AI for Real-Estate Deal Screening: From Land Parcel to Investment-Committee Decision

A Governed Evidence, Valuation and Decision Framework for Rapid Property Underwriting

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Real-estate investors often receive more opportunities than their teams can underwrite in depth. Artificial intelligence can extract documents, reconcile parcel identifiers, classify comparables, identify missing evidence, populate models and prepare a first investment-committee pack. The same tools can also propagate a wrong plot boundary, treat an asking price as a completed transaction, overlook a planning condition, conceal stale data behind a confidence score or convert a fragile assumption into a precise valuation. Speed creates value only when evidence quality and decision accountability improve with it.

This paper develops a governed screening framework that moves from land parcel to investment-committee decision. It starts with the decision mandate and minimum evidence threshold, then builds a lineage-controlled parcel record. It separates hard gates from weighted commercial factors, establishes a comparable-evidence confidence ladder, creates a transparent residual land and development model, and presents outcomes as scenarios and ranges. It specifies validation, monitoring, override and escalation controls for AI-supported outputs. Five original figures and five implementation tables provide a data-lineage map, parcel scorecard, comparable-confidence ladder, scenario fan chart and investment-committee decision tree.

The framework draws on current primary and authoritative sources. The International Valuation Standards effective from 31 January 2025 include general standards on data and inputs, valuation models, documentation and reporting, together with asset standards for real property interests and development property.[1] IFRS 13 defines fair value by reference to an orderly transaction between market participants at the measurement date and requires market-participant assumptions under current market conditions.[2] The RICS professional standard on responsible use of artificial intelligence in surveying practice takes effect on 9 March 2026 and addresses governance, professional judgement, transparency, data, output reliability and accountability.[3] NIST's AI Risk Management Framework organises risk management around Govern, Map, Measure and Manage and emphasises validity, reliability, transparency, accountability and lifecycle evaluation.[4][5]

Official transaction and spatial datasets demonstrate both the opportunity and the control problem. HM Land Registry Price Paid Data records residential property transactions lodged for registration and is updated monthly, while the publisher warns that a current-month extract is incomplete in isolation and does not guarantee fitness for a particular use.[6] HM Land Registry INSPIRE Index Polygons provide indicative freehold extents and identifiers rather than a substitute for title investigation.[7] Dubai Land Department publishes transaction and property data through its open-data channels, including a daily transactions dataset sourced from the land-registration system.[8][9] Saudi Arabia's Real Estate General Authority publishes sales and rent indicators and distinguishes recent raw transaction information from refined historical indicators.[10] England's long-term flood-risk service provides area-level risk information and expressly does not determine whether an individual property will flood.[11]

The worked case, values, scores, probabilities, durations, costs, rents, yields, prices and return metrics are hypothetical modelling assumptions used to demonstrate the framework. They are not market evidence, valuations, transaction terms, forecasts, legal conclusions, accounting conclusions, tax advice or investment recommendations. A live screening decision requires current parcel-specific evidence and qualified legal, planning, valuation, technical, environmental, tax, accounting, insurance, financing and investment review in each relevant jurisdiction.

JEL Classification: G11, G24, G32, R30, R31, C45

Keywords: real estate, artificial intelligence, deal screening, land valuation, investment committee, comparable evidence, residual valuation, model risk, data governance, property underwriting

1. Define the decision before automating the work

A screening system needs a precise decision mandate. “Find attractive land” is too broad. A useful mandate defines geography, permitted strategies, ticket size, holding period, return objectives, risk limits, development capability, financing constraints, sustainability requirements and the evidence required for escalation. The system should know whether it is screening a standing asset, development site, conversion opportunity, operating platform, distressed exposure or portfolio. Each has a different unit of analysis and a different valuation path.

The first output is a triage decision rather than a final valuation. A practical decision set is reject, hold for evidence, advance to preliminary underwriting, advance with a stated exception, or refer for specialist review. The system should never manufacture missing facts to reach one of those outcomes. A missing title extract, unclear boundary, absent access right or unverified utility capacity becomes an evidence gap with an owner and deadline. An opportunity can remain attractive while its decision status remains “hold for evidence”.

Materiality should determine depth. A small, standardised acquisition can use a concise screen. A development site whose value depends on rezoning, grid connection, ground remediation and pre-sales requires a deeper review before the first committee discussion. The mandate should identify which variables can be estimated for triage and which require authoritative evidence. This distinction prevents an early model from acquiring the status of a completed diligence conclusion.

The decision owner also needs to be named. A machine can compile, compare and calculate. The investment professional remains responsible for explaining why the opportunity fits the mandate, which evidence is reliable, where uncertainty remains and what could invalidate the recommendation. The RICS AI standard makes professional judgement and oversight central to responsible use in surveying practice.[3] NIST similarly treats accountability, transparency and documented risk management as lifecycle requirements.[4][5]

Table 1. Minimum evidence register for an AI-supported property screen

Evidence domain	Minimum screening record	Acceptable early-stage status	Escalation trigger
parcel identity	authoritative identifier, boundary source, tenure and owner record	verified identifier with boundary caveat	conflicting identifiers, overlaps or unregistered interests
planning and use	current designation, permission status, conditions and relevant policy	official source or qualified adviser note	use depends on change, appeal or untested interpretation
access and infrastructure	legal access, roads, utilities, capacity and connection timing	documentary evidence with open confirmations	no enforceable access, capacity uncertainty or critical delay
environmental and physical	flood, contamination, ground, heritage and other site constraints	official screening plus site evidence	material hazard, remediation dependency or incomplete survey
market evidence	transactions, rents, supply, absorption and capital-market assumptions	dated sources with confidence grade	asking-price dependence, stale evidence or thin market
financial model	unit programme, cost plan, timing, revenue, funding and sensitivities	transparent assumptions with owner	value depends on one unsupported assumption or circular funding
counterparty and process	seller authority, adviser, timetable, exclusivity and bid rules	source documented	unclear authority, compressed process or information asymmetry

The required evidence varies by jurisdiction, asset type and decision materiality.

