

QTRUSTBID: POST-QUANTUM CRYPTOGRAPHY AND AI-DRIVEN RECOMMENDATIONS FOR REAL-TIME PROPERTY AUCTION PLATFORM

¹Eima Nasir, ^{*2}Dr. Mahawish Fatima, ³Dr. Muhammad Ashraf, ¹Muhammad Ali, ¹Jawaria Hafeez,
⁴Dr. Muhammad Hassan Nasir, ⁵Shanila Azhar

¹Student, Department of Software Engineering Bahria University Karachi Campus, Pakistan

^{*2}Assistant Professor, Department of Software Engineering Bahria University Karachi Campus,
Pakistan

³Associate Professor, Department of Computer Engineering, BUITEMS, Quetta, Pakistan

⁴Department of Computer Science & IT, NED University of Engineering & Technology, Pakistan

⁵Lecturer, Department of Computer Engineering, BUITEMS, Quetta, Pakistan

¹eimanasir12@gmail.com ^{*2}mahwishfatima.bukc@bahria.edu.pk

³Muhammad.ashraf@buitms.edu.pk

¹Muhammad_Ali23@gmail.com ¹jawariahafeez86@gmail.com ⁴4mhassan.cse@gmail.com

⁵shanila.azhar@buitms.edu.pk

DOI: <https://doi.org/10.5281/zenodo.20245869>

Article History

Received: 17 April 2026

Accepted: 08 May 2026

Published: 22 May 2026

Copyright @Author Corresponding

Author: *

mahwishfatima.bukc@bahria.edu.pk

Abstract

Pakistan's real estate sector, valued at approximately USD 32 billion, is mainly dependent on the vague, broker-controlled transaction processes that have proven to have structural inefficiency and persistent fraud threats. Existing digital platforms, including Zameen.com, Bayut.com, and PropertyFinder.pk, function primarily as a passive listing platform. Therefore, there is no established mechanism that secures online bidding, coupled with AI-driven personalisation for property discovery and direct interaction between the buyer and seller. Simultaneously, the RSA and Elliptic Curve Cryptography (ECC) mechanisms protecting most digital platforms face an existential threat from quantum computing: Shor's algorithm renders classical key exchange fundamentally insecure in a post-quantum environment, and the harvest-now-decrypt-later (HNDL) threat is active today. This paper presents QTrustBid, a fully implemented and empirically evaluated secure and intelligent real estate bidding platform addressing these failures through three mutually reinforcing technologies: (1) a hybrid AI recommendation engine combining TF-IDF content vectors, collaborative filtering, and recency-decayed behavioural signals; (2) a real-time sealed-bid auction with an APScheduler-driven automated lifecycle; and (3) a post-quantum cryptography layer implementing ML-KEM (Kyber768), standardised as NIST FIPS 203 in August 2024, combined with HKDF-SHA256 key derivation. The system is built

on a six-service microservice architecture with a React frontend and a FAISS-indexed Retrieval-Augmented Generation chatbot. Empirical evaluation across 80 users, 300 listings, and 1,021 behavioural signals demonstrates strong performance. The hybrid engine achieves Precision@1 of 86.25%, Precision@5 of 83.75%, and nDCG@20 of 93.48%, exceeding published real estate recommendation baselines of 72–81% P@5. Cold-start users achieve 78.1% relevance — only 6.1 percentage points below warm users. The post-quantum bidding subsystem achieves 100% tamper detection across 135 adversarial trials with a per-bid encryption mean of 0.067 ms. ML-KEM key generation completes in 48 μ s 2,500 times faster than RSA-2048 — while providing 165 bits of post-quantum security, 37 bits above the NIST minimum threshold.

Keywords: Post-Quantum Cryptography, Recommender Systems, Collaborative Filtering, Real-Time Bidding

1. Introduction

One of the most economically significant sectors includes the real estate market in any emerging economy. The property sector in Pakistan, valued at nearly USD 32 billion and a CAGR record of 4.2% between the years 2011 and 2019 [14], continues to depend on transaction processes structurally unchanged from decades past: broker-mediated verbal negotiation, physical stamp paper, and partially-digitised government land records. In spite of digital property marketplace portals such as Zameen.com, Bayut.com, and PropertyFinder.pk that generate millions of listings, the buyer must move towards offline transactional process after discovering a property online. There is no existence of a Pakistani platform that enables competitive bidding along with personalised property discovery and post-quantum cryptography layer of protection for transaction data [1].

