Agent-Atom-Gradient: A Next-Gen AI Architecture for Evidence-Driven Outcomes Research

Acceptance code

A Livieratos[1], M Kudela[1,2], Y Zhao[1,2], A Chen[1,2], J Lin[1,3], D Zhang[1,4], X Luo[1,2], PA Ramos[1,2], S Dharmarajan[1,5], C Su[2], M Gaga[6], M Gamalo[1,2]

[1]SPAIML Scientific Working Group, USA; [2]Pfizer, USA; [3]Takeda Pharmaceuticals, USA; [4]Teva Pharmaceuticals, USA; [5]Sarepta Therapeutics, USA; [6]Hygeia Hospital, Greece

INTRODUCTION

Automated treatment optimization has emerged as a transformative approach in personalized medicine, leveraging computational methods to tailor therapy regimens to individual patient characteristics.[1] Early work demonstrated how pharmacokinetic/pharmacodynamic models, combined with Bayesian feedback loops, can optimize drug dosing for narrow-therapeutic-index agents, significantly improving efficacy and safety in routine clinical practice. Beyond dosing, reinforcement learning has been explored to derive dynamic treatment policies—for example, in sepsis management and mechanical ventilation—by framing clinical decision-making as a sequential decision problem, optimizing long-term patient outcomes in observational health data.[2] More recently, causal machine learning methods have been applied within precision medicine frameworks to estimate individual treatment effects and guide therapy selection, highlighting the potential for data-driven strategies to improve decision support in complex care pathways.[3]

Despite these advances, most existing automation frameworks focus on single-disease paradigms or finite decision points but are not readily adaptable to complex, multi-stage diseases. Immune-mediated conditions like ulcerative colitis involve heterogeneous patient trajectories, multiple lines of therapy, biomarker-defined subpopulations, and evolving real-world outcomes. Traditional rule-based or single-agent algorithms struggle to capture the interplay between induction, maintenance, and rescue therapy decisions, leading to suboptimal personalization and limited generalizability. There is a critical need for architectures that can decompose multifaceted treatment pathways into atomic decision units, aggregate insights across diverse model outputs, and refine recommendations through iterative feedback loops.

Concurrently, real-world evidence (RWE) has gained prominence as a complement to randomized controlled trials. Under the 21st Century Cures Act, the U.S. FDA established its RWE Program to evaluate how electronic health records, claims data, and registries can support regulatory decision-making for drugs, biologics, and medical devices.[4] Integrating RWE into automated treatment optimization pipelines is therefore essential for ensuring that AI-driven recommendations are grounded in both clinical trial data and real-world practice patterns.

OBJECTIVE

To design novel AI architecture that could seamlessly ingest and leverage RWE—including electronic health records, claims databases, and patient registries from various publications—within a multi-stage reasoning pipeline. This novel framework will:

1.Aggregate insights from heterogeneous LLM agents (Mixture-of-Agents, MoA) alongside data from publications.[5]
2.Decompose complex treatment pathways into atomic decision units via Atom-of-Thought (AoT) reasoning.[6]
3.Refine and personalize final recommendations through test-time textual differentiation (TextGrad).[7]

By integrating RWE alongside randomized trial results, the system will generate multi-faceted insights—covering guideline concordance, biomarker stratification, safety signals, and real-world utilization patterns—for treatment sequence optimization in ulcerative colitis (UC) and chronic pulmonary obstructive disease (COPD).

METHODS

Our treatment-sequence optimization pipeline (Figure 1) begins by harnessing a **MoA** architecture, based on each model's internal training data, to generate a broad set of candidate regimens. The clinical question—"What is the optimal, based on patient outcomes, first-, second-, and third-line sequence for ulcerative colitis, combining RCT, guidelines and real-world evidence?"—is issued in parallel to three proposer models (Meta-Llama-3.1-8B-Instruct-Turbo; Mistral-7B-Instruct-v0.3; Qwen2-72B-Instruct). Each model produces an independent rationale, which is then fed back into the same three proposers for two further refinement layers. Finally, all layer-3 outputs are synthesized by the DeepSeek-R1 aggregator under a system prompt that enforces critical, evidence-based synthesis. This multi-layered MoA stage yields an integrated MoA output that combines RCTs and guideline recommendations with emerging RWE insights.