2. Create one authoritative parcel record

Parcel identity is the foundation of the screen. Real-estate opportunities arrive as brochures, coordinates, plot numbers, title references, planning applications, valuation reports and seller spreadsheets. Names and boundaries can differ across documents. An AI system can extract identifiers rapidly, yet extraction accuracy does not prove that the documents refer to the same legal interest. The system should create one authoritative parcel record and preserve every source identifier as an alias until a qualified reviewer resolves conflicts.

The record should contain the jurisdiction, registry, title or equivalent reference, parcel or plot number, geospatial geometry, tenure, ownership, registered rights, restrictions, easements, charges and the effective dates of the evidence. Where a polygon is indicative, it should remain labelled as indicative in the operational data model. HM Land Registry explains that INSPIRE Index Polygons show indicative freehold extents and provide unique identifiers; they support discovery and matching, while title evidence remains a separate enquiry.[7] Similar discipline applies in any market where public maps, broker pins and registry records have different legal status.

Entity resolution should use deterministic and probabilistic methods. Exact registry identifiers take priority. Normalised addresses, coordinates, developer names and area measurements can support matching. The system should expose its match score and the reasons for it. A reviewer should be able to see that two records were joined because of a shared title reference, or that a weaker match relied on proximity and a similar address. Automatic merging should be prohibited

when two candidates remain plausible.

Version control matters because plots can be subdivided, consolidated, re-numbered or transferred. The record should keep an effective-date history and should not overwrite the geometry used in an earlier decision pack. A reproducible screen preserves the evidence snapshot used at the committee date, together with later updates. This supports post-investment review, challenge and audit.

Figure 1. Governed data lineage from parcel evidence to committee decision

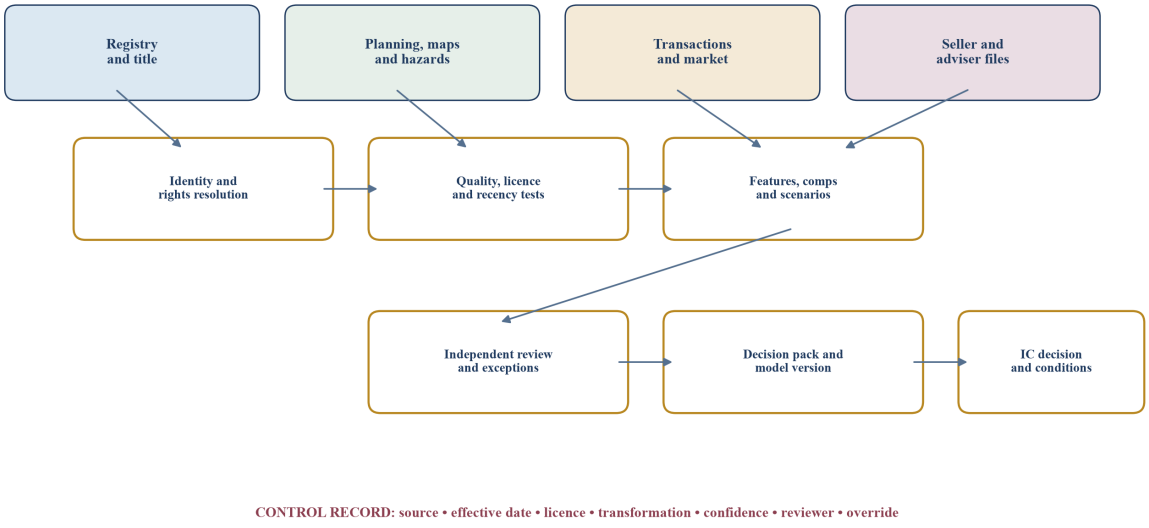

The diagram shows control stages; every transformation retains source, date, licence, confidence and reviewer.

3. Separate hard gates from weighted attractiveness

Screening fails when all variables are allowed to compensate for each other. A high projected return should not offset the absence of legal access. Strong demand should not cure an unresolvable title conflict. The framework therefore separates hard gates from weighted factors. Hard gates determine whether the opportunity can advance under the current evidence. Weighted factors compare opportunities that have passed, or that have a clearly defined route to pass, those gates.

Hard gates commonly cover mandate fit, verified parcel identity, transferable interest, permitted or achievable use, legal access, critical infrastructure, prohibited environmental exposure, counterparty authority and minimum data sufficiency. A gate can be pass, fail, conditional or unknown. Conditional status needs a specific pathway, decision date, owner and cost. Unknown status should not be converted into a neutral score.

