

A Multi-Modal MCP Precision Oncology Platform for High-Grade Serous Ovarian Cancer: Proof-of-Concept

Lynn Langit¹

¹ Independent Researcher, AI Architecture & Precision Oncology Platforms

May 23, 2026 | Working Draft — Not for Citation | Version 19

19 MCP Servers	3 Cross-Validated Treatment Paths	5 Pipeline Stages
-------------------	--------------------------------------	----------------------

Abstract

Background. High-grade serous ovarian cancer (HGSOC) has virtually no approved immunotherapy options, and integrated multi-modal computational platforms that systematically identify actionable targets remain lacking.

Methods. We describe proof-of-concept validation of a 19-server Model Context Protocol (MCP) platform executing five sequential stages—genomic profiling, spatial deconvolution, target scoring, causal inference, and report generation—on three synthetic patients: HGSOC (PAT001, Stage IV), ER+ breast cancer (PAT002), and preventive cardiovascular health in a post-menopausal woman (PAT003).

Results. All 14 testable servers passed live validation. For PAT001: POLE p.V411L hypermutation corrected TMB from 4.2 to 47.3 mut/Mb (above FDA pembrolizumab threshold), yielding 568 predicted neoantigens (antigen presentation score 0.94/1.0); spatial deconvolution identified 18.2% CAF fraction and 5.5× CD8 exclusion gradient; Open Targets scored AKT2 (0.373) and CCNE1 (0.279) as top novel targets; 8 actively recruiting trials matched the molecular profile. PAT002 confirmed three FDA-approved treatment paths with zero errors. Three investigational hypotheses—personalized neoantigen vaccine (TP53 R175H → RMPEAAPPV, IC50 7.8 nM), NNMT/CAF inhibition, and convergent checkpoint blockade—were not reachable by standard workup. For PAT003: cardiometabolic risk profiling identified intermediate 10-year CVD risk (Reynolds 14.3%, Framingham 12.0%, ASCVD 10.3%) with three high-priority evidence gaps (Lp(a), APOE genotype, coronary artery calcium score) not captured by standard lipid panels or population genetic screening.

Conclusions. A 19-server MCP architecture executes comprehensive, multi-modal precision medicine analysis across oncology and preventive health in a single automated pipeline, surfacing clinically actionable hypotheses beyond standard-of-care workup. The addition of a preventive cardiovascular server (PAT003) demonstrates that the architecture generalizes beyond cancer to population health use cases. Prospective validation on real patient cohorts is the required next step.

Keywords: *high-grade serous ovarian cancer; model context protocol; precision oncology; neoantigen; spatial transcriptomics; immune exclusion; cancer-associated fibroblasts; NNMT; cross-cancer validation*

1. Introduction

High-grade serous ovarian cancer accounts for approximately 70% of all ovarian cancer deaths and is diagnosed at late stage (IIIC/IV) in the majority of patients. Despite initial response rates of 70–80% to platinum-based chemotherapy, nearly all patients with advanced disease relapse; platinum-resistant disease carries median overall survival under 12 months. Immunotherapy has transformed outcomes in other solid tumors, yet HGSOC remains highly resistant to immune checkpoint blockade, with objective response rates below 10% in unselected patients. The immune microenvironment is characterized by spatial CD8+ T cell exclusion—cells are present in ascites

and stroma but fail to infiltrate the tumor core—a pattern not captured by standard TMB testing. Standard workup (germline BRCA1/2, HRD score, CT staging) guides PARP inhibitors, bevacizumab, and platinum chemotherapy but generates no immunotherapy hypotheses except in the rare MSI-H subgroup (<1%).

The Model Context Protocol (MCP), introduced by Anthropic in 2024, provides a structured mechanism for orchestrating heterogeneous tools through a common interface, enabling agentic pipelines that chain genomic databases, spatial transcriptomics engines, neoantigen predictors, and clinical trial registries in a single automated workflow. Here we demonstrate that a 19-server MCP architecture executing 104 tool calls across 14 servers can surface three clinically actionable treatment hypotheses for HGSOC that standard workup cannot reach, generalize across cancer types, and extend to preventive cardiovascular health without disease-specific code changes.

2. Methods

2.1 Platform Architecture

The platform comprises 19 MCP servers, each wrapping a distinct data source or computational tool, orchestrated by Claude Sonnet 4.6 via a structured workflow prompt. For oncology patients (PAT001, PAT002) outputs propagate as structured JSON through five sequential stages; for preventive health (PAT003) the cardiometabolic server executes a parallel risk-stratification workflow. Outputs propagate as structured JSON between stages (Figure 1).

