DCVC

002

Deep Tech Opportunities Report

2024

Investing for Resilience and Abundance

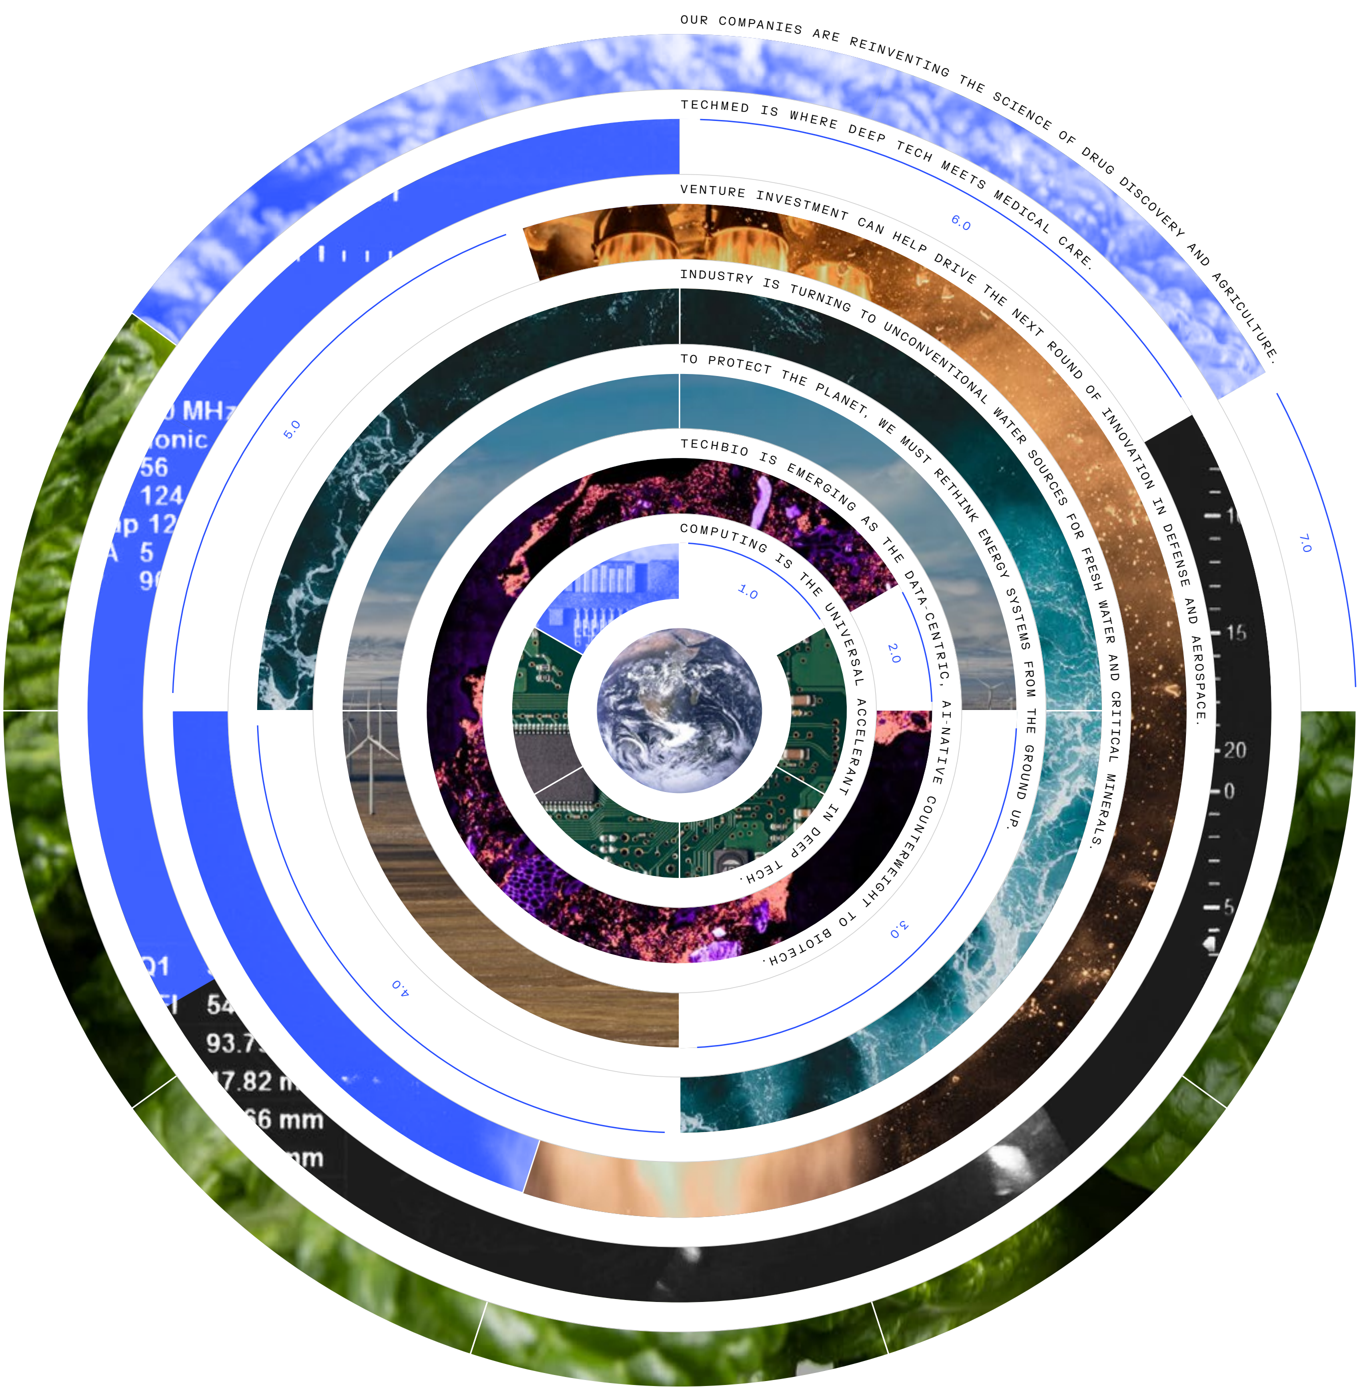

herapeutics ▸ Tidal Metals ▸ Rigetti Computing

Primer  Pivot Bio ▸ Mythic 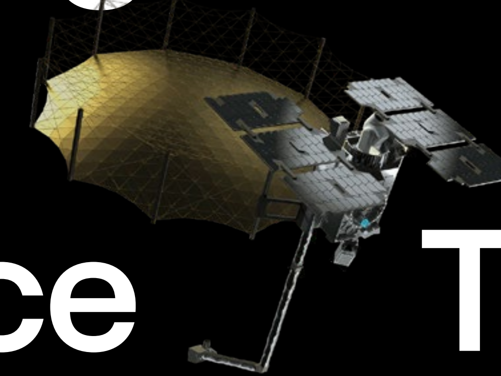 Creyon Bi

Freenome  Capella Space Twelve 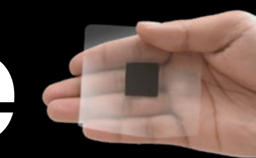

 Elo Life Agility Robotics ▸ Fourth Power

Alchemab Noetik 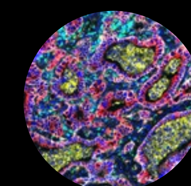 CH4 Global 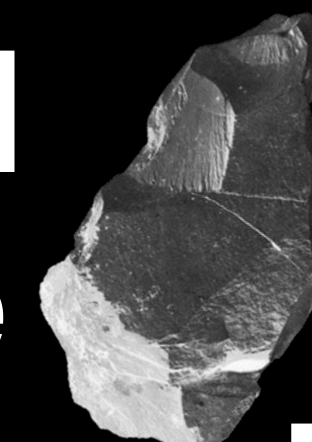 Rem

tics ▸ Oklo Valar Labs ▸ Brimstone Nium

Mosaic ML 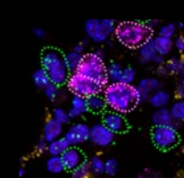 Kanvas Biosciences Planet

 Relation Therapeutics ▸ Recursion Pharm

g ▸ Caption Health  Databricks 

Deep Tech can transform the world for the better. In this report, we describe compelling investment opportunities in seven areas—highlighting the progress of some of the extraordinary companies it's our privilege to back.

ogies ▸ Rocket Lab Solu Therapeutics 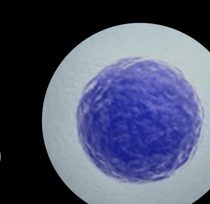 Q-

ing Chroma Medicine  Fervo Energy ▸ R

als Proprio  Umoja Biopharma ▸ Empiric

ch Fortem Technologies 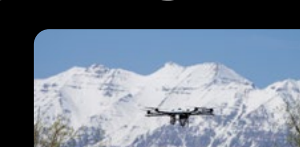 Radiant Nu

DCVC is deep tech venture capital.

Behind all of our investments there's a guiding agenda, an intellectual foundation, that we don't talk about much in public, but which differentiates the firm from most other venture funds. "Deep tech" is the shorthand, but the full argument is richer than any short phrase can capture.

↖ **The total solar eclipse** of April 8, 2024, reminded millions of the forces that shape our planet and the power of science to understand them.

↳ It goes something like this: we are not "techno-optimist" maximalists, but we do have hope. We see the possibility of a future—perhaps by the late twenty-first or early twenty-second century—when people will be living in unprecedented prosperity and equality. The citizens of this future world will enjoy greatly extended lifespans and healthspans. Water, food, and energy will be plentiful. And all this will be accomplished in harmony with a planetary climate that can ensure the long-term survival of our own species and all the others.

This isn't a utopian vision; it's the humane and hopeful alternative to apocalypse. But to support this prosperous and abundant future world, we'll need many new technologies. We love deep tech—and our investments range across diverse disciplines—because we understand the scope of the challenges and the rewards that will come from solving them.

After all, many of the technologies we'll need to decarbonize the power grid, implement CO_2-neutral industrial processes, make agriculture more sustainable, counter disease, and bring the benefits of AI and new computational approaches to everyone now exist in a nascent state. It will take decades to perfect and scale them until they're commonplace.

To build at the required pace for a sustained period, we will need consistent economic growth and a reliable surplus of people, talent, money, and resources. (Venture capital both flows from that surplus and replenishes it—which is why we see no trade-off between investing for profit and investing for the greater good.) The reality, however, is that we're beset *right now* by a range of existential threats including climate change, noncommunicable and emerging infectious diseases, the malicious application of artificial intelligence, and shortages of clean water and other critical resources. Unaddressed, these threats will devour any surplus, divide global polities, and make the prosperous future we envision feel like a faded dream.

Each of these big existential threats must be met by a coordinated counteroffensive.

01 We must decarbonize the global economy to stop adding to, and eventually lower, the concentration of greenhouse gases in the atmosphere. In part, that means finding drop-in, carbon-neutral replacements for nearly all fuels and materials currently derived from fossil hydrocarbons—where “drop-in” means delivering an identical product with the same or better performance, safety, and price, so that industry voluntarily adopts the better product. We think the key to changing the behavior of incumbent companies is to disrupt the system from *inside* the system—that is, to fix capitalism’s externality problems with more capitalism, in a way that delivers its benefits to more people, increasing its returns in the process.

02 Meanwhile, we’re cognizant that a devastating amount of warming is already baked into the future climate, which means we must adapt our infrastructure and physical surroundings to be more resilient to the coming changes in weather and sea level rise.

03 We must prepare for the demographic imbalances spreading across industrialized countries, which will sap industries of the young workers needed to power economies and fund social support systems.

04 We must keep the electrical grid stable and growing—which will mean, in part, using energy-efficient hardware and market arbitrage to limit the electricity needs of the information technology sector.

05 We must ease scarcities of clean water and healthy food.

06 We must find affordable, equitable ways to prevent and treat the noncommunicable diseases that will increasingly afflict an aging world population: cancer, metabolic disease, cardiovascular disease, and neurodegeneration.

07 We must prevent further pandemics like Covid-19.

08 We must ensure that artificial intelligence is deployed in ways that enhance, rather than diminish, our freedom.

09 We must prevent bad actors—whether individuals, organizations, or nation-states—from exploiting global uncertainty through acts of stochastic terrorism or asymmetric warfare.

10 We must accomplish all of this on a smaller, more efficient physical footprint to allow the existence and regeneration of nature, which provides us with assets beyond calculation.

Most of these problems are so big that solving them will require coordinated responses across governments and private sectors.

↳ *And we must do all of this at once,* since these threats are all interconnected and are likely to aggravate one another in a wicked tangle of negative feedback loops.

Most of these problems are so big that solving them will require coordinated responses across governments and private sectors. But we believe that venture capital can play a vital catalytic role, sparking chain reactions of ultimately societal-scale innovation and change, by commercializing ideas that show that the impossible can be solved both scalably and equitably. The technologies that most urgently require and merit deep-tech venture capital investments are those that will interrupt the scary negative feedback loops, boost the positive ones, and reduce the mismatch between our current capabilities and the existential problems we must eventually solve. Such technologies *buy us more time* to build an equitable, free, green, abundant economy.

To give two quick examples: We support advances in the chemistry to produce CO_2-negative Portland cement because they will allow us to build desperately needed infrastructure without adding to the climate crisis (see Opportunity 3.2). And we invest in error-suppression software for quantum computing because we think the world will need quantum computers to keep scaling up the data-heavy foundation models that power the new wave of artificial intelligence without draining electrical grids (see Opportunity 1.4).

Along the way, we gird ourselves against the temptation to invest in buzzy, cynical, parasitic ideas whose short-term payoffs often come at the expense of long-term survival. We think of these ideas as *shiny objects*. Because they contaminate the all-important ecosystem of investment and innovation, they add up to yet another existential threat.

Addressing climate change and similar challenges will require exquisite technical means, control of complex supply chains, chess-master concentration, rational policy support and regulatory change, and a massive shift in capital allocation—but none of these things are possible amidst fear, rage, or scarcity. Our investments range broadly, but they couple, resonate, and reinforce one another because they’re all meant to make our economy more bountiful and resilient. And we believe investors can get paid abundantly for doing all this, producing returns that can be put to use in the next round of a virtuous cycle.

Progress in each of the opportunity areas described in this report will help to counter existential threats, stabilize unsteady markets, and lay the groundwork for long-term prosperity. Founders and investors who aren’t afraid to get their hands dirty doing the risky, prolonged, difficult work of deep-tech innovation should join us. The reward is disproportionate societal *and* economic value.

2024

Deep Tech Opportunities Report

This second annual report surveys the technology applications and trends we're thinking about hardest in 2024. We're convinced that deep-tech companies, backed by patient and experienced investors, can build the new computing capabilities, energy systems, and drugs and medical devices we'll need to survive and thrive in this century

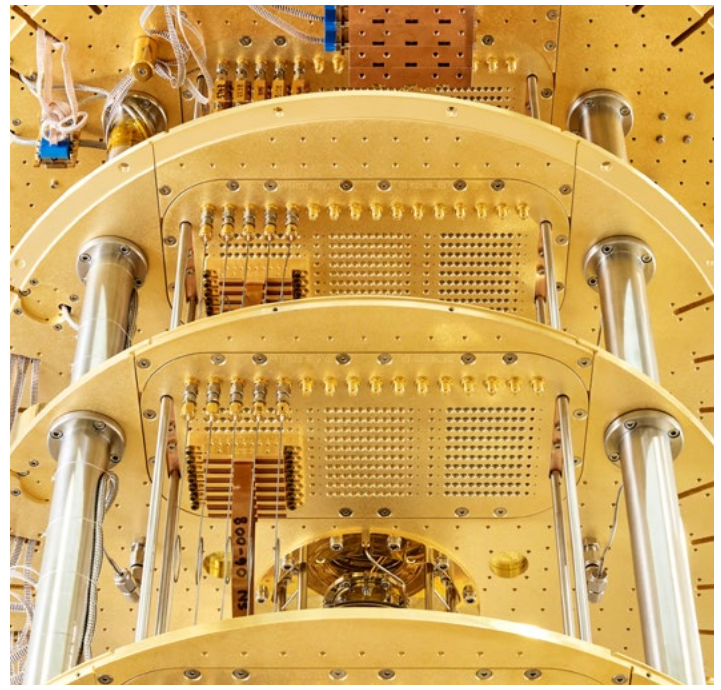

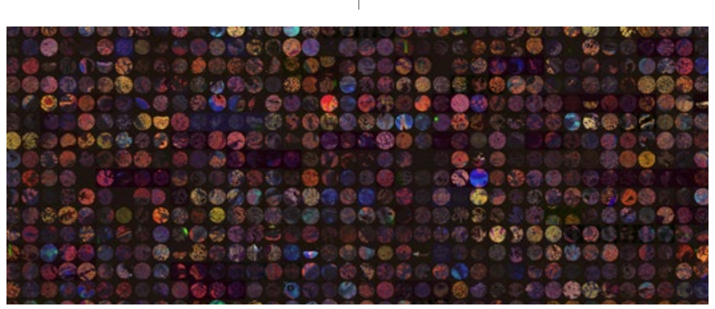

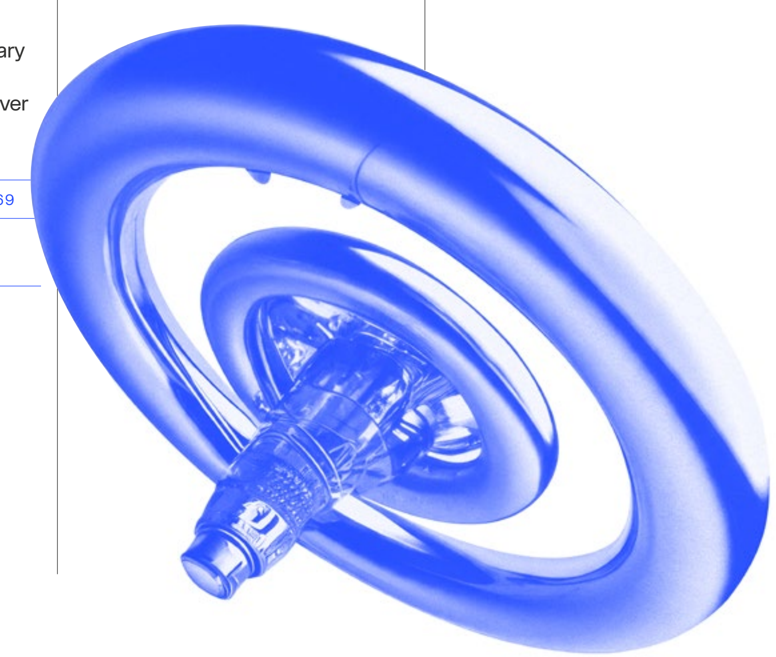

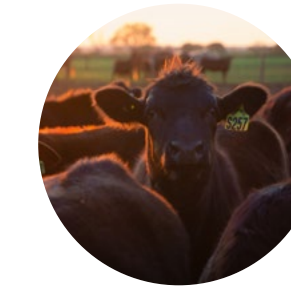

A few words on the investing climate

In today's economy, the companies with the clearest ideas, the best talent, the broadest connections, the most reliable access to capital, and the most immediate paths to market will win—and so will their investors.

HOW ARE WE putting our funds to work to realize the ambitious goals spelled out in this report? And why are we convinced, at an intellectual and a gut level, that DCVC's approach to deep-tech investing will bring venture-scale returns while nurturing companies that change the world?

Like the companies in our portfolio, we're closely attuned to data. So let's look at some numbers. Venture capital continues to be the highest-performing private capital asset class, as it has been almost every year since 2013 (the exceptions were 2016 and 2022). But venture activity has dropped off significantly since 2021, when U.S. venture firms invested a record $348 billion. Amid inflation, geopolitical tensions, and a sclerotic market for public exits and acquisitions, total U.S. deal value fell to $242 billion in 2022 and $171 billion in 2023.

In some ways, this pullback is both expected and welcome. At 2021 levels, too much money was pouring into companies that lacked the fundamentals they would need to thrive. Valuations soared to levels these firms couldn't support through revenue growth, making it harder for them to raise follow-on rounds. Now, with overall investment down, demand for capital is once again outstripping supply. That gives venture firms more leverage to keep valuations and deals to reasonable sizes, which is better over the long haul for both founders and investors (see Figure 0.1.3).

We do anticipate another kind of consolidation in the near future. The number of venture firms in existence has more than tripled since 2007, and total dollars they manage has sextupled (see Figure 0.1.1). A growing number of those firms now invest in deep tech; Boston Consulting Group counts about 800. We see that as a plus, since it means we have more potential co-investors for compelling deals. But when more venture firms chase a fixed supply of good ideas, their average returns plummet, in terms of multiples on invested capital (see Figure 0.1.2). So it wouldn't surprise us to see a dramatic contraction in the venture industry echoing those that followed the bursting of the dot-com bubble in 2000 and the Great Recession that began in 2007. Fortunately for us, times of contraction are exactly when the larger, more experienced, deeper-pocketed venture firms can shine—connecting portfolio companies with the capital, syndicate support, industry-veteran operating partners, marketing strategy, and early customers they need to succeed.

In 2024, one intriguing data point is that even as the amount of venture capital going into the deep-tech sector has fallen, the median post-money valuation for concluded deals continues to hold steady or rise slightly (see Figure 0.1.3). A big part of what's going on here is that funds are exercising more care in deal selection, meaning that the winning companies are stronger and merit larger investments per round. Here at DCVC, we're happy to lead or participate in a large round at a reasonable valuation—*if* our own techno-economic analysis shows that a startup has the team and the technology needed to deliver, as well as a clear path to profitability.

Which leads us to one last point. The companies we back are working on fundamental deep-tech breakthroughs that will *scale in a linear way* to take over existing markets and build new ones. They've found simple, cost-effective, climate-friendly ways to replace or expand resources vital to our economy or our health. (↪ Continued on page 15)

Figure 0.1.1 VC firms and their assets

The number of U.S.-based venture capital firms and the amount of capital they're managing have both increased dramatically since 2007.

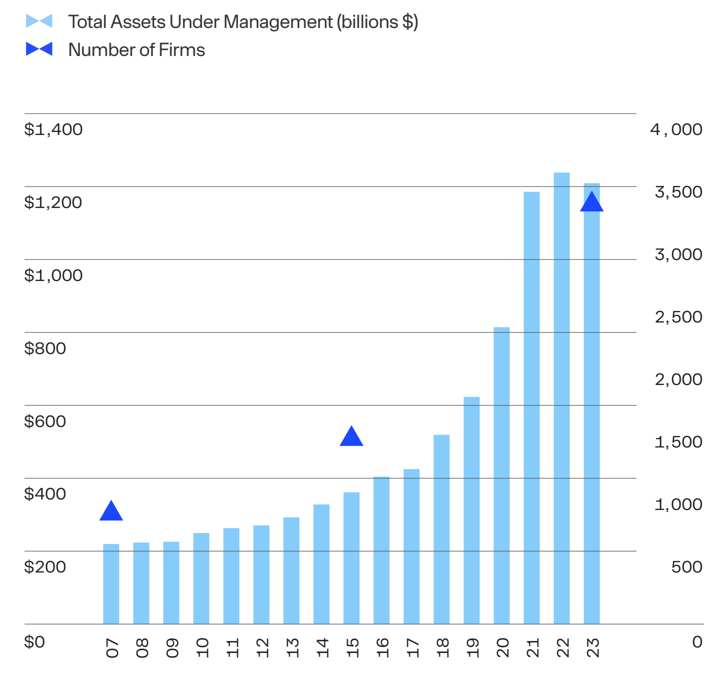

Figure 0.1.2 VC funds and returns: an inverse relationship?

Investment returns for individual venture capital funds, measured by the multiple on invested capital, seem to be higher when there are fewer firms in operation.

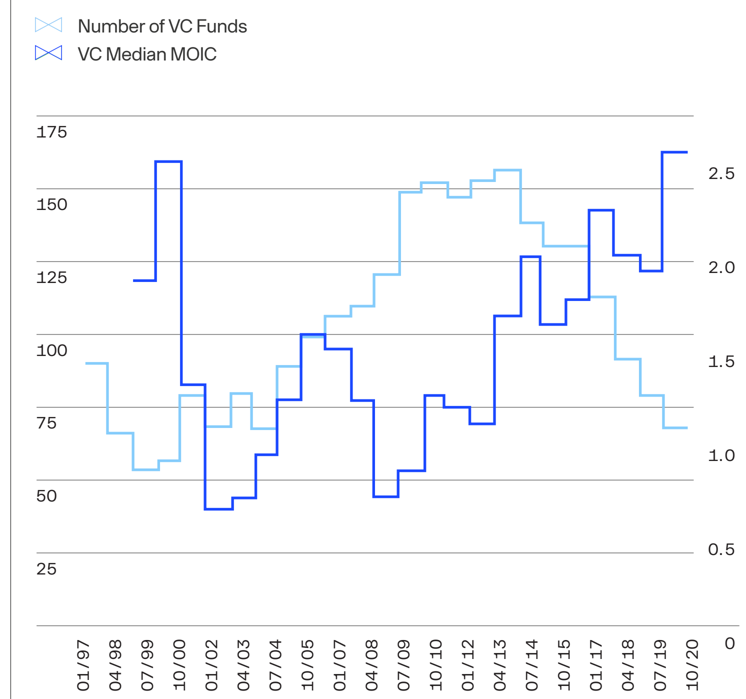

Figure 0.1.3 Valuations hold steady

Deep-tech companies' median post-money valuations at the Seed, Series A, and Series B stages stayed even or went up slightly between 2021 and 2023, even as VC firms poured less capital into startups overall.

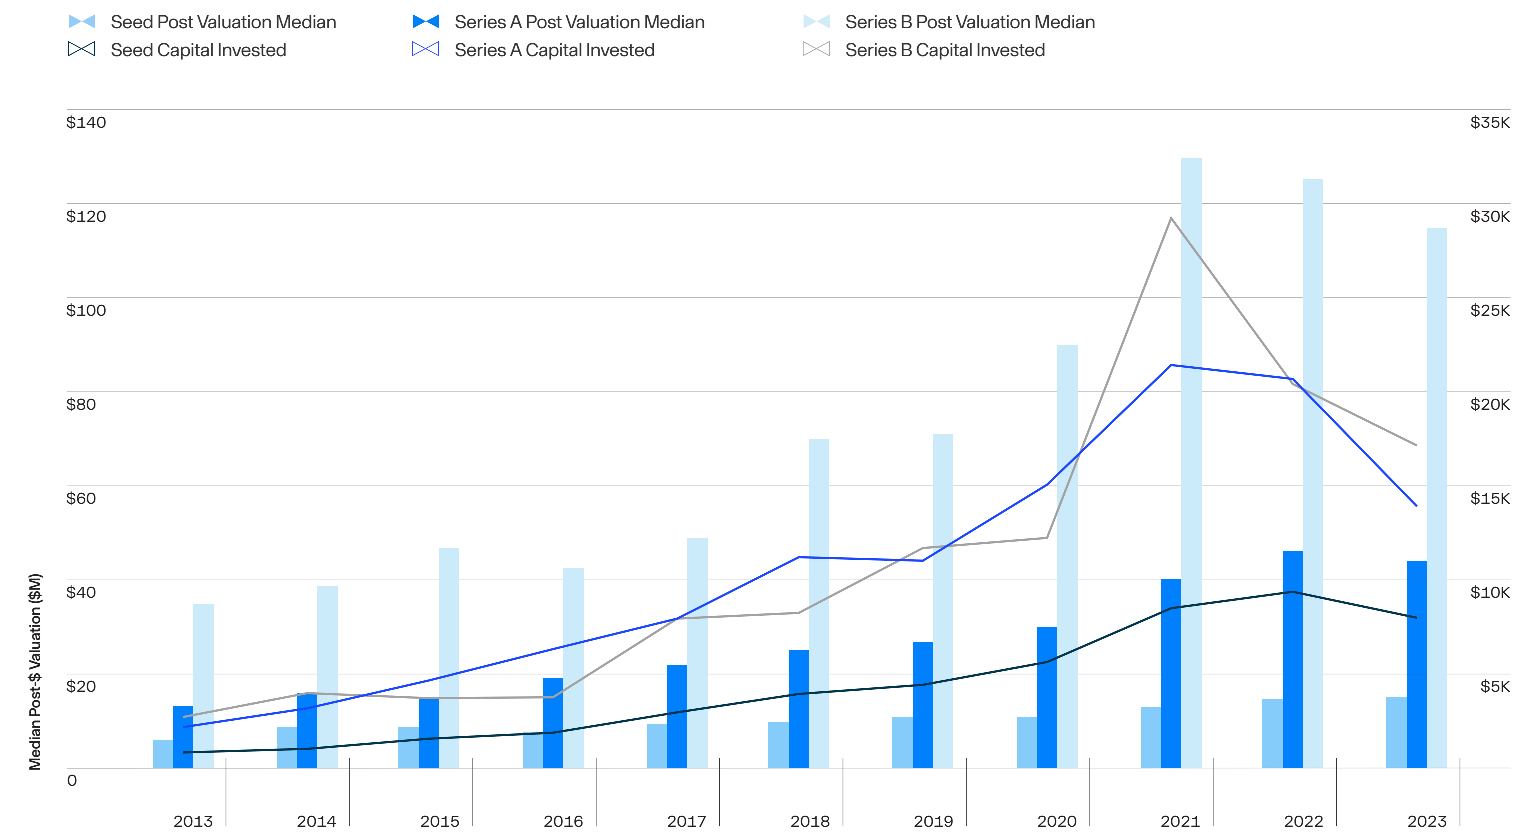

Sources (clockwise from top left): PitchBook, NVCA; Morgan Stanley Research; PitchBook, Dealroom, DCVC analysis

Figure 0.1.4 Fervo's learning curve

Each new geothermal well Fervo Energy drills is completed faster than the previous one. The "spud to total depth" time has decreased from 72 days for the company's first well to about 20 days as of Spring 2024.

Source: Fervo Energy

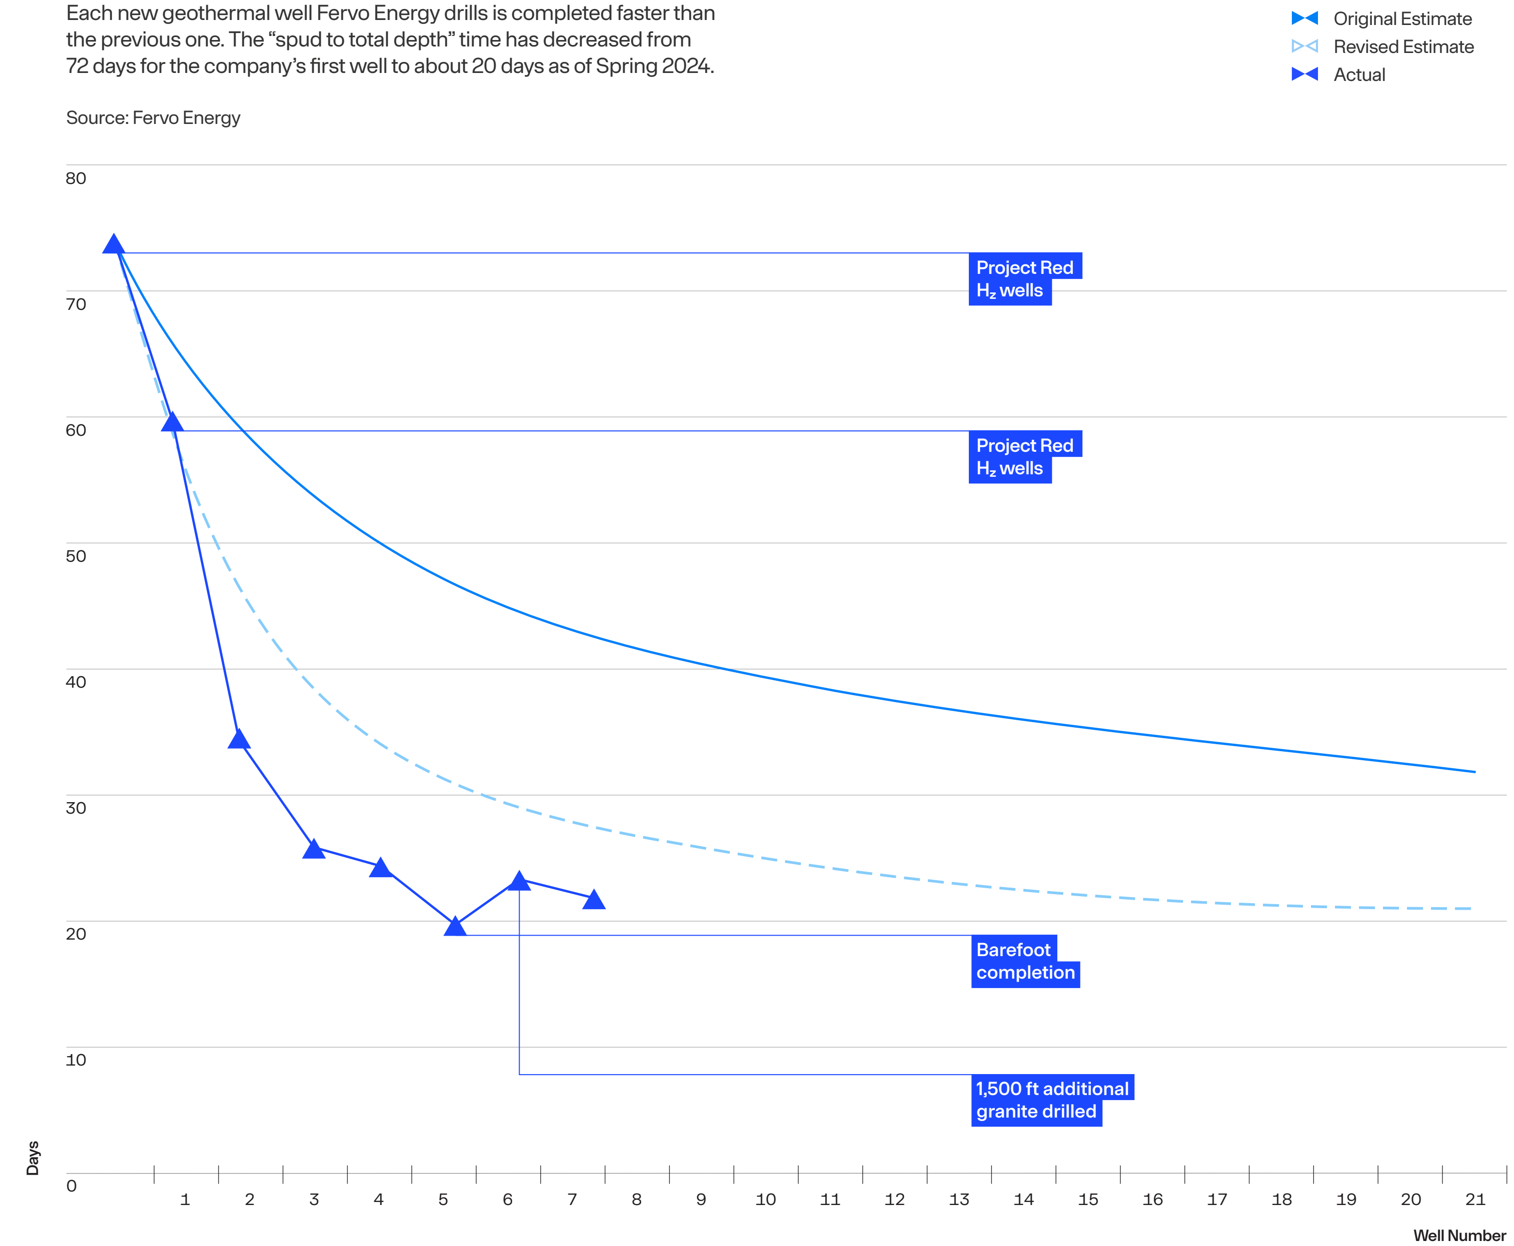

One great example is Fervo Energy, which is repurposing sideways-drilling and hydraulic fracturing technologies from the oil and gas industry to unlock previously inaccessible sources of geothermal energy (see Opportunity 3.1). After successfully bringing online its first commercial plant in Nevada, the company is now on a rapid learning curve, drilling each new well more quickly and cheaply than the previous one. The big technical hurdles have been cleared, and the company's job is now simply to replace as much of the nation's fossil-fueled electrical generating capacity as possible with renewable, firm geothermal power.

They're a typical DCVC company in that they have an **industry-leading team** applying a **brilliant but straightforward engineering insight** to create **a product that can immediately help to solve a trillion-dollar problem (or more)** whose solution will benefit the world. Our ability to discover and support founders like Fervo's is what gives us confidence, even in challenging and uncertain times, that we'll be able to deliver on DCVC's vision.

Zachary Bogue and Matt Ocko
Founders and Managing Partners, DCVC

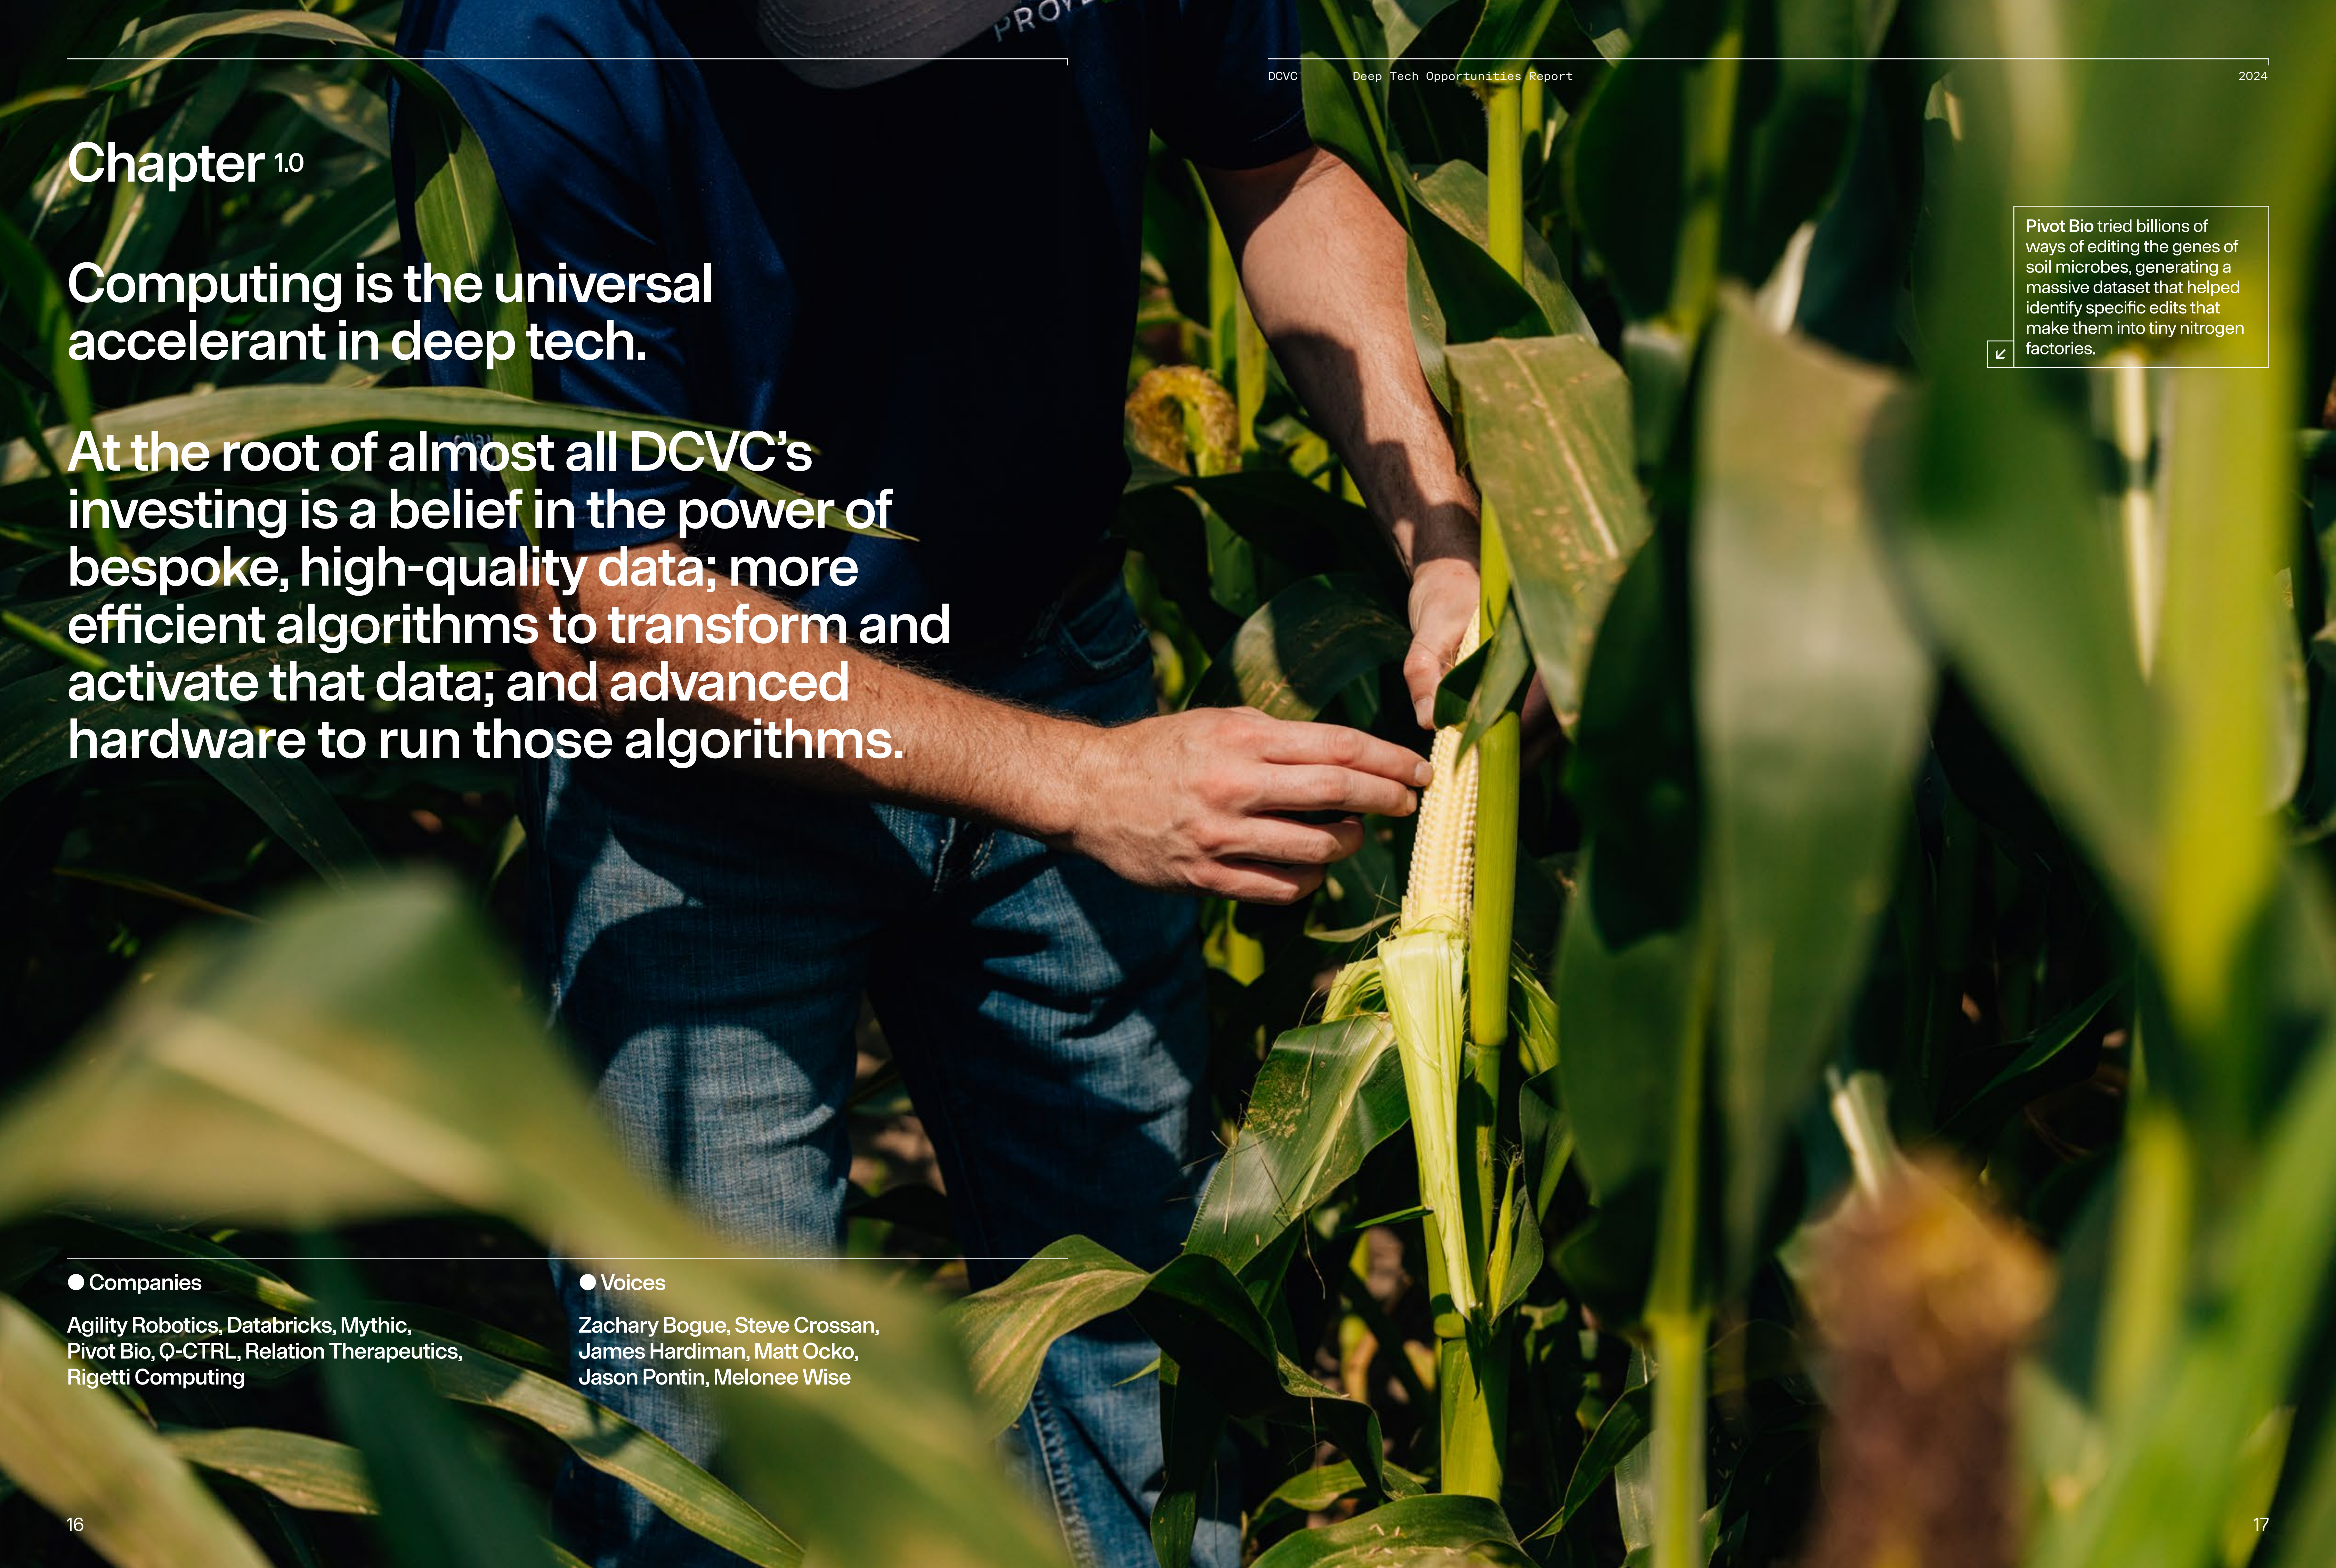

Chapter 1.0

Computing is the universal accelerant in deep tech.