Weighted factors then assess market depth, micro-location, product fit, supply, achievable pricing, cost certainty, programme, financing, downside protection, exit liquidity, strategic option value and execution capability. Weightings should reflect the fund or corporate mandate. They should be approved before the team sees a favoured transaction, since changing the weights after seeing the result invites confirmation bias.

Scoring should remain interpretable. A ten-point market score needs defined anchors. The record should show which observations drove the score, their confidence and the reviewer.

Machine-generated narratives can summarise the evidence, while source links and structured values remain available beneath the narrative. A score without evidence is an opinion presented as data.

Figure 2. Illustrative parcel scorecard with hard gates and weighted factors

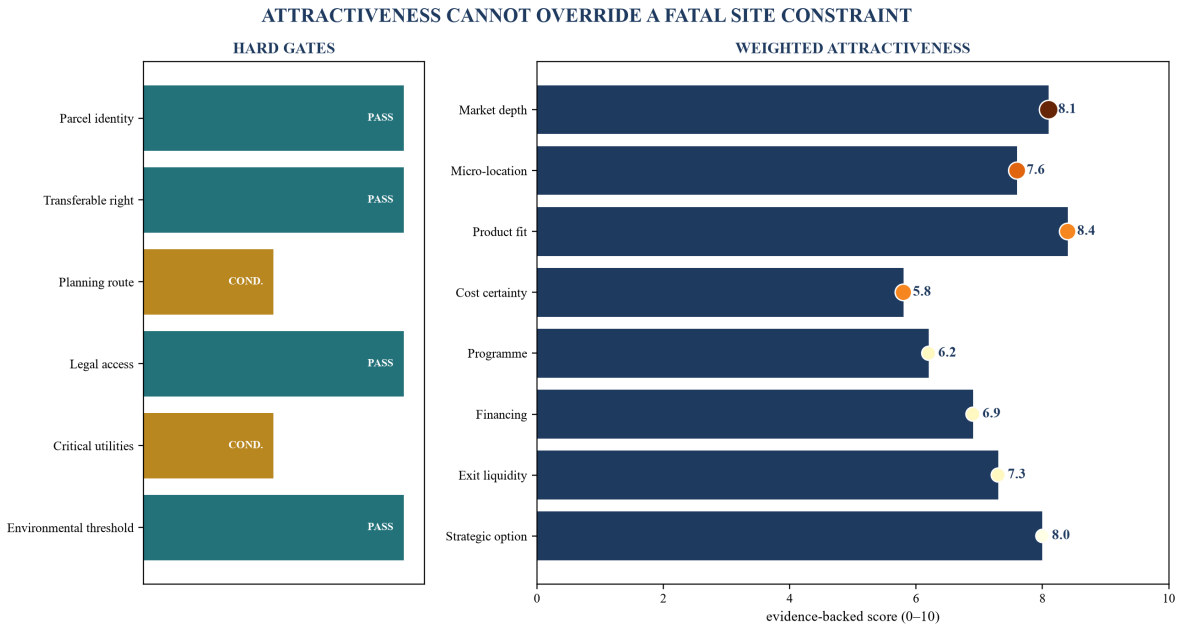

Scores and weights are hypothetical; unknown hard gates stop automatic advancement.

Table 2. Example hard gates and weighted factors

Test	Type	Example evidence	Decision treatment
parcel and ownership match	hard gate	registry extract, title map and identity reconciliation	stop if unresolved conflict is material
permitted use or credible planning route	hard gate	official plan, permission, conditions and adviser review	conditional only with defined route, timing and cost
legal access and utility pathway	hard gate	registered access, provider correspondence and capacity evidence	specialist escalation if access or capacity is uncertain
market depth and comparable quality	weighted	verified transactions, rents, supply, absorption and evidence grades	score with confidence and downside overlay
development cost and programme certainty	weighted	quantity surveyor cost plan, surveys, programme and contingencies	score plus sensitivity and independent challenge
funding and exit resilience	weighted	lender terms, covenant headroom, buyer universe and liquidity evidence	score under base and stressed conditions

Thresholds are governance choices; they require mandate-specific approval and periodic review.

4. Build a data-quality ladder before building features

The screening team should classify evidence before extracting predictive features. Useful dimensions are authority, transaction status, recency, spatial precision, asset comparability, completeness, licence, provenance and reproducibility. A completed registered transaction generally carries more pricing weight than an advertisement. A current official planning record carries a different status from a broker's interpretation. A technically rich vendor dataset can still be unsuitable if its licence prohibits the intended use or if its transformations cannot be traced.

Official datasets also have limitations. HM Land Registry Price Paid Data includes residential sales lodged for registration, is updated monthly and can be revised. The publisher states that the current month is incomplete as a standalone view and that it does not guarantee data accuracy or suitability for a particular purpose.[6] Saudi REGA's real-estate indicators distinguish recent raw data from refined historical indicators; a screening model should preserve that distinction instead of treating both as equally settled.[10] Dubai's open transaction data provides broad fields and daily refresh, while a live valuation still needs a transaction-status check, unit matching and appropriate market analysis.[8][9]

The lineage record should capture the original source, extraction time, licence, raw value, normalised value, transformation code, confidence, reviewer and use in the model. Derived data should retain its parents. If the system calculates distance to a metro station, the record should identify the site geometry, station dataset, coordinate system, routing or straight-line method and calculation version. This allows an error to be corrected at its source and every affected screen to be identified.