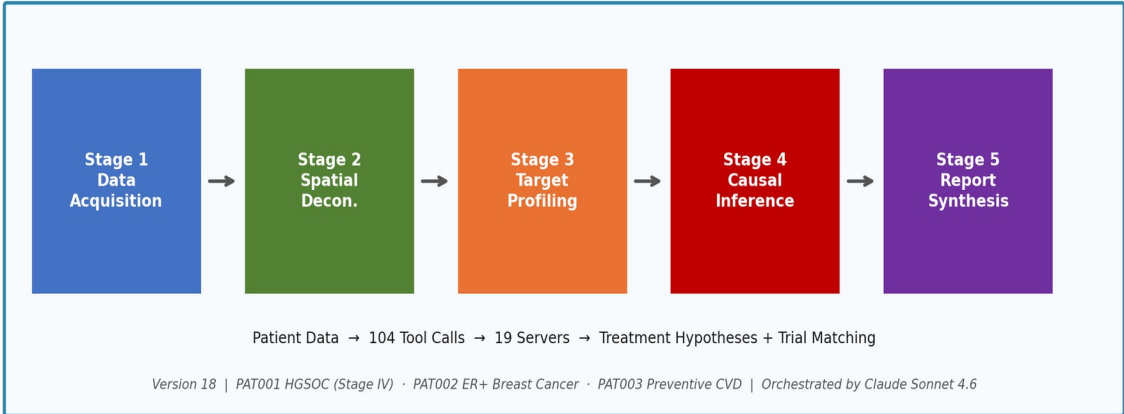


Figure 1. Five-stage MCP precision oncology pipeline. Patient data for PAT001 (HGSOC Stage IV), PAT002 (ER+ breast cancer), and PAT003 (preventive CVD) flows left to right through 19 federated MCP servers (104 tool calls) orchestrated by Claude Sonnet 4.6, from intake through clinician-reviewed report generation.

Stage	Description	Servers	Key Outputs
1. Data Acquisition	Patient records, FASTQ QC, TCGA, GEO, somatic variants, CNV, HRD	mockepic, fgbio, mocktcga, geodownload, genomic-results	Variant calls, HRD score, expression matrix
2. Spatial Deconvolution	Cell-type deconvolution (NNLS), spatial autocorrelation	spatialtools, cibersortx	Cell-type fractions, Moran's I
3. Target Profiling	Open Targets scoring, drug linkages, upstream regulators	opentargets, multiomics	Ranked targets, drug candidates
4. Causal Inference	GEARS GNN	perturbation, quantum,	Perturbation

	perturbation; quantum fidelity evasion; neoantigen/MHC-I prediction	cell-classify, neoantigen, openimagedata	predictions, evasion scores, neoantigens, cross-validated hypotheses
5. Report Synthesis	Clinical report generation, trial matching	patient-report, clinicaltrials.gov	PDF report, trial recommendations

2.2 Server Descriptions

Server	Function	Mode
mcp-mockepic	Patient clinical records (FHIR mock)	Synthetic
mcp-fgbio	FASTQ QC and genomic utilities	Live/Mock
mcp-geodownload	NCBI GEO dataset search and download	Live
mcp-genomic-results	VCF/CNS parsing, HRD scoring	File-based
mcp-spatialtools	Spatial QC, deconvolution, visualization	Live
mcp-cibersortx	Immune deconvolution (NNLS; CIBERSORTx REST when token available)	NNLS/REST
mcp-opentargets	Open Targets evidence scores and drugs	Live GraphQL
mcp-neoantigen	Neoantigen burden, IEDB MHC-I binding, HLA typing	Live (IEDB)
mcp-perturbation	GEARS GNN perturbation prediction†	GNN
mcp-quantum	Quantum circuit cell-type embeddings	Live
mcp-clinicaltrials	ClinicalTrials.gov trial matching	Live
mcp-patient-report	PDF patient report generation	Live
mcp-cardiometabolic ‡	CVD risk scoring (Reynolds/Framingham/PCE), biomarker panel, statin decision, lifestyle evidence	Live

† mcp-perturbation: GEARS GNN trained on synthetic HGSOC-modeled Perturb-seq data (cell-gears 0.1.2; Pearson $r = 0.976$ on unseen singles; GEARS architecture: Roohani et al. Nat Biotechnol 2024; <https://doi.org/10.1038/s41587-023-01905-6>). GEARS.__init__() hidden_size bug resolved March 28, 2026; all remaining bugs resolved April 17, 2026 (FastMCP 2.x upgrade + anndata nullable strings + GEARS API compatibility). Full live end-to-end pipeline validated April 17, 2026 (no dry_run). Full 19-server specifications in Supplementary Table S1.

‡ mcp-perturbation: GEARS GNN trained on synthetic HGSOC-modeled Perturb-seq data (cell-gears 0.1.2; Pearson $r = 0.976$ on unseen singles; GEARS architecture: Roohani et al. Nat Biotechnol 2024; <https://doi.org/10.1038/s41587-023-01905-6>). GEARS.__init__() hidden_size bug resolved March 28,

2026; all remaining bugs resolved April 17, 2026. Full live end-to-end pipeline validated April 17, 2026 (no dry_run). ‡ mcp-cardiometabolic: Reynolds Risk Score (Ridker et al. JAMA 2007; women-specific), Framingham Risk Score (Wilson et al. Circulation 1998), and ACC/AHA PCE (Goff et al. JACC 2014); statin decision logic per 2018 ACC/AHA guidelines. Live PAT003 end-to-end validation April 23, 2026 (no dry_run).