At the root of almost all DCVC's investing is a belief in the power of bespoke, high-quality data; more efficient algorithms to transform and activate that data; and advanced hardware to run those algorithms.

Pivot Bio tried billions of ways of editing the genes of soil microbes, generating a massive dataset that helped identify specific edits that make them into tiny nitrogen factories.

● Companies

Agility Robotics, Databricks, Mythic, Pivot Bio, Q-CTRL, Relation Therapeutics, Rigetti Computing

● Voices

Zachary Bogue, Steve Crossan, James Hardiman, Matt Ocko, Jason Pontin, Melonee Wise

OPPORTUNITY 1.1

IN THE AI RACE, DATA WILL EMERGE AS THE CRITICAL ASSET.

Generative AI models are utterly dependent on large-scale computing; that's why OpenAI's CEO, Sam Altman, aiming to make sure the company never runs up against hardware-imposed limits on AI's growth, says he thinks the company and its suppliers will have to raise and spend *trillions* of dollars on new semiconductor technology. But here's the thing: building better chips and designing the algorithms that run on them are engineering problems. Finding high-quality data is not—and an AI model is only as complete as the data it's trained on. What we're seeing is that in field after field, from robotics to energy to TechBio to traditional biotechnology, the most successful deep-learning or foundational models are those that have been trained using high-quality, curated, often proprietary data.

The story of AlphaFold, the AI program from DeepMind that transformed computational biology beginning in 2018, is usually framed as a victory for deep-learning algorithms. The problem it solved—predicting a protein's 3D structure based solely on its amino acid sequence—was long considered intractable, but the AlphaFold team showed that it could be cracked using a form of attention-based transformer. What's sometimes left out of the story, however, is that AlphaFold was trained to recognize plausible structures using public data banks of 170,000 proteins with known structures, derived through laborious crystallographic analysis and other methods over many decades. "Can you train AlphaFold on suppositions? No," observes DCVC co-founder and managing partner Matt Ocko. "A huge volume of X-ray crystallography had to exist before AlphaFold worked at all."

Steve Crossan was the first product leader of the AlphaFold team at DeepMind and is now an operating partner at DCVC. He says AlphaFold succeeded by combining the right AI models with the right data—but that in the deep-learning game, it's ultimately the data that's more precious. "The technique is not super hard to copy. It tends to be the case that the model, once it's published, becomes commoditized very quickly," Crossan says. "So if you do have a source of proprietary data—especially if you have a compounding source of proprietary data that is going to get better as your product does better in the world—then *that* is a sustainable advantage."

"The model, once it's published, becomes commoditized very quickly. So if you do have a source of proprietary data – especially if you have a compounding source of proprietary data that is going to get better as your product does better in the world – then *that* is a sustainable advantage."

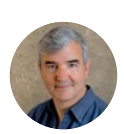

Steve Crossan
Operating Partner, DCVC

Relation Therapeutics gathers gene sequences and other data about human bone cells, uses that data to design experiments that test the roles of specific genes in osteoporosis, and uses the results to train machine-learning models that, in turn, inform the next round of experiments.

DCVC portfolio company Relation Therapeutics is a case study in that kind of compounding. The company's specialty is combining models and new data to find biological targets for drugs that treat osteoporosis and other polygenic conditions (those involving mutations across many genes). It does that using a circular workflow it calls "lab-in-the-loop." Relation sequences the genomes and RNA transcriptomes of single human bone cells from healthy and sick patients. That data feeds into active-graph machine-learning models, which predict which gene variants put people at highest risk for the disease. Then the company uses CRISPR to knock out those genes in new cell lines, singly and in pairs, and quantifies how the changes affect bone mineralization—a marker of osteoporosis. That, in turn, helps Relation's researchers zero in on interventions that might modify the course of the disease.

The company's AI models are smart, but it's the lab data Relation is gathering that makes them effective—and the same approach should help to untangle the causes of other disease that have traditionally been difficult to treat because they're polygenic. "Because Relation's AI models have been trained on all kinds of genetic data, not just osteoporosis data, they can be deployed quickly to many diseases," says DCVC general partner Jason Pontin, who sits on Relation's board. (For more on Relation, see Opportunity 2.1.)

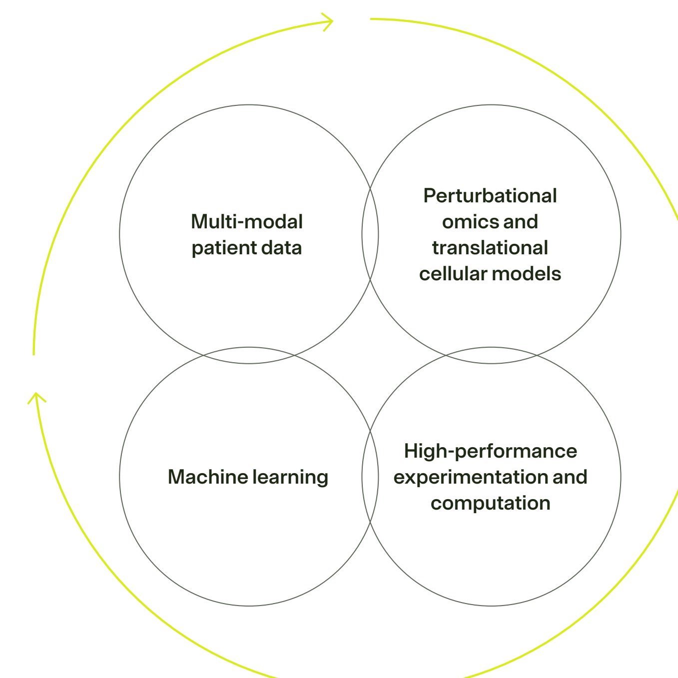

Pivot Bio's work shows how data reigns supreme in another field, agriculture. The company, which DCVC seeded and helped launch in 2014, sells nitrogen-fixing microbes for corn, wheat, and small grain crops called PROVEN 40 and RETURN. They're composed of microbes that, when applied at planting, find their way into the rhizomes of crop roots and begin turning atmospheric nitrogen into nitrogen accessible to the plant. This, in turn, drastically reduces a crop's need for synthetic nitrogen fertilizer—the manufacturing and use of which generates 5 percent of global CO_2 emissions, and whose use pollutes waterways.

Figure 1.1.1 Pivot Bio's breakthrough innovation

▸ Part 1
Mapping trillions of soil microbes, Pivot Bio scientists pinpoint one carrying the genetic code needed for nitrogen fixation, and evaluate billions of ways to edit the microbe's DNA to turn it on.

▸ Part 2
Applied at planting, this microbe naturally adheres to the root, fixing atmospheric nitrogen (N_2) into ammonia (NH_3).

▸ Part 3
These microbes continuously fix nitrogen, providing the plant with precise, steady NH_3 throughout the growing season.

The company created the additives by mapping trillions of natural soil microbes to find those with the genes needed to fix nitrogen directly. Then it tried billions of ways of editing those microbes' genomes, generating a massive dataset that helped isolate the specific edits that disabled natural braking systems and further enhanced nitrogen fixation. "These guys know more about the genomes of their species of soil bacteria than anyone on the planet," says Zachary Bogue, co-founder and managing partner at DCVC. "With computational breeding, they can pack millions of years of evolution into a few tweaks."

It's not just in the biological sciences that data is having its day, of course. As investors with a broad deep-tech lens, we're seeing some of the same advances in fields like chemistry, physics, materials science, energy, and even robotics (see Opportunity 1.2).

In sum: the surprising leaps forward in the power of AI models in 2022 and 2023 shouldn't lull anyone into forgetting the years of effort and billions of dollars that went into creating and curating the data used to train them. Ocko points out that DCVC helped to kick off the big-data revolution in the early 2010s (we began as Data Collective)—so the firm has been integral to the evolution of the field as organizations have gradually found ways to store, organize, clean up, and enrich their data, to the point that it finally could serve as the critical fuel for today's AI algorithms. Those algorithms understand only as much as we tell them—which is why proprietary data is and will remain king.

> "These large models are just things that translate human intent against billions of dollars of data. They're idiot savants that don't actually understand anything beyond their specific domain. Data is king, not models."

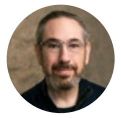

Matt Ocko
Managing Partner, DCVC

OPPORTUNITY 1.2

AGILITY ROBOTICS IS A REAL-WORLD CASE STUDY IN OPERATIONAL DATA.

DCVC-backed Agility Robotics is the world's leading robotics company on a path to mass production of bipedal robots. As it deploys its robots in workplaces such as warehouses and shipping centers, it's gathering the on-the-ground data that will be the key to making human-robot collaboration practical and efficient.

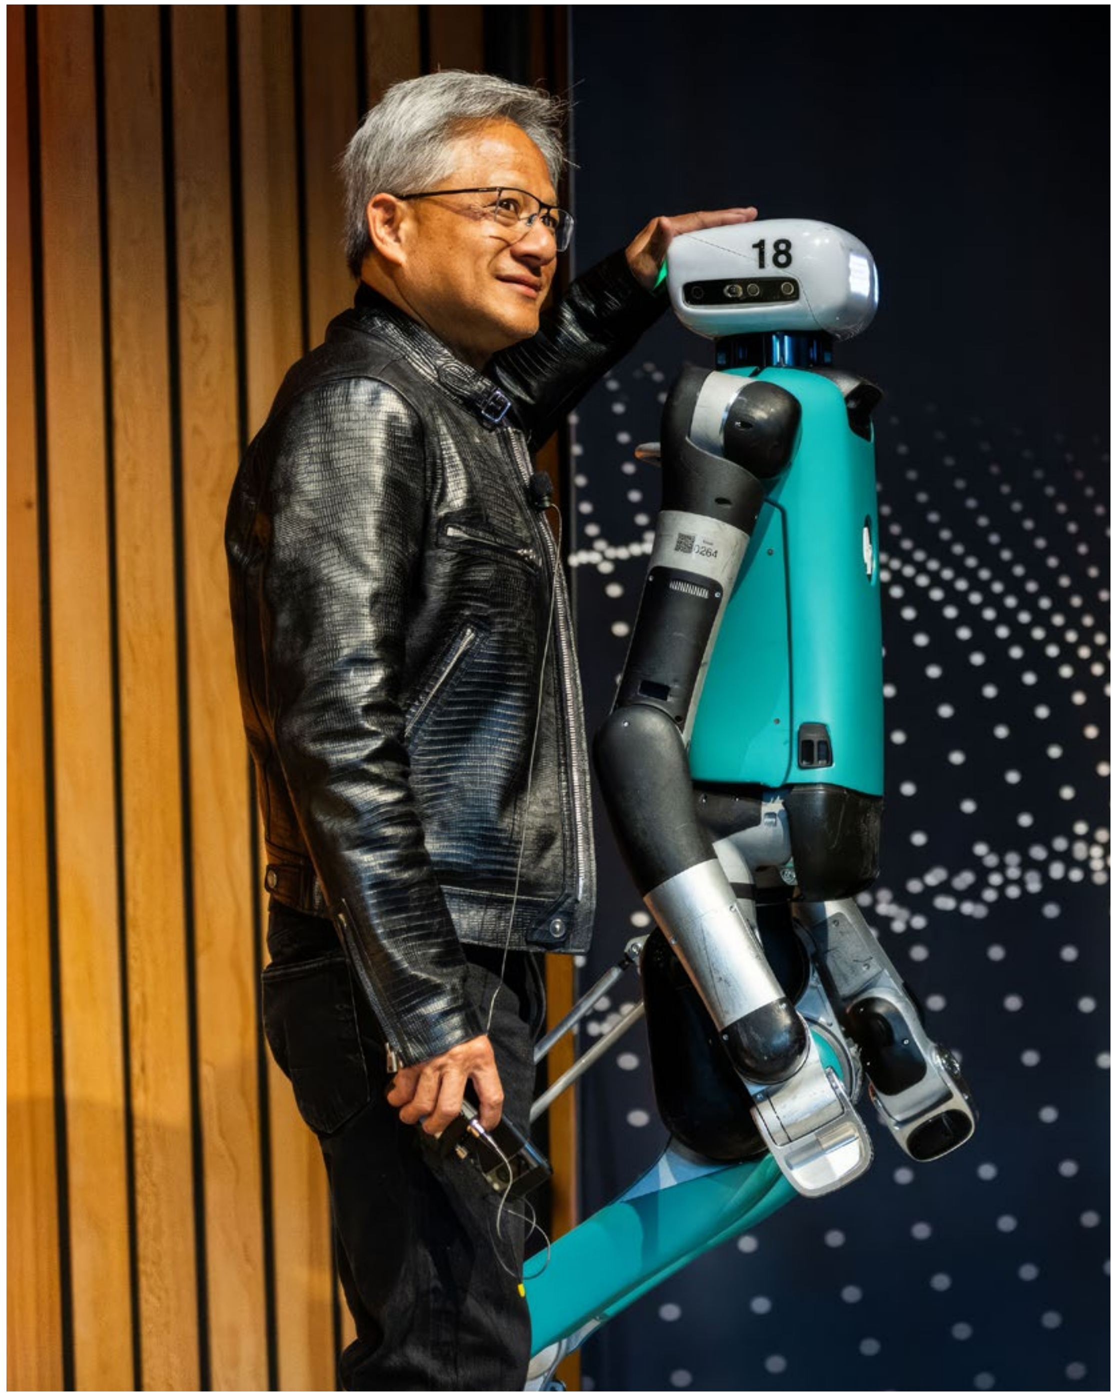

Agility's flagship humanoid robot, Digit, joined Nvidia CEO Jensen Huang on stage at the Oregon State University Global Futures Forum in April 2024.

There was a long period in the history of computing—decades before today's explosion in deep learning—when many researchers believed that building elaborate *expert systems* would be the key to unlocking artificial intelligence. These were programs painstakingly coded with rules and symbolic logic about the world; they amounted to massive cascades of if-then statements attempting to account for every decision point that programmers could foresee within a given domain. Expert systems ruled AI research into the 2000s. Eventually, however, it would become clear that *a*) as expert systems grow in scope, it's nearly impossible for programmers to anticipate every possibility or write every rule in advance, and *b*) advanced neural networks trained on huge amounts of real-world data can, in a sense, intuit the rules of the world on their own.

DIGIT SPECIFICATIONS

HEIGHT
5' 9"

WEIGHT
200 lbs

CARRY CAPACITY
35 lbs

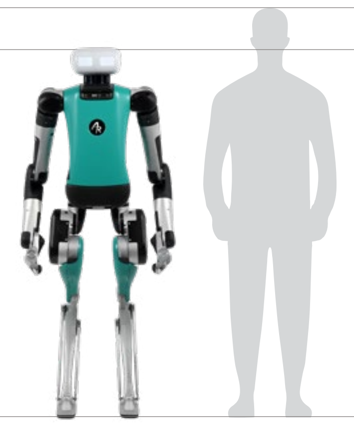

But just as it was a mistake in the 1980s to believe that human programmers could build omniscient expert systems, it's a mistake today to assume that deep learning, by itself, can cut through every real-world problem. Today's most powerful foundation models, such as GPT-4, learn mainly from the statistical patterns in vast amounts of data, but researchers are also experimenting with models that have been fine-tuned with question-answer pairs that help them reason more explicitly about things like the spatial world (what's usually next to a door? A wall), factoids (what was John Wayne's given name? Marion Morrison), and commonsense situations (where would you not want a fox? In a henhouse). And in deep tech, there are many domains where progress depends on a judicious mix of physics understanding, real-world data, and old-fashioned expert systems. One of those is robotics.

Agility Robotics, which we've backed since 2020, makes and sells Digit, a flexible bipedal robot designed to work alongside humans in factories and distribution centers designed for humans. The company is currently testing Digit in partnership with Amazon Robotics for tasks like carrying containers, called totes, from shelves to conveyors and recycling empty containers.

In recent demonstrations, Agility has shown that it's possible to use a large language model to translate natural-language requests from human users into plans Digit can carry out. If a user says, "Take the box that's the color of Darth Vader's lightsaber and move it to the tallest tower in the front row," Digit can do it.

But these same experiments show the limitations of AI models on their own, says Melonee Wise, Agility's chief product officer. As a simple example, she points to the unique jargon of logistics. If an automated pick-and-place system puts the wrong item into a tote, that tote is sent to the "hospital" for correction. If an item needs to get on a truck for shipping today, it's placed in a "hot" tote. "If we just threw all this data at the system, none of that would be contextualized," Wise says. "The minute a manager says, 'I have a load of hots coming in,' everything's going to break."

Agility Arc, the company's cloud-based automation management platform, is the place where customers can build in that kind of site-specific context. And it's also the level where machine learning becomes truly integral to Agility's system, Wise says, because Arc is a hybrid: a trainable expert system. Sensor data and other data flowing into the platform help Arc's AI models learn, and they can then make better recommendations to customers about how to set up their operations.

"At the lowest level, we can use reinforcement learning to do things like making the robot walk better," Wise says. "But if you want to improve the throughput of a facility, you need an actual expert system that has the context of the problem you're trying to solve. And the only way we get that context is, one, by measuring the real world, and, two, understanding why Digit did something in the real world."

Imagine, as a hypothetical scenario, that a Digit robot has been tasked with accepting totes from another robot and placing them on a stack. And imagine that every tenth tote, on average, is red, meaning it's a hot tote. "What we do with that, how we create an exception process to handle it, could be optimized" using the data Arc continuously collects, Wise says. "The expert system could make a recommendation that there should be a hot robot just to handle the red totes. Or maybe instead of placing that tote on a stack, Digit should put in on a special conveyor. Or the system could make a recommendation to do an A/B test where Digit does it for an hour one way and then an hour the other way."

Think of Arc as a learning system informed by the logic of a specific workplace—or as an expert system that's able to change its mind based on new data. "Because we're in the market today, we're starting to gather this kind of data already," Wise says.

To DCVC's Ocko, Agility is a prime example of a company being smart about the difference between data and models, and prioritizing the former. "There is a danger that people will believe that because AI can see beyond the horizon of immediate human perception, it magically solves hard limits in physics and engineering," he says. "There are robotics companies saying we can rig up automata with mechanical tendons and AI is going to bring them to life. But if you haven't done the engineering to make a robot actually work in human and industrial spaces—which Agility has—a large language model is not going to save you."

"If you want to improve the throughput of a facility, **you need an actual expert system that has the context of the problem you're trying to solve.** And the only way we get that context is, one, by measuring the real world, and, two, understanding why Digit did something in the real world."

Melonee Wise
Chief Product Officer
Agility Robotics

OPPORTUNITY 1.3

IT'S TIME TO START MAKING AI GREENER.

DCVC's core thesis is that new computational techniques continually power breakthroughs and present opportunities to solve hard problems in almost every field of innovation, from energy and climate to biotechnology and healthcare. But what powers computation? The literal answer is electricity—staggering quantities of it. Unless we find smarter ways to manage the burgeoning power requirements of data centers, including those supporting the newest machine-learning models in AI, deep-tech innovators could end up adding to humanity's net carbon footprint rather than shrinking it.

Data centers already eat up at least 1 percent of the global electricity supply. Training generative AI models takes so much power that the number could rise to 7.5 percent in the U.S. by 2030.

In 2020, data centers and data transmission networks were responsible for about 1.0 to 1.5 percent of global electricity use and 0.9 percent of global greenhouse gas emissions, according to the International Energy Agency. Given the enormous growth in global internet usage, streaming media, gaming, and blockchain applications over the last 15 years, these percentages could have risen much higher. But since 2010, the major data center operators have held power requirements roughly flat through energy-efficiency improvements and renewable energy purchases.

And then came AI's great breakthrough in the form of large transformer-based "foundation models" such as GPT-3, whose capabilities shocked the world in late 2022. Unfortunately, those improvements come with an equally shocking increase in electricity and cooling requirements, bringing an abrupt end to the era of flat power demand. By 2030, according to projections from Boston Consulting Group, data centers running generative AI models could be sucking up 7.5 percent of all electricity in the U.S.—as much power as is used by 40 million homes.

To understand how that could happen, we need to back up. Foundation models' seemingly magical ability to answer complex questions in natural language, analyze medical scans, predict how proteins will fold, and generate synthetic images and video is a product of the massive datasets they've been trained on, and they get smarter over time by ingesting ever more information. GPT-3 had 175 billion parameters in its model, a 100-fold leap over GPT-2. Google's Switch-C language model had 1.6 trillion parameters, and GPT-4 is rumored to have many more.

It's inefficient to train such large models using traditional central processing units (CPUs). AI-focused data centers use specialized graphics processing units (GPUs) and tensor processing units (TPUs), which have more processing cores and can therefore handle more matrix

Figure 1.3.1
Data centers have already implemented the "easy" fixes for energy waste.

Power use effectiveness (PUE) is the ratio of total energy used by a data center to the energy actually needed by the computing equipment. At data centers surveyed annually by the Uptime Institute, average PUE improved rapidly until 2014, but then leveled off—suggesting that all of the most cost-effective changes (upgrading to better cooling systems, for example) had already been performed.

multiplications—the foundation of deep-learning algorithms—in parallel. Unfortunately, because GPUs and TPUs perform so many more operations per second, they consume 10 to 15 times more electricity than CPUs, and have correspondingly large cooling requirements. Overall, U.S. data centers operated by Amazon, Google, Meta, Microsoft, and other "hyperscalers" needed 17 gigawatts of power in 2022, accounting for about 2 percent of electricity demand nationally. Newmark, a research firm that studies commercial property markets, says total data center energy demand is likely to rise to 35 gigawatts by 2030, and BCG says generative AI could push that number as high as 45 gigawatts.

Such estimates are driven in part by the assumption that AI models will continue to grow in scale, training time, and energy requirements. In reality, this assumption is hard to test. A team led by Alexandra Luccioni of the machine-learning startup Hugging Face estimated in 2022 that training GPT-3 required almost 1,300 megawatt-hours (MWh) of electricity—about the same as the annual energy consumption of 130 average U.S. homes. The equivalent figures for the latest AI models, such as GPT-4, Gemini, and Claude 3, aren't known, since OpenAI, Google, and Anthropic haven't released key details such as the number of parameters they use. But this January, Altman said at the World Economic Forum in Davos that future models will require even more energy than experts have imagined. "There's no way to get there without a breakthrough," Altman said. "We need fusion or we need radically cheaper solar plus storage or something at massive scale."

▸ **By 2030, according to projections from Boston Consulting Group, data centers running generative AI models could be sucking up 7.5 percent of all electricity in the U.S.—as much power as is used by 40 million homes.**

And in a world where nations are getting serious about decarbonization and electrifying everything (see section 2.2 of the 2023 Deep Tech Opportunities Report), AI will be competing for electricity with battery-powered vehicles, hydrogen production, and new industrial facilities. The situation demands innovation along two tracks. First, Altman is right: we must keep building clean new generating capacity, in the form of both variable, renewable sources like wind and solar power and 24/7 firm power from sources like nuclear, geothermal, and grid-scale storage (see Opportunity 3.1). Second, as part of any vision for the responsible, ethical rollout of new AI tools, we need a plan for limiting and eventually reducing the energy requirements of foundation models themselves.

Efficiency improvements could take many forms. Unfortunately, cloud computing providers have already implemented the easiest way to limit data center power consumption: moving most corporate computing from on-premises facilities to high-utilization cloud data centers mostly powered with renewable energy. One of the next goals should be to write new AI algorithms that reduce computational overhead. The AI companies' current strategy for improving foundation models is to keep feeding them larger and larger amounts of training data—including synthetic data, once all the real-world data has been exhausted. They should also invest more deeply in techniques that reduce their models' training requirements, such as transfer learning (using knowledge gained from one task or domain to perform better on another).

At the same time, we'd like to see semiconductor manufacturers work harder to optimize next-generation GPUs and TPUs for energy efficiency. The same goes for in-device chips, since we'll need to offload more AI inference tasks to edge devices such as laptops and phones. Here, a company in the DCVC portfolio called Mythic is one of the leading players. Its analog silicon architecture combines processing and memory on a single chip, making it thousands of times faster and more energy-efficient than traditional digital chips. The company's technology may ultimately be the only affordable option for running high-performance computer vision models or large language models on distributed devices such as smart-city cameras or smart-home appliances.

We'll also need new data center designs focused on making cooling systems more cost-effective, as well as ideas for offloading and systems for redistributing remaining data-center workloads to the locations with the cleanest energy at any given time of the day or the year. Amazon, Meta, Google, and Microsoft are already the four largest users of renewable power purchase agreements. Together, they've bought the rights to generation capacity equal to that of Sweden. The rest of the industry, including network and telecom operators, must follow suit. And the big cloud-scale computing providers themselves must find around-the-clock sources of carbon-free electricity, and stop relying on renewable energy certificates to offset their consumption of electricity generated from fossil fuels.

"Generative models in particular… need to become greener before they become more pervasive."

Thomas Davenport and Ajay Kumar
Information systems scholars

To date, companies have built larger and more powerful AI models without much consideration for their training costs or carbon footprints. Now it's time to work on making AI more inherently carbon-friendly. As information systems scholars Thomas Davenport and Ajay Kumar wrote in Harvard Business Review in 2023, "Generative models in particular…need to become greener before they become more pervasive." We'll look to all our portfolio companies, present and future, to do their part.

Figure 1.3.2 Foundation models and their increasing appetites

As the number of data points or "parameters" needed to train the latest generative AI models goes up exponentially, so does the amount of electricity needed to perform the required computations.

Source: Lisa Su/AMD

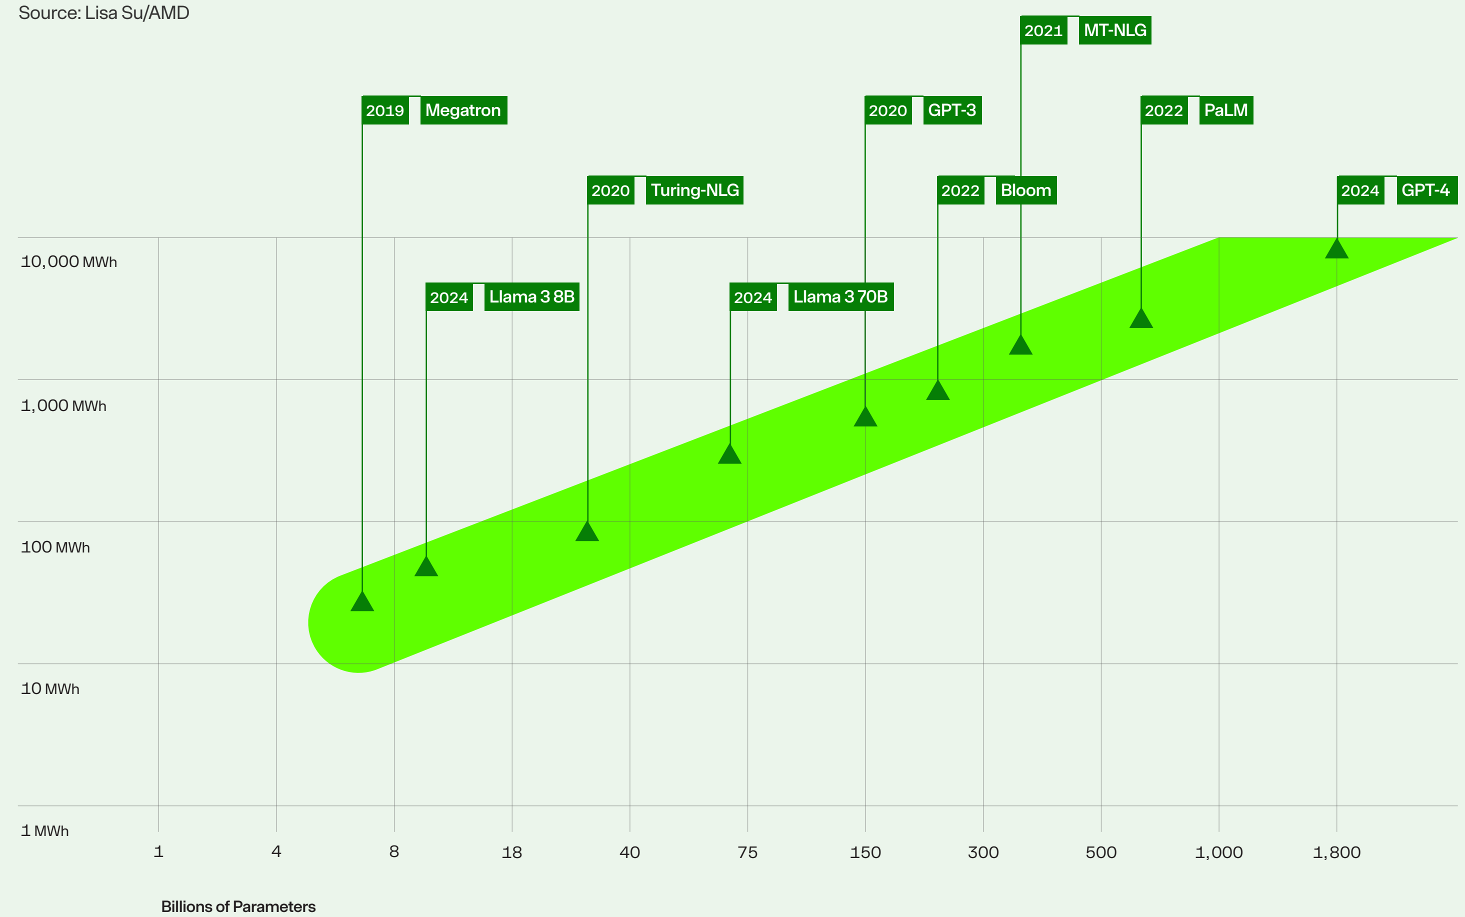

In the 2023 Deep Tech Opportunities Report, we spoke at length about technologies that are, in our view, "harder than people think," and that therefore may not be ready for investment from venture capital firms. (Like all private investors, we seek ideas that seem to have a direct path to commercialization.) After decades of research and engineering work, quantum computers—depending on what you're trying to do with them—are beginning to look like one of those technologies.

Though quantum computers are seen as successors to classical computers, irksome physics challenges have kept them from fulfilling that role. That said, we see compelling commercial opportunities in technologies that make today's glitchy quantum computers more usable, and in spinoff technologies that exploit the properties of individual qubits (quantum bits) to accomplish important tasks in new ways.

The problem, roughly, is that qubits don't behave the way we'd like, which means "recreating the entire computing stack around qubits has turned out to be a very, very difficult problem to solve," says James Hardiman, a general partner at DCVC. For one thing, qubits are exquisitely sensitive to quantum noise (random fluctuations arising from the fundamentally indeterminate state of matter) that can cause them to decohere. Whatever physical system is being used to embody qubits in a state of coherent superposition and keep multiple qubits entangled, it's likely to require elaborate and expensive controls. Even then, high error rates are a fact of life for quantum computing hardware builders.

We're excited about companies that are testing new ways of building qubits that cohere longer. Atom Computing, for example, encodes quantum information in the nuclear spin of electrically neutral atoms trapped by laser light. In late 2023 the company announced it had used ytterbium atoms to build a 1,225-qubit array, beating quantum giants IBM and Google to the more-than-1,000-qubit mark; the company has demonstrated that its nuclear-spin qubits can stay coherent for up to 40 seconds, which is a relative eternity in quantum terms. The error rate for Atom's system is still very high, but the company says it's working on measures that will increase fidelity.

At DCVC, our approach to quantum computing has been to invest in companies selling picks and shovels to businesses toiling to make quantum systems, rather than building the quantum computers themselves. (The exception is DCVC portfolio company Rigetti Computing, a pioneer in quantum computing that went public in 2021.) For example, we're ardent backers of Q-CTRL, a Sydney-based startup founded and led by quantum control researcher Michael Biercuk. The company initially specialized in quantum error suppression software, which pinpoints the sources of error in a quantum calculation and helps operators reduce that error by optimizing the control pulses used to manipulate qubits. IBM was so impressed by Q-CTRL's work that it decided to include the software in its Quantum Pay-As-You-Go Plan, which offers cloud-based access to its quantum processors. "There's a growing recognition in the industry that in order to get any kind of useful results from quantum computations, Q-CTRL's error suppression technology will be indispensable," says Hardiman.

Lately, though, Q-CTRL has also been surging ahead in another area: quantum sensing. Essentially, quantum sensors take the noise problem in quantum computing and turn it on its head. "The thinking is, if these things

are so sensitive, let's just make sensors out of them, so that the bug becomes a feature," Hardiman explains.

Q-CTRL starts with validated quantum sensor designs and adds its proprietary AI and quantum control algorithms to suppress noise. In this way it's developing prototype quantum-enabled inertial sensors that could provide position, navigation, and timing information in situations where GPS data is unavailable (say, due to a military conflict or a geomagnetic storm). Its ultra-sensitive gravimeters and magnetometers are also helping to create more detailed geophysical maps, which provide yet another backup form of navigation, as well as data about potential ore deposits or water tables.

We still believe in the promise of quantum computing. In fact, we see it as an absolute requirement for the future, since there are classes of problems that conventional computers simply can't solve given a reasonable amount of time or energy. Training future high-dimensional AI models with trillions or quadrillions of parameters may be one of these. "We need alternate architectures, or we are going to be stealing today's electricity in the service of tomorrow's fantasies," Ocko says.

But the path to an era of practical, useful quantum computing will remain rocky. We hope that fundamental efficiency improvements in conventional computing algorithms, together with the emergence of specialist AI chip architectures such as those being developed by Mythic (see Opportunity 1.3), can make that path smoother. And in the meantime, we continue to invest in adjacent, enabling technologies like quantum error control and quantum sensing that help innovators put the potential of quantum technology to practical use in the much nearer term.

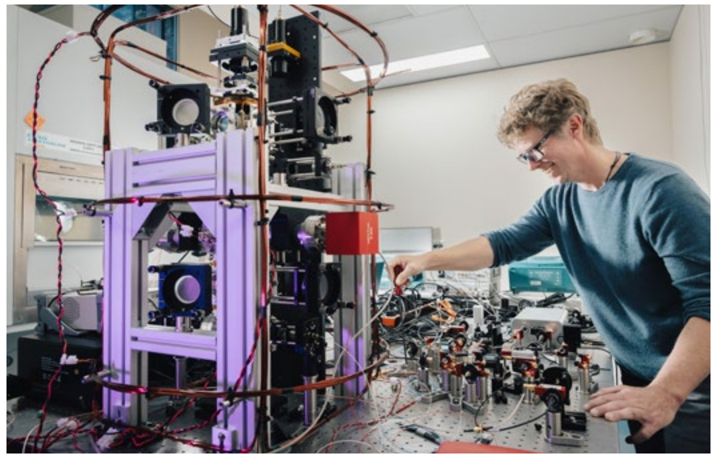

▸ Q-CRTL first developed software to suppress the effects of noise in quantum computations. Now it's applying its understanding of that noise to build ultra-sensitive instruments for navigation, gravimetry, and magnetometry.

Figure 1.4.1 Q-CTRL's quantum error correction cycle

The physical devices used for quantum computing are always subject to noise and errors, but Q-CTRL CEO Michael Biercuk has spent many years developing techniques for overcoming them. One is a cycle of quantum error correction in which "syndrome measurements" infer errors in individual qubits and apply corrections.

Source: Biercuk and Stace, "Quantum Error Correction: Time to Make It Work," IEEE Spectrum, 26 June 2022

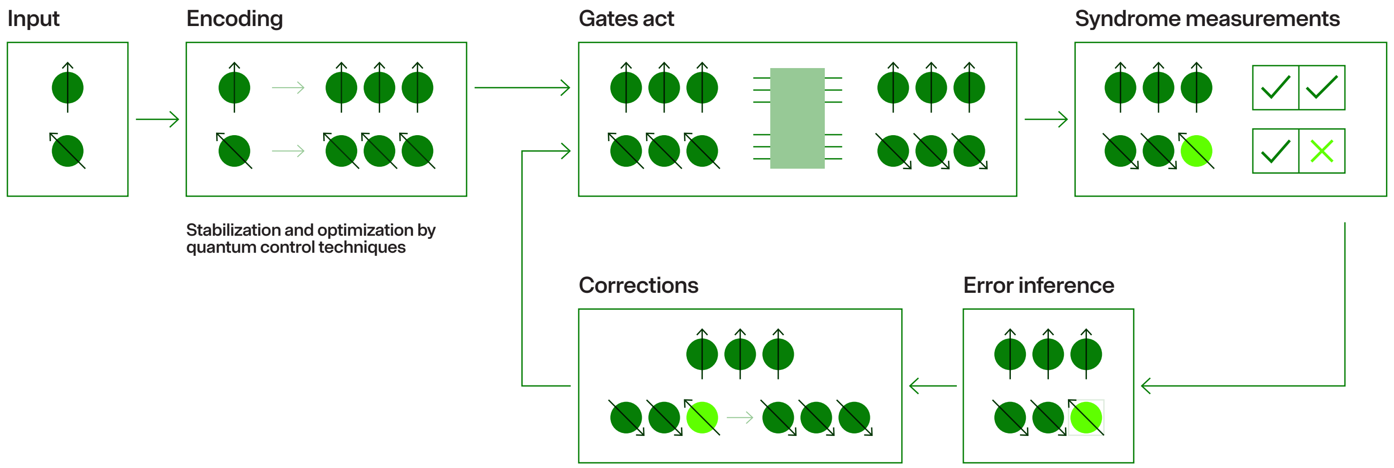

Occasionally in this report, we pause to consider shiny objects—technology ideas that offer understandable temptations to innovators and entrepreneurs, but that ultimately distract from more urgent and practical work. Perhaps the most discussed of these today is artificial general intelligence, or AGI. We don't see AGI as inherently impossible; but neither do we see it as imminent. A world with true AGI might be immeasurably better off—or worse. In any case, we think there are more immediate problems to be solved.

A world with true AGI might be immeasurably better off—or worse. In any case, we think there are more immediate problems to be solved.

↳ The prospect of a machine that genuinely thinks, or that at least causes humans to reconsider what we mean by "thinking," has been part of the teleology of computer science since the very beginning. Alan Turing's famous 1950 paper "Computing Machinery and Intelligence" argued that a computer capable of winning what he called the "imitation game," later known as the Turing Test, should be regarded as a thinking being. Turing's proposition was intriguing while it lasted, but the latest advances in AI have made it untenable. Large language models such as ChatGPT can convincingly simulate conversation—but no one would say that the underlying AI models display thinking or understanding. "They're blind watchmakers," says DCVC's Ocko.

(You can even ask ChatGPT about this. "The capacity to pass the Turing Test by generating responses that are indistinguishable from those a human might give does not equate to possessing understanding, consciousness, or sentience," the program told us in one interaction. Of course, a sentient AI *would* say that…)

Today's computer scientists sidestep intractable philosophical questions about the nature of mind by focusing on what AI can do in the world. "A system that is generally capable" is the definition of AGI offered by Demis Hassabis, leader of Google's DeepMind division. "Out of the box, it should be able to do pretty much any cognitive task that humans can do." Hassabis says he "would not be surprised if we saw systems nearing that kind of capability within the next decade or sooner." If building an AGI system is merely a matter of gluing together enough different cognitive skills to compete with a well-rounded human—completing a math problem, interpreting a visual scene, composing a sonnet or a melody—then it does seem that today's researchers are well on their way.

