

RISK MANAGEMENT REPORT / AI IN PRACTICE / AUGUST
2026



AI: Real Estate's Misunderstood Risk

*An Assessment of AI Risk Management in Real Estate and the
Deficiencies in Licensing Law and Broker Oversight*

MATTHEW RATHBUN

WWW.MATTHEWRATHBUN.COM



Artificial intelligence is, on balance, one of the more consequential productivity advances available to the real estate profession in a generation. Its benefits to agents, brokers, and the consumers they serve substantially exceed its risks. That statement should not be read as a qualification or a hedge. It is the premise this report proceeds from.

The risk side of that ledger, however, is largely unaddressed. Not lightly regulated. Not under-enforced. Largely unaddressed, by state real estate commissions, by brokerages, and to some extent, by the REALTOR® association infrastructure that trains and governs the profession. That gap, not the technology itself, is the subject of this report.

Artificial intelligence adoption in the real estate industry is not a trend subject to reversal. It is now embedded in the devices, platforms, and productivity tools agents and brokers use daily, often without their active choice, and its presence in the profession is best understood as a permanent condition to be governed well rather than a phenomenon to be resisted or wished away.

The rules will keep changing. Build practice and policy around durable principles rather than the statute of the moment, and neither the agent nor the brokerage will have to relearn the job every legislative session.

EXECUTIVE SUMMARY

Scope and Structure of This Report

PART I

Catalogs thirteen discrete categories of operational, consumer-facing, and market risk arising from the use of artificial intelligence in day-to-day real estate practice, ranging from Fair Housing exposure in AI-generated advertising to algorithmic bias in background screening and private-listing buyer targeting. Each category is presented with its primary exposure, its current regulatory status, and practical mitigation guidance.

PART II

Moves beyond discrete, transactional risk to examine three structural failures that make the risks in Part I more likely to occur and less likely to be caught: the near-total absence of AI governance at the brokerage level, the intrusion of wearable and ambient AI into the physical transaction environment, and the mismatch between how regulators think about AI oversight and how the technology needs to be governed. Part II closes with an affirmative policy argument, that the appropriate regulatory target is agent accountability for output, not futile attempts to prohibit or ration access to the technology itself.

PART III

Examines the National Association of REALTORS Code of Ethics and Standards of Practice against this risk landscape, identifying where existing Articles already reach AI-related conduct and proposing specific, narrowly drafted amendments where they do not.

A NOTE ON METHOD. The concerns, positions, and policy recommendations throughout are the author's own professional judgment, informed by two decades of teaching and advising within the real estate industry. Artificial intelligence tools assisted with research, citation verification, and formatting. They did not originate the opinions expressed. The full disclosure statement appears at the close of this report.

PART I

Operational, Consumer-Facing & Market Risk

Thirteen categories of risk that already exist in day-to-day practice, ranked by how much of the exposure is settled law versus a moving target.



OVERVIEW

The Risk Matrix

Thirteen categories, ranked by how much of the exposure is settled law versus moving target.

CATEGORY	PRIMARY EXPOSURE	STATUS
01. Fair Housing & Algorithmic Discrimination	Steering, disparate impact in ads and screening, biased marketing copy	Live enforcement risk
02. Data Privacy & Client Confidentiality	Client Personal Identifiable Information entered into consumer AI tools without governance	Growing, brokerage-level
03. Accuracy, Hallucination & Negligence	Fabricated facts, comps, or citations relied on without verification	Constant, underrated
04. Disclosure & Transparency Rules	AI-altered photos, synthetic media, consumer notice requirements	Fast-moving patchwork
05. Fraud, Deepfakes & Cybersecurity	Wire fraud, seller impersonation, voice cloning at closing	Live enforcement risk
06. Vendor Vetting, System Access & Insurance Gaps	E&O and cyber policies silent on AI, agents connecting unvetted tools to their own systems	Underexamined
07. Brokerage Supervision & Vicarious Liability	Agents adopting tools the brokerage never developed policies around or are aware of	Structural, ongoing
08. Intellectual Property & Content Rights	Generated images and copy that borrow style, likeness, or text	Low but rising

CATEGORY	PRIMARY EXPOSURE	STATUS
09. The Regulatory Wildcard	Federal preemption fight, stayed and rewritten state AI laws	Genuinely unsettled
10. Misinformation Cascade & Consumer Trust Gap	AI-hallucinated content cited as fact by other AI tools, agents trusted less than AI on factual questions	Live, underrated
11. Agentic AI Misuse & Regulatory Capacity Gap	Unsupervised agentic tools completing CE or license-required work, regulators unequipped to detect it	Live, largely undetected
12. Algorithmic Background Screening & ID Errors	Bulk-data reports treated as fact, automated screening on false-positive identity matches	Live, underrated
13. Algorithmic Buyer Targeting, Private Listings	AI-curated buyer exposure on off-MLS listings, narrowed further under dual agency	Emerging, watch closely

SECTION 01

Fair Housing & Algorithmic Discrimination

The one with an existing enforcement framework behind it. No new statute has to pass for this to bite.

The Fair Housing Act does not care whether a human or an algorithm made the call. HUD has been explicit since 2019, when it charged Meta over ad-targeting tools that let advertisers exclude users by drawing a line around neighborhoods, toggling off a gender, or excluding non-English speakers. That case set the pattern courts still use.

Three places this shows up in a normal brokerage, none of which feel like "AI" when they happen:

01. AI-generated marketing copy

Language a chatbot drafts in seconds can carry the same red flags a human writer would know to avoid. "Great for young professionals" reads as age discrimination. "Close to churches" signals religious preference. "Perfect for families" is familial-status steering. Volume is the problem: agents are generating this copy at a scale no one is proof-reading for Fair Housing exposure.

RISK NOTE. Run every AI-generated description through the same Fair Housing screen you would apply to a human draft. The tool did not do the writing, you did, as far as liability is concerned.

02. Ad delivery and targeting

Platform-side delivery optimization can quietly steer a listing away from protected classes based on predicted engagement, even when nobody set out to exclude anyone. This is disparate impact, not disparate treatment, and intent is not a defense.

03. Tenant and applicant screening

Third-party scoring models built on incomplete or historically biased data can reproduce the exact lending and rental patterns Fair Housing law exists to stop. If your brokerage or property management arm uses a screening vendor, know what the model weighs and ask for its bias-testing documentation.

WORTH FLAGGING

The federal posture is shifting under you

HUD proposed removing its disparate-impact regulations in January 2026, and the CFPB finalized a rule in April 2026 eliminating disparate-impact liability for lending decisions specifically. That narrows federal enforcement appetite. It does not repeal the Fair Housing Act itself, and it does nothing to state fair housing laws, private plaintiffs, or state attorneys general, several of whom have said they will fill the gap. Write your policy to the statute, not to whichever administration is currently enforcing it.

"Family" is not illegal to say. The 1995 HUD Achtenberg memo specifically protects property-descriptive terms like "master bedroom" or "family room" from being treated as violations. The line is between describing the property and describing who belongs in it. AI tools do not reliably know that line, so you have to.

SECTION 02

Data Privacy & Client Confidentiality

Every agent feeding client PII into a consumer AI tool is creating a data-governance question nobody signed off on.

Financial details, transaction terms, SSNs on loan paperwork, family circumstances behind a sale. Agents paste this into consumer-grade AI tools daily, often to draft a simple email, with no idea what happens to that data afterward or whether the tool's terms of service even allow real estate use.

01. Consumer tools are not built for client data

Free and consumer-tier AI products generally reserve the right to use submitted content to improve their models unless you are on an enterprise agreement with data controls turned on. "I pasted the contract in to summarize it" is a data-handling decision, not a convenience.