This disparity between online discovery and conventional transaction mechanisms is not only an inconvenience but also a mechanism that is the cause of persistence of vague, broker-driven reliance and fraud-prone transactional systems. The absence of verifiable bids, personalised recommendations, and cryptography auditable records leaves the most economically significant transaction steps in exactly

the informal domain where manipulation is most easily perpetrated.

Simultaneously, the cryptography mechanisms securing digital platforms — RSA-2048 and Elliptic Curve Cryptography (ECC) — face an existential long-term threat from quantum computing. Shor's algorithm, executable on a sufficiently powerful quantum computer, solves both integer factorisation (RSA) and the discrete logarithm problem (ECC) in polynomial time [15]. Credible institutional estimates suggest cryptography relevant quantum computers may emerge within a decade, and the active HNDL threat — archiving encrypted data today for retrospective decryption — means the transition to post-quantum cryptography is urgent [2][3].

This paper makes three principal contributions: (1) a hybrid AI recommendation engine tailored to sparse-interaction property markets, combining TF-IDF content vectorisation, user-similarity collaborative filtering [4], and time-weighted behavioural signals; (2) a real-time sealed-bid property auction mechanism with automated lifecycle management and wallet-integrated deposit handling [5]; and (3) the first reported deployment of ML-KEM (NIST FIPS 203, Kyber768) [2] within a live real-time sealed-bid auction context, providing post-quantum confidentiality and cryptography auditable bid integrity. The system is

built on a six-service microservice architecture with a React frontend, FAISS-indexed [6] RAG chatbot [7][12], and explicit alignment with Pakistan's property transaction legal framework.

2. Background and Related Work

2.1 Digital Real Estate Platforms

Existing property platforms are structurally constrained to the discovery stage. Zameen.com (~4M monthly users [14]), Bayut.com, PropertyFinder.pk, Zillow, and Rightmove share a critical limitation: they terminate at listing aggregation, with all transactional steps executed offline through intermediaries. None of the existing systems implement machine-learning personalisation beyond rule-based filters. Moreover, none of them implements a competitive bidding system that incorporates post-quantum cryptography security layer. However, some niche platforms have initiated attempts to integrate transactional mechanism into their system such as Opendoor's iBuying model and Purplebricks' fixed-fee agency but it has proven to require a carefully designed trust and verification layer instead of the elimination of offline legal procedures entirely [1].

2.2 Recommender Systems

Recommender systems literature comprises primarily of three paradigms [4]. Content-based filtering works by analyzing attributes of items that match to user profiles based on the interaction history. Collaborative filtering functions by identifying cross-user patterns that are further divided into memory-based and model-based approaches. Hybrid systems leverage both by using content features for cold-start conditions and collaborative signals when it comes to users with richer histories [4]. The field has been transformed significantly due to deep learning: Gao et al. [9] review neural architectures, highlighting state-of-the-art results in sequential behavior modeling. On the other hand, Wu et al. [13] focus on graph neural network approaches that capture higher-order collaborative signals. Real estate recommendation mechanism has presented challenges — infrequent transactions, geographic limitations, high price sensitivity — that lean towards hybrid content-collaborative approaches [8].

2.3 Real-Time Bidding Systems

Three essential properties are required for trustworthy bidding systems: bid integrity, bid confidentiality and audit transparency [11] Anderson [11] recognizes that sealed-bid protocols have proven to directly address manipulation concerns. Farajtabar et al. [5] survey manipulation resistance and cryptography protection for high-value bidding systems. Therefore, demonstrating seal-bid architectures with cryptography audit trails as the correct design response. All three properties aren't implemented by any existing Pakistani property platform.

