenario simulation applied to a representative AI-driven e-commerce platform. The methodological contribution lies in making explicit how computational techniques complement and extend doctrinal analysis by offering scalability, auditability, and reproducibility in digital consumer protection enforcement. Comparative benchmarks are drawn from the EU AI Act, OECD AI Principles, and European Law Institute model rules on digital assistants. The article concludes with targeted recommendations for Moroccan legislative reform and RegTech implementation.

Keywords: artificial intelligence, consumer contracts, computational law, Moroccan law, RegTech

1. Introduction

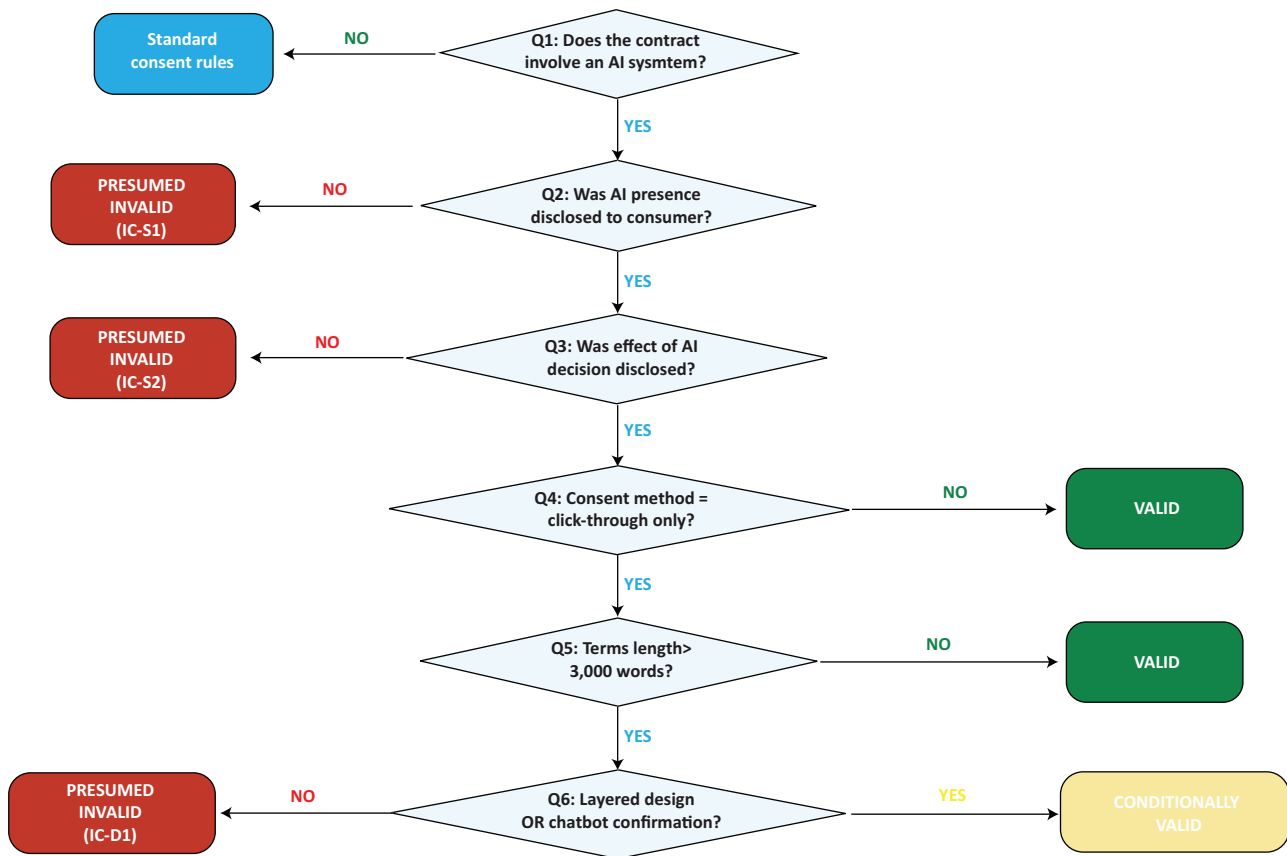
Artificial intelligence (AI) has become a transformative force across many sectors, from finance and healthcare to commerce and public services [1, 2]. In consumer markets specifically, AI systems now participate directly in contract formation and execution: they generate personalized offers, filter information environments, automate acceptance, and autonomously modify product behavior post-sale [3–5]. These developments challenge established principles of contract law, most acutely, informed consent, the balance of obligations, and the duty of conformity in ways that existing Moroccan and comparative doctrinal frameworks have not yet fully resolved [6, 7].

A critical methodological gap underlies existing scholarship: analyses of AI and consumer contracts have remained predominantly doctrinal, producing normative assessments of legal insufficiency without offering computational tools for operationalizing the obligations identified [8, 9]. Yet as AI systems increasingly act as contractual agents, the ability to encode, test, and audit legal obligations computationally is not merely a technical enhancement but it is a regulatory necessity. Scalable enforcement of consumer rights in AI-mediated environments requires frameworks that can be executed systematically, not only interpreted judicially.

This article addresses that gap by adopting a hybrid methodology: doctrinal analysis serves as the normative input, while computational legal methods constitute the primary analytical and constructive output. All proposed models and rule-sets are presented explicitly as illustrative frameworks, representing one possible operationalization of the relevant legal obligations; they are intended as demonstrations of potential, not as authoritative implementations. The central research question driving this inquiry is: how can defeasible logic and algorithmic compliance modeling operationalize consumer protection obligations in AI-mediated contractual environments, specifically under Moroccan Law 31-08 and Law 09-08? This contribution differs from prior work in the AI and Law formalization tradition, notably Governatori and Rotolo's [10] argumentation-based frameworks and Dell'Anna et al.'s [11] norm compliance models, in that it grounds formalization within the specific doctrinal constraints of an emerging-economy consumer law regime rather than abstractly modeling legal reasoning. Recent developments in AI governance, including the EU AI Act [12] and the OECD AI Principles [13], further underscore the urgency of operationalizing consumer protection obligations in algorithmic environments, providing additional comparative benchmarks for the reforms proposed here. The article is structured as follows. Section 2 develops the theoretical and computational framework, integrating doctrinal analysis with formal rule representations and figures. Section 3 presents the hybrid methodology in full. Section 4 applies the models to three simulation scenarios on an AI-driven e-commerce platform. Section 5 concludes with recommendations.

*Corresponding author: Abdelhakim Adnani, Faculty of Legal and Political Sciences, University of Hassan First, Morocco. Email: adnani.20006533@uhp.ac.ma

Figure 1
Visual decision tree for informed consent validity under rule-set IC-1 and Law 31-08/09-08. Leaf nodes indicate: VALID, CONDITIONALLY VALID, or PRESUMED INVALID



2. Theoretical and Computational Framework

2.1. Literature review

Artificial intelligence has attracted sustained scholarly attention across law, management science, and computer science. Mendoza-Caminade [14] provides a multidisciplinary analysis of legal frameworks governing AI in organizations, addressing liability, corporate governance, and compliance. Roder [15] examines the operational benefits of AI integration (automation, predictive analytics, and cost reduction) while emphasizing the importance of human oversight. El Badawi [16] offers a foundational overview of AI mechanisms and their ethical and organizational implications. These contributions share a predominantly descriptive orientation: they identify where legal frameworks fall short but do not formalize the obligations identified or model their computational implementation.

