


Navigating the Healthcare Contender of YC's Latest Batch: A Due Diligence Narrative for Entrepreneurs and Investors

SEP 10, 2025 • PAID



Share

Abstract

- This essay analyzes the thirteen healthcare startups from the most recent Y Combinator graduating batch.
- Each company is evaluated on business model, founder backgrounds, traction regulatory and technical risks, and total addressable market.
- Market analyses are grounded in healthcare system dynamics with emphasis payer, provider, and life sciences adoption cycles.
- Founder biographies are integrated to assess credibility, domain expertise, and likelihood of execution.
- The essay concludes with a prioritization and ranking of investment opportunities across the cohort.
- Word count: ~7,000

Disclaimer: These are my personal views and do not reflect those of my employer.

The latest Y Combinator graduating batch once again underscores the enduring tension between Silicon Valley's culture of rapid iteration and the healthcare sector's entrenched inertia. What is striking about this class is how heavily it leans into artificial intelligence as a horizontal enabler across nearly every healthcare niche. From AI-powered EMR agents to AI-driven chemistry platforms, the leitmotif is that healthcare's complexity is ripe for automation, pattern recognition, and scaling.

through computational leverage. Yet investors and entrepreneurs alike must resist temptation to be dazzled by AI gloss without grounding each proposition in the unforgiving reality of healthcare procurement, regulatory constraints, and entrepreneur stakeholder incentives. In the following pages, I will analyze each of the thirteen healthcare startups in this batch, not just at the surface level of their slides but with a deeper diligence lens, examining their business models, founding teams, technical feasibility, and the realistic contours of their addressable markets.

Louiza Labs presents itself as a company dedicated to simulating human biology by generating synthetic datasets. The premise echoes the Applied Intuition model, applied here to biomedicine rather than autonomous vehicles. The team, led by Georgia Witchel and Larissa Tyagi, aims to create digital twins of human anatomy that can be used to train medical devices, surgical robots, and large language models for biology. The business model revolves around selling simulated data as training fuel for companies constrained by scarce clinical datasets. The addressable market is immense, given that the medical device industry is a twenty-one billion dollar global market and surgical robotics alone exceeds two hundred billion in potential. The challenge lies in the fidelity of the simulation: unless their synthetic data achieves regulatory credibility, device manufacturers will hesitate to use it in pre-market submission. The opportunity, however, is profound. If Louiza Labs becomes the de facto provider of anatomical simulation, it can create a data moat that scales far faster than traditional clinical trials.

Locata, founded by Alejandro Salinas and Joshua Chang, is positioned at the provider workflow layer. Their AI agents automate referral management for primary care clinics. In the United States, referral leakage costs health systems billions annually, and primary care physicians lose hours per week navigating payer portals, fax machines, and EHR systems to coordinate specialist referrals. Locata's early traction—eighty-four thousand dollars in annual recurring revenue within two months—suggests strong latent demand. Their model is B2B SaaS, selling directly into clinics and health centers. The diligence risk lies in integration complexity. Every EHR handles referral workflows differently, and without seamless integration into Epic, Cerner, or athenahealth, adoption will stall. Yet the founders' technical pedigree

experience at Google, Meta, and Stanford, bodes well for their ability to solve technical scaling. The market opportunity is substantial: if Locata can automate a fraction of the fifty million annual specialist referrals in the U.S., its ARR could scale into the tens of millions.

Parachute, led by Aria Vikram and Tony Yamin, takes a different approach, building infrastructure to monitor and deploy clinical AI within hospital systems. Their thesis is essentially that health systems are going to be inundated with a proliferation of AI tools, and someone must provide governance, monitoring, and deployment infrastructure. Their team has experience in building AI tools for surgeons and military machine learning infrastructure, which signals credibility in high-stake environments. They cite deployment at a leading academic medical center and position themselves against an eleven billion dollar market. The business model is SaaS infrastructure sold to hospitals, charging enterprise contracts for safe AI deployment. The diligence here requires assessing regulatory alignment: the FDA has been explicit about requiring algorithmic transparency and monitoring in clinical AI. Parachute sits squarely in this compliance value chain. Their risk is that hospitals are notoriously slow buyers, and procurement cycles can take years. If they can ride the wave of regulatory mandates, however, they could become indispensable infrastructure akin to a Palantir for clinical AI.

