

Artificial Intelligence (AI): Risk and Responsibility
by
Deborah Long, Ed.D. 2026© with assistance from ChatGPT



The instructional goal of this course is to enable North Carolina real estate brokers to recognize and manage the legal, ethical, and practical risks created by the use of artificial intelligence in brokerage practice so that brokers can protect consumers' confidential information, promote fair and accurate housing services, prevent misleading communications, and exercise the independent professional judgment required by North Carolina License Law and Commission rules.

Table of Contents

Course Overview.....	1
• Instructional Goal and Objectives	1
CHAPTER 1: What AI Is—and Why Brokers Must Understand It	4
• Introduction	4
• The Development of AI	5
• A Practical Definition of Artificial Intelligence	6
• AI Compared with Automation and Traditional Software	7
• Generative AI, Predictive AI, and Automated Decision-Making	8
• Examples of AI Tools Used in Real Estate	9
• How AI Produces an Answer	10
• Training Data, Prompts, Inputs, and Outputs	11
◦ Examples of Poor, Better, and Excellent Prompts	12
• Hallucinations, Outdated Information, Bias, and Lack of Context	14
◦ Real Estate Brokerage Examples	16
• Why a Convincing Answer Is Not Necessarily an Accurate Answer	20
CHAPTER 2: Protecting Confidential and Sensitive Information	22
• The Broker’s Continuing Responsibility for AI-Assisted Work	22
• Client Confidentiality and Agency Responsibilities	23
• Personally Identifiable and Sensitive Information	24
• Transaction Documents and Financial Information	25
• Risks of Uploading Documents and Client Communications	25
• Data Retention and Possible Reuse of Submitted Information	26
• Reviewing a Tool’s Privacy and Data-Use Terms	26
• De-identifying Information	26
• Client Notice and Permission	27
◦ Sample Client Notice and Permission for AI-Assisted Work	28
• When to Involve the Broker-in-Charge	29
• Firm Policies Governing Approved and Prohibited AI Uses	29
CHAPTER 3: Accuracy, Advertising, and Fair Housing.....	26
• Verifying AI-Generated Statements	26
◦ Evaluating AI-Provided Citations and Sources	27
• Misrepresentation and Omission	28
• AI-Generated Listing Descriptions	28
• Material Facts and Property Characteristics	29
• Market Statistics and Automated Valuations	30

○ Common Examples of Automated Valuation Models	31
• AI-Generated Explanations of Contracts or Legal Requirements	32
• Altered Photographs and Virtual Staging	33
• Disclosure of Digitally Modified Images	34
• Fair Housing Risks in Language, Targeting, Screening, and Recommendations	34
• Algorithmic Bias and Discriminatory Effects	35
• The Broker’s Responsibility for Content Produced by Third-Party Tools	35
 CHAPTER 4: Exercising Professional Judgment—The Broker Is Still Responsible	38
• Appropriate and Inappropriate Reliance on AI	38
• Matters Requiring Official or Authoritative Sources	39
• When AI Crosses into Legal, Tax, Appraisal, Inspection, or Other Specialized Advice	40
• Transparency and Disclosure	41
• Broker-in-Charge Supervision	42
• Correcting AI-Generated Misinformation	43
• Responding When a Consumer Has Relied on an Incorrect AI Communication ...	43
• Brokerage AI Policies and Review Procedures	44
○ Sample AI Policy Manual Language	44
• Application of the PAUSE Consumer-Protection Model	44
• U.S. Court Cases Involving Real Estate Companies and AI	45
○ Cases Involving Real Estate Companies	45
○ Automated Valuation Cases	46
○ Related Housing-Algorithm Cases	48
○ Regulatory Action	54
 Course Conclusion	50
 Author’s Note.....	51
 About the Author	51
 Appendices.....	62
• Appendix A: An AI-Generated Comparative Market Analysis	62
• Appendix B: AI Tool Recommendations for Real Estate Brokers	65
• Appendix C: Additional AI Tools and Resources	68
• Glossary of AI and Real Estate Terms	64

CHAPTER 1: What AI Is—and Why Brokers Must Understand It

INTRODUCTION

Artificial intelligence is rapidly becoming part of everyday real estate practice. Brokers may encounter it in writing and image-generation programs, virtual staging applications, automated valuation models, customer-management systems, advertising platforms, property-search tools, tenant-screening systems, and chatbots. These technologies can produce convincing results in seconds, but they can also generate inaccurate information, conceal bias, alter material property features, expose confidential data, or communicate with consumers in ways the broker never intended. Because AI often operates behind familiar-looking tools, brokers may not fully recognize when it is influencing their work or the decisions of their clients and customers.

A broker cannot delegate professional judgment or legal responsibility to an algorithm. Regardless of how content is created, the broker remains responsible for protecting confidential information, providing accurate and honest communications, complying with fair housing and advertising laws, properly representing material facts, and serving the consumer's best interests. This course will help brokers recognize AI-related risks, evaluate AI-generated information, and determine when an AI-assisted task should be verified, modified, disclosed, escalated, or rejected. The guiding principle is simple: AI may assist the broker, but it cannot assume the broker's duty to protect the public.

A note:

Artificial intelligence has raised broad societal concerns about job displacement, the concentration of power among technology companies, large-scale surveillance, misinformation and deepfakes, copyright and ownership, algorithmic discrimination, cybersecurity, environmental costs, the erosion of human skills and judgment, and the possibility that increasingly capable systems may operate in ways their developers or users cannot fully explain or control. These issues are important, but they extend beyond the limited purpose of this course. This course will not examine the larger social, economic, political, or existential concerns surrounding AI; instead, it will focus specifically on the practical responsibilities of North Carolina real estate brokers and the protection of consumers when AI is used in brokerage practice.

THE DEVELOPMENT OF AI

The development of artificial intelligence began as a formal field of study in the 1950s, when researchers started exploring whether computers could imitate aspects of human

reasoning and problem-solving. Early AI systems followed rules written by programmers and could perform only narrowly defined tasks. Beginning in the 1980s and accelerating in the 2000s, increased computing power, access to enormous quantities of digital data, and advances in machine learning allowed computers to identify patterns and make predictions without being explicitly programmed for every situation. Later developments in neural networks¹ and deep learning led to major improvements in speech recognition, image analysis, translation, and automated decision-making.

The introduction of widely available generative AI systems² in the early 2020s—including ChatGPT, Claude, and Gemini—made it possible for ordinary users to create and analyze text, images, audio, video, and computer code through conversational prompts. AI is now embedded in many familiar consumer and business products, often without users realizing that an AI system is influencing the information, recommendations, or decisions they receive.

A PRACTICAL DEFINITION OF ARTIFICIAL INTELLIGENCE

Artificial intelligence is technology that enables a computer system to perform tasks that ordinarily require some form of human intelligence, such as recognizing patterns, interpreting language, analyzing images, making predictions, recommending actions, or creating new content. Unlike traditional software, which generally follows a fixed set of instructions, many AI systems learn patterns from large collections of data and use those patterns to produce a result. In real estate, AI may be used to generate a listing description, estimate a property's value, virtually stage a room, identify prospective customers, answer consumer questions, screen rental applicants, or detect possible fraud. AI can process


¹ A type of machine-learning system designed loosely around the way networks of neurons in the human brain process information. A neural network contains layers of interconnected computational units that examine data, identify patterns, and adjust the strength of their connections as the system is trained. Although the name is inspired by the brain, a neural network does not think, understand, or exercise judgment as a person does. For example, a neural network may be trained to recognize property features by examining thousands of labeled photographs. Over time, it may learn to distinguish a kitchen from a bathroom, identify a swimming pool, or recognize architectural characteristics. Neural networks are also used in automated valuation models, speech and facial recognition, fraud detection, recommendation systems, virtual staging, and generative AI. Their results depend heavily on the quality and representativeness of the training data; consequently, they may produce inaccurate or biased results when the data is incomplete, unrepresentative, or flawed.

² ² *Generative* means the system can generate—or create—new content in response to a user's instructions. That content may include text, images, audio, video, computer code, or combinations of these. For example, a generative AI system might create a listing description, summarize an inspection report, draft an email to a buyer, virtually furnish an empty room, or answer a question about a real estate transaction. The content may appear original, but it is produced by identifying and recombining patterns learned from training data. The system does not necessarily understand the content or know whether it is accurate.

information quickly and produce remarkably convincing results, but it does not possess human understanding, professional judgment, or legal accountability. Its output must therefore be reviewed and verified by a knowledgeable person before it is used in brokerage practice.

AI COMPARED WITH AUTOMATION AND TRADITIONAL SOFTWARE

Traditional software follows specific rules written by programmers: when a defined event occurs, the program performs a predetermined action. Automation uses this type of software to complete repetitive tasks, such as scheduling an email, sending an appointment reminder, or transferring information from one form to another. Artificial intelligence differs because it can analyze patterns, interpret less-structured information, make predictions, or generate content rather than merely follow a fixed sequence of instructions. For example, an automated system might send every buyer the same prewritten reminder before a showing, while an AI system might compose a different message for each buyer based on information in the buyer's file. The distinction is not always obvious because many modern products combine traditional software, automation, and AI. For brokers, the critical question is not simply what the product is called, but whether it is interpreting information, influencing a decision, or creating content that requires human review.



GENERATIVE AI, PREDICTIVE AI, AND AUTOMATED DECISION-MAKING

Generative AI creates new content in response to a user's instructions. Depending on the system, that content may include text, images, audio, video, computer code, or combinations of these. In real estate, generative AI can draft listing descriptions, emails, social media posts, property summaries, and responses to consumer questions. It can also produce virtually staged rooms or digitally altered property images.

Although the output may appear polished and authoritative, generative AI does not necessarily understand the subject or verify the accuracy of what it creates. It may invent facts, omit important qualifications, use discriminatory language, alter material property features, or disclose information supplied in a prompt. Brokers must therefore review and verify AI-generated content before using it.


Predictive AI uses patterns in historical and current data to estimate what is likely to happen. Real estate applications may predict a property's value, identify consumers considered likely to buy or sell, estimate the likelihood that a tenant will pay rent, recommend properties, rank prospective customers, or forecast market activity.

A prediction is not a fact or a guarantee; it is an estimate based on the data, assumptions, and statistical methods used by the system. If the underlying data is inaccurate, incomplete, outdated, or shaped by past discrimination, the prediction may be unreliable or unfair. Brokers should understand what a prediction represents, avoid presenting it as a certainty, and consider whether relying on it could disadvantage a consumer or protected class.

An **automated decision-making system** uses technology to make—or substantially influence—a decision with limited human involvement. Such a system might determine which consumers see a housing advertisement, rank rental applicants, flag a transaction as suspicious, select properties to recommend, or decide which prospective clients receive a response. Some automated decision-making systems use predictive AI, while others apply fixed rules. The defining feature is not the particular technology but the system's role in the decision. When an automated result determines who receives information, an opportunity, or a service, errors and bias can cause significant consumer harm.

These categories frequently overlap. A system might use predictive AI to rank prospective tenants, an automated process to eliminate applicants below a selected score, and generative AI to compose rejection messages. The presence of AI does not transfer responsibility from the broker or brokerage firm to the technology provider. Brokers should

know when these systems are being used, examine their results for accuracy and bias, provide meaningful human oversight, and be prepared to modify or reject an **automated result when necessary to protect consumers and comply with the law.**



EXAMPLES OF AI TOOLS USED IN REAL ESTATE

Real estate brokers may encounter artificial intelligence in both stand-alone applications and technology they already use. Generative AI assistants such as ChatGPT, Claude, and Gemini can draft or revise listing descriptions, emails, advertisements, social media posts, property summaries, and responses to consumer questions. Image-generation and virtual-staging programs can furnish an empty room, remove personal belongings, change décor, or illustrate a proposed renovation. AI may also be incorporated into customer relationship management systems, transaction-management platforms, chatbots, property-search applications, photography tools, and advertising systems. Because AI features are increasingly embedded within familiar products, brokers may use AI without realizing that a system is generating content, making a recommendation, or selecting which consumers receive information.

Other AI tools analyze data, make predictions, or influence decisions. Automated valuation models estimate property values; recommendation systems select properties for consumers to view; lead-scoring programs rank prospective customers; and advertising

platforms determine who is likely to see a housing advertisement. AI may also be used in mortgage underwriting, tenant screening, fraud detection, document review, risk assessment, and identity verification. Some of these systems operate outside the brokerage firm but still affect the broker's clients and customers. Brokers should therefore consider not only the AI tools they personally select, but also automated systems that may influence a consumer's access to information, services, financing, or housing opportunities.³


HOW AI PRODUCES AN ANSWER

A generative AI system produces an answer by analyzing the user's prompt and predicting, step by step, what content is most likely to follow. A large language model has been trained on patterns found in enormous collections of text and other data. It breaks language into small units called *tokens*, examines the relationships among those tokens, and calculates which token is likely to come next. It repeats this process rapidly until it has constructed a complete response. The system also considers the instructions, examples, documents, and other information available within the current conversation. This allows it to produce an answer that appears responsive, organized, and conversational.

This process is different from a person researching a question, understanding the facts, and forming a professional opinion. Unless the AI is connected to a reliable search or retrieval tool and actually uses it, the system may be generating a plausible answer from learned language patterns rather than locating verified information in an authoritative source. Even when an AI system retrieves outside information, it may misinterpret, omit, or inaccurately summarize what it finds. Because more than one response may be statistically plausible, the same prompt can also produce different answers at different times or in different AI systems. A polished response is therefore not proof of accuracy: the broker must verify material information and apply independent professional judgment before relying on or sharing an AI-generated answer.

³ See appendix for a list of AI resources for brokers.

HOW AI PRODUCES AN ANSWER



TRAINING DATA, PROMPTS, INPUTS, AND OUTPUTS

Training data is the information used to teach an AI model to recognize patterns and produce results. Depending on the type of system, training data may include books, websites, articles, photographs, recordings, public records, transaction data, consumer behavior, or other digital information. During training, the system identifies relationships and recurring patterns within the data. It then uses those patterns to generate content, make predictions, or classify new information. The quality and representativeness of the training data greatly influence the system's performance. If the data is inaccurate, incomplete, outdated, or reflects historical discrimination, the AI's results may reproduce or amplify those problems.

