

ARTIFICIAL INTELLIGENCE • DWELLING CLAIMS

When Algorithms Sit at the Table

What AI is quietly doing to appraisal, investigation, and settlement of dwelling claims — and what to watch for.

Most homeowners picture a human adjuster walking the loss, climbing the roof, stepping through water-damaged rooms, scoping the structural repairs, and deciding what the carrier will pay. For routine dwelling claims that picture is fading. Carriers now lean heavily on artificial intelligence (AI) software that learns patterns from large amounts of data and then makes predictions or recommendations, and on its narrower cousin, machine learning (ML), which is the technique that lets such software improve by example. These tools touch the file from the moment a claim is reported through the day a check is cut.

AI has real benefits in claims. It speeds up triage, flags suspicious activity, and standardizes routine valuations. But it also raises problems that homeowners, lawyers, neutrals, and adjusters need to understand. This article focuses on three places where those problems matter most for dwelling losses: appraisal, investigation, and settlement.

A note on acronyms before we begin. P&C means property and casualty insurance. FNOL, “first notice of loss” is the call, photo upload, or app submission that opens a claim. ACV stands for “actual cash value,” the depreciated value of damaged property; RCV is “replacement cost value,” what it costs to rebuild or repair without depreciation. EUO is “examination under oath,” a sworn statement carriers may

require during investigation. SIU is the carrier's "special investigations unit," which handles suspected fraud. A few more acronyms appear below; each is defined in context.

Appraisal: when the number on the page is a model's guess

Most homeowners policies contain an appraisal clause. Each side picks an appraiser; if the two disagree on the amount of loss, a neutral umpire decides. The process assumes two independent humans making independent judgments.

That assumption is wobbling. Many carrier appraisers now rely on vendor software that scores roof condition from aerial imagery, classifies damage from drone or ground-level photos, and generates a repair estimate using AI-assisted estimating platforms. The appraiser may not know which roof slopes the model flagged as "wear," why it called certain marks "mechanical" rather than "hail," or how it priced the underlayment, decking, flashing, and code-required upgrades.

Consider Robert and Patricia, who file a claim after a spring hailstorm. Their architectural asphalt shingle roof is ten years old. The carrier orders aerial imagery scored by an AI tool, which classifies the roof as showing "minimal storm impact, predominantly granule loss consistent with normal aging."

The repair estimate, generated by a photo-based estimating platform, comes back at \$4,800, a partial-slope shingle repair and a few interior drywall patches. Robert and Patricia's roofing contractor walks the roof, documents hail bruising on all four

slopes, soft hits on the metal vents and gutters, and a damaged ridge cap, and writes a scope under the local building code with matching-product requirements. His estimate is \$38,400.

The carrier's appraiser can produce the imagery report but cannot explain why specific slopes scored low, how the “wear” classification was applied, or whether the imagery actually analyzed was captured before the storm. The umpire is now being asked to weigh a documented manual estimate against a number whose underpinnings no one at the table can fully defend.

The concern is not that the model is necessarily wrong. It is that appraisal depends on independent reasoning, and a number produced by a black box software whose internal decisions cannot be readily explained, is not the same as an independent judgment.

A second worry is scope creep. Appraisal in most states is limited to the amount of loss, not coverage or causation. When an AI tool tags shingle damage as “wear” rather than “hail,” or labels certain marks “pre-existing,” it has quietly answered a causation question that may sit outside the appraisal panel's authority.

Investigation: synthetic evidence and the disappearing paper trail

Investigation is where generative AI or AI that creates new images, text, audio, or video rather than just analyzing existing data has changed the landscape most quickly. The same tools that let a small business design a logo on a phone can fabricate a photograph of hail strikes on a shingle, a contractor's invoice, or an engineer's wind-uplift report.

Picture Karen, a dwelling claims adjuster. She opens a homeowner's claim for wind and hail damage. The insured uploads close-up photographs of hail “strikes” on shingles, interior ceiling-stain photos attributed to a roof leak, and a PDF “engineer's letter” attesting to wind-driven shingle uplift across the south slope.

Each item looks clean; however, several are simply not true. The shingle photos were generated by a consumer image tool trained on real roofing imagery; the interior ceiling-stain photo was pulled from a public real estate listing of a different house; the engineer's letter was drafted by a chatbot using a real licensed engineer's name and license number scraped from a state board directory.

Karen's first clue is subtle: the metadata (the hidden information a camera writes into every photo, sometimes called EXIF data (short for Exchangeable Image File Format)) shows the shingle photos were created on a device that does not match the insured's phone, and the engineer's letter has a PDF creation date that postdates the supposed inspection by two weeks.

Unfortunately, synthetic claim evidence is no longer rare. Carrier SIUs and the National Insurance Crime Bureau (NICB) have reported steady increases in fabricated damage photos, contractor estimates, repair invoices, and engineer reports submitted in support of dwelling claims.