At the intersection of AI and law, Wachter et al. [17] demonstrate that fairness obligations cannot be fully automated without explicit legal rule encoding, a finding that motivates the present article's approach. Work in the AI and Law tradition particularly Governatori and Rotolo [10] on argumentation-based legal reasoning and Antoniou et al. [18] on defeasible logic provides the formal foundations for the rule-sets developed below. Dell'Anna et al. [11] demonstrate the applicability of defeasible logic to norm compliance in digital environments, directly informing the Algorithmic Consent Validation Engine (ACVE) model and Automated Conformity Assessment Model (ACAM). This article extends these traditions to the specific domain of consumer contract law in an emerging-economy context, making explicit a methodological contribution that the existing literature has not yet formalized for Moroccan law. Two further bodies of literature

warrant explicit engagement. First, the European Law Institute's Model Rules on Online Platforms¹ and Guiding Principles and Model Rules on Digital Assistants [19] offer authoritative soft-law instruments that address algorithmic governance in consumer contexts; these provide direct normative comparators for the rule-sets developed below. Second, recent work on algorithmic accountability, including Guidotti et al. [20] on interpretability and Kaminski and Malgieri [21] on algorithmic impact assessments, provides a theoretical bridge between computational compliance modeling and regulatory enforcement, thereby reinforcing the operationalization rationale of the present study.

2.1.1. AI and informed consent: doctrinal foundation and computational challenge

1) Doctrinal analysis

At the contract formation stage, AI operates in an invisible yet decisive manner, calling into question the scope of the consent expressed by the consumer. Three analytically distinct forms of harm are identifiable [22]. The first concerns silent algorithmic filtering: the consumer believes they enjoy freedom of choice, while what they see has been pre-selected according to criteria unknown to them. The second concerns manipulative interface design (calibrated notifications, dark patterns, and real-time behavioral adjustment) which creates a decision-making trajectory that appears free but is constrained.² The third is exploitation through personalized pricing: platforms use behavioral data to charge each consumer the maximum they are individually

¹ European Law Institute, *Model Rules on Online Platforms* (ELI Publications, 2020), <https://www.europeanlawinstitute.eu>.

² Guy Raymond, *Contrats de consommation*, in *JCI. Concurrence-Consommation*, fasc. 800, no. 47 (LexisNexis).

willing to pay, producing economic efficiency at the potential cost of legal fairness.

Consumer law responds to these risks through the obligation to provide clear, readable, and understandable information, an obligation extended to all consumer contracts under Law 31-08, Article 4, and reinforced with specific requirements of free, specific, informed, and unambiguous consent under Law 09-08³ [23]. The European Parliament has further emphasized that consumers must be informed of the existence of AI systems and their effects, and must be afforded the right to understand and challenge automated decisions [24, 25].

3) Formal rule representation: illustrative rule-set IC-1 (informed consent)

As a demonstration of how the doctrinal obligations identified above may be formalized, this study proposes the following illustrative rule-set encoded in defeasible logic. In this notation, strict rules ($\Rightarrow$) represent non-derogable obligations; defeasible rules ($\Rightarrow$) represent default conclusions rebuttable by exceptions; and defeat rules ($\rightarrow$) block specific defeasible conclusions. This is one possible formalization; alternative encodings are equally conceivable depending on jurisdiction and regulatory context.

Algorithm 1
Illustrative rule-set IC-1: informed consent obligations

```

STRICT RULE IC-S1 [Non-derogable AI Disclosure]:
  contract(C)  $\wedge$  consumer_contract(C)  $\wedge$  involves_ai(C)
   $\Rightarrow$  OBLIGATORY disclose_ai_presence(supplier(C), consumer(C))

STRICT RULE IC-S2 [AI Decision Effect Disclosure]:
  contract(C)  $\wedge$  automated_decision(C,D)  $\wedge$  materially_affects_consumer(D)
   $\Rightarrow$  OBLIGATORY disclose_decision_effect(supplier(C), consumer(C), D)

DEFEASIBLE RULE IC-D1 [Click-Through Inadequacy , IC-D1]:
  consent_method(C)='click_through'  $\wedge$  terms_length(C) > 3000
   $\Rightarrow$  consent_validity(C) = PRESUMED_INVALID
  UNLESS layered_design(C)=TRUE  $\vee$  chatbot_comprehension_confirmed(C)=TRUE

DEFEASIBLE RULE IC-D2 [Personalised Pricing Disclosure , IC-D2]:
  uses_personalised_pricing(C)  $\wedge$  ai_system(A)
   $\Rightarrow$  OBLIGATORY disclose_pricing_algorithm(supplier(C), consumer(C))
  UNLESS pricing_based_on_uniform_public_criteria(A)=TRUE

DEFEAT RULE IC-X1 [Chatbot Comprehension Defeats IC-D1]:
  interactive_chatbot_consent(C)=TRUE  $\wedge$  comprehension_confirmed(C)=TRUE
   $\rightarrow$   $\neg$ (consent_validity(C) = PRESUMED_INVALID)

OUTPUT STATES: { VALID | CONDITIONALLY_VALID | PRESUMED_INVALID }
NOTE: Illustrative , calibration to specific legal contexts required.

```

2) Illustrative consent validity decision tree

Figure 1 operationalizes the consent validity analysis as a decision tree, encoding the sequential conditional tests derived from Law 31-08 and Law 09-08: whether AI presence was disclosed, whether the effect of automated decisions was communicated, whether the consent mechanism was of sufficient quality, and whether compensatory design features are present. Reading from the root node, the tree branches at each mandatory checkpoint; any failure at a strict-rule node (IC-S1, IC-S2) produces an immediate INVALID verdict, while defeasible-rule failures produce CONDITIONALLY VALID outcomes that can be remediated. The decision tree renders the rule hierarchy visually explicit in a manner that prose analysis cannot, demonstrating how Rule IC-X1 (chatbot comprehension) operates as a defeat condition capable of reversing the IC-D1 penalty branch.

4) Algorithmic Consent Validation Engine (ACVE): illustrative model

Building on Rule-Set IC-1 and the decision tree of Figure 1, the ACVE model translates consent validation logic into an executable compliance module. The pseudo-code below specifies the engine's processing of a contract object through four sequential stages: disclosure checking (strict rules IC-S1, IC-S2), consent mechanism scoring, penalty application for structural inadequacies (defeasible rules IC-D1, IC-D2), and final classification. The ACVE model is designed as a gateway-layer tool: in a deployment scenario, it would intercept each contract formation event, run the four-stage assessment, and return a structured ConsentStatus object with a verdict and, where applicable, a machine-generated remediation plan. As discussed in Section 3.2.2, the scoring weights are heuristic in this study; empirical calibration against labelled contract corpora would be required before production deployment.

³ Morocco, Law No. 31-08 on Consumer Protection Measures, Dahir No. 1-11-03 of 18 February 2011, Official Bulletin No. 5932 (7 April 2011).