02. This is a brokerage vicarious-liability problem

The broker is responsible for tools agents adopted without telling anyone. State privacy regimes (CCPA-style amendments, biometric laws like Illinois BIPA) and frameworks like the Colorado AI Act's reasonable-care standard treat the deployer, not just the tool vendor, as accountable.

RISK NOTE. A brokerage AI-use policy that never mentions consumer chatbots is not a policy. It is a gap waiting to be revealed in discovery.

03. A tiered restriction on which tools may touch client data

A defensible brokerage policy, and arguably a defensible regulatory standard, would distinguish between AI tools by service tier rather than treating all AI use as equivalent. A workable rule: agents may not upload or review personal data, financial account information, dates of birth, Social Security numbers, or specific contract terms, in any AI model that is not under an enterprise-grade agreement with data controls appropriate to real estate use. General-purpose consumer tools remain suitable for drafting, brainstorming, and non-sensitive tasks. They are not suitable for anything touching identifiable client data absent a specific data-handling agreement.

RISK NOTE. This is not a restriction on using AI. It is a restriction on which tier of AI may see identifiable client information, the same distinction a firm would draw for any other vendor handling sensitive data.

BEFORE YOU PASTE IT IN

The five-second gut check

Would you be comfortable if this exact text showed up in a vendor's training set, a subpoena, or a headline. If not, redact the names and numbers before it goes into any AI tool that is not under a signed enterprise agreement with your brokerage.

SECTION 03

Accuracy, Hallucination & Professional Negligence

The risk nobody in the industry talks about enough, because it does not look like a scandal. It looks like a typo, until it is a lawsuit.

AI models generate plausible-sounding text whether or not the underlying fact is true. A fabricated square footage, an invented comp, a citation to a code section that does not say what the AI claims it says. Every one of these is a professional representation the moment you hand it to a client, whether you meant it to be or not.

01.

AI-assisted CMAs and valuations

A model can produce a clean, confident-looking comparative market analysis built on stale or mismatched comps. Presented to a seller as your professional opinion, an error like that is not a software bug, it is a negligent misrepresentation claim with your name on the signature line.

02.

Disclosure and compliance language

Asking a general-purpose AI tool to draft addendum language, disclosure text, or a summary of a code section without checking it against the actual statute or your state's approved form is how outdated or simply invented language ends up in a signed document.

03.

Client-facing communication

A drafted email that misstates a deadline, a financing term, or a contingency date reads exactly as authoritative as a correct one. The client cannot tell the difference. You have to be able to.

04. Most hallucination is a setup problem, not a technology problem

The instinct to categorically distrust AI output deserves a correction of its own. A meaningful share of what gets labeled a hallucination traces back to poor prompting and inadequate system setup by the user, vague instructions, missing context, no source material provided, no verification step built into the workflow, rather than an inherent unreliability in the tool itself. An agent who has never been taught how to construct a clear prompt, supply the right source documents, or configure a tool correctly will get worse output than one who has, using the identical underlying model.

RISK NOTE. This does not reduce the verification obligation above. It does mean that training agents to use AI tools competently is itself a risk-mitigation measure, not a separate, optional convenience.

TWO-STEP HABIT

Verify before it leaves your hands

Treat every AI-generated fact, figure, citation, or comp the way you would treat information from a source you have never worked with before. Cross-reference it against MLS data, the actual code section, or your broker's approved templates before it reaches a client. Document that you did.

If it isn't documented, it didn't happen.

SECTION 04

Disclosure, Transparency & the State Patchwork

Technology and innovation will always outpace regulation.

Several states have moved on AI-specific disclosure requirements for real estate marketing, and the patchwork has already started shifting under its own weight. What follows is the state of play as of this report. Verify current status before relying on any of it in a compliance conversation.

RULE	STATUS AS OF AUG 2026	WHAT IT REQUIRES
California AB 723, codified at Bus. & Prof. Code § 10140.8 (digitally altered listing images)	In effect	Digitally altered listing images must be disclosed, with online access to the clearly identified original. Full detail in Appendix A.
California SB 942 (AI transparency)	In effect	Broader AI content-labeling requirements, penalties up to \$5,000 per day.
Virginia HB 2094 (High-Risk AI Developer & Deployer Act)	Vetoed, not law	Would have required consumer disclosure for high-risk AI in consequential decisions including housing. Governor Youngkin vetoed it in March 2025. Do not train agents as if this is in force.
Colorado AI Act (SB 24-205)	Enforcement stayed	Original high-risk-AI framework is under federal court stay following an X.AI suit joined by DOJ. A narrower replacement (SB 189) pushed the effective date to January 1, 2027 with a scaled-back scope.
MLS "virtually staged" labeling rules	Widely in effect	Most MLS systems now require an explicit label on virtually staged or AI-enhanced images, separate from and often stricter than state law.
NAR Code of Ethics, Articles 2 & 12	In effect	Existing honesty-in-representation and honesty-in-advertising duties apply to AI-generated content with no carve-out.

Two things to take from this table. First, Virginia has no comprehensive AI-specific real estate disclosure statute in force right now, despite what got repeated in some 2025-era training material, including some of mine. Second, the state landscape is genuinely volatile, with laws passing, getting vetoed, getting stayed by federal courts, and getting rewritten inside the same calendar year. Build training and policy around the durable layer, which is MLS rules and the NAR Code of Ethics, and treat any specific state statute citation as something to re-verify before you repeat it in a classroom. Appendix A provides a jurisdiction-by-jurisdiction summary of enacted and proposed measures, including agency guidance and pending bills in New Jersey and New York.

01. A defensible disclosure standard requires both a global notice and per-asset labeling

A single blanket disclosure is not sufficient on its own, and neither is a policy of labeling only the assets someone happens to flag. The more defensible standard operates on two levels simultaneously. At the listing or marketing-set level, a general notice that some or all content associated with the listing may be AI-generated, placing a reasonable burden on the consumer to confirm material facts before acting. At the individual-asset level, a specific label, a watermark or an on-image line, on any photo or video that has been visually enhanced or materially altered by AI. The two are not substitutes for each other.

02. Not all AI-generated content requires the same disclosure

A property description drafted with AI assistance and then reviewed and edited by the agent is a different category of content than a virtually staged photo or an AI-enhanced video, and the disclosure standard should reflect that difference rather than applying uniformly to everything AI touched at any stage. Visual content that alters a consumer's perception of the physical property, staging, sky replacement, room enlargement, warrants the most specific, most visible disclosure, because it is the category most likely to be relied upon as a literal representation of what a buyer will see in person.

Fraud, Deepfakes & Cybersecurity

AI did not invent wire fraud. It made the con artists faster, cheaper, and far more convincing.

Title, settlement, and escrow sit at the center of this one, but every agent touching a transaction is a link in the chain. Seller impersonation fraud, synthetic identity, and deepfake voice cloning of wire instructions are no longer theoretical. Business email compromise, already the most common closing-stage fraud, is faster and more convincing with AI generating the pretext.

01. Signal verification beats voice verification

A cloned voice can now pass a casual phone check. It cannot forge an out-of-band confirmation through a channel the fraudster does not control. Verify wire instruction changes through a known phone number you dial yourself, never one provided in the email or call requesting the change.

02. Seller impersonation on remote or vacant listings

Synthetic identity and forged documentation are increasingly used to "sell" property the impersonator does not own, often targeting vacant land or absentee-owner properties where in-person verification is easy to skip.

RISK NOTE. Treat identity verification on any remote or unusually rushed transaction as a required step, not an optional courtesy.