Perspectives Health, founded by Eshan Dosani, Kyle Hyun Woo Jung, and Nicholas Kovalsky, is building AI agents for EMRs, automating administrative workflows. They are live in ten behavioral health clinics and claim to save clinicians two to three hours per day, with ninety-seven percent stickiness. Their first distribution deal gives them access to thirty thousand users, with a projected three point two million dollars ARR at a twenty percent conversion rate. The founders' ability to win early paid pilots indicates they have solved the distribution puzzle better than most. The business model is straightforward SaaS, charging per clinic or per seat. The total addressable market is enormous, as U.S. clinicians collectively spend hundreds of millions of hours on administrative work each year. The diligence questions include whether they can generalize beyond behavioral health into other specialties and whether they

fend off incumbent EHR vendors who may simply bolt on similar AI features. The risk of commoditization is real, but their early traction earns them a high diligence score.

Avelis Health, co-founded by Angel Onuoha and Ahmad Shehu, focuses on auditing medical claims for self-insured employers and health plans. They compare claim medical records to identify errors and fraud, citing fourteen thousand claims audited with over three hundred thousand dollars in errors detected. With a signed three-year contract worth three hundred and sixty thousand dollars and the founders' prior fintech exit, the company demonstrates early execution credibility. The model is SaaS with services layered on top, charging employers and administrators a percentage of identified savings. The market opportunity is vast: medical claims waste is estimated at over one hundred and fifty billion dollars annually in the U.S. alone. A primary risk is defensibility. Many companies have attempted claims auditing, and insurers often prefer to internalize this function. The question is whether Avelis Health can create proprietary machine learning models that materially outperform incumbents and whether employers will trust a startup with sensitive claims data.

Convexia, founded by Ayaan Parikh and Rahul Vijayan, brands itself as the world's first AI-maximalist pharma company. Their focus is using AI to identify overlooked drugs, initially in pediatric genetic diseases. They claim twenty-three thousand compounds in monthly recurring revenue within three weeks, with four pharma contracts signed. The founders, both young Stanford computer science graduates with prior acquisition experience, position themselves as disruptors of pharma R&D. The model is SaaS services, licensing AI-discovered compounds to pharmaceutical companies. The target is effectively the trillion-dollar pharma R&D market, though Convexia will initially target narrow rare disease niches. The diligence concern is credibility: can an eighteen-year-old founder team truly compete in the high-stakes world of drug discovery, where regulatory hurdles, wet lab validation, and clinical trials are inescapable? Their velocity is impressive, but investors must weigh whether pharma will trust critical R&D pipelines to a team without deep biology credentials.

CareSwift, with founders Brian Weigand, Jonathan Zero, and John Paul Mussalli, provides AI scribing for ambulance reports. Their pitch is niche but highly specific: ambulance technicians spend time on documentation that delays revenue capture.

creates compliance risk. CareSwift claims to power twenty percent of NYC emergency medical technicians' reports and has signed a first contract worth forty-five thousand dollars in ARR. Their market is estimated at three point eight billion dollars, encompassing EMS documentation across the U.S. The model is SaaS sold to EMS providers and municipalities. The diligence challenge is scale. EMS is a fragmented market with long procurement cycles, and the niche nature of the problem may be their upside. Nonetheless, their early traction in New York suggests they understand their users' pain points deeply.

Pharmie AI, founded by Neer Jain and Aniruddh Sriram, is pitching itself as the AI pharmacy technician. Their voice AI reduces call volume in pharmacies by seventy percent, while also driving customized upsells. They have one hundred and eighty thousand dollars in contracted ARR within four weeks and are live in ten pharmacies with two chains piloting. The pharmacy market is massive, with more locations than coffee shops in the United States, and an estimated eighteen billion dollar opportunity for workflow automation. The business model is SaaS per location. The diligence question is whether pharmacies will pay for yet another SaaS vendor, given razor margins. If Pharmie AI can prove a direct uplift in revenue from upsells, beyond savings, they may overcome this resistance.

Wedge, founded by Devraj Gopal and Steven Segawa, positions itself as a Palantir-like healthcare AI agents. Their model involves on-site engineers embedding at health systems to build AI agents customized for each institution. They cite one hundred and sixty thousand dollars in ARR across sixty locations. The founders are ex-NVIDIA and Rivian with published AI research, and they have established logos such as Brown Health. The model is hybrid services and SaaS, with the danger that it may become services-heavy and lack scalability. However, their positioning as trusted AI infrastructure providers could capture significant share if health systems embrace agent-based AI.

