003

Deep Tech Opportunities Report

2025

An American Industrial Renaissance

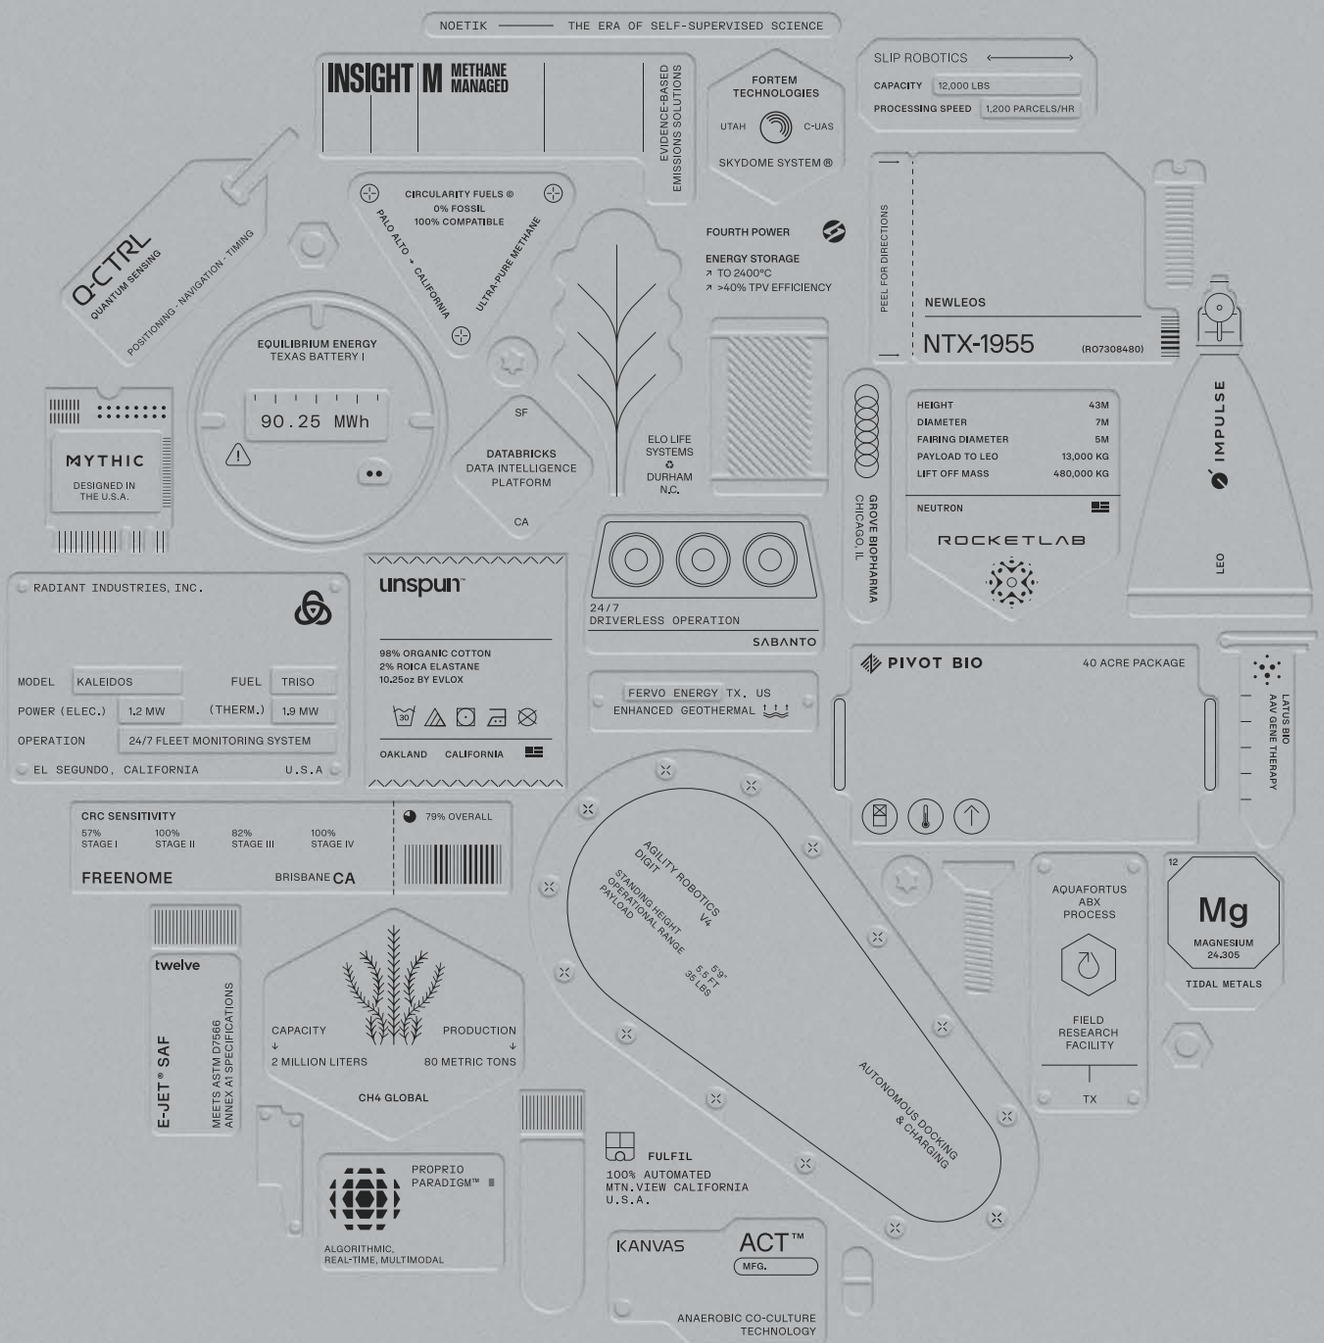

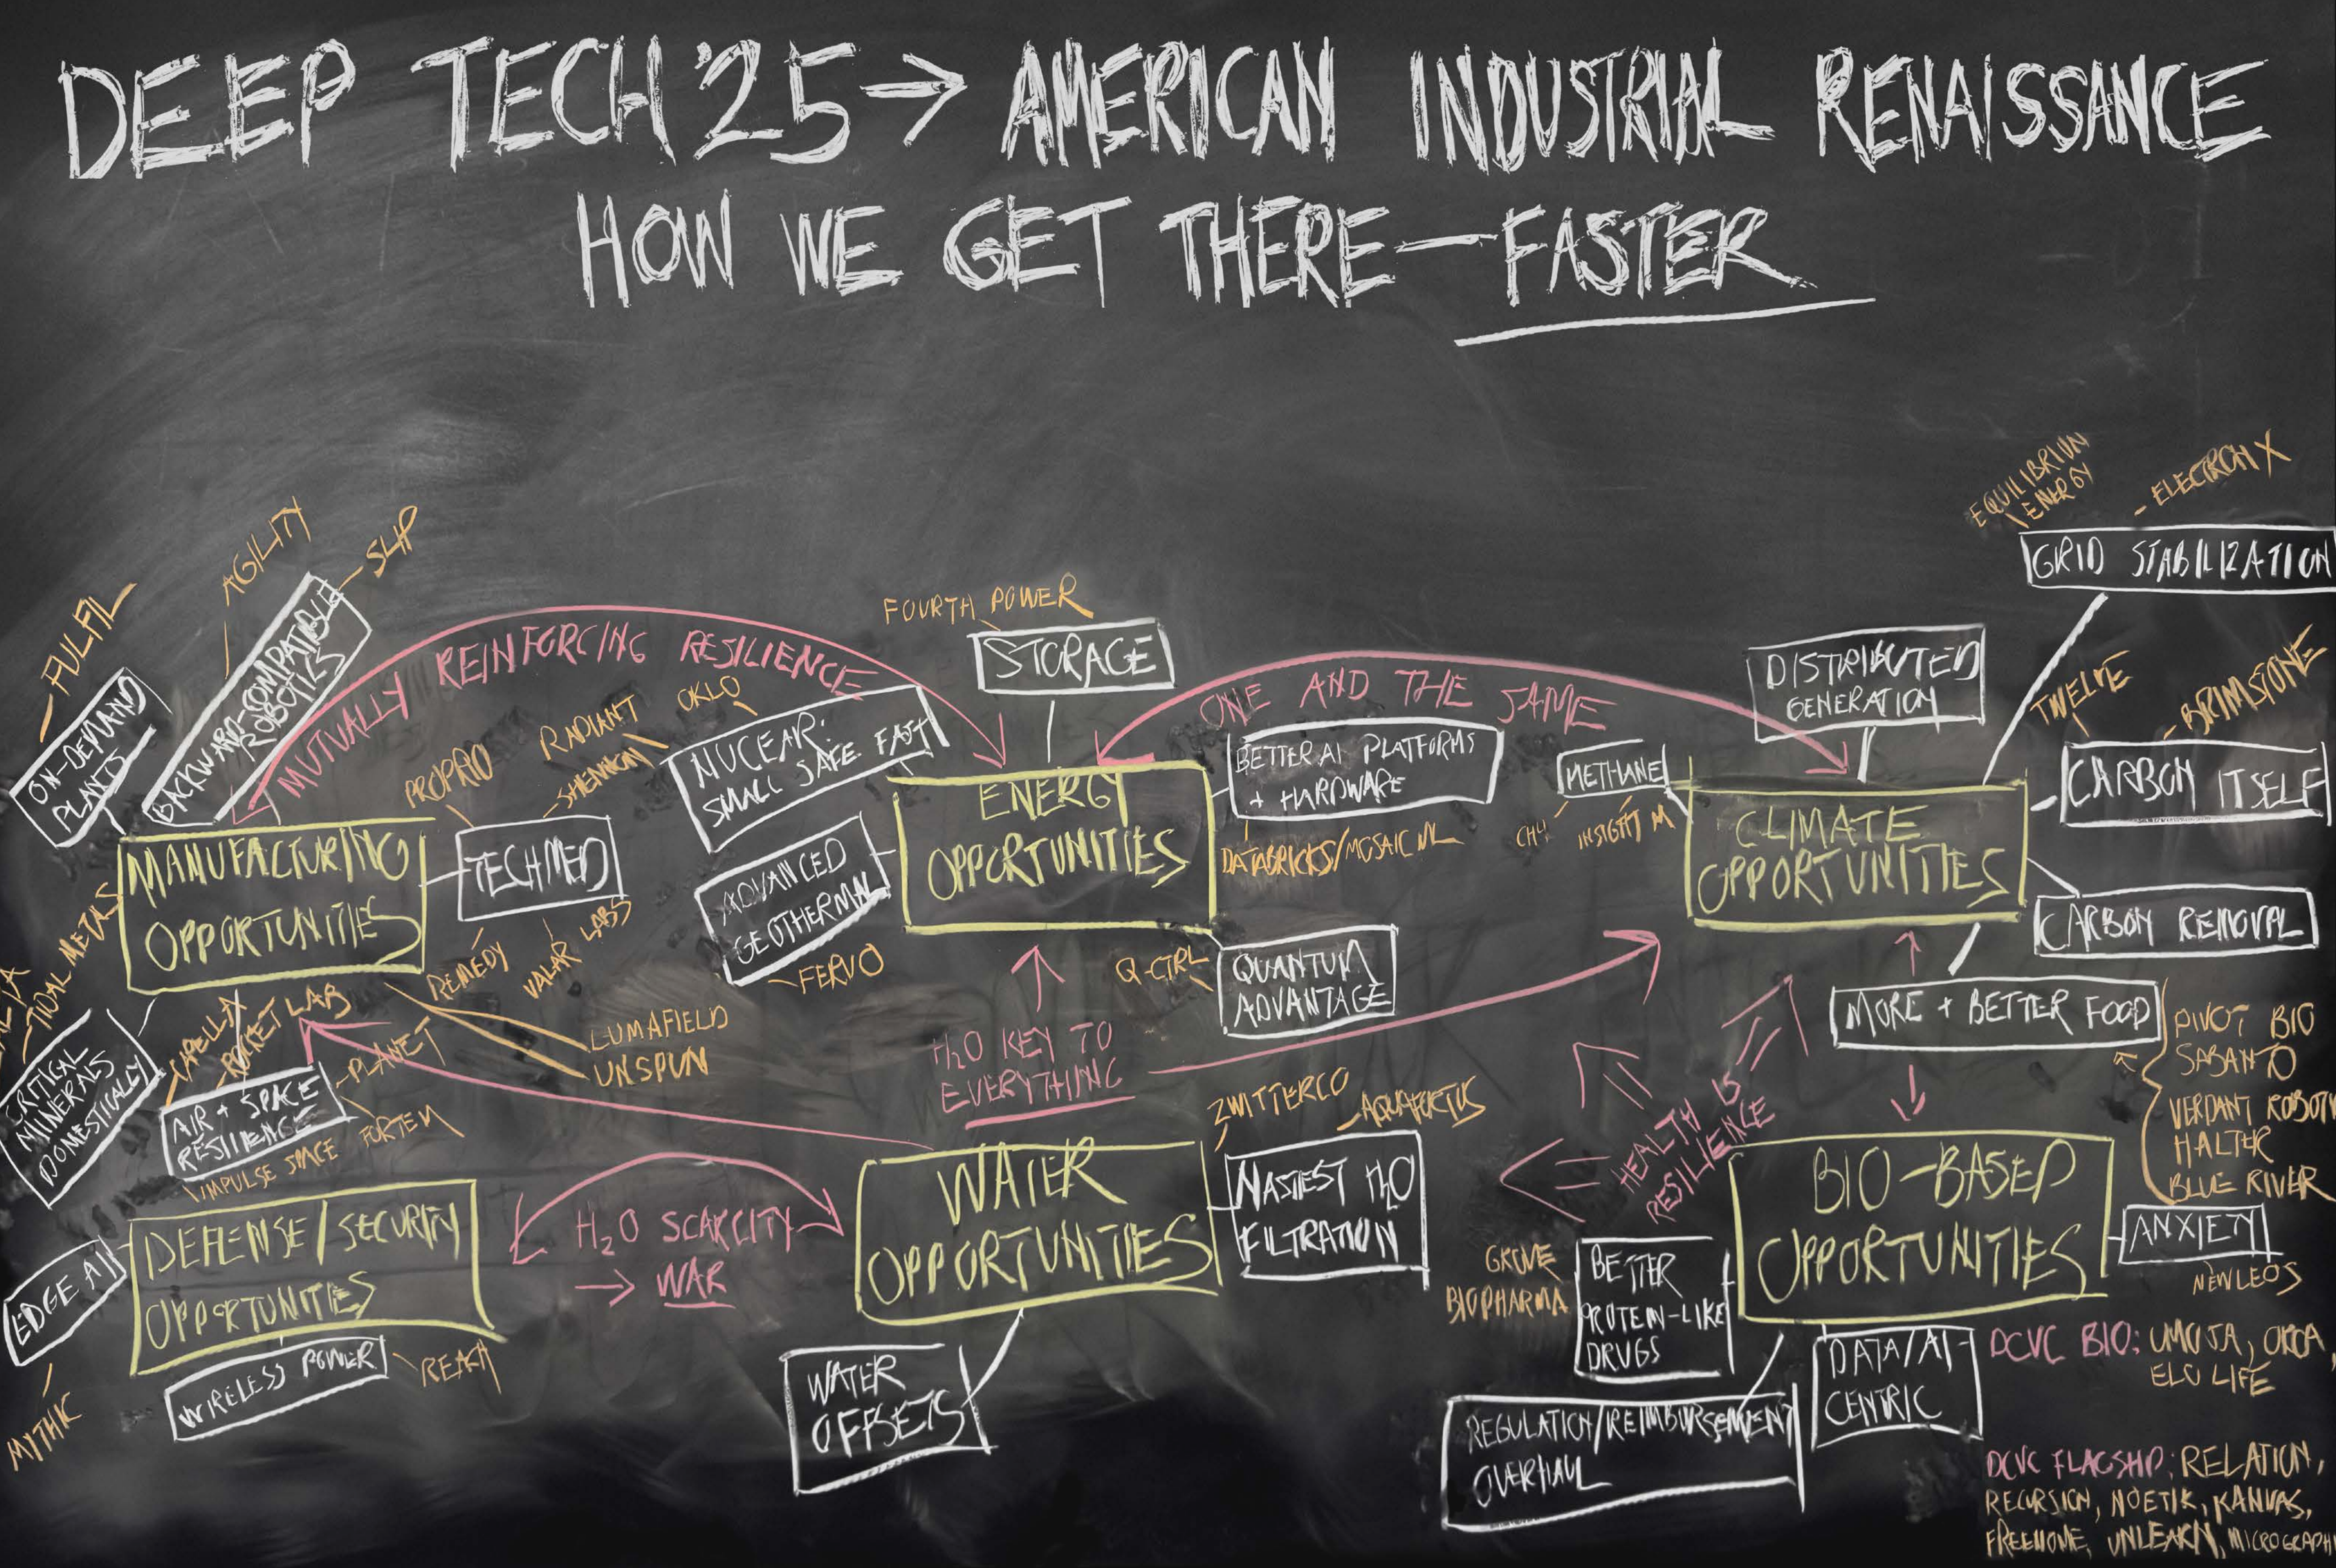

DEEP TECH '25 → AMERICAN INDUSTRIAL RENAISSANCE
HOW WE GET THERE — FASTER
