The YC W26 health tech field notes: what 22 companies at demo day tell u about where healthcare AI is actually going

MAR 26, 2026 • PAID

Abstract

- YC W26 Demo Day: March 24, 2026, approx. 196 companies total

- Healthcare/Biotech companies in batch: 22 (roughly 11% of batch)

- Batch-wide context: 64% B2B, sub-1% acceptance rate, Rebel Fund estimates 3 W26 companies score in top 20% of all YC companies ever evaluated

- Healthcare median seed: ~$4.6M vs. $3.1M batch-wide median

- Beacon Health notable outlier: Accel + Sequoia scout backing, Stanford/Harvar physician co-founder + ex-Amazon Alexa engineer co-founder

- Synthetic Sciences: $1.4M pre-YC + $500K standard deal

- Categories observed across the 22 companies: clinical operations/admin autom primary care AI, drug discovery/computational bio, surgical planning, practice-specific verticals (dental, medspa, infusion), revenue cycle/billing, life sciences infrastructure, language/translation access, wearables, and medico-legal

- Key structural thesis: the batch signals a move from “AI in healthcare” as a con to AI as the operational substrate of specific care workflows and discovery pipel

- Three investability tiers proposed: high conviction (Beacon Health, CellType, I Biosciences, Strand AI, Mango Medical), watch list (Origin, Overdrive Health, MochaCare, Eos AI, Prana, Scheduling Wizard, Rhizome AI, 10x Science), and early/niche (Ruma Care, Patientdesk.ai, Tepali, Mantis, O Health, Synthetic Sciences, Fort, OctaPulse, Docura Health, Wayco)

Table of Contents

Setting the Context: This Batch Is Different

Every YC batch gets called the strongest one yet. Usually that's marketing. For W the data is slightly harder to dismiss. Rebel Fund, which has been running a mac learning model against every YC batch since 2013, published something before a

single company presented: 35% of W26 startups scored in the top 20% of all YC companies ever evaluated. No prior batch has gotten close to that number. The distribution curve didn't just shift at the top end – it shifted across the board.

About 196 companies showed up to Demo Day on March 24th. Of those, 22 tagg themselves under Healthcare / Biotech in the YC system. That's roughly 11% of t batch, which is a meaningful concentration for an accelerator that historically sk B2B SaaS and developer tooling. For context, the batch overall was 64% B2B. Consumer barely registered at around 5%. Healthcare was one of the few categor punching above its historical weight.

That matters for a few reasons. YC acceptance is sub-1% from the application po The healthcare companies that make it through the filter are not random. They'v been pressure-tested by partners who have now seen hundreds of health tech pit across a decade-plus of batches. When YC starts concentrating health deals at 1(a batch, they're making a statement about where they think defensible businesse being built.

What's also worth noting before getting into the individual companies: the medi seed round for YC healthcare startups sits around $4.6M, compared to $3.1M ba wide. That gap is structural. Healthcare takes longer to build, costs more to distribute, and requires regulatory navigation that most other software categorie don't. When evaluating these 22 companies, keep that capital intensity reality in mind. A medspa operating system and a drug discovery company are not the sam business, even if both are checking the Healthcare / Biotech box on the YC webs

The Cluster Map: How to Read 22 Companies

Before going company by company, it helps to step back and look at what the 22 healthcare companies are actually doing as a group. There's a structure here that immediately obvious from the one-line descriptions.

Roughly a third of the batch is in clinical operations and administrative automat These are companies doing scheduling, prior auth, front-office workflow, billing care coordination. Another cluster is in primary care delivery itself, which overla with admin but extends into actual clinical decision support and autonomous ca smaller but very interesting cluster is in drug discovery, computational biology, life sciences research tooling. Then there are the practice-specific verticals cover dental, medspa, and biologic infusion. And at the edges you have a few companie doing things that are genuinely novel and harder to categorize, including digital twins, wearables, multimodal foundation models for patient biology, and AI for medico-legal work.