03. Automation is erasing the human fraud checkpoint

We are operating in what several federal and industry sources now describe as a pandemic-level volume of real estate fraud, and agents are increasingly letting AI automation triage and answer email at the exact moment more scrutiny is needed, not less. An AI model reading a request for a client's date of birth or account number may see nothing unusual about it in isolation. An agent who has closed forty transactions with that settlement company knows a payout inquiry from that vendor has never once looked like this. The AI has statistical pattern matching. The agent has relational memory, this is what an email from the title company normally looks like, this is not. The same gap shows up with a subtly wrong reply-to address or a domain one character off from the one used for years. An agent who has worked with the same loan officer for years may notice the writing voice changed. A general AI assistant running the inbox on autopilot has no baseline for that unless it is specifically built and configured to check for it.

RISK NOTE. The fix is not banning automation from a client-facing inbox. It is keeping a human in the loop on anything involving money movement, account numbers, or a request that deviates even slightly from a known relationship's pattern, and building that checkpoint into the automation's design rather than assuming the agent will remember to look.

WIRE FRAUD LANGUAGE, EVERY TIME

Say it before it matters

Wire instructions come from the title company only. If anyone calls or emails claiming to update wire information, hang up and call your title company back at a number you already had, before sending a dollar. Put this line in every closing-adjacent email and every buyer consultation. Wire fraud is real and irreversible.

SECTION 06

Vendor Vetting, System Access & Insurance Gaps

The gap between what your AI tools can do and what your insurance actually covers is wider than most brokerages have checked. So is the gap between what agents are connecting and what they understand about it.

AI fraud-detection tools, transaction-management AI features, and lead-scoring vendors all sit on top of your data and your client relationships. Most E&O and cyber policies were written before generative AI existed as a practice tool, and many have not been updated to address it explicitly.

01. Ask before you adopt, not after a claim

Where is data stored and for how long. Is it used to train the vendor's models. What happens to it if you cancel. Does the vendor carry its own cyber liability coverage. A vendor that cannot answer these in writing is a vendor you have not actually vetted.

02. Documentation closes the insurance gap

Where AI fraud-detection tools exist but coverage language is silent on AI-related losses, the paper trail of what the tool flagged, what a human reviewed, and what decision was made is often what determines whether a claim gets paid.

RISK NOTE. Ask your E&O carrier directly whether AI-assisted work product is covered under your current policy. Get the answer in writing, not a verbal assurance.

03. Most agents can't vet what they don't know exists

Every AI vendor pitch promises a shortcut. Very few agents have been trained to ask whether the shortcut is legal to use in the first place, because most agents have surprisingly thin working knowledge of the analog compliance framework that already governs them: state license law, RESPA, TCPA, data-safeguard rules, and the Code of Ethics itself. You cannot interview a vendor about whether their system meets regulatory requirements you cannot name. The gap here is not really AI literacy. It is baseline compliance literacy that AI adoption is now exposing at scale.

RISK NOTE. Vendor-vetting training that opens with "here are five questions to ask an AI vendor" skips a step. Start with what the law already requires of you. Teach the vetting questions as a way to verify a vendor helps you meet that, not the other way around.

04. Unsecured, unencrypted vendor connections into an agent's operating environment

Agents are increasingly granting AI tools and connected vendors direct access to email, calendars, transaction management systems, and file storage, often through consumer-grade integrations never built or audited for real estate data sensitivity. Few agents ask whether the connection is encrypted in transit and at rest, what the vendor's data retention policy is, or what happens to that access after they stop paying for the tool. Fewer still understand what permissions they actually granted when they clicked "allow" during setup.

RISK NOTE. A connected app with broad read and write access to an agent's inbox is functionally a data processor for every client communication that passes through it. Treat every new integration request the way you would treat a request to hand a stranger the key to your filing cabinet, because that is closer to what is actually happening.

05. A case for mandatory cyber insurance, modeled on the existing escrow requirement

Most states already require brokers to maintain escrow or trust accounts under defined safekeeping rules, on the theory that client funds passing through a brokerage warrant a mandatory financial safeguard. Given the documented rise in identity theft, deed fraud, and wire fraud specific to real estate closings, it would not be an unreasonable extension of that same logic for regulators to require every brokerage to carry cyber liability insurance, and related coverage addressing AI-assisted fraud specifically, as a condition of maintaining a brokerage license. The regulatory infrastructure for this kind of mandate already exists. It is simply not yet pointed at this exposure.

06. AI is a genuinely valuable tool for the brokerage side of this problem

It is worth stating plainly, in a section otherwise devoted to gaps and exposure, that AI tools are a legitimately strong resource for brokers specifically in drafting and reviewing internal policy, and in monitoring financial and escrow accounts for anomalies that a manual review might miss. None of the caution elsewhere in this report is an argument against that use.

RISK NOTE. The tools that help a broker write better policy and monitor escrow activity still need to be monitored and governed themselves. Value and oversight are not in tension. A tool worth using is a tool worth governing.

SECTION 07

Brokerage Supervision & Vicarious Liability

The structural risk underneath every other section in this report. A broker is responsible for tools the broker never approved.

Agents adopt AI tools the way they adopt anything else that makes their day easier: quietly, individually, and usually without asking. That is normal human behavior and also exactly how a brokerage ends up vicariously liable for a Fair Housing violation, a data breach, or a fabricated disclosure it never knew was happening.

01. A written AI-use policy is the floor, not the ceiling

It should name which tools are approved, what client data may never be entered into a consumer tool, what verification is required before AI-assisted work reaches a client, and who owns the answer when something goes wrong.

02. Training closes the gap policy alone cannot

A policy nobody has walked through in a training session is a document, not a practice. Pair every policy update with a session that uses real, close-to-home examples, not hypothetical ones.

03. Brokers are not required to demonstrate AI literacy, and arguably should be

Unlike agents, who face at least some continuing education exposure to the topic, most brokers face no requirement whatsoever to demonstrate an understanding of how AI systems work before deciding how, or whether, to introduce them into their company. A broker who does not understand what a tool does, what data it retains, or how it makes decisions is poorly positioned to write a policy governing its use, let alone to supervise agents using it. Testing broker comprehension, not merely broker awareness that AI exists, would be a meaningful improvement over the current standard, which is effectively none.

04. Independent contractor classification anxiety has become an excuse for having no standard at all

Many brokers who might otherwise write a reasonable AI-use policy decline to do so out of a generalized fear that any operational standard risks reclassifying their independent contractor agents as employees. That fear is frequently overstated and is being used as a reason to avoid governance entirely, rather than a reason to draft policy carefully. NAR and other industry organizations have already published usable AI policy templates, and a well-constructed AI-use policy, drafted with the same care already applied to other brokerage standards, does not by itself convert an independent contractor relationship into an employment relationship. The absence of any policy is not a neutral, risk-free position. It is its own form of exposure.

05. Brokers are deploying tools they never configure

A meaningful share of broker AI exposure is self-inflicted through simple disinterest. Many of the AI tools agents use daily, generative features inside Google Gemini, Microsoft 365 Copilot, and similar productivity suites, were provided to agents by the brokerage itself as part of a standard technology stack. Having provisioned the tool, the broker frequently never returns to configure its safeguards, define acceptable use, or limit its scope. The brokerage is not merely failing to police outside tools agents brought in on their own. It is failing to govern the tools it handed agents directly.

RISK NOTE. If the brokerage issued the tool, the brokerage owns the first configuration decision. That is a lower bar than writing a comprehensive AI policy from scratch, and there is no defensible reason for it to go undone.