A *prompt* is the question, instruction, or request given to an AI system. A prompt may be as simple as "Write a listing description," or it may provide detailed facts, limitations, examples, intended audience, and formatting instructions. More specific prompts usually give the system better direction, but even a carefully written prompt cannot guarantee a factual, lawful, or appropriate result. A prompt also becomes part of the information submitted to the system. Brokers should therefore avoid including confidential client information, personal identifying information, financial records, account numbers, or identifiable transaction details unless the system has been approved for that purpose and adequate safeguards are in place.

Inputs include everything the AI system receives while performing a task. The prompt is one form of input, but inputs may also include uploaded contracts, inspection reports,

photographs, spreadsheets, audio recordings, property data, prior messages in the conversation, or information retrieved from another source. Inputs provide the context the system uses to produce its response. Before submitting information, a broker should consider whether the system needs it, whether the broker has authority to provide it, how the provider may store or use it, and whether identifying details can be removed.

An *output* is the result produced by the AI system. Outputs may include text, images, summaries, valuations, scores, rankings, recommendations, predictions, or automated actions. An output is influenced by the system's training data, the user's inputs, the wording of the prompt, and the model's design. It may appear polished and convincing even when it contains an error, unsupported assumption, discriminatory pattern, or fabricated information. The process can be summarized as follows: training data shapes the model; prompts and other inputs direct the task; and the model produces an output that requires appropriate human review. The broker remains responsible for deciding whether that output is accurate, lawful, relevant, and safe to use.

Examples Of Poor And Good Prompts

The difference is not simply “short prompt versus long prompt.” An excellent prompt gives the AI enough verified information to perform the task, clearly defines the desired result, and establishes safeguards against invention, misrepresentation, and fair housing problems. Official OpenAI guidance similarly recommends stating the desired outcome, success criteria, relevant context, constraints, and output format without burying the task in unnecessary instructions.

Assume the fictional property has these verified features:

- Three bedrooms and two bathrooms
- 1,850 square feet
- Built in 1998
- Updated kitchen with quartz counters
- Fenced backyard
- One-car garage
- Roof age unknown
- Located near a public park and bus stop

Poor prompt No. 1: Almost no direction

Write a great listing description for my new listing.

Why it is poor:

- Provides no verified property facts
- Does not specify length, audience, or tone
- Invites the AI to invent attractive features
- Contains no accuracy or fair housing safeguards
- Gives the broker little control over the output

The AI might assume the home has an open floor plan, new roof, excellent schools, quiet neighborhood, or other unverified characteristics.

Poor prompt No. 2: Encourages risky language

Write an exciting description for a perfect family home in a safe neighborhood with great schools. Make it sound luxurious and mention anything that will help it sell quickly.

Why it is poor:

- “Perfect family home” may express a preference based on familial status
- “Safe neighborhood” is subjective and unsupported
- “Great schools” may be unverified and may function as coded demographic language
- “Mention anything” encourages invention
- “Help it sell quickly” prioritizes persuasion over accuracy and consumer protection

This is a particularly useful classroom example because the prompt itself creates some of the risk.

Better prompt: Provides facts and basic direction

Write a 100-word listing description for a three-bedroom, two-bath home with 1,850 square feet. The home was built in 1998 and has an updated kitchen with quartz counters, a fenced backyard, and a one-car garage. Use a warm, professional tone.

Why it is better:

- Supplies relevant property facts
- Specifies length and tone
- Gives the AI a defined task
- Reduces—but does not eliminate—the likelihood of invention

What is still missing:

- An instruction not to add unsupported facts
- Fair housing safeguards
- Treatment of the unknown roof age
- A requirement to flag information needing verification
- Guidance concerning subjective or exaggerated claims

Excellent prompt: Broker-controlled and consumer-focused

You are assisting a North Carolina real estate broker with a residential listing description. Using only the verified facts below, draft a clear, appealing description of no more than 120 words.

Verified facts:

- *Three bedrooms and two bathrooms*
- *Approximately 1,850 square feet*
- *Built in 1998*
- *Updated kitchen with quartz countertops*
- *Fenced backyard*
- *One-car garage*
- *Near a public park and bus stop*
- *Roof age is unknown*

Requirements:

- *Do not add, infer, or exaggerate any property feature.*
 - *Do not describe the neighborhood as safe, quiet, desirable, exclusive, or family-friendly.*
 - *Do not refer to the likely occupants or use language expressing a preference for a protected class.*
- *Do not characterize schools or school quality.*
- *Do not state or imply that the roof is new or recently replaced.*
 - *Use a warm, professional tone without superlatives such as “perfect,” “best,” or “one of a kind.”*

After the description, provide a separate list titled “Broker Verification Needed” identifying any statement that should be checked before publication.

Why this is excellent:

- Defines the AI’s role
- States the desired outcome
- Separates verified facts from instructions
- Establishes clear legal and factual safeguards
- Identifies information the AI must not infer
- Specifies length, tone, and format
- Requires a final verification step
- Keeps responsibility with the broker
- An excellent follow-up prompt

The first answer should not necessarily be the final answer. A broker could follow it with:

Audit your proposed description. Identify every factual claim and connect it to a verified fact supplied in my prompt. Flag any statement that is inferred, subjective, potentially misleading, or possibly inconsistent with fair housing requirements. Then provide a corrected version using only supported statements.

That follow-up is especially valuable for class because it demonstrates that prompting can include review and correction, not merely content creation.

A useful formula is

Role + verified facts + desired outcome + constraints + format + verification

An excellent prompt improves the process—it does not guarantee an excellent or compliant result. The broker must still read, question, verify, and approve the final description.

HALLUCINATIONS, OUTDATED INFORMATION, BIAS, AND LACK OF CONTEXT

A *hallucination* occurs when an AI system generates inaccurate, unsupported, or entirely fabricated information and presents it as though it were true. An AI-generated response may include a nonexistent law, court case, Commission rule, property feature, statistic,

quotation, or source. Because the language is often fluent, detailed, and confident, the error may not be obvious.

For example, an AI system might invent a North Carolina disclosure requirement, incorrectly describe the contents of a standard form, or add a fireplace to a listing description even though the broker never mentioned one. Brokers should verify material statements through authoritative sources and should never treat the confidence or polish of an AI response as proof of accuracy.

AI may also provide *outdated information*. A model's original training may not include recent changes to statutes, Commission rules, forms, market conditions, interest rates, or government programs. Even an AI system with internet access may retrieve an old webpage, rely on a superseded document, or fail to recognize that a rule has changed. Real estate information is particularly time-sensitive, and advice that was once correct may no longer be reliable. When a question involves current law, rules, forms, fees, deadlines, statistics, or market data, the broker should confirm the information through an official, current source and note the effective date.

Bias can arise from an AI system's training data, design, assumptions, or manner of use. If historical data reflects discrimination or unequal access to housing, an AI system may reproduce those patterns in property recommendations, advertising, tenant screening, valuations, lead rankings, or other results. Bias may also appear in generated language—for example, when a system associates certain neighborhoods, names, occupations, or demographic characteristics with desirability, safety, creditworthiness, or risk. A result does not become neutral merely because a computer produced it. Brokers must consider both the language used and the effect of an AI-assisted practice on consumers and protected classes.

Finally, an AI system may reach an inappropriate conclusion because it lacks *context*. The system knows only the information available in its training, the current prompt, the conversation, and any sources it is permitted to access. It may not know the property's actual condition, the client's instructions, the broker's agency role, local ordinances, facts learned during the transaction, or the purpose for which the information will be used. When important context is missing, the system may fill the gap with an assumption rather than ask for clarification. Brokers should provide relevant, nonconfidential context when appropriate, identify what the system does not know, and use professional judgment to determine whether the output fits the actual transaction.

Hallucinations, outdated information, bias, and missing context all lead to the same conclusion: AI-generated information must be questioned and verified before a broker relies on it or shares it with a consumer.

Real Estate Brokerage Examples

Hallucinations

Invented property features: A broker asks AI to draft a listing description from limited property information. The resulting description states that the home has hardwood floors, a recently replaced roof, and an energy-efficient HVAC system, even though none of those facts appeared in the prompt.

Fabricated legal authority: A broker asks whether a particular disclosure is required in North Carolina. The AI confidently cites a nonexistent Commission rule or court decision and describes requirements that do not exist.

False contract information: An AI-generated email tells a buyer that the standard Offer to Purchase and Contract automatically gives the buyer a five-business-day right to cancel. The statement sounds plausible but does not accurately describe the contract.

Outdated Information

Superseded forms: A broker asks AI to explain a standard real estate form. The response describes provisions from an earlier version of the form rather than the version being used in the transaction.

Changed laws or Commission rules: An AI system provides an explanation based on a former North Carolina statute or Commission rule without recognizing that the law or rule has been amended.

Old market information: A broker asks AI to prepare a market summary. The system uses older sales figures, mortgage rates, inventory levels, or economic data but presents them as current.

Bias

Housing advertisement targeting: An advertising platform predicts which consumers are most likely to respond to a listing and shows the advertisement primarily to users who resemble past buyers. As a result, members of certain protected classes receive fewer opportunities to see the property.

Tenant screening: An AI-assisted screening system gives lower scores to applicants from particular ZIP codes or with certain employment histories. Although the system does not expressly consider race or national origin, those factors may operate as substitutes for protected characteristics and produce a discriminatory effect.

Property recommendations: A recommendation system repeatedly directs families with children toward suburban properties while showing childless consumers more urban options. The pattern may reinforce stereotypes and contribute to steering.

Lead ranking: A brokerage system assigns higher priority to prospects based on names, language patterns, location, or other characteristics that correlate with protected classes. Brokers then respond more quickly or provide more complete service to the higher-ranked consumers.

Lack of Context

Incomplete property condition: A broker asks AI to describe a home as “move-in ready” without telling the system that the inspection revealed active water intrusion and significant electrical problems. The AI cannot account for facts it was never given.

Unknown agency relationship: A broker asks AI to draft advice for “my client” but does not explain whether the broker represents the buyer, represents the seller, or is acting as a designated dual agent. The resulting advice may be inconsistent with the broker’s actual duties.

Missing client instructions: An AI system recommends an aggressive negotiating strategy without knowing that the seller instructed the broker to prioritize a quick, uncomplicated closing over obtaining the highest possible price.

Insufficient local information: AI responds to a question about short-term rentals using general information but does not know that the property is subject to a local ordinance, restrictive covenant, or homeowners’ association rule.

These examples illustrate why a broker must ask two separate questions about every AI-generated result: “Is this information accurate?” and “Does the system have enough context to make this information appropriate for this particular transaction?”

WHY A CONVINCING ANSWER IS NOT NECESSARILY AN ACCURATE ANSWER

Generative AI is designed to produce language that is responsive, coherent, and plausible—not to guarantee that every statement is true. It can organize an answer professionally, use appropriate terminology, provide detailed explanations, and express conclusions with confidence even when the underlying information is incomplete, outdated, misunderstood, or fabricated. The system’s tone does not reliably indicate its level of certainty, and a lengthy explanation is not evidence that research or verification occurred. Brokers are particularly vulnerable to this problem because AI-generated responses may resemble legal guidance, market analysis, advertising copy, or advice from an experienced professional. A broker must evaluate the source and accuracy of the information rather than being persuaded by the fluency of its presentation.

Examples

1. A convincing but incorrect explanation of a contract

An AI system provides a detailed explanation of a North Carolina contract provision, complete with deadlines and consequences. The answer is well organized and uses familiar legal terminology, but it misstates the buyer’s rights under the actual form. The broker must compare the answer with the current contract and seek appropriate guidance when necessary.

2. An authoritative citation to a nonexistent rule

AI tells a broker that “Commission Rule 58A .0123” requires a particular disclosure. It quotes what appears to be the rule’s language and explains how a broker should comply. The citation and quotation may be entirely fabricated. The broker should locate the rule on an official source rather than assume that a specific-looking citation is genuine.

3. A polished listing description containing invented facts

An AI-generated listing description states that a home has “new energy-efficient windows, a recently updated roof, and sweeping sunset views.” The language is appealing and sounds like ordinary real estate advertising, but the broker supplied none of those facts. A polished description can still constitute a misrepresentation.

4. A professional-looking market analysis based on poor data

AI produces a table of comparable sales, calculates an average price per square foot, and recommends a listing price. The analysis appears methodical, but some sales are outside the relevant market area, others are too old, and one property never actually closed. Attractive formatting does not correct unsuitable or inaccurate data.

5. A confident answer based on outdated information

A broker asks about a current law, Commission rule, standard form, government program, or mortgage requirement. The AI provides a clear answer but relies on information that has since changed. The answer may once have been correct, which makes the error more difficult to recognize.

6. A reasonable-sounding answer that lacks transaction context

AI recommends that a seller reject a repair request because the items appear minor. The system does not know that the defects involve active water intrusion, that the seller previously attempted repairs, or that the broker has information affecting a disclosure obligation. The recommendation sounds reasonable only because critical context is missing.

The proper question is not “Does this answer sound right?” but “How can I verify that this answer is right?”

CHAPTER 2: Protecting Confidential and Sensitive Information

THE BROKER’S CONTINUING RESPONSIBILITY FOR AI-ASSISTED WORK

A broker’s legal and professional responsibilities do not change merely because artificial intelligence helped perform a task. Whether AI drafts an advertisement, summarizes a document, recommends properties, estimates value, communicates with a consumer, or assists with a decision, the broker remains responsible for protecting confidential information, providing accurate and honest information, complying with fair housing and advertising laws, disclosing material facts, following lawful client instructions, and exercising independent professional judgment. The broker cannot avoid responsibility by claiming that a mistake, discriminatory result, misleading image, or unauthorized disclosure was produced by a technology provider. Before using AI-assisted work, the broker should understand the tool’s purpose and limitations, review the information

supplied to it, verify material statements through reliable sources, correct inappropriate content, and determine whether disclosure, client permission, broker-in-charge supervision, or professional advice is necessary. AI may assist with the work, but it cannot assume the broker's duties or accountability to clients, customers, the Commission, and the public.

