

## BACKGROUND

- Biomedical researchers increasingly have access to multimodal data, including electronic health records, pathology imaging, mutation, and other omics data modalities. Turning these complex data into discovery insights still requires sophisticated computational analysis and domain knowledge.<sup>4</sup>
- Many existing informatics tools remain siloed and are difficult to use for basic and translational researchers.
- This gap motivated the development of agentic AI approaches that enable research and translational applications through flexible and simple natural language instructions<sup>3</sup>.
- Precision Medicine MCP<sup>2</sup> democratizes precision medicine discovery and clinical applications using Model Context Protocol<sup>1</sup> (MCP)-based AI agents.

## METHODS

- Precision Medicine MCP is a multimodal agentic AI framework implemented as a federation of 19 specialized MCP servers with a total of 100+ specialized tools, organized into a 5-stage pipeline. These specialized MCP servers retrieve, interpret, and synthesize heterogeneous biomedical data including EHR, omics, imaging, and disease knowledge sources.
- Heterogeneous data sources are connected into a shared AI-accessible working environment via MCP federation using Python and FastMCP and bioinformatics tools.
- Natural language interactions (templated prompts) allow researchers to analyze individual clinical cases or de-identified patient cohorts without disease-specific informatics pipelines.

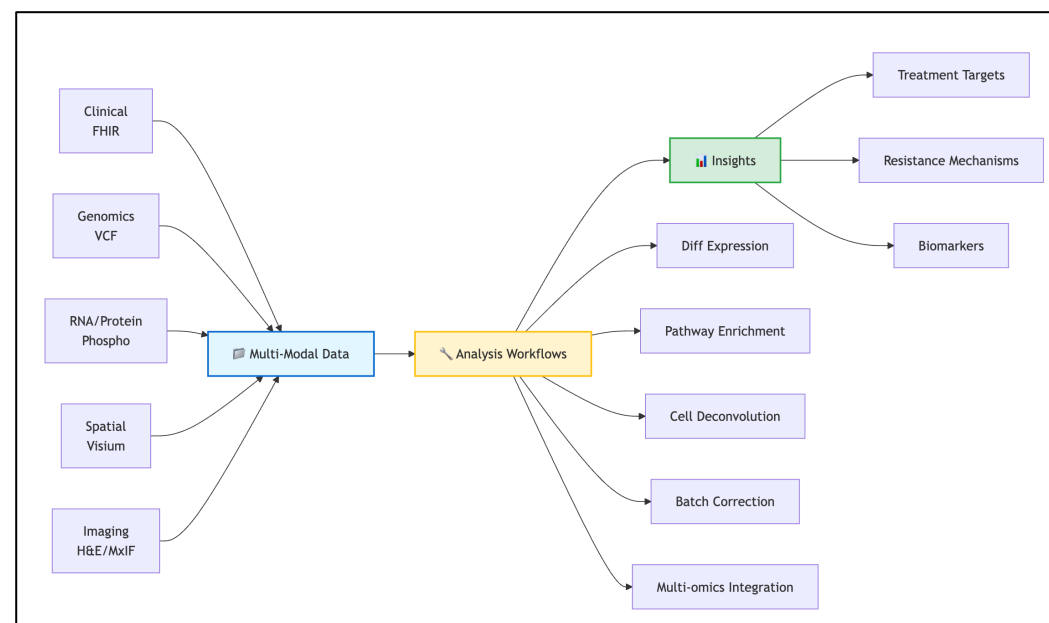


Figure 1. Precision Medicine MCP framework architecture: 19 federated MCP servers processing multi-modal data.

## RESULTS: OVARIAN CANCER TREATMENT DISCOVERY

- For synthetic PAT001 - Standard-of-care confirmed: Bevacizumab + olaparib (HRD+, BRCA1 germline mutation).
- Analysis done on synthetic data using templated prompts in a 5 stage AI pipeline which calls custom mcp servers and associated analysis tools. Final output is conciliated report for clinical team review
- Novel targets (not on standard NGS panels):** AKT2 amplification (log2 CN +1.15, OT score 0.373) → AKT inhibitor hypothesis; CCNE1 amplification (log2 CN +1.58, OT score 0.279) → CDK2 inhibition.
- Top upstream regulators: MYC (Stouffer Z = 4.2), STAT3 (Z = 3.8), CCND1 (Z = 2.9).