06. A hard deadline for having a policy at all

Given the preceding four points, the most direct regulatory intervention available is procedural rather than substantive: require every brokerage to maintain specific, written AI-use policy in its policy manual, effective immediately for new licensure and no later than the firm's next relicensing cycle for existing brokerages. This does not require regulators to solve every open question in this report. It requires brokers to have engaged with the questions in writing, on the record, before the next renewal deadline arrives.

AGENT-FACING DISCLAIMER LANGUAGE

For any AI-use policy or training handout

An internal working tool for licensed real estate professionals. Not a substitute for legal, brokerage, or compliance review. Adapt to your state's rules and your brokerage's policies.

SECTION 08

Intellectual Property & Content Rights

Lower probability than the other categories here, but rising fast as agents lean on AI for listing visuals and marketing copy.

01.

Generated images can carry someone else's fingerprints

Image generators trained on copyrighted photography, art, or branded material can output content that echoes a real photographer's style, a real building's protected design, or a real person's likeness without permission. "The AI made it" is not a defense against an infringement or right-of-publicity claim.

02.

Generated text can mirror a source too closely

Marketing copy that leans on an AI tool to "summarize this listing description from the last agent's flyer" can cross from inspiration into reproduction faster than it feels like it should. Treat AI-drafted copy the way you would treat a human intern's first draft, something to rewrite in your own voice, not to publish as-is.

PRACTICAL FLOOR. Never generate a likeness of a real, identifiable person without consent, never ask an image tool to imitate a specific living photographer's or artist's style by name, and keep the original unedited listing photos on file the way California law already requires.

The Regulatory Wildcard: Preemption and Volatility

Do not build a compliance program around any single statute surviving unchanged. Build it around the principle underneath the statute.

A December 2025 federal executive order directs federal agencies to challenge conflicting state AI laws and push toward a national framework. The Department of Justice has already intervened, for the first time, in a state AI law challenge, the X.AI suit against Colorado's original AI Act, which resulted in a federal court staying enforcement in April 2026 before Colorado replaced the law with a narrower version delayed to January 1, 2027.

The pattern to watch: state AI statutes are getting passed, then vetoed, delayed, sued over, stayed, and rewritten, sometimes all within the same year. Virginia's own comprehensive AI bill never became law at all. Meanwhile the EU AI Act's high-risk obligations, originally set for 2026, are being pushed toward 2027 through a provisional agreement that is not yet binding.

01.

Very little guidance exists, and what exists rarely touches daily practice

The volume of government guidance, regulatory rulemaking, and binding case law addressing AI use in ordinary business operations remains genuinely sparse. Where courts have ruled on AI-related disputes, the cases have concentrated overwhelmingly on intellectual property, large-scale copyright infringement claims arising from the bulk recording, scraping, or reproduction of books, articles, images, and video used to train models. That body of law says almost nothing about how an individual agent should use AI to draft a listing description or respond to a client inquiry. The profession is operating in a genuine guidance vacuum, not a lightly regulated space.

02.

Existing rules assume coverage they were never written to provide

Most current real estate laws, rules, and regulations were not drafted with any anticipation of AI-assisted practice, and the prevailing regulatory posture has been to assume existing rules apply by extension rather than to test that assumption deliberately. This produces a system

that appears to regulate the behavior in question while doing very little to actually constrain or clarify it, protection by assumption rather than protection by design.

Translation for a working brokerage: the statutes are the weather. Fair Housing, data stewardship, honest advertising, and documented verification are the climate. Build to the climate.

The Misinformation Cascade & Consumer Trust Gap

Two related problems. Bad AI content gets laundered into apparent fact, and consumers are increasingly right to trust the AI more than they trust you.

01. AI citing AI: the misinformation cascade

Consumers are turning up AI-generated blog posts on law firm and brokerage websites that contain outright errors, contingency timelines, disclosure rules, statutory citations that are wrong or invented. The problem does not stay contained. A separate AI assistant later crawls that same page, treats it as a published, seemingly authoritative source, and repeats the error back to a consumer with the appearance of independent confirmation. The mistake gets laundered through citation. Real estate content is a soft target for this because so much of it sounds procedural and authoritative and is cheap to mass-produce at scale.

RISK NOTE. A consumer who acts on a hallucinated deadline or disclosure rule sourced this way may reasonably expect their agent to have caught it, even though the agent never wrote, saw, or was consulted on the content.

02. Consumers trust the AI more than the agent, and the AI is often right

This deserves more candor than the industry usually gives it. Consumers increasingly treat an AI assistant's answer as more reliable than their agent's, right up until the AI is proven wrong in a specific instance. The uncomfortable part worth saying plainly in training: on a lot of factual, non-relationship-dependent questions, a code citation, a general market trend, a definition, the AI answer is frequently more accurate than what an under-informed or out-of-date agent would say off the cuff. The trust gap is not irrational. It is often earned by comparison.

RISK NOTE. The correct response is not defensiveness. An agent's differentiated value increasingly is not "I know things AI does not." It is judgment, negotiation, local relationship access, and accountability, since AI carries none of the last one. That only works if the agent is also verifying against current sources as a habit, not just recommending the habit to clients.

Agentic AI Misuse & the Regulatory Capacity Gap

Agentic tools can now complete multi-step work unattended. What happens when a licensee points one at a legal obligation instead of a marketing task, and no one is positioned to notice.

01. Unsupervised agentic AI completing CE and license-required work

Agentic tools capable of navigating a browser and completing multi-step online tasks, including tools like Claude Cowork, can complete continuing education coursework, license renewal modules, or other online compliance work on a licensee's behalf, unattended and undetected by the CE provider, the association, or the brokerage. This is materially different from using AI to help study or draft a note. It is a licensee not doing the work a state has decided a licensee must do, while a completion certificate says otherwise. Every CE hour is a proxy for a competency the state requires. Faking the proxy does not just create disciplinary exposure if discovered, it puts an under-informed licensee back into practice with clients who reasonably assume the credential means something.

RISK NOTE. Most CE providers and licensing boards have no technical control that distinguishes a human completing a module from an agentic tool completing it unattended. The exposure is not hypothetical. It is a detection gap, and detection gaps do not stay gaps forever.

02.

Regulators are not equipped to see any of this

The people responsible for catching the above are, in most states, working from license law written before generative AI existed, on internal case-management systems that predate it, without agency-level investment in AI tools, and without budgeted time for investigators to learn what these tools can do, let alone how a specific licensee might be misusing one. This is not a criticism of individual investigators. It is a resourcing and legislative-lag problem playing out across state real estate regulatory agencies nationally. The practical effect is that the profession is currently operating on an honor system for a category of misconduct almost no one is positioned to detect. State procurement and policy officials, often limit the use of AI or provide subpar tools to state employees which limits the exposure that state regulators would have to identify and fully understand how the agent or broker are utilizing a variety of AI resources.

RISK NOTE. A regulatory gap does not reduce a licensee's actual obligation. It delays the reckoning. The first real enforcement action against this kind of misuse, whenever it comes, will likely be made an example of.

03.

Why the incentive to fake CE exists in the first place

The temptation to let an agentic tool sit through required coursework is not only about laziness. It is a rational response to a system that gives agents little reason to engage. Most state CE approval processes evaluate a course on its content outline, not on whether it produces retention, application, or any measurable change in practice, and self-paced online CE in particular is rarely built around adult learning principles: relevance to the learner's actual work, spaced repetition, active application, or meaningful assessment beyond a low-stakes quiz. An agent renewing the same foundational material year over year, in a format that assumes no prior knowledge and asks for no real engagement, has little incentive to sit through it attentively, let alone at all. A likely result is that the agent will use a well drafted prompt on an AI Agent to then perform the task of completing CE online at a level that would be undetectable to any current self-paced learning product.