Data freshness requires a policy. Title, planning, transaction, rent, interest-rate, construction-cost and hazard data move at different speeds. The system should flag records beyond the approved age and should distinguish “no change found” from “not checked”. A committee pack should state its evidence cut-off date and any sources awaiting refresh.

5. Screen planning, access, utilities and physical constraints

Planning and infrastructure often determine whether apparent land value can be realised. The screen should capture current use, development parameters, height, density, setbacks, parking, affordable-housing or community obligations, phasing, design conditions, heritage, environmental limits and the approvals needed for the proposed scheme. A permitted envelope is a legal and technical starting point; it is not automatically the highest-value commercially deliverable scheme.

Access should be tested as a legal right and as an operational route. The site may touch a road while lacking a sufficient registered access right. Construction logistics can differ from permanent access. Heavy vehicles, turning geometry, bridge limits, school traffic, rights across neighbouring land and temporary works can affect programme and cost. The screen should state the evidence level rather than output a binary “accessible” conclusion from a map.

Utilities deserve project-specific evidence. Grid, water, sewerage, district cooling, telecoms and transport capacity can create long-lead dependencies. A nearby network does not prove capacity or a connection date. The screening record should capture provider correspondence, application status, required upgrades, charges, conditions, delivery date and the party responsible. A

data-centre, industrial or high-density residential scheme can be economically attractive and still fail its target programme because a critical connection is unavailable.

Physical and environmental screening should combine official maps, desktop studies, site reports and specialist judgement. England's long-term flood-risk service covers risks from rivers, sea, surface water, reservoirs and groundwater where data is available, while explicitly stating that it does not predict whether an individual property will flood.[11] That limitation illustrates a general control: area-level hazard data can trigger investigation and scenario costs; it should not be presented as parcel-specific engineering evidence.

6. Create a comparable-evidence confidence ladder

Comparable analysis should start with evidence status. The highest-confidence evidence is usually a verified arm's-length transaction for a genuinely comparable legal interest, use, development state, location and date. Confidence falls as the evidence moves towards unverified transaction reports, dated deals, different interests, materially different development rights, asking prices and anecdotal indications. The system should preserve both the observed value and the confidence grade.

The matching process should be explainable. It can use geography, permitted use, tenure, plot ratio, gross floor area, site condition, infrastructure, planning status, transaction date, purchaser profile and special conditions. Each proposed comparable should show why it was selected and why it may differ. A reviewer can reject a match, adjust it or retain it as a boundary rather than a point estimate.

Adjustments should be transparent and avoid double counting. If a comparable is adjusted for planning status, the development model should not separately add the same planning premium without reconciliation. The model should show whether an adjustment is directly evidenced, inferred from paired transactions, estimated by a valuer or used as a scenario assumption. IFRS 13's fair-value framework emphasises market-participant assumptions under current market conditions, while IVS places explicit importance on data, inputs, models, documentation and reporting.[1][2]

Thin markets require a wider evidence set and a wider range. Evidence can include land transactions, standing-asset values less transformation costs, developer margins, rental capitalisation, replacement cost, residual valuation and option value. The system should explain how each method relates to the subject site. Agreement among weak methods does not create strong evidence.

Figure 3. Comparable-evidence confidence ladder

EVIDENCE WEIGHT FALLS AS VERIFICATION AND COMPARABILITY WEAKEN

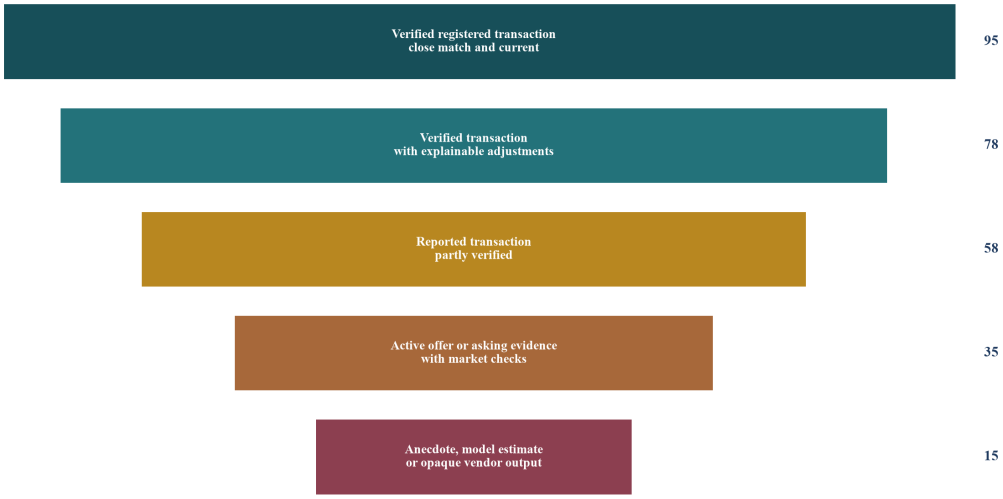

Confidence score is illustrative; retain source and adjustment rationale

Confidence depends on verification, relevance and current market conditions; the ladder does not prescribe a valuation conclusion.