However, this kind of problem-solving composite would not think the way a human does, if only because it would lack our sense organs, our emotion-racked nervous systems, and our networks of social relationships. "The key to a scientific theory of our intelligence lies in acknowledging the fact that humans are embodied, which is to say that we are living, biological creatures who are in constant interaction with the material, social, cultural and technological environment," writes Anthony Chemero, a professor of philosophy and psychology at the University of Cincinnati who's been studying the idea of "embodied cognitive science" for more than a decade. "This tight connection between experiencing and acting is a central feature of human intelligence and is something that LLMs lack entirely." (The idea of embodied computing comes with its own practical and philosophical problems, of course.)

Let's say Hassabis's forecast is right, and that Google or some other entity can, in the next decade, build a program that surpasses the threshold for general intelligence (whatever that might be). That future system might be more versatile than today's Gemini or ChatGPT, but it would not be more *interesting*, because it would be just as dead inside. Machine understanding, consciousness, sentience—all of these will likely require a fundamentally different approach to computing, if they can be achieved at all.

What we do know is that we face another, far more urgent task: making *today's* AI systems more accountable. AI safety is a fast-growing field of inquiry and policymaking, and of course we agree with its basic goals: protecting personal privacy and data security, eliminating algorithmic bias, preventing AI-assisted fraud, and the like. What worries us right now is something subtler: the possibility that AI models will be given responsibility for myriad real-world decisions in the absence of robust methods and mechanisms for *a*) understanding and explaining those decisions, and *b*) allowing humans to challenge and reverse them.

Long before the current AI explosion, organizations were using computer models to reduce the human element in business and legal decision-making—for example, by matching job candidates with openings, judging the risk of recidivism in sentencing and parole proceedings, or evaluating which borrowers are most likely to default on a home mortgage. Now there's a risk that AI could fully take over many professions. A 2023 study from OpenAI and the University of Pennsylvania found that 19 percent of U.S. workers are in medium-skill occupations where at least half of the tasks could be done by GPT-4 or similar large language models. Moreover, the study concluded that hundreds of occupations could be completely automated, including tax preparers, accountants, court reporters, legal secretaries, clinical data managers, mathematicians, and journalists. How quickly businesses would actually choose to automate is uncertain, but the threat is no longer theoretical; between May 2023 and February 2024, U.S. companies cited AI as the reason for 4,628 job cuts, according to Challenger, Gray & Christmas. (The firm specializes in helping "exiting employees" with their career transitions, which is certainly one business that seems poised to grow.)

On top of the heartbreak of unemployment, the rush to automate skilled jobs could have another chilling consequence. Moments of humanity and everyday mercy—the insurance claims adjuster who bends the rules, the traffic cop who waives a speeding ticket—are part of what make our interactions with bureaucracies tolerable. We fear a world where small decisions about our lives are made by a web of hundreds of invisible AI systems built or hosted by giant technology companies, producing possibly unfair or even hateful effects (depending on the biases inherent in their training data), with no practical means of appeal.

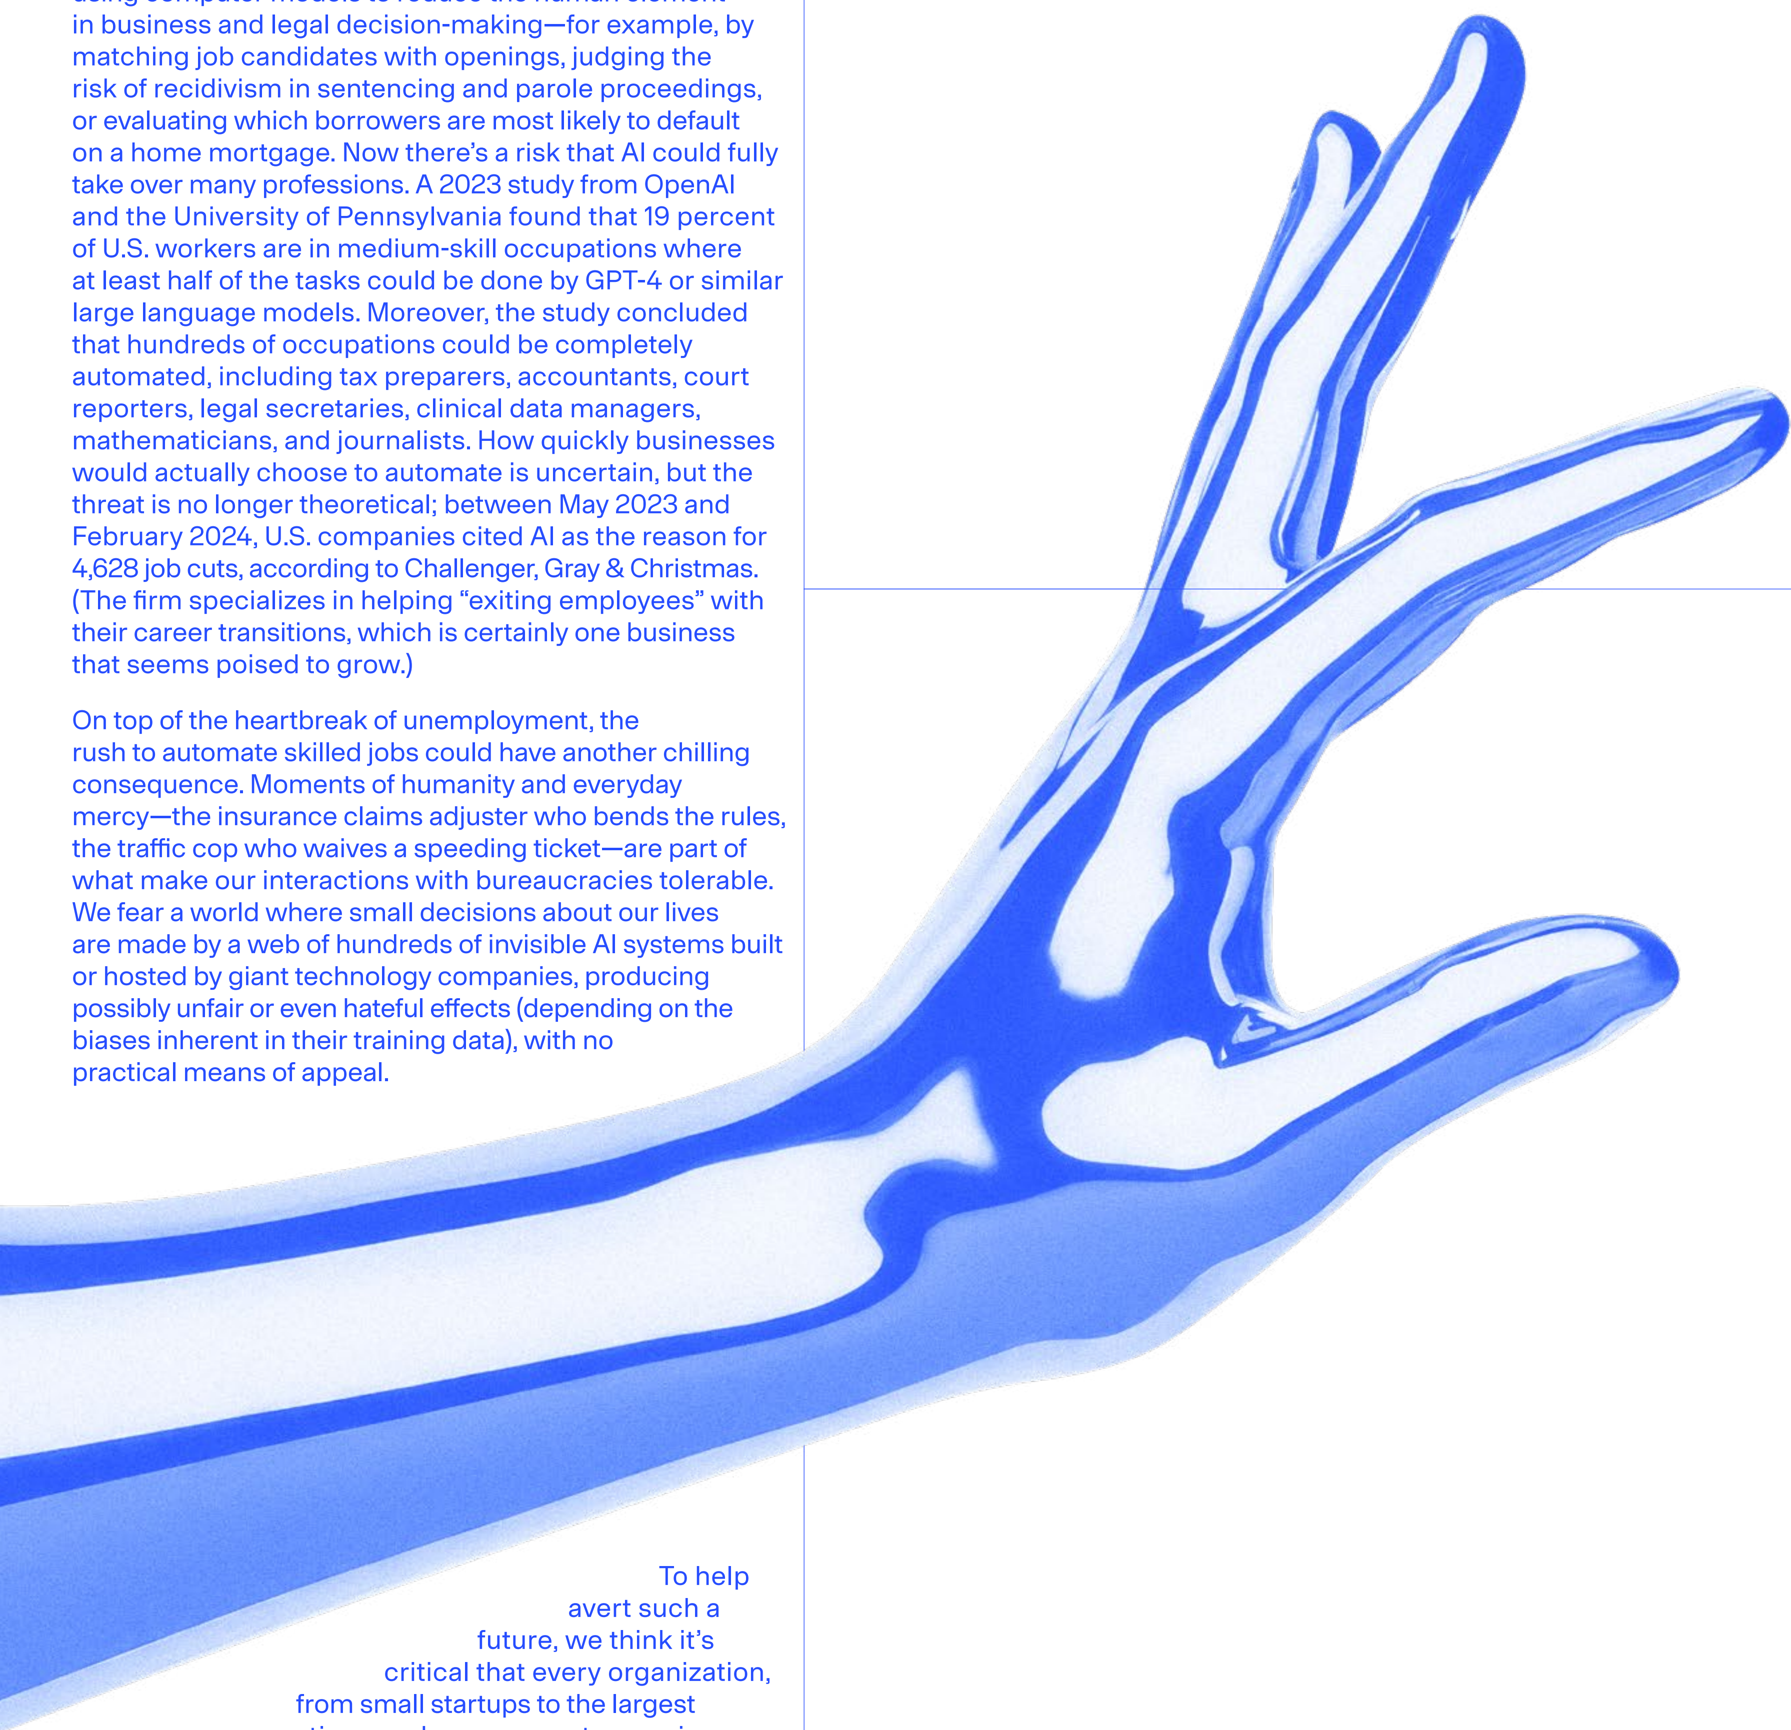

To help avert such a future, we think it's critical that every organization, from small startups to the largest corporations and government agencies, have the ability to build and run the machine-learning models and algorithms it needs for its operations, rather than ceding control to off-the-shelf models from the giant tech companies. Technology like that from DCVC portfolio company MosaicML, which was acquired in 2023 by another DCVC-backed company, Databricks, can help here. Databricks offers products that help developers deploy custom generative AI models quickly and easily, in their companies' own secure environments, and at a fraction of the cost of other comparable services.

"Forget AI doomerism; AGI is not the threat … What is the threat is a vast assortment of black-box, unappealable little AI gods that codify the vindictive, opaque policies of the gas company, the cable company, the parking enforcement division, the other oligopolies that we all endure. The ability to cost-effectively validate those models and provide the tools to call them to account—*that* is essential for the survival of human civilization."

Matt Ocko
Managing Partner, DCVC

We'd also like to see companies building and using AI sign on to a set of guidelines such as the "Blueprint for an AI Bill of Rights" proposed by the White House Office of Science and Technology Policy. Among the principles proposed in the OSTP document are "You should know how and why an outcome impacting you was determined by an automated system" and "You should have access to timely human consideration and remedy by a fallback and escalation process if an automated system fails, it produces an error, or you would like to appeal or contest its impact on you."

The OSTP proposes that these ideas be applied to content-moderation tools, criminal justice algorithms, voting systems, school admissions algorithms, housing-related systems, health and health-insurance technologies, and loan-allocation and credit-scoring systems. But we would go even farther. In principle these protections, and many others, could and should be hard-coded into any AI system that mediates access to opportunities, resources, or services.

In the end, making AI systems more interpretable, explainable, and reversible isn't just good social policy; it's good engineering practice that will guide the development of more effective AI models in the future. "Forget AI doomerism; AGI is not the threat," Ocko says. "What is the threat is a vast assortment of black-box, unappealable little AI gods that codify the vindictive, opaque policies of the gas company, the cable company, the parking enforcement division, the other oligopolies that we all endure. The ability to cost-effectively validate those models and provide the tools to call them to account—*that* is essential for the survival of human civilization."

Figure 1.5.1 **AI exposure by occupation**

Occupations that humans classified as having the highest exposure to AI and complementary technologies	% Exposure
Survey researchers	84.4
Writers and authors	82.5
Interpreters and translators	82.4
Public relations specialists	80.6
Animal scientists	77.8

* Under a broader definition, humans labeled 15 occupations as 100% exposed

Occupations that GPT-4 itself classified as having the highest level of exposure	% Exposure
Mathematicians	100.0
Blockchain engineers	97.1
Court reporters and simultaneous captioners	96.4
Proofreaders and copy markers	95.5
Correspondence clerks	95.2

* Under a broader definition, the AI model labeled 86 occupations as 100% exposed

Source: T. Eloundou et al., "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," working paper (2023)

Chapter 2.0

TechBio is emerging as the data-centric, AI-native counterweight to biotech.

● Companies

Freenome, Kanvas Biosciences, Noetik, Recursion Pharmaceuticals, Relation Therapeutics

● Voices

Zachary Bogue, James Hardiman, Jason Pontin

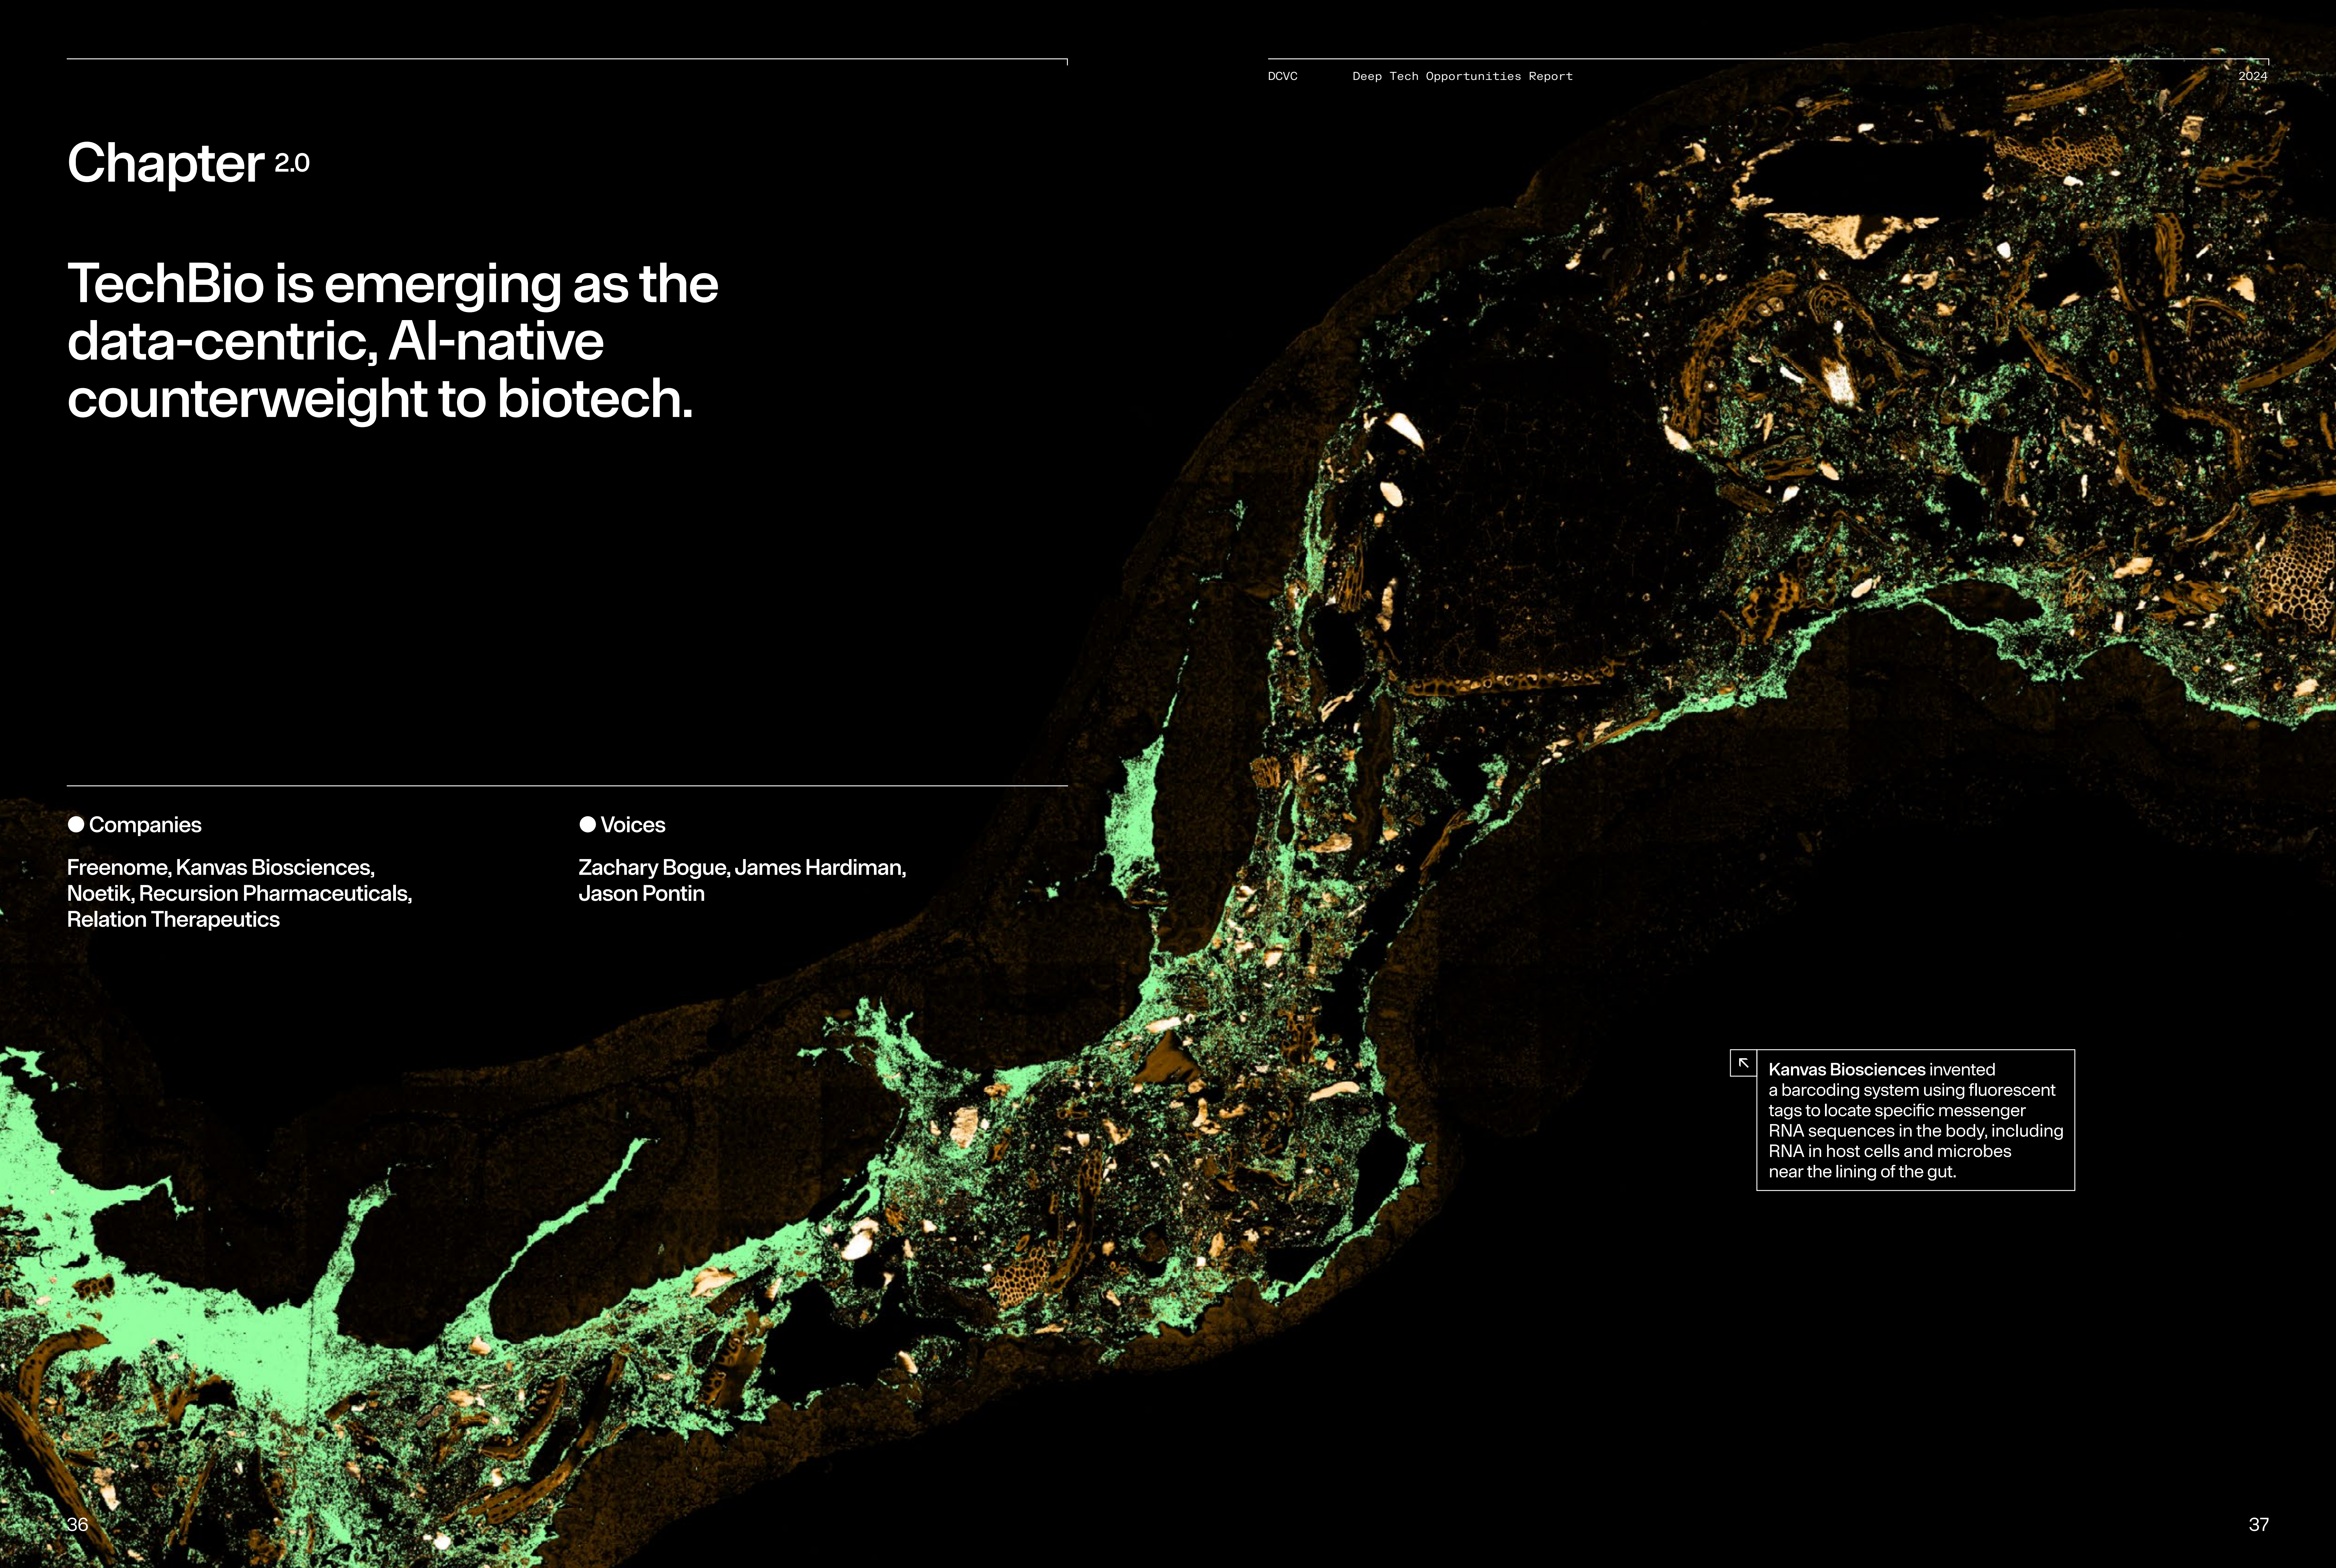

↖ **Kanvas Biosciences** invented a barcoding system using fluorescent tags to locate specific messenger RNA sequences in the body, including RNA in host cells and microbes near the lining of the gut.

TECHBIO IS EMERGING AS THE DATA-CENTRIC, AI-NATIVE COUNTERWEIGHT TO BIOTECH.

The ascension of data we examined in the previous section is overturning drug discovery and therapeutics. Historically, the biggest roadblock for biotechnology startups seeking new ways to interrupt disease processes has been the complexity of human biology itself. In other words: biopharma's hit rate is low because the number of potential starting points in the hunt for new drugs, and the skein of potential protein-protein interactions that need to be untangled, are both so huge. But companies are learning how to generate high-quality data on biological systems at industrial scale, then use that data for predictive modeling, and test those predictions in real cells. They can then cut through that barrier, narrow the search space, and drastically speed up the discovery phase of drug development.

We use TechBio, an inversion of biotech, as a name for the new branch of the AI industry that constructs proprietary biological and chemical datasets to fuel advanced computation and raise the hit rate in drug discovery. "We invest in AI-driven innovations that offer practical, real-world solutions, and that show the potential for substantial impact and scalability," says DCVC managing partner Zachary Bogue. "We see TechBio as one of the major opportunity areas for AI applications."

(Of course, there is no bright line between TechBio and biotech. Neither approach is inherently superior; that's why we range across the two. This section focuses on companies backed by DCVC's flagship funds. To learn about the broad scope of our biotech investing activity through DCVC Bio's funds, see Chapter 7.)

TechBio companies still do plenty of wet-lab work, but their fundamental business proposition is different from that of their biotech elders. TechBio companies aren't built to pick apart the molecular pathways behind specific forms of neurodegeneration, cancer, or arthritis and engineer potential cures—a high-risk strategy that fails far more often than it succeeds. Rather, they're built to collect the raw data that will reveal, through *in silico* modeling, which hypotheses and drug candidates are most worth testing. These tests then generate even more data that can be used to make even better predictions.

The end result: more leads at lower cost and lower risk. We think Jacob Kimmel, co-founder and head of research at the epigenetics startup NewLimit, hit the mark in a late-2023 essay when he called the construction of these datasets "a virtuous flywheel that can take off at the heart of these businesses."

We see DCVC portfolio company Relation Therapeutics, mentioned in the previous section, as a signal example of the TechBio sensibility. The company chose osteoporosis as its first target and built an "osteomics" platform to generate the world's largest atlas of genomic and transcriptomic data on bone cells from human patients. It applies machine-learning models to that data to predict which genes are most likely to contribute to the disease. It then knocks out those genes in other bone cells and measures the effects on bone mineralization. That assay data, finally, points the company toward possible disease mechanisms and suggests targets for intervention.

Given the right effector cells for disease modeling, the data-centric approach should work against many complex diseases, including neurodegenerative diseases, says Benjamin Swerner, chief operating officer of the London-based company. "The engine is that we're trying to identify which bits of human biology to prosecute and why," Swerner explains. With that *why* in hand, Relation can in theory avoid the classic biotech pitfall of spending tens of millions of dollars to get a single drug candidate into clinical trials, only to find out that it doesn't work. This March, NVentures, the venture arm of the giant chipmaker Nvidia, co-invested with DCVC in a new financing round for Relation, calling the company a leader in harnessing generative AI for drug discovery.

Another DCVC-backed company, Recursion Pharmaceuticals, in Salt Lake City, Utah, started doing TechBio long before there was a common term for it (the company was founded in 2013, and we first invested in 2016). Recursion's focus is on generating vast amounts of biological and chemical data—more than 50 petabytes and counting—including phenomics data showing how toxins, pathogens, genetic changes, and candidate drug molecules affect the morphology of individual cells. When the company's statistical analyses spot novel or unexpected effects or relationships, they can quickly design follow-up experiments to see how they relate to diseases of interest. The company calls the process a "virtuous cycle of atoms and bits." "They have one of the biggest datasets out there, which puts them at the vanguard of the data race in biopharma," says Bogue.

Recursion went public in 2021 and already has seven drug candidates in the preclinical or clinical testing phases, all targeting diseases with high unmet need and a lack of approved or effective therapies, such as ovarian tumors and neurofibromatosis. Last year the company acquired the Canadian startups Cyclica and Valence, makers of deep-learning engines for designing small-molecule drugs and predicting their effects; at the same time, Nvidia invested $50 million to help Recursion speed the training of its AI models. And in May, the company completed BioHive-2, the 35th-most powerful supercomputer in the world (according to TOP500.org) and the most powerful cluster wholly owned and operated by a pharmaceutical company. It comprises 504 separate Nvidia GPUs, and will be used to train Recursion's large AI models to predict how different drug candidates will affect the human body.

> "Because of advances like DNA and RNA sequencing and machine vision there's far more data coming from the lab. Now we have the capacity to understand that data, thanks to machine-learning models."

Jason Pontin
General Partner, DCVC

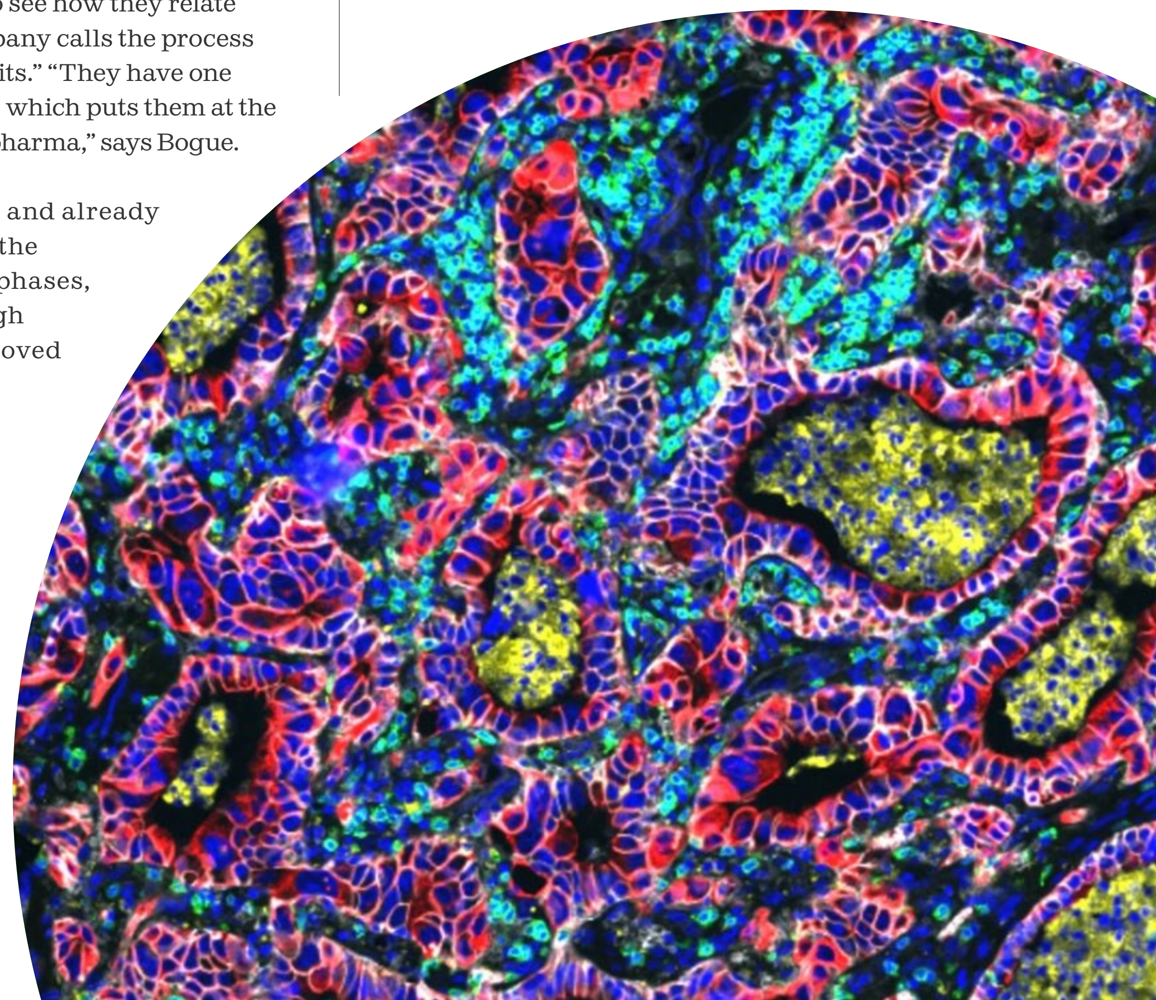

Noetik
By training machine-learning models on thousands of images of tumor slices, each stained to highlight the locations of different RNA molecules and proteins, the company can identify cancers that respond to specific immunotherapy drugs.

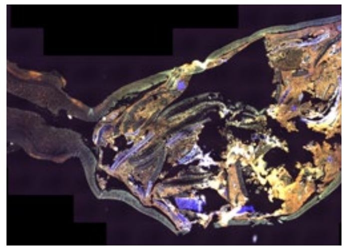

Kanvas Biosciences
A barcoding system for RNA in microbial and host cells could help explain how the gut microbiome mediates cancer immunotherapy.

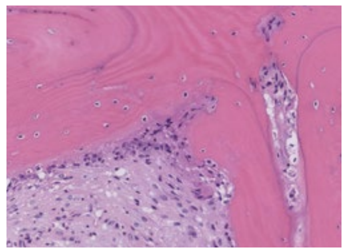

Relation Therapeutics
Culturing and analyzing human bone cells with various genetic modifications could lead to new drug targets for osteoporosis and other diseases.

Recursion Pharmaceuticals
Building a massive database chronicling the effect of toxins and other compounds on cell morphology could unlock new drug programs.

"Both Recursion and Relation are positioned to exploit the remarkable convergence we're seeing in biology," says DCVC's Pontin. "Because of advances like DNA and RNA sequencing and machine vision there's far more data coming from the lab. Now we have the capacity to understand that data, thanks to machine-learning models running on advanced hardware. Scientists working with machines can find medicines they would never have discovered on their own, cutting years from the drug development process."

Three other TechBio companies in our portfolio are marrying machine learning and biology in their own specialized corners of the biomedical world. Noetik, founded by two former Recursion scientists, studies the spatial locations of RNA and proteins in human tumor cells and feeds the data into AI models in an effort to identify cancer subtypes that could be treated with new immunotherapy drugs. Freenome is developing tests for early detection of tumors, using machine learning to sort through the complex mix of cell-free DNA, RNA, and proteins in the bloodstream and find patterns that may indicate cancer. Then there's Kanvas Biosciences, which has developed a multiplexed, AI-driven spectral barcoding system that provides massive amounts of data about messenger RNA molecules expressed by specific cells in the body, including microbes in the gut.

In its initial foray into drug development, Kanvas is using its spatial biology platform to give a boost to existing immunotherapies for cancer. Immune checkpoint inhibitors (ICIs)—which unleash the body's T cells to kill cancer cells—are highly effective, but fewer than 40 percent of patients respond to the treatment. Mysteriously, the response seems to depend in part on signals from the gut microbiome, and it's known that a fecal microbiota transplant (FMT) from a responding patient can convert a nonresponder into a responder. But fecal transplants can't be scaled up to treat thousands of patients, so Kanvas, working with the University of Texas MD Anderson Cancer Center, identified a handful of "superdonor" patients whose tumors melted away in response to ICIs and repackaged their gut microbes in the form of a precision microbiome therapeutic, delivered orally. Kanvas is using its spatial biology platform to detect precisely which superdonor microbes interact with gut epithelial cells, so it can optimize the microbial mix in future products. Without that platform, the company wouldn't have the data it needs to formulate effective drugs, says Kanvas co-founder and CEO Matt Cheng. "Now we think we will have the first precision microbiome therapeutic that's recognized by the FDA as an actual drug, as opposed to some version of an FMT," he says.

DCVC's Crossan predicts that the combination of AI and lab data at companies like Relation, Recursion, Noetik, Freenome, and Kanvas will have the same effect on drug discovery that DeepMind's AlphaFold program had on protein structure prediction. "Up to 2020, the whole of science had about 200,000 protein structures available based on X-ray crystallography, representing the work of about 50 to 60 years," Crossan says. "Today we have hundreds of millions of protein structures, thanks to AI and computational methods like AlphaFold"—which Crossan helped to pioneer at DeepMind—"but also other projects as well, like RoseTTAFold and ESMFold. That's something like a 5,000-fold change in productivity. And that same order of magnitude change is coming right across the sciences, in terms of the number of hypotheses we can explore and the number of experiments that we can do virtually. I often say to people who work in biology or chemistry, 'What problems would you think about if you could do a thousand times as many experiments as you can do now?' Because that's the world we're entering into."

Figure 2.1.1 Recursion data growth

Since 2018, Recursion has collected 50 petabytes of chemical and biological data to train its machine learning models.

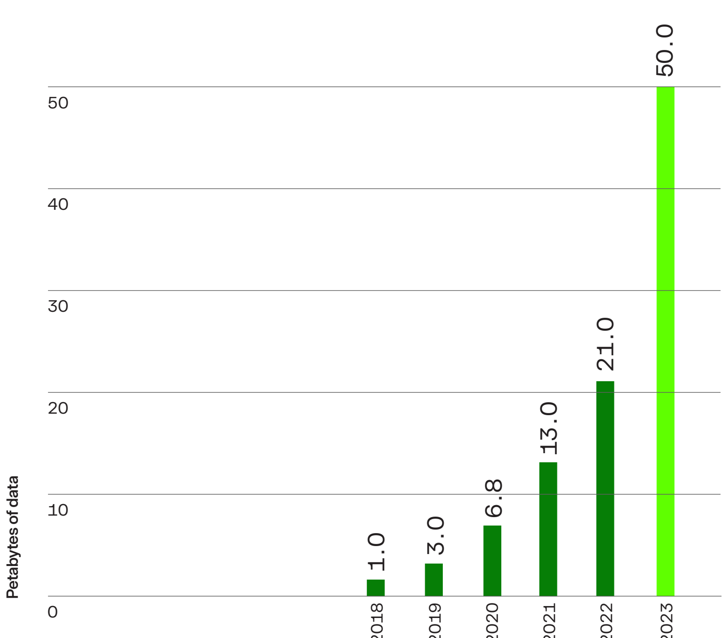

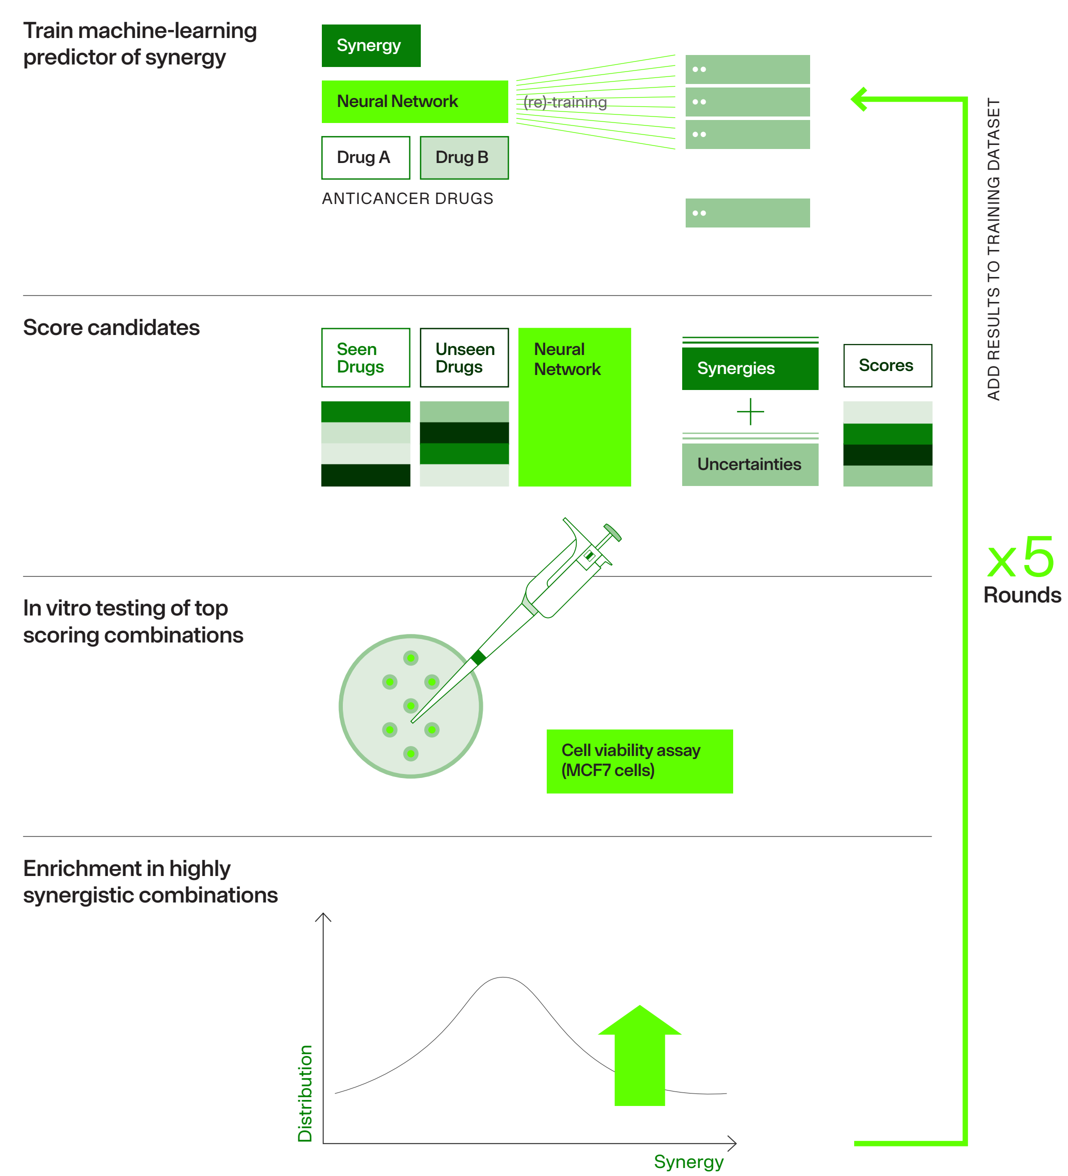

Figure 2.1.2 Relation

Under the lab-in-the-loop paradigm, Relation sequences the DNA and RNA of human bone cells and feeds this data into machine-learning models to identify gene variants that may be associated with disease risk. It then creates more cells with these genes knocked out to study the effects on bone mineralization, a marker for osteoporosis. That data, in turn, feeds back into the next round of experiments.