This structure isn't accidental. It maps almost perfectly to where the friction in healthcare is most acute right now. Administrative costs eat somewhere between 30% of total U.S. healthcare spend depending on whose estimate you use. Primar is in a documented access crisis. Drug discovery timelines are embarrassingly lo relative to what AI should be able to do. Practice management in specialty vertic historically fragmented. YC's batch, whether intentionally or just as a function o where talented founders are gravitating, has built a pretty coherent map of the highest-value pain points.

Clinical Operations and the Admin Automation Wave

The largest cluster in this healthcare batch is companies attacking the administ and operational layer of care delivery. This shouldn't surprise anyone who's beer watching health tech over the last few years. The administrative burden on healt organizations is staggering, and it's getting worse. Scheduling Wizard is building it describes as logistics infrastructure to modernize healthcare operations, whicl category that sounds boring until you realize that poor scheduling coordination of the single largest drivers of wasted clinical capacity in health systems. Getting scheduling right in a complex care environment means navigating provider availability, patient acuity, equipment, room availability, and payer requirements

simultaneously. The companies that have solved this for smaller verticals, like su scheduling, have built very sticky businesses.

MochaCare is positioning as an agentic management service for care organizatio targeting the broader category of care coordination infrastructure. Care organizations, particularly in the managed care and value-based care space, are drowning in the coordination work required to manage patient panels at scale. T traditional model is human-heavy and doesn't scale. Agentic systems that can ha outreach, follow-up, gap closure, and documentation represent a real unlock her assuming they can get EHR integration right.

Eos AI is calling itself an autonomous operating system for healthcare, which is most ambitious framing in this cluster. The OS metaphor suggests it's not just automating individual workflows but trying to become the orchestration layer ac workflows. That's a much harder product to build and sell, but if it works, it's al substantially more defensible. The challenge with this positioning is that health systems are notoriously difficult to sell horizontal platforms to. Most successful health IT companies have started vertical and expanded, not the other way arour Eos AI is either a few years early or perfectly timed, and that's probably the mos important diligence question for anyone looking at it.

Opalite Health is doing something that sounds narrow but has a real structural tailwind behind it: helping healthcare providers communicate in any language. T partly a compliance play given where CMS and OCR are heading on language ac requirements, and partly a quality-of-care play given the documented outcomes for patients with limited English proficiency. Multilingual care coordination has historically been solved with human interpreters at enormous cost and with significant scheduling friction. AI-native language support changes the economi substantially.

Primary Care as Platform: The Most Crowded and Most Important Bet

Primary care is the most contested category in this batch, which makes sense be
it's also the most broken. Over 200 million Americans nominally rely on primary
as their entry point to the healthcare system, and that system is visibly collapsin
Primary care physicians are burning out at record rates, patient panels are grow
while visit times shrink, and preventive care is falling through the cracks at scal
because no one has the capacity to close care gaps manually.

Beacon Health is the name that keeps coming up as the standout in this space. It
building AI employees specifically for primary care, automating the workflows t
eat physician time without adding clinical value: prior authorizations, preventiv
screening follow-ups, referral coordination, and risk adjustment documentation
What makes Beacon notable relative to other primary care AI plays is the combi
of the founding team and early traction signals. The co-founders include a Stanf
and Harvard-trained physician alongside an ex-Amazon Alexa engineer. That's a
pairing of clinical credibility and technical depth. Before Demo Day, Beacon hac
already secured backing from Accel and a Sequoia scout, making it reportedly th
largest healthcare raise in the batch. The EHR integration angle is critical here -
Beacon says it handles workflows directly in the EHR, which is the right distribu
strategy. Anything that requires providers to leave their EHR to interact with a t
will have serious adoption friction.