Next, we apply **AoT** reasoning to the MoA output, using GPT-4.1 via LangChain. In this phase, the composite recommendation is decomposed into a small directed acyclic graph of atomic (not further broken apart) sub-questions—such as "After TNF inhibitor failure, which treatment shows the best durability in RWE cohorts?"—and each node is independently solved. The answered sub-questions are then automatically contracted back into a simplified master question, with consolidated step-by-step logic. Two recursive AoT iterations focus the reasoning on key inflection points in the treatment pathway, sharpening clarity and clinical relevance.

Finally, we refine and personalize the AoT-vetted sequence via **TextGrad**, our textual backpropagation framework. We treat the blended MoA + AoT recommendation outputs as a trainable variable and fix the original clinical question. TextGrad's optimizer (TGD) applies three gradient-style update steps, guided by a DeepSeek backward engine. At each step, a TextLoss prompt evaluates alignment with guidelines, RCT evidence, and RWE patterns, while hard constraints enforce the required structure (First-line, Second-line, Third-line) and alignment of specific, non-predefined, patient subgroups to tailored sequences (can flexibly select the alignment inputs per user requirements). This iterative optimization sharpens both the accuracy and the explanatory depth of the final recommendation, yielding a clinically coherent, expert-grade treatment plan for ulcerative colitis. This pipeline was applied across diverse therapeutic areas, including UC and COPD, to demonstrate the flexibility and accuracy across different treatment sequence optimization problems.

Figure 1: Multi-Stage LLM Pipeline for Treatment Sequence Optimization: MoA → AoT → TextGrad

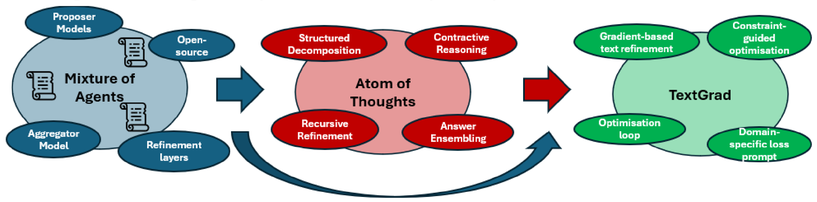

RESULTS

By benchmarking all three outputs side-by-side, we saw a clear evolution in both transparency and clinical utility. The UC MoA table (Table 1) crisply illustrated how three independent LLM proposer models could be combined via a DeepSeek-R1 aggregator to yield a unified first-, second-, and third-line treatment sequence—yet it remained largely conceptual, without trial data or subgroup stratification. The UC AoT (Table 2) then made the underlying reasoning explicit, decomposing the recommendation into six targeted clinical questions and re-contracting them into a coherent pathway. Finally, the UC TextGrad-refined table (Table 3) delivered the richest, most actionable output: each therapy line was paired with precise trial outcomes, real-world metrics, and explicit patient-subgroup guidance. This progressive refinement—from high-level model synthesis to fully evidence-anchored clinical decision support—highlights how successive reasoning layers can transform a conceptual LLM output into a dense, data-driven summary.