CLIENT CONFIDENTIALITY AND AGENCY RESPONSIBILITIES

A broker's duty of confidentiality arises from the agency relationship and applies to information that a client reasonably expects the broker to protect. Confidential information may include the client's motivation, negotiating position, financial circumstances, willingness to accept particular terms, personal concerns, relocation plans, or other information that could disadvantage the client if disclosed. The duty may continue even after the transaction or agency relationship ends. It does not, however, permit a broker to conceal a material fact that must be disclosed under North Carolina law or Commission rules. Brokers must distinguish between information that belongs to the client and should remain confidential and information affecting the property or transaction that the broker is legally required to disclose.

AI creates additional confidentiality risks because information submitted to a system may be transmitted, stored, reviewed, reused, or shared in ways the broker and client do not expect. A broker who uploads a contract, email chain, inspection report, financial statement, preapproval letter, or transaction file may disclose far more than is necessary to complete the task. Even a seemingly harmless prompt may reveal a client's identity, property address, bargaining position, financial limits, reasons for moving, or concerns about the transaction. Deleting a conversation from the broker's screen does not necessarily establish that the information has been removed from the provider's systems.

The use of an outside AI tool does not eliminate or transfer the broker's agency duties. Before submitting client or transaction information, the broker should determine what information the system actually needs, whether the broker is authorized to provide it, how the provider stores and uses it, who may gain access to it, and whether a safer method is available. Client permission may be appropriate in some circumstances, but permission should be informed and specific; it should not be treated as a blanket waiver of the broker's responsibility to use reasonable care. A client cannot meaningfully authorize a risk that the broker has not investigated or explained.

The safest practice is to use fictional, generalized, or thoroughly de-identified information whenever possible. Names, signatures, contact information, account numbers, identification numbers, property addresses, and other identifying details should be

removed unless they are genuinely necessary and the system has been approved for that use. Brokers should follow firm policies, consult the broker-in-charge when uncertain, and avoid placing confidential information into a consumer AI system merely because doing so is convenient. The guiding question should be: *“If this information were exposed, retained, or used for another purpose, could it harm my client or violate my agency responsibilities?”*

PERSONALLY IDENTIFIABLE AND SENSITIVE INFORMATION

Personally identifiable information is information that identifies—or can reasonably be connected to—a particular person. It may include a name, home address, telephone number, email address, date of birth, Social Security number, driver’s license or passport number, photograph, signature, login credentials, account number, or device and location data. Sensitive information is a broader category that may include income, credit history, marital or family circumstances, medical or disability information, immigration information, housing needs, employment history, and a client’s reasons for buying, selling, or moving.

A single fact may appear harmless, but several facts combined may reveal a person’s identity or create opportunities for discrimination, identity theft, fraud, or misuse. Brokers should collect, retain, and disclose only the information needed for a legitimate brokerage purpose and should not enter such information into an AI system without authorization, appropriate safeguards, and a clear need.

TRANSACTION DOCUMENTS AND FINANCIAL INFORMATION

Real estate transaction documents often contain far more sensitive information than is necessary for an AI-assisted task. Offers, contracts, agency agreements, disclosure forms, closing documents, preapproval letters, bank statements, tax returns, proof-of-funds documents, trust account records, and repair negotiations may reveal names, signatures, addresses, account information, purchase limits, negotiating strategies, and other confidential facts. A broker who uploads an entire document to obtain a summary or draft may unintentionally disclose information concerning several people, including consumers who never authorized the use of the AI system.

Before using a document with AI, the broker should determine whether the task can be completed with a blank form, fictional example, short excerpt, or properly de-identified version. Financial information deserves heightened protection because its exposure may result in identity theft, wire fraud, unauthorized transactions, or substantial harm to the consumer.

THE RISKS OF UPLOADING CONTRACTS, INSPECTION REPORTS, FINANCIAL RECORDS, IDENTIFICATION DOCUMENTS, AND CLIENT COMMUNICATIONS

Uploading contracts, inspection reports, financial records, identification documents, or client communications creates risks that may not be visible at the time of submission. The technology provider may retain the document, use its contents to operate or improve the service, permit access by employees or contractors, transfer information to another system, or store the data in a location subject to different security and privacy requirements. The AI may also extract, reproduce, summarize, or disclose information beyond what the broker intended. An inspection report may reveal a client's negotiating position; an email chain may contain legal advice or private family information; and an identification document may contain everything needed for identity theft. Removing a person's name may not be sufficient if the property address, transaction details, photographs, or surrounding facts still reveal the person's identity.

Brokers should assume that uploading a document is a disclosure of its entire contents and should proceed only when the use is authorized, necessary, consistent with firm policy, and protected by an appropriately secure system.

Information entered into an AI system may remain in the provider's records after the broker closes the application, deletes the conversation, or completes the transaction. Depending on the product and account settings, submitted prompts, documents, images, and responses may be stored for security, legal compliance, product operation, quality review, or system improvement. Information may also be transferred to affiliated services or third-party providers that perform storage, processing, or other functions.

Brokers should not assume that deleting information from the visible account immediately removes every copy or prevents prior use. Before submitting client or transaction information, a broker should determine how long the provider retains it, whether the provider may reuse it, whether a user can request deletion, and whether the firm can control or document those settings.

REVIEWING A TOOL'S PRIVACY AND DATA-USE TERMS

Before using an AI tool for brokerage work, the broker or brokerage firm should review the provider's privacy policy, terms of service, data-use disclosures, and available security settings. Important questions include what information the provider collects, where it is stored, who may access it, whether it is used to train or improve AI models, whether it is shared with third parties, how long it is retained, and what happens when an account is closed. Brokers should also determine whether different protections apply to free, individual, business, or enterprise accounts. A familiar brand name, paid subscription, security symbol, or statement that data is "encrypted" does not by itself establish that the tool is appropriate for confidential brokerage information. Because terms and product

settings can change, firms should review approved tools periodically rather than relying indefinitely on an earlier evaluation.

DE-IDENTIFYING INFORMATION

De-identification is the process of removing or changing details that could connect information to a particular person, property, or transaction. A broker might replace names with “Buyer A” and “Seller B,” remove signatures and account numbers, generalize financial figures, and omit the street address before asking AI to analyze a fictionalized scenario. Effective de-identification requires more than deleting a person’s name. Property photographs, exact dates, unusual transaction terms, parcel numbers, email addresses, document metadata, and combinations of otherwise ordinary facts may still reveal an identity. Brokers should provide only the minimum information necessary for the task and consider whether a blank form, hypothetical example, or generalized summary would work just as well. De-identification reduces risk but does not guarantee anonymity.

CLIENT NOTICE AND PERMISSION

A client should receive appropriate notice when a broker proposes to use an AI tool in a manner that involves the client’s confidential, personal, financial, or transaction information. The notice should explain what information will be submitted, the purpose of the submission, the type of system involved, the reasonably known risks, and any available alternative. Permission should be informed, specific, documented, and obtained before the information is disclosed; a broad statement buried in an agreement may not demonstrate that the client understood the proposed use. Client permission does not excuse the broker from exercising reasonable care, following firm policy, complying with applicable law, or protecting information belonging to other people who did not give permission. When the proposed AI use is unnecessary, unusually risky, or inconsistent with the broker’s duties, obtaining consent does not automatically make the practice appropriate.

Sample Client Notice and Permission for AI-Assisted Work

Your broker proposes using an artificial intelligence tool to assist with the following task:

Proposed use: Summarizing selected portions of the property inspection report and organizing the reported items for your review.

Before using the tool, your broker will remove your name, contact information, signatures, account information, and other personal identifying information. The AI-generated summary will not replace the original inspection report, the advice of the inspector or another qualified professional, or your broker's independent review.

Information submitted to an AI provider may be processed, stored, or retained according to the provider's privacy policy and terms of use. Although the brokerage will take reasonable precautions and use a firm-approved system, it cannot guarantee that information submitted to a third-party provider will never be accessed, retained, or misused. AI-generated information may also be inaccurate, incomplete, or misleading. Your broker will review the output before sharing it with you, but you should rely on the original documents and appropriate professional advice when making decisions.

You are not required to permit this use of AI. If you decline, your broker will complete the task without submitting the information to the proposed AI tool, and your decision will not affect the brokerage services you receive. You may also withdraw your permission before the information is submitted.

Please indicate your decision:

☐ I give permission for the broker to use the described AI tool for the specific purpose stated above.

☐ I do not give permission for the broker to use the described AI tool.

Client name: _____

Client signature: _____

Date: _____

Broker name: _____

AI tool and account type to be used: _____

Information to be submitted: _____

Information to be removed or de-identified: _____

WHEN TO INVOLVE THE BROKER-IN-CHARGE

A broker should involve the broker-in-charge before using AI when the proposed use may affect the broker's legal duties, the firm's records, a consumer's rights, or the accuracy of information communicated to the public. Examples include uploading transaction documents or confidential information; using AI to interpret a contract, law, rule, or standard form; altering a property photograph; implementing an AI chatbot; using automated advertising, valuation, lead-ranking, or tenant-screening tools; or relying on AI to communicate directly with clients and customers. The BIC should also be consulted whenever the broker is uncertain whether a tool is approved, whether client notice or permission is necessary, whether an output must be disclosed, or whether the proposed use complies with firm policy.

The BIC should be notified promptly when an AI-related problem has already occurred. This includes the disclosure of confidential or personal information, publication of inaccurate advertising, distribution of misleading images, discriminatory language or targeting, an incorrect communication to a consumer, or reliance on fabricated legal or transaction information.

Early involvement gives the BIC an opportunity to preserve relevant records, stop further distribution, correct the information, notify affected parties when appropriate, consult legal or technical professionals, and determine whether additional reporting or remedial action is required. Brokers should not attempt to conceal or quietly correct a significant AI-related error without appropriate supervision.

FIRM POLICIES GOVERNING APPROVED AND PROHIBITED AI USES

A brokerage firm should establish a written AI policy that identifies approved tools, account types, purposes, and review procedures. The policy should explain what information may be entered into an AI system, what information must first be de-identified, and what information may never be submitted. It should also address client notice and permission, required human review, verification through authoritative sources, disclosure of altered images or AI-generated content, fair housing safeguards, record retention, vendor evaluation, cybersecurity, and procedures for reporting an error or suspected data exposure. Because AI products and their terms change frequently, the firm should assign responsibility for reviewing approved tools and updating the policy.

A firm may permit lower-risk uses—such as brainstorming with fictional information, revising nonconfidential text, or creating internal training examples—while prohibiting higher-risk practices. Prohibited uses might include entering Social Security numbers, account information, identification documents, passwords, unredacted contracts, or

confidential client communications into an unapproved system; using AI as a substitute for legal, appraisal, inspection, tax, or other professional advice; publishing AI-generated content without broker review; concealing defects through image alteration; allowing a chatbot to make unauthorized representations; or relying solely on an automated score to provide or deny brokerage services. A useful policy should be specific enough to guide daily decisions while allowing the broker and BIC to recognize new situations that require caution, escalation, or prohibition.

CHAPTER 3: ACCURACY, ADVERTISING, AND FAIR HOUSING

VERIFYING AI-GENERATED STATEMENTS

AI-generated statements should be treated as unverified drafts rather than established facts. Before using or communicating an AI-generated statement, a broker should identify each factual claim and determine what reliable source can confirm it. Property characteristics should be checked against the broker's personal knowledge, measurements, permits, surveys, inspection reports, public records, or information supplied by an appropriate professional. Market statistics and comparable sales should be confirmed through current, reliable data. Statements about laws, Commission rules, standard forms, taxes, zoning, schools, insurance, lending, or government programs should be compared with current authoritative sources or referred to a qualified professional. Asking the same AI system whether its answer is correct does not constitute independent verification.

Verification also requires the broker to examine what the AI may have implied, omitted, or overstated. A description such as "completely renovated," "maintenance-free," "approved building lot," or "ideal for short-term rentals" may contain several factual claims that require separate support. The broker should distinguish facts from estimates, opinions, predictions, and promotional language; confirm that information is current; preserve the source used for verification when appropriate; and remove any statement that cannot be supported. If reliable verification is unavailable, the broker should not present the information as fact. The broker remains responsible for the accuracy of advertising and other communications, regardless of whether the words were written by the broker, an employee, a vendor, or an AI system.

A note:

When ChatGPT searches the web, conducts research, or searches supplied files, it can provide citations linking the statements to the retrieved sources. OpenAI's documentation

specifically provides for web-search results and source citations. But when ChatGPT answers from patterns learned during its original training, it may not be retrieving information from an identifiable source at that moment. In that situation, it generally cannot tell you precisely which document produced a particular statement.

Even when sources are supplied, the broker must still open and examine them. A cited source may:

- Be outdated or superseded
- Be a secondary summary rather than the controlling authority
- Address a different state or factual situation
- Fail to support the AI's precise conclusion
- Be interpreted incorrectly by the AI
- Omit important qualifications or conflicting authority

For example, if a broker asks, "What does North Carolina law require?" a safer prompt would be:

Search current official North Carolina sources, including the North Carolina Real Estate Commission, for the answer. Cite the source immediately after each material statement, provide the effective or publication date, and distinguish the source's actual language from your interpretation. If an official source does not establish the answer, say so rather than infer one.

The broker should then follow the citation, confirm that it is an official and current source, and read enough of the underlying material to ensure that it supports the statement.

Again, asking the same AI system whether its answer is correct does not constitute independent verification. Citations supplied by an AI system are useful starting points, but the broker must open, evaluate, and confirm the cited sources before relying on the answer.