Algorithm 2

Illustrative pseudo-code: ACVE v1.0 (Algorithmic Consent Validation Engine, one possible approach)

```

CLASS ConsentValidator:
    FUNCTION validate_consent(contract C) → ConsentStatus:

        // STAGE 1: Mandatory disclosure checks (strict rules)
        IF C.involves_ai AND NOT C.ai_disclosure_provided:
            RETURN ConsentStatus(valid=FALSE, rule='IC-S1',
                                fix='Insert pre-contractual AI notice before consent screen')
        IF C.involves_ai AND NOT C.ai_decision_effect_disclosed:
            RETURN ConsentStatus(valid=FALSE, rule='IC-S2',
                                fix='Disclose how AI decisions affect the consumer')

        // STAGE 2: Consent mechanism quality scoring
        score = 0
        IF C.consent_method == 'interactive_chatbot': score += 40
        ELIF C.consent_method == 'layered_design': score += 35
        ELIF C.consent_method == 'click_through': score += 10
        IF C.comprehension_check_passed: score += 25

        // STAGE 3: Structural penalties (defeasible rules)
        IF C.terms_word_count > 3000 AND C.consent_method == 'click_through':
            score -= 30 // IC-D1 penalty
        IF C.uses_personalised_pricing AND NOT C.pricing_algo_disclosed:
            score -= 20 // IC-D2 penalty

        // STAGE 4: Classify and return
        IF score >= 70: RETURN ConsentStatus(valid=TRUE, level='VALID')
        IF score >= 40: RETURN ConsentStatus(valid=TRUE, level='CONDITIONAL')
        ELSE: RETURN ConsentStatus(valid=FALSE, level='INVALID',
                                fix=generate_remediation_plan(C))

        // NOTE: Score weights are illustrative; empirical calibration required.

```

5) Thoughtful consent: AI as facilitative and disruptive force

Beyond informed consent, Moroccan law provides a second protective layer through the withdrawal period (7 or 30 days depending on contract type) and the unenforceability of unfair terms under Title II of Law 31-08.⁴ AI-powered legaltechs, chatbots, legal design tools, interactive contract visualization, offer genuine potential to augment consent quality by making contractual terms more legible and accessible. In business-to-business (B2B) contexts, digital datarooms and AI-assisted drafting already approximate genuine negotiation. However, where structural imbalance exists between a professional and a consumer, the contract remains adhesive, and AI does not compensate for this asymmetry [26]. Algorithmic biases, whether explicit (deliberate design choices) or implicit (arising from machine learning (ML) training data), further compound this imbalance, as illustrated by dynamic pricing systems that exploit repeated search behavior to maximize supplier profit [27]. Rule IC-D2 in the illustrative Rule-Set IC-1 encodes this risk as a default disclosure obligation.

2.1.2. Data as consideration: doctrinal analysis and formal representation

1) Doctrinal foundation

Traditional consumer contract law positions monetary payment as the consumer's primary obligation.⁵ The digital economy has challenged this assumption fundamentally: consumers routinely access platforms presented as 'free' while in practice providing personal data, identity details, location, browsing history, behavioral traces, as the true consideration for services received. This data is collected, cross-referenced, and monetized through targeted advertising and algorithmic optimization, constituting an exchange with undeniable economic value.

From a civil law perspective, if a consumer's data has measurable economic value, denying the existence of consideration becomes difficult to sustain. This recognition entails three legal consequences: the need to characterize data as a form of consideration equivalent to price; the requirement that consent to data exchange satisfies the same standards (free, specific, informed, and unambiguous) imposed by Law 09-08; and the obligation to afford data-providers protections comparable to those traditionally accorded to monetary payers, including transparency regarding purpose, retention, and third-party sharing.⁶ This doctrinal

⁴ Morocco, Law No. 31-08 on consumer protection, promulgated by Dahir No. 1-11-03, Official Bulletin No. 5932 (2011).

⁵ European Commission, *Digital Contract Rules*, https://commission.europa.eu/business-economy-euro/doing-business-eu/contract-rules/digital-contracts/digital-contract-rules_en.

⁶ European Parliament and Council, *Directive 1999/44/EC on Consumer Goods and Associated*

position engages directly with ongoing debates in contract theory. The classical consideration doctrine, rooted in benefit-detriment analysis [28], supports the recognition of data as consideration: the consumer suffers a detriment (loss of personal information and privacy) while the platform receives a benefit (monetizable data asset). The EU Digital Content Directive 2019/770⁷ has already operationalized this logic by explicitly bringing contracts “where the consumer provides personal data” within its scope, a legislative precedent that Moroccan law has not yet matched. The principal counterargument, that data provision is merely incidental to service delivery, not a bargained-for exchange, is difficult to sustain where platforms systematically withhold access to users who decline data sharing, confirming the bargained character of the exchange. A second objection, grounded in the non-rivalrous and non-excludable properties of data (unlike money, data provision does not deplete the consumer's assets), calls for a calibrated rather than wholesale application of monetary-consideration doctrine: the key protective obligations (transparency, purpose limitation, and withdrawal right) are directly transferable; the financial remedy structure requires adaptation. Rule-Set DC-1 formalizes the transferable obligations while leaving the remedy architecture to future legislative and judicial elaboration.

2) Illustrative rule-set DC-1: data-as-consideration

Algorithm 3

Illustrative rule-set DC-1: data as consideration

```
STRICT RULE DC-S1 [Data Qualifies as Consideration]:
  service_provided_for_free(C) ∧ consumer_provides_data(C,D)
  ∧ data_has_economic_value(D)
  ⇒ data_is_consideration(C,D)=TRUE ∧ consumer_protection_applies(C)=TRUE

STRICT RULE DC-S2 [Transparency Obligations , Non-Derogable]:
  data_is_consideration(C,D)=TRUE
  ⇒ OBLIGATORY disclose_data_purpose(supplier(C), consumer(C))
    ∧ OBLIGATORY disclose_retention_period(supplier(C), consumer(C))
    ∧ OBLIGATORY disclose_third_party_sharing(supplier(C), consumer(C))

DEFEASIBLE RULE DC-D1 [Inferred Non-Awareness of Data Payment]:
  data_is_consideration(C,D)=TRUE ∧ terms_explicitly_mention_data_value(C)=FALSE
  ⇒ consumer_aware_of_payment(C)=FALSE
  UNLESS consumer_independently_knowledgeable(C)=TRUE

DEFEASIBLE RULE DC-D2 [Right to Withdraw Data Consent]:
  data_is_consideration(C,D)=TRUE ∧ consumer_withdraws_data_consent(C)
  ⇒ contract_must_terminate_or_renegotiate(C)=TRUE
  UNLESS withdrawal_economically_disproportionate(C)=TRUE

OUTPUT: { COMPLIANT | NON_COMPLIANT | REQUIRES_REVIEW }
NOTE: Illustrative , one possible encoding. Regulatory calibration required.
```