The marginal cost of producing a convincing-looking damage file has collapsed.

Three specific risks deserve attention.

The first is image fabrication or ‘diffusion models’ which is a type of AI that builds images from random noise. It can produce hail-strike close-ups, water-stain ceilings, and storm-damaged siding that look authentic to the human eye. Detection tools exist, but they produce false positives often enough that no single tool should be treated as proof.

The second is document forgery. A fraud ring no longer needs a graphic designer when a chatbot can draft a plausible roofing contractor estimate, an itemized water-mitigation invoice, or an engineer's wind-uplift report in seconds. As a result, cross-referencing contractor and roofer licenses, Professional Engineer (PE) numbers against state boards, and business tax IDs is more important than ever.

The third is what technologists call “Prompt Injection”. Some carriers now run AI or agentic tools that read incoming PDFs, repair estimates, engineer reports, supplemental scopes and write a summary for the adjuster. A bad actor can hide instructions inside a submitted document, text the adjuster cannot see but the AI reads, that change how the carrier's own software describes the claim. The adjuster reads a summary that has been quietly manipulated.

Investigation in the AI era therefore depends less on any single document and more on triangulation or comparing what the claim file says against National Oceanic and Atmospheric Administration (NOAA) hail and wind records for the date of loss, dated pre-loss aerial imagery from providers such as EagleView or Nearmap, local building permits, neighboring claim activity, and prior-loss databases like ISO ClaimSearch. When documents are cheap to fabricate, independent data sources are what carry the weight.

Settlement: the invisible authority range

Settlement is where AI's influence is hardest to see from the outside. Many carriers now use algorithmic tools to set “reserves” (the money the insurer sets aside for an expected payout) and to recommend a settlement range to the adjuster on first-party dwelling claims. The adjuster has discretion within that range but typically needs supervisor approval to exceed it.

Consider a mediation: The homeowner's dwelling sustained extensive wind and water damage from a covered storm. The contractor's RCV estimate is \$187,000. The carrier has paid roughly \$62,000 over six months of supplements and partial denials. At mediation, the adjuster's settlement authority tops out at \$94,500. The number was not set by a claims manager weighing the file; it was generated by a property claims evaluation model that scored the photographs, line-item estimate, the dwelling's age and ZIP code, and prior claim history. The adjuster cannot explain it because she did not set it, and the model's reasoning is not exposed to her. The mediator senses a wall but cannot tell whether it is policy, posture, or software.

Two concerns flow from this. The first is “anchoring”, the well-documented tendency of negotiators to settle near the first credible number on the table. When that number is algorithmic and the other side does not know it, the negotiation is asymmetric in a way the policyholder cannot see.

The second is the duty of good faith. Most states require an insurer's first-party claim decisions to rest on a reasonable, individualized evaluation of the loss. A range generated by a model the adjuster cannot explain may not satisfy that duty if the case proceeds to litigation, especially in jurisdictions that recognize bad-faith damages for unreasonable handling of dwelling claims.

What can be done

None of this argues for banning AI from dwelling claims. Used well, these tools speed payment to honest homeowners and catch fraud that would otherwise slip through. The question is how to keep the benefits while limiting the harms.

For carriers, the emerging standard, reflected in the National Association of Insurance Commissioners' (NAIC) Model Bulletin on the Use of Artificial Intelligence Systems by Insurers, now adopted by roughly half the states, is straightforward in concept. Inventory the AI tools in use, including roof-scoring imagery, photo-based estimating, and claim evaluation models. Test them for accuracy and bias. Then document how decisions are made. Train adjusters to treat AI outputs as inputs, not verdicts. Keep records that allow the carrier to produce a specific, individualized explanation for any adverse decision, not “the system said so.”

For neutrals, appraisers, umpires, mediators, arbitrators, a few habits help. Ask, on the record, whether either side's position relies on an algorithmic output, and which one. Require the proponent to be able to explain it, including the capture date of any aerial imagery and the methodology behind any roof-condition score. Decide foundation before substance when authenticity of evidence is challenged. Adopt a short AI-use disclosure at the preliminary conference so both sides know what is in play.

For homeowners and their counsel, the most useful step is to insist on the methodology behind any AI-driven valuation or denial. Ask for the model name, version, inputs, image-capture dates, and any human review. Pair a field expert, a licensed contractor, a roofer, or an engineer who actually walks the property, with someone who can critique the data. Preserve original storm photographs with their metadata and request the same from the carrier when imagery is the basis of the decision.

For everyone, the underlying point is the same. AI in dwelling claims is not a future problem. It is in the file today, sometimes openly and sometimes invisibly. Honest claims and honest defenses both benefit from a process where the role of the algorithm is on the table, where humans retain meaningful judgment, and where the evidence, on both sides, can be tested.