2.4 Post-Quantum Cryptography

Focus of post-quantum cryptography (PQC) algorithms is to design algorithms that remain secure against both, classical and quantum adversaries. Following an eight-year standardisation process by NIST, ML-KEM (based on CRYSTALS-Kyber) was standardised as FIPS 203 in August 2024 [2]. Its security is based on the Module Learning with Errors (MLWE) problem, for which there is no known efficient quantum attack [2] [10]. Bos et al. [10] report that Kyber768 achieves key generation in approximately 48 μ s, encapsulation in 62 μ s, and decapsulation in 64 μ s on x86-64 AVX2 hardware. Therefore, outperforming ECDH in runtime while offering qualitatively stronger protection. Aquina et al. [3] argue that NIST-standardised PQC is the most practical approach currently for the protection of large-scale digital infrastructure, and highlight that NHDL threats make PQC deployment a critical requirement rather than a precautionary measure.

2.5 Research Gap

There are no existing systems that implement the integration of AI-driven property recommendation, real-time sealed-bid auctions, along with post-quantum security in a single platform that is designed for a geographic specific market with its particular legal framework. There have been extensive studies for each component individually [2][4][5]. QTrustBid contributes by the principled integration of these components into a deployable system architecture.

3. System Architecture

QTrustBid’s system architecture comprises of six independently deployable FastAPI microservices: Authentication (port 8000), User Management (8001), Property Management (8002), Administration (8003), Recommendation (8004), and Bidding/Quantum(8005). The communication between the microservices is made possible through REST APIs over a shared PostgreSQL 15 database,

ensuring ACID consistency. Unified user interface, with WebSocket connections to allow real-time auction updates, is provided by a React 18 single-page application. A FAISS-indexed [6] retrieval augmented generation (RAG) chatbot [7][12] is also integrated to support natural language queries over the verified property listing corpus.




Figure 1: QTrustBid System Architecture

Table 1: QTrustBid System Architecture Overview

Component	Technology	Role
Auth Service	FastAPI + JWT (HS256)	CNIC-linked identity; 30-min token expiry; refresh rotation
User Service	FastAPI + SQLAlchemy	Profile management; preference vector storage
Property Service	FastAPI + Cloudinary	Listing CRUD; image CDN; admin verification workflow

Admin Service	FastAPI	Listing review queue; RERA-aligned verification
Recommendation Service	FastAPI + scikit-learn	Hybrid TF-IDF + CF + behavioural scoring; RAG chatbot
Bidding Service	FastAPI + liboqs + APScheduler	ML-KEM encryption; sealed-bid lifecycle; wallet management
Frontend	React 18 + TailwindCSS + WebSocket	SPA with live auction updates and buyer/seller dashboards
Database	PostgreSQL 15	13 tables; full ACID; indexed on slot status and city

4. Hybrid Recommendation Engine

4.1 Property Vectorisation

A 2,007-dimensional combined property vector is built by the Property Vectoriser through concatenating TF-IDF-based content features along with scaled numeric attributes. The TF-IDF component [16] makes use of 2,000 features derived from unigram and bigram representations, where city and property type tokens are repeated three times to increase the primary geographic and type signals. In addition, seven numeric features — price bucket index, number of bedrooms, number of bathrooms, property size in Marla, and one-hot encoded property type — are scaled using MinMax normalization to [0, 1] and concatenated. Through this design, geographic preference is ensured, which is a dominant constraint for Pakistani property buyers, while still allowing numerical attributes to refine distinctions within localized geographic clusters [1].

4.2 User Embedding Construction

A user's embedding is constructed as a weighted average of property vectors for all interacted properties, weighted by interaction type and recency. Interaction weights reflect engagement intensity: favourite (5), interested (4), bid_registered (3), browsed (1). The recency decay function is:

$$w(t) = \exp(-\ln(2) \times \text{age_days} / \mathbf{T}), \text{ where } \mathbf{T} = 30 \text{ days} \quad [4]$$

This yields a weight of 1.0 at interaction time, 0.5 at 30 days, and ~0.25 at 60 days. The interaction-weighted embedding is blended 70/30 with an explicit preference vector derived from user-stated city, area, property type, and price range preferences — ensuring stated preferences remain influential even as interaction history grows [4].

4.3 Hybrid Scoring

For each candidate property, the hybrid scorer computes three component scores and combines them as:

$$\text{final_score} = (0.45 \times \text{content} + 0.30 \times \text{collaborative} + 0.25 \times \text{behavioural}) \times \text{recency_boost}$$

The content score is the cosine similarity between the user embedding and the candidate property vector. The collaborative score is the normalised weighted sum of interaction signals from users with similar preference vectors [13]. The behavioural score is the normalised recency-decayed personal interaction count per property. The recency_boost term provides up to a 15% uplift for listings posted within the preceding seven days. The weight configuration (0.45/0.30/0.25) was selected via sensitivity analysis to occupy a high-performance region robust to distribution shift as platform interaction density increases.