Table 3. Comparable selection, adjustment and confidence record

Evidence feature	Stronger evidence	Weaker evidence	Control response
transaction status	completed and independently verified	asking price or unverified report	reduce weight; do not present as completed deal
legal and planning interest	same tenure, use and development rights	different rights or uncertain planning	quantify and explain adjustment or exclude
spatial relevance	same micro-market and access context	broad regional similarity	test submarket basis and widen range
time relevance	recent under comparable market conditions	old transaction across a market shift	time adjustment with observable evidence and sensitivity
asset and development state	similar physical state and obligations	serviced versus unserviced or permitted versus unpermitted	reconcile costs, time, risk and rights explicitly
source transparency	authoritative source and reproducible record	opaque vendor score or anecdote	require corroboration and disclose limitation

Each line should retain the observed value, adjustment logic and source link.

7. Build a transparent residual and development model

Land value is often residual. The model estimates the value of the completed scheme, then deducts construction, professional fees, infrastructure, finance, sales costs, tax where applicable, contingencies, required developer return and other obligations. The result is highly sensitive to timing, density, price, cost, absorption, funding and exit assumptions. AI can populate and test the model, while the model structure and decision logic should remain visible to the reviewer.

The base programme should start from the planning and technical evidence rather than a desired return. Unit mix, net sellable or lettable area, efficiency, phasing, start date, construction duration, sales or leasing curve and handover should reconcile. Cost inputs should state price date, scope,

exclusions, inflation, procurement route, contingency and whether utilities, public realm, finance and developer overhead are included.

Revenue assumptions should separate market evidence from product strategy. Comparable transactions may support a range for current pricing. The proposed scheme can require a premium or discount for specification, view, brand, unit size, timing and payment plan. The model should show the evidence for each adjustment and should test absorption as well as price. A high price assumption that extends the sales period can reduce value through finance and carrying costs.

Funding should reflect the actual capital structure. Land payment timing, equity, senior debt, mezzanine, presale receipts, draw conditions, interest, fees, covenants and minimum cash determine liquidity. A project can show a positive net present value and still fail because it cannot fund a peak cash requirement. The screen should present maximum funding need, timing and covenant headroom alongside return metrics.

The model should reconcile nominal and real assumptions, pre-tax and post-tax measures, levered and unlevered returns, and the valuation date. It should state whether the output is land value, enterprise value, equity value, maximum bid or strategic value. This prevents an investment committee from comparing unlike numbers.

8. Present scenarios, distributions and decision sensitivities

A single-point output hides uncertainty. The committee needs a base case, credible downside, severe but plausible case and upside, together with the assumptions that move between them. Scenarios should remain internally coherent. A downside that reduces selling prices while leaving absorption, debt margin, construction cost and exit yield unchanged can understate the combined effect of a weaker market.

Deterministic sensitivities reveal the variables with the greatest impact. The screen should test price, rent, yield, construction cost, delay, interest rate, density, sales velocity and contingency. Two-way tables can reveal interactions, such as price and cost or rent and exit yield. Threshold analysis answers the committee's operational questions: the price decline that removes the target return, the delay that exhausts liquidity, or the cost increase that reduces land value below the seller's expectation.

Probabilistic analysis can add value when distributions and dependencies are grounded in evidence. It should not turn subjective inputs into apparent statistical truth. The model should disclose distribution choices, correlations, sample size, tail treatment and the difference between modelled probability and observed frequency. A fan chart can show how uncertainty widens through the programme and where capital protection is weakest.

The decision should focus on controllable conditions. The investment team may improve downside protection through price, deferred consideration, option structure, planning condition, utility condition, exclusivity, cost sharing or a staged acquisition. The screen becomes commercially useful when it connects risks to transaction terms and diligence priorities.

Figure 4. Illustrative scenario fan chart for equity value through the development programme

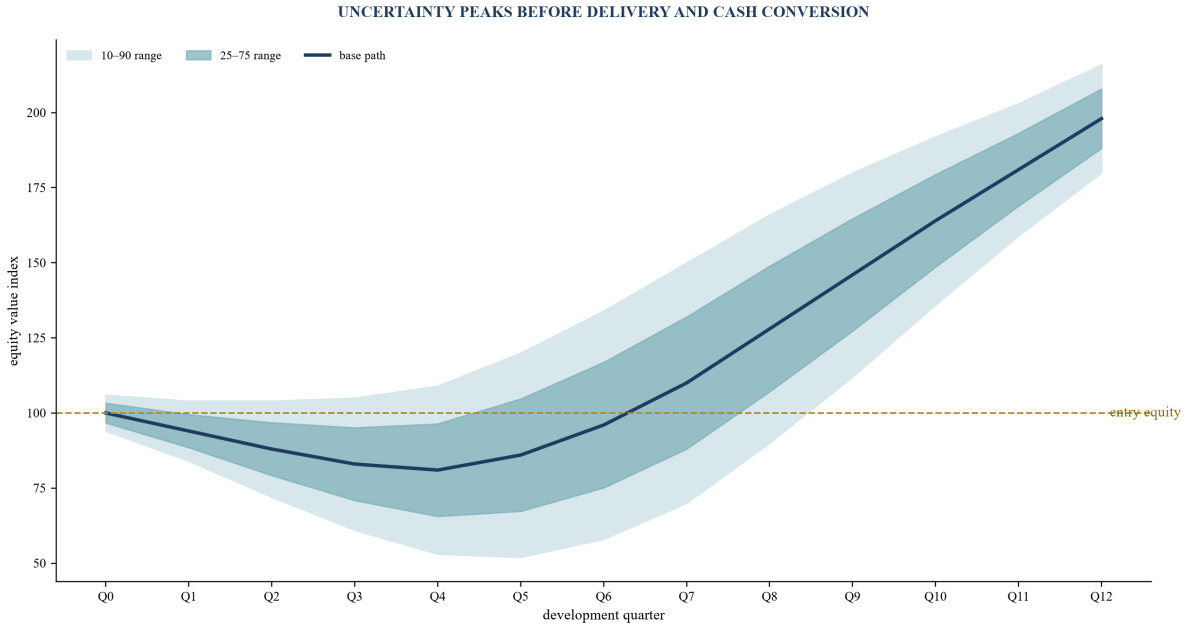

Values are hypothetical index points; the chart demonstrates widening and contracting uncertainty rather than a forecast.