Prana is going at this from a different angle entirely. Instead of augmenting exis
primary care capacity, it's positioning as an AI primary care doctor in the consu
pocket. That's a fundamentally different business model: consumer-direct rathe
provider-sold. The tricky thing about consumer-direct primary care AI is that it
simultaneously huge as an opportunity and genuinely hard to execute compliant
Scope of practice, liability, licensing, and reimbursement are all live minefields.
companies that have succeeded in direct-to-consumer clinical AI have generally
it through very careful regulatory positioning or by operating in a narrow enoug
clinical lane that the liability surface is manageable.

The tension between Beacon's provider-facing model and Prana's consumer mod
actually one of the more interesting strategic questions in health tech right now.
Provider-facing distribution is slower and more expensive, but it comes with cli

trust, EHR integration, and institutional relationships. Consumer-direct is faste scale but runs directly into the credentialing and reimbursement wall. Where pr care AI ultimately wins is probably in some hybrid – using consumer engagemer to power the clinical intelligence, but monetizing through the provider or payer

Drug Discovery Goes Computational: T Science-Heavy End of the Batch

The most technically dense corner of this healthcare batch is the drug discovery computational biology cluster. These are companies doing things that require se scientific depth to evaluate, and the one-line YC descriptions don't come close t capturing what's actually being built.

CellType is calling itself the agentic drug company, with the thesis that it can simulate human biology at enough fidelity to meaningfully accelerate drug development. The "we simulate human biology" framing is bold, but it points at something real. In vitro and animal models fail to predict human outcomes at a r that is embarrassing and expensive. The ability to run virtual experiments agains better human biology simulations doesn't eliminate the need for wet lab validati but it should dramatically compress the hypothesis testing cycle. CellType is on the harder companies in this batch to evaluate quickly because the science is genuinely complex, but it's also one of the higher-potential outcomes if the simu approach holds up.

Ditto Biosciences is doing something clever and genuinely novel. The thesis is th parasites, viruses, ticks, and worms have spent millions of years evolving protein suppress and manipulate the human immune system, and those proteins are ther a largely untapped library of drug candidates for autoimmune disease. Using AI protein structure prediction and target identification, Ditto is essentially mining evolutionary biology for immune modulation molecules. The founding team is th PhD scientists from UCSF, UCSD, and UC Berkeley with backgrounds in parasi biology, computational biology, and protein structure prediction. These are peop who worked together for three years before founding the company. That's a team

serious scientific credibility and real domain depth, which matters a lot in a cate where most of the risk is scientific rather than market.

Origin is doing AI and data for cancer therapeutics, which covers a lot of ground depending on what the actual product is. The cancer therapeutics space is one w AI has been promised to deliver breakthroughs for years, and progress has been but slower than early hype suggested. The companies making genuine progress a generally the ones with proprietary data assets that give their models an advanta over what you can build with public datasets alone.

Strand AI is building multimodal foundation models to predict uncollected pati biology. This one is worth unpacking. The core idea appears to be using multim data, presumably some combination of clinical records, imaging, genomics, and behavioral signals, to predict biological information that hasn't been directly measured. This is the kind of capability that could be enormously valuable for ri stratification, early disease detection, or treatment selection, depending on what predictions it's actually making and how well it validates. Strand is probably the highest-uncertainty, highest-potential company in the bio cluster.

Synthetic Sciences is taking the tooling angle rather than the direct discovery an It's building what it calls Claude Code for scientific research – a platform where researchers delegate complex research tasks to AI co-scientists. The founders se biggest unlock in ML research, computational biology, and proteomics, where th tooling infrastructure hasn't kept pace with the models. Having raised $1.4M pr with another $500K from the standard deal, Synthetic Sciences had more capitalization coming into Demo Day than most of its cohort.

Surgical AI and the Procedural Frontier

Mango Medical is building foundation models specifically for planning orthoped surgery. Surgical planning is a genuinely interesting wedge. It's a high-stakes, hi value decision point where errors are expensive in both human and financial ter It's also a workflow that currently relies heavily on manual review of imaging an implant specifications, which is exactly the kind of task where AI has shown

consistent performance improvements. The orthopedic market specifically is attractive because implant selection and surgical planning are tightly linked to a expensive supply chain. A company that can demonstrate it reduces revision rate improves implant fit precision is going to get attention from both payers and the implant manufacturers who would love to take share from each other. The SaaS here also has a natural integration point with surgical scheduling systems and th EHR, giving Mango Health several potential distribution paths.