4.4 Cold-Start Strategy

New users with no interaction history are served a trending-based fallback:

$$\text{trending_score} = 0.35 \times \text{favourites} + 0.25 \times \text{interests} + 0.25 \times \text{bid_registrations} + 0.15 \times \text{active_slot}$$

This aggregates platform-wide engagement signals to surface properties with demonstrated broad appeal. A preference profile collected at registration (city, property type, price range) seeds the explicit preference vector, allowing first-session recommendations to be filtered to the user's geographic market — addressing the cold-start challenge documented extensively in the recommender systems literature [4][9].

5. Post-Quantum Cryptography Bidding System

5.1 Cryptography Architecture

A three-layer pipeline is followed by the “BidEncryptionSystem” at the point of bid submission. In the first layer, ML-KEM encapsulation, using the system's Kyber768 public key, produces a 1,088-byte

KEM ciphertext and a 32-byte shared secret. A fresh encapsulation is triggered by each bid and therefore, produces a unique shared secret and providing per-bid semantic security [2]. The second layer applies HKDF-SHA256 (RFC 5869 [17]) along with the info string “columnar-cipher-key” to derive an 8-character alphabetic columnar key from the shared secret. In layer three, the bid payload (property ID, bid amount, and user ID) is encrypted using a Columnar Transposition Cipher with the previously derived key. Both, the encrypted bid amount and the KEM ciphertext are stored along with the plaintext amount for post-auction dual verification.

5.2 Auction Lifecycle and Bid Integrity

The full auction slot lifecycle is managed by APScheduler through polling every 60 seconds, progressing slots through upcoming, live, ended, and result-processing states [5]. When a slot closes, the system decrypts all valid bids through ML-KEM decapsulation and the corresponding derived columnar key, identifies the highest valid bid, records the bid_result, and executes wallet transactions accordingly. Therefore, returning surplus deposits to unsuccessful bidders and securing the winner's deposit. Bid integrity maintained through dual verification: the decrypted amount must be the same as the stored plaintext amount with a tolerance of 0.01. Any mismatch is considered as a tampering event, resulting in bid invalidation. To prevent duplicate execution, the process is idempotent via a NOT EXISTS subquery on bid_results, ensuring safe recovery after scheduler starts.

5.3 Security Analysis

ML-KEM (Kyber768) security is based on the MLWE problem, for which no practical classical or quantum attack is currently known [2][10]. Kyber768 is estimated to provide about 180 bits of classical security

and 165 bits of post-quantum security, which is 37 bits above the NIST minimum 128-bit threshold [2][10]. Defense in depth is implemented through this margin against algorithmic improvements against lattice problems (bounded within ~20 bits by the CRYSTALS-Kyber design team's analysis [10]) and improvements in BKZ lattice reduction algorithms [10]. In context of QTrustBid, the harvest-now-decrypt-later (NHDL) threat is commercially significant as real estate bidding data — including the maximum prices buyers were prepared to pay for specific properties — can retain commercial sensitivity years after the transactions are processed [2][3].

6. Empirical Evaluation

6.1 Experimental Setup

The recommendation evaluation was performed using 300 synthetic property listings from Karachi, Lahore, and Islamabad. It covered three property categories (house, plot, and commercial). The user base consisted of 80 synthetic users, each assigned a primary city and one or two property-type


preferences. 1,021 behavioural interactions were generated in total using an exponential decay distribution with a mean age of 45 days. Five users were assigned as cold-start cases with no prior interaction history. Relevance was defined using a conservative binary rule, where a recommendation is considered relevant if it aligns with either the user's preferred city or property type.

For the post-quantum cryptography evaluation, 135 test cases were executed. This was carried out across three attack classes (45 each): ciphertext bit-flip attacks that target the encrypted amount field, amount substitution attacks modifying stored plaintext_amount values, and KEM ciphertext corruption via replacement of the base64encoded ciphertext. In addition, the false positive rate was estimated using 165 valid bid executions. ML-KEM performance benchmarks are taken from published loboqs eBACS results on x84-64 AVX2 hardware, using the same implementation as the deployed system [10].