Table 1: MoA Example Output for UC

Stage	Output (Summary)
Proposer Layer	**Models:** Meta-Llama-3.1-8B, Mistral-7B, Qwen2-72B • **1st-line:** 5-ASA (mesalamine) + short-course steroids • **2nd-line:** TNF inhibitors (e.g. infliximab / adalimumab) or vedolizumab or JAK inhibitors (e.g. tofacitinib / upadacitinib) • **3rd-line:** IL-23 inhibitors (e.g. ustekinumab / risankizumab) or novel agent (S1P) • **Notes:** safety-driven stratification, real-world response rates
Final Aggregate	**Model:** DeepSeek-R1 **Recommended Sequence:** 1. Mesalamine + induction steroids 2. Infliximab (rapid) or vedolizumab (safety) or tofacitinib / upadacitinib (oral) 3. Risankizumab / ustekinumab (durable) or ozanimod (oral S1P) • Integrates RCT + RWE to tailor by infection/malignancy risk, patient preference, and guideline placement

Table 2: AoT Example Output for UC

Atom question (de-compose step)	Condensed summary output
1. What is the initial approach to inducing remission in moderate-to-severe UC?	Use **systemic corticosteroids** (IV for acute-severe, oral for moderate). 5-ASA may be continued if already prescribed but is not adequate for induction at this severity.
2. When and how should advanced therapies be introduced if patients are steroid-refractory or steroid-dependent?	**Start an advanced agent during the steroid taper** when response is incomplete or taper fails. Choice depends on disease urgency, comorbidities, infection/CV risk, prior drugs, and patient preference (oral vs injectable).
3. How do TNF-α inhibitors, vedolizumab, and JAK inhibitors compare for second-line use?	**TNF-α (infliximab/adalimumab):** fastest biologic, higher infection risk. **Vedolizumab:** slower onset, gut-selective, safest for elderly/high-risk. **JAK (upadacitinib > tofacitinib):** fastest oral, effective post-biologic; watch VTE/zoster.
4. What is the recommended approach if there is inadequate response or intolerance to the second-line agent?	**Switch to a new mechanism (MA)** rather than "cycling" within the same class: e.g., move from TNF or Vedo to **IL-23 (risankizumab/ustekinumab)** or to **ozanimod (S1P).** MoA switching outperforms within-class cycling.
5. How should therapy be tailored for special patient subgroups?	• **Elderly / infection-risk:** vedolizumab or IL-23.• **Hospitalised acute-severe:** infliximab IV.• **Needle-averse / oral preference:** JAK or ozanimod.• **CV/VTE risk:** avoid JAK; use vedolizumab or IL-23.• **Young, high inflammatory load:** TNF or JAK for speed.
6. What does RWE say about sequencing and outcomes?	Early use of advanced therapies lowers colectomy rates; vedolizumab shows fewer serious infections in older adults; MA switching after failure outperforms within-class cycling; JAKs work post-biologic but need risk screening; IL-23s give durable remission as third-line.

Table 3: TextGrad Example Output for UC

Treatment Step	When to Use / Escalation Trigger	Preferred Agents	Patient Subgroups	Key Evidence / Rationale	Monitoring & Safety
Induction (1st-line)	• New moderate–severe flare ▶ start immediately.• Re-evaluate at ≤ 4 weeks	Systemic steroids ± ongoing 5-ASA	• Applies to all patients at presentation	IV/PO steroids remain fastest inducer of remission	• Limit to ≤ 8-12 wk taper • 5-ASA kept only for maintenance benefit
Early Advanced Tx (1st-line biologic / small-molecule)	Steroid-refractory / steroid-dependent after ≤ 4 wk	Inflix. (TNFi)	• Hospitalised acute-severe • Young, high inflammatory burden	ACT-1: 60 % wk-8 remission; real-world 35 % colectomy reduction	TB screen; combine ± azathioprine to curb antibodies
		Vedoliz.	• Elderly / infection or malignancy risk • Latent TB / C. difficile history	EU registry: 42 % fewer hospitalisations vs TNFi	Slower onset (6-10 wk); but lowest systemic infection risk
		Upadaciti nib (JAK1)	• Needle-averse / strict oral preference • Prior biologic failure needing rapid oral option	U-ACHIEVE: 63 % wk-8 remission	CV/VTE & zoster risk; check lipids, vaccinate, avoid if high thrombotic risk
Second-line (after 1st agent fails)	Partial Mayo > 4 at wk 8 or loss of response	Risankizu mab (IL-23)	• Biologic-naïve needing durable control • Multi-exposed but infection-concerned	SURPASS-UC: 65 % wk-52 remission; 30 % ↓ relapse vs TNFi (RWE)	Slower onset (8-12 wk); favourable long-term safety
		Ozanimod (S1P)	• Prior anti-TNF exposure • Wants oral but JAK contraindicated	True North: 27 % wk-10 remission; oral convenience	7-day dose titration; baseline ECG, ocular exam
Third-line / Rescue	Failure of ≥ 2 advanced classes or high-grade dysplasia	• Unused IL-23 (risankizumab/ustekinumab) • Ozanimod (if not yet tried)	• Multi-refractory: IL-23 performs best • Young, QoL focused: consider early trial enrolment or surgery	AGA 2023: recommend colectomy discussion after two classes fail	Annual retinal check for ozanimod; reinforce shared decision-making on surgery
Definitive Surgery	≥ 2 advanced therapy failures or complications	Total collect. with IPAA	• Persistent steroid or biologic dependence • Severe adverse drug events	Modern laparoscopic IPAA offers > 90 % long-term satisfaction	Pre-op fertility counselling; pouchitis education