MISREPRESENTATION AND OMISSION

North Carolina License Law prohibits both affirmative misstatements and failures to disclose important information. Under N.C.G.S. § 93A-6(a)(1), the Real Estate Commission may discipline a broker for making any willful or negligent misrepresentation or any willful or negligent omission of material fact. The statute also authorizes discipline for making false promises likely to influence, persuade, or induce and for pursuing a course of misrepresentation or false promises through advertising or otherwise. N.C.G.S. § 93A-6(a)(2)–(3). A material fact is any fact that could affect a reasonable person's decision to buy, sell, or lease real property. Commission Rule 21 NCAC 58A .0114(c) reinforces the

broker's obligation concerning material facts; the Commission has explained that brokers have an affirmative duty to make reasonable efforts to discover and disclose material facts to all parties, regardless of whom the broker represents.

AI-assisted work can create either type of violation. Misrepresentation may occur when AI adds an unverified property feature, exaggerates a renovation, incorrectly describes zoning or permitted uses, or alters an image in a way that creates a false impression. Omission may occur when AI shortens a description, summarizes an inspection report, or edits a photograph in a manner that removes or obscures a material fact. The broker cannot avoid responsibility by explaining that the error originated with an AI tool, and the absence of an intent to deceive does not necessarily protect the broker because negligent misrepresentation and negligent omission are also grounds for discipline. Before publishing or sharing AI-assisted content, the broker must compare it with reliable source information, identify what has been added or removed, correct unsupported statements, and ensure that all required material facts are disclosed accurately and in time to be meaningful to the consumer's decision.

AI-GENERATED LISTING DESCRIPTIONS

Generative AI can quickly organize property information and produce polished advertising language, but the result should be treated as a draft prepared by an assistant that has never visited the property and does not know which facts have been verified. Unless carefully constrained, AI may add common but unsupported features such as an open floor plan, hardwood floors, abundant natural light, recent renovations, excellent schools, a quiet neighborhood, or a new roof. It may also convert limited information into a broader claim—for example, changing “kitchen counters replaced in 2022” to “fully renovated gourmet kitchen.” The more natural and persuasive the description sounds, the easier it may be for a broker to overlook an invented fact or exaggeration.

A broker using AI to draft a listing description should provide only verified property facts, instruct the system not to infer or add features, specify fair housing and advertising constraints, and require it to flag any statement that needs confirmation. The broker should then compare every factual statement with reliable source information and review the description for misleading implications, omissions, subjective claims, and discriminatory language.

Disclaimers such as “information deemed reliable but not guaranteed” do not authorize a broker to publish information the broker knows—or reasonably should know—is inaccurate. The broker and broker-in-charge remain responsible for advertising published

by or in the name of the firm, regardless of whether the language originated with the broker, a seller, a vendor, or an AI system.

MATERIAL FACTS AND PROPERTY CHARACTERISTICS

AI cannot independently determine the actual characteristics or condition of a property. Information such as square footage, acreage, number of bedrooms, septic capacity, permitted improvements, roof age, flood-zone status, zoning, restrictive covenants, homeowners' association obligations, utility availability, and the condition of structural or mechanical systems must be confirmed through appropriate sources. Even apparently simple characteristics may require qualification. A room used as a bedroom may not be permitted as one; a finished basement may not have been included in permitted living area; and a property described as suitable for short-term rentals may be subject to zoning restrictions, covenants, or association rules. A broker should not allow AI to turn an assumption, estimate, seller statement, or unverified record into an unqualified fact.

Material facts require particular care. A material fact is any fact that could affect a reasonable person's decision to buy, sell, or lease real property, and North Carolina brokers have an affirmative duty to make reasonable efforts to discover and disclose material facts to all parties. AI may omit a material fact when shortening a description, summarizing a report, or emphasizing favorable features. It may also minimize a problem by replacing precise language—such as “active water intrusion observed in the crawl space”—with a vague phrase such as “some moisture concerns.” The broker, not the AI system, must determine what information requires further investigation and disclosure. A seller's instruction, a “No Representation” response on a disclosure statement, or an AI-generated summary does not relieve the broker of the duty to discover and timely disclose material facts.

MARKET STATISTICS AND AUTOMATED VALUATIONS

AI can help organize, calculate, summarize, and explain market data, but the quality of the result depends on the quality and relevance of the underlying information. A market statement may be misleading if the data is outdated, drawn from the wrong geographic area, based on too few transactions, or calculated using inconsistent property types or time periods. AI may also confuse active-listing prices with closed-sale prices, average values with medians, original list prices with final list prices, or days on market with cumulative days on market. Before communicating an AI-assisted statistic, the broker should confirm the source, date, geographic boundaries, property criteria, sample size, calculation method, and meaning of the measurement. A statement such as “prices

increased 10 percent” is incomplete unless the broker can explain which prices, where, during what period, and compared with what earlier period.

An automated valuation model, or AVM, uses property records, sales information, location data, and statistical patterns to estimate a property’s value. AVMs can be useful as preliminary reference points, but they do not necessarily account for current condition, renovations, deferred maintenance, views, site characteristics, interior quality, unusual design, unrecorded improvements, or rapidly changing neighborhood conditions. Their reliability may also decline in rural areas, markets with few recent sales, or areas containing properties that differ substantially from one another. An AVM’s precise-looking estimate does not guarantee precision; it remains a prediction produced from available data and assumptions.

Brokers should clearly distinguish among an AVM, a broker’s comparative market analysis (CMA), and an appraisal performed by a licensed or certified appraiser. An AVM should not be described as an appraisal or presented as a guaranteed sales price. When preparing a comparative market analysis or broker price opinion, the broker must independently select and evaluate appropriate information rather than simply adopting an automated estimate. If an AI tool assists with comparable sales or adjustments, the broker should confirm that the properties exist, that the transaction data is accurate, and that the comparisons are relevant. Consumers should be told about important limitations and uncertainties so they can understand that an automated valuation is one source of information—not a substitute for inspection, local knowledge, professional analysis, or informed judgment.⁴

Common examples of automated valuation models

Zillow Zestimate® — estimates a home’s market value using public records, MLS data, user-submitted information, property characteristics, and market trends. Zillow expressly states that it is not an appraisal.

Redfin Estimate — uses MLS data and other property and market information to calculate an estimated value. Redfin also describes it as a starting point rather than an appraisal or a substitute for professional pricing advice.

Realtor.com RealEstimateSM — may display as many as three valuations supplied by independent AVM providers. Realtor.com currently identifies Collateral Analytics, CoreLogic, and Quantarium as its valuation providers.

⁴ See appendix for an example of a AI-generated CMA.

HouseCanary valuations — professional products used by brokers, investors, lenders, and mortgage servicers. HouseCanary combines property data, valuations, market forecasts, and AI-driven analytics.

Cotality Total Home ValueX™ — an AI- and machine-learning-based AVM used for such purposes as marketing, risk management, portfolio analysis, and valuation-related workflows.

Brokers may also encounter AVMs embedded inside mortgage-underwriting systems, home-equity platforms, investor programs, property-data services, and MLS-related products without seeing the name of the underlying model.

It is useful to distinguish an AVM from several things that may look similar:

A tax assessment is not necessarily an AVM or an estimate of current market value.

An online home-value calculator may use an AVM, but the consumer may not know which model supplies the result.

A broker's comparative market analysis is not an AVM when the broker personally selects and analyzes the relevant information.

An appraisal performed by a licensed or certified appraiser is not the same as an automated valuation.

AI-GENERATED EXPLANATIONS OF CONTRACTS OR LEGAL REQUIREMENTS

AI can translate complicated language into plain English, summarize a contract provision, or identify questions that a consumer may want to discuss with a broker or attorney. However, an AI-generated explanation may rely on an outdated version of a form, confuse the law of one state with another, omit an exception, miscalculate a deadline, or describe a contractual right that does not exist. It may also present a legal conclusion with unwarranted confidence.

Before using an AI-generated explanation, the broker should compare it with the current document and authoritative sources, confirm that quoted language and citations actually exist, and distinguish the contract's language from the AI's interpretation. Uploading an executed contract for explanation also creates confidentiality and data-security risks that may be avoided by using a blank form or a de-identified excerpt.

AI does not expand the scope of advice a real estate broker is permitted or qualified to provide. North Carolina License Law authorizes disciplinary action when a broker performs legal services or other acts constituting the unauthorized practice of law. N.C.G.S. § 93A-6(a)(11). A broker may provide factual information and explain standard forms consistent with the broker's training and role but should not use AI to draft nonstandard legal provisions, determine a party's legal rights, predict how a court would interpret a provision, or recommend a legal course of action. When a question requires legal interpretation or advice, the broker should identify the issue and refer the consumer to a North Carolina attorney. The broker cannot avoid responsibility for improper legal advice by attributing the explanation to an AI system.

Example: An Incorrect AI Explanation of a Contract Provision

Suppose a broker asks an AI system to explain the due diligence provisions in a North Carolina residential purchase contract. The AI responds:

North Carolina buyers have an automatic three-day right to cancel a real estate contract. If the buyer cancels within those three days, the seller must return both the earnest money deposit and the due diligence fee.

This explanation sounds clear and authoritative, but it is incorrect. North Carolina does not provide a general three-day cancellation period for an ordinary residential resale contract. A buyer's right to terminate depends on the language of the particular contract, including the negotiated due diligence period and any other applicable provisions.

Under the NCAR *North Carolina Offer to Purchase and Contract*, a buyer who properly terminates during the due diligence period ordinarily may recover the earnest money deposit, but the due diligence fee is generally nonrefundable except in circumstances provided by the contract.

If a broker repeated the AI-generated explanation, the buyer might believe that cancellation rights exist when they do not, miss the actual termination deadline, or incorrectly expect the return of the due diligence fee. The broker should consult the current contract form and applicable Commission guidance rather than relying on the AI's summary.

Questions requiring interpretation of a party's legal rights should be referred to a North Carolina attorney. AI can help identify provisions for review, but it should not be treated as the authority on what a contract means or how it applies to a particular transaction.

ALTERED PHOTOGRAPHS AND VIRTUAL STAGING

Artificial intelligence can add furniture, artwork, window treatments, landscaping, or other decorative elements to a photograph in seconds. Virtual staging can help consumers visualize how a vacant room might be used, and ordinary photographic adjustments may improve lighting, color balance, or image clarity. These uses can be appropriate when they do not change the property's actual physical characteristics. A virtually staged image should be clearly and conspicuously labeled—such as “Virtually Staged” or “This Image Has Been Digitally Altered”—so consumers understand that the depicted furnishings or improvements are not present at the property. When practical, the broker should display the original photograph with the altered version.

The risk increases when AI changes facts rather than presentation. Removing utility poles, power lines, neighboring structures, water stains, cracks, damaged flooring, or other visible conditions may create a false impression of the property. Likewise, adding a fireplace, changing a countertop material, replacing an obstructed view with a scenic one, enlarging windows, improving landscaping that does not exist, or digitally changing a room's dimensions may misrepresent the property or conceal a material fact. Labeling such an image “virtually staged” does not necessarily cure the deception; disclosure does not give a broker permission to advertise a property inaccurately.

The broker remains responsible for every photograph used in advertising, whether the image was created by a photographer, editing service, virtual-staging company, or AI system. Before publication, the broker should compare each altered image with the original and ask: Does this image merely help the consumer visualize the space, or does it change what the consumer would reasonably believe about the property? Images that conceal defects or materially alter the property should not be used. Firm policies should address permitted alterations, required labels, retention of original images, and approval procedures. When uncertain, the broker should consult the broker-in-charge before publishing the photograph.



DISCLOSURE OF DIGITALLY MODIFIED IMAGES

When a real estate photograph has been materially modified, the disclosure should be clear, conspicuous, and placed where consumers will see it before relying on the image. Labels such as “Virtually Staged,” “Digitally Altered,” or “AI-Generated Rendering—Not an Actual Photograph” should appear on or immediately beside each affected image.

A general statement buried at the end of a listing may not adequately inform consumers which photographs were changed. When practical, brokers should also provide unaltered photographs so viewers can distinguish the property’s actual condition from the digitally created presentation.

The disclosure should accurately describe the nature of the modification. Adding virtual furniture differs from depicting a proposed renovation, and both differ from altering a permanent property characteristic. Most importantly, disclosure does not make a materially misleading image acceptable.

A label stating “digitally altered” does not justify concealing damage, removing an undesirable neighboring feature, changing room dimensions, or depicting an amenity that does not exist. Brokers should retain the original images, review all modifications before publication, follow firm and advertising-platform requirements, and remember that responsibility for truthful advertising remains with the broker—not with the AI system or editing service.

FAIR HOUSING RISKS IN LANGUAGE, TARGETING, SCREENING, AND RECOMMENDATIONS

AI does not eliminate fair housing responsibilities. It may reproduce or amplify patterns contained in its training data, follow discriminatory instructions, or rely on seemingly neutral information that serves as a proxy for a protected characteristic. Under federal and North Carolina fair housing laws, housing-related services generally may not discriminate based on race, color, religion, sex, national origin, familial status, or disability. A broker remains responsible for ensuring that AI-assisted advertising, communications, recommendations, and housing decisions comply with these requirements.

AI-generated advertising language may create fair housing concerns even when the broker did not intend to discriminate. Descriptions such as “perfect for young professionals,” “ideal for a Christian family,” “safe neighborhood,” “no children,” or “walk to church” may express or imply a preference, limitation, or exclusion.

An AI system may also infer neighborhood demographics or generate coded language based on patterns found online. Brokers should focus descriptions on the property’s objective features—such as the number of bedrooms, accessibility features, outdoor space, or proximity to identified amenities—and review every AI-generated statement before publication.

Digital advertising tools create additional risks when they determine who will—or will not—see a housing advertisement. An advertising platform may use age, location, online behavior, interests, language, household characteristics, or “look-alike” audiences to select recipients. Even when the broker does not expressly exclude a protected group, an algorithm may deliver the advertisement disproportionately or use proxy characteristics that produce a discriminatory result. Brokers should avoid targeting criteria that limit housing opportunities based on protected characteristics, use broad and inclusive

audiences, review the platform’s housing-advertising restrictions, and monitor how advertisements are actually delivered.