MUTUALLY REINFORCING RESILIENCE
ONE AND THE SAME
MANUFACTURING OPPORTUNITIES
ON-DEMAND PLANTS
FULFIL
BACKWARD-COMPATIBLE ROBOTICS
AGILITY
SLIP
TECHMED
PROPRIO
SHENMON
REMEDY
VALAR LABS
CRITICAL MINERALS DOMESTICALLY
TIDAL METALS
AIR + SPACE RESILIENCE
CAPELLA
ROCKET LAB
PLANET
FORTEM
IMPULSE SPACE
LUMAFIELD
UNSPUN
ENERGY OPPORTUNITIES
STORAGE
FOURTH POWER
NUCLEAR: SMALL SAFE FAST
RADIANT
OKLO
ADVANCED GEOTHERMAL
FERVO
BETTER AI PLATFORMS + HARDWARE
DATABRICKS/MOSAIC ML
QUANTUM ADVANTAGE
Q-CTRL
H_2O KEY TO EVERYTHING
CLIMATE OPPORTUNITIES
DISTRIBUTED GENERATION
GRID STABILIZATION
EQUILIBRIUM ENERGY
ELECTRON X
METHANE
CH_4
INSIGHT M
CARBON ITSELF
TWELVE
BRIMSTONE
CARBON REMOVAL
MORE + BETTER FOOD
PIVOT BIO
SABANTO
VERDANT ROBOTICS
HALTER