Chapter 3.0

To protect the planet, we must rethink energy systems from the ground up.

In our energy and climate-tech investing, we think about generating clean power, supplementing intermittent sources, and finding zero-carbon replacements for polluting fuels and materials.

Fervo Energy uses a combination of subsurface imaging and horizontal drilling to reach previously inaccessible sources of geothermal heat. ↗

● Companies

Brimstone, CH4 Global, Fervo Energy, Fourth Power, Oklo, Pivot Bio, Radiant Nuclear, Twelve

● Voices

Zachary Bogue, Matt Ocko, Dr. Rachel Slaybaugh

OPPORTUNITY 3.1

NEW STORAGE AND GENERATION TECHNOLOGIES WILL MAKE THE ELECTRICAL GRID CLEANER AND MORE RESILIENT.

In the United States, about 22 percent of the dispatchable, or "firm," power feeding into the electrical grid comes from nuclear plants, while most of the rest comes from coal- or gas-fired facilities. To meet growing electricity needs, increase the supply of carbon-free firm power, and clear the way for faster decommissioning of fossil-fuel plants, the grid will need more energy from deep within the Earth itself, in the form of enhanced geothermal energy, as well as a profusion of small, next-generation nuclear fission plants. At the same time, there's an acute need for technologies that can absorb and store electrical power during peak hours of solar and wind production and feed it back into the grid when needed.

Building carbon-free renewable energy facilities is so cheap nowadays that it has become the default way to expand generation capacity. Wind and solar farms account for about 80 percent of the megawatts of capacity added to the U.S. grid since 2016. But one thing solar and wind can't deliver is constancy: they don't work when the sun isn't shining or the wind isn't blowing. To make the most efficient use of existing transmission capacity—and to compensate for extreme weather events, which are growing in frequency—nations will always need "clean firm" power, meaning carbon-free electricity that's available on demand when solar and wind aren't. Here at DCVC, we're investing in several forms of clean firm power, as well as long-term energy storage to round out the bumps in power supplies.

First, there's next-generation nuclear power, which we covered extensively in 2023's Deep Tech Opportunities Report. In theory, it should be possible for the nuclear industry to converge on a standardized design for light-water reactors, like the Vogtle Unit 3 and 4 reactors that recently entered commercial operation in Georgia, and make new gigawatt-scale plants more economical. But in practice, the Vogtle project's huge cost overruns and schedule delays (the units came online seven years late and $17 billion over budget) will likely scare away most utilities. We think microreactors, on the scale of 1–20 megawatts of thermal energy each, are more likely to achieve economic viability, and so we've been backing microreactor pioneer Oklo since 2018.

The company is developing an advanced fast neutron reactor that runs on waste fuel from conventional reactors. It's an updated version of an experimental breeder reactor that operated safely for 30 years at what's now Idaho National Laboratory. In May, the company went public through a merger with AltC, a special purpose acquisition company co-founded by OpenAI CEO Sam Altman, raising more than $300 million in new capital. Oklo plans to build plants for about $70 million per project and sell power directly to individual customers who need clean firm power, such as data centers and factories.

Radiant Nuclear, another DCVC-backed company, is developing an even smaller microreactor, one designed to fit on the back of a semitrailer and generate about 1.2 megawatts of electricity—enough to replace the polluting diesel generators used by expeditionary military forces or as backup power in remote village hospitals. Radiant's Kaleidos reactor is powered by a fuel called TRISO (for TRI-structural ISO-tropic particle) that's made of poppyseed-sized uranium-carbon-oxygen kernels encased in three layers of carbon- and ceramic-based materials, making it meltdown-proof. And it uses supercritical CO_2 for power conversion, allowing for smaller and more efficient turbines that don't need access to water. "I think this will ultimately be a game changer," says DCVC partner Dr. Rachel Slaybaugh. "We'll be able to bring reliable zero-emissions power to remote locations and emergency situations, creating tremendous flexibility that's never been possible before."

Radiant Nuclear's Kaleidos fission reactor could replace diesel generators at hospitals, data centers, and military installations.

Then there's a new wave of innovative projects that extract naturally occurring heat from the depths of the earth and use it to generate electricity. There are a variety of approaches to geothermal energy, and we've placed our bet on Fervo Energy, which in 2023 became the first company to demonstrate commercial-scale flow rates for a pair of wells—one for injecting water and another for extracting it. Fervo has adapted hydraulic fracturing, horizontal drilling, and subsurface-imaging techniques pioneered by the oil and gas industry for geothermal drilling, unlocking previously inaccessible heat sources. (In its demonstration, the company drilled 8,000 feet down, turned sideways, and drilled 4,000 feet over.) Water pumped at high pressure through the injection well emerges into the hot surrounding rock across multiple zones and is sucked back up into a production well, which returns it to the surface to power a 3.5-megawatt steam generator, a new record for an enhanced geothermal system, or EGS.

Most excitingly, Fervo is on a rapid learning curve, completing each new well faster than the one before. Scaling up its generating capacity 100-fold won't be a matter of inventing new technology, but will demand drilling 100 identical wells, with significant speed and cost savings expected along the way as the company gathers more data about the best way to drill and to tap different geothermal reservoirs. Already Fervo has broken ground on a project in Utah that will deliver 400 megawatts of power when it reaches full capacity in 2028. Fervo CEO Tim Latimer believes that EGS could eventually produce several hundred *gigawatts* of power, enough to cover 20 to 30 percent of U.S. electrical demand.

There's one more technology that could help supply clean firm power—which will be necessary even after we've saturated the country with wind and solar farms. It's storage: both short-term (0–4 hours), where lithium-ion and other types of chemical batteries suffice, and long-term (5–100 hours), where other technologies like

> "All these companies are dealing in different ways with climate-forcing functions and the expansion of variability on the grid, and all in complementary ways."

Dr. Rachel Slaybaugh
Partner, DCVC

Figure 3.1.1 Tapping the Earth's heat

Fervo Energy uses precision directional drilling technology to access hard-to-reach geologies. Water from the company's injection wells reaches superhot temperatures as it travels through fractures in deep rock. A production well brings it back to the surface, where it powers electricity-generating turbines.

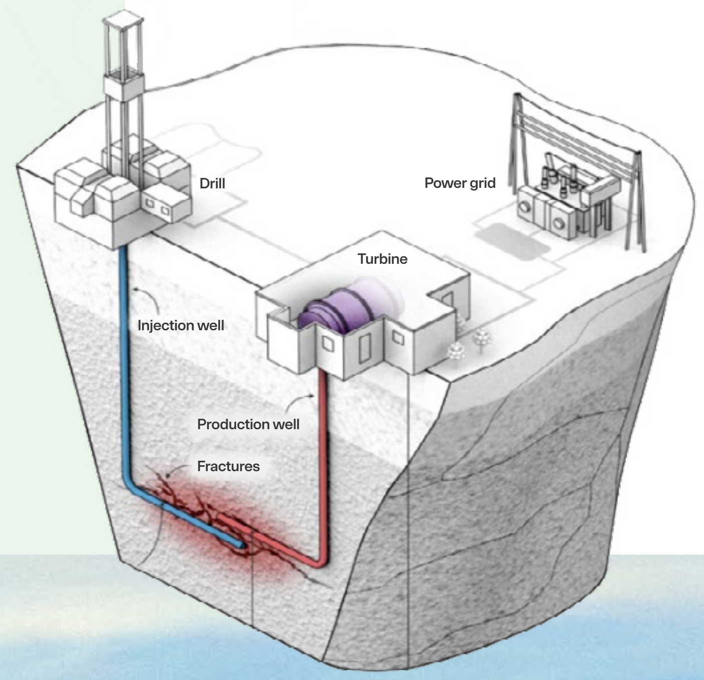

Note: This is a general representation of Fervo's process and doesn't represent the actual positioning or drilling of wells. Diagram not to scale.

▸ **Nearly unlimited power**
There's so much thermal energy in rocks two to six miles below the surface that capturing just 2 percent of it would satisfy U.S. power needs 2,000 times over.

thermal storage will be needed. Last year we helped finance a Massachusetts company called Fourth Power, founded by MIT mechanical engineering professor and heat-transfer expert Asegun Henry. The company's pioneering technology can draw electricity from the grid during periods of peak production, transform it into heat, and store it in large graphite blocks. When the grid needs more power, the heat is turned back into electricity using thermophotovoltaic cells. The life blood of the whole system is a unique heat-circulation system using molten tin propelled by all-graphite pumps.

"They're storing heat at 2,200°C, and that makes a bunch of things possible," says Slaybaugh. "Radiative heat transfer from the graphite to the thermophotovoltaic cells goes as a function of temperature to the fourth power—hence the company's name—so temperature makes a really big difference in terms of how much power you can move. And because it's so hot, you can also address any temperature people might need for industrial processes." In other words, a Fourth Power facility could also supply heat for hard-to-decarbonize industries such as steelmaking or cement making. We believe Fourth Power's unique long-duration storage capability, coupled with the ability to return stored power as electricity, is a game-changer for the grid.

In all of the energy projects DCVC backs, we focus on technologies that provide multiple advantages. They could be more versatile, powering many applications at once. They could be smaller or less costly, facilitating faster and wider deployment. They could be located closer to where they're needed; less likely to disrupt communities or provoke NIMBYism; or more distributed and therefore more resilient against outages and failures. "All these companies are dealing in different ways with climate-forcing functions and the expansion of variability on the grid, and all in complementary ways," says Slaybaugh. "Which is good, because these are issues we're going to be dealing with for a long time."

"I think Fervo and Oklo are well-positioned to be big stories in the coming year," says DCVC's Bogue. "Climate change is the gravest existential threat we face, and nuclear and geothermal energy have enormous proven potential to make an immediate-term impact."

Figure 3.1.2 U.S. electricity generation by major energy source, 1950–2021

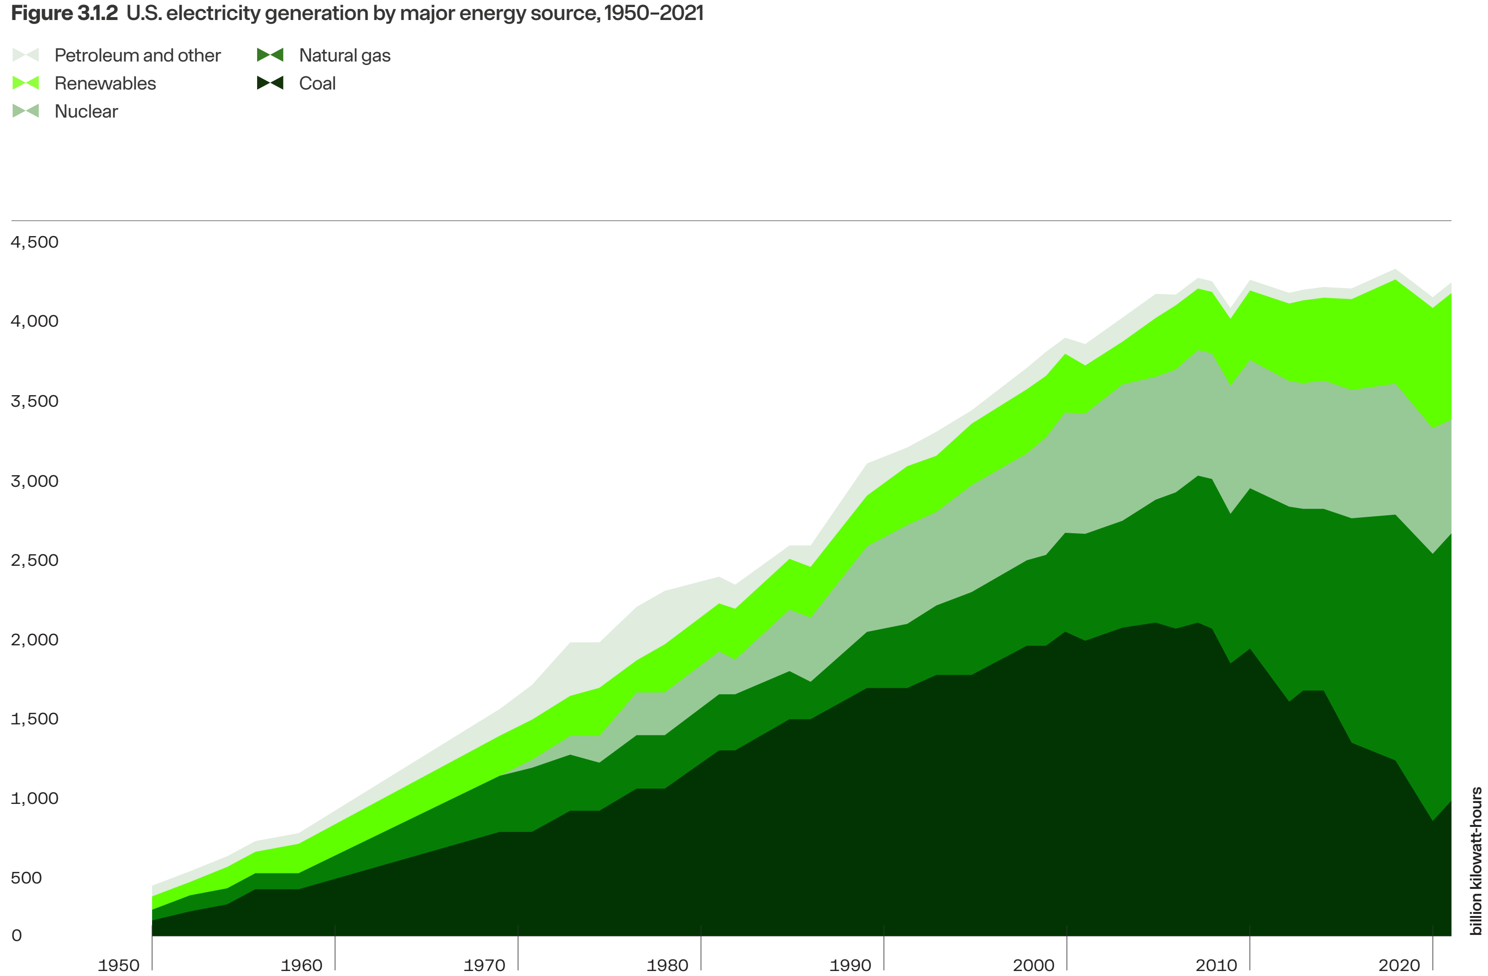

> "Climate change is the gravest existential threat we face, and nuclear and geothermal energy have enormous proven potential to make an immediate-term impact."

Zachary Bogue
Managing Partner, DCVC

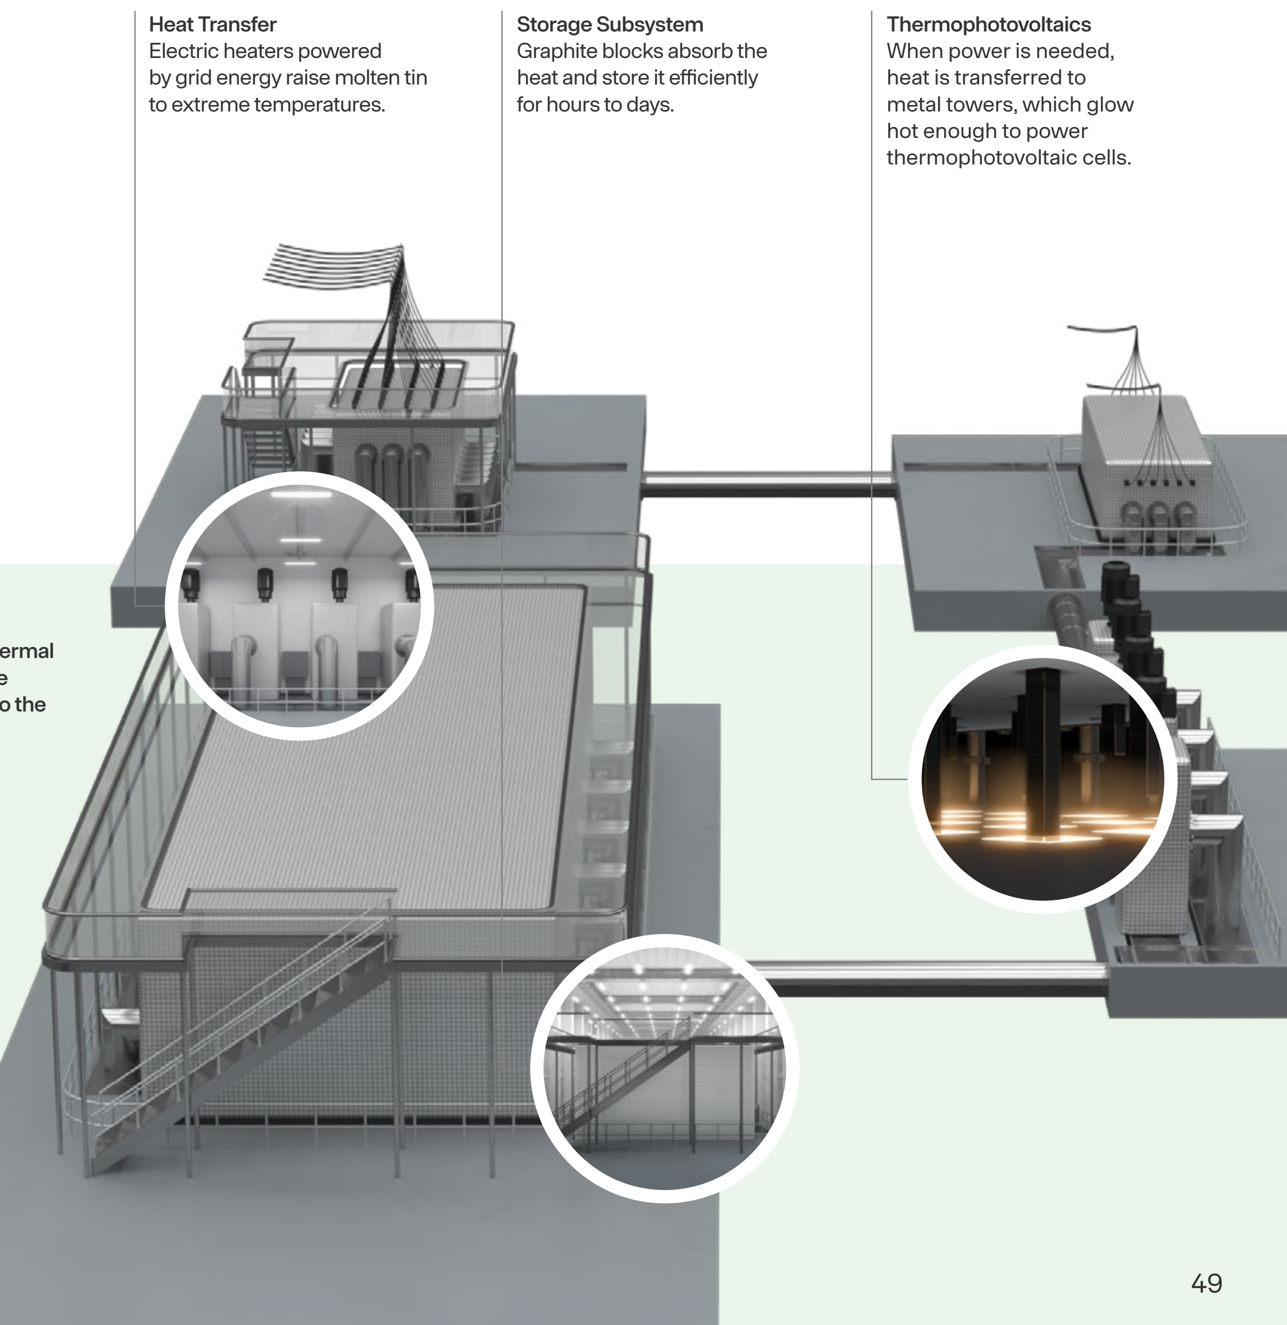

Figure 3.1.3
Fourth Power stores thermal energy in large graphite blocks and releases it to the grid when needed.

OPPORTUNITY 3.2

DROP-IN REPLACEMENTS ARE A WAY TO START DECARBONIZING THE ECONOMY, NOW.

Building all-electric vehicle fleets and adding wind, solar, geothermal, nuclear, and storage capacity to the grid are the easy ways to lower carbon emissions. It will be much harder, and will likely take much longer, to decarbonize high-emissions industries such as aviation, iron and steel production, cement making, fertilizer production and use, and livestock agriculture. (Globally, these five sectors were responsible for about 21.4 percent of greenhouse gas emissions in 2020, according to the World Resources Institute.) But there is a way to shrink the carbon footprints of these industries in the near term. It's to use new low-carbon or zero-carbon chemistries as substitutes for traditional feedstocks and fuels.

Certain realities seem inescapable. Planes burn fossil-derived jet fuel. Cows burp methane, and bacteria release more of it as they break down manure. Steel plants emit massive amounts of CO_2 as the oxygen is stripped from iron ore at high heat. Cement plants release more CO_2 in the process of breaking down limestone into lime. Both steelmaking and cement-making require huge amounts of energy (usually from burning fossil fuels), and so does capturing nitrogen from the air to make synthetic fertilizer.

But there is, in fact, an escape hatch. With the right kind of deep-tech innovation in chemistry and biology, it's possible to implement versions of all these industrial processes that emit far lower volumes of greenhouse gases, and sometimes produce no net emissions at all. These innovations often involve new materials or fuels that work exactly the same way and perform to the same specifications as the traditional versions, and can therefore function as literal drop-in replacements—meaning customers don't have to retool their businesses to use them. The idea is to "meet industry where it is" in order to buy time for the harder, longer transitions, says DCVC's Bogue. "We're building the technologies that we need for the future. But we also need to do something for the interim."

One of the most compelling examples of a drop-in replacement is E-Jet, a form of sustainable aviation fuel (SAF) made using waste CO_2. E-Jet's manufacturer is a DCVC-backed company named Twelve, after the atomic weight of carbon. Its founders invented a form of low-temperature electrolysis that can take in CO_2 captured from industrial sources (and, one day, from the atmosphere), combine it with H_2O, and release oxygen as well as syngas. A mixture of hydrogen and carbon monoxide, syngas can easily be converted into long-chain hydrocarbons like kerosene, which powers jet engines. Electrolysis isn't the only way to make SAF, but it leads to far lower carbon emissions and uses far less water than the competing processes, including hydrodeoxygenation of vegetable oils or waste oils, the gasification of municipal waste, and the dehydration, oligomerization, hydrogenation, and fractionation of ethanol derived from corn.

Last summer, Twelve broke ground on a plant in Washington State that will produce up to a million gallons of E-Jet per year. This year the European airline group IAG—an umbrella for British Airways, Iberia, Aer Lingus, Vueling, and LEVEL—signed a 14-year purchase agreement with Twelve for 785,000 metric tons of the fuel. Microsoft and Alaska Airlines are partnering with the

Sustainable aviation fuel is one of the main products at Twelve. The company's electrolysis technology converts water and carbon dioxide (from point industrial sources or, eventually, from the air) into carbon monoxide and hydrogen—a combination known as syngas, which can be a feedstock for everything from jet fuel to plastic.

> "We will have the quiver of arrows that we need to fix [the climate]. But some of them may not come in time. Ergo, we're also investing in these replacement and adaptation technologies."

Zachary Bogue
Managing Partner, DCVC

company to use the fuel on demonstration flights, and eventually to reduce the climate impact of business travel on Alaska by Microsoft employees. "The aviation industry is serious about getting to net-zero carbon emissions," Bogue says. "Twelve offers a way to get there with a product that, when produced at scale, can fundamentally change aviation's carbon footprint without requiring changes to its infrastructure."

A similar story is playing out in a very different industry: the manufacturing of cement, the key binding ingredient in concrete. To make ordinary Portland cement (OPC)—the world's standard cement used in more than 90 percent of cement applications—limestone (calcium carbonate) is heated in large kilns, where it breaks down into quicklime (calcium oxide) and huge amounts of CO_2: over 900 kilograms for every 1,000 kilograms of cement produced. Combined with fly ash and slag from ore-smelting, quicklime forms clinker, which is ground up into cement powder. Brimstone, a DCVC portfolio company, has invented a new cement-making process that also results in OPC. But it starts with calcium silicate rocks such as basalt that are already rich in calcium oxide, meaning the company can skip the CO_2-producing step; if all OPC were made this way, the industry's process emissions would

▸ **Brimstone** made this cube of concrete using cement from carbon-free calcium silicate rocks.

go down by 60 percent. Not only that, but Brimstone's rocks also contain magnesium, which binds with and permanently sequesters CO_2 from the atmosphere. That makes Brimstone's overall process carbon-negative. And the more that the kilns and the remaining processes are powered with clean electricity, rather than coal, the more carbon-negative it can be.

In the agriculture business, one of the steps innovators would like to skip is fixing ammonia—the main ingredient in synthetic fertilizer—from nitrogen in the air. Modern farming wouldn't work without vast amounts of ammonia, and yet our main method for making it, the Haber-Bosch process, requires high pressure and temperature and is incredibly energy-intensive, using more than 1 percent of global energy production. (And that's just from the energy that goes into *making* fertilizer. The full nitrogen supply chain, from manufacturing to use, is responsible for about 5 percent of global greenhouse gas emissions.) Cambridge, U.K.-based Nium, where we invested in 2023, has built a reactor filled with a nanocatalyst material that can synthesize ammonia from nitrogen and hydrogen at low temperature and low pressure. This clean ammonia—produced at a fraction of the price and pollution of the Haber-Bosch process—can act as an energy-dense, convenient vector for transporting hydrogen fuel, or it can be used in fertilizer.

In the production of ammonia, we're also aided by evolution, which has produced microorganisms capable of making ammonia on their own. Pivot Bio's nitrogen-fixing microbes (see Opportunity 1.1) originally came from the soil, and have been modified to fix even more nitrogen than their natural cousins. Pivot's process for fermenting the microbes in the factory emits 98 percent less CO_2 than the Haber-Bosch process, and corn farmers who apply the microbes during planting have found that they can reduce synthetic fertilizer use by up to 40 pounds per acre, even as the plants themselves experience a 7 percent increase in nitrogen levels.

Unfortunately, we don't yet have a drop-in replacement for the beef from cattle (for more on why, see the "Cellular Agriculture" section of our 2023 Deep Tech Opportunities Report). But there is a way to cut the amount of methane cattle burp up as the result of fermentation in their guts (estimated at up to 500 liters per day for every cow, bull, steer, and heifer). In 2021, we invested in CH4 Global, a Nevada-based company that makes a cattle-feed supplement from a species of seaweed called *Asparagopsis*. Bromoform, a chemical produced in the seaweed to make it less tasty to marine herbivores, blocks methane production by gut bacteria, while leaving more calories available for the animal's own metabolism. Co-founder and CEO Steve Meller says

reaching 150 million cattle—just 10 percent of the world population—with CH4's Methane Tamer™ product would reduce methane emissions by 1.5 gigatons per year. The company announced in June that 70 head of cattle fed Methane Tamer™ at a feedlot in South Australia had been processed, making it the first and only reduced-methane beef selling into the Australian domestic market.

It's critical to note that progress on drop-in replacements is not a substitute for, or an excuse for delaying, the decarbonization of the broader economy. The 2015 Paris Agreement called on the nations of the world to cut overall greenhouse gas emissions by about 60 percent by 2035 and to reach net zero emissions by 2050, all toward the goal of keeping global temperature increases below 1.5 degrees Celsius. The lack of progress on emissions cuts in most nations, including the U.S., means that the 1.5-degree goal is now "deader than a doornail," in the words of legendary climate change researcher James Hansen, and that even a less ambitious 2-degree goal may already be out of reach. That makes it all the more important that we ultimately phase out or abate *all* fossil fuels and reinvent every carbon-emitting industrial

▸ **CH4 Global**
Asparagopsis seaweed contains a chemical called bromoform that, when added to cattle feed, blocks methane production in the gut.

process. Finding smart substitutes for traditional fuels and feedstocks will be a key part of that push, but only a part.

"In interviews I'm often asked the question, 'Are you pessimistic or optimistic about the climate?'" says DCVC's Bogue. "My answer is that I'm both. I'm optimistic because all these amazing technologies are coming down the pike and we will have the quiver of arrows that we need to fix this. But some of them may not come in time, so there's going to be a mismatch, which will potentially lead to some rough decades or more while we scale them. Ergo, we're also investing in these replacement and adaptation technologies."

Figure 3.2.1 Traditional cement's carbon footprint

The most energy- and carbon-intensive parts of traditional cement production are heating the kiln and breaking down limestone (calcium carbonate) into quicklime (calcium oxide).

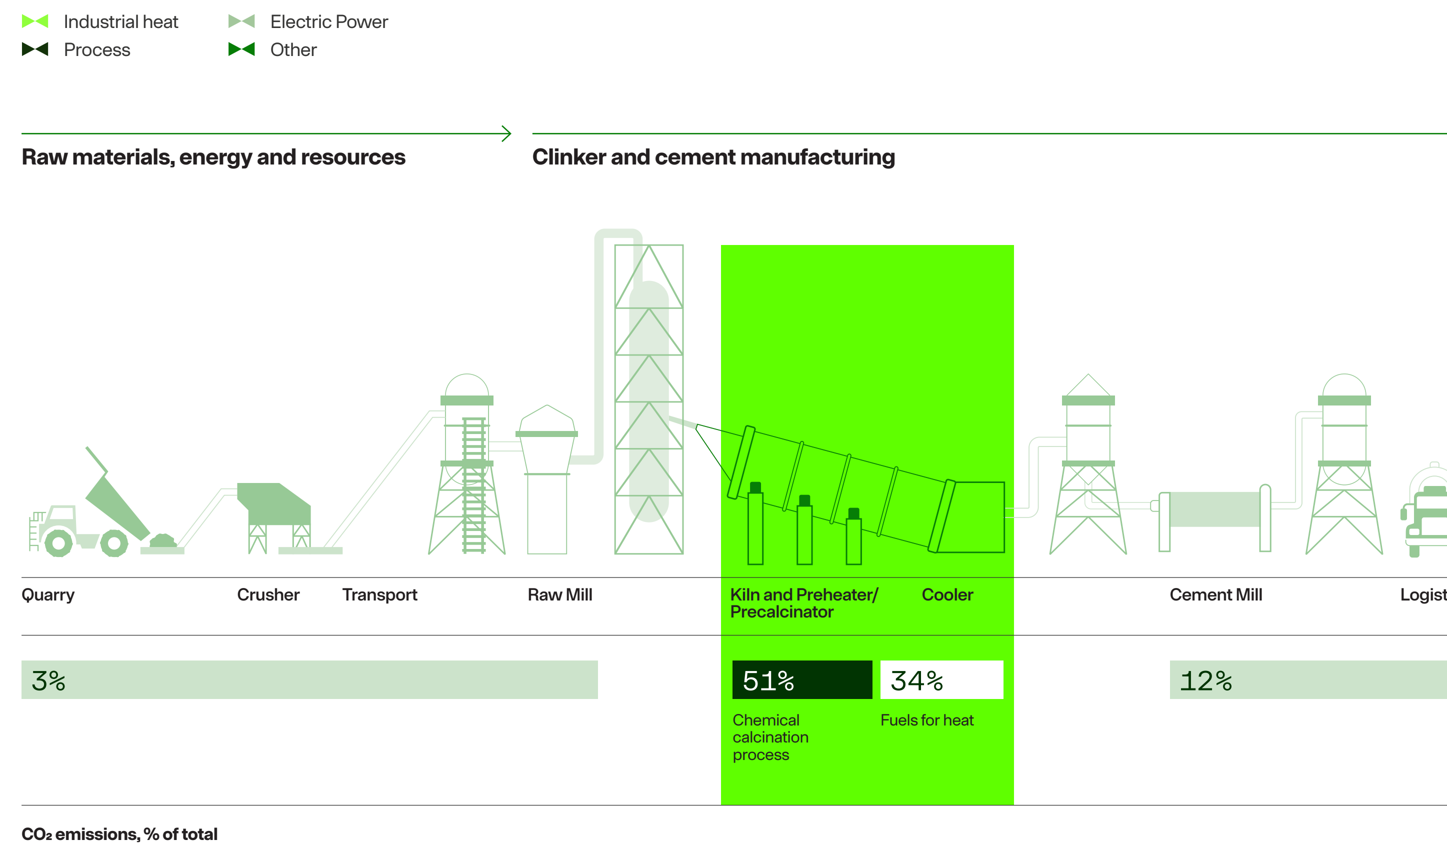

Greenhouse gas emissions from solar farms, wind turbines, and nuclear plants are extremely low compared to those from natural gas, oil, and coal, but they're not zero (see Figure 3.3.1). Until we decarbonize the global economy, every made thing will have small associated GHG emissions, usually upstream in its lifecycle, when the resources for its components were extracted and those components were assembled. The question is how small—and that depends, in part, on where the mines and factories are located.

To understand this nuance, look at the solar industry. Solar photovoltaic panels are made using a high-purity form of silicon called polysilicon. The unfortunate reality is that half of the polysilicon consumed by the solar industry globally comes from the Xinjiang province of China.

That's a problem for two reasons. First, China's forced-labor system for Uyghurs and other Muslim minorities in the western province means that workers in polysilicon mines and factories "are forced to work under guard and constant threats," according to the U.S. Department of Labor. Second, about 63 percent of China's electricity, including the power for the factories that make solar cells, comes from coal-fired power plants, according to the International Energy Administration. The percentage is likely even higher in the country's far west, where there are no nuclear or natural gas power plants.

And China is *adding* to its coal generating capacity, not shrinking it. The country approved 114 gigawatts of new coal power plants in 2023, up 10 percent from 2022. While China signed on to the 2015 Paris Agreement, which calls for the almost complete phase-out of coal energy by 2040, a new report from Norwegian risk assessment firm DNV projects that China will manage to cut its coal consumption by only one-third by 2040.

The truth is that solar, wind, and other renewable and storage projects using equipment sourced from China have a notable and indelible carbon footprint—which is part of the reason DCVC continues to invest in clean, onshore energy projects like Fervo Energy's geothermal wells and Oklo's next-generation fission microreactors (see Opportunity 3.1). "PRC energy is so dirty, and the supply chain and fab for this stuff is so energy-intensive, that the true 'net-zero GHG' payback period for this gear is longer than people think," says DCVC's Ocko.

Figure 3.3.1 Life cycle greenhouse gas emissions from the energy infrastructure

Even renewable energy technologies such as solar and wind power have a small but indelible carbon footprint.

Source: National Renewable Energy Laboratory

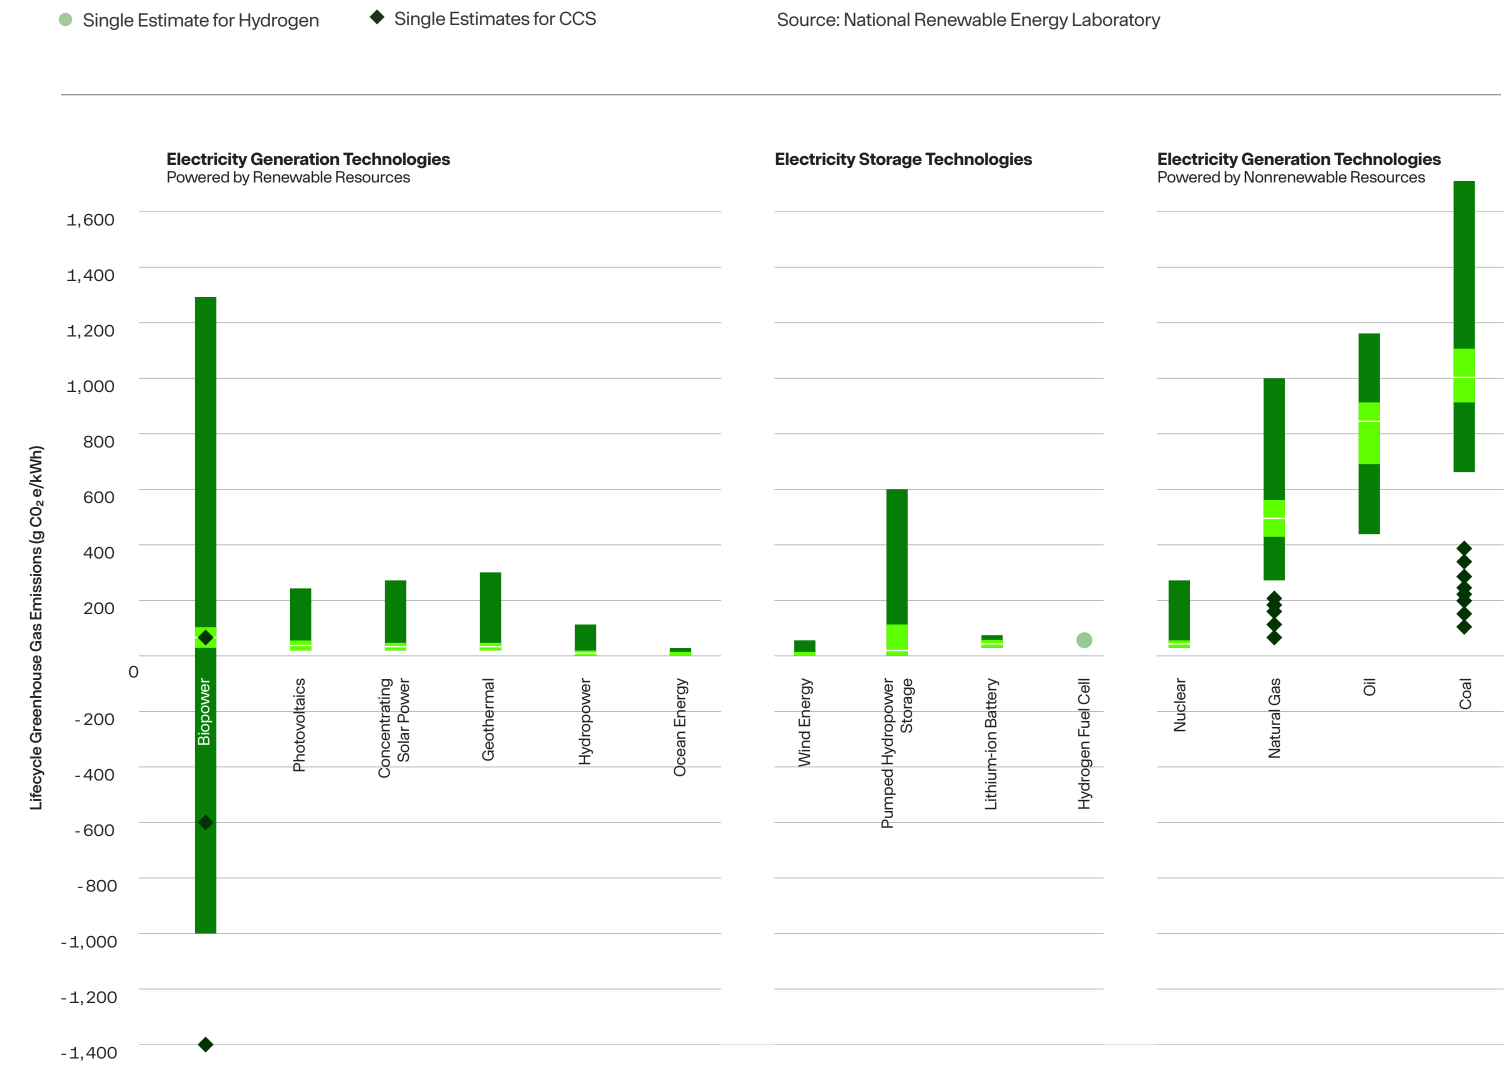

"PRC energy is so dirty, and the supply chain and fab for this stuff is so energy intensive, that the true 'net-zero GHG' payback period for this gear is longer than people think."