If CE is to be considered a series level of education for minimum competency then in-person live training is the only current learning methodology that would be out of reach of AI automation.

RISK NOTE. This is a design problem regulators and CE providers created, not an excuse for agents to offload the requirement to an unsupervised tool. But any conversation about agentic CE misuse that skips this cause will keep treating the symptom instead of the disease.

04. The same agentic capability can fabricate retroactive compliance records

Several state commissions permit self-reporting on certain matters and require licensees to retain supporting files for a period of years. An agent or broker facing a commission inquiry could, in principle, prompt an AI system with a single short instruction to reverse-engineer a plausible, internally consistent set of meeting minutes, policy documents, or self-reporting records covering the retention period, generated in minutes rather than created contemporaneously as required. The output can be sophisticated enough to withstand casual review, which is precisely what makes it dangerous: it exploits the assumption underlying self-reporting regimes, that retained files were produced in good faith at the time, not manufactured afterward to answer an investigation.

RISK NOTE. This is not a hypothetical scenario invented for effect. It is a foreseeable use of a capability that already exists and is already accessible to anyone with a basic subscription.

05. A proficiency requirement, modeled on how states already regulate assistants

Most states impose at least some regulatory structure on an agent's use of human virtual or in-person assistants. No comparable structure currently exists for AI. A more coherent long-term approach would require agents to demonstrate proficiency in setting up and using AI tools correctly, prompt construction, source verification, appropriate scope of delegation, as a component of both initial licensing and continuing education. Given that a meaningful share of AI error traces back to poor setup rather than the technology itself, this requirement would plausibly reduce the very hallucination and misuse risk documented throughout this report, rather than simply adding another box to check.

BROKERAGE POLICY LANGUAGE

Name what cannot be delegated

None of this argues against agentic AI tools, which are legitimate productivity tools for the work they are meant to do. It argues for a brokerage AI-use policy that names what can never be delegated to an unsupervised agent: continuing education, license renewal coursework, examinations, and anything requiring a personal attestation of completion. Say it in the policy once, in writing, rather than assuming it goes without saying.

Algorithmic Background Screening & Identity-Matching Errors

False information in a listing description is a training-deck risk. This one is where agents are already making exclusion decisions about real people, often without knowing it. Marketing copy is the visible, narrow slice of AI exposure. This is closer to where the real weight sits.

01. Bulk-data screening tools treated as verified fact

Tools like Websluth and Forewarn compile information from bulk public-record and third-party data sources and present it as a screening report. Agents frequently treat that output as verified fact rather than what it actually is: a probabilistic match built on names, approximate birth years, and address history, run against records never designed to serve as definitive identity verification. The report reads like a background check. It functions more like a fuzzy search. These tools are important for agent safety and provide a needed service in the industry, but agents must be trained to be discerning in their use and practical in their approach. Some agents are using these tools for marketing purposes and applying an AI automation to segment potential clients, without discretion or investigation could lead to further disparate impact on consumers. In some cases this negative impact may be solely because a consumer shares the same name and date of birth range.

RISK NOTE. A report is not a fact until a human confirms the identity behind it. Treating it as gospel is the failure mode, not the existence of the tool.

02. Automating the report turns a research aid into an exclusion engine

When an agent runs those same tools through an AI assistant on an automated basis, screening every prospective buyer, tenant, or lead without individually reviewing the match, a name-and-approximate-birth-year collision can silently disqualify someone who never did anything the report alleges. The person flagged for an alleged criminal history may simply share a name and birth year with someone who has one. Automating the lookup does not automate the judgment required to catch that. It removes the moment a human would have caught it.

RISK NOTE. An agent who never meets the person being screened, and never applies independent judgment to a flagged match, has effectively let a bulk-data algorithm make a housing decision. That is the same theory HUD has used against automated tenant screening tools since 2019, applied here to something the agent chose to run.

03. This is a Fair Housing problem, not just an accuracy problem

A false-positive match that disproportionately affects people with common names, which itself correlates with race, ethnicity, and national origin in measurable ways, is not a neutral error. It is exactly the kind of facially neutral process that can produce disparate impact liability without any intent to discriminate. An agent who would never consciously deny someone housing based on their name is still accountable for what an automated screening process did in their name. Improper use of AI, but poorly trained agents will be a contributing factor to access to service and representation in the future.

Algorithmic Buyer Targeting in Private & Off-MLS Listings

As pocket listings and pre-MLS marketing grow, who gets shown a home is increasingly an algorithm's decision before it is an agent's.

01. AI-curated buyer lists for exclusive and coming-soon listings

As non-MLS, pocket, and coming-soon listing strategies grow, more of a property's actual buyer exposure happens before it ever reaches the open market, often through an agent-curated or AI-curated list of "likely" buyers. When that curation is done by an algorithm trained to predict who is statistically likely to buy in a given neighborhood or price point, it can quietly reproduce the exact demographic patterns Fair Housing law exists to interrupt, without anyone deciding to discriminate.

RISK NOTE. Dual agency compounds this. An agent representing both sides of a private-exposure transaction has both the access and the incentive to keep the buyer pool small, and an algorithm narrowing that pool on the agent's behalf makes the exclusion easier not to notice.

02.

The bias-mitigation guardrails in some AI tools may be narrower than agents assume

Worth being precise about a claim that circulates in training conversations. A July 2025 federal executive order, "Preventing Woke AI in the Federal Government" (EO 14319), requires federal agencies to procure only large language models that meet "Unbiased AI Principles," and it names DEI-related concepts among the things it treats as ideological bias to be avoided in federally procured tools. That order governs federal government procurement. It does not directly regulate the AI tools a brokerage licenses for its own marketing or lead-scoring use. The more plausible, still unproven concern is a downstream one: AI vendors competing for federal contracts have a commercial incentive to build models that satisfy EO 14319 broadly, not only in a federal-facing version, and some of those same base models power the consumer and enterprise tools brokerages use. Whether that produces measurably less bias-mitigation in a real estate lead-scoring tool is a theory worth watching as of this report, not a documented fact.

RISK NOTE. Ask any vendor providing buyer-matching, lead-scoring, or listing-exposure tools directly whether and how their system tests for disparate impact across protected classes, and get the answer in writing. Do not assume the tool has fairness guardrails built in, and do not assume it does not. Verify.

PART II

Structural and Systemic Risk

The thirteen categories in Part I are symptoms. The three sections that follow address why they are more likely to occur, and less likely to be caught, than they should be. These are not tabulated in the risk matrix because they describe governance conditions rather than discrete transactional events.



Wearable AI & the Physical Transaction Environment

The next frontier of AI exposure in real estate is not on a screen. It walks into the property with the buyer.

01. Ambient AI capturing property data no one agreed to release

A buyer touring a property while wearing AI-enabled glasses or a similar ambient device may use that device, deliberately or as a background function, to capture images, ask the embedded assistant about a well or septic system's service history, or query the make and age of a visible appliance. The device is, in effect, introducing an additional party into the showing that neither the seller, the listing agent, the buyer's agent, nor the seller's own service providers agreed to admit. Whatever the device captures, records, and potentially transmits to a third-party AI system is data about a private property that its owner never authorized anyone to collect.

RISK NOTE. This is not a hypothetical for the next decade. Consumer AI-enabled wearables with camera and voice-query capability are already commercially available and already being carried into showings.