ChemBio-12, founded by Andres Bran and Zlatko Jončev, is a chemical copilot for pharmaceutical biotech. Their team includes PhDs in computer science and chemistry, with pilots signed with Roche and Takeda. They claim tenfold time savings in molecule design and robotics, with publications in Nature demonstrating their AI's ability to per-

chemical reasoning. The market opportunity is estimated at nine billion dollars, though if validated their impact could expand far beyond. The business model is licensing to pharma R&D labs, charging for access to their copilot platform. The diligence risk is adoption: pharma companies are notoriously slow to adopt new tools, and validation cycles can take years. Yet with marquee logos already piloting, 12 has one of the most credible paths in the cohort.

Blank Bio, led by Ian Shi, Jonny Hsu, and Philip Fradkin, applies RNA-based AI to drug design and smarter clinical trials. Their claim of building the world's best model for drug discovery situates them in a four hundred billion dollar market. With twenty thousand dollars in pilot revenue and six hundred thousand dollars in signed letters of intent, they are already generating momentum. The team's PhD credentials and track record of high-impact publications lend credibility. The business model is a blend of SaaS and licensing. The diligence risk is whether RNA-specific models can outcompete broader foundation models in biomedicine, and whether pharma will contract deeply with an early-stage startup.

Novaflow, founded by Aman Agarwal and Amulya Balakrishnan, serves as an AI analyst for biology labs. Their early traction includes thirty-six thousand dollars in ARR across nine research labs. They emphasize the twelve billion dollars spent annually by labs on data analysis. The founders' academic pedigree includes multiple Nature publications, and their model is SaaS per lab seat. Their risk is academic chronic underfunding and limited software budgets. To succeed, Novaflow must expand from academic customers to pharma and biotech R&D groups with deep pockets.

Phases, co-founded by James Wall, Jonathan van Wersch, and Anton Zaliznyi, automates clinical trial operations with voice AI for patient recruitment. They raised two hundred and forty-six thousand dollars in revenue in eight weeks, spanning twenty-two clinical trials. Their model is SaaS and services for sponsors and CROs. The clinical trial recruitment market is a ten billion dollar opportunity, with persistent bottlenecks that slow drug approvals. The diligence question is scale: Can Phases standardize their AI across trials with vastly different inclusion criteria? Their early traction is impressive, suggesting strong founder-market fit.

When prioritizing across the thirteen, my ranking balances near-term traction with long-term defensibility. At the top, Perspectives Health stands out for its combination of strong early ARR, sticky users, and a clear wedge into EMR workflows. Close behind is Parachute, which could become indispensable as regulatory mandates around clinical AI monitoring accelerate. Louiza Labs earns the third slot as a long duration bet with enormous TAM if their synthetic datasets achieve credibility. Health and Pharmie AI rank next for their clear market pain points and solid early traction, though each faces defensibility questions. b-12 and Blank Bio represent moonshots in biopharma AI, requiring patient capital but potentially transformative. Locata deserves a mid-tier slot for addressing referral inefficiency, though integration risk looms. Convexia, despite its hype and early revenue, ranks lower due to four experience gaps and the brutal realities of pharma adoption. Wedge, CareSwift, Novaflow sit lower due to scalability and budgetary constraints, though each could carve out profitable niches. Phases, with strong early revenue but uncertain scalability, sits in the middle, representing a reasonable but risky bet.

This batch, taken as a whole, reveals YC's deepening bet on AI as the universal solution for healthcare's inefficiencies. Investors must discriminate sharply between surface level AI branding and real traction grounded in clinical workflows, payer incentives or pharma R&D. The best investments will be those that align with the system's but inevitable drift toward automation while building moats based on data, distribution, or regulatory compliance. Entrepreneurs should take note: this batch demonstrates that early ARR and credible pilots with established institutions are the surest way to cut through investor skepticism. In healthcare, as always, adoption is won not in the demo day pitch but in the grinding realities of integration, procurement, and trust-building with clinicians and scientists.

If you have interest investing in these companies, feel free to reach out to treylawles@gmail.com and we can discuss an angel group that I am starting to put together.



2 Likes • 1 Restack

← Previous

Next →

Discussion about this post

Comments

Restacks



Write a comment...