Figure 2. PAT001 (Ovarian cancer synthetic patient data and: multi-omics integration via MCP agents)

## RESULTS: BREAST CANCER TRIAL MATCHING &amp; CV RISK

- For synthetic PAT002 (ER+/PR+/HER2-IDC, Stage IIA, BRCA2 germline) — Standard-of-care confirmed (olaparib, letrozole + palbociclib, alpelisib + fulvestrant). **Three beyond-SOC hypotheses unreachable by standard workup:** inavolisib over alpelisib (2024 FDA approval); MYC triple therapy — ribociclib + inavolisib + fulvestrant (MYC Z=26.9); cold-TME conversion via YSAPLSSSL vaccine + CAF depletion + anti-PD-1 (CD8 T cells exhausted, not absent — viable for priming). Same architecture as PAT001. Zero disease-specific code changes.
- For synthetic PAT003 (female) - Cardiovascular prevention: Standard lipid panel was missing three key risk inputs — Lp(a) unmeasured, APOE genotype unknown, CAC score absent. The multi-agent framework identified these gaps and modelled synthetic values for each. Three CVD risk algorithms incorporating those synthetic inputs all converged above the 7.5% statin threshold (Reynolds: 14.3%, Framingham: 12.0%, ACC/AHA ASCVD: 10.3%), suggesting PAT003 was **undertreated due to incomplete workup**.