RESULTS

Similarly to UC, this treatment-sequence optimization pipeline was applied to COPD as per tables 4,5, and 6.

Table 4: MoA Output for COPD

Stage	Output (Summary)
Proposer Layer	**Models:** Meta-Llama-3.1-8B, Mistral-7B, Qwen2-72B • **1st-line (Non-pharmacologic):** Smoking cessation; pulmonary rehabilitation • **1st-line (Pharmacologic):** LAMA/LABA combination inhalers (e.g., umeclidinium/vilanterol, tiotropium/olodaterol) • **2nd-line:** Add ICS to dual bronchodilation (ICS/LABA/LAMA triple therapy) for frequent exacerbations with eosinophils > 300 cells/μL (or ICS/LABA if LAMA-intolerant) • **3rd-line:** PDE4 inhibitor (roflumilast) for chronic-bronchitis phenotype with recurrent exacerbations • **4th-line:** Biologic therapies (anti-IL-5 agents, e.g., mepolizumab) in severe eosinophilic COPD
Final Aggregate	**Model:** DeepSeek -R1 **Recommended Sequence:** 1. Smoking cessation & pulmonary rehabilitation 2. LAMA/LABA combination inhaler 3. Add ICS (triple therapy) for exacerbations + eosinophilia 4. PDE4 inhibitor (roflumilast) in chronic-bronchitis phenotype 5. Anti-IL-5 biologics (mepolizumab) for severe eosinophilic disease • **Integration:** Tailors RCT + RWE by exacerbation risk, eosinophil count, patient preference & GOLD 2024 guideline placement