2.1.3. The duty of conformity for AI-driven products: doctrine and ACAM

1) Doctrinal analysis

The duty of conformity, long central to consumer law, acquires renewed complexity when applied to AI-based products [27]. Unlike

traditional goods, AI systems may evolve through post-sale updates, learn from user interactions, and exhibit emergent behaviors not anticipated at the point of sale. European Directive 1999/44/CE defined consumer goods as ‘any tangible movable property’,⁸ which initially excluded intangible software from the guarantee of conformity regime. Moroccan Law 31-08 is potentially more adaptive: its use of deliberately broad language ‘products, goods, or services’ without distinguishing tangible from intangible creates interpretive space for including software and AI applications within its scope. This interpretive position is not without counterarguments. A first objection holds that the conformity regime presupposes a fixed reference point, the product as described at sale, that is structurally incompatible with inherently adaptive AI systems whose performance is designed to evolve. On this view, applying conformity obligations to self-improving AI products would paradoxically penalize beneficial adaptation. Rule CF-D3 addresses this objection directly by limiting the non-conformity finding to cases of performance degradation (not improvement) that was not disclosed at sale, thus preserving space for beneficial evolution while protecting against undisclosed deterioration. A second objection concerns the causal attribution of post-sale performance changes: where

AI performance depends on consumer-provided data, degradation may be attributable to data quality rather than to the product itself. Rule CF-D4 encodes this exception, ensuring that suppliers are not held liable for non-conformity attributable to the consumer's own inputs, provided the data quality requirement was disclosed at sale. These internal limits to the conformity regime, formalized in the defeat conditions of CF-D1 through CF-D4, demonstrate that the framework is not a blanket

⁷ Guarantees, Official Journal of the European Communities L171, 12–16 (1999).

⁸ European Commission, Digital Contract Rules.

⁸ Directive 1999/44/EC, OJ L171, 12–16.

imposition of liability on AI suppliers, but a calibrated allocation of responsibility between platform capabilities and consumer obligations.

2) Illustrative rule-set CF-1: conformity for AI products

quality dependency (CF-D4); Stage 3 aggregates findings into a severity-weighted verdict. This architecture maps the legal rule hierarchy directly onto an audit workflow, making the relationship between legal obligation and computational check explicit and auditable at each step.

Algorithm 4

Illustrative rule-set CF-1: conformity obligations for AI products

```
STRICT RULE CF-S1 [AI Products Within Scope]:
  product(P) ∧ ai_based(P) ∧ supplied_under_consumer_contract(P)
  ⇒ conformity_obligations_apply(P)=TRUE

STRICT RULE CF-S2 [Core Fitness-for-Purpose]:
  conformity_obligations_apply(P)=TRUE
  ⇒ OBLIGATORY performs_stated_function(P)
    ∧ OBLIGATORY meets_reasonable_consumer_expectations(P)
    ∧ OBLIGATORY safe_for_intended_use(P)

DEFEASIBLE RULE CF-D1 [Undisclosed Autonomous Behaviour]:
  ai_product(P) ∧ exhibits_autonomous_behaviour(P,B)
  ∧ behaviour_not_disclosed_in_contract(B)
  ⇒ product_non_conforming(P)=TRUE
  UNLESS behaviour_within_disclosed_operational_range(B)=TRUE

DEFEASIBLE RULE CF-D2 [Update-Induced Degradation]:
  ai_product(P) ∧ updated_post_sale(P) ∧ update_degrades_functionality(P)
  ⇒ product_non_conforming(P)=TRUE
  UNLESS consumer_consented_to_update(P)=TRUE
    ∧ functionality_restored_within_reasonable_time(P)=TRUE

DEFEASIBLE RULE CF-D3 [Machine Learning Drift]:
  ai_product(P) ∧ uses_machine_learning(P)
  ∧ performance_degrades_significantly_over_time(P)
  ⇒ product_non_conforming(P)=TRUE
  UNLESS performance_degradation_disclosed_at_sale(P)=TRUE

DEFEASIBLE RULE CF-D4 [Consumer Data Quality Dependency]:
  ai_product(P) ∧ performance_depends_on_consumer_data(P)
  ∧ consumer_data_of_insufficient_quality(P)
  ⇒ non_conformity_attributable_to_consumer(P)=TRUE
  UNLESS consumer_not_informed_of_data_quality_requirement(P)=TRUE

VERDICTS: { CONFORMING | CONDITIONALLY | PARTIALLY | NON_CONFORMING }
```

3) ACAM workflow and figure integration

As one possible operationalization of the duty of conformity for AI products, this study proposes the ACAM: a three-stage automated assessment model. As shown in Algorithm 5 below, ACAM processes an AI product through three stages: product characterization (Stage 1), rule application against CF-1 (Stage 2), and verdict generation with severity-weighted violation scoring (Stage 3). The three-stage architecture reflects the logical structure of CF-1: Stage 1 determines whether the product falls within scope (CF-S1); Stage 2 applies each rule in the defeasible hierarchy, checking for undisclosed autonomous behavior (CF-D1), update-induced degradation (CF-D2), ML drift (CF-D3), and data

2.2. Illustrative knowledge graph: linking core obligations

Figure 2 visualizes the relational structure among the three principal legal obligations formalized in this study. Reading the graph from left to right: the contract formation event activates both the IC-1 consent obligation and the DC-1 data-as-consideration obligation simultaneously, reflecting the doctrinal finding that these two rule-sets are jointly triggered by the same transactional event. The CF-1 conformity obligation is connected to the contract execution node and to the AI product node, capturing the post-sale dimension of the duty of conformity. The computational models (ACVE, ACAM) are

Algorithm 5

Illustrative pseudo-code: ACAM v1.0 (Automated Conformity Assessment Model, for example purposes)

```

CLASS ConformityAssessor:
    FUNCTION assess_conformity(product P, contract C) → ConformityReport:

        // STAGE 1: Product characterisation
        profile.is_ai      = detect_ai_components(P)
        profile.uses_ml    = detect_ml_system(P)
        profile.updates    = detect_update_mechanism(P)
        profile.data_dep   = detect_data_dependency(P)
        profile.autonomy   = detect_autonomous_behaviour(P)
        violations = []

        // CF-S1: Scope check
        IF NOT profile.is_ai OR NOT C.is_consumer_contract:
            RETURN ConformityReport(scope='OUT_OF_SCOPE')

        // CF-S2: Fitness for purpose
        IF NOT P.performs_stated_function():
            violations.add(Violation('CF-S2', 'HIGH', 'Stated function not performed'))

        // CF-D1: Undisclosed autonomous behaviour
        FOR b IN P.autonomous_behaviours:
            IF b NOT IN C.disclosed_range:
                violations.add(Violation('CF-D1', 'MEDIUM', 'Undisclosed behaviour: '+b))

        // CF-D2: Update-induced degradation
        IF profile.updates:
            FOR u IN P.post_sale_updates:
                IF u.degrades AND NOT C.update_consented:
                    violations.add(Violation('CF-D2', 'HIGH', 'Degrading update: '+u.id))

        // CF-D3: ML performance drift
        IF profile.uses_ml:
            drift = P.measure_drift(baseline=C.sale_date)
            IF drift > 15 AND NOT C.drift_disclosed:
                violations.add(Violation('CF-D3', 'MEDIUM', 'ML drift: '+drift+'%'))

        // CF-D4: Consumer data quality
        IF profile.data_dep AND P.data_quality() < P.min_quality():
            IF NOT C.data_req_disclosed:
                violations.add(Violation('CF-D4', 'LOW', 'Data quality undisclosed'))

        // STAGE 3: Verdict
        high = COUNT(v WHERE v.severity='HIGH')
        med  = COUNT(v WHERE v.severity='MEDIUM')
        IF high>0: verdict='NON_CONFORMING'
        ELIF med>1: verdict='PARTIALLY_CONFORMING'
        ELIF med=1: verdict='CONDITIONALLY_CONFORMING'
        ELSE:      verdict='CONFORMING'
        RETURN ConformityReport(verdict, violations,
            generate_remediation(violations))

```

represented as processing nodes that receive rule-set inputs and return compliance verdicts. The Moroccan legal sources (Law 31-08, Law 09-08) anchor each obligation node, making the normative provenance of each computational rule explicit. This knowledge graph representation, drawing on the RDF/OWL tradition in the AI and Law field, reveals structural interdependencies that a linear doctrinal exposition necessarily flattens: in particular, it makes explicit that a single consumer data-sharing event simultaneously engages IC-1 (consent quality), DC-1


(data as consideration), and potentially CF-1 (if data quality affects AI product performance under CF-D4).