02. Recording and wiretap consent law was not written with AI in mind, but it still applies

State recording and wiretap statutes vary considerably, some require only one party's consent to record a conversation, others require the consent of everyone present, and many of these statutes are decades old. They were not drafted with AI transcription or ambient wearable capture in view. That does not exempt AI-based recording from their reach. An agent using an AI notetaker or a wearable device to transcribe a showing, a listing consultation, or any consumer engagement in a two-party consent state is bound by the same informed-consent requirement that would apply to a handheld recorder, whether or not the agent has drawn that connection for themselves.

RISK NOTE. An agent who believes AI transcription occupies some undefined exception to existing recording law is operating on an assumption, not a legal conclusion. Verify your state's consent requirement and disclose the AI tool's presence the same way you would disclose any other recording device.

Contract Review & the Scope-of-Practice Blind Spot

AI contract review answers a narrower question than the one most brokers think they are asking.

01. AI can validate language. It cannot validate the agent

Brokers frequently ask how AI can assist with contract review, and the premise is reasonable on its face. In practice, an AI system reviewing a contract is generally checking whether non-standard language is present, whether required initials and signatures appear where expected, and whether the language used is facially enforceable. That is a meaningfully narrower question than whether the agent who drafted or modified the contract was operating within their scope of expertise. An agent can insert contradictory or client-disadvantaging language with enough internal sophistication that it reads as deliberate, as though an attorney had constructed it to quietly undercut a weaker provision elsewhere in the document, when in fact the agent simply drafted carelessly or overreached beyond their competence. AI contract review, as currently deployed, will often validate that language as enforceable without ever surfacing that it works against the client's interest.

RISK NOTE. A broker who relies on AI contract review as a substitute for actually knowing their agents' individual competence and habits has not solved the supervision problem. They have automated the part of it that was never the hard part to begin with.

Toward Accountability-Based Regulation

Prohibition is not a viable regulatory strategy for a technology this deeply embedded. Accountability for output is.

Attempting to restrict or prohibit agent access to AI tools is very likely to fail as a matter of practical enforcement, and this report has made that argument more than once with respect to individual risk categories. It is worth stating as a general proposition here. Artificial intelligence is now integrated into phones, computers, productivity software, and increasingly into consumer vehicles. There is no realistic mechanism by which a state regulator could isolate and prohibit a licensee's access to AI-assisted drafting, research, or analysis, and an enforcement strategy built on that premise would be attempting to hold back a condition of the current technological environment rather than regulating conduct within it.

01. Regulate the output, not the tool

A more workable regulatory target is holding the agent unambiguously responsible for anything that leaves their keyboard or their name, regardless of whether the underlying draft originated with the agent personally, a friend, a virtual assistant, an in-person assistant, or an AI system. This is not a novel legal theory. It is the same accountability standard that already applies, at least in principle, to work product an agent delegates to a human assistant. The most direct way to extend that standard to AI is to treat AI-assisted work exactly as if it had been produced by an unlicensed personal assistant, work the agent must review, correct, and stand behind before it reaches a client or the public.

RISK NOTE. The honest caveat here is that most states do not currently have particularly robust rules governing agents' use of human virtual or personal assistants either. Extending a weak existing framework to a new technology does not by itself produce a strong one. The recommendation is to build both frameworks deliberately, not to treat the analogy as a finished solution.

02. The demands on agents keep expanding, and AI dependence is a downstream effect

Agents are increasingly expected to function simultaneously as marketing specialists, small-business operators, technologists, contract drafters, and, informally and without the credential, quasi-attorneys, quasi-accountants, and quasi-financial advisors. That accumulation of expectation is not primarily a technology story, but it is the context within which AI dependence has grown so quickly. An agent facing that scope of responsibility has little practical choice but to lean on AI tools or human assistance at a level well beyond historical norms simply to remain functional. Any regulatory or brokerage response that treats rising AI use as a discipline problem, rather than a predictable response to rising role demands, is diagnosing the symptom and missing the cause.

THE THESIS OF THIS SECTION

Authorize, then hold accountable

The affirmative regulatory step this report recommends is not restriction. It is clear, explicit authorization for agents and brokers to use modern AI technology in their practice, paired with an equally clear and enforceable standard holding them accountable for the output of that technology exactly as they would be held accountable for the output of a staff member or personal assistant. Ambiguity is the current default. Ambiguity protects no one, least of all the consumer this framework exists to serve.

PART III

The Code of Ethics as a Regulatory Instrument

Of every governance mechanism already available to the profession, the Code of Ethics is the one closest to being usable today, with the narrowest set of amendments.



SECTION 17

Where the Code of Ethics Already Reaches This

None of the sixteen sections above sit outside the Code's existing language. AI is new. The duties are not.

The word "AI" does not appear anywhere in the 2026 Code of Ethics and Standards of Practice. That does not mean the Code is silent on any of this. Four Articles already impose obligations that reach straight through to AI-generated work product, they just were not written with a chatbot in mind. Sections 10 through 16, covering the misinformation cascade, agentic CE misuse and record fabrication, automated background screening, algorithmic buyer targeting, wearable AI, contract review, and the accountability framework proposed in Section 16, sit further outside the Code's current reach and are addressed separately below.

ARTICLE / SOP	EXISTING OBLIGATION	HOW IT ALREADY REACHES AI
Article 1, SoP 1-9 (Confidentiality)	Client confidential information may not be revealed or used for the REALTOR's or a third party's advantage without consent.	Pasting client financials, deal terms, or PII into a consumer AI tool arguably discloses that information to a third party, the AI vendor, whose retention and use of it is outside the REALTOR's control.
Article 2 (Misrepresentation)	REALTORS shall avoid exaggeration, misrepresentation, or concealment of pertinent facts.	A hallucinated comp, square footage, or code citation that reaches a client is a straightforward Article 2 problem. Authorship by a model does not change whose name is on it.

ARTICLE / SOP	EXISTING OBLIGATION	HOW IT ALREADY REACHES AI
Article 10, SoP 10-1 & 10-3 (Non-discrimination)	No discriminatory statements, advertising, or steering, in any medium.	AI-generated marketing copy, AI-optimized ad delivery, automated screening, and algorithmic buyer targeting that steer on a protected characteristic all violate Article 10 exactly as a human decision would. The Article does not carve out automated decisions.
Article 11, SoP 11-1 (Competence, opinions of value)	Opinions of value require property-type knowledge, access to the right information, area familiarity, and disclosed limitations.	An AI-assisted CMA presented as the REALTOR's professional opinion, without the underlying diligence 11-1 requires or without disclosed limitations, fails on its face regardless of the tool used to draft it.
Article 12, SoP 12-8 & 12-10 (True picture)	Advertising and website content must be current and accurate. Manipulating content, using others' work without attribution, and misleading images are explicitly prohibited.	This is the closest the Code already comes to naming the problem directly. An unlabeled AI-altered photo or AI-drafted description presented as human-original is arguably already a 12-10 violation under the current text.

Read this way, four of the sixteen substantive sections in this report already have a direct textual hook: Fair Housing, steering, screening, and buyer targeting (Article 10), accuracy and hallucination (Articles 2 and 11), disclosure and true picture (Article 12), and client confidentiality (Article 1). What is missing for those four is not obligation, it is specificity. Enforcement currently depends on a hearing panel reading these Articles expansively enough to cover a tool that did not exist when the language was last amended, rather than on language written with that tool in view. The misinformation cascade and agentic CE misuse in Sections 10 and 11 are a different problem: less a gap in the Code's text than a gap in what any Board is currently positioned to detect.