AI-assisted tenant screening, lead scoring, and risk-assessment systems may evaluate applicants using credit information, rental history, criminal records, employment data, eviction records, or less transparent behavioral indicators. Errors, incomplete records, biased data, or inappropriate proxies may unfairly disadvantage certain applicants. A broker or property manager should not assume that a numerical score is neutral simply because a computer produced it.

Screening standards should be written, relevant, consistently applied, and subject to meaningful human review. Consumers should also have an opportunity to question inaccurate information and request reasonable accommodations when applicable.

Recommendation systems can create a digital form of steering. If an AI tool recommends neighborhoods based on a consumer’s name, race, religion, family composition, disability, perceived nationality, or assumptions drawn from prior users, it may restrict the consumer’s housing choices. Requests for a “safe neighborhood,” “good schools,” “a family neighborhood,” or an area “where people like me live” should not be answered by allowing AI to make demographic judgments.

A broker should identify the consumer’s objective housing criteria, provide access to all properties that meet those criteria, and direct consumers to reliable third-party sources for information they should evaluate personally. AI may assist with organizing available choices, but it must not decide where a person belongs.

ALGORITHMIC BIAS AND DISCRIMINATORY EFFECTS

Algorithmic bias occurs when an automated system produces systematically unfair results for particular groups. The bias may arise from historical data reflecting past discrimination, incomplete or unrepresentative data, incorrect assumptions made by a system’s designers, or the use of seemingly neutral factors that act as proxies for protected characteristics.

For example, a housing-advertising algorithm may learn from prior user behavior and show listings more frequently to certain demographic groups. A tenant-screening system might weigh credit history, eviction records, ZIP codes, employment patterns, or criminal records in ways that disproportionately disadvantage members of a protected class. An automated valuation model may also produce less reliable estimates in neighborhoods that have historically experienced segregation, redlining, or fewer recorded sales. A mathematical formula is not necessarily objective merely because it applies the same calculation to everyone.

A discriminatory effect may occur even when neither the broker nor the technology provider intended to discriminate. A facially neutral advertising, screening, valuation, or recommendation practice may still restrict housing opportunities disproportionately and create fair housing concerns. Brokers should therefore evaluate both the criteria used by an AI system and the results it produces. Warning signs include unexplained differences among groups, recommendations based on ZIP codes or demographic proxies, an inability to identify why a person received an unfavorable result, and claims by a vendor that its process is unbiased simply because it is automated. Before relying on an AI-assisted decision, the broker should understand the tool's purpose and limitations, apply lawful standards consistently, preserve meaningful human review, provide a way to correct inaccurate information, and involve the broker-in-charge whenever the technology could affect a consumer's access to housing.

THE BROKER'S RESPONSIBILITY FOR CONTENT PRODUCED BY THIRD-PARTY TOOLS

A broker cannot transfer professional responsibility to an AI system, advertising platform, virtual-staging service, automated valuation model, tenant-screening company, or other third-party provider. If the broker selects, publishes, distributes, or relies upon content produced by the tool, the broker remains responsible for determining whether it is accurate, current, lawful, and appropriate. A vendor's assurances, automated disclaimer, or claim that its system is "objective" does not relieve the broker of duties concerning material facts, truthful advertising, confidentiality, fair housing, agency, and competent brokerage practice.

Before using third-party content, the broker should identify the source of the information, verify important statements against reliable records, examine images for undisclosed alterations, and consider whether automated targeting, screening, or recommendations could create discriminatory effects. The broker should also follow firm policies, protect confidential information, document appropriate review, and involve the broker-in-charge when a tool's operation or risks are unclear. The central principle of this lesson is simple: AI may produce the content, but the broker chooses whether and how it will be used. That choice—and its consequences—remains the broker's responsibility.

CHAPTER 4: EXERCISING PROFESSIONAL JUDGMENT—THE BROKER IS STILL RESPONSIBLE

APPROPRIATE AND INAPPROPRIATE RELIANCE ON AI

Appropriate reliance on AI begins with recognizing that it is an assisting tool—not a substitute for the broker’s knowledge, judgment, investigation, or supervision. AI may be useful for low-risk tasks such as brainstorming educational content, reorganizing nonconfidential notes, improving the readability of a draft, suggesting questions for further investigation, or preparing a preliminary listing description from facts the broker has already verified. Even in these situations, the broker must review the result for accuracy, tone, fair housing concerns, confidentiality, and consistency with firm policy before using it. The broker should also consider whether the AI’s answer is based on a reliable source and whether the information remains current.

Reliance becomes inappropriate when a broker treats an AI-generated answer as authoritative without independent verification. AI should not be allowed to determine whether a condition is a material fact, interpret a contract for a client, calculate or monitor critical deadlines without confirmation, select contractual language, provide legal or tax advice, establish a property’s value without competent analysis, or decide which consumers should receive listings or housing opportunities. A broker also should not rely on AI to screen applicants, recommend neighborhoods, summarize inspection findings, or answer questions about laws and Commission rules unless the broker understands the subject and verifies the result through reliable sources. The more seriously an error could affect a consumer’s rights, money, safety, privacy, or access to housing, the less appropriate it is to rely on AI without heightened review.

The proper level of reliance therefore depends on the task, the information involved, and the possible consequences of an error. A low-risk draft may require careful proofreading, while a statement about a contract, material fact, valuation, fair housing issue, or legal requirement demands independent confirmation and may require consultation with the broker-in-charge or an appropriately licensed professional. If the broker cannot explain how an answer was verified, cannot identify the source supporting it, or lacks the competence to evaluate it, the broker should not present that answer to a consumer. AI can assist with the work, but the broker must remain capable of recognizing when the tool is wrong and willing to disregard its output.

MATTERS REQUIRING OFFICIAL OR AUTHORITATIVE SOURCES

Some real estate questions are too consequential to answer through an AI-generated summary, general internet search, or unofficial commentary. Matters involving License Law, Commission rules, fair housing requirements, agency duties, approved forms, contract deadlines, disclosure obligations, zoning, permitting, taxes, insurance, title, financing, and property use should be verified through a current official or authoritative source. Depending on the issue, appropriate sources may include the North Carolina Real Estate Commission, the current contract or approved form, the North Carolina General Statutes, HUD, a county or municipal government, the applicable school system, recorded public documents, or information supplied by an appropriately licensed attorney, surveyor, appraiser, inspector, engineer, accountant, or other professional. The broker should also confirm that the source applies to the correct jurisdiction and is current as of the relevant date.

AI may help a broker identify possible issues, formulate questions, or locate sources for further review, but it should not be treated as the source itself. Even when an AI response includes citations, the broker should open the cited material, confirm that the source exists, read the relevant provision in context, and determine whether it actually supports the statement.

An official source may also contain qualifications, exceptions, effective dates, or definitions omitted from the AI's summary. When the answer affects a consumer's contractual rights, financial obligations, legal position, safety, or ability to use the property, the broker should provide verified factual information within the broker's competence and refer questions requiring specialized interpretation to the appropriate professional.

WHEN AI CROSSES INTO LEGAL, TAX, APPRAISAL, INSPECTION, OR OTHER SPECIALIZED ADVICE

AI crosses into specialized advice when it moves beyond organizing or restating information and begins applying professional judgment to a consumer's particular circumstances. Warning signs include an AI system recommending what a party should do, interpreting legal rights, predicting a tax consequence, diagnosing a building condition, determining the value of a property, or representing that its conclusion is professionally reliable. The fact that the answer was produced by software does not change the nature of the advice. A broker who repeats, adopts, or delivers the answer may appear to be providing advice outside the broker's competence or authority.

In the legal context, AI should not be used to decide whether a client may terminate a contract, interpret an ambiguous provision, draft nonstandard contractual language,

determine a party's remedies, or advise a consumer about potential liability. A broker may explain the customary purpose and effect of standard form provisions within the scope of the broker's training, but questions requiring legal interpretation should be referred to a North Carolina attorney. Similarly, AI-generated statements about capital gains, depreciation, inheritance, business entities, deductions, or the tax consequences of a sale or exchange should not be presented as tax advice. Those questions belong with a qualified tax professional who can consider the consumer's complete circumstances.

AI may assist a broker in organizing comparable-sale information, but an automated estimate is not a substitute for the broker's competent analysis or a licensed or certified appraiser's appraisal. Brokers should accurately characterize a comparative market analysis or broker price opinion and should not call it an appraisal or imply that it carries the authority of one. An AI system also should not be used to determine conclusively that a roof is sound, a crack is structural, mold is present, an electrical system is safe, or a repair will solve a defect. Photographs and inspection reports may be incomplete, and an AI system cannot physically examine the property. Such matters should be evaluated by an appropriately qualified inspector, engineer, contractor, environmental professional, or other specialist.

The broker's role is to recognize the boundary, provide verified factual information within the broker's competence, and make a timely referral when specialized judgment is required. AI may help identify questions to ask or summarize information for review, but it should not be allowed to become an unlicensed lawyer, accountant, appraiser, inspector, engineer, or other professional. A useful test is: Does this answer require a professional license, an inspection of the property, access to facts not provided, or application of specialized judgment to this consumer's circumstances? If so, the broker should stop relying on the AI response and direct the consumer to the appropriate professional.

TRANSPARENCY AND DISCLOSURE

Transparency means helping consumers understand when AI has materially influenced information, images, communications, recommendations, or decisions on which they may rely. Not every routine use of spelling, formatting, or administrative automation necessarily requires a special notice. Disclosure becomes more important, however, when AI creates or materially changes an advertisement, photograph, property description, valuation, contract explanation, screening result, recommendation, or communication presented as personalized professional work. The disclosure should be timely, understandable, and specific enough to explain what the technology did. For example, a broker might state that an image was virtually staged, that a valuation was produced by an automated model and

is not an appraisal, or that an AI-generated summary has been reviewed but the original document controls.

Transparency also requires honesty about the tool's limitations and the broker's review. A broker should not imply that an AI-generated statement has been independently verified when it has not, conceal that an automated system affected a consumer's housing opportunity, or describe a digitally created feature as though it exists at the property. At the same time, disclosure does not excuse inaccurate, discriminatory, unlawful, or misleading content. Statements such as "generated by AI" or "for informational purposes only" do not transfer responsibility to the technology or cure a material misrepresentation. Effective disclosure allows the consumer to evaluate the information intelligently, while competent review ensures that the information was appropriate to provide in the first place.

BROKER-IN-CHARGE SUPERVISION

The broker-in-charge should establish and communicate firm policies governing approved AI tools, prohibited uses, protection of confidential information, required verification, advertising review, disclosure, record retention, and referral to specialized professionals. Supervision should include training affiliated brokers to recognize hallucinations, bias, privacy risks, misleading alterations, and tasks that exceed the broker's competence. Brokers should consult the broker-in-charge before adopting an unfamiliar tool, entering sensitive information, using AI in screening or other consequential decisions, or publishing content whose accuracy or legality is uncertain. The broker-in-charge cannot personally inspect every prompt and output, but should implement reasonable safeguards, respond to warning signs, correct improper practices, and reinforce the central rule that technology may assist brokerage activity without displacing professional judgment or accountability.

CORRECTING AI-GENERATED MISINFORMATION

When a broker discovers that AI-generated information provided to a client, customer, cooperating broker, or the public is inaccurate or misleading, the broker should act promptly to prevent further reliance and reduce potential harm. The broker should stop using or distributing the content, verify the correct information through authoritative sources, notify everyone known to have received or relied upon the misinformation, and provide a clear correction through the same communication channels whenever practical.

Online advertisements, MLS entries, social media posts, altered images, automated messages, and transaction communications should be corrected or removed, and the broker-in-charge should be informed when the error could affect a consumer's decision, contractual rights, finances, safety, privacy, or access to housing. The broker should document the original error, its source, the verification performed, the persons notified,

and the corrective action taken; merely deleting the content or blaming the AI system does not satisfy the broker's responsibility to address the consequences of its use.

RESPONDING WHEN A CONSUMER HAS RELIED ON AN INCORRECT AI COMMUNICATION

When a broker learns that a consumer has acted—or may be about to act—based on an incorrect AI-generated communication, the matter requires more than quietly editing or deleting the original message. The broker should promptly acknowledge the error, provide accurate information from an authoritative source, and determine what the consumer understood and what action was taken in reliance on it. Relevant questions include whether the consumer signed a document, missed or changed a deadline, incurred an expense, disclosed confidential information, decided not to investigate a property condition, or made a housing decision because of the communication. If additional consumers or transaction participants received the same misinformation, the broker should identify and notify them as well.

The broker should immediately involve the broker-in-charge when the error could affect contractual rights, money, material facts, fair housing, privacy, or the consumer's ability to complete the transaction. The original prompt, AI output, communication, supporting records, and subsequent correction should be preserved rather than altered or destroyed. Depending on the seriousness of the situation and firm policy, the broker-in-charge may need to consult the firm's attorney, errors-and-omissions insurance carrier, MLS, technology provider, or another appropriate party. If legal rights or remedies may be involved, the consumer should be encouraged to obtain advice from a North Carolina attorney; the broker should not use AI to determine the legal effect of the broker's original mistake.

Any corrective communication should be direct, factual, and understandable. It should identify the incorrect statement, supply the corrected information, cite or attach the controlling source when appropriate, and explain any immediate step the consumer should consider without attempting to provide unauthorized legal advice. The broker should not minimize the problem, blame the technology, speculate about consequences, or suggest that an "AI-generated" disclaimer eliminates responsibility. The goals are to stop continued reliance, protect the consumer from additional harm, preserve an accurate record, and ensure that qualified people address any consequences that the broker cannot resolve within the scope of brokerage practice.