3. Research Methodology: A Hybrid Computational Legal Framework

This study adopts a hybrid research design that positions computational legal methods as the primary analytical engine, with

Figure 2

Knowledge graph linking core consumer protection obligations, illustrative rule-sets (IC-1, DC-1, CF-1), computational models (ACVE, ACAM), and Moroccan legal foundations (Law 31-08, Law 09-08)



doctrinal analysis as the normative foundation. The rationale for this methodological choice is explicitly substantive: as AI systems increasingly function as contractual agents, traditional doctrinal analysis, however rigorous, cannot alone produce the scalable, auditable, and reproducible compliance tools that consumer protection enforcement in digital environments requires.

3.1. Doctrinal foundation (input layer)

This study employs doctrinal-analytical research as its normative input layer. Doctrinal research involves systematic review and analysis of legal sources (statutes, regulations, case law, and scholarly commentary) to interpret and evaluate existing legal rules [29]. The primary sources analyzed are Moroccan Law 31-08 (consumer protection), Law 09-08 (personal data protection), and the DAHIR of obligations and contracts. European Union sources, Directive 1999/44/CE on consumer goods⁹ and Directive 2019/770 on digital content (comparative reference),¹⁰ are drawn upon to identify reform benchmarks. No primary human-subject data was collected. Crucially, doctrinal analysis is used not only as an end in itself, but also as the normative input from which computational rules are derived: it answers ‘what does the law require?’; the computational layer translates that answer into ‘how can that requirement be systematically checked?’

⁹ Directive 1999/44/EC, OJ L171, 12–16.

¹⁰ European Commission, *Digital Contract Rules*.

3.2. Computational legal methods (core analytical layer)

3.2.1. Defeasible logic formalization

Defeasible logic [30] is chosen as the formalization framework because it accurately mirrors the structure of legal reasoning. Legal rules are characteristically defeasible: they hold as defaults unless overridden by more specific rules, exceptions, or contrary evidence. This choice merits explicit justification relative to alternative formalisms. Deontic logic, the dominant tradition in formal legal modeling, captures normative operators (obligation, permission, prohibition) but struggles with norm conflicts and the closed-world assumption that characterizes static rule systems, limitations that are particularly acute in consumer law, where protective defaults interact dynamically with exceptions [31]. First-order logic offers greater expressive power but lacks the built-in exception-handling mechanisms that legal reasoning requires. Defeasible logic occupies a pragmatically optimal position: it accommodates norm conflicts through explicit priority orderings, supports open-world reasoning, and has been demonstrated to scale to realistic regulatory corpora [30]. The question of rule consistency is addressed structurally: strict rules (non-derogable) take unconditional priority over defeasible rules; within the defeasible layer, defeat rules function as explicit conflict resolvers, ensuring that the rule-sets produce determinate outputs without requiring external conflict-resolution mechanisms. This tripartite structure, strict rules (non-derogable obligations), defeasible rules (rebuttable defaults), and defeat rules (explicit exception triggers), maps directly onto consumer law, where non-derogable protective obligations (Law 31-08, Art. 4) coexist with specific exceptions (e.g.,

Rule IC-X1). The rule-sets developed in Section 2 represent one possible operationalization of informed consent, data-as-consideration, and conformity obligations; alternative encodings are equally conceivable, and empirical calibration would be required before any deployment.

3.2.2. Algorithmic compliance modeling

The ACVE model and ACAM translate the defeasible rule-sets into procedural compliance algorithms expressed as pseudo-code. This translation demonstrates how legal rules could be embedded in software compliance modules, for example, within a platform's onboarding flow (ACVE) or a product liability audit system (ACAM). The pseudo-code level of specification is deliberately chosen: it is sufficiently precise to demonstrate implementability while remaining accessible to legal scholars without software engineering backgrounds. The computational enhancements over purely doctrinal analysis are threefold: scalability (hundreds of contracts assessed against the same rule-set simultaneously), auditability (every decision path explicit and reviewable), and reproducibility (the same rule-set applied to the same fact pattern always produces the same outcome). Regarding implementation environments, the ACVE model is architecturally suited for integration into platform Application Programming Interface (API) gateway layers, where contract formation events are intercepted and assessed prior to acceptance, or as a standalone RegTech module queried by compliance officers. ACAM is designed for periodic audit cycles, operating on product telemetry and update logs. Both models are compatible with standard rule-engine frameworks (e.g., Drools and OpenRules) that support defeasible-style priority-based rule execution. Concerning calibration and validation, the scoring thresholds in the ACVE model (e.g., the 70/40 score boundaries for VALID/CONDITIONAL classifications, and the 3,000-word penalty trigger in IC-D1) are assigned heuristically in this study to demonstrate the model's logic. Empirical calibration would proceed through two stages: (i) ground-truth labelling of a corpus of real consumer contracts by legal experts, establishing benchmark consent validity classifications; and (ii) optimization of threshold parameters to minimize classification error against that benchmark, potentially using supervised learning techniques. The 15% ML drift threshold in CF-D3 is similarly illustrative; calibration against sector-specific performance data (e.g., smart home device false-alarm baselines) would be required for deployment.

3.2.3. Knowledge graph representation

Legal knowledge modeling through directed graphs (Figure 2) provides a relational view of the normative structure that linear prose cannot convey. By representing obligations, rules, models, and legal sources as nodes and their relationships as typed directed edges, the knowledge graph reveals structural interdependencies for example, that consent validity (IC-1) and data-as-consideration transparency (DC-1) are jointly activated by the same contract formation event. This representation draws on the RDF/OWL tradition in the AI and Law field.

3.3. Simulation and scenario-based validation

The computational frameworks developed in Section 2 are evaluated through scenario-based simulation (Section 4). Three concrete

scenarios are constructed around a hypothetical AI-driven e-commerce platform ('SmartShop AI'), each designed to exercise distinct aspects of the rule-sets and models: AI-mediated consent validation, data-as-consideration compliance auditing, and AI product conformity assessment for a machine-learning-based product with post-sale degradation. Scenario-based validation is used here as a demonstration of potential operationalization rather than an empirical test of predictive accuracy, which would require deployment and real-world calibration beyond the scope of the present study.

3.4. Limitations

Four principal limitations are acknowledged. First, the rule-sets and models are illustrative: they encode one reasonable interpretation of Moroccan and comparative law, but alternative encodings are equally legitimate. Second, the scoring weights in the ACVE model are assigned heuristically for demonstration; empirical calibration would be required for deployment. Third, the study does not address enforcement mechanisms or the procedural law governing algorithmic compliance evidence. Fourth, as Moroccan AI-specific legislation remains nascent, the rule-sets rely on purposive interpretation of existing law that has not yet been judicially confirmed (Table 1).