Recommended Modifications to the Code of Ethics

Five amendments would close the specificity gap without touching the Code's structure, philosophy, or enforcement mechanism. Each slots under an Article that already carries the underlying duty.

01. New Standard of Practice under Article 2 (Misrepresentation)

Proposed text: "REALTORS shall not present information generated by artificial intelligence tools to clients, customers, or the public without independently verifying its accuracy against reliable source data. Reliance on an artificial intelligence tool's output does not relieve a REALTOR of the obligations imposed by Article 2." This closes off "the AI said it, not me" before a hearing panel has to invent the answer on its own.

02. New Standard of Practice under Article 11 (Competence)

Proposed text: "When REALTORS use artificial intelligence tools to assist in preparing opinions of value, comparative market analyses, or other real property valuations, the requirements of Standard of Practice 11-1 continue to apply in full. REALTORS shall disclose the use of artificial intelligence in preparing such opinions upon request, and shall not represent an AI-assisted valuation as the product of independent professional judgment unless that judgment was in fact exercised." This targets the AI-assisted CMA risk in Section 03 directly.

03. Amendment adding a sixth item to Standard of Practice 12-10

SoP 12-10 already lists five prohibited internet and advertising practices: deceptive framing, manipulating others' content, deceptive metatags, using others' content without attribution, and misleading images. *Proposed sixth item: "6) presenting AI-generated or AI-altered images, video, or written content as unaltered or human-original without clear and readily apparent disclosure of the use of artificial intelligence."* This is the most surgical fix available, and it elevates the labeling standard most MLS systems already enforce locally to a national Code obligation.

04. New interpretive note under Standard of Practice 1-9 (Confidentiality)

Proposed text: "For purposes of this Standard, providing a client's confidential information to a third-party artificial intelligence tool or platform, where that information may be retained, processed, or used by the platform beyond the REALTOR's control, constitutes use of confidential information for the advantage of a third party under Article 1, absent the client's informed consent." This gives the consumer-AI-tool data risk in Section 02 a name it does not currently have.

05. New Standard of Practice under Article 10 (Non-discrimination)

Proposed text: "The obligations of Article 10 apply regardless of whether discriminatory statements, advertising, steering, screening, or buyer or tenant targeting result from a REALTOR's direct decision or from an artificial intelligence tool, automated system, or algorithm employed by or on behalf of the REALTOR. REALTORS remain responsible for reviewing AI-generated or AI-assisted advertising, marketing, screening, and lead-targeting practices for compliance with this Article before use." This removes any ambiguity that an algorithm's decision is somehow not the REALTOR's decision, which is precisely the theory HUD has used in enforcement since 2019, and it is written broadly enough to reach the screening and buyer-targeting risk in Sections 12 and 13 as well as the advertising risk in Section 01.

Sections 10 and 11 are deliberately left out of the amendment list above for a different reason than Sections 12 and 13, which the broadened Article 10 proposal already covers. The misinformation cascade is a consumer-education and marketing-literacy problem, not a Code violation. Agentic misuse of continuing education is a state licensing law and CE-provider integrity issue, not an ethics issue the Code of Ethics governs. DPOR and equivalent state agencies, not NAR's Professional Standards Committee, would need to close that gap through statute, regulation, or CE-provider technical controls.

WORTH FLAGGING

These are proposed drafting, not NAR policy

The five items above are proposed language for discussion in training, brokerage policy conversations, and association-level advocacy. They have not been reviewed, drafted, or adopted by the National Association of REALTORS Professional Standards Committee and should not be represented as official NAR guidance. Treat them as a starting point for the conversation your local Board or state association would need to have, not as a citable amendment.

None of these five require a new Article, a new enforcement body, or a philosophical shift in what the Code already asks of REALTORS. Each one takes an obligation the Code already imposes and names the technology specifically enough that a hearing panel, a new licensee, and an AI vendor's terms-of-service reviewer could all read it the same way.

ONE LAST THING

The tool changed. The obligations did not.

Every risk in this report existed before generative AI showed up. Fair Housing liability, professional negligence, wire fraud, vendor due diligence, brokerage supervision. AI did not create new categories of exposure so much as it sped up the ones that were already there and put them in the hands of every agent at once, whether or not the brokerage signed off.

That is the actual argument for a written policy, regular training, and a verification habit that does not bend under deadline pressure. Not fear of the technology. Respect for what happens when speed outruns judgment.

Two rules. Verify before it reaches a client. Document that you did.

An educational reference for licensed real estate professionals and instructors. The practices described are general guidance and are not a substitute for legal, brokerage, or compliance review. State, MLS, and brokerage rules vary and change quickly. Verify current statutory status before relying on any specific citation in this report, particularly in Section 04 and Section 09.

APPENDIX A

Current State Licensing Law & AI-Specific Regulation

A jurisdiction-by-jurisdiction summary of enacted and proposed AI-specific rules bearing directly on real estate licensees, current as of August 2026.

The regulatory record on AI in real estate practice remains thin, which is itself a finding. As of the date of this report, only one state, California, has enacted a statute directly regulating AI use by real estate brokers and salespersons in the context of real-property advertising. One additional state, Colorado, has enacted a cross-industry AI statute with direct implications for residential housing decisions. Two states have considered but not enacted listing-image proposals. Everything else in this appendix is either agency guidance, general-application law that reaches brokerages incidentally, or proposals that should not be treated as binding until enacted and effective.

JURISDICTION / INSTRUMENT	STATUS	REACH INTO REAL ESTATE PRACTICE
California, Bus. & Prof. Code § 10140.8 (Real Estate Digitally Altered Images Disclosures, enacted via AB 723)	In effect Jan. 1, 2026	The only enacted state law directly regulating licensee advertising. Requires disclosure of digitally altered listing images and online access to the original.
California DRE Licensee Advisory (March 17, 2026)	Agency guidance	Confirms that AI use does not shield a licensee from existing real estate, fair housing, lending, advertising, or supervision obligations.

JURISDICTION / INSTRUMENT	STATUS	REACH INTO REAL ESTATE PRACTICE
Colorado SB 24-205 as amended by SB 26-189 (Colorado AI Act)	Effective Jan. 1, 2027	Cross-industry law treating residential lease or purchase as a consequential decision. Reaches screening, ranking, and scoring systems, not routine generative use.
California Bot Act, Bus. & Prof. Code §§ 17940 to 17943	In effect	Requires disclosure of automated identity for online bots used to induce a commercial transaction. Reaches lead-generation chatbots.
California AI voice and robocall disclosure rules	In effect	Disclosure and consent requirements for automated calls using artificial voice. Reaches AI voice agents and prerecorded prospecting.
California CCPA automated decision-making technology rules	In effect	Consumer rights around ADMT used in significant decisions including housing. Reaches tenant screening and applicant scoring.
New Jersey Assembly Bill 4728	Proposed, not enacted	Would limit certain AI-generated or edited listing imagery, permit disclosed virtual staging, and impose penalties.
New York Assembly Bill A.11635	Proposed, not enacted	Would require disclosure of AI or digital alteration beyond routine editing and online access to originals.
California Assembly Bill 2025	Proposed, not enacted	Would extend altered-image disclosure to rental-property promotional materials.

A.1. California: direct regulation of listing images

Business and Professions Code section 10140.8, effective January 1, 2026, applies to California real estate brokers, salespersons, and persons acting on their behalf. When a covered party advertises or promotes real property for sale using a digitally altered image, the party must disclose that the image has been altered. For internet-based advertising, the material must also provide a link, URL, or QR code giving the consumer access to the clearly identified original, unaltered image. The requirement reaches AI virtual staging and any alteration that materially changes a consumer's impression of the property, including adding or removing furniture, changing landscaping, altering views, modifying appliances or finishes, or otherwise creating an inaccurate representation.