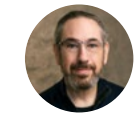

Matt Ocko
Managing Partner, DCVC

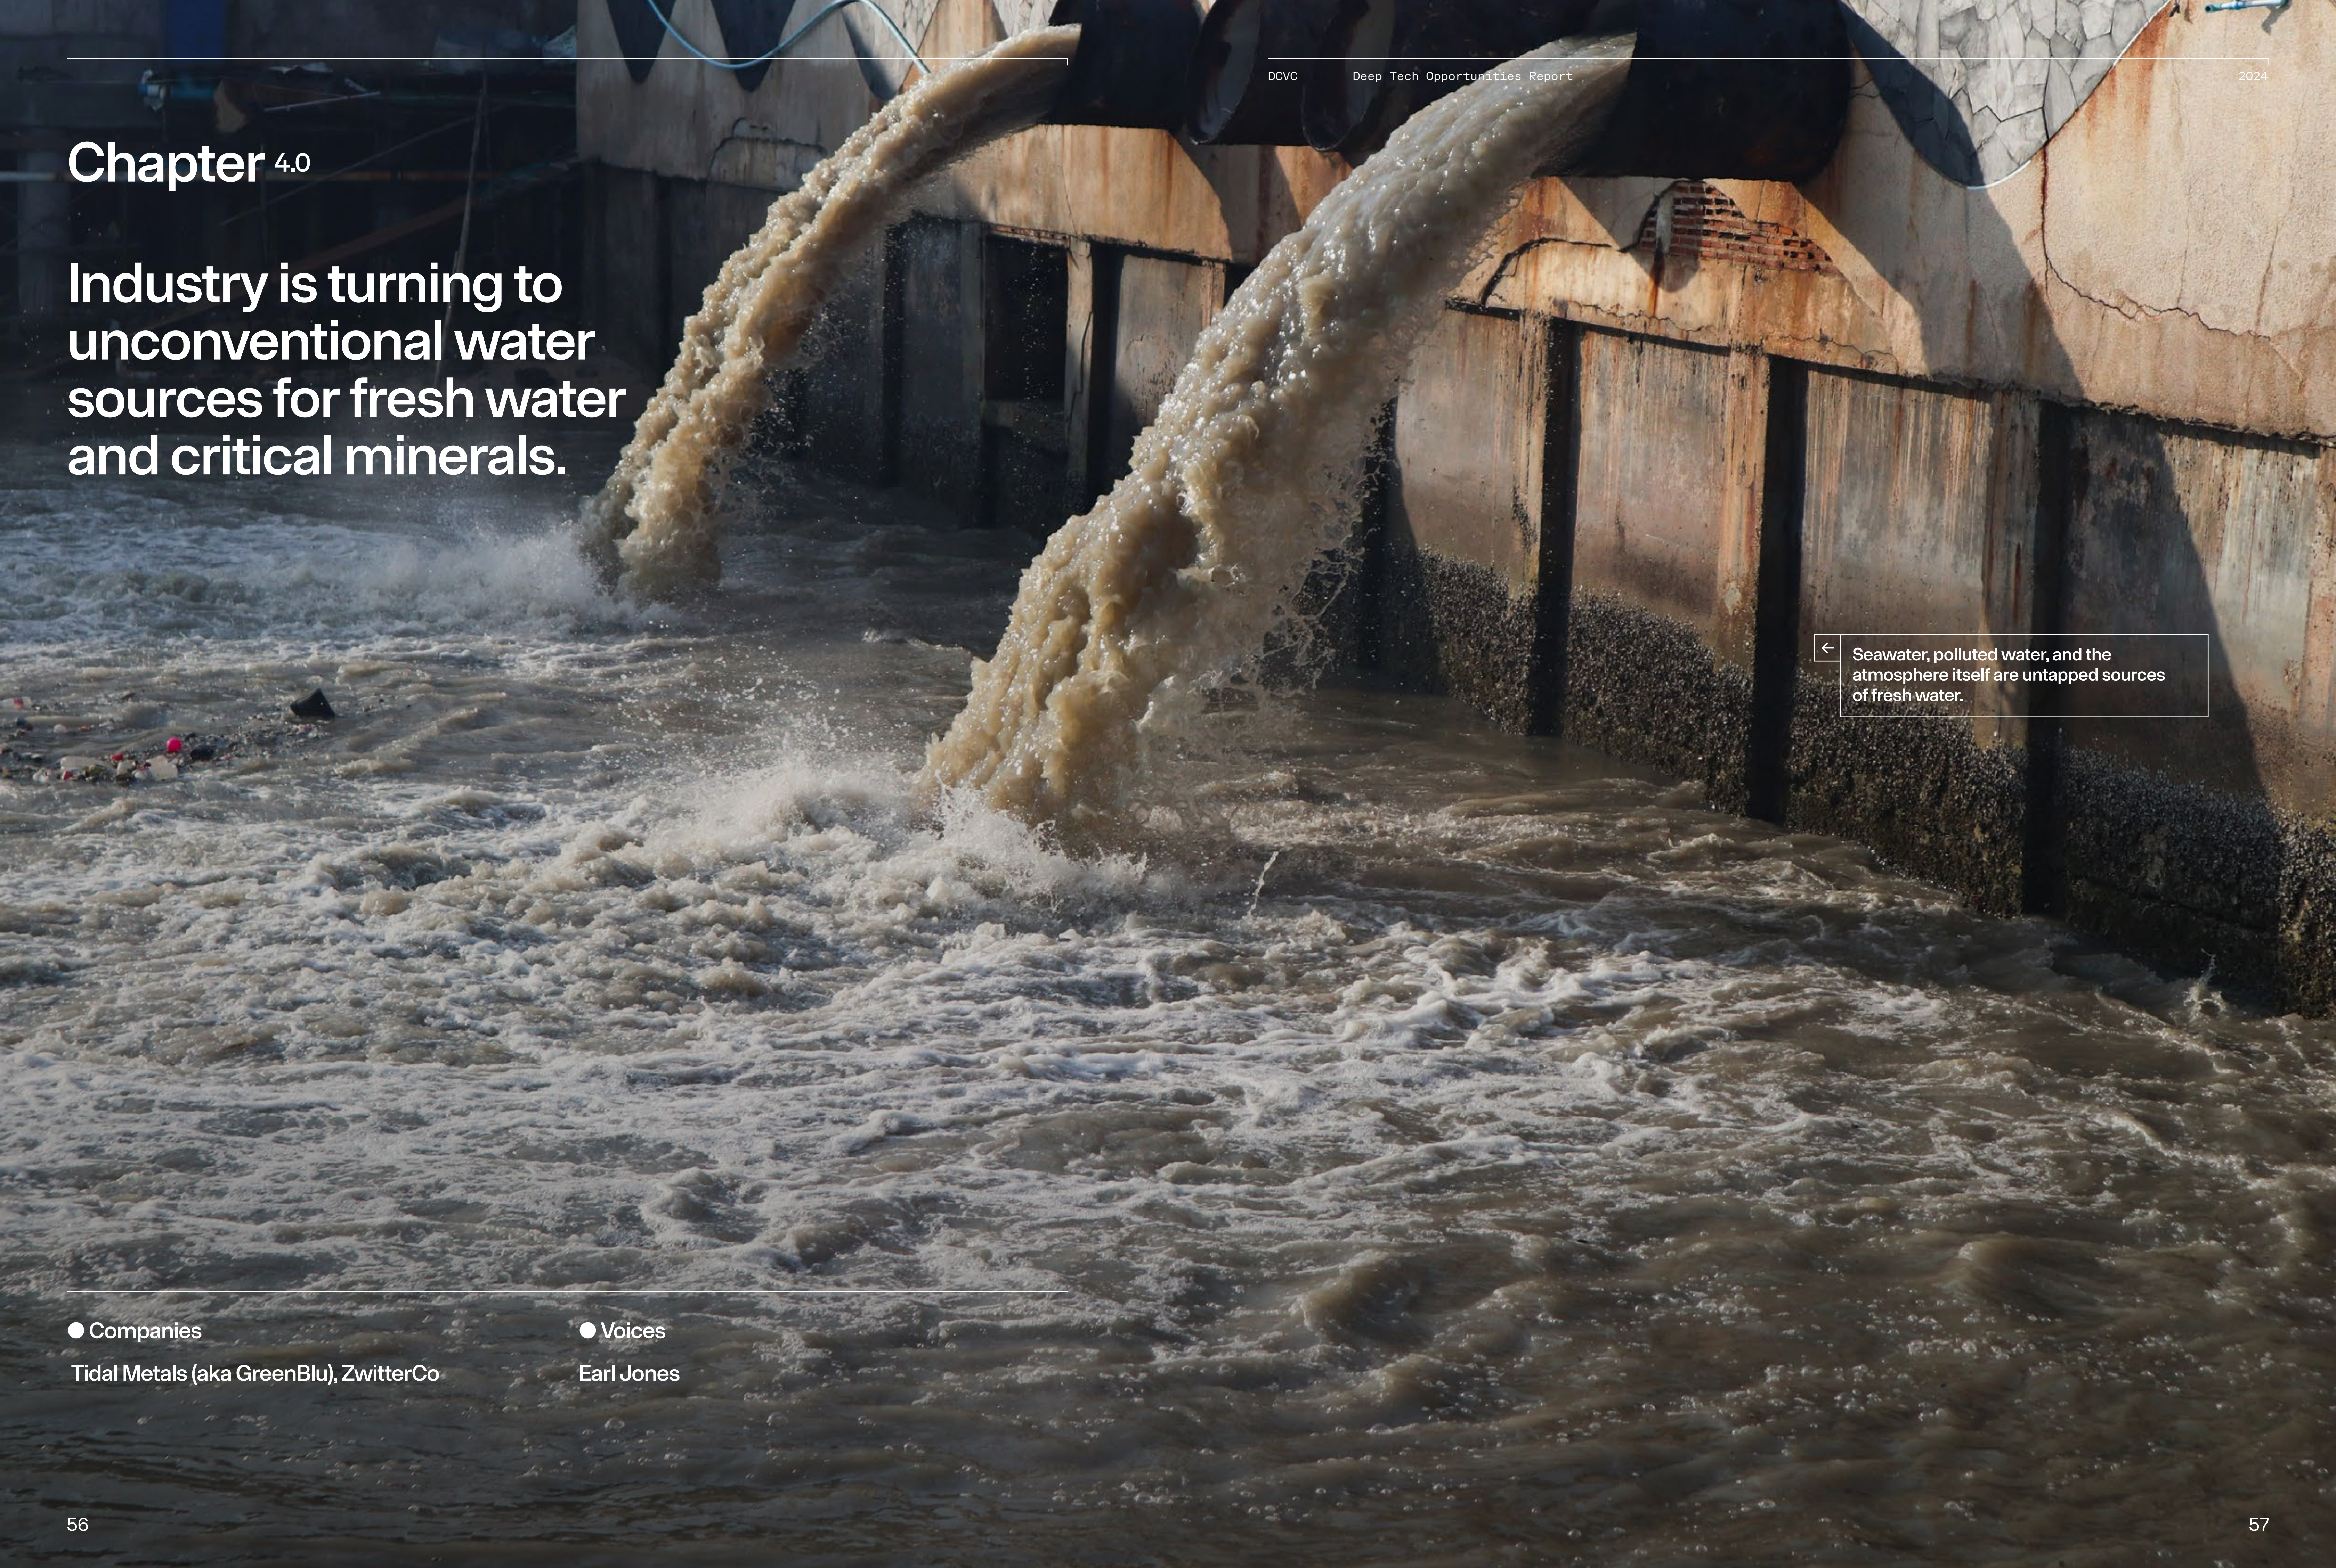

Chapter 4.0

Industry is turning to unconventional water sources for fresh water and critical minerals.

← Seawater, polluted water, and the atmosphere itself are untapped sources of fresh water.

● Companies

Tidal Metals (aka GreenBlu), ZwitterCo

● Voices

Earl Jones

OPPORTUNITY 4.1

INDUSTRY IS TURNING TO UNCONVENTIONAL WATER SOURCES FOR FRESH WATER AND CRITICAL MINERALS.

Almost every goal outlined in the other parts of this report would be impossible, or far more difficult, to meet without a steady supply of clean water and critical minerals. Indeed, in the absence of these basic resources, the economic stability and the surplus of ideas and capital that undergird innovation itself would quickly disappear. Yet the reality is that the world's most easily accessed supplies of fresh water and key minerals are being exhausted.

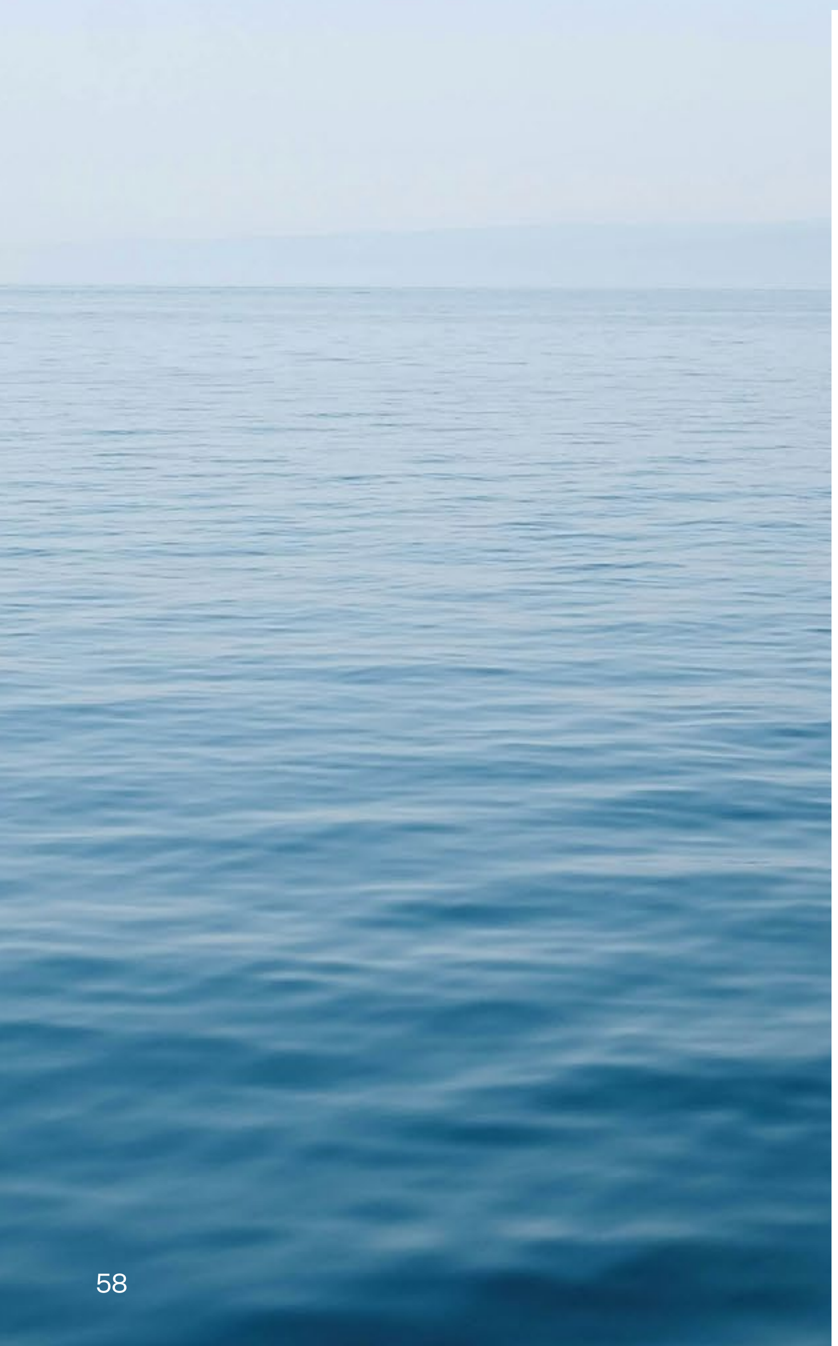

Thanks to population growth, *conventional* water sources—lakes, rivers, reservoirs, and aquifers—are being used up faster than they can be restored by the natural water cycle. At the same time, due to changing precipitation patterns, many locations on Earth are receiving far less water than they need, and others far more. Over the next 10 years, water-driven events (storms, floods, droughts, and wildfires) will likely cause $1.8 trillion in damage, or about 69 percent of all expected losses from climate change, according to a 2023 forecast from Global Water Intelligence and XPV Water Partners. So not only is the era of cheap, abundant water over, but we must now account for—and find ways to compensate for—global warming's impact on the water cycle.

That will certainly mean conserving existing water resources and using them more efficiently. But it will also mean aggressively harnessing *unconventional* water resources, such as seawater, industrial and municipal water, polluted water, and the atmosphere itself. "We use the term 'unconventional' to describe currently untapped water resources," says Earl Jones, an operating partner at DCVC. "It's a nice way of saying 'dirty, expensive, and hard to access.' When problems are expensive and hard, and the value of solving them is large, innovators get to work."

One unconventional source of water is the air. Earth's atmosphere holds about 13,000 cubic kilometers of water, or 13 quadrillion liters. That's tiny compared to the

Figure 4.1.1	Where is the Earth's Water?
Most of the world's water is locked up in places where we can't easily or cheaply access it: the oceans, the icecaps, and the atmosphere, as well as polluted industrial and municipal sources. Source: U.S. Geological Survey	

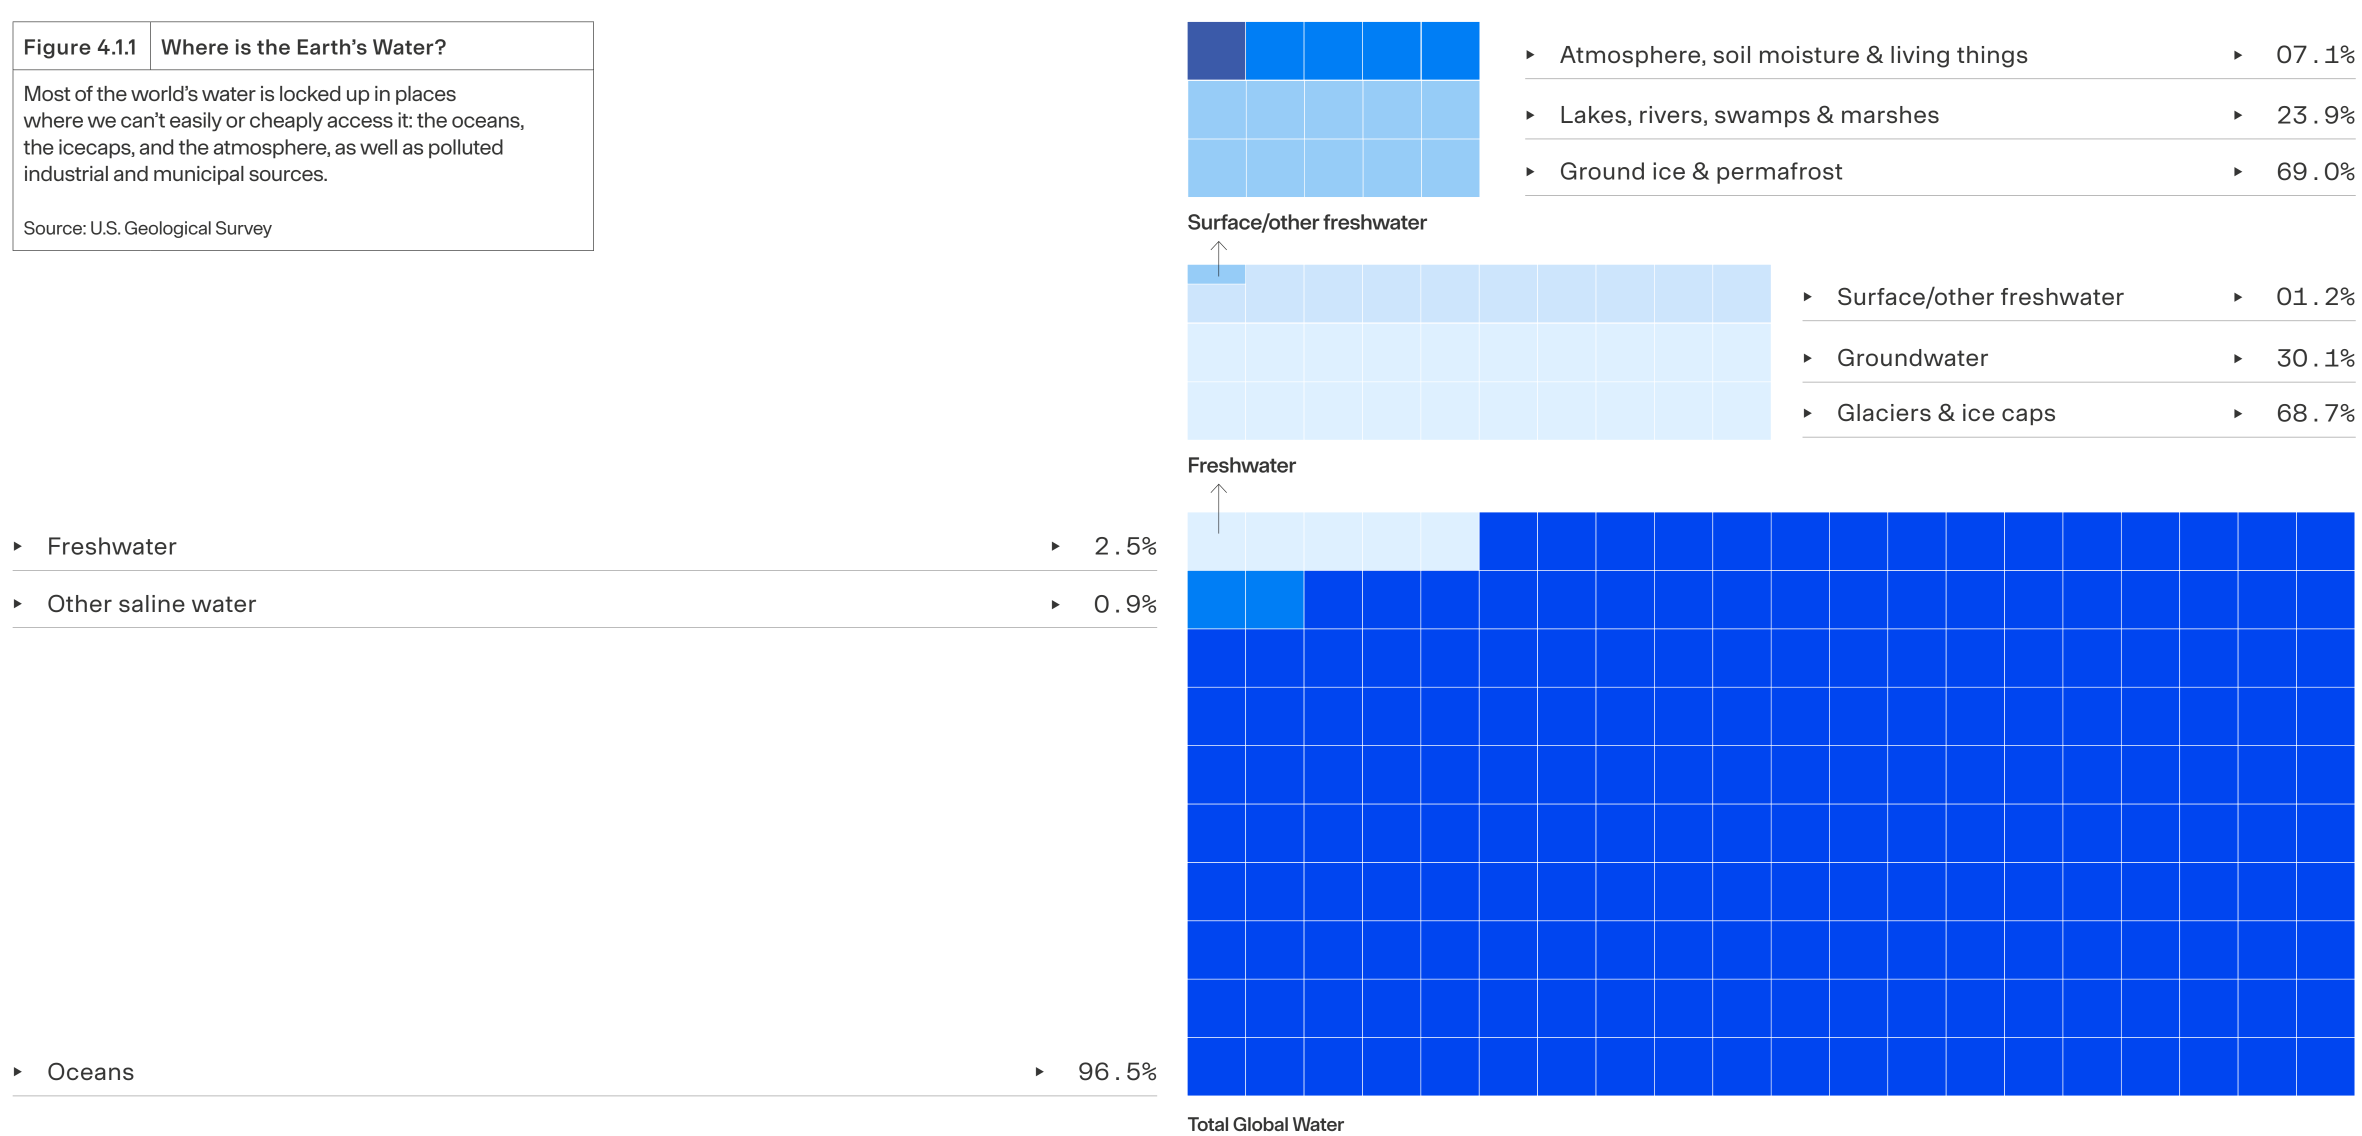

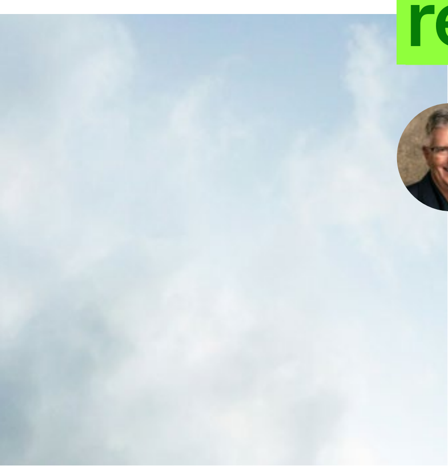

"There's no inexpensive conventional water left, and in fact, we are depleting our conventional water faster than we can sustainably replenish it."

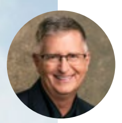

Earl Jones
Operating Partner, DCVC

volume of the oceans; that said, the exchange of water through solar evaporation and rain is a whopping 600,000 cubic kilometers per year. The only catch is that rain often falls far from where we need it. If we could extract H_2O from the atmosphere on demand, inexpensively, we'd have an essentially inexhaustible source of freshwater.

The purposeful extraction of water from air is called atmospheric water harvesting. Technically, it's simple: we do it all the time with home dehumidifiers. Such mechanical condensing systems have the advantage of being able to produce large quantities of water—to the tune of tens of thousands of liters per day. Unfortunately, they use so much energy that they're orders of magnitude more expensive than other approaches to water generation, such as seawater desalination. And thermodynamic constraints may limit their cost-reduction potential.

There's another approach to atmospheric water harvesting that may hold more promise: sorption systems that use novel materials, such as salt-impregnated polymers, zeolites, or metal-organics frameworks to suck water out of the air. Here, the energy penalty comes from desorption, since it takes a lot of heat to get the water back *out* of the materials. But there's the potential to reduce the energy cost below that of condensing systems, with the advantage that sorption systems can be run completely off-grid, using solar photovoltaic and solar thermal energy.

DCVC is supporting the development of a computationally designed light-switchable polymer that could fundamentally alter the energy requirements for sorption systems. An example of a light-switchable polymer is the photochromic layer in transition eyeglasses. When exposed to sunlight, the polymer changes color, giving the tint; indoors, the lenses revert to being clear. A DCVC-backed incubation will use a similar technology for water harvesting. "When the light is on, water is captured, and when the light is off, the water is released," Jones says. "If we can bring this incubation to market and grab a lot of water cost-effectively, it has the potential to be world-changing."

Another way to improve the economics of accessing unconventional water sources is to extract impurities or contaminants that may themselves be valuable. Seawater, for example, represents a nearly inexhaustible supply of magnesium—a special metal that has structural properties similar to steel and aluminum, but is 75 percent lighter than steel and 33 percent lighter than aluminum. If we had access to larger, cheaper supplies of magnesium, we could build lightweight magnesium vehicles that could travel much farther on a single tank of fuel or a single battery charge. (Magnesium is also used in many aluminum alloys, and is essential in the purification of steel and the refining of titanium.)

Herbert H. Dow invented a process for extracting magnesium from seawater in 1916. However, the cost of the process was high. Today, nearly all magnesium comes from mined ores such as dolomite and magnesite, 70 percent of which are mined and refined in China.

The refining method used to isolate magnesium, known as the Pidgeon process, is incredibly dirty, using more than 18 metric tons of coal and emitting 28 tons of CO_2 for every ton of magnesium produced. Clearly, we need a cleaner source of magnesium, and a DCVC-backed company, Tidal Metals, is turning to the oceans.

Every metric ton of seawater contains 1.3 kilograms of magnesium chloride. Tidal, headed by three physicists from MIT, Princeton, and the University of Wisconsin, has patented a technology that can economically process seawater to extract anhydrous magnesium chloride, which can then be separated into chlorine and magnesium metal using traditional electrolysis.

The cost-efficient Tidal Metals process runs on electricity, and in the limit, can be powered entirely by renewable electricity. By eliminating the need for mined ores and dirty, energy-intensive thermal refining,

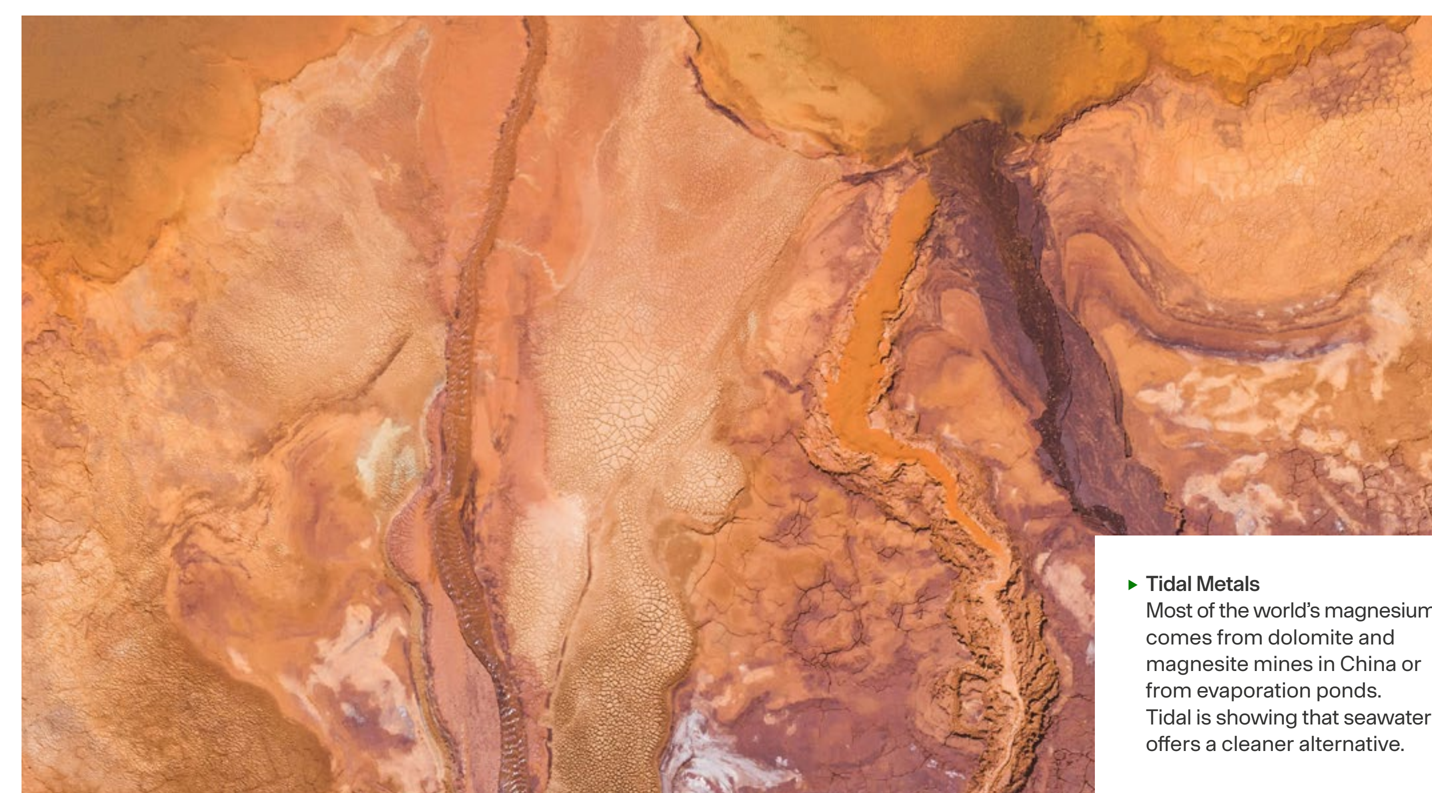

▸ Tidal Metals
Most of the world's magnesium comes from dolomite and magnesite mines in China or from evaporation ponds. Tidal is showing that seawater offers a cleaner alternative.

Tidal Metals will deliver "green magnesium" for industries like auto manufacturing without a green-premium price. (Replacing one ton of steel in a car with one ton of Tidal's magnesium would save 28 tons of CO_2 from steel production and 67 tons of CO_2 from lower lifetime vehicle emissions, according to Tidal's chief technology officer, Ethan Schartman.)

In last year's Deep Tech Opportunities Report we profiled ZwitterCo, another DCVC-backed company using new materials to rethink water treatment. The company started out selling a non-fouling membrane designed to recapture high-value biological materials from the wastewater at facilities like dairies and pharmaceutical fermentation plants. Now it's being pulled into an even bigger market, with a reverse osmosis membrane for brackish water, developed at its new R&D center in Woburn, Mass.

R&D investments ZwitterCo built a 30,000-square-foot innovation center in Woburn, Mass., to scale up new water filtration products.

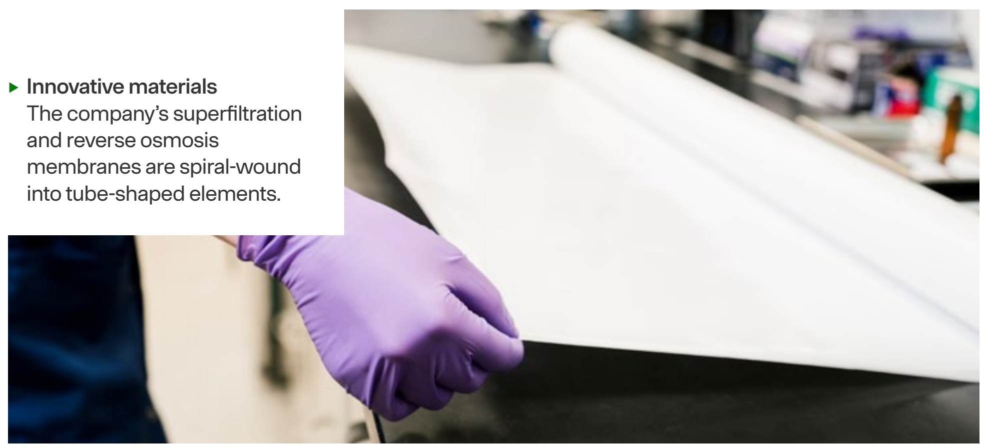

Innovative materials The company's superfiltration and reverse osmosis membranes are spiral-wound into tube-shaped elements.

Figure 4.1.2 The power of zwitterions

Unlike traditional filtration products, ZwitterCo membranes are built from zwitterions that attract water while seamlessly repelling organic compounds that normally stick to membranes and ruin their filtration capacity. The membranes are designed to industry-standard specifications, allowing them to work with existing filtration equipment.

01 **Zwitterions attract to water molecules.**

02 **Copolymers bind zwitterions with hydrophobic molecules.**

03 **Zwitterions create foul-immune, water-loving channels.**

04 **Zwitterionic channels retain contaminants and produce clean water.**

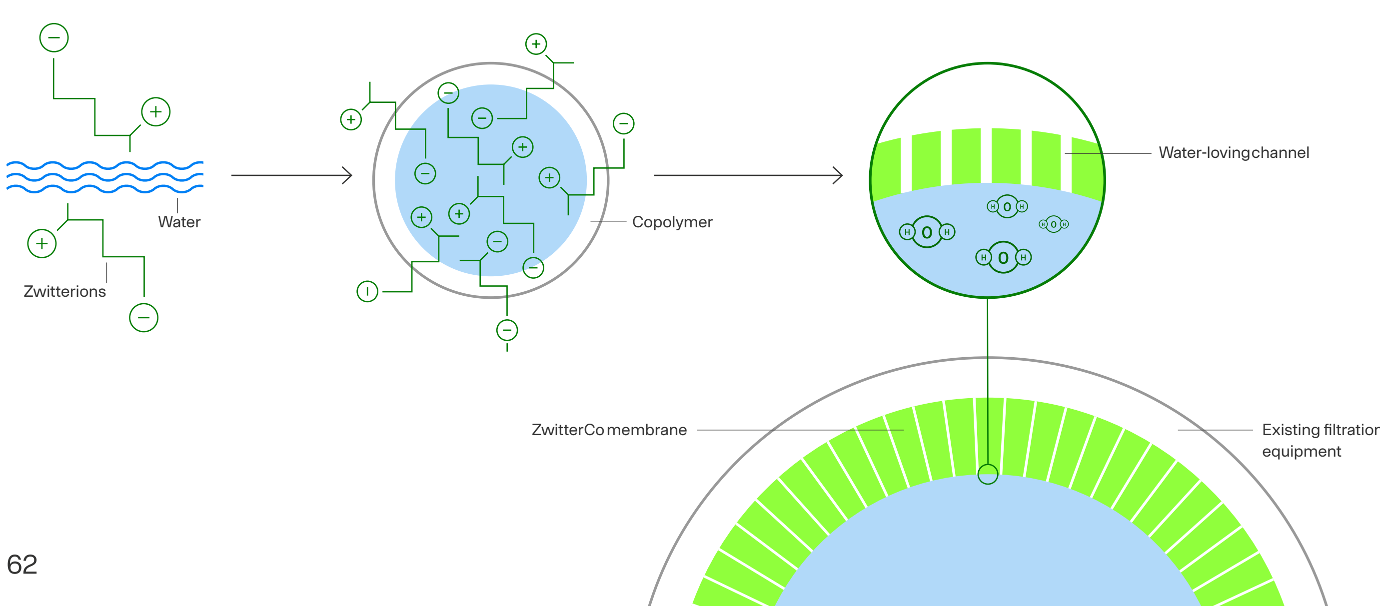

Reverse osmosis Stacks of ZwitterCo filtration elements are a longer-lasting, easier-to-maintain replacement for existing brackish water reverse osmosis (BWRO) membranes.

Brackish water contains dissolved solids like sodium chloride but isn't quite as salty as seawater. Facilities such as power plants and wastewater treatment plants use trillions of gallons of it each year, and to help keep it clean they generally use reverse osmosis. In RO, pressure is applied on one side of a selective membrane, forcing the smaller molecules of the solvent (water, in this case) through the membrane while retaining the larger molecules of solute (salts and other solids) on the pressurized side. The problem is that most brackish waters contain biological contaminants that quickly clog RO membranes, meaning it takes more pressure and more energy to push the H_2O molecules through. The electrical properties of ZwitterCo's zwitterionic membranes make them immune to this fouling, which minimizes downtime, since the filters don't have to be cleaned as frequently.

"If you can keep the membranes cleaner, you don't need the same kind of harsh cleaning chemicals, and you don't need as much energy," says Jones. "And that's where ZwitterCo has found receptive customers, because the RO membranes that are in use today are just not working well."

If the U.S. intends to extend water, power, and mineral resources, rebuild its industrial base, and repatriate more of the manufacturing supply chain from other countries, we don't need to build more polluting factories or dig new mines. To a large extent, we simply need to get smarter about finding the needed resources in unexpected places. And we think that's one mission where venture-backed deep-tech innovation can help.

Chapter 5.0

Venture investment can help drive the next round of innovation in defense and aerospace.

Smart capital can steer these industries toward technologies that will help the U.S. and its allies reestablish their advantage in national security and global intelligence.