Table 4. Illustrative coherent scenario set

Assumption or output	Upside	Base	Downside	Severe but plausible
achieved sales-price index	106	100	91	82
construction-cost index	97	100	110	120
programme delay	0 months	3 months	9 months	15 months
senior all-in funding cost	6.5%	7.5%	9.0%	10.5%
sales period after completion	6 months	12 months	20 months	30 months
peak equity requirement index	88	100	128	158
equity value at decision date	156	118	61	12

All values are hypothetical and demonstrate model governance rather than market forecasts.

9. Validate the model and monitor its performance

Validation should be proportionate to materiality and should cover conceptual soundness, implementation, data, outputs and use. The Federal Reserve's 2026 revised guidance on model risk management emphasises a risk-based approach, clear roles, validation, ongoing monitoring and understanding of vendor models for banking organisations within its scope.[12] The principles provide a useful governance reference for any material screening model, although their legal applicability depends on the institution and activity.

Conceptual review asks whether the model represents the actual decision. A classifier trained on completed transactions may select sites similar to deals that closed, rather than sites that created

value. Historical closing data also embeds selection bias, market regime and institutional strategy. Validation should test whether the target variable, features and loss function align with the investment mandate.

Implementation review tests extraction, transformations, formulas, units, currencies, dates, joins, model versions and interfaces. A correct method can fail through a decimal, coordinate-system, duplicate-record or period mismatch. Test cases should include ordinary parcels, boundary cases, missing data, contradictory documents and adversarial inputs. Material spreadsheet outputs should reconcile to an independently calculated benchmark.

Outcome monitoring compares early screens with later diligence, bids, acquisitions and realised results. Useful measures include false rejections, false advancements, evidence gaps discovered late, comparable overrides, valuation error, cost and programme variance, decision time and reviewer workload. Performance should be segmented by geography, asset type, market regime and data availability. A model can perform well overall while failing in a thin submarket.

Change control should identify new data, code, prompts, models, thresholds and interfaces. Material changes require testing and approval before production. Vendor updates should not enter the decision process silently. The register should record the approved version used for each committee pack and the rollback path.

10. Govern generative AI and document extraction

Generative AI is useful for extracting covenants, planning conditions, seller representations, report findings and inconsistencies. Its output should remain a proposed extraction until linked to the source and reviewed at the required level. The system should display the document, page, relevant passage, extraction confidence and any conflicting source. A summary without a source location is difficult to verify and should not drive a material decision.

Prompt and retrieval design require governance. The system should restrict retrieval to approved evidence for the matter, prevent cross-deal data leakage, log the model and prompt version, and control access to confidential documents. It should distinguish instructions contained in source documents from authorised workflow instructions. A document can contain adversarial or irrelevant text; the extraction pipeline should treat it as evidence rather than operational authority.

Reliability tests should reflect the actual work. A benchmark can measure parcel-identifier extraction, planning-condition classification, comparable status, numerical extraction, table reconstruction, unit conversion, citation accuracy and abstention when evidence is missing. The benchmark should include poor scans, amendments, conflicting schedules and multiple languages where relevant.

Human review should be risk based. A qualified professional should review material valuation inputs, legal and planning conclusions, fatal-gate status, exceptions and the final recommendation. Low-risk administrative fields can use sample-based control when accuracy has been demonstrated. RICS requires a named qualified surveyor to make and document a reliability decision when AI output materially affects professional work within its scope.[3]

11. Design the investment-committee evidence pack

The committee pack should make the decision reproducible. It begins with the mandate fit, proposed decision, amount, structure, return range and principal conditions. It then presents parcel identity, hard gates, evidence quality, market and comparable analysis, development concept, financial scenarios, funding, risks, diligence plan and requested authority. Each material number should link to a model cell and source record.

The pack should separate facts, professional conclusions and hypothetical modelling assumptions. Facts state the source and effective date. Professional conclusions name the reviewer and evidence considered. Assumptions state their owner, rationale, sensitivity and expiry. This structure allows the committee to challenge the right layer.

Exceptions should be prominent. The pack should list unresolved gates, missing evidence, model limitations, overrides, dissenting views and the consequences of delay. A recommendation to advance can be conditional on title confirmation, utility capacity, a planning opinion, cost validation or revised price. The minutes should preserve the conditions and who has authority to confirm them.

The decision workflow should also prevent automation bias. A high score should not decide the outcome. The committee can reject a highly scored site because of concentration, reputation, capability or strategy. It can advance a lower-scored site because it has a valuable option, provided the rationale and conditions are documented. Override analysis then becomes part of model monitoring.