4. Applied Simulations: The SmartShop AI Platform

This section applies the compliance models and rule-sets developed above to three concrete simulation scenarios on a hypothetical AI-driven e-commerce platform, designated 'SmartShop AI'. As shown in Figure 3, the platform employs five AI components: a collaborative-filtering recommendation engine, a dynamic real-time pricing algorithm, an automated contract formation module, an AI-powered customer service chatbot, and a personalized behavioral analytics pipeline. The simulations demonstrate how the ACVE model, ACAM, and DC rule-engine would function in practice. It is important to note explicitly that these are designed scenarios, not empirical observations from a real platform: the input parameters are constructed to exercise each rule-set, not drawn from field data. The scenarios therefore demonstrate logical operationalization, showing that the models produce legally coherent outputs, rather than empirical validation of predictive accuracy. Full empirical validation would require deployment on real contracts, comparison against expert-labelled ground truth, and statistical performance assessment, as discussed in Section 3.4. This limitation is consistent with the study's scope as a conceptual legal engineering contribution; the simulation findings are valid as logical demonstrations of the framework, not as data-driven predictions.


4.1. Scenario 1: ACVE model application, consent validation for AI-mediated onboarding

Consumer A visits SmartShop AI and is shown personalized product recommendations. Prices vary dynamically based on browsing frequency. At checkout, A is presented with a 5,200-word terms-of-service document and clicks 'Agree'. No explicit AI disclosure is

Table 1
Hybrid methodology overview. Each method's role, primary output, and computational value added

Method	Role in study	Primary output	Computational value added
Doctrinal analysis	Normative input layer	Legal obligation identification	Provides rule content
Comparative legal analysis	Reform benchmarking	Regulatory gap mapping	Identifies missing norms
Defeasible logic formalization	Core analytical engine	Rule-Sets IC-1, DC-1, CF-1	Scalability, auditability
Algorithmic compliance modeling	Core analytical engine	ACVE v1.0, ACAM v1.0	Reproducibility, automation
Scenario-based simulation	Illustrative validation	Section 4 scenarios	Demonstrates operationalization

Figure 3
SmartShop AI compliance schema illustrating AI components, consumer data processing, applicable legal frameworks, and regulatory objectives



provided; no layered design or chatbot confirmation is used. Applying Rule-Set IC-1 and the ACVE model to this scenario, the following illustrative assessment results (Algorithm 6):

Under Law 31-08, Article 4, the supplier is obligated to inform the consumer of the essential characteristics of the service by ‘any appropriate means’. Failure to disclose the AI recommendation

Algorithm 6

Scenario 1 simulation — ACVE model execution: consent validation for AI-mediated onboarding (SmartShop AI, Consumer A)

```

INPUT:
  contract.involves_ai           = TRUE
  contract.ai_disclosure_provided = FALSE
  contract.ai_decision_disclosed  = FALSE
  contract.consent_method        = 'click_through'
  contract.terms_word_count      = 5200
  contract.uses_personalised_pric = TRUE
  contract.pricing_algo_disclosed = FALSE
  contract.layered_design         = FALSE
  contract.comprehension_check    = FALSE

ACVE EXECUTION:
  STAGE 1: IC-S1 check → ai_disclosure_provided = FALSE
  ► EXIT: ConsentStatus(valid=FALSE, rule='IC-S1',
    fix='Insert pre-contractual AI notice before consent screen')

VERDICT: INVALID
  Primary violation: IC-S1 (strict , non-derogable)
  Secondary violations (if IC-S1 remediated): IC-S2, IC-D1, IC-D2

REMEDIATION PLAN (illustrative):
  R1: Add visible AI disclosure notice prior to consent screen [IC-S1]
  R2: Disclose dynamic pricing algorithm in plain language [IC-S2, IC-D2]
  R3: Implement layered consent design (summary + detail layers) [IC-D1]
  R4: Add interactive comprehension check or chatbot onboarding [IC-X1]
  
```

and dynamic pricing systems constitutes a direct breach of this pre-contractual obligation. The ACVE model exits at Stage 1 because Rule IC-S1 is strict (non-derogable); no amount of consent mechanism quality can remedy undisclosed AI involvement. This scenario illustrates how the computational model reflects the legal hierarchy between strict and defeasible obligations, a hierarchy that purely textual analysis identifies but cannot enforce algorithmically.

4.2. Scenario 2: DC rule-engine, data-as-consideration compliance audit

Consumer B subscribes to SmartShop AI's 'free' premium feature tier. In exchange, B's behavioral data is collected and sold to third-party advertisers. The terms mention data collection but do not quantify its economic value, do not specify third-party recipients, and do not provide a withdrawal mechanism. Applying Rule-Set DC-1 (Algorithm 7):

4.3. Scenario 3: ACAM application, conformity assessment for ML drift

Consumer C purchases an AI-powered home security system ('SecureBot AI') from SmartShop AI. Six months post-sale, SecureBot AI receives an automatic software update that retrain its anomaly-detection ML model. Following the update, the product's false-alarm rate increases by 38% compared to performance at the point of sale. Neither the possibility of post-update performance degradation nor the ML drift threshold was disclosed at sale. Applying ACAM (Algorithm 8):

This scenario illustrates the particular challenge posed by self-updating AI products to classical conformity doctrine. A product that performed correctly at sale becomes non-conforming through an autonomous update, a situation without clear precedent in Moroccan consumer law but directly addressable through the CF-D2 defeasible rule. The ACAM's verdict of 'NON_CONFORMING' triggers the consumer's conformity remedy rights under Law 31-08, Title V.

Algorithm 7

Scenario 2 simulation — DC rule-engine execution: data-as-consideration compliance audit (SmartShop AI, Consumer B)

```

INPUT:
  contract.service_free           = TRUE
  contract.consumer_provides_data = TRUE
  contract.data_economic_value    = TRUE
  contract.terms_mention_value    = FALSE
  contract.purpose_disclosed        = FALSE
  contract.retention_disclosed     = FALSE
  contract.third_party_disclosed   = FALSE
  contract.withdrawal_provided    = FALSE

DC ENGINE EXECUTION:
  DC-S1: data_is_consideration = TRUE (all conditions met)
  DC-S2: STRICT obligations triggered:
    ► VIOLATION: disclose_data_purpose      → NOT MET
    ► VIOLATION: disclose_retention_period → NOT MET
    ► VIOLATION: disclose_third_party     → NOT MET
  DC-D1: consumer_aware_of_payment = FALSE
  DC-D2: withdrawal_mechanism_provided = FALSE → VIOLATION

VERDICT: NON_COMPLIANT
  Violations: DC-S2 (×3 sub-violations, strict), DC-D1, DC-D2

REMEDIATION (illustrative):
  R1: Add 'data-as-payment' notice with economic value indication [DC-S1]
  R2: Specify all third-party data recipients and purposes [DC-S2]
  R3: State data retention periods clearly [DC-S2]
  R4: Implement one-click data consent withdrawal [DC-D2]

```

The DC engine correctly identifies that the three DC-S2 violations are strict and therefore non-derogable: no contextual exception can excuse failure to disclose data purpose, retention, and third-party sharing, because these obligations derive directly from Law 09-08's consent requirements. Rule DC-D1 adds the inference that consumer awareness of data payment cannot be presumed from mere terms-of-service acknowledgment, directly addressing the 'apparent gratuity' phenomenon identified in the doctrinal analysis.

