

Rockport's approach to Responsible AI in Commercial Real Estate Tech

By Will Trepp, Co-Founder and Co-CEO

Rockport's approach to Responsible AI in Commercial Real Estate Tech

By Will Trepp, Co-Founder and Co-CEO

Across commercial real estate (CRE), the same internal discussions are playing out in lending institutions, valuation & advisory firms, asset managers, and property owners. The pull of AI is obvious. The frontier models are wildly capable, and the potential leverage across underwriting, valuation, servicing, and asset management is tantalizing.

But the worries are real too.

The following paper discusses some of these worries and separates true legitimate concerns from the ghosts that don't really exist.

PICKING THE LLM MODEL IS THE EASY PART

There is a natural, false assumption that picking the right LLM (the large language model behind tools like ChatGPT and Claude) or Agent provider is the hard part. It's not.

To use a car as an analogy: the frontier model is the engine. The product that your end users eventually use is the car. Buyers buy a car for many reasons—safety, reliability, brand, track record, high-quality service, the list goes on. Very few people select it based on the engine manufacturer.

The same is true here. Anyone can license the same frontier model we do. What you're actually buying is everything built around it, and that is where the real work lives.

Part I: How we build with AI

Our engineers work with AI coding tools every day to develop the code that underlies our platform. The gains this has given us and will give us in the years to come are real.

To get technical for a second: LLMs are probabilistic, not deterministic. This means that they will function like a human and give you a slightly different result each time you provide it

with a prompt. This is not a bug; this is how they are designed.

As a result, to write reliable, repeatable, auditable math at scale, a reliable method is to use the LLMs to write normal, ordinary (deterministic) code. In summary, the LLMs are helping our engineers produce the same code they always have, but orders of magnitude faster over time.

This ordinary code then needs to be analyzed as we have always analyzed human-developed code: reviewing it, grading it for architectural and upgradability soundness, security reviewing it, testing it for correctness and alignment with the specs, analyzing for bound cases, performance testing it, version controlling it, etc. This discipline is exactly what an internal build would need to recreate from nothing, which is a topic worth its own paper (click [here](#) and [here](#))

The above is meant to address several misconceptions:

“If you are building software with LLMs that are inherently probabilistic, aren’t your calculations fundamentally unstable?”

False. The LLMs produce deterministic code—the same as always.

“Since “AI does everything”, can’t I just vibe code your product?”

False. LLMs produce work like humans always have. Theoretically, you can build an army of product agents, coding agents, testing agents, and QA agents in the same way we are. But this work is not easy, it is not quick, it is not cheap, and it is not your business. As someone who has run a 100+ person (and now a lot more agents) software company, it is very, very far from simple.

“But I’ve seen an impressive demo from my internal IT department or a start-up; they must be close since it looks good?”

False. A vibe-coded prototype is probably 2% or less of the actual work to develop a system, let alone the post-development carry cost. Click [here](#) to read more.

Part II: How we deploy AI in the product

The other discipline is the AI our clients actually touch. This is the capability now starting to appear inside the product. Our instinct here has been restraint until we mastered it for our own use. If we don’t eat our own food first, we don’t serve it to clients. We have not scattered AI across every screen because it would let us “AI-wash” our product (i.e., the deceptive practice of making a product look heavily AI-driven and falsely cutting-edge).

Instead, we are adding AI intelligently, where it earns its place, in the workflows where it measurably helps our users better perform their jobs.

We are also making sure these features come with clear opt-ins and opt-outs, because you, as a customer, should get to choose where and when you use AI.

Part III: How we protect your data

Long before AI, we adopted a simple principle: We are custodians of your data, not owners of it.

That principle predates AI at Rockport for decades. As a leadership team with great experience in data, we decided to take the path less traveled. It was important to us to keep the client as the client, not as the product in a play to provide software to extract data for resale. This approach is highly unusual today, but what we believe is right.

This is why we do not benchmark or re-sell client data.

We build our AI practices with the same ethical model. We keep your data in our protected ecosystem and work with models/security architectures that prevent 3rd-party learning on your data. The opt-in and opt-out mentioned above govern whether you use the AI features,

not whether your data is used improperly. The most common question we hear is whether our AI trains on client data. The answer is no. Inference, the term for using a finished model to do work rather than building it, runs inside firewalled enterprise channels under our controls. The model reads your information in place and produces its output there, and nothing leaves to be absorbed elsewhere. Picture the difference between handing a brilliant outside contractor your files to take home and having that same contractor work through the files in your office, under supervision, carrying nothing out when they leave.

It delivers the same capability without the same risk.

In regulated markets, that difference is important, and we engineered for it on purpose.

WHAT'S THE GOAL: PEOPLE MAKE THE DECISIONS, AGENTS DO THE LEGWORK

As AI shifts from writing you an answer to going and taking action, the line that matters most is between legwork and consequential decisions.

We believe an agent can and should make procedural choices in service of a legwork task, such as gathering and extracting data from a loan file, cross-checking different sources, and putting the results on screen for review and approval.

This is genuinely useful and saves users real time and clients real money.

The decision on what to do with that data, and its consequences, belongs to the person who reviews and approves the path forward.

We are not building toward autonomous decision-making tools as we don't think it's the right path forward. We say so plainly because serious buyers are right to ask, and different companies will have different answers.

SAME OLD DISCIPLINE, NEW TECHNOLOGY

Our products are workflow and analytic systems based on decades of CRE experience, trillions of dollars of real closed loans and properties analyzed, and the ability to build math at scale.

The reason the AI transition does not rattle us is that the obligations behind our systems are the same ones we have always carried.

Rockport has been building secure, scalable CRE technology for decades, through many market cycles, with a high bar for client service and data ethics the whole time. Our founder, Rick Trepp, has been a pioneer in the industry for more than forty years.

The technology in front of us is new and exciting—but in reality, changes nothing.

The speed of the game gets faster. The expectations are higher. But the tooling is likewise more capable. Same game, just at lightning speed. These kinds of environments favor the experienced players who have seen it before and don't have to make it up as they go.



Will Trepp

Co-Founder and Co-CEO

Will Trepp is the Co-Founder & Co-CEO of The Rockport Group, where he leads SaaS solutions managing over \$2T in CRE assets. With a background in Computer Science and CRE finance, he focuses on the intersection of technology, real estate, and data.