The surgical AI category more broadly has been heating up since the FDA starte getting more comfortable with AI-assisted imaging analysis and decision suppor Companies like Mango Medical are entering a market where the regulatory path getting incrementally clearer, even if it's still not simple. That's a better situatio it was even two or three years ago.

Practice Verticals: Dental, Medspa, and Infusion

Three companies in this batch are going after specific specialty practice verticals the logic is sound even if the individual markets are smaller than primary care o pharma.

Patientdesk.ai is building an AI front and back office age dental practices. Dental is a great vertical for this approach for several reasons. practices are predominantly small and independent, their administrative workfl are standardized enough to automate at scale, and they're historically underserv software vendors relative to their medical counterparts. DSOs, the dental service organizations consolidating the industry, are also a natural enterprise distributi channel if Patientdesk can show unit economics that hold up across a managed g

Tepali is doing an AI-native operating system for medspas, and the timing here i interesting. The medspa market has grown fast and fragmented fast. There are n somewhere around 8,000 to 10,000 medspas operating in the U.S., most of them owner-operated businesses with revenue in the $500K to $3M range. They need scheduling, inventory management, consent management, treatment documenta

billing, and increasingly sophisticated CRM. The fact that Tepali also tags Finte its categories suggests it's probably looking at payment and financing as part of stack, which makes sense given how much medspa revenue flows through consu financing.

Ruma Care is building the operations stack for biologic infusion clinics. This is maybe the most interesting vertical choice of the three because biologic infusion rapidly growing category driven by the GLP-1 and autoimmune drug markets, it clinically complex enough that generic practice management software doesn't ha it well, and the revenue per patient is high enough that operators care deeply ab operational efficiency. A clinic that infuses biologics at $10K to $50K per patien year has a very different operational profile than a primary care practice, and Ru appears to be building to that specific complexity.

Revenue Cycle and Billing AI: Boring Name, Massive Market

Overdrive Health is building AI-native medical billing services. This is a categor where the opportunity is massive, the pain is real and well-documented, and the execution path is very hard. Medical billing is not a new problem. It's been the graveyard of many well-funded health IT companies. What's different now is tha combination of large language models for claims interpretation and autonomous agents for follow-up and correction workflows gets the unit economics to a plac wasn't achievable before. The companies that fail in this space usually fail becau they underestimate the complexity of payer-specific claim rules, the pace of paye policy change, and the cost of errors. The ones that succeed build a very strong p intelligence layer that gets better with volume.

The medico-legal angle is covered by two companies that sit at the edge of healt and legal. Docura Health is calling itself an AI-native med-legal firm. Wayco is building an AI operator for medico-legal cases. These are different angles on the underlying workflow: the process of using medical records to support or defend claims is extremely document-heavy, time-consuming, and prone to errors from

manual review. Both companies are betting that AI can accelerate the record rev summarization, and documentation chain substantially. Given that medical-lega is typically billed at several hundred dollars per hour for highly trained professi the unit economics of automation here are compelling. The market is smaller tha primary care but the margin structure is attractive.

Life Sciences Infrastructure and Research Tooling

Rhizome AI is building an agent platform for life sciences, which puts it in the increasingly crowded category of agentic infrastructure specifically for biopharr research workflows. Life sciences companies spend enormous amounts on knowl management, regulatory documentation, clinical trial coordination, and R&D pr management. These workflows have historically been managed with a combinati expensive specialized software and an enormous amount of human labor. Agenti systems that can handle the coordination, documentation, and search tasks with these workflows have a real enterprise sales opportunity, though the sales cycle a implementation complexity are both significant.