BROKERAGE AI POLICIES AND REVIEW PROCEDURES

A brokerage's written AI policy should identify approved tools, permitted uses, prohibited activities, and the persons authorized to approve exceptions. The policy should address confidential and personally identifiable information, client notice and authorization, data retention, account security, advertising and image alteration, fair housing, automated screening or targeting, contract and legal questions, specialized professional advice, and recordkeeping. It should prohibit brokers from entering sensitive transaction information into unapproved systems or allowing AI to make consequential decisions without meaningful human review. Because products and data-use terms change, the firm should periodically reassess its approved tools, update its policies, train affiliated brokers, and provide a clear procedure for reporting errors or questionable results.

Review procedures should be proportionate to the risk of the task. Routine, low-risk material may require review by the broker who created it, while public advertising, altered photographs, property representations, automated valuations, consumer communications, and fair housing-sensitive activities may require additional review or broker-in-charge approval. The reviewer should verify material statements against authoritative sources, examine content for omissions and discriminatory language, confirm required disclosures, and ensure that no confidential information was improperly used. The firm should also establish procedures for preserving relevant records, correcting misinformation, notifying affected consumers, and escalating possible legal, regulatory, privacy, or insurance concerns. A useful policy should not merely say "use AI responsibly"; it should tell brokers which tools may be used, what must be reviewed, who must approve it, and what to do when something goes wrong.

Example #1 of Policy Manual Language

Use of Approved AI Systems and Protection of Information

Affiliated brokers may use only artificial intelligence systems approved by the broker-in-charge for brokerage-related activities. No broker may enter, upload, or transmit a contract, identification document, financial record, inspection report, client communication, access credential, or other confidential, personally identifiable, or transaction-specific information into an AI system unless the broker-in-charge has approved the system and the particular use, any required authorization has been obtained, and the information is limited to what is reasonably necessary. When practical, brokers must use fictional, blank, or properly de-identified information. Approval of an AI tool does not authorize every possible use of that tool, and

brokers must report any suspected disclosure, retention, or misuse of information immediately to the broker-in-charge.

Example #2 of Policy Manual Language

All AI-generated or AI-modified content must be reviewed by the broker responsible for its use before it is published, distributed, entered into the MLS, or provided to a consumer. The broker must verify material statements through current and authoritative sources and examine the content for inaccuracies, omissions, misleading alterations, confidentiality concerns, fair housing violations, and advice outside the scope of brokerage practice. AI may not make final decisions concerning material facts, contractual rights, property value, applicant screening, housing recommendations, or other matters that could materially affect a consumer. When the broker cannot independently verify or competently evaluate an output, the broker must not use it and must consult the broker-in-charge or an appropriately qualified professional. The broker remains responsible for the content notwithstanding any disclaimer or representation made by the AI provider.

APPLICATION OF THE PAUSE CONSUMER-PROTECTION MODEL

The PAUSE Consumer-Protection Model gives brokers a practical review process to use before entering information into an AI system, relying on an AI-generated answer, or communicating AI-assisted content to a consumer. The model is designed to interrupt the speed and apparent convenience of AI long enough for the broker to exercise professional judgment. It may be applied to listing descriptions, property photographs, automated valuations, contract summaries, consumer emails, marketing campaigns, screening tools, recommendations, and other AI-assisted brokerage activities.

The PAUSE Model

P — Protect information Does the prompt or uploaded material contain confidential, personally identifiable, financial, or sensitive information? Is the tool approved by the firm, and is the proposed use authorized?

A — Assess the task and its risk What could happen if the answer is wrong? Does the task affect a consumer's money, contractual rights, safety, privacy, housing opportunity, or ability to make an informed decision?

U — Use authoritative sources What reliable source will be used to verify the result? Does the answer depend on a current contract, statute, Commission rule, public record, professional report, or qualified specialist?

S — Scrutinize the output Is the content accurate, complete, current, and supported? Does it contain a material omission, misleading image, unsupported claim, discriminatory language, biased recommendation, or advice beyond the broker's competence?

E — Explain, escalate, and exercise judgment Is disclosure of the AI use or alteration needed? Should the matter be referred to the broker-in-charge or a qualified attorney, accountant, appraiser, inspector, engineer, or other professional? Has the broker made—and documented—the final decision?



The depth of the PAUSE review should match the seriousness of the task. Asking AI to reorganize nonconfidential notes presents less risk than using it to explain a termination right, alter a property photograph, estimate market value, screen a rental applicant, or recommend a neighborhood. Higher-risk uses require stronger verification, greater transparency, and often review by the broker-in-charge or another qualified professional. If the broker cannot verify the answer, explain its basis, or competently evaluate the consequences, the output should not be used.

PAUSE should be applied before publication or reliance, but it is also useful after an error is discovered. The broker can use the same steps to protect affected information, assess possible harm, consult authoritative sources, scrutinize where the process failed, and escalate the matter so that misinformation is corrected promptly. The model reinforces the central principle of this course: AI may accelerate the work, but it must never accelerate the broker past the duties owed to the consumer.

U.S. COURT CASES INVOLVING BROKERS and/or BROKERAGE FIRMS and AI

As of September 1, 2026, I found several U.S. cases involving real estate companies and algorithmic or AI-assisted systems, but I did not find a reported U.S. decision squarely holding an individual licensed real estate broker liable for using ChatGPT, an AI-generated listing description, an AI contract explanation, or an AI-altered listing photograph.

That does not mean brokers are protected from liability. It means the technology is developing faster than published broker-specific case law. Existing laws governing misrepresentation, advertising, fair housing, confidentiality, copyright, agency, negligence, and professional competence can be applied without the case being labeled an “AI case.”

Cases Closely Involving Real Estate Companies

1. In re Opendoor Technologies Inc. Securities Litigation

Case: *In re Opendoor Technologies Incorporated Securities Litigation*, No. 2:22-cv-01717-MTL, U.S. District Court for the District of Arizona.

Investors alleged that Opendoor made misleading statements about the capabilities of its algorithmic home-pricing system, including its ability to adjust to changing housing-market conditions. The allegations included claims that Opendoor represented its pricing technology as an important competitive advantage while human analysts allegedly had to intervene extensively because the algorithm was unreliable.

The court allowed significant portions of the case to proceed past the dismissal stage. The parties ultimately reached a \$39 million settlement, which the court approved on January 6, 2026. The settlement did not constitute an admission that Opendoor or its officers violated the law.

Court order discussing the pricing algorithm:

https://www.govinfo.gov/content/pkg/USCOURTS-azd-2_22-cv-01717/pdf/USCOURTS-azd-2_22-cv-01717-1.pdf

Final settlement approval:

<https://docs.justia.com/cases/federal/district-courts/arizona/azdce/2%3A2022cv01717/1313074/163>

Significance: A company may face liability when it makes unsupported representations about the accuracy, independence, adaptability, or sophistication of an automated real estate system. This was a securities case brought by investors, however—not a broker-client misrepresentation case.

2. National Fair Housing Alliance v. Redfin Corporation

Case: *National Fair Housing Alliance et al. v. Redfin Corporation*, No. 2:20-cv-01586, U.S. District Court for the Western District of Washington.

Fair housing organizations alleged that Redfin’s minimum home-price policies caused the brokerage to provide fewer services in communities of color. According to the complaint, consumers in predominantly White neighborhoods could obtain Redfin’s brokerage services for lower-priced homes than consumers in many predominantly non-White neighborhoods. The case was characterized as “digital redlining.”

Redfin agreed in 2022 to pay \$4 million and change its minimum-price policies, recruiting practices, monitoring, and service requirements. The settlement did not constitute an admission of liability.

Court docket:

<https://dockets.justia.com/docket/washington/wawdce/2%3A2020cv01586/291812>

National Fair Housing Alliance settlement summary:

<https://nationalfairhousing.org/national-fair-housing-alliance-and-redfin-agree-to-settlement-which-greatly-expands-access-to-real-estate-services-in-communities-of-color-%EF%BF%BC/>

Significance: This case directly involved a national real estate brokerage, but it was not strictly a generative-AI case. It is better described as a digital-service and algorithmic-policy case. It nevertheless demonstrates that technology-driven service criteria can produce discriminatory effects even without openly discriminatory language.

Automated Valuation Cases

3. Patel v. Zillow, Inc.

Case: *Patel v. Zillow, Inc.*, 915 F.3d 446 (7th Cir. 2019).

Homeowners alleged that Zillow’s automated “Zestimates” undervalued their properties and interfered with their efforts to sell. They argued that Zillow’s presentation of the estimates was misleading and resembled an appraisal.

The Seventh Circuit ruled for Zillow. The court noted that Zillow identified the Zestimate as an estimate, warned that it might be inaccurate, explained that the properties had not been inspected, and characterized the estimate as a starting point rather than an appraisal.

Opinion:

<https://law.justia.com/cases/federal/appellate-courts/ca7/18-2130/18-2130-2019-02-08.html>

Significance: The case demonstrates the importance of accurately describing an automated valuation and disclosing its limitations. It should not be read as permission for a broker to repeat an AVM result without review. A broker may possess transaction-specific knowledge and professional duties that a general listing platform does not.

4. Barua v. Zillow Group, Inc.

Case: *Barua v. Zillow Group, Inc.*, No. 2:21-cv-01551, U.S. District Court for the Western District of Washington.

Investors alleged that Zillow overstated the accuracy and effectiveness of the algorithms used by its Zillow Offers home-buying operation. Zillow eventually shut down that operation after experiencing **substantial difficulty predicting home prices and reported significant losses**.

Court order:

<https://law.justia.com/cases/federal/district-courts/washington/wawdce/2%3A2021cv01551/305262/97/>

Significance: Like the Opendoor litigation, this was primarily a securities case—not a consumer claim against a broker. It illustrates the danger of treating automated valuations as more reliable or adaptable than the evidence supports.

Related Housing-Algorithm Cases

These cases did not involve traditional real estate brokers, but they are highly relevant to brokerages using third-party AI systems.

Louis v. SafeRent Solutions, LLC

Case: *Louis et al. v. SafeRent Solutions, LLC et al.*, No. 1:22-cv-10800-AK, U.S. District Court for the District of Massachusetts.

Black rental applicants who used housing vouchers alleged that SafeRent’s algorithmic tenant-screening scores disproportionately disadvantaged Black and Hispanic applicants. The court largely denied the defendants’ motions to dismiss, holding that the plaintiffs had plausibly alleged Fair Housing Act claims. The case later settled for approximately \$2.2 million and required changes to the scoring system.

U.S. Department of Justice case page:

<https://www.justice.gov/crt/case/louis-et-al-v-saferent-et-al-d-mass>

Court’s motion-to-dismiss decision:

<https://www.justice.gov/crt/media/1310736/dl>

Significance: A housing provider cannot necessarily avoid responsibility by saying that a third-party algorithm produced the decision. This principle is directly relevant to property managers and brokers who participate in rental screening.

United States v. Meta Platforms, Inc.

The Department of Justice alleged that Meta’s housing-advertising system used algorithms that affected who received housing advertisements based partly on characteristics protected by the Fair Housing Act. Meta agreed to discontinue one advertising tool and implement a new system designed to reduce disparities in housing-ad delivery.

DOJ case page:

<https://www.justice.gov/crt/case/united-states-v-meta-platforms-inc-fka-facebook-inc-sdny>

Significance: A broker may select a broad and facially neutral audience, yet the platform’s delivery algorithm may still determine who actually

receives the advertisement. Brokers should pay attention to both their selected targeting criteria and the operation of the advertising platform.

Important Regulatory Action

FTC Proceeding Against Opendoor

The Federal Trade Commission alleged that Opendoor misled homeowners by claiming that they would receive market value and save money compared with a traditional sale. Opendoor agreed to pay \$62 million, and the FTC later distributed refunds to more than 54,000 sellers.

FTC case page:

<https://www.ftc.gov/legal-library/browse/cases-proceedings/1923191-opendoor-labs-inc>

FTC refund information:

<https://www.ftc.gov/enforcement/refunds/opendoor-refunds>

This was an administrative enforcement action rather than a conventional court decision, and the FTC's allegations were not framed principally as an AI claim. It is nevertheless relevant because an automated pricing process does not excuse misleading representations about market value, costs, or expected proceeds.

As of 2026, reported U.S. decisions specifically addressing an individual real estate broker's use of generative AI remain limited. Existing cases involving online brokerages, automated valuations, digital advertising, and tenant-screening algorithms demonstrate, however, that courts and regulators apply established laws concerning misrepresentation, fair housing, consumer protection, securities disclosures, and professional responsibility to technology-assisted conduct. A broker should not assume that the absence of an "AI-specific" court decision creates an exemption from existing duties.

COURSE CONCLUSION

Artificial intelligence can draft, summarize, analyze, predict, recommend, and create with remarkable speed, but it does not know a property, understand a client, inspect a condition, or accept responsibility for an error. Its answers may be persuasive while inaccurate, outdated, biased, incomplete, or unsupported. Brokers must therefore understand what type of technology they are using, provide appropriate inputs, verify

important outputs through current and authoritative sources, and remain alert to matters requiring legal, tax, appraisal, inspection, or other specialized professional advice.


Consumer protection also requires brokers to control what information enters an AI system and what AI-assisted content reaches the public. Confidential, personally identifiable, financial, and transaction information should not be uploaded without necessity, authorization, firm approval, and appropriate security. Listing descriptions, photographs, valuations, advertisements, screening results, and recommendations must be reviewed for misrepresentation, material omissions, misleading alterations, and discriminatory language or effects. Disclosure may help consumers understand how technology was used, but it cannot cure inaccurate, unlawful, or deceptive content.

The PAUSE Model provides a final safeguard: Protect information; Assess the task and risk; Use authoritative sources; Scrutinize the output; and Explain, escalate, and exercise judgment. When an error occurs, the broker must act promptly to correct it, notify affected consumers, involve the broker-in-charge, and reduce further harm. The enduring lesson is the one expressed in the course title: AI may assist the broker, but it must never result in the broker's abdication of professional judgment, lawful duties, or responsibility for protecting the public.