Table 5: AoT Output for COPD

Sub-Questions	Output Summary
1. Foundational non-pharmacologic interventions?	• **Smoking cessation**: single most effective to slow progression; offer behavioural support + pharmacotherapy (varenicline, bupropion, NRT). • **Pulmonary rehabilitation**: improves exercise tolerance, symptoms, QoL; reduces hospitalizations. • **Vaccinations** (influenza, pneumococcal, COVID-19): lower infection-triggered exacerbations.
2. Initial pharmacologic therapy?	• **LAMA/LABA combination inhaler** (e.g., tiotropium/olodaterol): dual bronchodilation superior to monotherapy for symptom control & 15–25% greater exacerbation reduction; once-daily, single-inhaler → better adherence. • **LAMA monotherapy**: may suffice for low-symptom, infrequent-exacerbation patients.
3. Escalation for persistent symptoms/exacerbations?	• **Add ICS → triple therapy (ICS/LAMA/LABA)** if ≥2 exacerbations/yr **and** eosinophils ≥100–300 cells/μL • **Risks**: ↑ pneumonia (especially older/low BMI). • **ICS/LABA** if LAMA intolerant.
4. Further escalation for refractory cases?	• **PDE4 inhibitor (roflumilast)** for chronic-bronchitis phenotype with frequent exacerbations despite triple therapy: ↓exacerbations 15–20%; GI side-effects limit use. • **Anti-IL-5 biologics** (mepolizumab) for severe eosinophilic COPD (eosinophils ≥300/μL) with persistent exacerbations: ↓exacerbations 20–40% in targeted subgroup.
5. Tailoring to patient subgroups?	• **Frequent exacerbators**: escalate early to triple therapy, then PDE4/biologic if needed. • **Eosinophilic phenotype**: consider earlier ICS or biologic. • **Chronic bronchitis**: add PDE4 inhibitor. • **High pneumonia/CV risk**: avoid ICS if possible; choose safest devices. • **Adherence/device factors**: favour once-daily, single-inhaler regimens.
6. In-class & between-class drug comparisons?	• **Within-class** (e.g., tiotropium vs. umeclidinium): similar efficacy/safety → device, dosing, cost decide. • **Between-class**: LAMA/LABA > monotherapy; ICS added only for exacerbation-prone/eosinophilic patients; PDE4 inhibitors & biologics reserved for specific phenotypes. • **Real-world**: adherence highest with simple, once-daily single-inhaler regimens.

Table 6: TextGrad Output for COPD

Stage	Agents & Interventions	Key Trial Data	Real-World Insights	Patient Considerations
1st-Line	**Non-pharmacologic** • Smoking cessation (varenicline + counselling) • Pulmonary rehab (8 wk home/hybrid) • Vaccinations (COVID-19, PCV20, flu) **Pharmacologic** • LAMA/LABA combo (tiotropium/olodaterol, umeclidinium/vilanterol)	• EAGLES: 34 % 1-yr quit rate (varenicline) • TONADO: +120 mL FEV_1 vs mono	• 30 % ↓ hospitalization post-rehab;40 % real-world dropout • 75 % adherence once-daily vs 50 % twice-daily (COPDBREATHE)	• Universal for all patients • Preferred in GOLD; once-daily improves adherence
2nd-Line	**Frequent exacerbators (≥2/yr)** • **Eosinophilic (eos ≥300):** ICS/LAMA/LABA triple (fluticasone/umeclidinium/vilanterol) • **Non-eosinophilic (eos <100):** LAMA/LABA + azithromycin (250 mg 3×/wk) • **Chronic bronchitis:** Roflumilast or erdosteine	• IMPACT: 25 % ↓ exacerbations with triple • MACRO: 27 % ↓ with azithromycin • REACT: 17 % ↓ with roflumilast	• 40 % clinicians start ICS at eos ≥150 (COLDEN) • Limit azithro to 1 yr for stewardship	• Eosinophils guide ICS • Monitor for ↑pneumonia (ICS) and GI AEs (roflumilast)
3rd-Line	**Refractory eos (eos ≥300):** Mepolizumab (anti-IL-5) **Persistent exacerbators:** LAMA/LABA + ICS/LABA (e.g., salmeterol/fluticasone) if triple not tolerated	• METREX: 20 % ↓ exacerbations (mepolizumab) • BOREAS: emerging data for dupilumab	• Only ~8 % of eligible receive biologics (cost)	• Biologics for severe eosinophilic COPD • Cost and access limit use

CONCLUSIONS

Our multi-stage AI pipeline—combining MoA, AoT reasoning, and TextGrad test-time refinement—demonstrates a clear progression from broad, conceptual regimen proposals to fully evidence-anchored, patient-stratified treatment sequences for UC and COPD. The MoA stage swiftly synthesizes diverse guideline and RWE perspectives into a cohesive sequence; AoT then decomposes this into atomic clinical questions and recontracts them into a logically transparent pathway; and finally, TextGrad backpropagates FDA-endorsed trial data, real-world outcomes, and hard structural constraints (first-, second-, third-line) into the recommendation, yielding an expert-grade plan with explicit subgroup guidance (Table 7).