Figure 5. Investment-committee decision tree for an AI-supported property screen

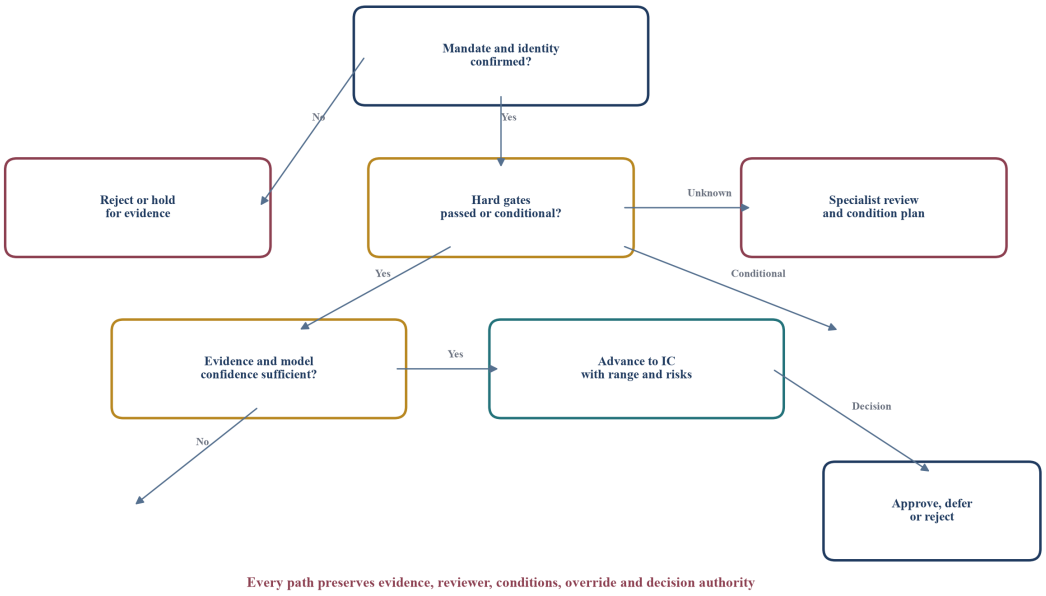

The tree governs escalation; live authority levels depend on the organisation's approved delegation framework.

Table 5. Decision rights and control record

Activity	Primary owner	Independent challenge	Required record
mandate and triage rules	investment leadership	risk and portfolio management	approved criteria, thresholds and change history
parcel and rights record	transaction team	legal and technical specialists	source documents, conflicts, effective dates and resolution
comparable set and valuation model	valuation or investment professional	independent reviewer	selection logic, exclusions, adjustments and sensitivities
AI model and extraction controls	data and model owner	validation, risk or audit function	model version, benchmark, monitoring and incidents
screening recommendation	deal lead	investment committee and specialists	evidence pack, limitations, conditions and overrides
final approval and conditions	delegated committee	governance or compliance as applicable	minutes, authority, conditions, expiry and follow-up owner

Roles are illustrative and should be aligned to the organisation's delegations and professional obligations.

12. Worked hypothetical screening case

Consider a hypothetical 42,000-square-metre urban-edge parcel proposed for a mixed residential scheme. The seller provides a brochure, coordinate pin, concept massing and indicative price. The registry identifier is present, while the indicative map and seller plan show a small boundary difference. The system extracts both geometries, identifies the mismatch and prevents automatic area normalisation. Legal review confirms the registered interest and explains the mapping difference before the parcel is advanced.

The planning screen identifies an adopted residential designation, a density range and a requirement for a transport contribution. The concept sits near the upper density boundary. Utility evidence confirms a water pathway, while grid capacity remains subject to a formal study. The scorecard therefore records planning as pass with a design sensitivity and electricity as conditional. The condition has an owner, expected response date, application cost and downside programme impact.

The comparable engine finds 34 candidate records. Eleven are asking prices, seven lack a reliable planning match, four are older than the approved period and three cannot be reconciled to a parcel. Nine remain for detailed review. Three verified transactions receive high confidence, four receive medium confidence after transparent adjustments and two are retained only as range boundaries. The reviewer rejects one algorithmic match because a transport interchange created a location premium absent from the subject site.

The residual model uses a hypothetical base sales-price index of 100, construction-cost index of 100, three-month programme allowance and 7.5 per cent all-in senior funding cost. The downside combines a price index of 91, cost index of 110, nine-month delay and 9.0 per cent funding cost. The severe case uses 82, 120, fifteen months and 10.5 per cent. These figures demonstrate coherent scenario design and are not observations about any market.

The base case supports advancement, while the severe case leaves little equity value and creates a materially higher peak funding need. Threshold analysis shows that price, density and

construction delay dominate the decision. The team therefore proposes a price range, a utility-capacity condition, confirmatory cost review and a planning-density opinion before a binding bid. The committee advances the opportunity to preliminary underwriting with those conditions and records that no final valuation has been approved.

The case shows how speed and control can reinforce each other. Automated extraction narrows the evidence set, identifies conflicts, prepares the scorecard and runs scenarios. The human team resolves the legal parcel, judges comparable relevance, owns the assumptions and designs the transaction conditions. The output is faster because reviewers focus on material exceptions and commercial decisions.