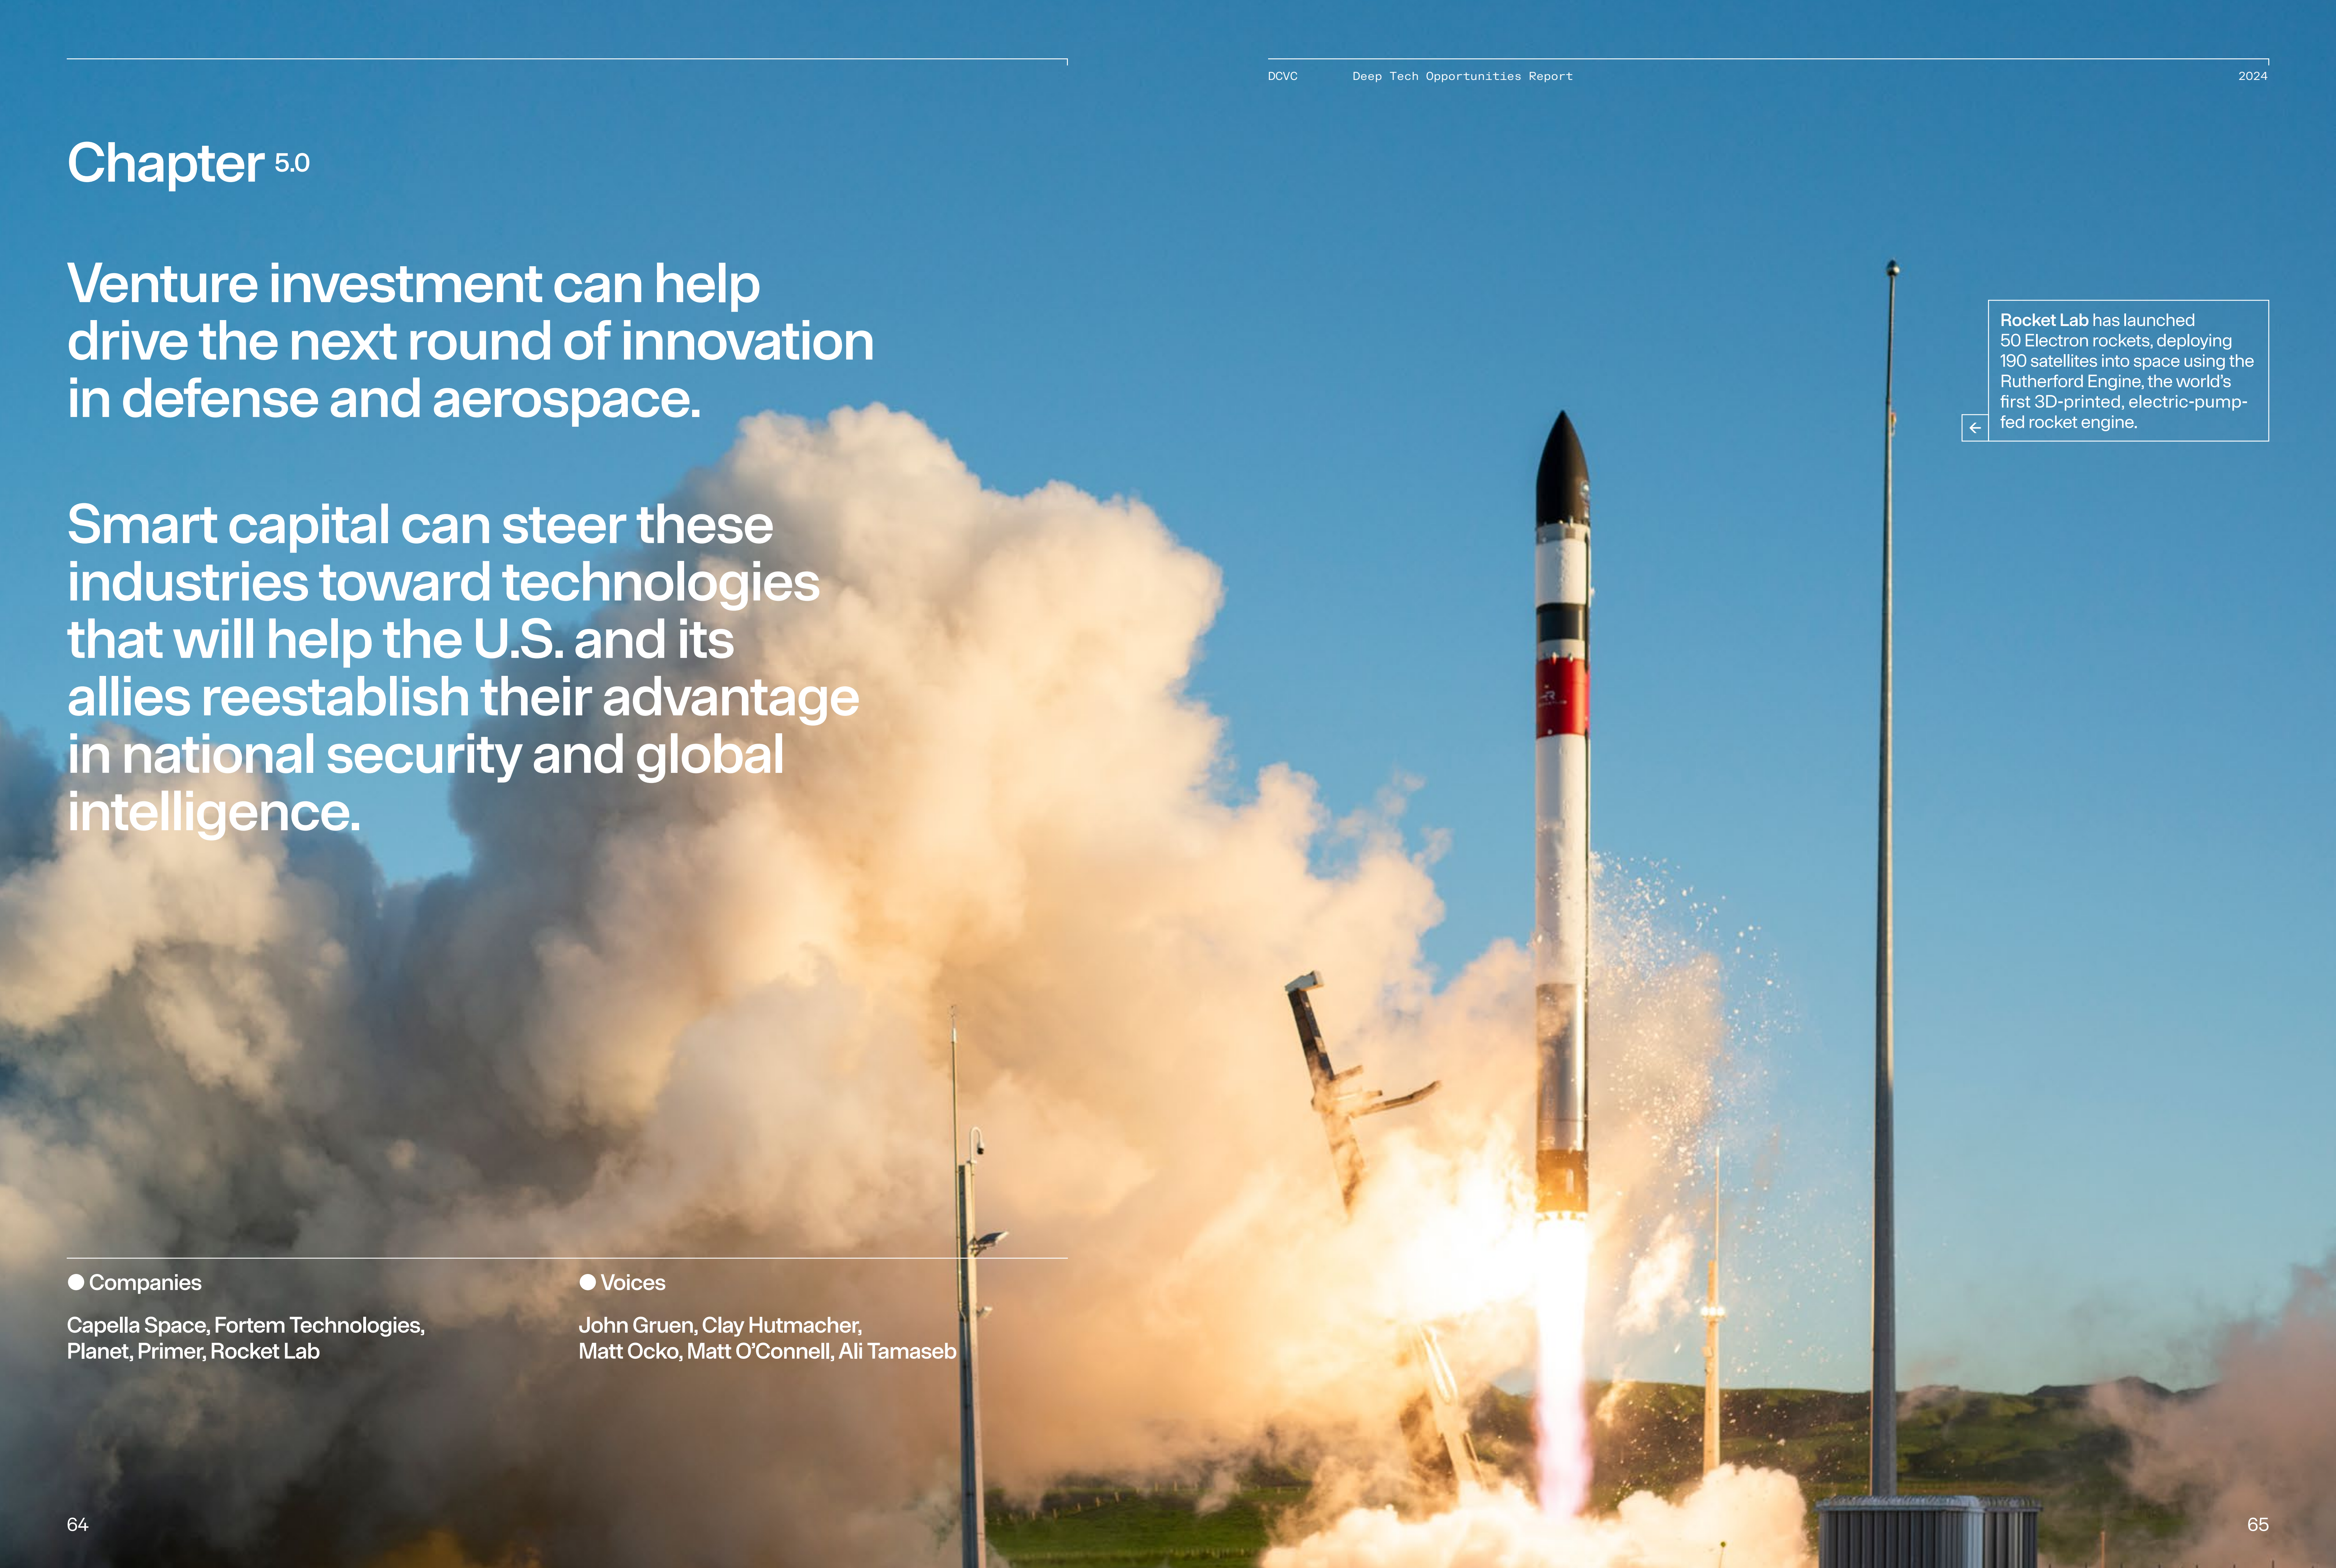

Rocket Lab has launched 50 Electron rockets, deploying 190 satellites into space using the Rutherford Engine, the world's first 3D-printed, electric-pump-fed rocket engine.

● Companies

Capella Space, Fortem Technologies, Planet, Primer, Rocket Lab

● Voices

John Gruen, Clay Hutmacher, Matt Ocko, Matt O'Connell, Ali Tamaseb

OPPORTUNITY 5.1

STARTUP THINKING WILL HELP THE DEFENSE AND SECURITY SECTORS PREPARE FOR NEW THREATS.

To counter emerging threats, the U.S. must do more to balance its big, expensive, and increasingly vulnerable defense systems with platforms that are "small, smart, cheap, and many," in the words of Deputy Defense Secretary Kathleen Hicks. This is a realm where venture-backed entrepreneurs can make a key difference.

There was a time when individual scientists, inventors, and entrepreneurs were able to contribute world-changing ideas to U.S. and allied military establishments. Private inventors built the first combat submarines and machine guns in the 1860s and the first tanks in the 1910s. Tinkerers in the hinterlands—the Wright brothers in Dayton, Ohio, and Robert Goddard in Worcester, Mass.—pioneered the ideas of powered, controlled flight and liquid-fueled rocketry, now cornerstones of defense. What President Dwight Eisenhower called the "military-industrial complex," dominated by a handful of consolidated defense contractors, only emerged during and after World War II. And for good reasons: projecting world-leading military power takes coordinated, continued investment on a scale only national governments and big companies can afford. But today the landscape of threats and vulnerabilities has shifted, and smaller players, including venture-scale startups, can once again contribute to maintaining and modernizing our defense infrastructure.

First, there's a new domain where offensive and defensive capabilities take a very different form: cyberspace. When computer systems are globally interconnected, they can be globally disrupted. "The next series of [cyber] attacks are going to be AI-generated," says DCVC general partner Ali Tamaseb. "And we need to be able to catch that with AI."

Second, we've entered an era when small, inexpensive weapons pose a growing threat. Today drones and other commercial-off-the-shelf (COTS) technologies can quickly, cheaply, and stealthily carry out high-precision single-use strikes. That makes our traditional "big iron" platforms such as tanks, planes, and surface ships look increasingly lumbering. The Ukrainian military, for example, has been able to use airborne drones, some costing only a few hundred dollars, to monitor the movement of Russian forces, identify artillery targets, drop munitions on tanks and trenches, and damage military installations. At sea, Ukraine has used remote-controlled drone boats to hit an amphibious landing ship, a minesweeper, and a civilian Russian tanker.

The world's largest militaries are not adjusting quickly to these asymmetric threats. "It's still important to build systems that can deliver a lot of throw weight against our adversaries," says DCVC managing partner Matt Ocko. "Rockets, artillery systems, ICBMs aren't going away. But today someone with relatively small financial resources and few connections could put together a private army, air force, or navy and degrade the U.S. Sixth Fleet [which operates around Europe and Africa] as a combat operational group. The technology is there to do this right now. And frankly, that should concern us tremendously."

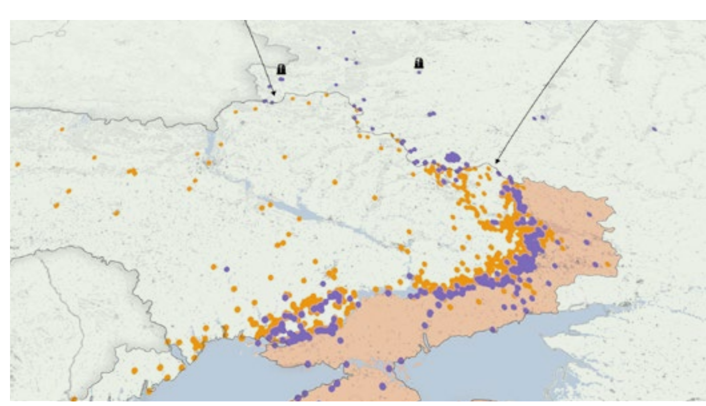

▸ **The Drone Wars** A map of Ukrainian drone attacks on Russia (purple) and Russian drone attacks on Ukraine (orange) between February 24, 2022, and March 12, 2024. (Source: Reuters)

At DCVC, we believe entrepreneurs can help counter asymmetric threats and regain a technology advantage over our real and potential adversaries. Fortem Technologies, the leader in AI-powered drone-hunting systems, is a prime example. The company builds small, high-definition radar systems—from the size of a pencil box to that of a laptop—that enable drone-on-drone defense. Its DroneHunter drones can detect other fixed-wing and rotary drones and fire countermeasures such as tether nets that bring the enemy craft to the ground without destroying them. "The Ukrainians have also been particularly successful in reprogramming the captured drones and using them against the Russians, which saves the Ukrainians considerable time and money that would otherwise be spent on purchasing and programming their own drones," says Fortem CEO Jon Gruen. The same net drones have been used to protect the World Economic Forum in Davos and the World Cup in Qatar.

Geospatial intelligence is another key to anticipating and defending against attacks. In this area we've invested in Capella Space, whose constellation of synthetic aperture radar (SAR) satellites offers 24/7, all-weather imagery of Earth's surface. Together with Planet, DCVC's first space investment in 2013, Capella is imaging Earth every day and helping security officials detect changes on the ground and anticipate threats. And thanks to SAR's high resolution and high contrast, as well as Capella's high revisit rate, buyers of the company's unclassified images can better monitor sea lanes, ports, and other facilities to track the arrivals and departures of ships and aircraft. (Through our investment in Rocket Lab, maker of the world's first reusable orbital small rocket, we've also helped open up cheaper access to space for a range of customers, including Planet and Capella.)

▸ **Fortem Technologies** DroneHunter F700 can intercept small and large drones—from quadcopters to fixed-wing aircraft—and ensnare them using rapidly expanding nets fired from its NetGuns attachment.

AI technology from a DCVC portfolio company called Primer is another compelling example of a dual-use technology that can aid in national defense as well as commercial applications. The company monitors both proprietary and publicly available data, such as text, audio, and video, and uses AI models to organize and summarize it and help analysts identify changes linked to critical entities, locations, or events. Primer can also help make sense of the data held *inside* a customer's organization, says DCVC operating partner and Primer board member Clay Hutmacher, who is a retired major general and former director of operations for the U.S. Special Operations Command. The Department of Defense is "applying [Primer's] AI to questions like, How would the Air Force and the Navy quickly disperse ships in the event of a conflict?" Hutmacher says. "How can they improve maintenance procedures on aircraft and ships and to give them more operational availability?"

The reality, we think, is that many of the ideas the U.S. military needs to stay ahead of both near-peer adversaries and smaller foes can and should come from the startup world. In Washington, D.C., we've advocated for increased investment in efforts such as the Defense Innovation Unit (DIU), which is working to accelerate the adoption of commercial technology throughout the military, and Replicator, a Pentagon-wide program intended to "accelerate the delivery of innovative capabilities to warfighters at speed and scale." (The program's goal is to field thousands of expendable, autonomous systems in the next 18 months.) We've also pushed for more spending within the combatant commands to support testing and adoption of such technologies. And together with other defense-tech investors, we intend to back more Silicon Valley–style companies creating a new backbone of distributed, autonomous, intelligent, affordable systems. "You need an atmosphere where people can test new things, big things, things that might fail, but that could also succeed in a game-changing way," Deputy Secretary Hicks said in a 2023 address. We agree.

We're proud to live in a nation that celebrates adventurers and risk-takers. And we're excited to watch NASA inch closer to its goal of returning Americans to the moon—a goal that will be accomplished, just as it was during the Apollo era, through an intensive collaboration between government agencies and aerospace contractors. But...

...we see purely private-sector efforts to establish a permanent human presence in space, or spin up mining and manufacturing activities in space, as premature—and, arguably, profligate.

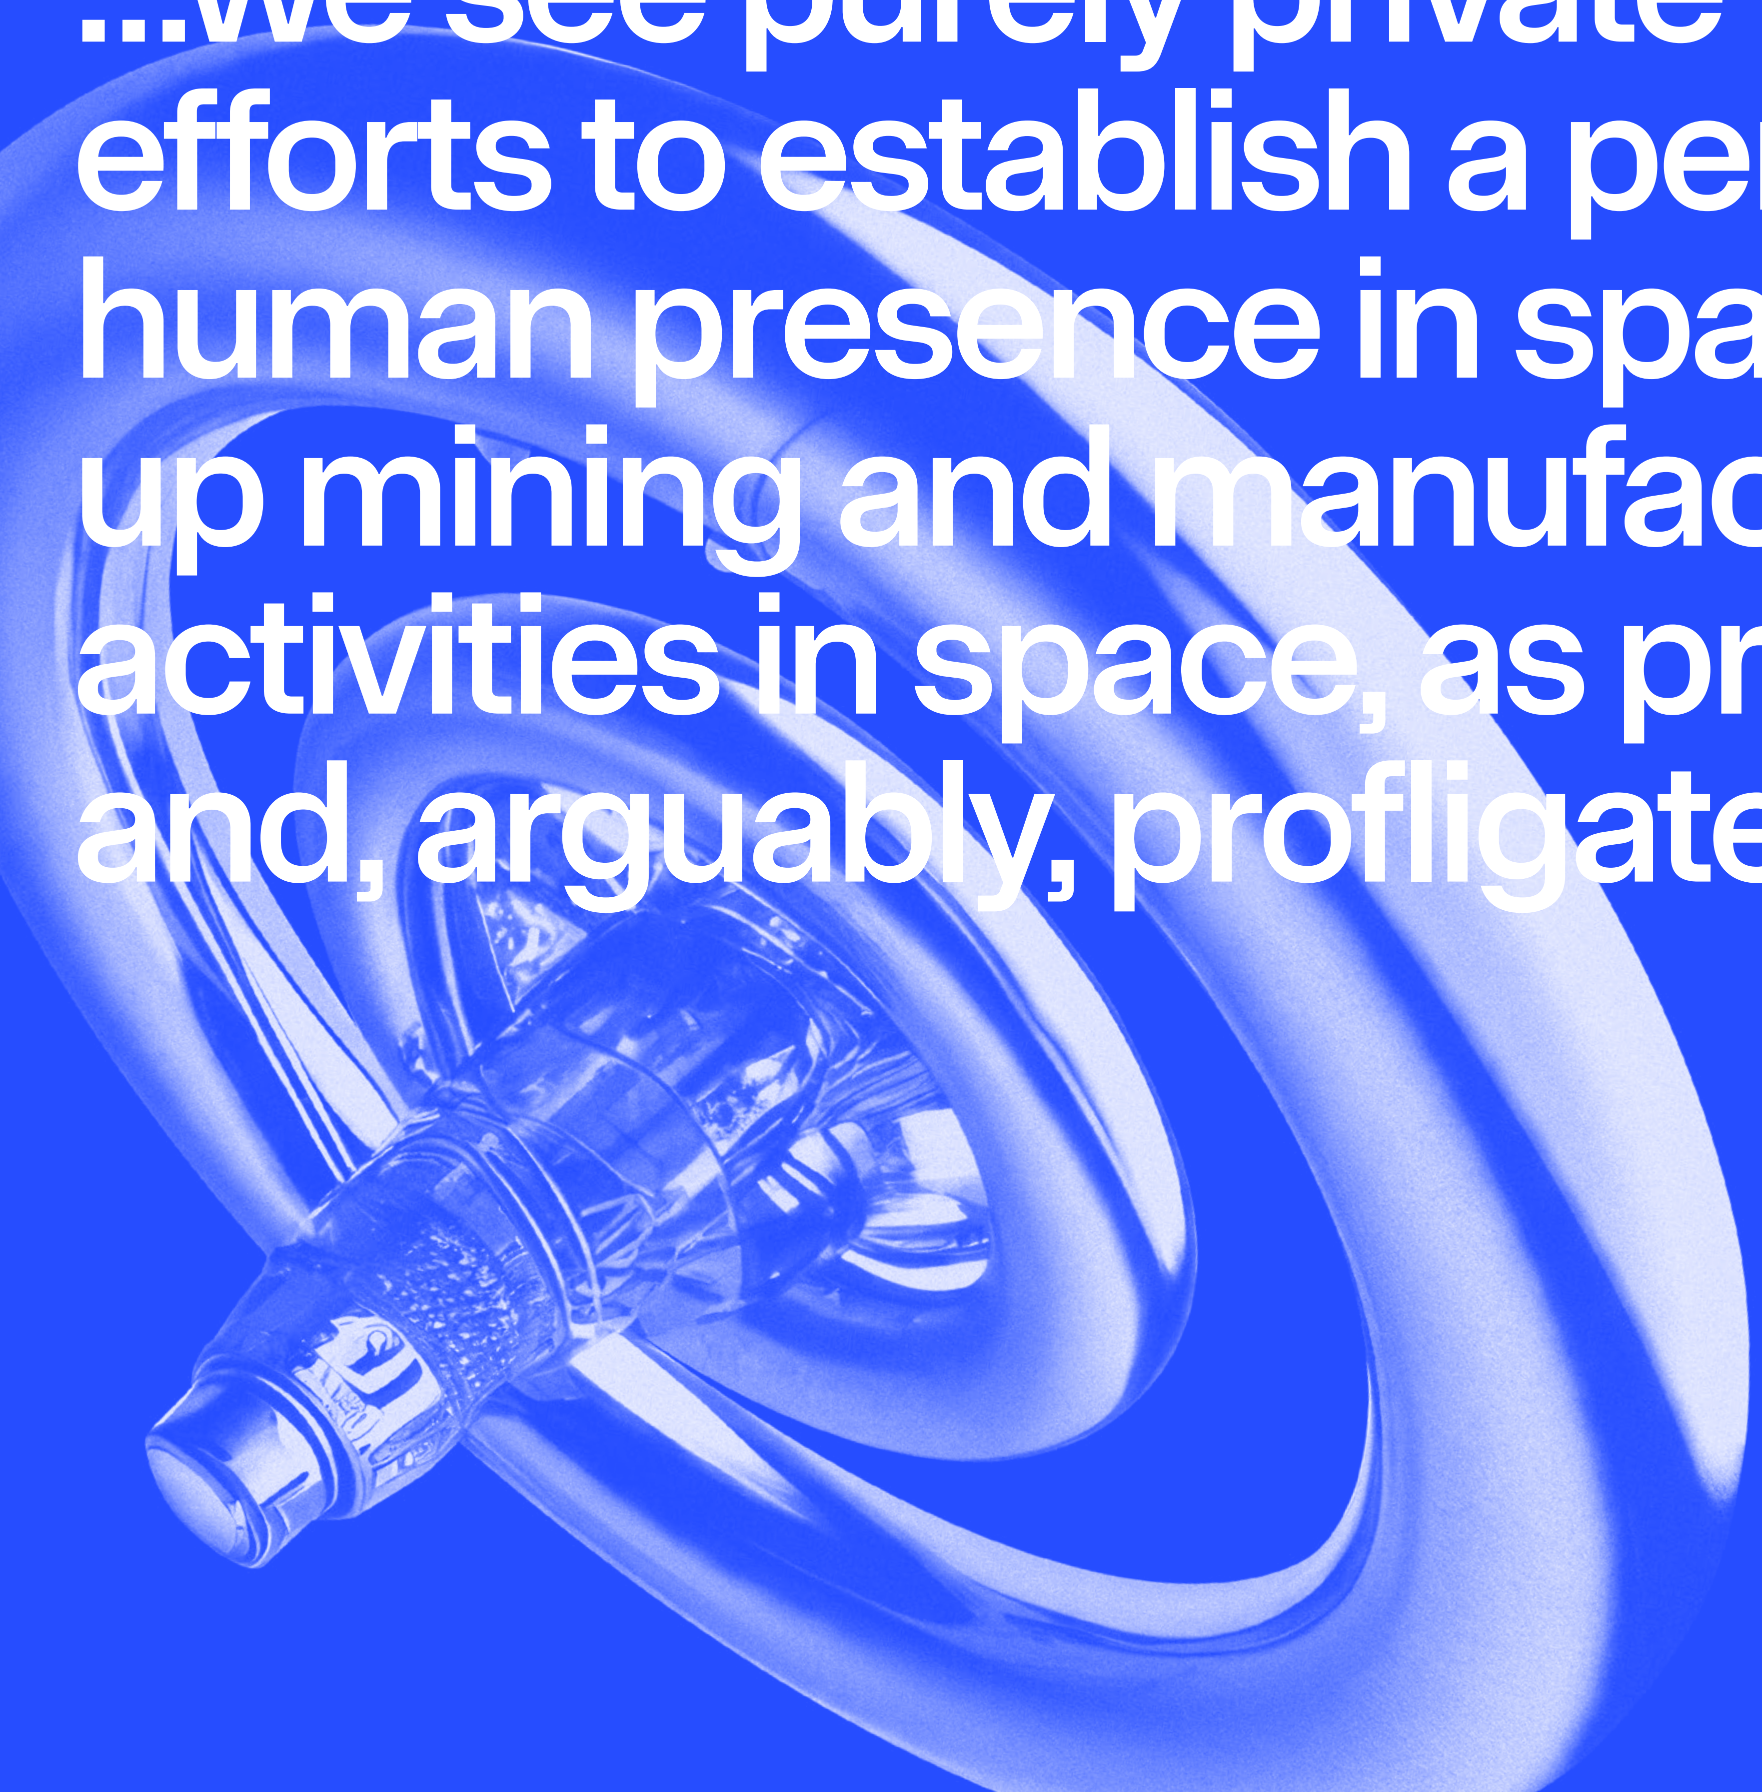

↳ In 2023, a California company called AstroForge orbited a satellite designed to test its technology for refining platinum from asteroids. This year, another startup called Inversion Space will launch a prototype capsule designed to show it can store cargo in space for rapid, on-demand delivery to any location on Earth. Another company called ABOVE Space aims to build a giant, habitable space station with a crew of 400—resembling nothing so much as the wheel-like Space Station V in *2001: A Space Odyssey*—and says it will start this year by launching a boxcar-like, unmanned station called Prometheus. Elon Musk's SpaceX is building its Starship spacecraft and Super Heavy rocket partly in order to take astronauts and refueling equipment to the surface of Mars.

We wish all of these ventures well, and we admire their daring. But right now, most of the things entrepreneurs propose to do in space can be done for far lower cost right here on Earth. (Why launch cargo into space when you can simply deliver it by plane?) Space hotels, Mars bases, asteroid mining, and space-based storage and manufacturing may be part of humanity's long-term future, but at the moment they feel like luxuries. And on a planet racked by climate-related disruption, displacement, scarcity, and conflict, they could become unaffordable. "A true in-space economy, one that thrives in space without going back and forth to Earth, is interesting, but it's probably not a venture bet yet," says DCVC operating partner Matt O'Connell, founder of satellite imaging company GeoEye.

By far the most compelling reason to put things in space, O'Connell believes, is to protect and preserve the planet itself, through real-time observation and communication. If you don't know what's happening on the surface, you can't see storms and other climate threats approaching, and you can't anticipate security threats. If you don't have satellite phones to communicate when cell towers are damaged or jammed, you can't coordinate disaster response or military movements.

Those are the principles guiding DCVC's investments in companies like Planet and Capella Space, which photograph huge swaths of the Earth daily at different frequencies and use machine learning and other forms of AI to spot the differences from one image to the next (see Opportunity 5.1). "The holy grail in Earth observation is change detection," says O'Connell, who chairs Capella's board. If a tank shows up on a battlefield, if a submarine docks at a port, if a storm surge compromises a road network, satellite images will quickly show it.

"I was in the satellite business, and I'm the first guy to say, 'Satellites are cool,'" O'Connell says. "But really, what I care about is the data."

Chapter 6.0

TechMed is where deep tech meets medical care.

By integrating advanced technologies like AI, machine learning, and computer vision with devices, TechMed companies are distinguishing themselves from their MedTech brethren—and are poised to improve patient outcomes worldwide.

● Companies

Caption Health, Proprio, Remedy Robotics, Valar Labs

● Voices

Zachary Bogue, Alan Cohen, Dr. Linda Maxwell

↙ **Proprio**
Proprio uses computer vision and AI to help surgeons make more informed decisions in the operating room.

Better computer-vision, image-processing, and predictive technologies are augmenting the capabilities of medical professionals at all levels of patient care. Thanks to a flood of data about patients themselves—coming from medical images, wearable devices, electronic health records, and genomic and proteomic studies— there's a growing opportunity to bring computing-driven innovations, or what we call TechMed, to all parts of medicine.

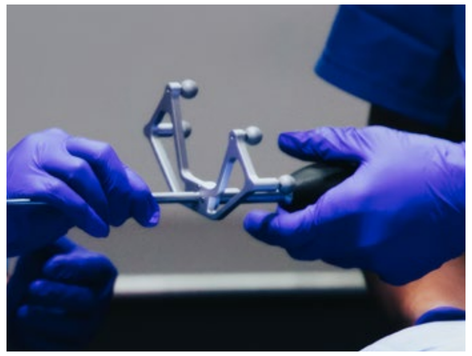

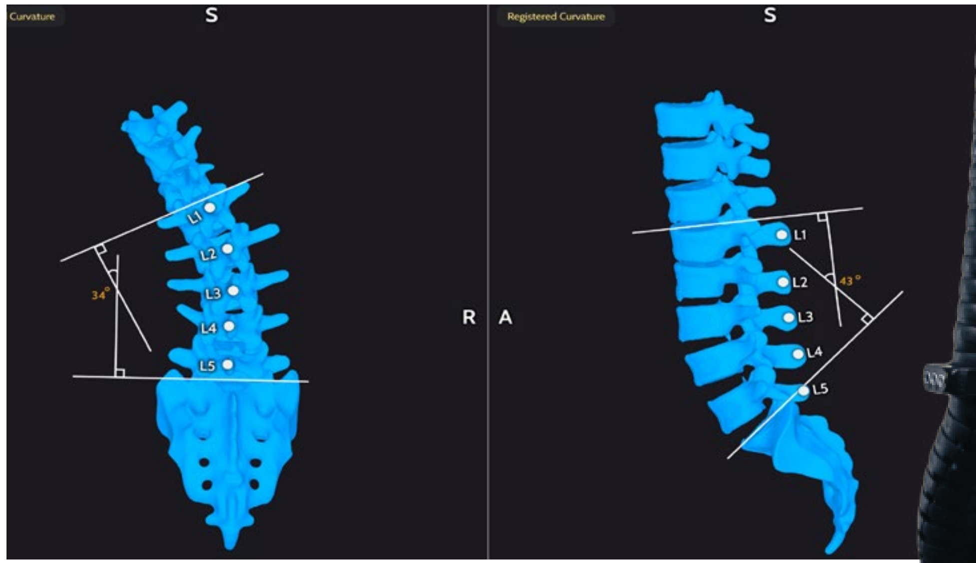

▸ Proprio
Proprio's imaging and guidance system for surgeons uses light-field technology, video cameras, and proprietary algorithms to provide a real-time 3D model of the surgical area, all without radiation.

A commercial airline pilot wouldn't take to the skies without multiple layers of machine assistance at every step—from the digital preflight checklist on her tablet, to the plane's own flight management system, to the conflict-detection system that air traffic controllers rely on to keep planes from colliding in the air, to the autopilot that handles the landing. So why is it still common for a surgeon to operate on your heart, brain, or abdomen guided by little more than a static X-ray, a CT scan, or a laparoscopy camera?

In fact, computing has been slowly infiltrating medicine for decades—for example, in the form of imaging in the 1960s and '70s (CT stands for *computed* tomography) and in electronic health records in the 2010s. But there's still a dearth of technology in healthcare relative to other critical fields, for complex reasons. Medicine is a highly regulated field. Errors can cost lives. Doctors aren't trained as engineers. Government and private insurance systems historically haven't provided clear reimbursement pathways or incentives for spending on new technology. And there are aspects of patient care that can't, or shouldn't, be fully automated.

But today computer vision, AI, robotics, and new kinds of medical devices are finding their way into operating rooms, emergency departments, cancer wards, and even primary care encounters at a faster pace. "No one is saying technology can replace the expert humans who deliver medical care," says DCVC's Bogue. "TechMed companies are building systems that can guide, support, and augment them."

"No one is saying technology can replace the expert humans who deliver medical care. TechMed companies are building systems that can guide, support, and augment them."

Zachary Bogue
Managing Partner, DCVC

Two DCVC-backed companies, for example, are building computerized systems designed to help surgeons save lives and reduce complications. One is called Remedy Robotics, and its mission is to get help to stroke victims faster. When a clot blocks one of the larger blood vessels of the brain, one of the best ways to prevent death and disability is to thread a catheter through the arteries—from the groin all the way to the brain—and remove the blockage; it's called an endovascular thrombectomy. But if a patient doesn't live near a specialized hospital with a trained surgeon, precious hours can go by before the procedure happens, increasing the chances of irreversible brain damage. Remedy is building a robot that can allow expert surgeons to perform endovascular surgery remotely. As Remedy puts it, "No one in the world should suffer an inferior outcome because of where they live."

Another company called Proprio has invented a guidance system for a range of spinal, orthopedic, and cranial operations. When surgeons are correcting spinal deformities—say, fusing two vertebrae together, or attaching rods or other implants—they use preoperative CT scans to plan the operation down to the millimeter. But "no plan of operations extends with any certainty beyond the first encounter with the enemy," as Prussian military leader Helmuth von Moltke wrote.
The difficulty of matching up the static CT images with a patient's actual spinal anatomy during surgery means that the complication and revision rates for spinal deformity surgery are very high.

Proprio has built a multi-eyed sensor for the operating room that uses light-field technology, video cameras, and proprietary algorithms to build a 3D model of the surgical area in real time—showing surgeons how much reality diverges from the preoperative plan, and giving them the data they need to adjust on the fly without having to expose the patient to more X-rays or fluoroscopy. "They've gotten FDA clearance, and they're actually doing human surgeries that are being guided by these new forms of sensing," DCVC general partner Alan Cohen says of Proprio. The data the company is gathering before, during, and after operations could dramatically improve the surgical experience for patients *and* doctors, and greatly reduce complication rates.

DCVC was also proud to back Caption Health, another TechMed company using data to improve the quality of care. Caption, which was acquired by GE HealthCare in 2023, used tens of thousands of recorded ultrasound exams to train software that guides ultrasound technicians, helping them obtain high-quality images even if they don't have extensive training. The idea is to make ultrasound exams such as echocardiograms—which are crucial in the early detection of heart disease—cheaper and more widely available. "Whether it's Remedy, Proprio, or Caption, computing power is dealing with these nearly intractable issues—providing expertise in areas where it wasn't available or increasing the capability of practitioners," Cohen says.

Figure 6.1.1 FDA approvals of AI-enabled medical devices

Manufacturers are marketing—and the FDA is approving—a skyrocketing number of radiology devices and other systems that incorporate machine learning algorithms and other forms of AI.

Source: U.S. Food & Drug Administration

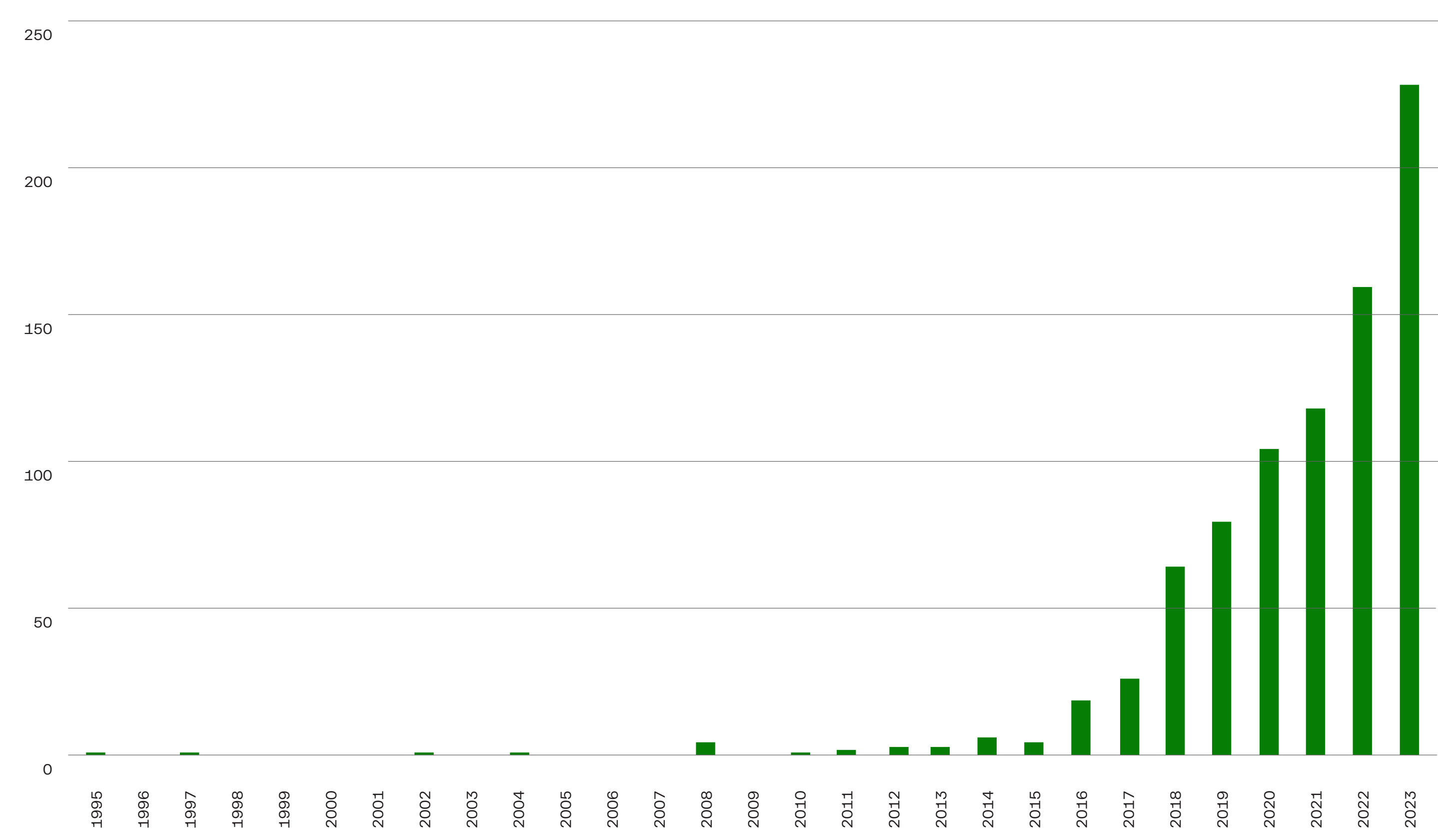

Figure 6.1.2
Guiding precision cancer care

For a few types of cancer, such as colorectal and endometrial cancer, doctors can usually identify genetic mutations that help guide precision treatments. In other cancer types, actionable mutations are rarely found, creating a need for other ways to tailor treatments, such as Valar Labs' AI-driven digital pathology tests.

Source: Journal of Translational Medicine

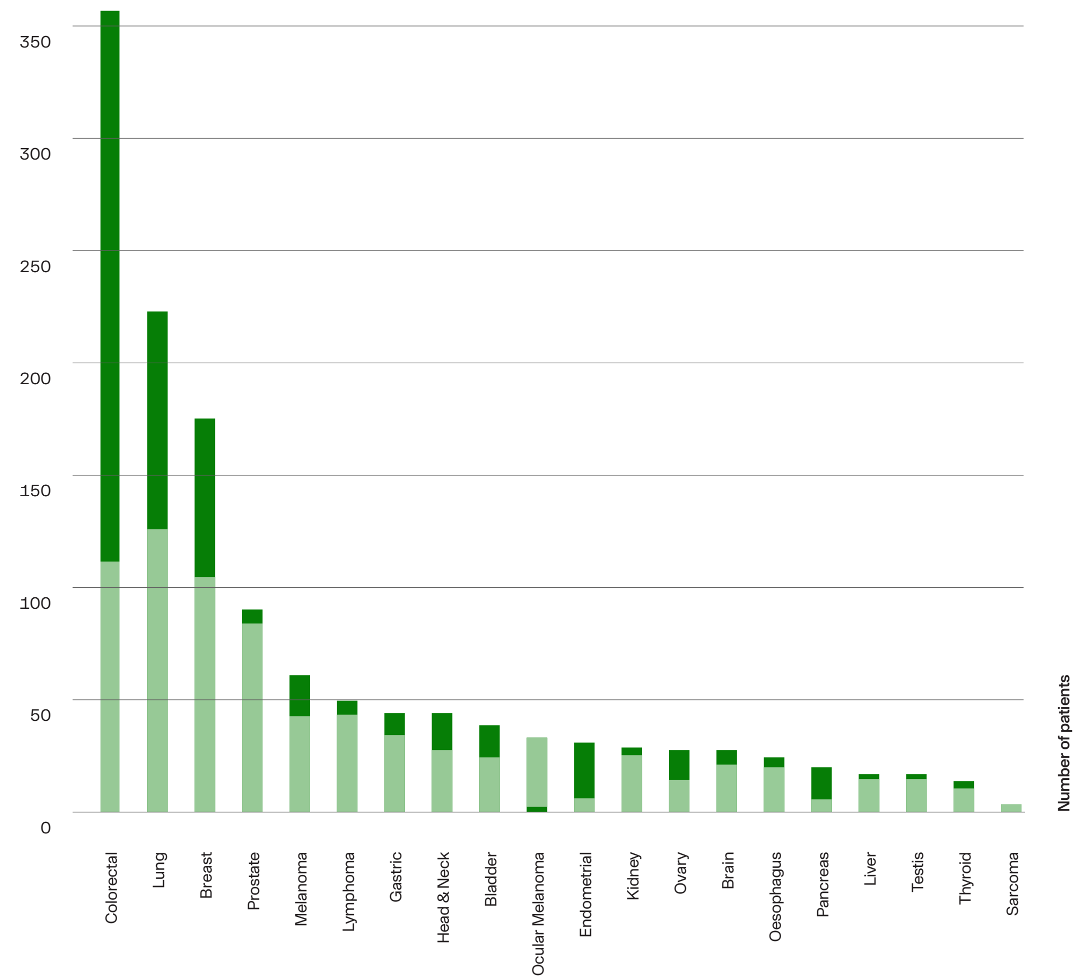

Mutation detected
No mutation detected

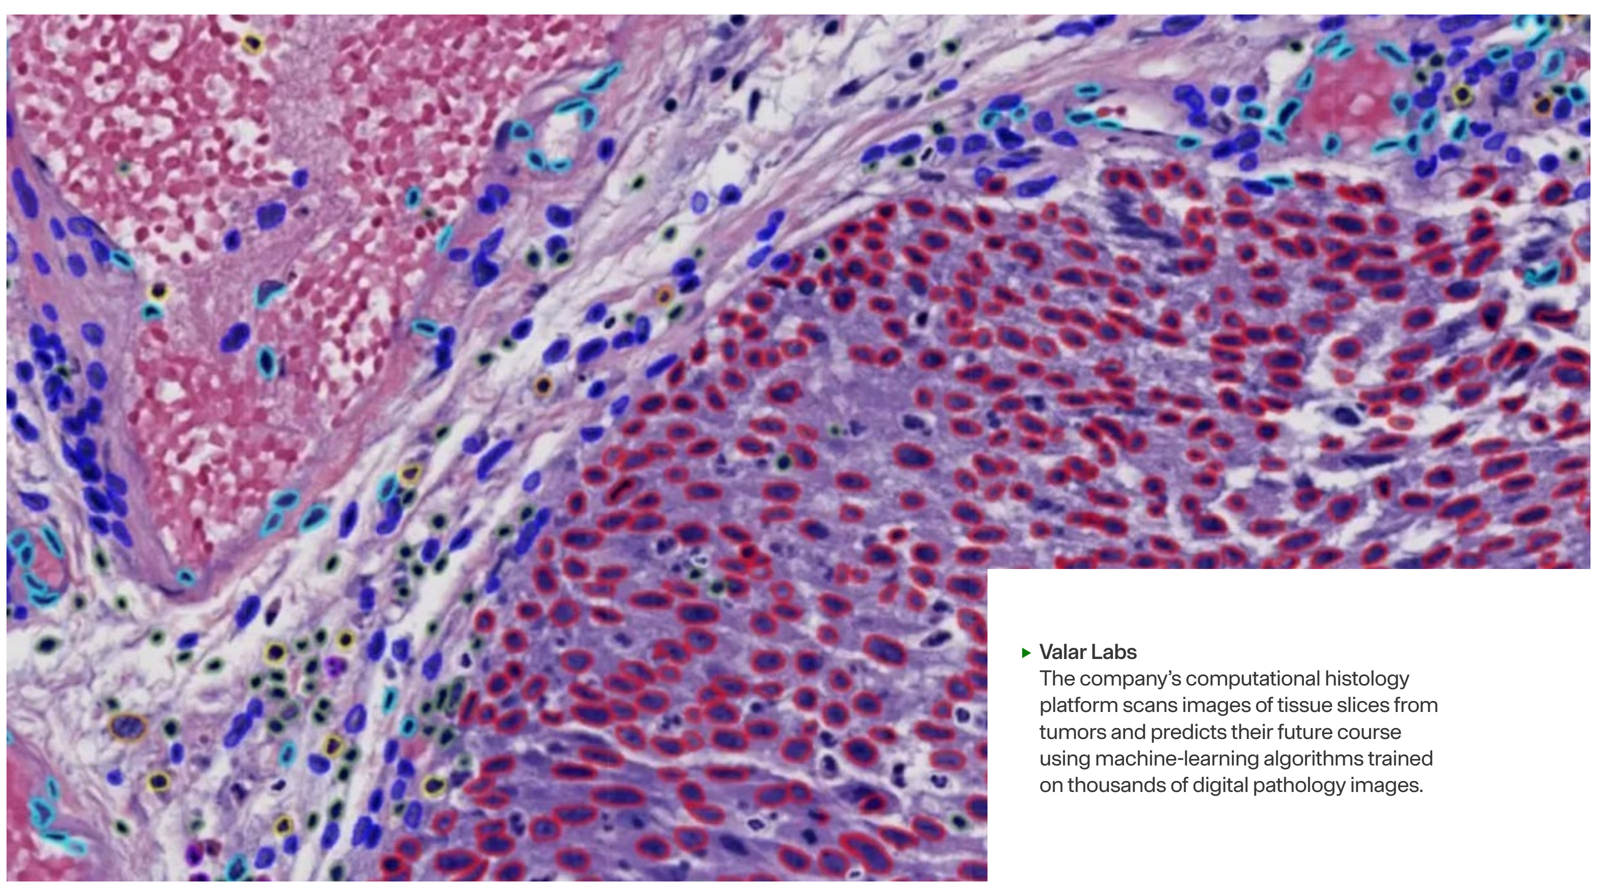

▸ Valar Labs
The company's computational histology platform scans images of tissue slices from tumors and predicts their future course using machine-learning algorithms trained on thousands of digital pathology images.