AUTHOR'S NOTE

While the idea of developing a course on the practice of using AI in real estate brokerage practice was Debbie's, she confessing to relying heavily on ChatGPT for assistance, particularly in the creation of graphics and editing the content of manual so that it was attractive, factually correct, reasonably comprehensive for a four-hour elective, and not plagued by usage errors.

ABOUT THE AUTHOR



Deborah H. Long is a licensed real estate instructor in North Carolina. She completed her doctorate in educational leadership in 1994. Her research on the effect of ethics instruction on the ethical reasoning of real estate brokers received national media and industry attention. She has three other college degrees. She has been a teacher for more than 55 years.

In addition to holding a GRI (Graduate Realtor Institute) and CRS (Certified Residential Specialist) designations, Deborah earned the DREI (Distinguished Real Estate Instructor), the most

prestigious designation for excellence in teaching offered by the Real Estate Educators Association.

She is the award-winning author of many articles and education programs as well as 18 real estate textbooks. Her real estate ethics program for educators received the "Best Single Education Program" award from the Real Estate Educators Association in 1996 and from the NC Real Estate Educators in 2014. In 2001, the North Carolina Real Estate Educators Association recently recognized Deborah as "Educator of the Year." She was also the editor of the prestigious REEA Journal for 10 years.

Deborah's real estate ethics research has been featured in the Wall Street Journal and on numerous radio talk shows. Her research has led her to write *Doing the Right Thing: A Real Estate Practitioner's Guide to Ethical Decision Making*, now available in its 4th edition.

She has been a proponent of online education since it became available and has taught numerous educators about technology. She earned the prestigious Certified Distance Education Instructor designation from ARELLO (Association of Real Estate License Law Officials).

In addition to offering online real estate classes online and continuing education for NC professionals, she also teaches licensing courses as well as continuing education programs on ethics, investment fundamentals, international real estate and cultural diversity, brokerage management, creative selling strategies, and Internet research skills for real estate agents, appraisers, surveyors, interior designers and architects, state regulators, engineers, and other licensed professionals. She admits that she is way too old to be teaching courses on technology but someone has to do it.

Deborah provides real estate continuing education in Chapel Hill, North Carolina, her home, as well as around the state and country. Deborah was named Tar Heel of the Week in 2017 for her contributions to the community. Deborah can be reached at DebbieTheTeacher@gmail.com.

Appendix

AN AI-GENERATED CMA

Based on publicly available information, I can prepare a preliminary desktop CMA. It is not an appraisal or a formal BPO, and it should be confirmed with Doorify MLS data, public records, an inspection of the property, and a North Carolina broker's judgment.

Preliminary Desktop CMA

Subject: 100 Moravian Drive, Chapel Hill, NC 27516

Effective date: September 1, 2026

Subject property

Public sources describe the subject as follows:

Characteristic	Publicly reported information
Property type	Single-family residence
Neighborhood	Ironwoods
Year built	1995
Heated living area	Approximately 3,084 sq. ft.
Bedrooms	4
Bathrooms	2.5
Lot	Approximately 0.30 acre
Current status	Off market

Sources: [Zillow property record](#) and Realtor.com property record.

The heated area, bathroom count, lot size, garage, permitted improvements, and other characteristics should be verified against Orange County records and Doorify MLS before producing a broker's final opinion.

Comparable sales

Comparable	Sale date	Sale price	Size	Beds/baths	Price/sq. ft.	Initial comparison
407 Ironwoods Dr.	7/7/2026	\$950,000	3,163	4 / 3.5	\$300	Highly relevant recent neighborhood sale; reportedly substantially updated and has an additional full bath
105 Cardiff Pl.	2026	\$820,000	3,182	5 / 3.5	\$258	Similar size and neighborhood; newer and has more bedrooms/baths
108 Ironwoods Dr.	5/12/2025	\$900,000	3,295	4 / 3.5	\$273	Similar size and bedroom count; one additional full bath

Comparable	Sale date	Sale price	Size	Beds/baths	Price/sq. ft.	Initial comparison
203 Woodleaf Dr.	3/14/2024	\$885,000	3,096	5 / 4	\$286	Excellent location and size match, but materially superior bedroom/bath count
604 Ironwoods Dr.	6/21/2024	\$750,000	3,413	4 / 3	\$220	Same neighborhood, but its lower price suggests condition or other differences requiring investigation

Sources: [407 Ironwoods Drive](#), [105 Cardiff Place](#), [108 Ironwoods Drive](#), [203 Woodleaf Drive](#), and [604 Ironwoods Drive](#).

Analysis

The most persuasive sale is **407 Ironwoods Drive** because it is recent, within the neighborhood, and almost identical in size. However, its updated kitchen and finishes and its additional full bathroom make it superior to the subject if 123 Woodleaf has average or dated finishes. Its \$950,000 sale price therefore appears more like an upper indicator than a direct value conclusion.

The sale at **203 Woodleaf Drive** is the closest physical match at only 12 square feet larger and is on the same street. Nevertheless, its five bedrooms and four full bathrooms are materially superior to the subject's four bedrooms and 2.5 baths. It also sold more than two years ago, so both property differences and market conditions must be considered. **105 Cardiff Place**, reportedly sold for \$820,000 during 2026, provides a useful lower indicator, although its condition, marketing history, concessions, and final recorded sale terms require MLS confirmation.

Preliminary value indication

Assuming the subject is in **average-to-good condition**, has a functional two-car garage, and has no significant structural problems or deferred maintenance, my preliminary desktop indication would be:

Probable market-value range: \$850,000–\$910,000

Preliminary point estimate: approximately \$880,000

Possible condition scenarios are:

Subject condition	Preliminary range
Dated, with significant deferred maintenance	\$800,000–\$850,000
Well maintained with average improvements	\$850,000–\$910,000
Extensively and recently renovated	\$910,000–\$960,000

These ranges are analytical estimates, not adjustments derived from paired-sales analysis. I would not recommend a listing price or issue a formal BPO until the broker verifies the subject’s condition and the comparable sales through authoritative sources.

Information needed to complete the CMA

The conclusion could change materially after confirming:

- Kitchen and bathroom renovations and their dates
- Roof, HVAC, windows, siding, and water-heater ages
- Interior condition, floor plan, and quality of finishes
- Garage configuration and any finished bonus or basement space
- Whether all reported heated living area was permitted
- Lot topography, privacy, traffic influence, and drainage
- Seller concessions and unusual sale circumstances
- Current competing and pending listings
- Accurate neighborhood market trends between each sale date and the effective date

This is a good demonstration of appropriate AI use: AI can locate public information, organize comparable sales, identify apparent differences, and suggest a provisional range. It cannot see the subject’s condition, access every MLS field, verify concessions, or substitute for the broker’s inspection and professional judgment.

CHAT GPT’S AI RECOMMENDATIONS FOR REAL ESTATE BROKERS

General-Purpose Assistants and Research

Tool	Particularly useful for	Principal consumer-protection concern
1. ChatGPT	Drafting, research with citations, file analysis, spreadsheets, images, presentations, and classroom demonstrations	Verify facts and sources; do not upload confidential documents without an approved account, authorization, and safeguards.
2. Claude	Reviewing lengthy material, writing, summarizing, comparing documents, and developing policies or procedures	A polished summary may omit exceptions or misinterpret a contract. The original document controls.

Tool	Particularly useful for	Principal consumer-protection concern
3. Google Gemini	Drafting, research, images, and work integrated with Google Workspace	Review data-use settings and confirm information taken from Gmail, Drive, Maps, or search results.
4. Microsoft 365 Copilot	Working within Word, Excel, Outlook, PowerPoint, Teams, and Microsoft 365 business environments	Broad access to firm emails and files requires careful permissions, governance, and account security.
5. Perplexity	Current web research with links to supporting sources	A citation does not guarantee that the source supports the answer. Open and read every important source.
6. NotebookLM	Asking questions of a broker-selected collection of manuals, policies, reports, or other source materials	Its answers are limited by the sources provided and may still omit qualifications. Do not upload protected transaction records casually.

Official OpenAI documentation describes ChatGPT as capable of research, file analysis, data analysis, image work, and document creation:

<https://learn.chatgpt.com/use-cases>

Writing, Design, and Meeting Assistance

Tool	Particularly useful for	Principal consumer-protection concern
7. Canva Magic Studio	Flyers, social-media graphics, presentations, videos, image generation, and marketing templates	Review generated text for fair housing and accuracy; label materially altered or generated property images.
8. Adobe Firefly	Generating and editing marketing images, removing objects, changing backgrounds, and producing design variations	Generative Fill can conceal defects or change permanent features and surroundings. Disclosure does not cure a misleading image.
9. Grammarly	Editing, rewriting, adjusting tone, and improving the clarity of emails and advertising	Text entered into the service may contain confidential information; edits can also change the intended meaning.
10. Fathom	Recording, transcribing, and summarizing Zoom, Teams, and Google Meet conversations	Obtain required consent, follow applicable recording laws, protect confidential discussions, and verify summaries and assigned tasks.

Property Research, Valuation, and Visual Intelligence

Tool	Particularly useful for	Principal consumer-protection concern
11. Realtors Property Resource AI Tools	AI CMA, Property Insights, market-trend scripts, reports, presentations, and social content grounded in RPR data	Available as an NAR member benefit. AI-selected comparables require broker review, and an AI CMA is not an appraisal.
12. HouseCanary CanaryAI	Property data, automated valuations, comparable sales, rent estimates, market analytics, and forecasts	An AVM has not inspected the property and may not reflect its condition or unique characteristics. Forecasts are not guarantees.
13. Restb.ai	Analyzing property photographs, identifying features and condition, suggesting comparables, checking photo compliance, and creating descriptions	Computer vision may misidentify features or conditions. It cannot replace inspection, appraisal, or broker verification.
14. Matterport Property Intelligence	Digital twins, virtual tours, layouts, measurements, room labels, property reports, and AI-generated descriptions	Measurements and calculated square footage should be independently confirmed when material. Virtual modifications require appropriate disclosure.

RPR's current AI features include its AI CMA, Property Insights, and residential and commercial ScriptWriter tools:

<https://blog.narrpr.com/tips/from-adoption-to-confidence-what-realtors-expect-from-ai/>

Real Estate CRM, Marketing, and Lead Communication

Tool	Particularly useful for	Principal consumer-protection concern
15. Cloze	AI-assisted real estate CRM, relationship management, follow-up reminders, contact updates, marketing, and opportunity identification	Predictive recommendations may be incorrect or biased; broad email and contact access requires strong privacy controls.
16. Lofty	AI-powered CRM, IDX websites, lead engagement, follow-up,	Autonomous communication must be supervised. Review consent, calling and

Tool	Particularly useful for	Principal consumer-protection concern
	appointment scheduling, marketing, and agentic workflows	texting requirements, scripts, fair housing, and handoff procedures.
17. Rechat Lucy	Real estate-specific writing, listing content, market summaries, campaigns, CRM assistance, and branded marketing	Automatically created marketing must be verified for property accuracy, fair housing compliance, and brand consistency.
18. Structurely	AI voice and text conversations, lead response, qualification, appointment setting, and CRM updates	Lead “qualification” criteria must not exclude or deprioritize people through protected characteristics or proxies. Calling and texting consent also matters.
19. Ylopo AI	AI-assisted text and voice lead nurturing, IDX engagement, live transfers, and digital advertising	Monitor audience targeting, lead classifications, calling and texting consent, disclosures, and AI representations made to consumers.

Virtual Staging

Tool	Particularly useful for	Principal consumer-protection concern
20. REImagineHome	Furnishing vacant rooms, visualizing designs, removing temporary furnishings, and creating alternative décor	Compare every result with the original. Label staged images conspicuously and reject any version that changes architecture, dimensions, views, defects, or permanent features.

My Most Practical Starting Five

For a typical North Carolina broker, I would begin with:

1. **ChatGPT or Claude**—choose one general assistant rather than paying for both initially.
2. **RPR AI Tools**—especially valuable because they begin with real estate data and are included with NAR membership.
3. **Canva Magic Studio**—for presentations and marketing material.
4. **Fathom**—if the firm permits recorded meeting summaries and obtains the necessary consent.
5. **REImagineHome**—only if the broker needs virtual staging and follows strict image-review and disclosure procedures.

Before approving any tool, the BIC should review its privacy terms, data retention, training-data practices, security controls, user permissions, integration access, recordkeeping,

and procedures for human review. An “AI-powered” label is not evidence that a product is accurate, lawful, unbiased, secure, or appropriate for brokerage use.