6.2 Recommendation Engine Results

Table 2: Recommendation Engine Retrieval Quality (80 users, 300 listings, 1,021 signals)

K	Precision@K	Recall@K	nDCG@K	vs. Literature Baseline [8][9]
1	0.8625	0.0862	0.8625	Above (lit: 0.72–0.81 P@5)
3	0.8250	0.2475	0.8825	Above
5	0.8375	0.4188	0.9021	Above (target: $P@5 \geq 0.80$)
10	0.8313	0.8313	0.9178	Above
15	0.8292	0.9958	0.9247	Above
20	0.8363	1.0000	0.9348	Above (target: $nDCG \geq 0.90$)



Precision@K remains largely stable across $K=1$ to $K=20$, ranging from 0.8250 to 0.8625. In a property discovery setting, this plateau effect is desirable. This is because it suggests that the hybrid scoring mechanism is capable of preserving relevance throughout the ranked list instead of concentrating it only at the top ranks. This behavior is caused due to the relatively high content weighting (45%), where TF-IDF effectively captures the primary user preference signal—geographic location—with strong fidelity [1].

nDCG@K [18] increases from 0.8625 at $K = 1$ to 0.9348 at $K = 20$, denoting that the top-20 ranked results achieve 93.48% of the ideal ranking quality. The Precision@1 score of 86.25% is profound in an applied setting: in real estate, where each decision may be coupled with significant financial commitment, presenting a relevant first result directly influences user trust and conversion behavior.

In comparison to prior work, where real estate recommender systems usually report $P@5$ values in the

range of 0.72-0.84 [4][8], QTrustBid has been capable of achieving a $P@5$ of 0.8375, placing it at the upper end of reported performance under a strict binary relevance assumption.

6.3 Cold-Start Analysis

Cold-start users achieve a relevance rate of 78.1%, in comparison to 84.2% for warm users resulting in a 6.1 percentage-point gap. This gap is considerably smaller than the 15–25 point cold-start penalties typical of collaborative filtering-heavy systems reported in the literature [9]. The reduced disparity largely due to the higher content weighting (45%) and use of the trending fallback strategy, which together ensure new users receive geographically appropriate recommendations based on their stated preferences. A 78.1% first-session relevance rate substantially outperforms the ~45–50% expected from a random ranker for a well-specified preference profile [4], thereby reducing the risk of first-session churn.



Figure 3: Cold-Start vs. Warm User Recommendation Analysis

6.4 Weight Sensitivity Analysis

Evaluation against 16 weight configurations within the feasible space ($W_c + W_{co} + W_b = 1.0$) shows that $nDCG@5$ increases consistently as W_c increases, across all values of W_b . This shows that content features are the primary driver of recommendation quality in sparse-interaction property

markets. The selected production configuration (0.45/0.30/0.25) lies within a high-performance region of the parameter space. It maintains appropriate distance from sensitivity boundaries. This placement ensures stability under changing interaction densities, avoiding configurations at the extremes that are more susceptible to performance degradation under distribution shift [4].

6.5 Post-Quantum Cryptography Results

Table 3: Post-Quantum Cryptography Bidding System Evaluation Results

Metric	Result	Comparison / Target
Tamper detection — ciphertext bit-flip (45 trials)	100%	Target: 100%
Tamper detection — amount substitution (45 trials)	100%	Target: 100%
Tamper detection — KEM ciphertext corruption (45 trials)	100%	Target: 100%
False positive rate (165 valid bids)	0%	Target: 0%
Per-bid encrypt latency (mean, 200 trials)	0.067 ms	Target: < 100 ms

Per-bid decrypt latency (mean)	0.038 ms	
ML-KEM KeyGen (Kyber768)	48 μ s	RSA-2048: 120,000 μ s (2,500 $\times$ faster) [10]
ML-KEM Encapsulation (Kyber768)	62 μ s	ECDH-P256: 185 μ s (3 $\times$ faster) [10]
Post-quantum security level	$\sim$ 165 bits	NIST minimum: 128 bits (+37 bits margin) [2]
Classical security level	$\sim$ 180 bits	NIST Level 3 [2]