4.4. Cross-scenario comparative assessment

Figure 4 presents a comparative bar chart summarizing violation severity distribution across the three simulation scenarios. This visualization makes explicit a pattern that doctrinal prose tends to obscure: strict rule violations (HIGH, non-derogable) require no contextual analysis, while the severity profile of defeasible rule violations varies significantly across scenarios, informing compliance prioritization.

Algorithm 8

Scenario 3 simulation — ACAM execution: conformity assessment for post-sale ML drift (SmartShop AI, Consumer C / SecureBot AI)

```

INPUT:
product.is_ai           = TRUE
product.uses_ml         = TRUE
product.self Updating   = TRUE
contract.is_consumer    = TRUE
product.stated_fn_intact = TRUE    // basic function still operates
update.degrades_fn      = TRUE    // false-alarm rate +38%
contract.update_consented = FALSE
drift_vs_baseline       = 38%    // > 15% illustrative threshold
contract.drift_disclosed = FALSE

ACAM STAGE 2:
CF-S1: In scope ✓ | CF-S2: Stated function intact ✓
CF-D1: No undisclosed pre-update autonomous behaviour ✓
CF-D2: Update post-sale → degrades → NOT consented
      ► VIOLATION HIGH: CF-D2
CF-D3: ML drift = 38% > 15% threshold → NOT disclosed at sale
      ► VIOLATION MEDIUM: CF-D3
CF-D4: N/A

VERDICT: NON_CONFORMING (1 HIGH + 1 MEDIUM violation)


REMEDIATION (illustrative):
R1: Roll back update OR restore equivalent performance [CF-D2]
R2: Offer conformity remedies: repair / replacement / price
    reduction / termination , Law 31-08, Title V [CF-D2]
R3: Implement update consent mechanism for future ML retraining
R4: Add pre-sale disclosure of ML performance drift range [CF-D3]

```

As demonstrated in Figure 4, Scenario 2 (Data-as-Consideration) produces the highest number of strict-rule violations, reflecting the non-derogable nature of the Law 09-08 transparency obligations encoded in DC-S2. This finding resonates with Wachter et al.'s [17] argument

that fairness obligations in AI systems cannot be automated without explicit legal rule encoding: the DC-S2 violations are precisely the obligations that existing 'apparent gratuity' platform designs fail to encode, rendering them invisible to purely doctrinal analysis. Scenario 3 (Conformity) produces fewer violations in total but includes a HIGH-severity finding from Rule CF-D2, indicating that the consumer's conformity remedy rights are directly triggered, a result that extends Dell'Anna et al.'s [11] norm compliance framework to the novel domain of self-updating AI products. Scenario 1 (Consent) shows the starkest outcome: a single strict-rule failure (IC-S1) terminates the ACVE model at Stage 1, confirming that disclosure obligations are structurally prior to consent mechanism quality, a hierarchy that doctrinal analysis identifies textually but that the computational model enforces procedurally. This cross-scenario comparison illustrates a further advantage of computational legal analysis: the ability to aggregate, compare, and rank compliance findings systematically across multiple contracts or scenarios, enabling prioritization of regulatory intervention in ways that purely textual scholarship cannot produce [21].

Figure 4
Comparative violation severity distribution across three simulation scenarios (SmartShop AI platform). HIGH = strict rule violations; MEDIUM = defeasible rule violations. Violation counts derive from simulation outputs above; not empirical data



5. Conclusion

The evolution of consumer contracts in the digital era reveals a profound transformation in both the structure of contractual obligations and the scope of consumer protection. This article has advanced existing scholarship in two complementary and integrated directions, making

explicit the methodological premise that computational legal methods and doctrinal analysis are not alternatives but complements.

At the doctrinal level, the study has demonstrated that three legal challenges, the validity of AI-mediated consent, the recognition of data as contractual consideration, and the application of conformity obligations to AI-driven products, are tractable within Moroccan legal frameworks (Law 31-08, Law 09-08) while requiring regulatory clarification. The broad language of Law 31-08 offers valuable interpretive flexibility for including AI-based applications within its scope; Law 09-08's consent requirements provide the normative basis for addressing data-as-consideration transparency failures.

At the computational level, the article has introduced three defeasible logic rule-sets (IC-1, DC-1, CF-1), two algorithmic compliance models (ACVE v1.0, ACAM v1.0), figures depicting model architecture and knowledge structures, and three scenario simulations demonstrating how these frameworks operationalize legal obligations in practice. The framing throughout is deliberately demonstrative: the models illustrate the formalization process and the potential of computational legal methods; they do not constitute regulatory tools ready for deployment without empirical calibration and institutional validation.

The methodological contribution of the hybrid approach is explicit: computational techniques add scalability, auditability, and reproducibility, advantages that purely doctrinal analysis cannot offer and that are essential to consumer protection enforcement in AI-mediated digital environments. The article's theoretical positioning also warrants reflection in relation to competing frameworks. The rights-based approach dominant in European AI governance scholarship, exemplified by the EU AI Act's [12] high-risk classification system and the ELI's Model Rules on Digital Assistants [19], addresses algorithmic accountability through ex ante risk categorization rather than ex post compliance assessment. The present article's rule-based operationalization approach is complementary rather than competing: it provides the specific compliance-checking logic that risk-categorization frameworks presuppose but do not themselves supply. The OECD AI Principles [13] similarly articulate high-level obligations (transparency, explainability, robustness) without formalizing the compliance assessment mechanisms needed for enforcement, a gap the ACVE model and ACAM are designed to illustrate how to fill. Future research should pursue empirical calibration of the models against real contract corpora, extension of the framework to other consumer law domains (e.g., financial services AI under Moroccan Law 103-12), examination of the procedural law questions raised by algorithmic compliance evidence in judicial proceedings, and further comparative analysis across jurisdictions at comparable stages of digital market regulation.

Recommendations

Legislative reform (Law 31-08/09-08): legislators should amend Law 31-08 to include an explicit provision on pre-contractual AI disclosure, and Law 09-08 to expressly recognize the provision of personal data as triggering consumer protection obligations equivalent to monetary payment. The rule-sets formalized in this article represent one illustrative technical annex that could inform such reform.

RegTech implementation: AI platform operators should consider implementing automated compliance engines based on the ACVE model and ACAM frameworks, adapted to their specific platform architectures and in consultation with data protection and consumer protection authorities. The growing literature on automated unfair-clause detection [32] and RegTech compliance systems [33] provides practical precedents for this implementation pathway, demonstrating that rule-based computational compliance is already operationally feasible in comparable consumer protection domains.

Regulatory enforcement: consumer protection authorities should explore the adoption of algorithmic conformity and consent assessment tools, such as those illustrated by the ACAM and visualized in Figure 4, as part of their market surveillance functions.