"We're going to see a new class of digital-native doctors and other practitioners who are empowered and comfortable practicing at the data-driven bleeding edge of medicine and science."

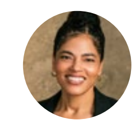

Dr. Linda Maxwell
Operating Partner, DCVC

Another company backed by DCVC, meanwhile, is bringing TechMed's data-centric philosophy to the treatment of cancer. Sometimes, cancer patients have a telltale mutation that steers oncologists toward specific treatment. (People with HER2-positive breast cancer, for example, almost always undergo a combination of surgery, chemotherapy, and treatment with the antibody drug Herceptin.) But 90 percent of patients don't have one of these actionable mutations. Valar Labs is filling that gap by building AI models trained on patient records and digital pathology slides from thousands of cancer patients. When shown new pathology slides, the models can predict how the cancers will progress, and whether tumors will respond to standard first-line treatments.

"We almost serve as the co-pilot for the oncologist," says Valar's CEO, Anirudh Joshi. "There are a lot of 'digital pathology 1.0' companies that build AI to make grading slides faster and more reliable. What we do goes beyond that. We call it 'digital pathology 2.0,' where it's predicting the downstream outcome of treatment, which is something we could never do until AI got more sophisticated."

This spring, Valar introduced its first product, a test to evaluate bladder cancer patients, and collected $22 million in new funding, in a funding round co-led by DCVC and Andreessen Horowitz. Dr. Linda Maxwell, an operating partner at DCVC, says she expects that Valar's technology, and similar AI capabilities, will see enthusiastic uptake, especially from next-generation physicians who expect to be empowered by technology, not burdened by it.

"Most oncologists are general oncologists serving a broad population of cancer patients. They're not necessarily trained in the latest breakthroughs in genomics for specific cancers," Maxwell says. "I think we're going to see a new class of digital-native doctors and other practitioners who are empowered and comfortable practicing at the data-driven bleeding edge of medicine and science."

What's new and exciting about TechMed companies, says Cohen, is that they're not just building medical devices, but "are focused on creating intelligent software systems that can augment and optimize the human technique involved in the practice of medicine itself." Soon these systems will begin to be interconnected, ingesting multimodal data from clinic and hospital environments and providing support across different medical specialties and use cases. At that point, Cohen says, there will be "an opportunity for TechMed companies to develop into multibillion-dollar, all-encompassing platforms monetizing at a much greater scale than the product silos of the legacy medical industry."

Artificial general intelligence is one of the allurements distracting some deep-tech innovators from more practical and achievable goals (see Section 1.5). We think brain-computer interfaces are another shiny object. Early-stage studies of devices that read a patient's neural signals and translate them into actions in the outside world are exciting and appropriate, especially in cases where the reward, such as restoring function for people with brain or spinal injuries or other disorders, might outweigh the medical risks. But…

…talk of neural implants that merely augment normal human abilities feels rash—and we think it will take decades to work out the scientific, ethical, and business challenges involved.

↳ In March 2024, Neuralink broadcast a livestream on X, viewed by more than 19 million people, in which a human patient—a 29-year-old quadriplegic man named Noland Arbaugh—apparently used a wireless brain implant to play chess and activate other cursor-driven functions on his computer. It wasn't the first time a company had used a permanent brain implant to convert brain signals into cursor movements. What was new was the implant itself, which may herald a new era of tiny, unobtrusive brain-computer interfaces (BCIs). Neuralink's 1,024-electrode N1 implant is threaded into place by a surgical robot, is not externally visible, and is recharged wirelessly via an inductive charger.

Neuralink says its mission is to "create a generalized brain interface to restore autonomy to those with unmet medical needs today and unlock human potential tomorrow." We're proponents of the first part of that statement, and skeptics about the second. Arbaugh, who said moving the cursor with his mind was like "using the Force" from *Star Wars*, volunteered to receive the Neuralink implant because of a 2016 diving accident that left him paralyzed below the shoulders. Subjects of other studies have received implants that helped them mentally operate and receive sensory feedback from prosthetic limbs.

But Neuralink's co-founder and principal backer, Elon Musk, has also spoken of building brain-computer interfaces that would allow healthy people to operate their smartphones or summon their Teslas telepathically, and even says he wants to test an implant on himself. Meanwhile, the idea that humans might one day be able to control machines with the power of thought is filtering out into a range of non-medical fields. The Defense Advanced Research Projects Agency's Next-Generation Nonsurgical Neurotechnology (N3) program, for example, aims to develop brain-machine interfaces that would allow able-bodied soldiers to control unmanned aerial vehicles or team up with computers "to successfully multitask during complex military missions."

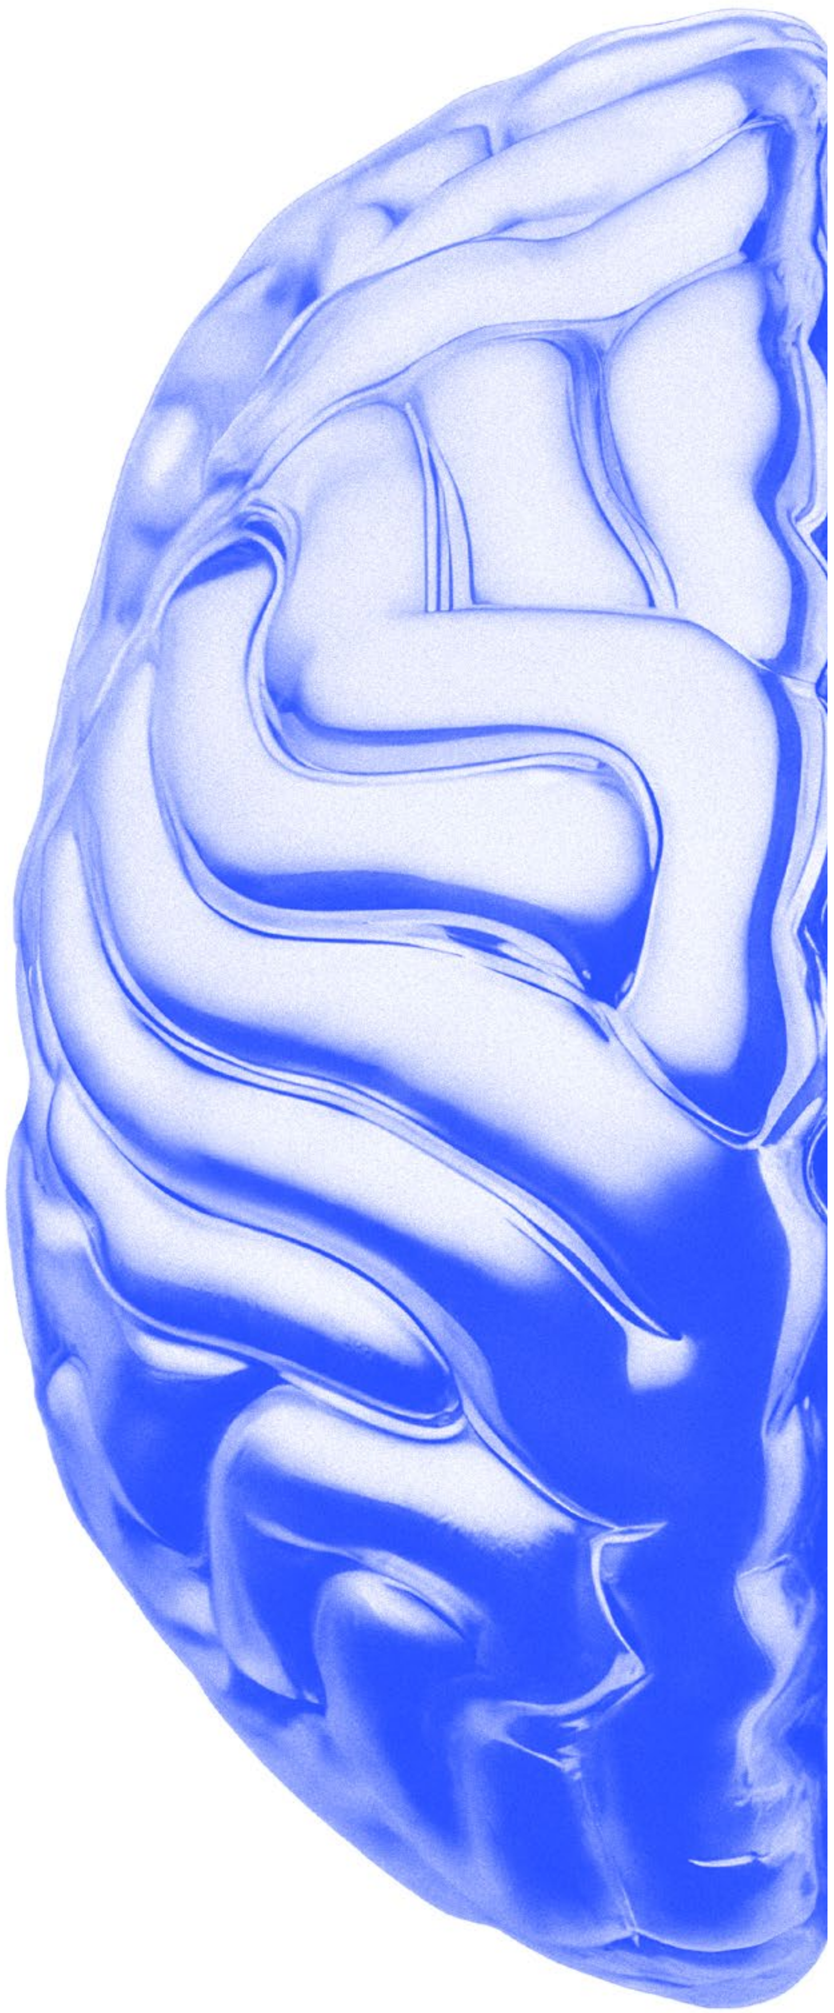

We think that's getting ahead of the game. In the Neuralink demo video, electrodes monitoring a small subset of neurons in Arbaugh's motor cortex picked up the brain signals that occurred when he imagined an intended action, like swiping a touchscreen or moving a mouse. Software learned to interpret these signals and translate them into computer commands. That's certainly impressive technology, and we're sure it will get better. But scientists' understanding of the deeper workings of the mind remains so primitive that more sophisticated applications of brain-computer interfaces, like piloting a UAV, will take many years to operationalize, let alone commercialize. "I believe there is deep value and human benefit to be unlocked by brain-computer interface technology," says Dr. Linda Maxwell, an operating partner at DCVC with years of experience as a surgeon. "However, there is still a great deal of research to be done before BCIs will be able to interpret and act out the neural correlates of consciousness."

"I believe there is deep value and human benefit to be unlocked by brain-computer interface technology. However, there is still a great deal of research to be done before BCIs will be able to interpret and act out the neural correlates of consciousness."

Dr. Linda Maxwell
Operating Partner, DCVC

Meanwhile, surgical implantation is a highly invasive procedure with a risk of complications. (Neuralink said in May that many of the threads connecting its implant to Arbaugh's motor cortex had already wriggled free.) The idea of implanting a BCI to allow a subject to operate a mobile device "may be appealing at first blush, but must be considered in light of the attendant short-term risks as well as the associated and yet-to-be-articulated long-term risks of such devices," Maxwell says.

Again, helping people with motor disabilities is a valuable goal. But before we build cyborg soldiers, give healthy people productivity-enhancing cognitive prosthetics, or, in Neuralink's words, "expand how we experience the world," we should engage in a thorough, culture-wide discussion about the implications for personal safety, autonomy, and privacy. Meanwhile, we'd like to see company founders focusing on nearer-term problems that will be prerequisites for building the next generation of brain-computer interfaces—such as innovations that would help us understand, describe, and quantify human cognition. It's becoming possible, for example, to analyze ocular scans using computer vision and machine learning to diagnose mental health conditions such as PTSD. We can also measure brain activity up close without implanting invasive Neuralink-style implants: an Australian company called Synchron has developed a sensor-equipped stent for the blood vessels of the brain. When placed near the motor cortex, it allows patients with neurodegenerative diseases or quadriplegia to operate computers. To build better BCIs and ultimately support safe human use, we'll need many engineering advances like these—along with a lot of old-fashioned scientific puzzle-solving.

Chapter 7.0

Reinventing the science of drug discovery and agriculture

Transformative progress in today's pharmaceutical and agricultural worlds comes from companies that use computation and new lab techniques to translate basic science insights into platforms that generate a myriad of product leads and predict which ones are winners. Look at the companies in the DCVC Bio portfolio, and you'll find world-leading scientific teams building data-driven businesses that are helping patients with unmet needs.

Elo Life engineered new varieties of watermelons and sugar beets that carry zero-calorie sweetener molecules called mogrosides. See Opportunity 7.4.

● Companies

Alchemab, CH4 Global, Chroma Medicine, Creyon Bio, Elo Life, Empirico, Nium, Pivot Bio, Solu Therapeutics, Umoja Biopharma

● Voices

Dr. John Hamer, Dr. Kiersten Stead

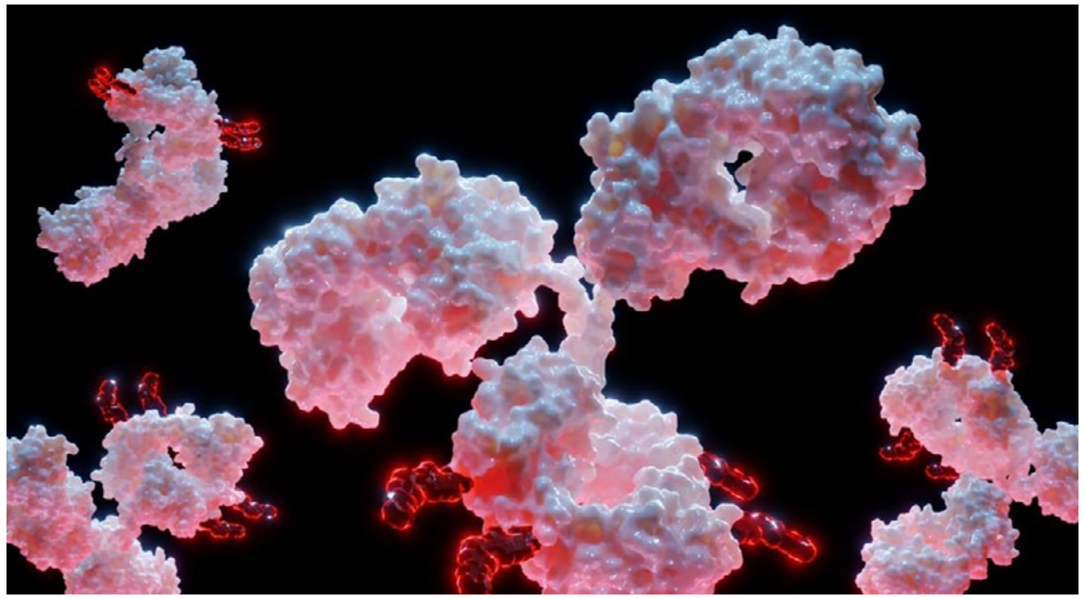

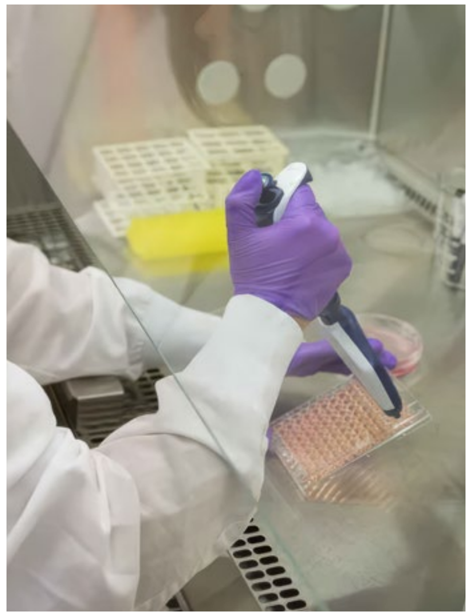

↳ DCVC Bio is different from most other biotechnology venture funds in that we look to invest not in specific drugs or molecules, but in scientific breakthroughs by top innovators, plus computation-based platforms that could generate whole new arrays of treatments or products. For most biotech investors, "the starting point is a drug that's in the clinic, and they want to make a bet that a drug that's in Phase 2 will make it to Phase 3 and then make it to approval," observes Dr. John Hamer, one of DCVC Bio's four managing partners, along with Dr. Kiersten Stead, Zachary Bogue, and Matt Ocko. "But that's not what we do. We tend to bet on science and teams, because if we back the right scientific approach, we can break into whole new areas of medicine. We can produce not just one drug but a whole range of drugs across different indications."

One of the keys to unlocking new approaches to treating disease is a heavy reliance on data and computation, especially machine learning and other forms of AI. We're excited by teams who bring both a unique biological insight *and* the ability to generate enormous amounts of data around that insight. The algorithms these companies build help them sort through the data to find the most promising drug candidates or agricultural products. Only then—having significantly increased their own chances of success—do they proceed to the clinic or the farm.

In Chapter 2 of this report, we talked about DCVC's investments in TechBio, the emerging sector of life-science companies defined by their commitment to building huge biological and chemical datasets to fuel drug discovery. In this section, we'll discuss four more areas where novel science insights, often backed by powerful computational platforms, are accelerating drug discovery and therapeutics.

A drug hunter's fantasy molecule would combine the best qualities of protein drugs, which bind very selectively to their target cells or tissues, and small molecules, which are easier to make, manage, and deliver. In a sense, that's what a new DCVC Bio portfolio company called Solu Therapeutics is creating.

While biologic drugs such as monoclonal antibodies can recruit the body's own immune cells to help achieve their goals, they have their downsides. They're usually difficult and expensive to manufacture, are temperature-sensitive (and therefore difficult to package and transport), and are poorly absorbed by the body. Meanwhile, small molecules can travel to more places in the body, get inside cells, and can be packaged in pill form. But they can also bind to unexpected targets, meaning they often have unwanted side effects—and they often don't last very long in the body.

When Brandon Turunen was head of U.S. chemical biology at drug giant GlaxoSmithKline, he proposed cooking up hybrid drugs that could recruit the patient's own immune system, like biologics, but bind to many targets, like small molecules. His central idea was to design a two-sided small molecule that could bind to a cell-surface target receptor on one side and to an engineered antibody on the other. He called this new chemical platform CyTaC, for Cytotoxicity Targeting Chimera—and hoped it would translate into new treatments for cancer, inflammation, and other disorders.

GSK ultimately opted not to take the idea to the clinic. But it licensed the technology to Solu, a Boston-based startup founded by Longwood Fund, with additional financing from us at DCVC Bio. Now, with Turunen as co-founder and chief technology officer, Solu is moving toward Phase 1 clinical trials of its first drug candidates.

"What's really exciting is that they can change the pharmacology of a small molecule to give it the pharmacology of an antibody," Hamer says. "Once they're bound to a target, antibodies can trigger a number of different events to destroy it—for example, by attracting killer cells that engulf and phagocytize it. So, lots of interesting things can be done." For example, CyTaCs can bind to targets that were previously off-limits to biologics, including ion channels and G-protein-coupled receptors, one of the most ubiquitous types of cell-surface receptor molecules. (More than 800 human genes, or about 4 percent of the protein-coding genome, code for GPCRs.)

With such a large choice of antibody-binding spots, drug developers could, in theory, expose almost any type of cell, including tumor cells, to the immune system's wrath. CyTaCs also have a longer half-life in the body than small molecules, meaning their effects will last longer. "Solu is combining what we've learned in computational biology and protein engineering to come up with a modular functionality, and be able to drug things that we've never been able to drug before," says Stead.

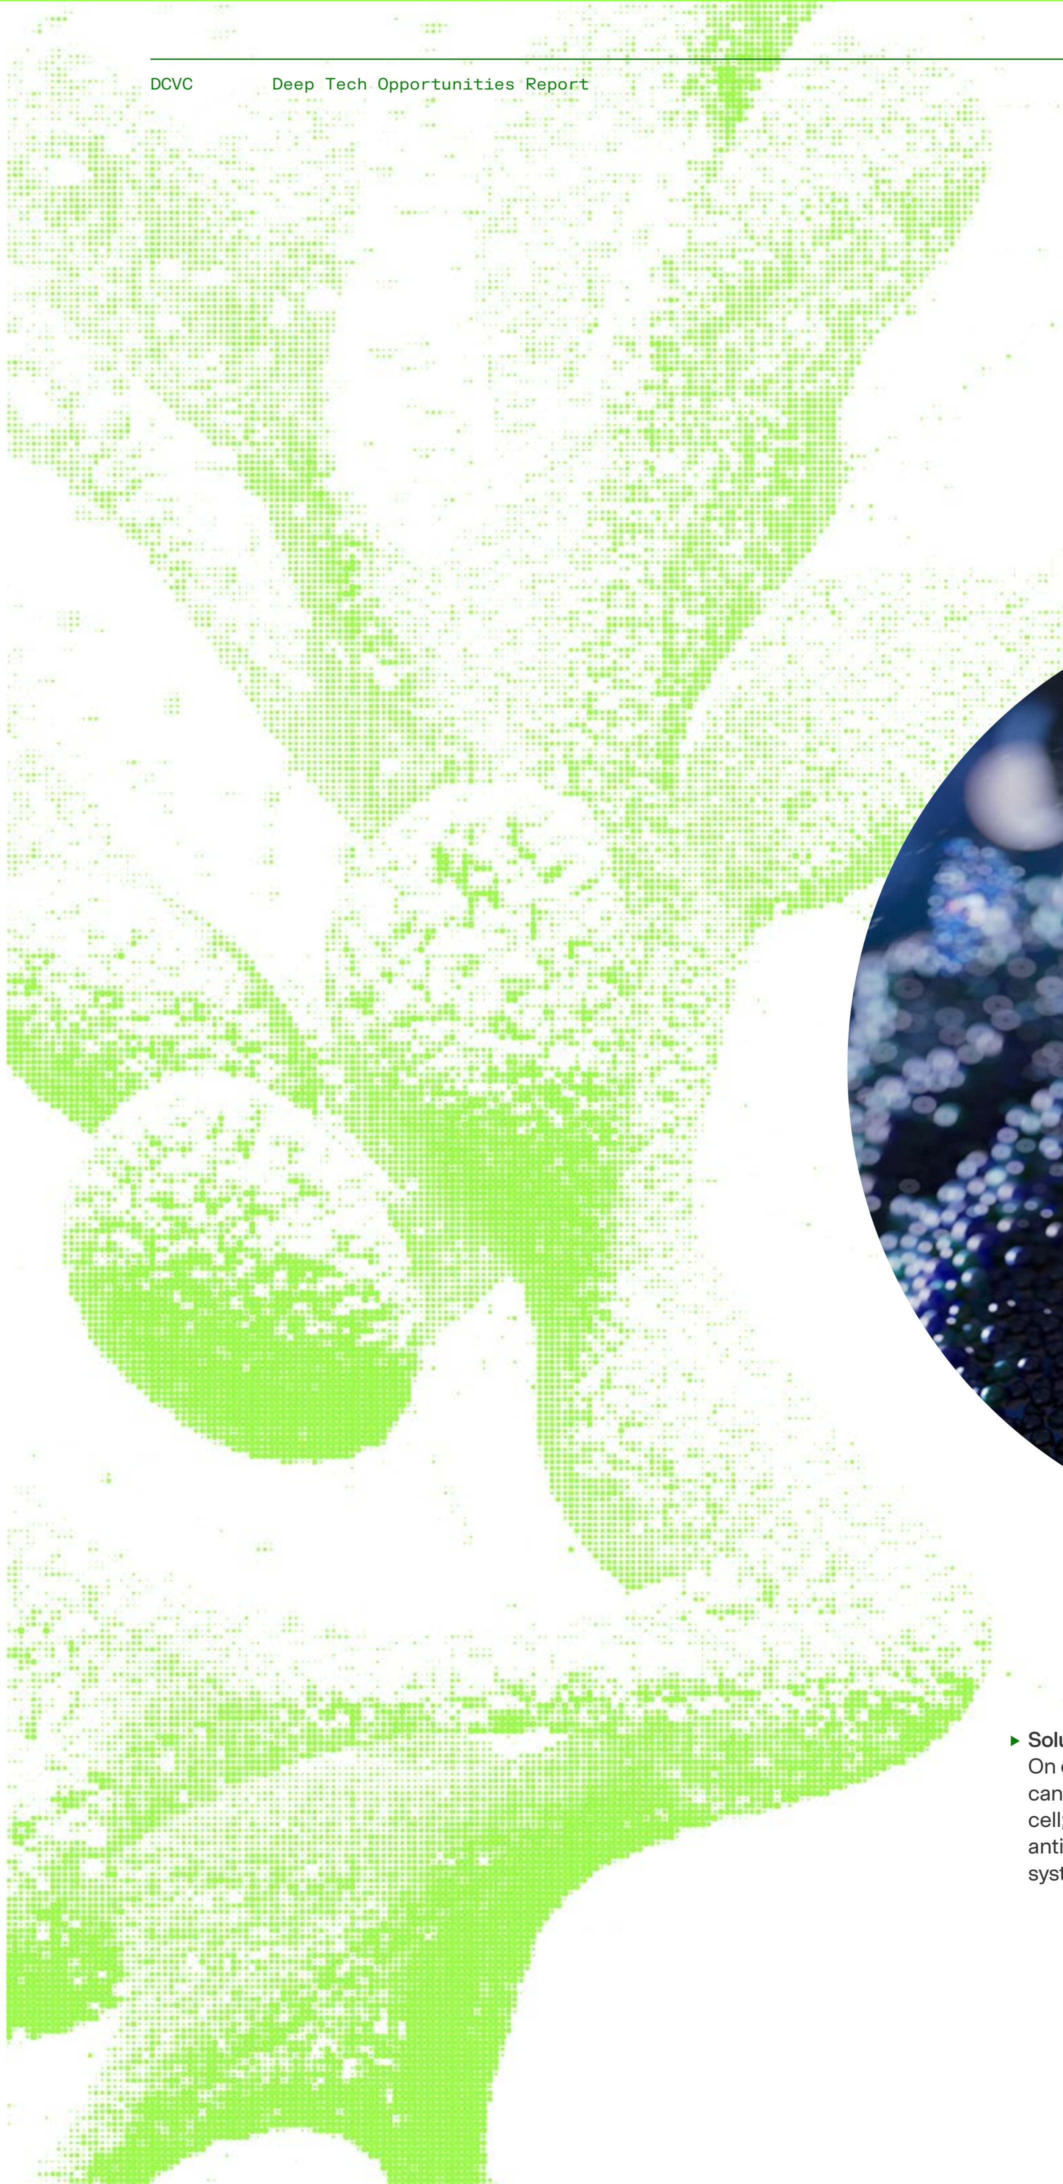

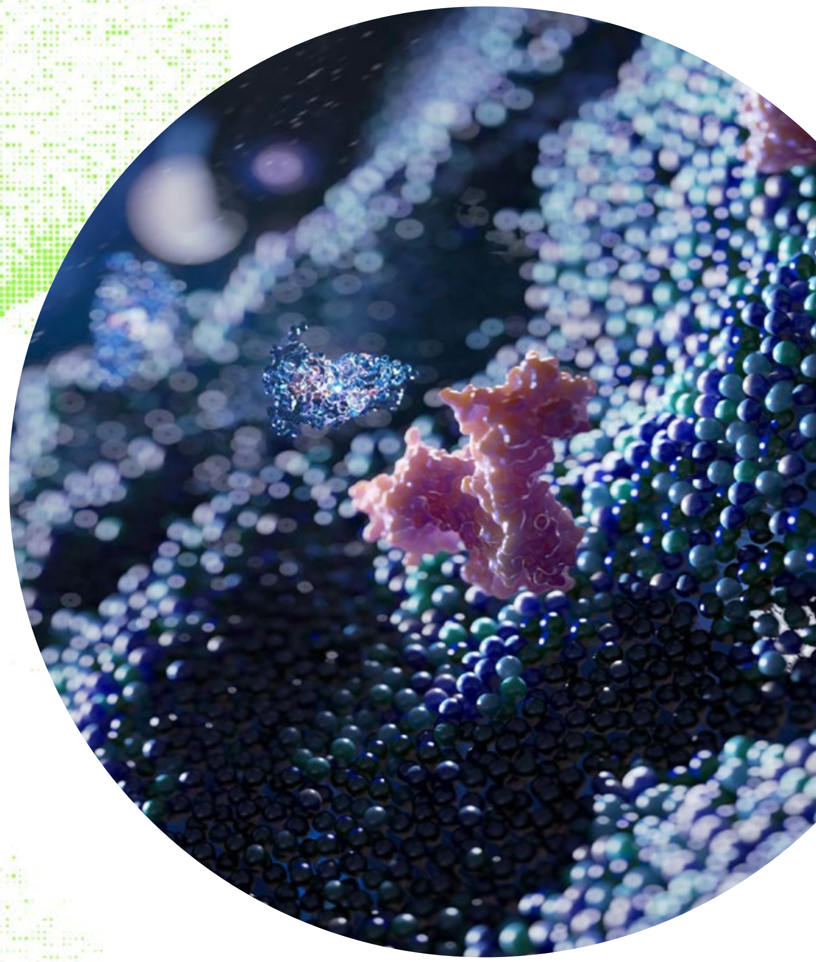

Solu Therapeutics
On one end, the company's hybrid "CyTaC" molecules can bind to a receptor molecule on the surface of a tumor cell; on the other end, they can bind to an engineered antibody intended to recruit the patient's own immune system to attack the cell.

FINDING NEW WAYS TO REPROGRAM IMMUNE CELLS

→ Solu is going after cancer and other conditions by giving small molecules some of the superpowers of protein drugs. Another approach to reinventing cancer immunotherapy is to reprogram the body's own white blood cells, called T lymphocytes, to attack and kill specific tumor cells. That's already being done *ex vivo*, by extracting patients' stem cells and modifying them in the lab, but a DCVC Bio portfolio company called Umoja Biopharma believes it's better done *in vivo*, inside the body.

Figure 7.2.1 Finding new ways to reprogram immune cells

DCVC-backed Umoja has a way to transform immune cells into tumor killers, inside the body.

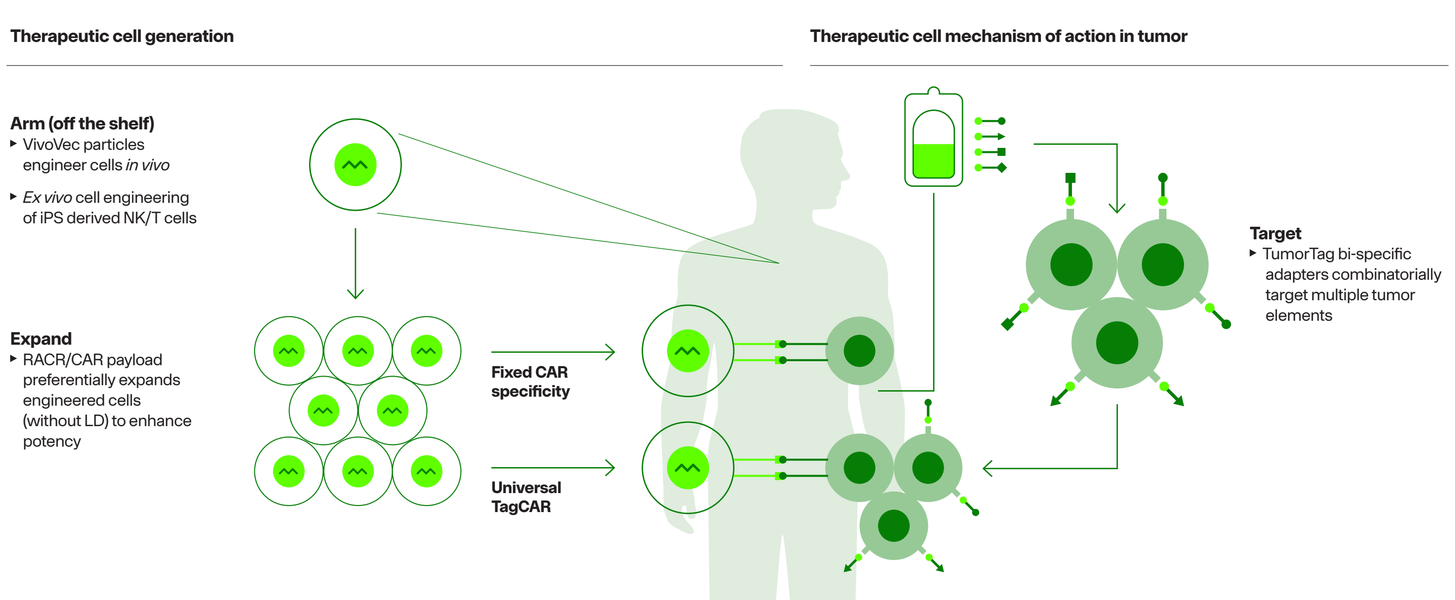

▸ **iPS** induced pluripotent stem cells ▸ **NK/T** natural killer/T cells ▸ **LD** lymphodepletion ▸ **RACR/CAR** rapamycin activated cytokine receptor/chimeric antigen receptor

In the 2023 Deep Tech Opportunities Report we wrote extensively about Umoja's work to deliver custom genes to a cancer patient's T cells (see pages 72–74 of that report). Since then, the company has made regular progress, and it has just received FDA clearance to begin recruiting for Phase 1 clinical trials this year.

Umoja's cellular engineering technique builds on the success of an older approach to cancer treatment called CAR-T therapy. CAR-T involves extracting a patient's T cells and inserting DNA that encodes a desired chimeric antigen receptor (CAR), that is, a receptor protein that recognizes antigens found on the surfaces of cancer cells. Once the cells are placed back in the patient's body, they have the ability to seek and destroy the cancer cells.

CAR-T therapy is highly effective, but it can also be debilitating, not to mention costly. Expanding the population of modified cells can take weeks. Before the cells can be transfused back into the patient, the patient's own remaining white blood cells must be destroyed through chemotherapy, a process called lymphodepletion. Introducing the new T cells can provoke a systemic inflammatory response called cytokine release syndrome, requiring additional treatment—adding to the astronomical base cost of CAR-T (more than $400,000 per infusion).

Umoja's approach is to enable the body to make its own CAR-T cells. It invented an efficient viral delivery system, modeled after the human immunodeficiency virus, that puts new genes directly into T cells inside the body, with no extraction, transfection, expansion, or lymphodepletion. The goal is to help cancer patients beat back cancer, especially hard-to-treat solid tumors, with less time lag and fewer side effects.

Late last year, Umoja announced that the technology successfully generates CAR-T cells in both mice and non-human primates, and that treated mice engrafted with human tumor tissue saw their tumors shrink. These promising preclinical results helped the company land a lucrative deal under which pharma giant AbbVie will pay up to $1.44 billion for the right to market Umoja's first five cell therapy drugs.

"Before the end of the year, we'll have treated the first patients, which will be a big, big breakthrough," Hamer says. Over the longer term, Hamer and Stead say, *in vivo* T-cell engineering could completely replace older, more laborious forms of cancer therapy meant to restart or redirect the immune system, including both allogeneic and autologous stem cell transplants. Tens of thousands of American cancer patients receive stem cell transplants every year, and almost all of them undergo lymphodepletion first. By harnessing the body's own immune system, Umoja aims to reduce this burden and provide broader access to lifesaving forms of cancer therapy.

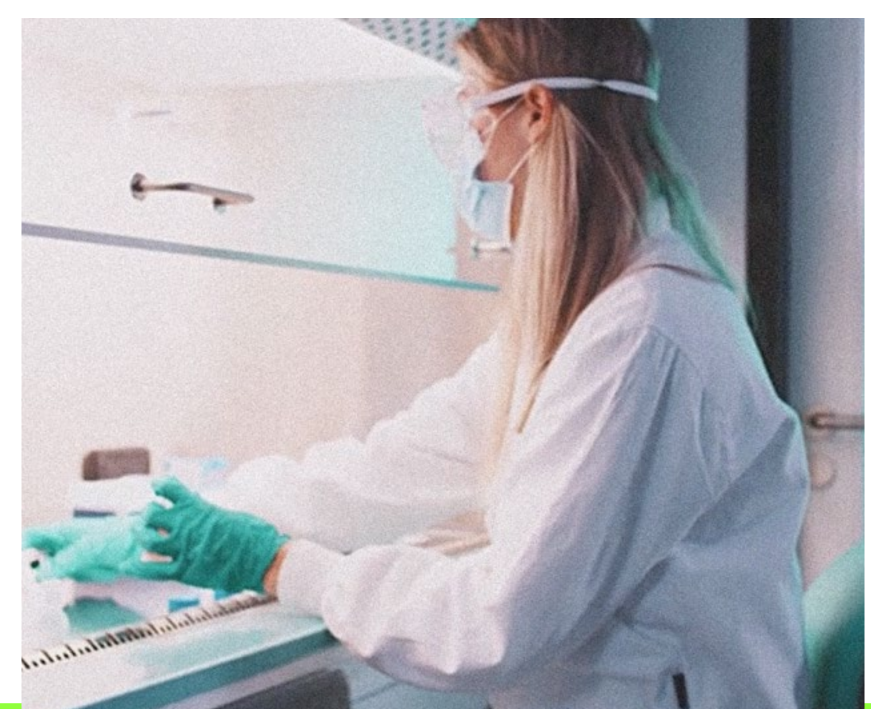

▸ **Umoja**
The company appropriated a delivery system used by the human immunodeficiency virus to put new genes into human T cells that enable them to seek out and destroy cancer cells.

The DCVC Bio companies we've been talking about so far are giving small-molecule drugs the properties of antibody drugs, or reprogramming immune cells themselves at the genetic level to attack new targets. Empirico and Creyon Bio, both based in the San Diego area, are going even deeper into the machinery of the cell, using computation to find new ways to intervene at the level of DNA transcription, RNA splicing, and RNA transcription and degradation.

Many genetic diseases are the result of small, often single-nucleotide mutations in human DNA that result in the production of defective proteins. Sometimes it's possible to block the messenger RNA molecules these genes encode, and prevent them from being translated into proteins, using short segments of complementary RNA called small interfering RNAs (siRNA). The challenge is knowing which diseases are the result of these small mutations; which genes can be interrupted this way; and which siRNAs will target the rogue mRNA most effectively.

Empirico is using a combination of big data, computation, and wet-lab work to crack all of these problems. The company starts at the level of genetics, combing through the health records of millions of patients and using machine-learning algorithms to see which genetic varations seem to be associated with health problems. That helps Empirico's scientists make educated guesses about which gene-protein pathways are decisive in a given disease, and where intervening with an siRNA might work. Then the company synthesizes and tests many siRNAs in high-throughput screens, feeding the data back into its machine-learning platform to predict which of the molecules will prove to be safe and effective in actual patients.

The appeal of siRNA drugs is that they could potentially slow diseases with no known cures, by targeting gene-protein pathways in places where nobody had thought to intervene. So far Empirico's techniques have yielded six candidate medicines. One of them—an siRNA that knocks out a gene target in the liver in order to treat a common, devastating, but as yet undisclosed disease—could move into Phase 1 clinical trials this year, with the FDA's approval. The molecule interrupts the disease through "a completely novel mode of action," Hamer says. "Empirico's big thing is drugging targets that have never been drugged before."

About 25 miles up the coast from Empirico is Creyon Bio, another DCVC Bio portfolio company exploring smarter ways to design siRNAs, antisense oligonucleotides (ASOs), and other oligonucleotide-based medicines (OBMs, which are usually short stretches of RNA or DNA). As noted above, most OBMs are designed to interrupt or modify gene-protein translation in patients with genetic diseases. But that's not a simple proposition: given a known mutation, the number of possible OBMs that could reduce the expression of a gene or fix an RNA splicing problem is still astronomical. "Most of the time, the drug is either inactive or toxic," says Creyon CEO Chris Hart.

So the company put most of its seed funding into building a comprehensive, proprietary database showing what happens to human cells and systems when they're exposed to OBM-like compounds. It uses the data to train machine-learning models to identify the design rules and engineering principles behind OBMs. These rules, in turn, help predict whether a given OBM will do its intended job, without unintended side effects. "What Creyon has really put forward is a process for creating the right dataset to leverage best-in-class AI to extract the design rules and principles that then allow you to engineer knowably safe and effective compounds," Hart summarizes.

Creyon's ultimate goal is to speed up the development of OBM-based drugs and make it possible to treat the thousands of rare genetic diseases for which there's no available drug. This spring, Creyon shared encouraging results in a study of one baby born with a mutated copy of the gene for a receptor protein called Transportin-2. Normally, this protein helps move molecular cargo between the nucleus and the cytoplasm of cells, but the mutated gene produced a toxic variant that resulted in epileptic seizures and developmental delay. Creyon used its models to design a custom antisense oligonucleotide that would safely bind to and block the messenger RNA encoding the mutant Transportin-2 receptor.

> "Empirico's big thing is drugging targets that have never been drugged before... [their siRNA molecule for knocking out a gene target in the liver uses] a completely novel mode of action."

Dr. John Hamer
Managing Partner, DCVC Bio

After just two doses, the baby's epileptic seizures dramatically slowed, and after a third dose "he regained developmental milestones that were lost and developed new skills," according to an abstract of research findings the company presented in May. The whole project, from initiation to first dose, took just one year. Says Stead: "That is transformational, because this was a drug that cost hundreds of thousands of dollars to develop, not billions."