The correctness of the dual-verification implementation is confirmed by the system’s 100% tamper detection rate across all three attack categories. From a security perspective, the KEM ciphertext corruption scenario is particularly significant as it models a complex adversary with database write access attempting to replace a ciphertext with one tied to a known shared secret. Detection is ensured by ML-KEM’s binding property: any ciphertext substitution results in either decapsulation failure of the generation

of a different shared secret, both of which signify the trigger of verification failure [2][10]. 0.067 ms of processing overhead is introduced from per-bid cryptography pipeline—less than 0.1% of a standard HTTP bid submission request—making post-quantum integration effectively transparent to end users. In addition, Kyber768’s about 2,500 times faster key generation, in comparison to RSA-2048, removes the traditional performance-based objection to PQC adoption [10].






Figure 5: ~~Quantum Security Levels~~

7. Comparison with Existing Platforms

Table 4: Feature Comparison — QTrustBid vs. Existing Platforms

Feature	Zameen.com	Bayut.com	Zillow	eBay	QTrustBid
AI Recommendations	✗	✗	Rule-based	✗	✓ Hybrid ML
Online Bidding	✗	✗	✗	✓	✓ Sealed-bid
Post-Quantum Crypto	✗	✗	✗	✗	✓ FIPS 203 [2]
Bid Confidentiality	N/A	N/A	N/A	✗	✓ Sealed bids
Tamper-Evident Audit	✗	✗	✗	Partial	✓ 100% [10]
Natural Language Search	✗	✗	✗	✗	✓ RAG Chatbot [6][7][12]
Listing Verification	Basic	Basic	Agent-linked	✗	✓ Admin review
Direct Buyer-Seller	Via broker	Via	Via	✓	✓ Platform-mediated

		broker	agent		
--	--	--------	-------	--	--

Table 5: Cryptography Scheme Comparison — Security and Performance [2][10]

Algorithm	Classical Security	PQ Security	KeyGen (μs)	Encap (μs)	NIST Status
ML-KEM Kyber768	~180-bit	~165-bit	48	62	FIPS 203 ✓
ML-KEM Kyber512	~128-bit	~118-bit	27	36	FIPS 203 ✓
ML-KEM Kyber1024	~256-bit	~230-bit	80	100	FIPS 203 ✓
ECDH P-256	128-bit	0-bit (broken)	98	185	Deprecated (PQ)
ECDH P-384	192-bit	0-bit (broken)	210	350	Deprecated (PQ)
RSA-2048	112-bit	0-bit (broken)	120,000	2,800	Deprecated (PQ)
RSA-4096	140-bit	0-bit (broken)	900,000	4,100	Deprecated (PQ)
X25519	128-bit	0-bit (broken)	62	95	Deprecated (PQ)

Table 6: Recommendation System Comparison with Literature

System / Study	Approach	Domain	Cold-Start Strategy	P@5
Amazon	CF + content [1]	E-commerce	Bestseller fallback	~0.78 (est.)

Netflix Cinematch	Matrix factorisation [9]	Streaming	Popularity + genre	~0.76
Bobadilla et al. [4]	Survey of CF methods	General	Various	0.72–0.84
Zhang et al. [8]	Deep learning rec.	General	Embedding-based	0.81–0.93
Gao et al. [9]	Neural recommendation	General	Meta-learning	0.79–0.92
QTrustBid (this work)	Hybrid TF-IDF + CF + behavioural	Real estate (PK)	Trending score fallback	0.8375

8. Pakistan Legal Framework Alignment

QTrustBid's design is structured around the existing property transaction legal framework in Pakistan, digitizing and securing processes that currently rely on informal trust while preserving the statutory legal procedures required by law. The digital

bid records are treated by the platform as pre-contractual evidence and verified identity documentation rather than substitutes for mandatory legal steps such as stamp paper execution, Sub-Registrar registration, or mutation (Intiqal) within provincial land record systems.

