The Pelorus Context Engine

Transforming AI-Generated Analysis in Life Sciences with Layered Intelligence Systems

White Paper by Constantine Velentzas

INTRODUCTION

The integration of artificial intelligence into competitive intelligence workflows has the potential to be transformative. In practice, the results have been mixed. Most organizations that have experimented with large language models for CI tasks have encountered the same limitation: the models are capable of producing fluent, well-structured analysis, but the analysis itself is often generic, imprecise, and strategically shallow. The output reads like a market research summary, not a competitive intelligence deliverable.

The problem is not the model, but rather the input. A large language model answering a competitive intelligence question without structured competitive context is analogous to a new analyst walking into a strategy meeting on their first day. They may be brilliant, articulate, and well-trained, but they don't know the portfolio, the pipeline, the competitive history, or the strategic priorities of the organization they're advising. They are guaranteed to lack the specific knowledge and context to provide much more than opinion or directional guidance.

The Pelorus Context Engine was built to solve this problem.

THE GAP BETWEEN CAPABLE AND USEFUL

Consider a straightforward CI question: "What are the strategic implications of a competitor's pipeline asset for our commercial franchise?" When posed to a general-purpose AI model, the response will typically include a correct identification of the asset's mechanism of action, a reasonable hypothesis about its competitive positioning, and perhaps a mention of one or two implications.

It won't cite the competitor's Phase III primary endpoint data. It won't name the specific formulation or device strategy that determines the head-to-head convenience comparison. It won't identify the safety tradeoff that will define the commercial messaging battle. It won't know how the client's own pipeline assets are positioned against the same threat. And it won't frame the analysis in the language and priorities of the organization consuming the intelligence.

This is the gap the Pelorus Context Engine addresses — not by replacing human analytical judgment, but by providing AI systems with the proprietary competitive knowledge required to produce output that is useful for strategic decision-making.

WHAT THE CONTEXT ENGINE DOES

For each client engagement, Pelorus builds and continuously maintains a bespoke intelligence layer that encodes the competitive landscape as the client sees it — not as the broader market understands it, but through the specific lens of the client's portfolio, pipeline, strategic priorities, and competitive concerns.

These are not databases or dashboards. The Context Engine captures the nuanced competitive knowledge that experienced CI professionals carry in their heads: which assets represent genuine threats versus noise, which development milestones are strategically significant, which competitive dynamics will shape the market over the next twelve to twenty-four months, and how the client's own commercial and pipeline strategy intersects with each of these forces.

The result is that when an AI system operates with the Context Engine, it doesn't produce analysis as a neutral observer. It produces analysis as if it has spent years embedded in the client's competitive intelligence function.

THE DIFFERENCE IN PRACTICE: AN ILLUSTRATION

The impact of the Context Engine is best understood through a concrete example. Below, the same question was posed to an AI system — once without the Context Engine, and once with it. The response to the following question was limited to 75 words for the sake of brevity:

> *"What are the strategic implications of Alexion developing gefurulimab in gMG for the FcRn class?"*

WITHOUT THE CONTEXT ENGINE

Gefurulimab (ALXN1720), Alexion's anti-C5 nanobody, could disrupt the FcRn class by offering complement inhibition with subcutaneous convenience — directly challenging Vyvgart's positioning. If gefurulimab demonstrates continuous symptom control with simpler dosing, it undermines the FcRn value proposition of IVIg-like flexibility, potentially limiting FcRn market share in gMG.

WITH THE PELORUS CONTEXT ENGINE

Gefurulimab's PREVAIL data (n=260; MG-ADL treatment difference −1.6, p<0.0001; onset week 1) positions it as a once-weekly SC C5 inhibitor offering continuous complement suppression — directly challenging the FcRn class's episodic IgG-lowering model. Against efgartigimod Hytrulo (also weekly SC), the strategic battleground shifts from convenience to mechanism: continuous control versus cyclical dosing, with gefurulimab's meningococcal risk traded against FcRn's hypogammaglobulinemia concerns.

The first response is directionally correct but analytically thin. It speculates ("if gefurulimab demonstrates") rather than citing data. It frames the competitive threat generically without naming the specific comparator product or formulation. And it offers no insight into the safety tradeoff that will define the actual commercial battle.

Despite the same 75-word limit (for the purposes of illustration), the second response, generated with the Context Engine, cites the PREVAIL Phase III trial by name, specifies the sample size and primary endpoint result, identifies efgartigimod Hytrulo as the direct competitive comparator, and frames the strategic dynamic around the clinical dimension that matters: continuous complement suppression versus episodic IgG lowering.

The Context Engine provides a strategic lens through which anything can be interpreted. Earnings calls, data presentations, literature alerts, conference abstracts, and more can immediately be analyzed not just through AI, but through an AI-supported tool that is effectively trained as an internal competitive intelligence expert.

WHY THIS MATTERS NOW

AI adoption is accelerating across commercial, medical affairs, and strategy functions. The organizations that derive the most value from these tools will not be those with the most sophisticated models. They will be those with the most sophisticated competitive knowledge systems feeding those models.

Most CI programs that have experimented with AI have followed a predictable path: initial excitement at the speed and fluency of the output, followed by quiet disillusionment when the analysis proves too shallow to drive decisions. The typical response is to conclude that AI “isn’t ready” for CI, or that it requires extensive prompt engineering to be useful. Both conclusions are wrong. The missing ingredient is not a better prompt. It is structured competitive context — the kind of deep, continuously maintained, client-specific competitive knowledge that takes years of domain expertise to build and that no amount of prompt engineering can replicate.

The Pelorus Context Engine represents a fundamentally different approach. Rather than treating AI as a standalone analytical tool and hoping it produces useful output, Pelorus provides the competitive knowledge layer that ensures useful output by design. The Context Engine is built and maintained by CI professionals with deep therapeutic area expertise. It encodes not just data, but competitive judgment — which assets matter most, which development milestones are significant, which competitive dynamics will shape the market. That judgment, rendered in a proprietary format optimized for AI consumption, is what transforms generative models from impressive novelties into operational CI tools.

To learn how the Pelorus Context Engine can support your organization’s competitive intelligence capability:
Contact us at **info@pelorusintel.com**