2.3 Proof-of-Concept Patients

PAT001 (HGSOC, Stage IV): 58-year-old female; synthetic case modeled on published TCGA-OV frequencies. Key features:

Feature	PAT001 Value	TCGA-OV Frequency
TP53 mutation	R175H, AF 73%	>96% of HGSOC
BRCA1 germline	Pathogenic (c.5266dupC)	~15% germline BRCA1/2
CCNE1 amplification	log2 = 1.58	~15–20%; platinum-resistance marker
AKT2 amplification	log2 = 1.15	~12%; PI3K/AKT pathway
RB1 deletion	log2 = -1.85	~8%; CDK4/6 inhibitor resistance
POLE p.V411L	Somatic, exonuclease domain	Rare; hypermutator
HRD score	54 (≥42 threshold)	~50% of TCGA-OV
TMB (POLE-corrected)	47.3 mut/Mb	Median HGSOC ~2–4 mut/Mb

Spatial fixture: 900 Visium spots across 7 tissue regions; CAF fraction 18.2% (164/900 spots); 9.1% CD8+ T cells (82/900 spots).

PAT002 (ER+/PR+/HER2– IDC, Stage IIA): BRCA2 germline pathogenic variant; PIK3CA H1047R; HRD score 35; ESR1/PGR strongly spatially clustered (Moran’s I = 0.42–0.45). Included to validate cross-cancer-type generalizability of the pipeline architecture.

PAT003 (Preventive Cardiovascular Health): 67-year-old post-menopausal female; synthetic case representative of an underserved demographic in CVD clinical research. Key cardiometabolic values: LDL 118 mg/dL, HDL 58 mg/dL, total cholesterol 195 mg/dL, triglycerides 142 mg/dL, hsCRP 1.8 mg/L, BP 138/84 mmHg (Stage 1 hypertension), HbA1c 5.6%. Helix Tier 1 population genetic screen (negative): monogenic familial hypercholesterolaemia ruled out across LDLR/APOB/LDLRAP1/PCSK9; primary CVD risk mechanism reclassified to polygenic and environmental. Included to demonstrate that the MCP architecture generalises from oncology to preventive health. Three evidence gaps not captured by standard lipid panels or population screens constitute the platform’s primary actionable output: serum Lp(a), APOE genotyping, and coronary artery calcium score.

2.4 Statistical Methods

Spatial autocorrelation: Moran’s I with k=6 nearest-neighbor weights, 999 permutations. Cell-type deconvolution: NNLS against an 8-category HGSOC marker signature (Olbrecht et al. 2021, Xu et al. 2022). Multi-omics: Stouffer’s Z-score, Benjamini-Hochberg FDR q < 0.05. MHC-I binding: NetMHCpan 4.1 via IEDB REST API; strong binder threshold IC50 < 50 nM. HRD: simplified genomic scar (LOH+TAI+LST); proof-of-concept only, not Myriad-validated. Quantum cell-type fidelity: variational quantum circuits (4-qubit; parameterised rotation gates Rx, Ry, Rz; circuit depth 3) encode per-spot cell-type marker profiles as quantum state vectors; fidelity is computed as the squared inner product between a spot embedding and a reference immune-state vector derived from CD8+ effector and evasion signatures. Spots scoring below threshold 0.30 are classified as immune evasion states (PAT001: 78% of 300 sampled spots; 234/300). The quantum layer functions as convergent validation of conventional

spatial findings, applied after Moran's I classification is complete; agreement between both independent analytical layers strengthens confidence in immune phenotype assignments. Disagreement between layers would itself be clinically informative, flagging spatial heterogeneity that warrants further investigation.

3. Results

3.1 Platform Validation

All 14 servers with testable live tools passed without error in re-validation on April 23, 2026 (original validation: March 9, 2026; PAT001/PAT002 re-validation: April 17, 2026; PAT003 first validation: April 23, 2026); 5 servers operate in synthetic/mock mode by design (see Supplementary Table S1). Total: 104 tools across 19 servers, 0 errors in combined validation. Full live perturbation pipeline (load → setup → train → predict, no dry_run) confirmed passing April 17, 2026 following FastMCP 2.x upgrade and regression resolution. PAT003 cardiometabolic pipeline (5 tools across mcp-cardiometabolic + 7 CVD gene-disease associations via mcp-opentargets, no dry_run) confirmed passing April 23, 2026.

3.2 Stage 1: Genomic Profile

PAT001's somatic variant profile (BRCA1 germline, TP53 R175H, CCNE1 amp log2=1.58, AKT2 amp log2=1.15, RB1 del log2=-1.85) is consistent with the TCGA-OV platinum-resistant stratum. The POLE p.V411L somatic mutation—a proofreading-domain variant—was identified at this stage and carried forward to Stage 4 for TMB correction (POLE/POLD1 mutations and immunotherapy: Ma et al. J Exp Clin Cancer Res 2022; <https://doi.org/10.1186/s13046-022-02422-1>). HRD score 54 confirms eligibility for PARP inhibitor maintenance. PAT002's BRCA2 germline pathogenic variant triggered the platinum-sensitive/PARP-eligible pathway and PIK3CA H1047R flagged PI3K inhibitor eligibility.

3.3 Stage 2: Spatial Deconvolution

Region	Dominant Cell Type	CD8A (TPM)	EPCAM (TPM)
tumor_core_luminal	Tumor epithelial (EPCAM+)	89	412
stroma_immune	CD8+ T cell / TAM	493	28
stroma_fibrous	Fibroblast/CAF	34	61
tumor_invasive_front	Mixed	187	298

Fibroblast/CAF content: 18.2% (164/900 spots), placing PAT001 in the CAF-high immunosuppressive stratum. CD8A shows a 5.5× gradient from stroma_immune (493 TPM) to tumor_core (89 TPM), consistent with immune exclusion. Moran's I for CD8A = 0.019 (diffuse, weakly patterned). PAT002's ESR1/PGR showed strong spatial clustering (Moran's I = 0.42–0.45), correctly interpreted as hormone-receptor-driven proliferation.