Table 7: Pakistan Property Law Alignment Matrix

Legal Instrument	Relevant Provision	QTrustBid Implementation
Transfer of Property Act 1882 [22]	Requires formal written agreement for immovable property transfers	Verified digital bid record serves as pre-contractual evidence; does not replace formal conveyance
Registration Act 1908 [23]	Mandates Sub-Registrar office registration of transfers	Winning bid record provides basis for subsequent offline Bayana; registration step preserved
Stamp Act 1899 [24]	Requires stamp duty on transfer instruments	Platform provides verified price basis for stamp duty calculation; does not generate

		stamp instruments
Land Revenue Act 1967 [25]	Governs mutation (Intiqal) process for updating provincial land records	CNIC-linked user accounts provide verifiable buyer identity documentation for mutation
RERA (Emerging) [26]	Listing registration and agent certification requirements	Administrative verification workflow structurally aligned with RERA listing registration mandate

9. Limitations and Future Work

9.1 Acknowledged Limitations

The evaluation of the recommendation engine was conducted using synthetic interaction data that was specifically designed to reflect realistic engagement funnel proportions. However, synthetic datasets are not capable of fully capturing the temporal clustering, geographic correlation, or long-tail behavioural distributions found in real user interactions, making live production validation through A/B testing against click-through and bid conversion metrics essential [9]. ML-KEM performance benchmarks were sourced from loboqs eBACS testing on x86-64 AVX2 hardware. ARM-based servers (AWS Graviton, Apple M-series) have the capability to yield 20-40% higher latency due to AVX2 vs. NEON SIMD differences, although sub-millisecond performance remains achievable [10]. The Columnar Transposition Cipher is strengthened by a unique HKDF-derived key for each bid but it does not satisfy formal semantic security requirements independently. For production deployment, stronger authentication methods like AES-256-GCM or ChaCha20-Poly1305 should replace it at the cipher layer [21]. Moreover, formal legal assessment by qualified Pakistani property law practitioners would be necessary before the commercial implementation.

9.2 Future Work

There are four main directions for future improvement emerge from this work. First of all, TF-IDF could be replaced with fine-tuned multilingual sentence-BERT embeddings specifically trained on Pakistani property descriptions in both English and Urdu transliteration, allowing stronger semantic understanding beyond simple lexical matching [8][19]. Secondly, brute-force pairwise collaborative filtering could be upgraded to a Hierarchical Navigable Small World (HNSW) approximate nearest-neighbour index over user embedding vectors using FAISS [6]. Therefore, significantly improving scalability through the reduction of recommendation complexity from $O(n^2)$ to $O(n \log n)$. Thirdly, integrating ML-DSA (CRYSTALS-Dilithium, NIST FIPS 204) digital signatures into bid submission would introduce quantum resistant non-repudiation. This leads to strengthening legal and transactional assurance under Pakistan's Electronic Transactions Ordinance 2002 [2][20]. Finally, bid result hashes stored on a consortium blockchain could establish immutable audit trails, reducing sole reliance on platform operator trust while enhancing transparency [5].

10. Conclusion

This paper is a representation of a fully implemented and empirically evaluated secure and intelligent real estate bidding platform designed specifically with Pakistan's USD 32 billion property market in mind [14]. The system provides solutions for the limitations of existing systems: broker dependence, no existing competitive bidding mechanisms, limited AI personalization for recommendations, and cryptography vulnerability to quantum computing. The solutions to these problems are approached by the structured integration of machine learning, real-time bidding mechanism and post-quantum cryptography through a six-service microservice architecture.

The recommendation engine has proven to achieve Precision@1 of 86.25%, Precision@5 of 83.75% nDCG@20 of 93.48%. These evaluation results exceed published real-estate recommendation baselines of 72-81% P@5 [4][8] under a conservative

binary relevance criterion. Moreover, Cold-start users are able to achieve 78.1% first-session relevance. This means only 6.1% below warm users and considerably above the 15-25 point penalties typical in collaborative filtering-heavy systems [9]. The post-quantum cryptography module achieves 100% tamper detection with respect to 135 adversarial trials with zero false positives, with per-bid encryption implemented in a mean time of 0.067 ms. This is lesser than 0.1% overhead on a typical bid submission request. ML-KEM Kyber768 key generation has proven to work 2,500 times faster as compared to RSA-2048, while also providing 165 bits of post-quantum security [2][10]. The combination of sealed-bid confidentiality (using ML-KEM) along with post-hoc auditable integrity (using dual verification) provides features that neither classical cryptography nor simple sealed-bid designs can offer at the same time. This is a combination that is not previously documented in the academic literature for a real-time property bidding context.