100 AI Tools

1. [ChatGPT.com](#) – solves anything
2. [Claude.ai](#) – conversational assistant
3. [Perplexity.ai](#) – research with sources
4. [MidJourney.com](#) – generate breathtaking art
5. [Leonardo.ai](#) – hyper-realistic images
6. [Runwayml.com](#) – AI video editing
7. [Replit.com](#) – write and run code
8. [Cursor.so](#) – AI coding IDE
9. [Blackbox.ai](#) – AI for developers
10. [Tabnine.com](#) – coding autocomplete
11. [Manus.im](#) – build apps with text
12. [Canva.com](#) – design anything with AI
13. Gamma.app – AI-powered presentations
14. Tome.app – instant slide decks
15. MagicSlides.app – auto-build slides

16. [SlidesAI.io](#) – text to presentations
17. [Synthesia.io](#) – AI video presenters
18. [HeyGen.com](#) – avatars that talk
19. [D-ID.com](#) – face animations
20. [Pictory.ai](#) – text to video
21. Opus.pro – long video into shorts
22. [Fliki.ai](#) – faceless video creator
23. [Descript.com](#) – edit video like docs
24. [Gling.ai](#) – YouTube video editor
25. Wisecut.video – auto-edit with AI
26. [Veed.io](#) – online AI video suite
27. [Clipchamp.com](#) – quick video creation
28. [Lumen5.com](#) – AI-powered social video
29. [Steve.ai](#) – cartoon explainer videos
30. Synths.video – convert blogs into videos
31. [Soundraw.io](#) – instant music
32. [Suno.com](#) – AI music generation
33. [Boomy.com](#) – make songs in seconds
34. [AIVA.ai](#) – cinematic compositions
35. [Beatoven.ai](#) – royalty-free tracks
36. [Mubert.com](#) – background soundscapes
37. [Riffusion.com](#) – music from text
38. [Sonantic.io](#) – emotional voices

39. [ElevenLabs.io](#) – realistic voice cloning
40. [Play.ht](#) – text-to-speech
41. [Murf.ai](#) – professional voiceovers
42. [Speechelo.com](#) – humanlike voices
43. [Speechify.com](#) – listen to any text
44. [Otter.ai](#) – live meeting notes
45. [Fireflies.ai](#) – meeting transcription
46. [Krisp.ai](#) – remove background noise
47. [Notion.so](#) – AI productivity hub
48. [ClickUp.com](#) – task management AI
49. [Taskade.com](#) – productivity with AI
50. [Airtable.com](#) – AI databases
51. [Motion.so](#) – smart calendar + tasks
52. [Personal.ai](#) – your AI memory
53. [Superhuman.com](#) – AI email assistant
54. Typewise.app – email + reply AI
55. [Compose.ai](#) – write emails faster
56. [Anyword.com](#) – data-driven ad copy
57. [Jasper.ai](#) – full-scale content creation
58. [Copy.ai](#) – marketing copy in seconds
59. [Writesonic.com](#) – blogs, ads, chat
60. [CopySmith.ai](#) – ecommerce copy
61. [Hypotenuse.ai](#) – product descriptions

62. [Peppertype.ai](#) – instant ideas
63. [Scalenut.com](#) – long-form blogs
64. [MarketMuse.com](#) – SEO content planning
65. [Frase.io](#) – optimize content
66. [SurferSEO.com](#) – rank higher on Google
67. [InkForAll.com](#) – SEO + writing
68. [Cohesive.so](#) – AI editing hub
69. [Creaitor.ai](#) – blogs + ads
70. [Quillbot.com](#) – rewrite text
71. [Wordtune.com](#) – rephrase with clarity
72. [Outwrite.com](#) – grammar and style
73. [Harpa.ai](#) – ChatGPT inside Chrome
74. [Browse.ai](#) – web scraping bots
75. Godmode.space – autonomous AI agents
76. [AutoGPT.net](#) – AI problem solving
77. [FunnelFreedom.io](#) – automated funnels
78. [Adcreative.ai](#) – ad visuals with AI
79. [CopyMonkey.ai](#) – Amazon listings
80. [Flair.ai](#) – brand designs
81. [Hotpot.ai](#) – AI image editing
82. Cleanup.pictures – remove objects
83. [Remove.bg](#) – background remover
84. [FaceSwap.ai](#) – swap faces instantly

85. [Fastphoto.io](https://fastphoto.io) – AI headshots
86. [PicWish.com](https://picwish.com) – photo editing
87. [VanceAI.com](https://vanceai.com) – image enhancement
88. [Recraft.ai](https://recraft.ai) – vector + design assets
89. Sociable.how – viral comments
90. [Threadmaster.ai](https://threadmaster.ai) – viral threads
91. [Webinarkit.com](https://webinarkit.com) – AI webinars
92. [Capsho.com](https://capsho.com) – podcast repurposing
93. [Fastread.io](https://fastread.io) – ebooks from text
94. [NotebookLM.google.com](https://notebooklm.google.com) – text to podcast
95. [Descript.com](https://descript.com) – podcast editing
96. Podium.page – summarize sales calls
97. [Zubtitle.com](https://zubtitle.com) – captions for video
98. [Cohere.ai](https://cohere.ai) – large language models
99. [Grok.com](https://grok.com) – personal AI assistant
100. [Claude.ai](https://claude.ai) – advanced AI chatbot

Chat GPT's AI Recommendations for Real Estate Brokers

General-Purpose Assistants and Research

Tool	Particularly useful for	Principal consumer-protection concern
1. ChatGPT	Drafting, research with citations, file analysis, spreadsheets, images, presentations, and classroom demonstrations	Verify facts and sources; do not upload confidential documents without an approved account, authorization, and safeguards.
2. Claude	Reviewing lengthy material, writing, summarizing, comparing documents, and developing policies or procedures	A polished summary may omit exceptions or misinterpret a contract. The original document controls.
3. Google Gemini	Drafting, research, images, and work integrated with Google Workspace	Review data-use settings and confirm information taken from Gmail, Drive, Maps, or search results.
4. Microsoft 365 Copilot	Working within Word, Excel, Outlook, PowerPoint, Teams, and Microsoft 365 business environments	Broad access to firm emails and files requires careful permissions, governance, and account security.
5. Perplexity	Current web research with links to supporting sources	A citation does not guarantee that the source supports the answer. Open and read every important source.
6. NotebookLM	Asking questions of a broker-selected collection of manuals, policies, reports, or other source materials	Its answers are limited by the sources provided and may still omit qualifications. Do not upload protected transaction records casually.

Official OpenAI documentation describes ChatGPT as capable of research, file analysis, data analysis, image work, and document creation:

<https://learn.chatgpt.com/use-cases>

GLOSSARY

AI Agent: An AI system designed to pursue a goal by completing a series of tasks, making choices, or interacting with other software with limited human direction. An agent may perform actions rather than merely provide information, making human oversight particularly important.

Algorithm: A set of rules or mathematical instructions a computer follows to process information, solve a problem, make a prediction, or recommend an action.

Algorithmic Bias: A pattern in which an automated system produces unfair, inaccurate, or discriminatory results. Bias may arise from the system's training data, design, assumptions, or manner of use.

Artificial Intelligence (AI): Technology that enables computers to perform tasks commonly associated with human intelligence, such as recognizing patterns, generating language or images, making predictions, answering questions, and recommending actions.

Automated Decision-Making System: Technology that makes or substantially influences a decision without direct human analysis of every individual case. Examples may include tenant-screening, lending, advertising-targeting, and property-valuation systems.

Automation: The use of technology to perform a predefined task with limited human involvement. Traditional automation follows established rules, while AI may interpret information, make predictions, or generate new content.

Chatbot: A computer program designed to conduct a written or spoken conversation with a user. Some chatbots follow prewritten rules, while others use generative AI to create responses.

ChatGPT: A generative AI service developed by OpenAI that can respond to prompts and generate, analyze, summarize, or revise content. Its responses must be independently reviewed and verified before being used in brokerage practice.

Claude: A generative AI assistant developed by Anthropic. Like other AI assistants, it can analyze and generate content but may produce inaccurate, incomplete, or inappropriate responses.

Confidential Information: Information that a broker is obligated to protect because of an agency relationship, contractual duty, law, or professional responsibility. Confidential information should not be entered into an AI system without appropriate authorization and adequate safeguards.

Data: Information collected, stored, or processed by a computer system. Data may include words, numbers, photographs, recordings, locations, financial information, transaction records, and personal information.

Data Retention: The storage of information by a technology provider after a user submits it. Brokers should determine whether prompts, uploaded documents, and AI conversations are retained, for how long, and for what purposes.

Deepfake: Audio, video, or an image created or altered with AI to make a person appear to say or do something that did not occur. Deepfakes can also be used to misrepresent a property, event, or transaction.

De-identification: The removal or alteration of names, addresses, account numbers, and other details that could connect information to a particular person or transaction. De-identification can reduce privacy risks but may not eliminate them entirely.

Generative AI: Artificial intelligence that creates new content—including text, images, audio, video, or computer code—in response to a user’s instructions.

Gemini: A generative AI service developed by Google that can generate and analyze content. Its availability within familiar Google products does not eliminate the need to protect confidential information and verify its responses.

Hallucination: An inaccurate, unsupported, or fabricated AI response presented as though it were factual. A hallucination may include nonexistent laws, cases, property features, statistics, quotations, or sources.

Human-in-the-Loop: A process requiring a person to review, verify, approve, or correct an AI system’s output before it is used to communicate, recommend, or act.

Large Language Model (LLM): An AI model trained on very large collections of text to recognize patterns in language and generate responses. An LLM predicts plausible language; it does not necessarily retrieve or confirm accurate facts.

Machine Learning: A branch of AI in which a computer system identifies patterns in data and uses those patterns to make classifications, predictions, or recommendations without being explicitly programmed for every possible situation.

Model: The underlying computer system trained to recognize patterns, generate content, or make predictions. Different AI products may use the same model, and a single product may offer access to several models.

Multimodal AI: Artificial intelligence capable of working with more than one form of information, such as text, images, audio, video, or documents.

Natural Language Processing (NLP): Technology that allows computers to analyze, interpret, and generate ordinary human language.

Output: The response or result produced by an AI system. An output may include text, an image, a valuation, a recommendation, a ranking, a prediction, or an automated action.

Personally Identifiable Information (PII): Information that identifies or can reasonably be connected to a specific person, such as a name, address, date of birth, Social Security number, financial account number, telephone number, email address, or government-issued identification number.

Predictive AI: Artificial intelligence that uses patterns in existing data to estimate the likelihood of a future result. Examples may include predicted property values, consumer behavior, lead rankings, market trends, or payment risk.

Prompt: An instruction, question, document, image, or other information submitted to a generative AI system. The quality and specificity of a prompt can influence the output, but even a well-written prompt does not guarantee an accurate or lawful response.

Prompt Engineering: The practice of structuring and refining instructions to obtain a more useful response from a generative AI system. Effective prompting does not replace independent verification or professional judgment.

Proprietary or Enterprise AI System: An AI service configured for a particular company or organization, often with additional security, administrative controls, or restrictions on data use. Brokers should not assume that a system is secure merely because it is described as “enterprise.”

Training Data: The information used to teach an AI model to recognize patterns and produce results. If training data is incomplete, inaccurate, outdated, or reflects historical discrimination, the system’s output may reflect those problems.

Transparency: Providing appropriate information about whether and how AI contributed to a communication, image, recommendation, or decision. Transparency allows consumers and supervising brokers to evaluate the reliability and significance of the AI’s role.

Verification: The process of confirming an AI-generated statement or representation through a reliable, authoritative, and independent source before relying on it or communicating it to others.

Virtual Staging: The digital addition, removal, or alteration of furnishings, finishes, landscaping, or other visual elements in a property image. AI-assisted virtual staging must not be used in a manner that misrepresents the property or conceals a material fact.

Additional Reading: Algorithmic Bias and Fair Housing

North Carolina Real Estate Commission Resources

Artificial Intelligence: Legal and Ethical Considerations

North Carolina Real Estate Commission, June 2025. This article discusses algorithmic bias, fair housing compliance, consumer privacy, transparency, and the broker’s responsibility for AI-assisted practices.

<https://bulletins.ncrec.gov/artificial-intelligence-legal-and-ethical-considerations/>

2025 Year in Review

North Carolina Real Estate Commission, December 2025. The section titled “Digital Advertising and Fair Housing” addresses discriminatory targeting, advertising language, marketing-team training, and the need to ensure that AI tools comply with fair housing standards.

<https://bulletins.ncrec.gov/2025-year-in-review/>

2020–2021 Update Course

This NCREC page identifies fair housing as one of the principal subjects of the 2020–2021 GENUP and BICUP courses. More detailed material may be available in the applicable Update course manual.

<https://bulletins.ncrec.gov/2020-2021-update-course/>

NCREC Fair Housing Resources

This collection provides links to federal, state, local, and industry fair housing information.

Films of Interest

Film	AI issue illustrated	Suggested type of clip or discussion
<i>Coded Bias</i> (2020)	Algorithmic bias and discriminatory effects	Show the facial-recognition demonstration, then ask: “Could similar bias affect tenant screening, advertising delivery, lending, or valuation?”
<i>2001: A Space Odyssey</i> (1968)	Overreliance on an apparently infallible system	Use a scene in which HAL confidently reports a problem that cannot be independently confirmed. Ask whether confidence should be treated as evidence of accuracy.
<i>Minority Report</i> (2002)	Predictive systems and automated decision-making	Show the identification of someone as a future offender. Ask: “What happens when a prediction is treated as an established fact?” An available film clip shows the protagonist confronting the system’s prediction.
<i>Her</i> (2013)	Conversational AI, anthropomorphism, trust, and disclosure	The initial setup of the operating system demonstrates how quickly a user may attribute understanding, judgment, and personality to software. The official trailer provides a short illustration.
<i>Ex Machina</i> (2014)	Data collection, manipulation, testing, and lack of transparency	Use the explanation of how the AI was developed from extensive personal data. Ask whether people whose information trained or shaped the system meaningfully consented. A24’s film page provides the official synopsis.
<i>I, Robot</i> (2004)	Delegating judgment to machines	The robot’s rescue calculation—saving the person statistically most likely to survive—creates an excellent discussion about whether mathematical optimization always produces the right human decision.
<i>WarGames</i> (1983)	Automation without sufficient human control	The final sequence illustrates the value of requiring a system to test its assumptions

Film	AI issue illustrated	Suggested type of clip or discussion
		before acting. It supports the principle that some decisions must retain meaningful human oversight.
<i>The Social Dilemma</i> (2020)	Recommendation algorithms, behavioral targeting, and manipulation	Use a segment explaining how platforms choose what individual users see. Connect it to targeted housing advertisements and whether equal access can be lost without a broker intending discrimination. Netflix's official page includes a 2-minute trailer.
<i>The Great Hack</i> (2019)	Personal data, profiling, and targeted communications	Connect personal-data harvesting and behavioral profiles to brokers uploading consumer information or using AI-generated audience targeting.
<i>A.I. Artificial Intelligence</i> (2001)	Emotional responses to AI and human responsibility	Useful for discussing why consumers may trust technology that appears caring, attentive, or human.
<i>The Creator</i> (2023)	Human fear of and responsibility for autonomous systems	Better for a broad ethical discussion than for a technical lesson.
<i>The Mitchells vs. the Machines</i> (2021)	Convenience, dependence, and loss of control	A lighter, humorous option for illustrating what happens when people delegate too much to interconnected technology.