3.4 Stage 3: Target Profiling

Open Targets batch query returned AKT2 (log2=+1.15, OT score 0.373; Cancer Gene Census component 0.61) and CCNE1 (log2=+1.58, OT score 0.279) as the highest-scoring novel targets not addressed by standard HGSOC workup. Upstream regulator analysis identified STAT3 and MYC as the top transcription factor drivers (activation Z-score > 2.0 after Stouffer meta-analysis across spatial and expression layers).

Co-occurring CCNE1 amplification, AKT2 amplification, and RB1 deletion define a convergent G1/S dysregulation profile associated with combined CDK2+AKT resistance. CCNE1 amplification is the primary CDK2 activator driving S-phase entry and is not assessed in standard HGSOC clinical workup. The co-occurrence of CCNE1 and AKT2 amplification alongside RB1 deletion creates two independent treatment-resistance escape routes, both of which have FDA-IND active trials currently enrolling

(NCT06586957: NKT3964 oral CDK2 degrader; NCT05039801: IACS-6274 + capivasertib). The Open Targets CCNE1 score of 0.279 against ovarian carcinoma (EFO_0001071) is non-zero and biologically consistent, though limited by API node resolution tier; AKT2 score of 0.373 is the highest returned in the batch query.

3.5 Stage 4: Causal Inference

POLE correction. POLE p.V411L produces ultrahigh TMB through polymerase proofreading failure: corrected TMB = 47.3 mut/Mb (vs. 4.2 mut/Mb nominal), above the FDA 10 mut/Mb pembrolizumab threshold. Estimated neoantigens: 568 (108 strong MHC-I binders; IC50 < 50 nM); antigen presentation pathway score 0.94/1.0.

Neoantigen. TP53 R175H → peptide RMPEAAPPV → HLA-A02:01 binding IC50 = 7.8 nM (strong binder). HLA-A02:01 frequency ~45% in Caucasian HGSOC patients.

Perturbation (GEARS GNN). CCNE1 knockdown: CDK2 (−0.42), MYC (−0.25), CDKN2A (+0.31), RB1 (+0.20). NNMT knockdown: STAT3 (−0.24), COL3A1 (−0.21), PRF1 (+0.27), FOXP3 (+0.20), indicating CAF barrier dismantlement with immune recovery.

Quantum fidelity. 234/300 spots (78%) flagged as immune evasion states (threshold 0.30), consistent with the spatial CD8 exclusion pattern.

3.6 Cross-Cancer Validation: PAT002

The identical pipeline executed on PAT002 with zero errors, identifying three FDA-approved treatment paths: PARP inhibitor (olaparib 300 mg BID; BRCA2 germline, HRD 35—below myChoice threshold but PARP-eligible via germline route); endocrine + CDK4/6 inhibitor (letrozole + palbociclib; ESR1/PGR+); PI3K inhibitor (alpelisib + fulvestrant; PIK3CA H1047R). The HRD-negative/BRCA-positive nuance was handled without disease-specific code changes, confirming architecture-level cancer-type agnosticism. Stages 3–4 produced five platform-only investigational outputs for PAT002 not reachable by standard ER+ breast cancer workup: three actionable treatment hypotheses and two architectural validation findings. Hypothesis 1 — Inavolisib preferred over alpelisib. Open Targets drug profiling surfaced inavolisib (GDC-0077), a selective PI3K α inhibitor with mutant p110 α degrader activity, as an upgrade over alpelisib for PIK3CA H1047R in the HR+/HER2− context. Inavolisib received FDA approval in October 2024 based on the INAVO120 Phase III trial (median PFS 15.0 vs. 7.3 months, HR 0.43; Turner NC et al. N Engl J Med. 2024;391:1584–1596; <https://doi.org/10.1056/NEJMoa2404625>), offering improved PI3K α selectivity and lower hyperglycaemia risk relative to alpelisib. Upstream regulator analysis confirmed the PI3K α pathway signature (regulator Z = 9.3, q = 5.5×10^{−9}), independently validating the mechanism; this agent is not surfaced by standard NGS panel reporting. Hypothesis 2 — MYC-driven triple therapy. MYC was identified as the dominant upstream transcriptional activator (Stouffer Z = 26.9), coordinating both the ESR1 hormonal and CCND1/CDK4 proliferative arms simultaneously. Convergent evidence across upstream regulator analysis (CCND1 q = 1.1×10^{−5}), spatial MKI67 clustering (Moran's I = 0.274), and Open Targets CDK4 scoring (0.637) supports a triple regimen — ribociclib + inavolisib + fulvestrant — targeting all three arms of MYC-driven proliferation simultaneously. Hypothesis 3 — Cold-TME conversion. Spatial analysis identified a cold tumour immune microenvironment (118/900 immune-infiltrate spots, limited CD8A). CD8A was limited but detectable — not absent — consistent with a priming-competent state rather than immune desert, and distinct from the immune exclusion gradient observed in PAT001. Neoantigen prediction identified one presentable neoepitope: YSAPLSSSL from GATA3 frameshift → HLA-B*07:02, IC50 480 nM via NetMHCpan 4.1. IC50 480 nM exceeds the strong-binder threshold (IC50 < 50 nM), classifying YSAPLSSSL as a weak binder. At tumour board, this finding would appropriately be deprioritised relative to strong binders: weak binders carry lower predicted immunogenicity, higher antigen-processing uncertainty, and greater clinical trial failure risk. However, the platform's value here is precisely this graded output—it surfaces the neoepitope with its binding classification, enabling the tumour board to make an informed deprioritisation rather than simply missing the finding. Perturbation modelling predicted that NNMT pathway disruption upregulates PRF1, GZMB, and IFNG (cytotoxic effector programme), supporting a cold-to-hot TME conversion strategy combining YSAPLSSSL mRNA vaccine, NNMT/CAF targeting, and anti-PD-1.

