Scaling Multilayer Network Medicine with LLMs: A Structured Prompt-Engineering and RAG-Driven Case Study in ER+ and HER2+ Breast Cancer

Ahmed Abdeen Hamed[1] and Tamer E. Fandy[2] and Luis M. Rocha[1,3]

[1] School of Systems Science & Industrial Engineering, Binghamton University, Binghamton, NY, USA
ahamed1@binghamton.edu,rocha@binghamton.edu,
[2] Department of Pharmacology, Texas Tech University Health Sciences Center El Paso, El Paso, TX, USA
[3] Católica Biomedical Research Centre, Universidade Católica Portuguesa, Lisbon, Portugal

Abstract. We demonstrate how LLMs can scale construction of biomedical multilayer network models for drug repurposing, specialized in breast cancer types ER+ and HER2+. We built three layers: drug combinations from clinical trials, drug–target interactions from the biomedical literature, and protein interactions from KEGG signaling pathways. Structured prompt engineering extracted drug combinations, while a Retrieval-Augmented Generation (RAG) approach mined drug–target interactions. Protein interactions were parsed from the XML representation of the ER+ and HER2+ pathways in KEGG. Drug combinations were validated for self-consistency across (0, 1, and n-shot) prompts. Drug–target interactions were cross-checked using five open-source LLMs: Aya, Gemma, Phi4, Mistral, and Qwen2.5. Results show that LLMs can scale network medicine construction but vary greatly in reliability, underscoring the need for rigorous verification.

Keywords: LLMs, Multilayer Networks, Network Medicine, Drug Repurposing, Breast Cancer, Estrogen Receptor+

Introduction

Since the emergence of the era of ChatGPT and genAI models, there have been various efforts in utilizing such tools to accelerate the development of specialized models for network medicine. Pre-trained large language models (LLMs), offer a natural language interface for users to interact with, known as prompt-engineering. The ease in the use of natural language inspired scientists to creatively design prompts that perform tasks such as identification of biological entities and various relationships in the biomedical literature, and produce the results in a format that can be further processed algorithmically. Such tasks were known to be bottlenecks for the construction of multilayer network models for a

given disease, treatment options, and digital twins [2]. Here, we present an LLM-driven approach that scales the construction of multilayer network models for breast cancer types: ER+ and HER2+, using 5 underlying open-source LLMs: (Aya, Gemma3, Phi4, Mistral, and Qwen2.5.) [10].

Results:

We searched the ("clinicaltrials.gov") [1] using the search query ("drug combinations") to extract drugs that are approved or currently under the investigation for combinatorial potentials. This provides a starting point not only for what drugs can be combined as a therapy, but also offers individual drug members of the combinations to be explored for potential repurposing. The search query produced 2,496 unique trials, with a textual description that provided the actual details on the drug names and how they are combined. Performing a series of prompt engineering, with a self-consistency mechanism using Phi4, yielded 4,169 drugs and 2,621 combinations, and contributed the combination layer for our network medicine where the drugs were the nodes and the combinations were the edges. The next task performed was to collecting drug-target evidence from the biomedical literature. Specifically, we designed a RAG-driven information extraction approach against a dataset of 10,000 pubmed abstracts, which were retrieved by searching the pubmed database [12] for the search query ("breast cancer"). When analyzed by the LLMs, the task produced a varying number of drug-target pairs for the following models: Aya (4640), Mistral (6577), Phi4 (1155), Qwen2.5 (374), Gemma3 (3883). The outcome of this task offered the drug-pairs layer needed for our network where the drugs and targets were the nodes and the connection among the pairs formed the edges. To further investigate the novelty of the connections between drug-target pairs, we constructed a third layer from the signaling pathways for ER+ and HER2+ from the KEGG database [11]. This process produced layer of 79 entities (genes, proteins, and compounds), and 68 interactions, which offered the third layer for the network. The summary of outcome of these tasks is summarized in Table 1, while Figure 1 shows the percentage of yield.

Table 1: Summary of the information extraction tasks needed to construct the drug combination layer, drug-target layer, and the signaling pathways driven layer.

Self-Consistency Extraction		
Method / Model	# Drugs / Entities	# Combinations / Interactions
Phi4	4,169	2,621
KEGG signaling pathway	79	68
RAG-IE Extraction Across LLMs		
Method / Model	# Drugs / Entities	# Combinations / Interactions
Aya	4,640	—
Mistral	6,577	—
Phi4	1,155	—
Qwen2.5	374	—
Gemma3	3,883	—

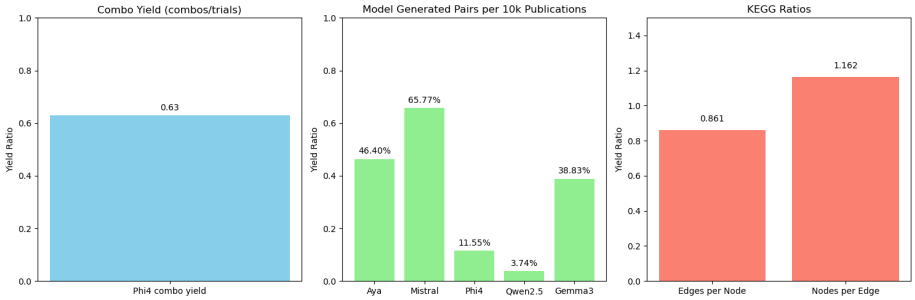

Fig. 1: Summary of LLM-driven drug repurposing and KEGG pathway network statistics showing edges per node and nodes per edge and network connectivity.