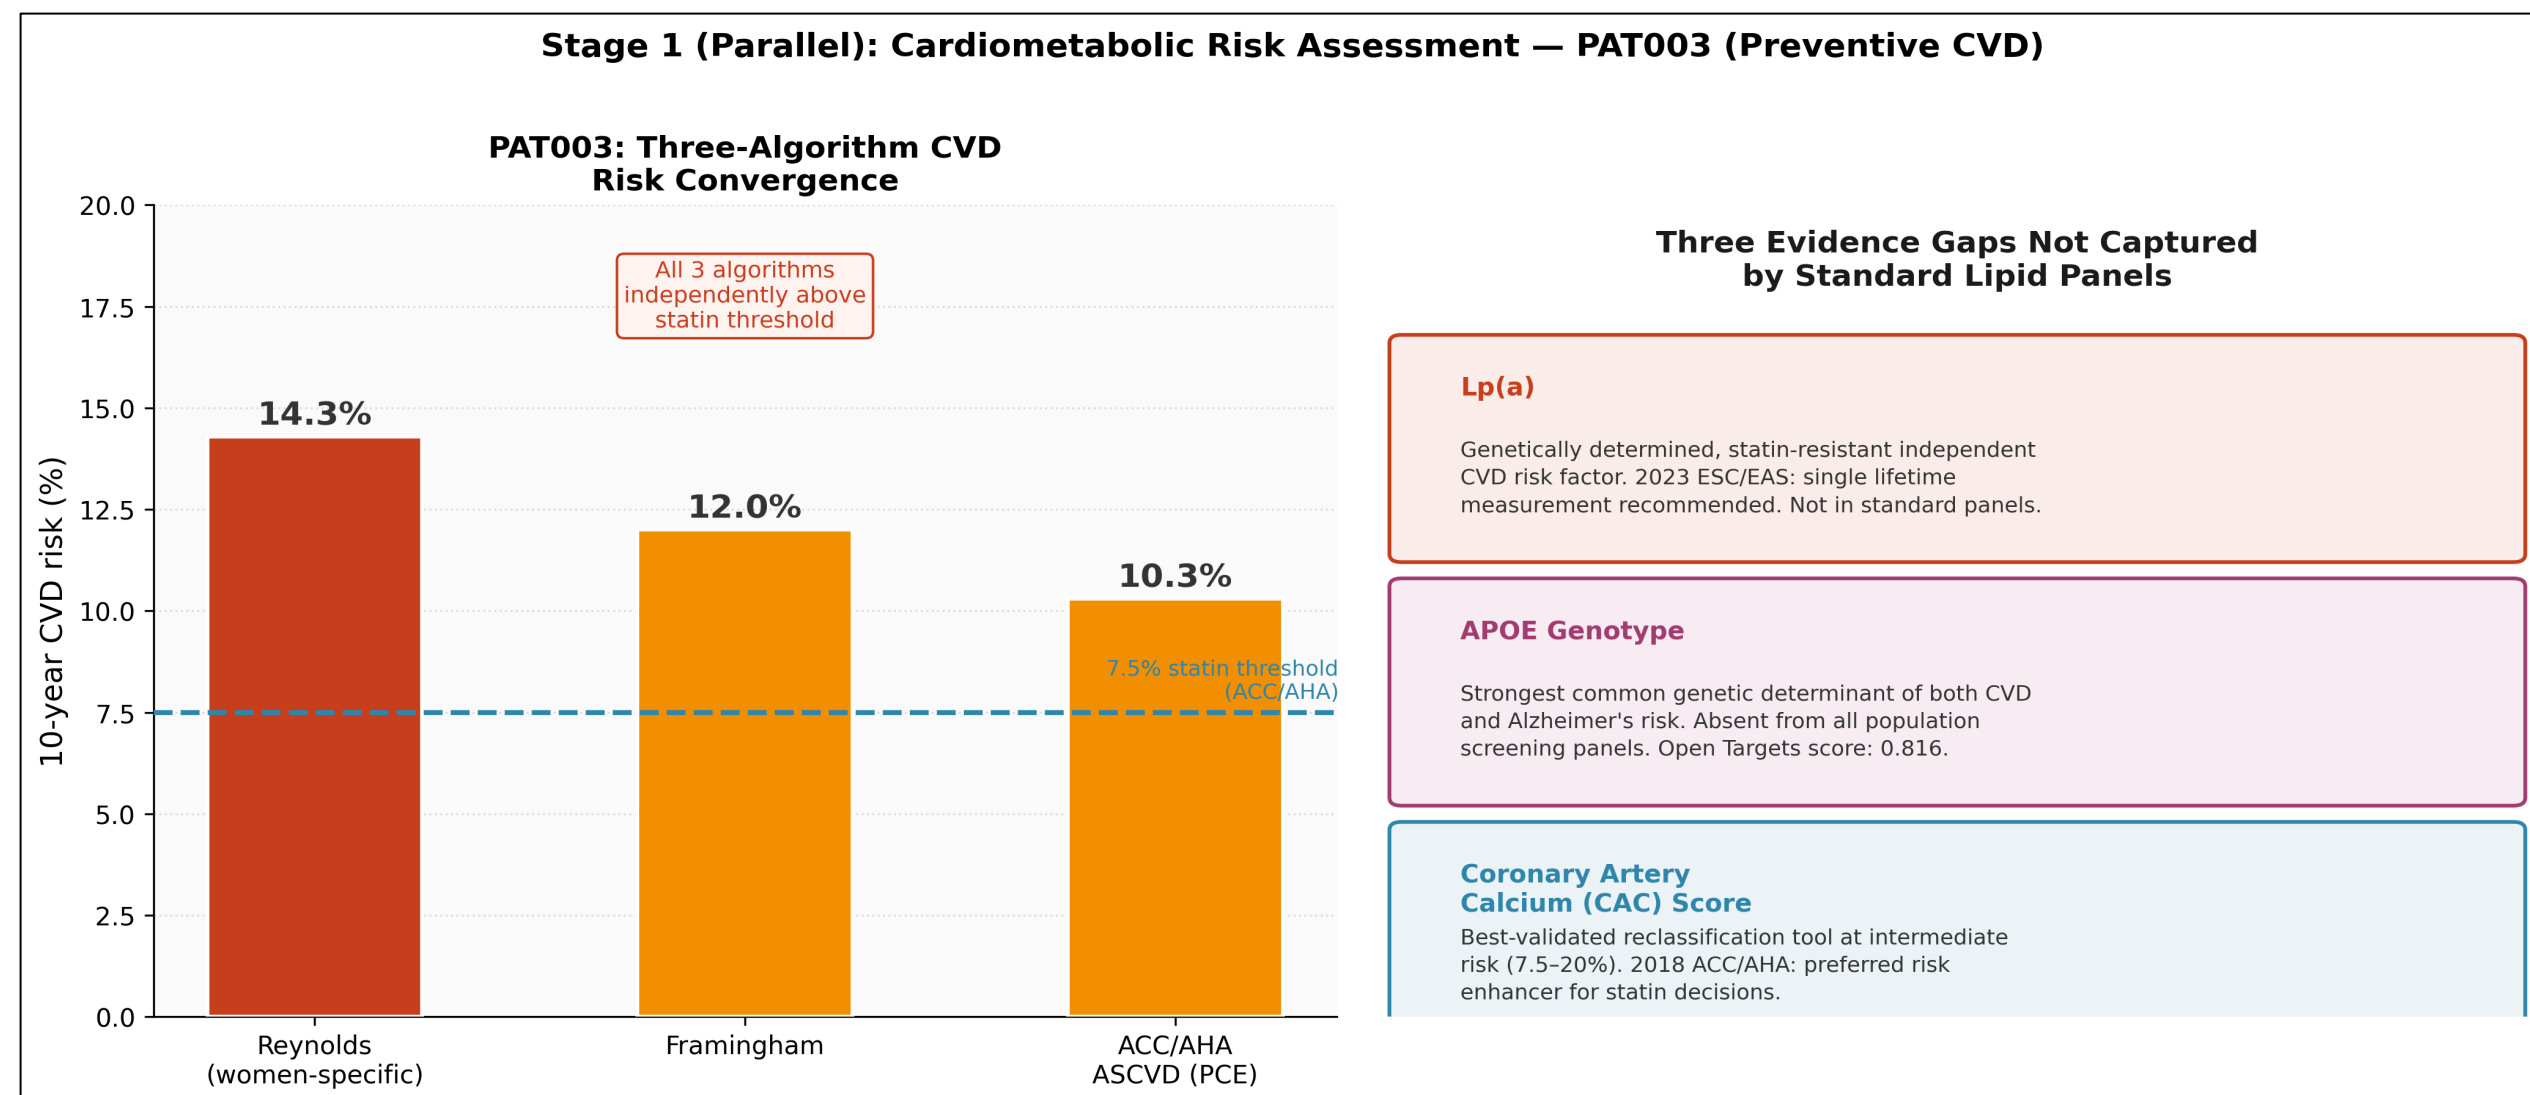


Figure 3. PAT003 (Preventive CVD): Three-algorithm CVD risk convergence above 7.5% statin threshold and three evidence gaps absent from standard lipid panels.

## DISCUSSION

- Architectural generalizability:** The platform traverses from raw sequencing data to clinical recommendations across oncology (HGSOC, ER+ BC) and preventive medicine (CVD) without pipeline redesign.
- Investigational hypothesis generalizability:** PAT002 generated three beyond-SOC hypotheses for ER+ BC via the identical Stage 3–4 architecture used for PAT001 — confirming that hypothesis-generation capacity is a property of the multi-modal integration architecture, not of the specific cancer biology.
- Confidence calibration:** Platform explicitly reports confidence gaps (e.g., YSAPLSSSL IC50 480 nM = weak binder; vaccine-context priming rationale required).

## FUTURE DIRECTIONS

- Expand MCP server federation to support additional disease areas and data modalities and analysis methodologies (complex comorbidities).
- Prospective validation in real (non-synthetic) patient cohorts with IRB-compliant de-identification workflows.
- Integrate clinical decision support interfaces for translational researchers and clinicians at the point of care.

## REFERENCES &amp; RESOURCES

- Anthropic (2024). Model Context Protocol (MCP): Open standard for AI-tool interoperability. [anthropic.com/mcp](https://anthropic.com/mcp)
- Langit L, Wang C. Precision Medicine MCP. [github.com/lynnlangit/precision-medicine-mcp](https://github.com/lynnlangit/precision-medicine-mcp) (2025).
- Singhal, K., et al., (2023). Large language models encode clinical knowledge. *Nature*, 620(7972), 172-180.
- Topol, E. J. (2023). As artificial intelligence goes multimodal, medical applications multiply. *Science*, 381, adk6139

## ACKNOWLEDGEMENTS



This work was supported by the Department of Quantitative Health Sciences, Mayo Clinic. QR code to GitHub Repo.

Contact: [langit.lynn@mayo.edu](mailto:langit.lynn@mayo.edu) | [wang.chen@mayo.edu](mailto:wang.chen@mayo.edu)