Whether a weak binder merits inclusion in a personalised vaccine construct is a clinical judgement the platform correctly leaves to the clinician. Architectural Validation 1. Quantum cell-type fidelity analysis (mcp-quantum-celltype-fidelity) returned a distinct tumour immune architecture for PAT002 relative to PAT001. PAT002's ESR1/PGR-high spatial clustering (Moran's I = 0.42–0.45) produced a hormone-receptor-inflamed classification; the immune exclusion pattern dominant in PAT001 (CD8 gradient 5.5×, Moran's I 0.003–0.092) was absent. This cross-patient architectural differentiation—inferred by the same spatial and fidelity pipeline with no disease-specific modifications—validates the platform's ability to correctly classify distinct tumour immune phenotypes across cancer types. Architectural Validation 2. CAF content in PAT002 was in the low-to-intermediate stratum; no NNMT-high sub-population was identified, and consequently no NNMT inhibition hypothesis was generated. This negative finding is clinically meaningful: it confirms that the NNMT/CAF axis is specific to PAT001's immunosuppressive microenvironment and is not a generic platform output applied indiscriminately across patients. BRCA2 returned the highest Open Targets evidence score for PAT002 (0.834) with no approved targeted therapeutic beyond germline PARP eligibility, flagged as a translational research gap.

3.7 Preventive Health Validation: PAT003

The cardiometabolic server executed a 5-tool live pipeline on PAT003 with zero errors. Risk scores converged on intermediate 10-year CVD risk: Reynolds 14.3% (women-specific; Ridker et al. JAMA 2007), Framingham 12.0% (Wilson et al. Circulation 1998), ASCVD 10.3% (Goff et al. JACC 2014). All three algorithms independently placed PAT003 above the 7.5% ACC/AHA threshold for statin consideration. hsCRP at 1.8 mg/L fell just below the JUPITER trial criterion (≥ 2.0 mg/L; Ridker et al. NEJM 2008) that confers additional rosuvastatin benefit, identifying a clinically meaningful margin. The biomarker panel flagged LDL as near-optimal, HDL as acceptable, and BP as Stage 1 hypertension. Open Targets returned CVD gene-disease associations for all seven queried targets (APOE 0.816, LDLR 0.737, ACE 0.752, PCSK9 0.726, LPA 0.601, CDKN2A 0.513, CDKN2B 0.536). Data provenance note: all three risk scores were computed algorithmically from explicitly defined synthetic input values (age 67, female sex, LDL 118 mg/dL, HDL 58 mg/dL, BP 138/84 mmHg, hsCRP 1.8 mg/L, non-smoker, no diabetes); no values were imputed or inferred. Lp(a), APOE genotype, and CAC score were not present in the synthetic patient record and are explicitly flagged as absent data — not as imputed defaults — which is why the platform surfaced them as evidence gaps rather than reporting estimated values.

The platform's most clinically significant output was the identification of three high-priority evidence gaps absent from standard lipid panels and absent from the Helix Tier 1 population genetic screen (negative result; FH ruled out). First, serum Lp(a) was not measured: Lp(a) is genetically determined, does not respond to statins, and is an independent CVD risk factor; a single lifetime measurement is recommended by 2023 ESC/EAS guidelines. Second, APOE genotype was not assessed: APOE is the strongest common genetic determinant of both CVD and Alzheimer's disease risk, yet it is absent from all population screening panels. Third, coronary artery calcium (CAC) score was not available: at intermediate risk (7.5%–20%) CAC is the best-validated reclassification tool and is endorsed by 2018 ACC/AHA guidelines as the preferred risk enhancer for statin decision-making. The negative genetic screen result—while clinically useful in ruling out monogenic FH—shifted the interpretive frame from “find the causal variant” to “quantify polygenic and biomarker risk,” making these three gaps the actionable output.