References

- [1] Ricci, F., Rokach, L., & Shapira, B. (Eds.). (2015). Recommender systems handbook (2nd ed.). Springer. <https://doi.org/10.1007/978-1-4899-7637-6>
- [2] National Institute of Standards and Technology. (2024). Module-lattice-based key-encapsulation mechanism standard (Federal Information Processing Standard 203). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.FIPS.203>
- [3] Aquina, N., Kaur, H., Bhattacharya, A., & Chakraborty, C. (2025). A critical analysis of deployed use cases for quantum key distribution and comparison with post-quantum cryptography. EPJ Quantum Technology, 12. <https://doi.org/10.1140/epjqt/s40507-024-00299-5>
- [4] Bobadilla, J., Ortega, F., Hernando, A., & Gutierrez, A. (2013). Recommender systems survey. Knowledge-Based Systems, 46, 109–132. <https://doi.org/10.1016/j.knosys.2013.03.012>
- [5] Farajtabar, M., Wang, Y., Coates, M., & Zha, H. (2024). Towards trustworthy AI-empowered real-time bidding for online advertisement auctions. ACM Computing Surveys. <https://doi.org/10.1145/3701741>
- [6] Douze, M., Guzhva, A., Deng, C., Johnson, J., Szilvasy, G., Mazaré, P.-E., Lomeli, M., Hosseini, L., & Jégou, H. (2024). The FAISS

- library. arXiv preprint.
<https://arxiv.org/abs/2401.08281>
- [7] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474. <https://arxiv.org/abs/2005.11401>
- [8] Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), Article 5. <https://doi.org/10.1145/3285029>
- [9] Gao, C., Lei, W., He, X., de Rijke, M., & Chua, T.-S. (2023). A comprehensive survey of deep learning based recommender systems. *Applied Sciences*, 13(20), 11378. <https://doi.org/10.3390/app132011378>
- [10] Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J. M., Schwabe, P., Seiler, G., & Stehlé, D. (2018). CRYSTALS-Kyber: A CCA-secure module-lattice-based KEM. 2018 IEEE European Symposium on Security and Privacy (EuroS&P), 353–367. <https://doi.org/10.1109/EuroSP.2018.00032>
- [11] Anderson, R. (2020). *Security engineering: A guide to building dependable distributed systems* (3rd ed.). Wiley.
- [12] Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., & Wang, H. (2024). Retrieval-augmented generation for large language models: A survey. arXiv preprint. <https://arxiv.org/abs/2312.10997>
- [13] Wu, S., Sun, F., Zhang, W., Xie, X., & Cui, B. (2022). Graph neural networks in recommender systems: A survey. *ACM Transactions on Information Systems*, 41(1), Article 21. <https://doi.org/10.1145/3568022>
- [14] State Bank of Pakistan. (2021). Financial stability review 2020–21. State Bank of Pakistan. <https://www.sbp.org.pk/FSR/2021/index.asp>
- [15] Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509. <https://doi.org/10.1137/S0097539795293172>
- [16] Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511809071>
- [17] Krawczyk, H., & Eronen, P. (2010). HMAC-based extract-and-expand key derivation function (HKDF) (RFC 5869). Internet Engineering Task Force. <https://doi.org/10.17487/RFC5869>
- [18] Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, 20(4), 422–446. <https://doi.org/10.1145/582415.582418>
- [19] Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
- [20] National Institute of Standards and Technology. (2024). Module-lattice-based digital signature

- standard (Federal Information Processing Standard 204). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.FIPS.204>
- [21] Dworkin, M. (2007). Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC (NIST Special Publication 800-38D). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-38D>
- [22] Transfer of Property Act 1882 (Act IV of 1882). Government of Pakistan. <https://www.pakistancode.gov.pk>
- [23] Registration Act 1908 (Act XVI of 1908). Government of Pakistan. <https://www.pakistancode.gov.pk>
- [24] Stamp Act 1899 (Act II of 1899). Government of Pakistan. <https://www.pakistancode.gov.pk>
- [25] Land Revenue Act 1967 (West Pakistan Act XVII of 1967). Government of Pakistan. <https://www.pakistancode.gov.pk>
- [26] Real Estate Regulatory Authority Act (Draft). Ministry of Housing & Works, Government of Pakistan. <https://www.mhw.gov.pk>

