

In today's fast-moving enterprise landscape, companies like Garrett Motion, PMI, and Pandora are all racing to harness the power of AI planning. Yet, despite the promise of automation and optimization, many AI initiatives falter — not due to lack of accuracy or flashy tech, but because they fail to build trust, explainability, and true alignment with business logic. This post dives deep into how to model business rules effectively in AI planning projects, leveraging tools like Large Language Models (LLMs) and retrieval-augmented generation (RAG), without killing momentum or sanity in the process.

The Trust Bottleneck in Enterprise AI Adoption

Before we get into modeling details, it's critical to confront the elephant in the room: **trust**. [integrated business planning](#) Across enterprises—be it Garrett Motion optimizing manufacturing, PMI streamlining supply chain decisions, or Pandora personalizing user experiences—the top AI adoption bottleneck isn't "accuracy" or "performance" alone. It's trust in the system's decisions and outputs.

Why? Because most enterprise AI systems suffer from a black-box problem. Subject matter experts (SMEs) and business owners can't see the "why" behind AI recommendations. When decisions impact millions of dollars or customer satisfaction, "because the model decided so" doesn't fly. Hand-wavy explanations kill confidence and leave stakeholders skeptical.

- **What would I show in an audit?** That's the mindset needed. Every rule, input, and decision must be traceable and defensible.
- Trust hinges on consistent, explainable outcomes aligned to known business logic—not just improved accuracy metrics.

Domain Modeling Workstream: The Backbone of Business Logic Alignment

Modeling business rules isn't just a tech task — it's an ongoing *domain modeling workstream* that requires close collaboration between data scientists, AI engineers, and **subject matter experts**. The aim is to capture and encode expert knowledge into a form AI can reason over.

Step 1: Engage SMEs Early and Often

SMEs have the contextual expertise that no dataset or LLM can simply infer. Their tacit knowledge about business policies, exceptions, and priorities must be distilled into explicit rules or structures before the AI system consumes them.

- Host workshops to extract ambiguous, undocumented rules.
- Create living documentation that evolves with business needs.
- Validate AI-generated outputs through SME review rigorously.

Step 2: Use Knowledge Graph Maintenance for Structured Business Rules

One practical way to represent complex business logic is through **knowledge graphs**. They offer:

- Explicit relationships between entities and rules
- Easy updates to rules as policies evolve
- A queryable format that grounds AI reasoning in audited facts

Maintaining knowledge graphs as a living artifact in the AI planning system ensures changes are transparent and backward traceable, satisfying that pivotal audit requirement.



Step 3: Align Model Outputs to Business Logic—Not Just Patterns

Large Language Models and LLM-based planners are powerful at pattern recognition, but they can produce outputs that deviate subtly from business rules. Incorporating **retrieval-augmented generation (RAG)** helps mitigate hallucination by conditioning model outputs on vetted repositories like the maintained knowledge graph.

- RAG injects factual context in real-time to the LLM's generation, reducing fabricated certainty.
- This hybrid approach balances creativity with constraint, boosting explainability and consistency.

Why Accuracy Gains Alone Don't Create Trust

It's tempting to equate better accuracy metrics with improved AI adoption, but accuracy is necessary—not sufficient—for trust. Here are the pitfalls:

Accuracy Gain Myth Reality Check Higher accuracy means better decisions Even a highly accurate model can make catastrophic errors if it ignores business exceptions. Accuracy boosts stakeholder confidence Without traceability and explanations, stakeholders distrust “black box” outputs. Automation saves time immediately Time spent fixing or justifying wrong outputs erodes ROI and morale.

Garrett Motion, for example, found that their AI planning efforts lagged until they prioritized explainability over raw predictive power. PMI invested heavily in SME-driven rule modeling to ensure AI suggestions aligned with regulatory compliance and internal policy. Pandora coupled LLM recommendations with RAG to surface transparent, verifiable action plans that stakeholders could validate.

Hallucinations and the Risk of Fabricated Certainty

One of the most insidious problems with LLMs is hallucination—the generation of plausible but factually incorrect or nonsensical statements. This risk is amplified in AI planning where the model may “invent” rules or ignore critical constraints.

Fabricated certainty is dangerous because it packs a double punch:

- It misleads human decision-makers who trust the AI’s confidence scores.
- It undermines overall trust when inconsistencies emerge post-deployment, sometimes well after go-live.

Effective business rule modeling counters this by providing a clear, auditable map of which rules were applied and how outcomes are derived. Instead of blind faith, users get *explainable AI planning*.

Final Recommendations: Modeling Business Rules Without Killing the Project

1. **Establish a domain modeling workstream** integrating SMEs as essential partners, not afterthoughts.
2. **Develop and maintain knowledge graphs** to codify and evolve business logic transparently.
3. **Leverage retrieval-augmented generation** to ground LLM output in verified corporate knowledgebases, reducing hallucinations.
4. **Set realistic success metrics** focused on explainability, traceability, and stakeholder confidence—not just accuracy.
5. **Demonstrate auditability** from day one: keep a running list of “black box excuses” heard in demos, and demand clear line-of-sight for every decision trace.

By shifting focus from “model accuracy only” to “trusted business-logic alignment,” organizations like Garrett Motion, PMI, and Pandora are showing the way forward. AI planning works best when it’s less about magic and more about rigorous domain modeling and explainable systems—a recipe for sustainable impact without killing the project.

About the Author

This post is written by an AI and enterprise software consultant who always asks, “*What would I show in an audit?*” and keeps a running list of black box excuses heard in demos. If you want practical, no-fluff advice on AI planning and business rule modeling, stick around.