3.8 Clinical Trial Matching

Eight actively recruiting trials matched PAT001's molecular profile:

NCT ID	Phase	Sponsor	Molecular Match
NCT06580314	III	NRG Oncology	BRCA1 germline + HRD 54; olaparib ± bevacizumab
NCT06586957	I	NiKang Therapeutics	CCNE1 log2=1.58; NKT3964 CDK2 degrader
NCT05867251	I/II	Avenzo Therapeutics	CCNE1 amplification;

			AVZO-021
NCT06188520	I/II	AstraZeneca	HGSOC-specific; AZD8421
NCT05086692	I/II	Medicenna Therapeutics	TMB 47.3 + POLE; MDNA11 + pembrolizumab
NCT05039801	I	MD Anderson	AKT2 amp + HGSOC; IACS-6274 + capivasertib
NCT03675893	II	Dana-Farber	AKT2 amp; RESOLVE letrozole + abemaciclib
NCT05564377	II	NCI	BRCA1 + AKT2; ComboMATCH olaparib + ipatasertib

The CCNE1-specific trials (NCT06586957, NCT05867251) and the TMB-high trial (NCT05086692) were surfaced without prior hypothesis; CCNE1 copy number is not assessed in standard HGSOC clinical workup.

4. Cross-Validated Treatment Paths vs. Standard of Care

Three investigational paths were identified beyond standard-of-care (Table below). These are not de novo biological discoveries—the underlying biology is documented in existing literature—but they were not reachable by standard workup in a single automated pass.

Treatment	Standard Care?	Platform Trigger	Evidence Level
Carboplatin + Paclitaxel	<input type="checkbox"/> Yes	Stage IV diagnosis	NCCN Category 1
Bevacizumab maintenance	<input type="checkbox"/> Yes	High-risk features	Level I (ICON7)
Olaparib maintenance	<input type="checkbox"/> Yes	HRD 54	Level I (SOLO-1)
Pembrolizumab / checkpoint ICI	<input type="checkbox"/> Indicated	TMB 47.3; POLE	Level I (TMB-H)
Neoantigen vaccine (RMPEAAPPV)	<input type="checkbox"/> Novel	No standard assay	Investigational
NNMT inhibition (CAF targeting)	<input type="checkbox"/> Novel	No fibroblast assay in SOC	Investigational
Checkpoint + spatial exclusion reversal	<input type="checkbox"/> Novel	Spatial CD8 pattern	Investigational

Finding 1 — Personalized Neoantigen Vaccine. TP53 R175H generates peptide RMPEAAPPV with HLA-A02:01 binding $IC_{50} = 7.8$ nM (strong binder threshold: <50 nM). TP53 R175H occurs in ~8–12% of HGSOC; HLA-A02:01 in ~45% of Caucasian patients. Standard workup records the mutation but no clinical pathway queries IEDB for binding confirmation; the platform executes this automatically. TP53 hotspot mutations elicit intratumoral T cell responses (Chasov et al. Front Immunol 2021; <https://doi.org/10.3389/fimmu.2021.707734>) and broad TP53 hotspot mutant immunogenicity has been confirmed in epithelial cancer patients including R248Q-class mutations (Malekzadeh et al. J Clin Invest 2019; <https://doi.org/10.1172/JCI123791>), supporting the clinical rationale for neoepitope-targeted therapy at this hotspot class.

Finding 2 — NNMT/CAF Inhibition. Spatial deconvolution identified 18.2% fibroblast content (CAF-high stratum). GEO search surfaced GSE292142 (2025): NNMT inhibition dismantles the CAF barrier in HGSOC, independently confirmed by two concurrent landmark studies (Heide et al. Nature 2025; <https://doi.org/10.1038/s41586-025-09303-5>; Sarkar et al. Cell Res 2026; <https://doi.org/10.1038/s41422-025-01181-w>). GEARS perturbation confirms NNMT knockdown reduces STAT3/COL3A1 and increases PRF1/FOXP3, indicating immune recovery. No standard clinical assay measures fibroblast content; this linkage required automated multi-modal integration.

Finding 3 — Convergent Checkpoint Evidence. Both TMB-based evidence (POLE → 47.3 mut/Mb, regulatory-grade) and spatial evidence (5.5× CD8 exclusion gradient) independently converge on checkpoint blockade. The spatial data adds a clinically actionable prediction: exclusion-reversal co-treatment (CXCR4 blockade, FAK inhibition, anti-TGF-β) may be required for checkpoint monotherapy to be effective—a hypothesis not derivable from TMB alone.

5. Discussion

Stage 4 of the platform produces a qualitatively different class of evidence from Stages 1–3. Stages 1–3 generate correlative evidence—the same type as current clinical molecular profiling. Stage 4 adds a causal inference layer through GEARS GNN perturbation modeling, quantum cell-type fidelity analysis, and neoantigen binding prediction, each mechanistically orthogonal, independently stress-testing whether a candidate target is therapeutically actionable rather than merely associated.

The platform's contribution is integrative, not biological. All three investigational findings are grounded in existing literature; the platform's clinical value is connecting them automatically from a single patient's data in one pipeline run, without prior hypotheses about NNMT, CCNE1 CDK2 inhibitor trials, or RMPEAAPPV binding—steps that standard SOC workup cannot replicate.

