Evolutionary System-Prompt Optimization for LLM Agents in Competitive Market Simulations

Ossama Zaroual, Federico Carucci, Leonardo Mascagni, Francesco Bertolotti
School of Industrial Engineering, LIUC - Carlo Cattaneo University, Castellanza, Italy
Intelligence, Complexity and Technology Lab (ICT Lab), LIUC - Carlo Cattaneo University

Large Language Models (LLMs) are increasingly deployed in interactive multi-agent environments, including in complex systems simulations. In such settings, agent behavior is highly sensitive to system prompt design, yet prompt design is typically manual and heuristic-based. This creates a methodological issue, especially when developing LLM-enhanced agent-based models.

We propose an evolutionary framework that automates system-prompt design for LLM-driven agents. System prompts are modeled as modular discrete objects: a prompt is assembled from short components (instructions, constraints, contextual facts), and a binary genome encodes which components are active. This generates an interpretable combinatorial search space and enables contextual behavior optimization without fine-tuning or access to internal model states.The prompt optimize prompts with a binary Genetic Algorithm (GA) using realized cumulative profit as the only fitness signal.

We evaluate the method in an agent-based market simulation, where firms repeatedly choose price and product quality while competing for probabilistic consumers. Consumer choice follows a bounded-rational softmax over quality-to-price utility, and production costs scale linearly with quality. Each GA generation is evaluated via a co-evolutionary market: the genome population is instantiated as a population of firms competing simultaneously, mirroring competitive selection and exposing strategies to an evolving set of opponents.

Across controlled experiments, the GA reliably converges to stable and interpretable prompt configurations in a compact 8-component space and remains robust in an augmented 14-component space containing neutral and adversarial instructions. Harmful components are rapidly eliminated, while neutral content shows transient drift and is pruned as selection concentrates near the optimum. Increasing LLM sampling temperature degrades convergence stability, consistent with stochastic model outputs corrupting the fitness signal. Optimized LLM firms achieve substantial profit gains over unoptimized prompt baselines and outperform transparent rule-based competitors under identical conditions.

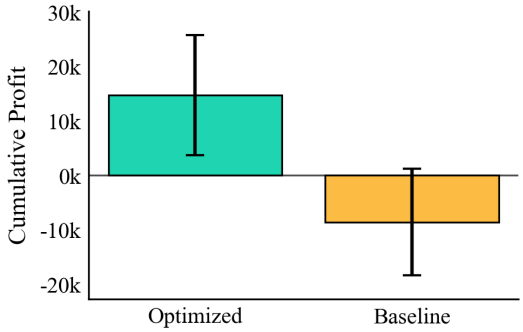

Figure 1: GA-optimized system prompts yield higher cumulative profit than unoptimized baselines.

References

[1] J. H. Holland, *Adaptation in Natural and Artificial Systems.* MIT Press, 1992.
[2] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley, 1989.
[3] M. Tesfatsion and K. L. Judd (eds.), *Handbook of Computational Economics, Vol. 2.* Elsevier, 2006.
[4] J. R. Epstein and R. Axtell, *Growing Artificial Societies.* Brookings Institution Press, 1996.