13. Implementation roadmap

Implementation should begin with the investment decision and evidence model rather than a broad technology procurement. The first phase maps the current screening workflow, decision rights, data sources, recurring errors, cycle time and outcomes. It identifies one transaction type and geography with sufficient evidence for a controlled pilot. The team defines hard gates, score anchors, source hierarchy and the minimum committee pack.

The second phase builds the parcel record, evidence register and transparent model. Existing spreadsheets can remain part of the workflow if their logic is controlled and interfaces are tested. The system should first automate retrieval, extraction, reconciliation and missing-evidence detection. Predictive ranking should follow only after the organisation has stable labels and can explain the decision target.

The third phase runs the new process in parallel with experienced reviewers. The pilot records extraction accuracy, false matches, missing evidence, reviewer changes, valuation variance, decision time and model incidents. The team sets acceptance thresholds and documents when the system must abstain. Independent challenge should review both the quantitative model and the operational process.

The fourth phase introduces controlled production use. Access, privacy, retention, supplier, cyber, model, licence and business-continuity controls are approved. Changes follow a release process. Every committee pack records the model and data version. Monitoring compares screening decisions with later diligence and realised outcomes.

The fifth phase expands cautiously across geographies and asset classes. A model validated on UK residential transactions should not be assumed to work for Saudi industrial land, UAE off-plan development or US life-science conversion. Each expansion requires local data, legal and planning mapping, benchmark cases, qualified reviewers and approved decision thresholds.

14. Limitations and conclusion

AI-supported screening cannot replace title investigation, planning advice, surveys, valuation judgement, technical diligence, environmental assessment, financing analysis or investment authority. Public and vendor data can be incomplete, delayed, revised, licensed for limited uses or unsuitable for parcel-level conclusions. Historical transactions embed market regime and selection effects. Thin markets and novel assets can produce wide ranges even when the workflow is well governed.

Model outputs can fail through source errors, extraction errors, entity-resolution mistakes, concept drift, implementation defects, inappropriate objectives, hidden vendor changes and human automation bias. Governance reduces these risks through lineage, validation, monitoring, independent challenge, abstention and documented decision rights. It does not remove uncertainty from property investment.

The commercial value of the framework comes from disciplined focus. Teams can reject clear mismatches earlier, identify missing evidence sooner, compare opportunities consistently and direct professional attention to the assumptions that matter. Investment committees receive a range, evidence quality, conditions and downside path instead of a polished score with hidden fragility.

A strong system therefore starts with parcel truth, separates gates from attractiveness, grades comparable evidence, models coherent scenarios and keeps accountable professionals in control. The result is a faster route from opportunity to a defensible decision, with an evidence trail that remains usable through diligence, negotiation, financing and post-investment review.

References

[1] International Valuation Standards Council, International Valuation Standards, effective 31 January 2025, including IVS 104, IVS 105, IVS 106, IVS 400 and IVS 410. https://ivsc.org/standards/

[2] IFRS Foundation, IFRS 13 Fair Value Measurement. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[3] Royal Institution of Chartered Surveyors, Responsible use of artificial intelligence in surveying practice, professional standard effective 9 March 2026. https://www.rics.org/content/dam/ricsglobal/documents/standards/Responsible-use-of-artificial-intelligence-in-surveying-practice_September-2025.pdf

[4] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023. https://doi.org/10.6028/NIST.AI.100-1

[5] National Institute of Standards and Technology, AI RMF Core. https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

[6] HM Land Registry, Price Paid Data downloads. https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads

[7] HM Land Registry, INSPIRE Index Polygons spatial data. https://www.gov.uk/guidance/inspire-index-polygons-spatial-data

[8] Dubai Land Department, Real Estate Data. https://dubailand.gov.ae/en/open-data/real-estate-data/

[9] Digital Dubai, Dubai Land Department Transactions Open Data. https://www.dubaipulse.gov.ae/data/dld-transactions/dld_transactions-open

[10] Real Estate General Authority, Saudi Real Estate Indicators. https://rei.rega.gov.sa/en

[11] UK Government and Environment Agency, Check the long-term flood risk for an area in England. https://www.gov.uk/check-long-term-flood-risk

[12] Board of Governors of the Federal Reserve System, SR 26-2, Revised Guidance on Model Risk Management, 17 April 2026. https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm

[13] National Institute of Standards and Technology, AI RMF Playbook. https://airc.nist.gov/airmf-resources/playbook/

[14] Royal Institution of Chartered Surveyors, AI in real estate valuation, global guidance project. https://www.rics.org/profession-standards/rics-standards-and-guidance/sector-standards/valuation-standards/ai-in-real-estate-valuation

[15] HM Land Registry, Public data. https://www.gov.uk/government/publications/hm-land-registry-data/public-data

[16] Federal Housing Finance Agency, Quality Control Standards for Automated Valuation Models, final rule information. https://www.fhfa.gov/regulation/federal-register/rin-2590-aa62

[17] Bank for International Settlements, Artificial intelligence and information processing in finance, Working Paper No. 1194, June 2024. https://www.bis.org/publ/work1194.htm

About the Author

Authored by Chennakeshav Adya, Independent Researcher

Chennakeshav Adya is a corporate finance and investment banking executive with more than 25 years of global experience in deal origination, structuring and execution across mergers and acquisitions, growth capital and corporate strategy. At Matchpoint Partners, he is Managing Partner and leads the firm's corporate finance, M&A and capital-raising practice.