RISK NOTE. A vendor's use of AI does not shift compliance responsibility away from the broker or licensee. The practical controls are an approval workflow for every altered image, retention of both the original and the published image in the marketing file, disclosure added before publication rather than after a complaint, and training that extends to photographers, marketing vendors, and virtual-staging providers.

A.2. California DRE guidance on AI generally

In a March 17, 2026 licensee advisory, the California Department of Real Estate stated that artificial intelligence does not shield a licensee from responsibility under real estate law, fair housing law, lending law, advertising rules, or supervision obligations. Licensees remain responsible for reviewing AI-supported work, ensuring factual accuracy, avoiding misleading advertising, protecting client information, and complying with fair housing requirements. This is the clearest statement to date from a state licensing regulator that the accountability standard argued for in Section 16 of this report is already the operative expectation, at least in one jurisdiction.

A.3.

Colorado: AI in consequential housing decisions

Colorado Senate Bill 24-205, as modified by the 2026 amendment legislation SB 26-189, takes effect January 1, 2027, subject to implementing rules and statutory timing. It applies to developers and deployers of covered high-risk or automated decision-making systems. Colorado treats the lease or purchase of residential real estate as a consequential decision. When covered AI materially influences that decision, the law requires consumer notice and creates related consumer rights. Where an adverse decision occurs, obligations may include an explanation of the AI system's role, correction rights for relevant personal data, and an opportunity for meaningful human review or reconsideration where commercially reasonable.

RISK NOTE. The primary exposure is not routine generative AI use. It is AI used for tenant screening, application ranking, eligibility scoring, automated housing-access decisions, or setting materially different terms for consumers, which is precisely the risk profile described in Sections 12 and 13. Brokerages with Colorado exposure should inventory every tool that screens, ranks, scores, approves, or denies consumers, obtain vendor documentation on model purpose, training data, limitations, and anti-discrimination safeguards, build documented adverse-action and human-review procedures, and monitor Colorado Attorney General rulemaking through the implementation date.

A.4.

Other California rules a brokerage may encounter

Three general-application California rules reach AI-powered brokerage operations without being real-estate-specific. The California Bot Act, Business and Professions Code sections 17940 through 17943, restricts online bots that mislead consumers about their artificial identity for the purpose of inducing a commercial transaction, so a publicly facing lead-generation chatbot should be clearly disclosed as automated where the statute applies. California's AI voice disclosure requirements govern certain automated calls using an artificial voice, so any AI voice agent or prerecorded prospecting workflow requires confirmation of disclosure, consent, and telemarketing compliance before deployment. California's CCPA automated decision-making technology rules create consumer rights around covered ADMT used in significant decisions, including housing, which is most relevant to tenant screening and applicant scoring rather than ordinary listing-copy generation.

A.5.

Proposed measures that are not yet law

Three proposals should be tracked but not treated as effective obligations. New Jersey Assembly Bill 4728 would govern photos used in real estate listings, limiting certain AI-generated or photo-edited imagery while permitting disclosed virtual staging, and would establish financial penalties. New York Assembly Bill A.11635 would require disclosure when AI or other digital alteration changes real-property imagery beyond routine editing and would require online access to original imagery. California Assembly Bill 2025 would extend altered-image disclosure to rental-property promotional materials. None had been enacted as of August 2026.

RISK NOTE. The pattern across all three is instructive. Every pending state proposal identified here concerns listing imagery, not data privacy, not fraud, not screening, not supervision. The legislative attention is concentrated on the most visible and least consequential slice of the exposure described in this report.

A National Compliance Baseline for Brokerages

A scalable policy floor that does not depend on which state statutes survive, pass, or get struck down next year.

A brokerage policy that waits for additional state statutes will always be behind. The following baseline, drawn from the same source material as the regulatory summary above and consistent with the recommendations throughout this report, reduces exposure across advertising, fair housing, cybersecurity, privacy, and professional liability regardless of jurisdiction.

CONTROL	WHAT IT REQUIRES
Human review before publication	Documented human review and approval before publishing AI-generated listing copy, images, video, social posts, price opinions, or client communications.
No undisclosed material alteration	AI tools may not materially change property features, condition, location, views, or surroundings without clear disclosure and retention of source imagery.
Version-controlled archive	Retain original images, altered images, material AI prompts or vendor instructions, and final approval records in a version-controlled file.
No AI final housing-access decisions	AI may not make final tenant-screening or applicant-ranking decisions without legal review, bias controls, consumer notices, and meaningful human oversight.
No confidential data in unapproved tools	Confidential client data, transaction details, identification documents, financial data, and protected personal information stay out of unapproved public AI tools.
Vendor vetting	Vet AI vendors for privacy commitments, security controls, fair-housing testing, data-use limitations, audit rights, record retention, and indemnification.
Bot and voice disclosure	Clearly disclose customer-facing bots and any artificial or prerecorded voice use where required by law.
Verify every factual claim	Train licensees that AI does not excuse inaccurate advertising, misrepresentation, discriminatory steering, privacy failures, or unauthorized practice of law.

Two features of this baseline deserve emphasis. First, none of it depends on a state having passed an AI-specific statute, which means it is equally defensible in California, Virginia, and a state with no AI legislation on the books at all. Second, every item in it is something a brokerage can implement today with existing tools and existing policy mechanisms. The barrier is not legal uncertainty. It is the decision to act before being required to.

SOURCES FOR APPENDIX A

- California Department of Real Estate. Licensee Advisory: Artificial Intelligence in California Real Estate, Opportunities, Risks, and Compliance Considerations for Licensees. March 17, 2026.
- California Business and Professions Code section 10140.8, Real Estate Digitally Altered Images Disclosures.
- Colorado Attorney General. Colorado Anti-Discrimination in AI Law (ADAI) Rulemaking. coag.gov/ai.
- Orrick AI Law Center. U.S. State AI Law Tracker, All States.
- National Conference of State Legislatures. Artificial Intelligence Legislation Database. July 1, 2026.
- New Jersey Assembly Bill 4728; reporting in Homes.com, New Jersey Moves to Rein In AI Use for Property Listings, June 22, 2026.
- New York Assembly Bill A.11635; reporting in The Post-Journal, Limits on AI in Real Estate Photos Introduced, August 24, 2026.

DISCLOSURE STATEMENT

Authorship, Method, and Professional Limitations

This report was researched and formatted with the assistance of artificial intelligence tools, which were used to organize source material, verify current regulatory and legislative status, structure the document, and support drafting. The concerns, interpretations, policy positions, and recommendations presented throughout this report are my own professional opinion, developed from approximately two decades of experience teaching, advising, and consulting within the real estate industry at the state, national, and international level. Where an AI tool's assistance shaped a specific factual claim, that claim was independently verified against primary or secondary sources before inclusion, and remaining errors are mine.

I am not an attorney, and nothing in this report should be construed as legal advice or as a substitute for review by qualified legal counsel licensed in the relevant jurisdiction. Statutory citations, case references, regulatory status, and the status of pending or enacted legislation described in this report reflect the state of publicly available information as of the date noted on the cover and are subject to change without notice. Readers, including brokers, agents, association staff, and regulators, should independently verify any statute, regulation, case, or Code of Ethics provision referenced here against current primary sources before relying on it, citing it, or acting upon it in any professional, legal, educational, or regulatory context.

This report reflects my judgment as of the date it was prepared. It does not reflect the official position of any brokerage, association, or regulatory body with which I am affiliated, unless separately and expressly adopted by that organization.