Beyond UC and COPD, this architecture is readily generalizable to any condition characterized by multi-stage decision pathways and rich RWE availability—for example:

- Oncology therapy sequencing: Integrating trial protocols and registry data to optimize induction, consolidation, and maintenance regimens in leukaemia or lymphoma.

Table 7: Summary of novel AI methodology utilized for treatment sequence optimization

Stage	Core Method
MoA (Mixture-of-Agents)	• Layered LLM proposers & aggregator • Parallel generation → synthesis
AoT (Atom-of-Thought)	• Decompose MoA output into a DAG of atomic clinical questions • Independently solve and re-contract
TextGrad	• Textual backpropagation across LLM feedback loops • Enforces trial/RWE constraints & structural prompts

Limitations:

- **Susceptibility to bias** – Because outputs reflect the quality and biases of the source LLMs and real-world evidence feeds, recommendations may propagate unseen errors unless confirmed in prospective clinical studies.
- **Limited scope**– Development relied mainly on just two conditions (UC & COPD); how well the approach applies to other disease areas is still unknown.

Key Messages

Layered Collaboration: *MoA harnesses heterogeneous LLM proposers and an aggregator to blend RCT guidelines with real-world usage patterns, yielding a robust initial regimen.*

Transparent Decomposition: *AoT's Markov-style DAG (Directed Acyclic Graph) extraction and contraction phases expose each clinical sub-question, ensuring interpretability and logical rigor.*

Evidence-Anchored Refinement: *TextGrad applies textual gradients against prespecified trial and RWE metrics, enforcing structure and subgroup alignment to produce a fully validated, actionable treatment plan.*

Scalable Generalization: *The MoA→AoT→TextGrad framework can be applied to other complex, multi-stage diseases—such as sepsis protocols, oncology treatment algorithms, and precision dosing pipelines—where sequential decision-making and RWE integration are critical.*

REFERENCES

1. Poweleit, Ethan A.; Vinks, Alexander A.; Mizuno, Tomoyuki (2023). "Artificial Intelligence and Machine Learning Approaches to Facilitate Therapeutic Drug Management and Model-Informed Precision Dosing". Therapeutic Drug Monitoring. **45** (2): 143–150. doi:10.1097/FTD.0000000000001078.
2. Gottesman O, Johansson FD, Meier J, Dent J, Lee D, Srinivasan S, Zhang L, Ding Y, Wihl D, Peng X, Yao J, Lage I, Mosch C, Lehman LWH, Komorowski M, Faisal A, Celi LAC, Sontag D, Doshi-Velez F.
Evaluating Reinforcement Learning Algorithms in Observational Health Settings.
arXiv preprint arXiv:1805.12298. 2018. doi:10.48550/arXiv.1805.12298
3. Sanchez P, Voisey JP, Xia T, Watson HI, O'Neil AQ, Tsaftaris SA. Causal Machine Learning for Healthcare and Precision Medicine. arXiv preprint arXiv:2205.11402; 2022. https://doi.org/10.48550/arXiv.2205.11402
4. U.S. Food and Drug Administration. Real-World Evidence. Science & Research Special Topics. FDA. Available at: https://www.fda.gov/science-research/science-and-research-special-topics/real-world-evidence. Accessed July 5, 2025.
5. Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2024. Mixture-of-Agents Enhances Large Language Model Capabilities. arXiv preprint arXiv:2406.04692
6. Fengwei Teng, Zhaoyang Yu, Quan Shi, Jiayi Zhang, Chenglin Wu, and Yuyu Luo. 2025. Atom of Thoughts for Markov LLM Test-Time Scaling. arXiv preprint arXiv:2502.12018
7. Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. 2024. Automatic "Differentiation" via Text (TEXTGRAD). arXiv preprint arXiv:2406.0749

Email: maria.kudela@pfizer.com