Figure 7.3.1 Creyon's arsenal of precision-engineered drugs

In the cell nucleus, the genes in DNA are transcribed into RNA molecules, which are then used as the templates to build proteins. Creyon's oligonucleotide-based medicines allow it to intervene at any point in that process.

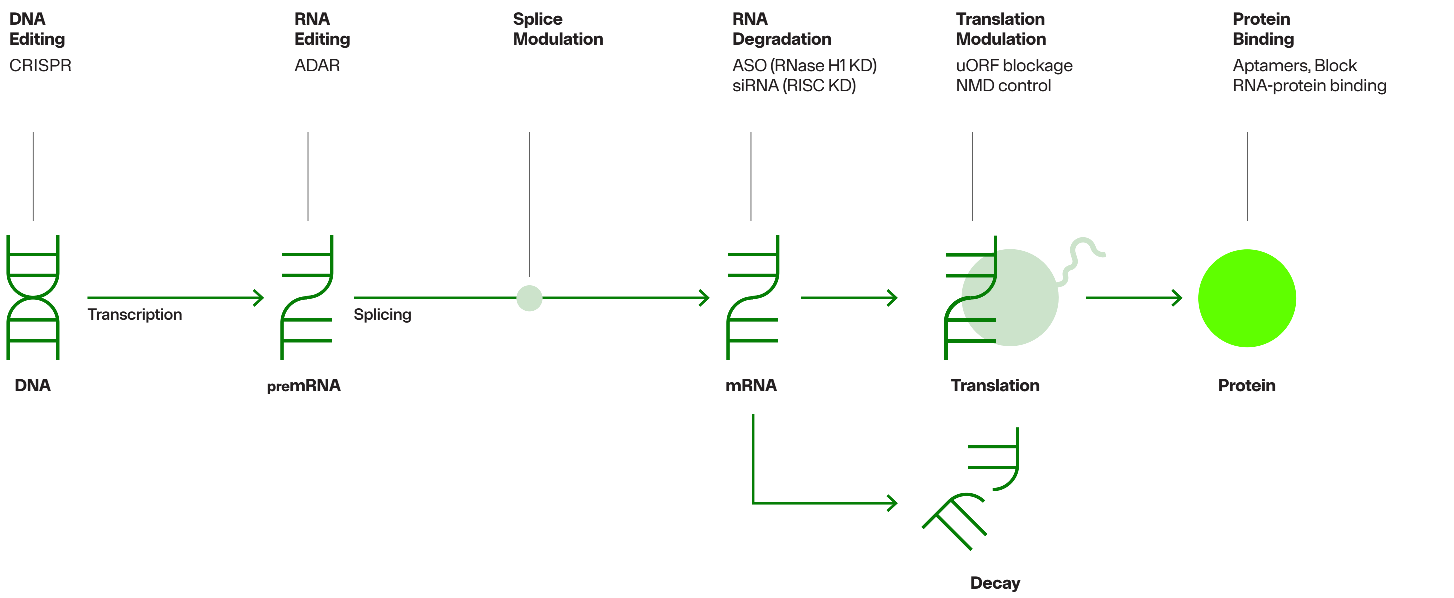

▸ **CRISPR** clustered regularly interspaced short palindromic repeats ▸ **ADAR** adenosine deaminases acting on RNA ▸ **ASO** antisense oligonucleotide ▸ **siRNA** short interfering RNA ▸ **uORF** upstream open reading frame ▸ **NMD** nonsense-mediated decay

One final way to prevent the expression of a mutated gene is to keep the cell's gene-transcription machinery from reaching the gene at all. That's becoming possible thanks to new techniques for precisely manipulating a process called DNA methylation.

In 2012, the development of the CRISPR-Cas9 system—genetic scissors borrowed from bacteria and archaea—gave biologists the ability to cut DNA in specific places. That made it far easier for researchers to knock entire genes out of a cell's genome. Later, researchers figured out how to knock other genes back in, and even how to edit the individual nucleotide bases inside genes. Yet while these techniques can be immensely powerful in the lab, and even in humans, they can also have unwanted off-target effects, cutting DNA in places researchers didn't anticipate.

But there are other, more reliable and more flexible ways to modify gene expression. One is to harness DNA methylation—the technique being perfected in preclinical work by a DCVC Bio company called Chroma Medicine.

When DNA is methylated, it's tightly wound up around spool proteins called histones, which puts the genes on that stretch of DNA out of reach. When the methyl groups are removed, the DNA unspools and can be accessed by the RNA polymerase molecules responsible for gene transcription. (Histones and DNA together are known as chromatin, hence the company's name.)

Chroma engineered enzymes that can target specific DNA sequences and either deactivate them through methylation or activate them through demethylation. The Chroma technology does all this without cutting or damaging the DNA itself. And while the effect is long-lasting, it's also completely reversible. "It's the next evolution of genome editing," says Stead.

Figure 7.4.1 A system for silencing and activating genes

Chroma built enzymes that can demethylate "silent" DNA that's spooled up into histones, thus exposing it to the cell's transcription machinery, or the reverse, methylating it to isolate it from that machinery.

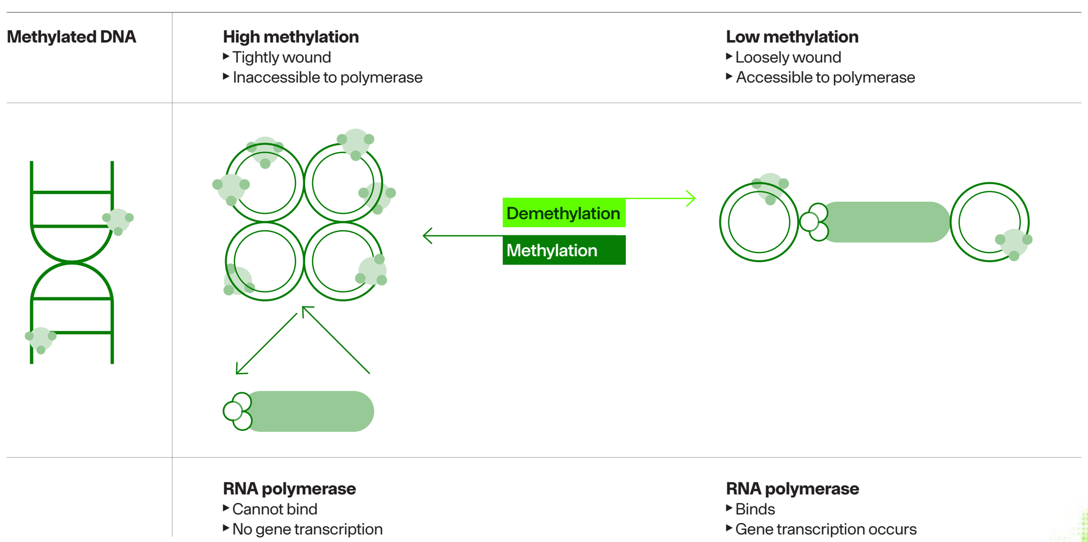

Now Chroma is putting its gene silencing system to work to go after the world's leading killer, cardiovascular disease. Humans carry many genes responsible for trafficking and storing lipids—a level of redundancy that served us well in our deep evolutionary past. But in the context of our modern lifestyles, it can predispose us to high cholesterol and lipid levels, one of the biggest risk factors for atherosclerosis.

It turns out we can live without some of these genes. For example, there's a gene called *PCSK9* that encodes an enzyme that prevents the breakdown of LDL cholesterol in the liver. If the gene is absent or artificially blocked, LDL is digested more effectively and LDL levels in the bloodstream go down. In fact, that's exactly how two costly cholesterol-lowering biologic drugs, Praluent and Repatha, work. "People who have knocked-out *PCSK9* live quite nicely and also have lower chance of getting cardiovascular disease," Stead explains.

Chroma announced in late 2023 that it can block the enzyme in a new way: by demethylating and silencing *PCSK9*. In studies of mice, a silencing molecule carried into the liver inside a lipid nanoparticle lowered *PCSK9* expression by 98 percent, and the effect lasted throughout the 10-month study period. Nonhuman primates showed a similar response.

Since *PCSK9* is already a safe and validated target, it makes sense for Chroma to pursue it as one of the first applications of its epigenome-editing technology, "because you don't necessarily want to pair a new modality with a risky gene," Stead says. "But then they have a whole pipeline of other things to go after."

One target is the hepatitis B virus, a devastating microbe that causes cirrhosis and liver cancer, impacts nearly 300 million people worldwide, and is almost impossible to eliminate from the body because of the unusual way the virus hijacks host RNA to replicate itself. "There are many, many copies of it inside of your liver cells, and so people have been trying to figure out how can we get rid of all this virus that's accumulating in the liver," Hamer says. "Once you stop taking your antiviral drugs, it just comes roaring back." Adds Stead: "A really good way to get rid of that is to just methylate it and turn it off. It's the very definition of precision genomics."

> "A really good way to get rid of [Hepatitis B] is just to methylate [the viral DNA] and turn it off. It's the very definition of precision genomics."

Dr. Kiersten Stead
Managing Partner, DCVC Bio

The developed world's food system was designed during and immediately after World War II—an era when the predominant challenge, for most people in most countries, was getting enough calories, protein, and vitamins and minerals every day. Now, as a side effect of the Green Revolution and the commoditization of food—itself a miraculous achievement—we have the opposite problem. Today's typical diets include so much high-calorie, high-fat food that they're causing widespread metabolic and cardiovascular disease—while the agribusiness sector producing all this food generates enormous amounts of CO_2 and methane (CH_4). Modern farming and eating patterns can't be changed overnight, but we can begin to use deep tech to design ways around both the calorie problem and the carbon problem.

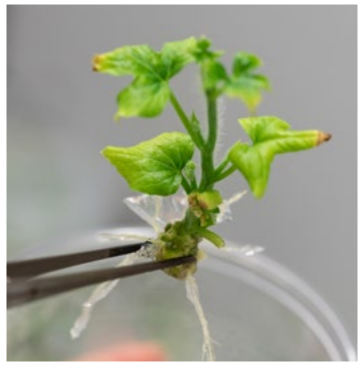

▸ Elo Life
The mogrosides in a rare Asian plant called monk fruit are 300 times sweeter than cane sugar. Elo developed watermelons and sugar beets that make monk fruit mogrosides, creating a natural, zero-calorie sweetener.

When we talk about deep-tech answers to our fraught relations with food, we're not making the facile argument that every problem created by technology can be fixed with more technology. Rather, we're conscious that "Just choose healthier foods" isn't always an actionable piece of advice for people already dealing with diabetes, obesity, or heart disease, or whose eating habits are deeply ingrained in their culture or environment. To overcome the problems created by our existing agricultural and food production system, innovators are intervening on at least three fronts.

The first is the human body itself. GLP-1 receptor agonists such as semaglutide and tirzepatide are proving to be astonishingly powerful tools for reducing high blood sugar in people with diabetes and reducing appetite in people prone to obesity. Now we need to make these drugs more affordable and accessible (see Opportunity 7.6). And we've already mentioned another form of treatment that works at an even deeper level to activate or deactivate specific genes involved in metabolism. One of the first drugs in the pipeline at DCVC Bio portfolio company Chroma Medicine is designed to reduce the amount of LDL ("bad" cholesterol) in the bloodstream by blocking expression of a fat-transport gene called *PCSK9*. Such drugs could be especially helpful for people whose cholesterol levels don't go down in response to dietary changes or statins. Two biologic drugs already approved for this purpose, Praluent and Repatha, have been shown to lower LDL levels by 40–60 percent (see Opportunity 7.4).

The second front is food production. We will forever crave sweet, fatty, and salty foods, as they are tied to our biology. However, it's possible to engineer foods that provide, for example, the sweetness consumers crave while helping them reduce their intake of refined sugars. Researchers at Elo Life, a North Carolina company backed by DCVC Bio, knew about an Asian plant, the

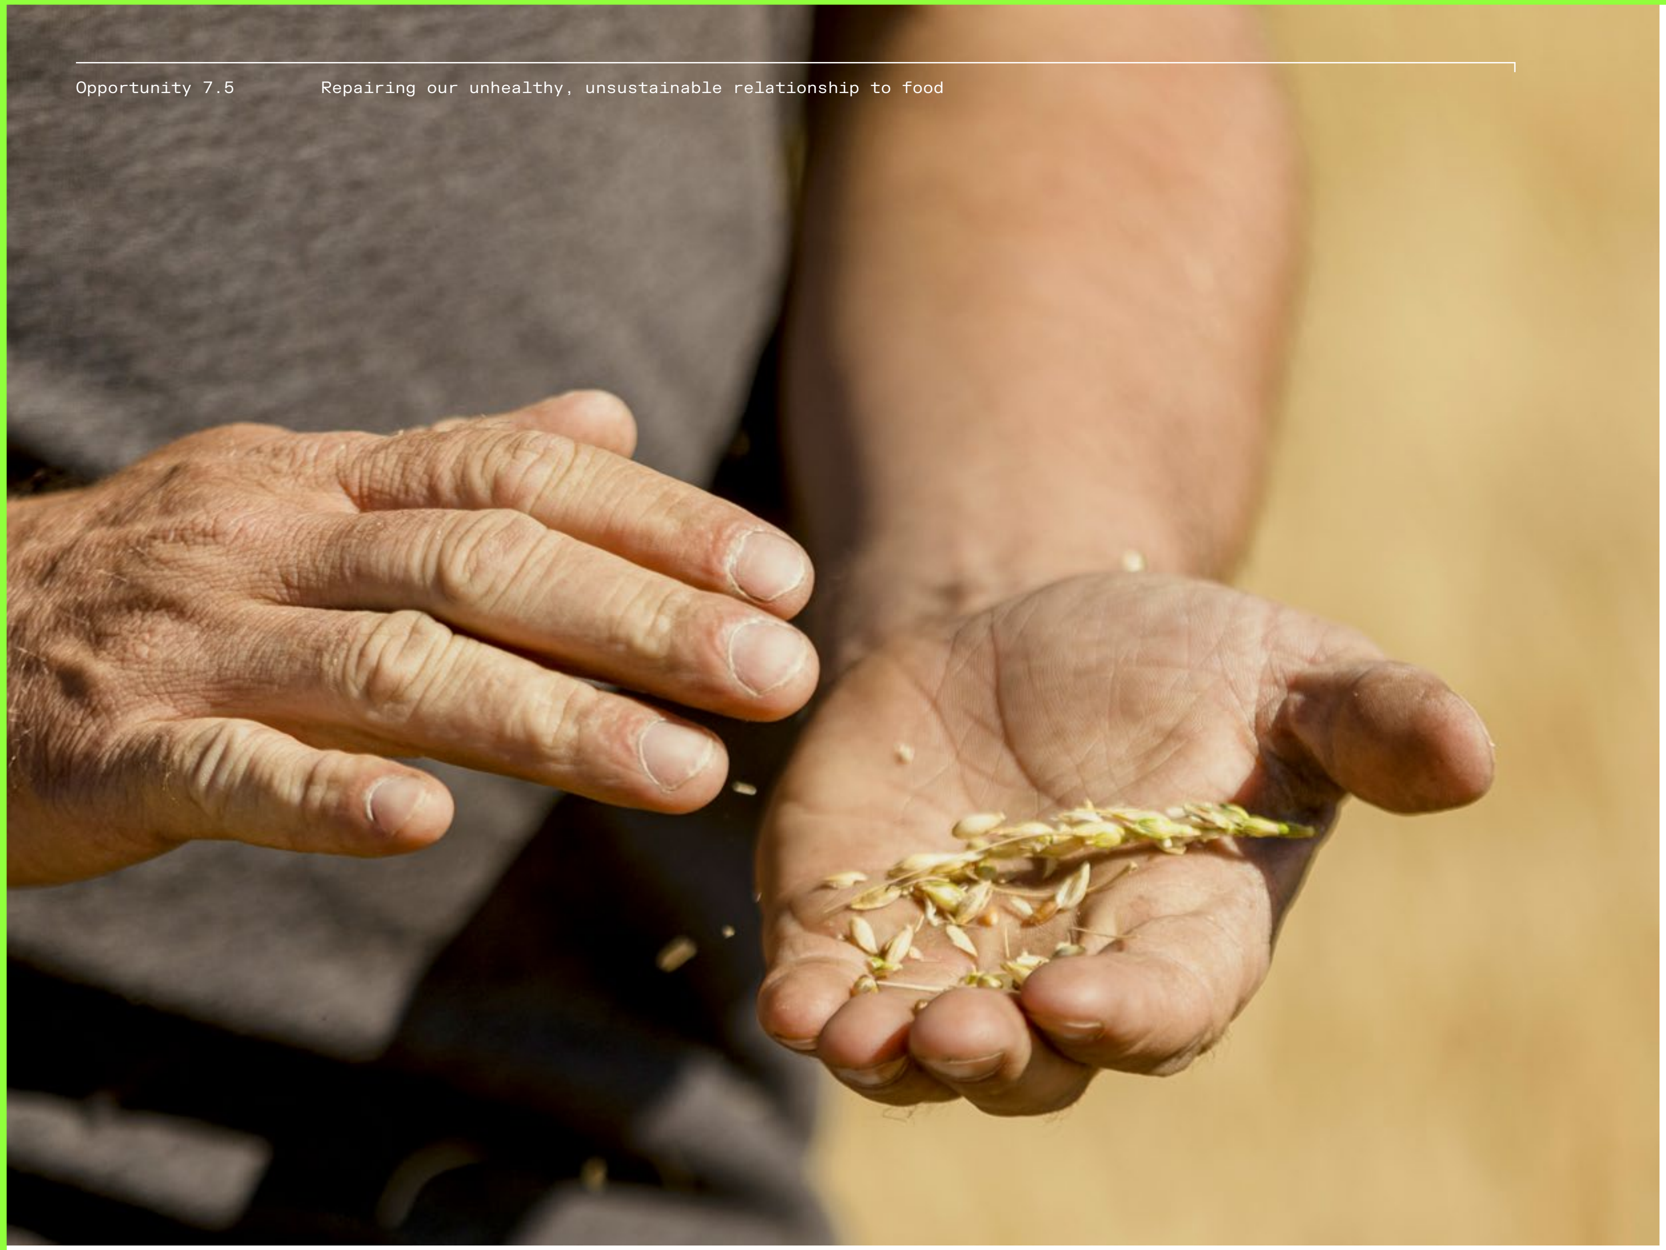

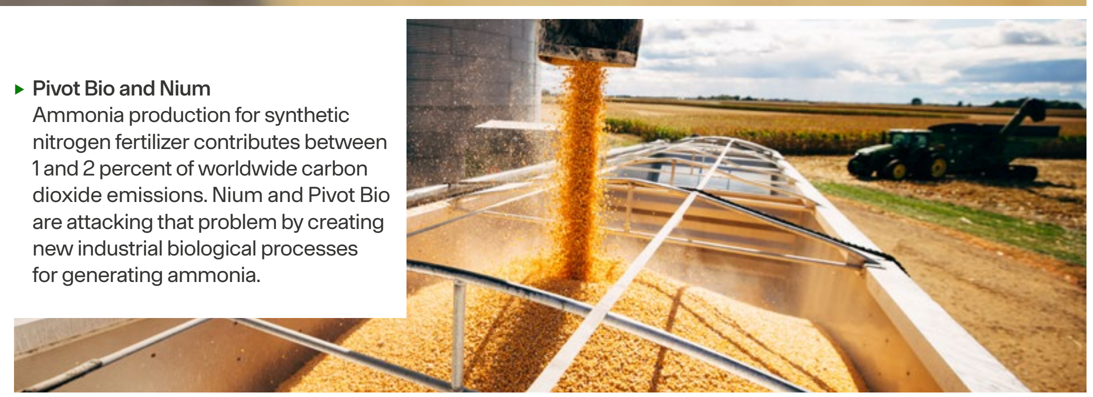

▸ **Pivot Bio and Nium**
Ammonia production for synthetic nitrogen fertilizer contributes between 1 and 2 percent of worldwide carbon dioxide emissions. Nium and Pivot Bio are attacking that problem by creating new industrial biological processes for generating ammonia.

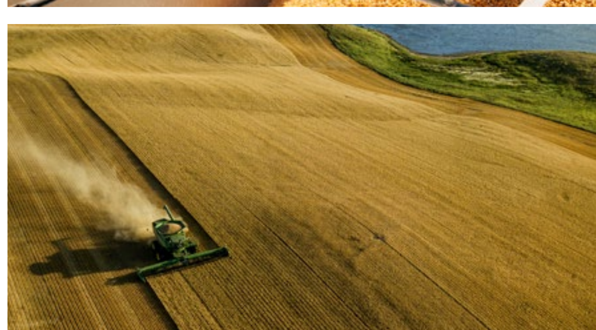

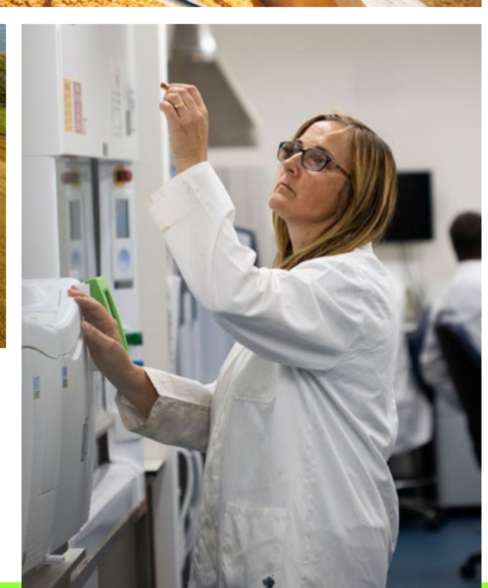

monk fruit, whose juice contains natural, zero-calorie chemicals called mogrosides that are 300 times sweeter than sugar. The problem was that monk fruit is grown only in remote valleys of southern China and northern Thailand. Elo figured out how to grow watermelons and sugar beets that manufacture their own mogrosides. Because these crops can be grown almost anywhere, Elo now has a convenient supply of natural sweeteners that can replace sugar and artificial sweeteners in thousands of foods and beverages. The company plans to bring its monk fruit–derived sweetener to market in 2026 and will have samples in late 2024. “If we can reduce sugar in people’s diets and get to a place where you get all the sweetness, but none of the calories and the other harmful effects that sugar causes, that would be the single most important thing we could do to improve human health,” Elo CEO Todd Rands has said.

The third front for deep-tech intervention is at the level of the specific agricultural practices responsible for our food system’s huge carbon footprint. (About 23 percent of all human-induced greenhouse gas emissions come from agriculture, according to the U.N. Environmental Program.) As we explained in Opportunity 3.2, the DCVC portfolio companies Nium and Pivot Bio are working to reduce CO_2 emissions from fertilizer manufacturing by coming up with new industrial and biological processes for creating fertilizer’s main ingredient, ammonia. Meanwhile, DCVC and DCVC Bio have invested in CH4 Global, which is producing a cattle-feed supplement that virtually eliminates methane production in the guts of beef cattle and dairy cows—a major contributor to agriculture’s greenhouse gas emissions.

Healthy bodies, healthy eating, and a healthy climate are interrelated problems—so it will take multiple approaches, like those being pursued by our portfolio companies at DCVC and DCVC Bio, to solve them. “CH4 Global is getting some really fantastic data,” says Stead. “Obesity is a curable disease now that you can modify how much people want to eat. We can modify how much fat their blood is carrying, with the *PCSK9* knockouts. The other piece that needs solving is sugar intake itself. And we’ve got companies working in all of these spaces.”

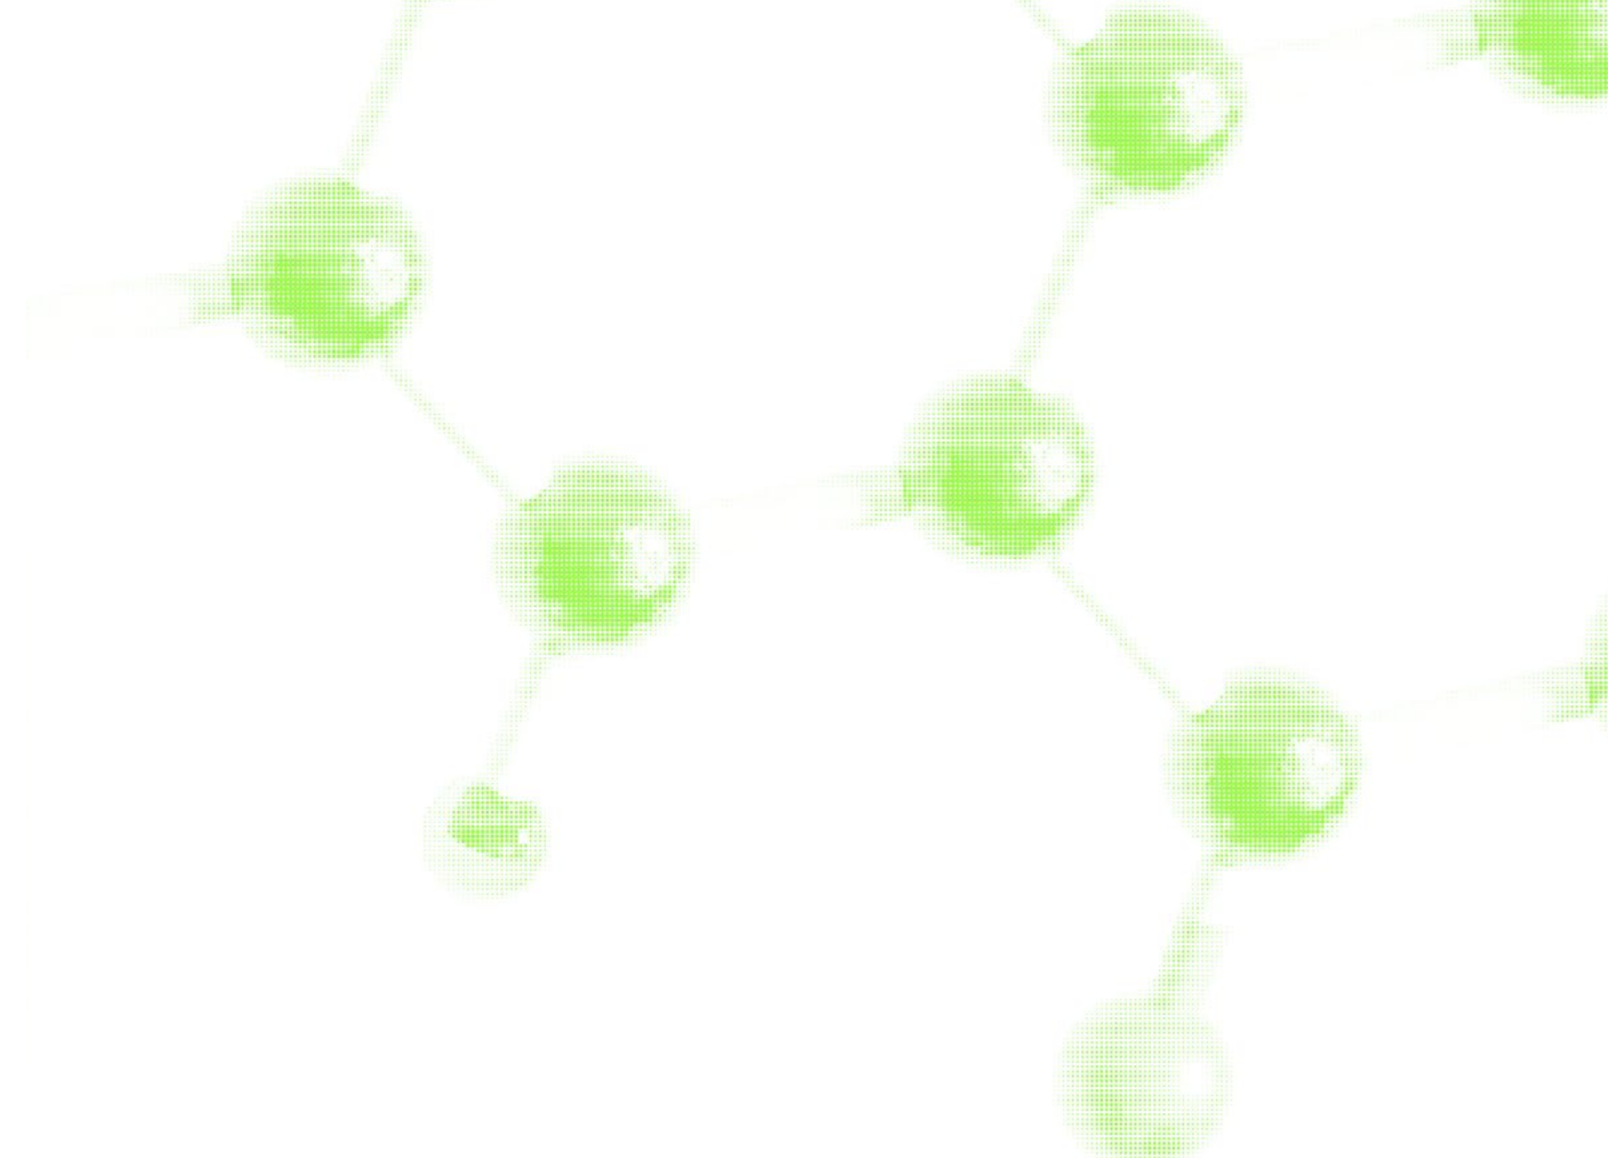

OPENING UP ACCESS TO GLP-1 RECEPTOR AGONISTS

Sometimes, a new medicine seems so broadly beneficial that there's a strong public-health argument for making it cheap or even free. When mRNA vaccines against Covid-19 became available in late 2020, for example, nations devised government-funded programs to buy millions of doses from manufacturers and distribute them free to everyone. But when it comes to GLP-1 receptor agonists, which can be life-changing for people with obesity and diabetes, a very different story is playing out, at least in the United States. The drugs are extremely costly here (around $1,200 per monthly dose) and they aren't covered by most insurance plans. We'd like to see some real innovation, both at the policy level and within the drug business itself, aimed at changing this situation.

Figure 7.6.1 Heavy out-of-pocket costs for weight loss medications

As the number of Americans taking Ozempic and other prescribed GLP-1 agonists grows, the drugs' high costs are thinning their wallets as well as their waistlines.

Source: KFF analysis

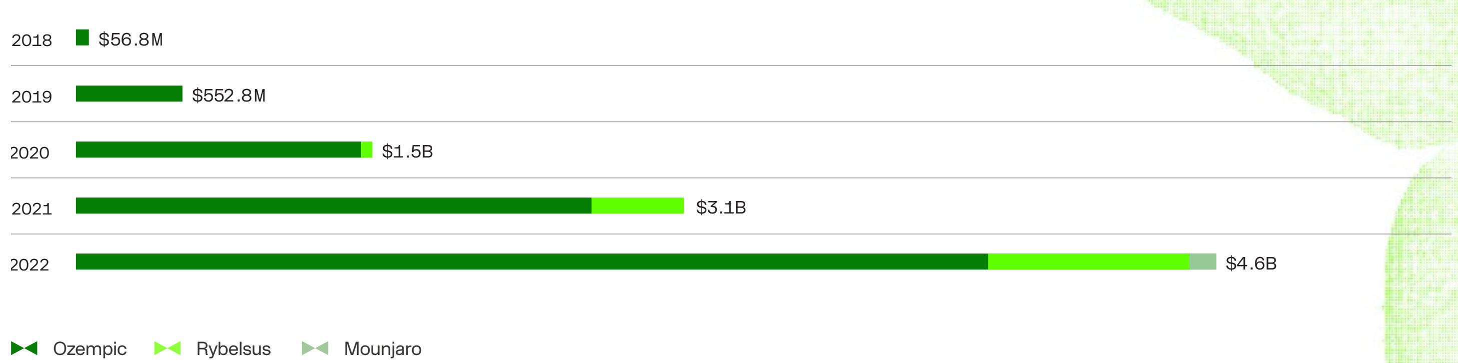

The GLP-1 receptor agonists mimic the effects of glucagon-like peptide-1, a hormone normally produced in the gut that tells the brain it's okay to stop eating. The difference is that they last much longer than the natural hormone in the bloodstream. They're so effective at curbing high blood sugar that they're now recommended as a first-line treatment for Type 2 diabetes, even before insulin. Obese patients taking semaglutide (branded as Ozempic or Wegovy) lose an average of 15 percent of their body weight, and those taking tirzepatide (Mounjaro, Zepbound) lose up to 21 percent.

These are profound benefits. "If you talk to doctors about this, it's not an evolution, it's a revolution in medicine, where morbidity decreases across every dimension," says DCVC's Bogue. For individuals, broad access to GLP-1 receptor agonists would likely lead to lower rates of heart disease, stroke, and cancer, according to a 2023 study from the University of Texas Health Sciences Center. Minority and underserved populations, who are disproportionately affected by obesity and diabetes and can least afford the new drugs at their current prices, would see a particular boost. The larger healthcare system would benefit too, with fewer expensive treatments and hospitalizations needed down the road.

Unfortunately, current policy makes it harder for people to get access to these drugs, not easier. By law, the Centers for Medicare and Medicaid Services is barred from paying for weight-loss drugs (though the agency announced in March that it will cover Wegovy for patients who also have heart disease). As of the beginning of the year, three-quarters of employer-provided healthcare plans didn't cover GLP-1 agonists. Those that do are trying to control costs by discouraging doctors from prescribing the drugs. Some observers argue that patients should be steered toward older weight-loss drugs, which aren't as effective as the GLP-1 agonists but are usually far less expensive.

We see no reason for the U.S. healthcare system to fight obesity and diabetes with one hand voluntarily tied behind its back. It's time for insurers and Congress to listen to pressure from groups such as the American Medical Association, which adopted a policy last November urging insurers to provide coverage of all FDA-approved weight-loss medications, including GLP-1 agonists. And Congress should okay broader Medicare coverage of the drugs, as the Washington Post editorial board urged last fall.

We understand the thinking behind the high list prices for GLP-1 receptor agonists. Drug development is a long and risky process; somebody has to cover the cost of a robust drug R&D sector; the GLP-1 agonists do represent a significant innovation with potentially huge public-health payoffs down the road. For all these reasons, the drug makers should be allowed to capture a good deal of the value of what they've created. The question is *how much* value. A study published in JAMA Network Open in March showed that if manufacturing costs were the only consideration, generic GLP-1 agonists could be made for less than $75 per month.

Figure 7.6.2 Monthly costs of GLP-1 agonists

Med	Monthly cost
Ozempic	$1,202
Wegovy	$1,664
Mounjaro	$1,222
Zepbound	$1,268

Source: SingleCare

At DCVC Bio, we believe in preserving and augmenting health rather than waiting to treat the already sick (see Section 7.7). And it's important to remember that prescription drugs are the most cost-effective tool we have in our health-preserving arsenal, representing only 11 percent of total personal healthcare spending. GLP-1 agonists, while currently too expensive in the U.S., are still cheaper than cancer therapy, a heart bypass, or other invasive interventions. We need to develop a way to resolve the big socioeconomic questions around cutting-edge treatments like these, because Ozempic and Zepbound certainly won't be the last major new drugs with the potential to transform our health. If their benefits can't be widely shared, what is innovation really for?

In our reading and our conversations with entrepreneurs, DCVC Bio occasionally hears of molecules, diets, or devices that (their advocates promise) will reverse aging. But a pill or regimen that will let you stay young forever is a shiny object if we ever saw one. Unfortunately, much of what's called "longevity science" today is built around ideas that are either impossible to replicate, impossible to falsify (since they would take a lifetime to test properly), or have overly broad side effects. Rather than pursuing elusive anti-aging treatments, we believe innovators should be focused on removing the barriers to disease prevention created by today's healthcare system.

Rather than pursuing elusive anti-aging treatments, we believe innovators should be focused on removing the barriers to disease prevention created by today's healthcare system.

↳ Resveratrol, a plant compound found in found in red wine and other plant products, was once alleged by scientists to prolong life spans. Online, you can buy "telomere-lengthening" supplements meant to protect the chromosomal end caps that supposedly act as internal clocks ticking down to cell death. Project Blueprint—a multimillion-dollar effort by tech CEO Bryan Johnson to reverse his own measured biological age—is now transmuting into a business selling olive oil and supplements to Johnson's acolytes.

Most of these proposed strategies for living longer have ultimately amounted to fads or popular myths, unsupported by evidence. The word *longevity* itself, notes physician and author Peter Attia, "has been hopelessly tainted by a centuries-long parade of quacks and charlatans who have claimed to possess the secret elixir to a longer life." One of the only lifestyle changes that *does* seem to extend lifespan—severely restricting caloric intake, whether by eating less or fasting intermittently—is also one the most arduous for people to sustain.

Of course, there's nothing wrong with wanting more years of healthy life. But now there's a long-overdue switch underway in the medical world, away from a focus on lifespan, the sheer number of years a person survives, toward *healthspan*, the years in later life that are relatively free from disability and disease. Researchers call the idea "compression of morbidity." We wholeheartedly agree with this shift. And when it comes to extending healthspan, biotechnology innovations like those described in this report certainly have a big role to play.

The biggest problem is that all healthcare economies are built around fixing disease after symptoms arise, rather than preventing it.

But we also see a need to reexamine the way we practice and pay for medical care. The biggest problem is that all healthcare economies are built around fixing disease after symptoms arise, rather than preventing it. An individual patient can take a proactive, participatory stance toward their own health and adopt tactics like eating less, sleeping and exercising more, and seeking an optimal, personalized combination of medicines. But they'd still be thwarted in every medical encounter by a system that's designed to treat everyone the same way, limit access to preventive services, and minimize or overlook the risks of doing nothing.

For example:

- Doctors and PAs see so many people every day that they often can't take the time to talk with patients about diet, lifestyle, and other preventive measures.
- Diagnostic tools are underutilized by physicians, and hard to get reimbursed for, even after approvals.
- Many health plans cover basic annual screenings, but restrict or deny payments for advanced health screenings that are seen as nonessential or elective.
- People with high-deductible plans may need to pay significant out-of-pocket costs before their insurance will cover vaccinations, lifestyle counseling, and other preventive care.
- Providers are reimbursed at lower rates for preventive services compared to treatment services; often, there aren't even clear billing codes for preventive services.
- Doctors too often wait until a patient has obvious symptoms to begin any kind of treatment, even though many patients are at elevated health risk long before their biomarkers cross widely accepted thresholds (e.g. an HbA1c level of 6.5 for Type 2 diabetes).
- When disease-preventing drugs such as statins do get prescribed, they're often dosed to bring patients into "normal" ranges for key biomarkers, instead of the ranges that may be optimal *for them*—which dooms patients to dying from an average set of diseases, on an average timeline.
- Physicians remain averse to prescribing diagnostic tests that have a high rate of false positives, even when additional tests are available that could rule them out and the consequences of missing a diagnosis would be severe, as in cancer.
- In our regulatory framework for drug trials, side effects that may be seen as worth suffering for patients who are sick are often seen as intolerable in people who are healthy; this extreme risk aversion makes it difficult to even test the longevity benefits of promising molecules with known side effects, such as rapamycin or metformin.

Clearly, lengthening healthspans will require reducing the death toll from the so-called Four Horsemen: heart disease, cancer, neurodegenerative disease, and metabolic dysfunction such as Type 2 diabetes. But this, in turn, will require a fundamental overhaul of the medical system's approach to disease prevention.

Figure 7.7.1 Life expectancy beyond age 60, by region

The average 60-year-old in Europe can expect to live another 21.8 years, and to spend 17.0 of those years in relative health. These numbers are lower in other parts of the world. Source: U.S. Census Bureau

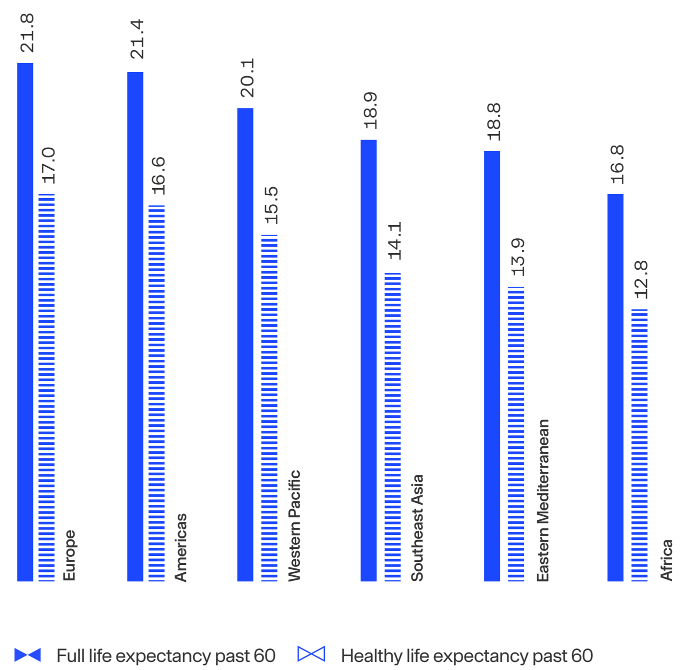

For each of these disease complexes, we should be testing for disease indicators before frank disease emerges, intervening earlier, and taking a more nuanced and flexible approach to risk. It may also be time to tweak the very idea of value-based care, in which payments to providers are tied to patients' health outcomes. That's a welcome challenge to the dominant fee-for-service health insurance model, and it's gradually infiltrating the healthcare system, with encouragement and incentives from the federal Centers for Medicare and Medicaid Services. But the most positive health outcome is to prevent disease in the first place, and only a few healthcare organizations so far see this as their job.

In all of this, as we noted above, biomedical innovation will also be critical. We'd like to see earlier and more widespread deployment of new diagnostic tests such as liquid biopsies for cancer. Health therapies such as statins and GLP-1 agonists should be more widely prescribed and easier to afford (see Opportunity 7.6). Advanced therapies such as Chroma Medicine's DNA methylation technology, which is already being tested as a way to eliminate harmful expression of the *PCSK9* gene and lower cholesterol levels, should be tested and advanced to the market as quickly as possible (see Opportunity 7.4). We should continue to study the immune systems of people who seem naturally resistant to cancer or neurodegenerative disease, as DCVC Bio–backed Alchemab is doing, to see whether they carry antibodies whose protective effects could be replicated. Ultimately, it's an evidence-based approach, together with innovations in the way we practice and pay for medical care, that will beat back the Four Horsemen and lengthen healthspans.

So why, in the end, do we invest in deep tech?

Because we see a world confronted by many urgent problems, from climate change to inequality to the diseases of aging to new threats around national security.

Because we believe that science and technology are still the most powerful tools we have for solving these problems.

Because we love working with the company founders who throw their minds, hearts, and lives into the solutions they're building.

And because we think we can do our small part to help these creators accomplish great things, while, in the process, earning a venture-scale return for our investors.

It's time to get moving. We hope you'll join us.

Thank you to the following DCVC partners

Zachary Bogue
Alan Cohen
Steve Crossan
Dr. John Hamer
James Hardiman
Clay Hutmacher
Earl Jones
Dr. Linda Maxwell
Matt Ocko
Matt O'Connell
Jason Pontin
Dr. Rachel Slaybaugh
Dr. Kiersten Stead
Ali Tamaseb

Editing Nate Nickerson, Jason Pontin
Writing Wade Roush
Data, Graphics, and Fact Checking Matt Mahoney
Copy Editing Tim Farrington
Photo Research Hailey Hiss
Design One Design Company, Chicago, IL

Typefaces FK Grotesk Mono, Kleukens-Egyptienne, Onsite
First Published September 2024

Deep Tech
Opportunities Report

2024

270 University Avenue
Palo Alto, CA 94301

500 2nd Street, Floor 2
San Francisco, CA 94107

dcvc.com