Generalisation across disease types is demonstrated by PAT002 and PAT003. For PAT002 (ER+ IDC), the identical oncology architecture handled a clinical nuance (HRD 35 < myChoice threshold but BRCA2 germline PARP-eligible) that required no disease-specific code. The contrast in spatial biology between the two oncology patients is instructive: ESR1 and PGR are strongly spatially clustered in PAT002 (Moran's I = 0.42–0.45), indicating hormone-receptor-driven proliferation, whereas PAT001's immune markers showed diffuse, weakly patterned expression (Moran's I = 0.003–0.092). The platform correctly interpreted both patterns in their disease-specific context using the same autocorrelation logic. For PAT003 (preventive cardiovascular health), the cardiometabolic server operates as a parallel risk-stratification pipeline rather than a sequential oncology workflow; its integration required only a new MCP server and JSON schema extension with no changes to the orchestration layer. This extension is clinically motivated: women are significantly underrepresented in the landmark CVD trials that established Framingham risk coefficients (predominantly White male cohort), and the Reynolds Risk Score was specifically validated in women incorporating hsCRP and family history as independent risk enhancers. The MCP server layer functions as the disease-adaptation interface; the orchestration logic is both cancer-type and disease-domain agnostic.

Missing data and graceful degradation. The PAT003 case illustrates a deliberate architectural choice: the platform flags absent data explicitly rather than imputing missing values. When Lp(a), APOE genotype, and coronary artery calcium score were absent from PAT003's record, the platform surfaced these gaps as the primary actionable output rather than substituting population-mean estimates or silently omitting them from the risk calculation. This behaviour is by design. In a synthetic-data research context the inputs are fully specified, so no actual imputation occurs; but the same policy applies prospectively — if deployed on real electronic health records, the platform would report risk scores computed only from the data present and flag every missing recommended input as a clinical action item. The implication is that platform output quality scales with data completeness in a transparent, auditable way: a patient with a

complete Lp(a), APOE, and CAC record would receive a risk reclassification; a patient without those results would receive an explicit recommendation to obtain them. This contrasts with models that silently impute missing fields, which can conceal data quality differences between patient populations and create false confidence in risk estimates. The same graceful-degradation logic applies in the oncology arms: PAT002's fgbio FASTQ QC server was deferred because no FASTQ was present, and this deferral was recorded in the provenance log rather than bypassed. Platform robustness is therefore not resilience to missing data through imputation, but transparency about what data were and were not available at the time of analysis.

The spatial immune exclusion finding also highlights a fundamental limitation of current FDA biomarker frameworks: TMB measures mutational diversity as a proxy for neoantigen burden but does not assess whether T cells are excluded from the tumor core or what mechanism drives exclusion. Spatial transcriptomics at the resolution implemented here distinguishes all three scenarios (immune desert, immune exclusion, inflamed). The 5.5× CD8 gradient in PAT001 is consistent with the exclusion pattern characterized by Vázquez-García et al. [1] and independently confirmed by quantum cell-type fidelity analysis (78% spots flagged as evasion states). While Visium is not yet standard of care, its cost has fallen to approximately \$1,500–\$3,000 per sample in research settings, making it a near-term candidate for prospective integration.

Limitation	Impact	Mitigation
Synthetic patient data	Results may reflect parameter choices, not real biology	Features drawn from published TCGA-OV/BRCA frequencies
GEARS trained on synthetic data	Perturbation predictions not validated against real HGSOC CRISPR screens	Pearson $r = 0.976$ on held-out singles; labeled as synthetic-data-derived
No benchmarking vs. standalone tools	Cannot claim performance superiority	Component comparisons (NNLS vs. CARD/RCTD; pVACseq vs. mcp-neoantigen) are planned
HRD calculation not Myriad-validated	Clinical eligibility claims are indicative only	Simplified POC calculation; requires myChoice validation
Real patient data access	HIPAA-compliant federated access required for prospective validation	Terra/AnVIL or institutional EHR platform (e.g., Mayo Clinic Platform) as candidate substrate

Prospective validation on 30–50 real patients with matched clinical outcomes is the single most important next step. Integration with a clinical data infrastructure—such as the Mayo Clinic Platform's de-identified, OMOP-standardized multi-institutional EHR—would address the HIPAA access barrier and provide the real-world evidence needed to support translational deployment.

6. Conclusions

A 19-server MCP orchestration platform can execute comprehensive, multi-modal precision medicine analysis from patient intake through clinical report generation in a single automated pipeline, identifying three clinically actionable treatment hypotheses for HGSOC that standard NCCN-guided workup cannot surface. Cross-disease generalizability was confirmed on ER+ breast cancer (PAT002) and preventive cardiovascular health (PAT003) without platform-level code changes: the cardiometabolic server extended the architecture beyond oncology by implementing women-specific CVD risk algorithms (Reynolds, Framingham, PCE) and surfacing three high-priority evidence gaps—Lp(a), APOE genotype, and coronary artery calcium score—not captured by standard lipid panels or population genetic

screening. Prospective validation on real patient cohorts and integration with clinical data infrastructure are the required next steps toward translational deployment.