10x Science is building what it calls the AI-native platform for next-generation protein characterization. Protein characterization is a bottleneck in both drug discovery and industrial biotech. The ability to accelerate the throughput and qu of protein characterization workflows directly translates into faster developmen cycles. This one sits at the infrastructure layer of life sciences research, which m it has a long sales cycle but potentially very high lifetime value per customer.

The Sleepers and the Long Shots

Fort is building a strength tracking wearable that tags both Consumer and Healt in its categories. The consumer health wearable space is notoriously difficult to a durable business in. The hardware economics are brutal, the competition from Apple and Garmin is real, and clinical reimbursement for wearable data is still v limited. Fort's only clear path to being a healthcare business rather than a consu

gadget is if the strength tracking data can be credentialed enough to matter in c or insurance contexts, which is a long road.

OctaPulse is using computer vision and robotics to automate quality inspection farms, which is tagged Healthcare / Biotech but is more accurately an aquacultu food safety play. It's an interesting use case for computer vision applied to biolo systems, and there's a real food supply chain argument for why fish health monit matters from a public health perspective, but it's an unusual fit with the healthc bucket.

Mantis is calling itself digital twins of humans, which is the most ambitious frar in the entire batch. Building credible physiological digital twins is one of the genuinely hard unsolved problems at the intersection of AI and biology. The cat has attracted serious academic and industrial investment. Whether a YC-stage company is actually building digital twins in any clinically meaningful sense, or whether it's using that language to describe something narrower, is the key dilig question.

How to Think About This as an Angel

For the angel and early-stage investor trying to figure out where to put attention this batch, a few frameworks are worth applying.

First, distribution clarity matters as much as technology. The best technology pl this batch is probably CellType or Ditto Biosciences, but the companies with th legible near-term distribution are Beacon Health, Overdrive Health, and Patientdesk.ai. That's not a knock on the science plays – i recognition that for a pre-seed or seed check, the question isn't just "can this wo but "can this get to a Series A on the capital they're raising." Companies with cle EHR integration strategies and specific provider customer segments are going to an easier time showing the traction metrics that Series A investors want to see.

Second, the admin automation cluster is probably more crowded than it looks. B Health, MochaCare, Eos AI, and Scheduling Wizard are all, in different ways,

attacking the operational layer of care delivery. These are not identical businesse they're competing for the same IT budget and the same limited pool of health sy champions who are willing to sponsor new vendor relationships. Differentiation through specific workflow ownership, rather than general intelligence claims, is to separate the winners from the also-rans.

Third, the science-heavy companies (Ditto Biosciences, CellType, Strand AI, 10x Science) are playing a very different risk/return game. They have longer timeline higher capital intensity, and more binary outcomes. For angel syndicates, the qu is whether you have the portfolio construction to hold a position through multip rounds over a seven-to-ten-year horizon, because that's realistically what the tim looks like for a drug discovery company that is pre-clinical at YC stage.

Fourth and finally, the practice vertical companies are probably the most underappreciated segment of this batch from an angel return perspective. Ruma Tepali, and Patientdesk.ai are all going after markets that large enough to build real businesses in but small enough that the Epics and Cer of the world aren't fighting hard for them. Vertical SaaS with high switching cos growing specialty market is one of the more reliable archetypes in health tech investing. The exit multiples aren't always unicorn-scale, but the path to profital and M&A exits from strategic buyers is genuinely clear. Health systems, DSOs, a private equity-backed specialty groups are all active acquirers of well-built prac management software.

The W26 healthcare batch is a microcosm of where the entire sector is heading: from "AI as a feature" and toward AI as the core operational architecture of spec high-value workflows. The companies that win from this batch will be the ones t picked a narrow enough beachhead to own completely, integrated deep enough i existing clinical systems to be hard to rip out, and built a data asset along the wa makes their models meaningfully better than anything a new entrant could build scratch. That's the playbook. Watch who's executing it.

1 Like

← Previous

Discussion about this post

Comments Restacks

Write a comment...