Consumer education: because consumers frequently underestimate the economic value of their data and the implications of AI-mediated contract formation, targeted digital literacy initiatives should be developed, with particular attention to the rights afforded by Law 31-08 and Law 09-08.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by any of the authors.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Author Contribution Statement

Abdelhakim Adnani: Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization.

References

- [1] European Parliament. (2020). *Artificial intelligence: How does it work, why does it matter, and what can we do about it?* Retrieved from: [https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU\(2020\)641547](https://www.europarl.europa.eu/thinktank/en/document/EPRS_STU(2020)641547)
- [2] Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., ..., & Williams, M. D. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- [3] Mantelero, A. (2024). The Fundamental Rights Impact Assessment (FRIA) in the AI Act: Roots, legal obligations and key elements for a model template. *Computer Law & Security Review*, 54, 106020. <https://doi.org/10.1016/j.clsr.2024.106020>
- [4] Surden, H. (2020). Ethics of AI in law: Basic questions. In M. D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI* (pp. 719–736). Oxford University Press.
- [5] Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- [6] Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ..., & Vayena, E. (2021). An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. In L. Floridi (Ed.), *Ethics, governance, and policies in artificial intelligence* (pp. 19–39). Springer. https://doi.org/10.1007/978-3-030-81907-1_3
- [7] Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU Artificial Intelligence Act - analysing the good, the bad, and the unclear elements of the proposed approach.

- Computer Law Review International*, 22(4), 97–112. <https://doi.org/10.9785/cr-2021-220402>
- [8] Francesconi, E., & Governatori, G. (2023). Patterns for legal compliance checking in a decidable framework of linked open data. *Artificial Intelligence and Law*, 31, 445–464. <https://doi.org/10.1007/s10506-022-09317-8>
- [9] Ashley, K. D. (2017). *Artificial intelligence and legal analytics: new tools for law practice in the digital age*. UK: Cambridge University Press. <https://doi.org/10.1017/9781316761380>
- [10] Governatori, G., Rotolo, A., & Sartor, G. (2021). Logic and the law: Philosophical foundations, deontics, and defeasible reasoning. In D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden & L. van der Torre (Eds.), *Handbook of deontic logic and normative reasoning* (pp. 655–760). UK: College Publications.
- [11] Dell'Anna, D., Alechina, N., Dalpiaz, F., Dastani, M., & Logan, B. (2022). Data-driven revision of conditional norms in multi-agent systems. *Journal of Artificial Intelligence Research*, 75, 1549–1593. <https://doi.org/10.1613/jair.1.13683>
- [12] EUR-Lex. (2024). *Regulation (EU) 2024/1689 of the European Parliament and if the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)*. Official Journal of the European Union, L 1689.
- [13] OECD.AI. (2024). *OECD AI Principles overview*. <https://oecd.ai/en/ai-principles>
- [14] Mendoza-Caminade, A. (2023). *L'entreprise et l'intelligence artificielle: Les réponses du droit* [Business and artificial intelligence: Legal responses]. France: Presses de l'Université Toulouse Capitole.
- [15] Roder, S. (2024). *Guide pratique de l'intelligence artificielle dans l'entreprise (2e éd.)* [A practical guide to artificial intelligence in business (2nd ed.)]. France: Éditions Eyrolles.
- [16] El Badawi, L. (2025). *L'intelligence artificielle et ses applications* [Artificial intelligence and its applications]. France: Éditions Ellipses.
- [17] Wachter, S., Mittelstadt, B., & Russell, C. (2021). Why fairness cannot be automated: bridging the gap between EU non-discrimination law and AI. *Computer Law & Security Review*, 41, 105567. <https://doi.org/10.1016/j.clsr.2021.105567>
- [18] Antoniou, G., Billington, D., Governatori, G., & Maher, M. J. (2001). Representation results for defeasible logic. *ACM Transactions on Computational Logic*, 2(2), 255–287. <https://doi.org/10.1145/371316.371517>
- [19] European Law Institute. (2025). *Guiding principles and model rules on digital assistants for consumer contracts*. Retrieved from: <https://www.europeanlawinstitute.eu/projects-instruments/instruments/eli-guiding-principles-and-model-rules-on-digital-assistants-for-consumer-contracts/>
- [20] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 1–42. <https://doi.org/10.1145/3236009>
- [21] Kaminski, M. E., & Malgieri, G. (2021). Algorithmic impact assessments under the GDPR: Producing multi-layered governance. *International Data Privacy Law*, 11(2), 125–142. <https://doi.org/10.1093/idpl/ipaa020>
- [22] Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online manipulation: Hidden influences in a digital world. *Georgetown Law Technology Review*, 4(1), 1–45. <https://georgetownlawtechreview.org/online-manipulation-hidden-influences-in-a-digital-world/GLTR-01-2020/>
- [23] Wiśniewska, K., & Pałka, P. (2023). The impact of the Digital Content Directive on online platforms' Terms of Service. *Yearbook of European Law*, 42, 388–406. <https://doi.org/10.1093/yel/yead004>
- [24] De Gregorio, G. (2021). The rise of digital constitutionalism in the European Union. *International Journal of Constitutional Law*, 19(1), 41–70. <https://doi.org/10.1093/icon/moab001>
- [25] Crawford, K. (2022). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. USA: Yale University Press.
- [26] Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492. <https://doi.org/10.1257/jel.54.2.442>
- [27] Sarah, C. (2022). *Le contrat de consommation et l'intelligence artificielle (IA)* [Consumer contracts and artificial intelligence (AI)]. In A. Mendoza-Caminade (Ed.), *L'entreprise et l'intelligence artificielle - Les réponses du droit* [Business and Artificial Intelligence - Legal Responses] (pp. 349–361). France: Presses de l'Université Toulouse Capitole. <https://doi.org/10.4000/books.putc.15447>
- [28] Peel, E. (2020). *Treitel: the law of contract (15th ed.)*. UK: Sweet & Maxwell.
- [29] McConville, M., & Chui, W. H. (2021). *Research methods for law (3rd ed.)*. UK: Edinburgh University Press
- [30] Governatori, G., Romeu, P. C., & de Koker, L. (2020). On the formal representation of the Australian spent conviction scheme. In V. Gutiérrez-Basulto, T. Kliegr, A. Soylu, M. Giese & D. Roman (Eds.), *Rules and reasoning* (pp. 177–185). Germany: Springer. https://doi.org/10.1007/978-3-030-57977-7_14
- [31] Boella, G., van der Torre, L., & Verhagen, H. (2006). Introduction to normative multiagent systems. *Computational & Mathematical Organization Theory*, 12, 71–79. <https://doi.org/10.1007/s10588-006-9537-7>
- [32] Micklitz, H. W., Pałka, P., & Panagis, Y. (2017). The empire strikes back: Digital control of unfair terms of online services. *Journal of Consumer Policy*, 40, 367–388. <https://doi.org/10.1007/s10603-017-9353-0>
- [33] McNulty, D., Miglionico, A., & Milne, A. (2023). Data access technologies and the 'New Governance' techniques of financial regulation. *Journal of Financial Regulation*, 9(2), 225–248. <https://doi.org/10.1093/jfr/fjad008>

How to Cite: Adnani, A. (2026). Reassessing Consumer Contract Formation and Execution in Autonomous and AI-Driven Contexts. *Journal of Computational Law and Legal Technology*. <https://doi.org/10.47852/bonviewJCLLT62029386>