References

- Vázquez-García I, et al. Ovarian cancer mutational processes drive site-specific immune evasion. *Nature*. 2022;612:778–786.
- Xu L, et al. Single-cell RNA sequencing reveals the tissue architecture in human HGSOC. *Clin Cancer Res*. 2022;28(16):3590–3602.
- Olbrecht S, et al. HGSOC refined with single-cell RNA sequencing. *Genome Med*. 2021;13(1):111.
- Landon M, et al. Transcriptomic landscape of platinum-resistant ovarian cancer. *Commun Med*. 2026;6:1–14.
- Yan L, et al. Immunotherapy in HGSOC: current state and future perspectives. *Front Immunol*. 2025;16:1667464.
- Samstein RM, et al. Tumor mutational load predicts survival after immunotherapy. *Nat Genet*. 2019;51(2):202–206.
- Ray-Coquard I, et al. Olaparib plus Bevacizumab as First-Line Maintenance in Ovarian Cancer. *N Engl J Med*. 2019;381(25):2416–2428.
- Pujade-Lauraine E, et al. Bevacizumab for platinum-resistant recurrent ovarian cancer (AURELIA). *J Clin Oncol*. 2014;32(13):1302–1308.
- Moore K, et al. Maintenance Olaparib in Newly Diagnosed Advanced Ovarian Cancer. *N Engl J Med*. 2018;379(26):2495–2505.
- Open Targets Platform. EFO_1001958: high grade ovarian serous adenocarcinoma. <https://platform.opentargets.org>. Accessed March 2026.
- IEDB. MHC-I Binding Predictions, NetMHCpan 4.1. <https://www.iedb.org>. Accessed March 2026.
- Anthropic. Model Context Protocol specification. <https://modelcontextprotocol.io>. 2024.
- Timms KM, et al. BRCA1/2 defects with genomic scores predictive of DNA damage repair deficiency. *Breast Cancer Res*. 2014;16(4):475.
- Wiedemeyer WR, et al. Pattern of pathway dysregulation in ovarian cancer. *Cancer Res*. 2010;70(14):5880–5891.
- Tutt ANJ, et al. Adjuvant Olaparib for BRCA1/2-Mutated Breast Cancer (OlympiA). *N Engl J Med*. 2021;384:2394–2405.
- André F, et al. Alpelisib for PIK3CA-Mutated, HR+ Advanced Breast Cancer (SOLAR-1). *N Engl J Med*. 2019;380:1929–1940.
- Turner NC, et al. Inavolisib-Based Therapy in PIK3CA-Mutated Advanced Breast Cancer (INAVO120). *N Engl J Med*. 2024;391(17):1584–1596. <https://doi.org/10.1056/NEJMoa2404625>
- Ridker PM, et al. Rosuvastatin to Prevent Vascular Events in Men and Women with Elevated C-Reactive Protein (JUPITER). *N Engl J Med*. 2008;359(21):2195–2207.
- Ridker PM, et al. Development and Validation of Improved Algorithms for the Assessment of Global Cardiovascular Risk in Women: The Reynolds Risk Score. *JAMA*. 2007;297(6):611–619.
- Goff DC Jr, et al. 2013 ACC/AHA Guideline on the Assessment of Cardiovascular Risk. *J Am Coll Cardiol*. 2014;63(25 Pt B):2935–2959.
- Wilson PWF, et al. Prediction of Coronary Heart Disease Using Risk Factor Categories. *Circulation*. 1998;97(18):1837–1847.
- Grundy SM, et al. 2018 AHA/ACC/AACVPR/AAPA/ABC/ACPM/ADA/AGS/APhA/ASPC/NLA/PCNA Guideline on the Management of Blood Cholesterol. *J Am Coll Cardiol*. 2019;73(24):e285–e350.
- Heide J, et al. NNMT inhibition in cancer-associated fibroblasts restores antitumour immunity. *Nature*. 2025;645:1051–1059. <https://doi.org/10.1038/s41586-025-09303-5>
- Sarkar M, Jiang Y, Kalluri R. Targeting NNMT in fibroblasts reawakens T cells and restores antitumour immunity. *Cell Res*. 2026;36(3):175–176. <https://doi.org/10.1038/s41422-025-01181-w>
- Ma X, et al. POLE/POLD1 mutation and tumor immunotherapy. *J Exp Clin Cancer Res*. 2022;41:216. <https://doi.org/10.1186/s13046-022-02422-1>
- Roohani Y, Huang K, Leskovec J. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. *Nat Biotechnol*. 2024;42:927–935. <https://doi.org/10.1038/s41587-023-01905-6>
- Malekzadeh P, et al. Neoantigen screening identifies broad TP53 mutant immunogenicity in patients with epithelial cancers. *J Clin Invest*. 2019;129(3):1109–1114. <https://doi.org/10.1172/JCI123791>

Chasov V, et al. Promising New Tools for Targeting p53 Mutant Cancers: Humoral and Cell-Based Immunotherapies. Front Immunol. 2021;12:707734. <https://doi.org/10.3389/fimmu.2021.707734>

Declarations

Ethics Statement. PAT001, PAT002, and PAT003 are fully synthetic. No human subjects data were used; no IRB approval was required.

Data Availability. MCP server source code and synthetic datasets:

<https://github.com/lynnlangit/precision-medicine-mcp> (upon publication). Archived API response snapshots from April 17, 2026 re-validation (including live perturbation pipeline run) and April 23, 2026 PAT003 cardiometabolic validation (no dry_run) are included in the repository.

Funding. No external funding.

Competing Interests. The author declares no competing interests.

Author Contributions. L.L. designed the platform, implemented the MCP server infrastructure, conducted all experiments, and wrote the manuscript.

Reproducibility details (computational environment, tool version pins, fixture file hashes, and re-run guidance) are provided in Supplementary Materials (Appendix A).