Methods

Our approach to identifying repurposable drug combinations builds on the principle that drugs, while designed for specific protein targets, they may also interact with unknown targets [4, 6, 8, 3, 7]. We systematically mine the biomedical literature to capture both known and potential drug–protein interactions. The resulting multilayered network medicine presents a platform that can be interrogated algorithmically to assess repurposing potential.

1. **Extraction of Drug Combinations from Clinical Trials.** Three prompts are executed: a zero-shot prompt, a one-shot prompt, and an n-shot prompt. Majority voting is then employed to determine the final set of drug combinations.
2. **Identification of Drug–Target Interactions.** A Retrieval-Augmented Generation (RAG) pipeline searches the biomedical literature drugs and their biological targets. The output is a set of high-confidence drug–protein associations for subsequent analysis.
3. **Construction of ER+ and HER2+ Signaling Pathways Network.** To accomplish this, we traversed the designated pathways. By parsing which proteins and interacting with others, we formed a directed protein-protein interaction network layer. The main task is to discriminate the novel drugs-target connections from already known ones, which signals the repurposing potential.
4. **AgenticAI for Drug Combination Repurposing Opportunities in Network Medicine** The agenticAI framework is under active development.

Conclusions and Future Direction

Our LLMs-driven networks offered a platform to study drug repurposing and combination therapy opportunities for ER+ and HER2+. The self-consistency usage of shot-based prompt engineering, combined with a majority vote criterion, offered a more reliable approach to identify drug combinations in the textual descriptions of clinical trials. We continue to follow prior verification protocols [5, 9] and implement agentic-AI agents to enable combination discoveries.

References

1. ClinicalTrials.gov: Clinicaltrials.gov. https://clinicaltrials.gov/. U.S. National Library of Medicine, accessed August 25, 2025
2. De Domenico, M., Allegri, L., Caldarelli, G., d'Andrea, V., Di Camillo, B., Rocha, L.M., Rozum, J., Sbarbati, R., Zambelli, F.: Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. npj Digital Medicine **8**(1), 37 (2025)
3. Gates, A.J., Correia, R.B., Wang, X., Rocha, L.M.: The effective graph reveals redundancy, canalization, and control pathways in biochemical regulation and signaling. Proceedings of the National Academy of Sciences **118**(12), e2022598,118 (2021). DOI 10.1073/pnas.2022598118. URL https://www.pnas.org/doi/abs/10.1073/pnas.2022598118
4. Gates, L.E., Hamed, A.A.: The anatomy of the sars-cov-2 biomedical literature: Introducing the covidx network algorithm for drug repurposing recommendation. Journal of medical Internet research **22**(8), e21,169 (2020)
5. Hamed, A.A., Crimi, A., Misiak, M.M., Lee, B.S.: From knowledge generation to knowledge verification: examining the biomedical generative capabilities of chatgpt. iScience **28**(6) (2025)
6. Hamed, A.A., Fandy, T.E., Tkaczuk, K.L., Verspoor, K., Lee, B.S.: Covid-19 drug repurposing: A network-based framework for exploring biomedical literature and clinical trials for possible treatments. Pharmaceutics **14**(3), 567 (2022)
7. Hamed, A.A., Fandy, T.E., Xindong, W.: Accelerating complex disease treatment through network medicine and genai: A case study on drug repurposing for breast cancer. In: Accepted with IEEE MedAI'24 (2024)
8. Hamed, A.A., Jonczyk, J., Alam, M.Z., Deelman, E., Lee, B.S.: Mining literature-based knowledge graph for predicting combination therapeutics: A covid-19 use case. In: 2022 IEEE International Conference on Knowledge Graph (ICKG), pp. 79–86. IEEE (2022)
9. Hamed, A.A., Rocha, L.M.: Protocol for assessing the utility of chatgpt in generating and verifying biomedical associations. Available at SSRN 5366821 (2025)
10. Inc., O.: Ollama: Open-source llm inference engine (2025). URL https://github.com/ollama/ollama
11. Kanehisa, M., Goto, S.: Kegg: Kyoto encyclopedia of genes and genomes. Nucleic acids research **28**(1), 27–30 (2000)
12. PubMed: Pubmed. https://pubmed.ncbi.nlm.nih.gov/. U.S. National Library of Medicine, accessed August 25, 2025