BLUE RIVER
WATER OPPORTUNITIES
NASTIEST H_2O FILTRATION
ZWITTERCO
AQUAFORTUS
WATER OFFSETS
H_2O SCARCITY → WAR
DEFENSE/SECURITY OPPORTUNITIES
EDGE AI
MYTHIC
WIRELESS POWER
REACH
HEALTH IS RESILIENCE
BIO-BASED OPPORTUNITIES
BETTER PROTEIN-LIKE DRUGS
GROVE BIOPHARMA
ANXIETY
NEWLEOS
DATA/AI-CENTRIC
REGULATION/REIMBURSEMENT OVERHAUL
DCVC BIO: UMOJA, ORCA, ELO LIFE
DCVC FLAGSHIP: RELATION, RECURSION, NOETIK, KANVAS, FREENOME, UNLEARN, MICROGRAPHIA

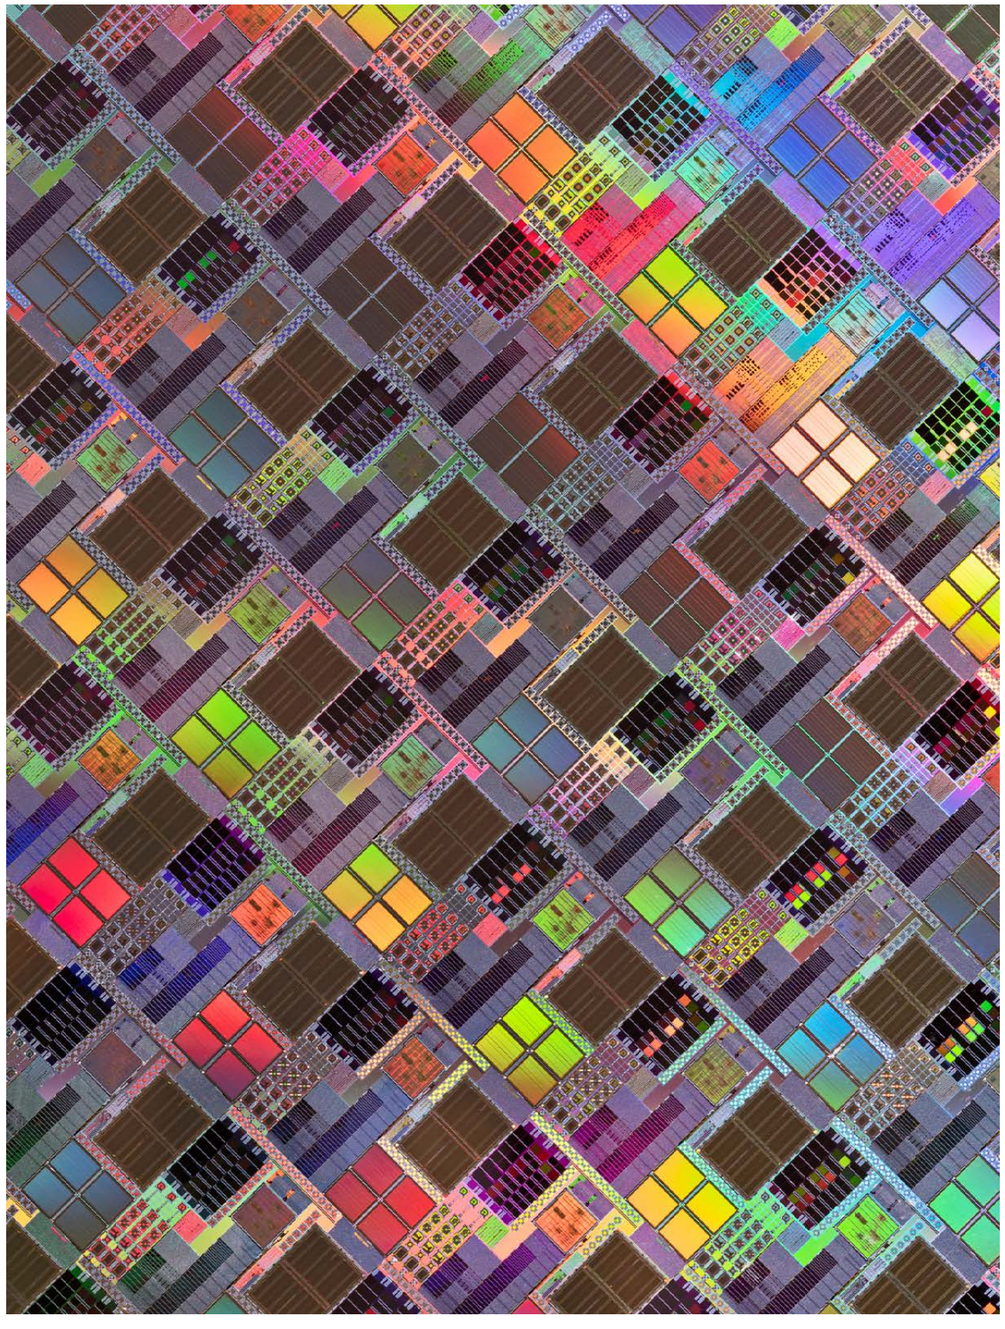

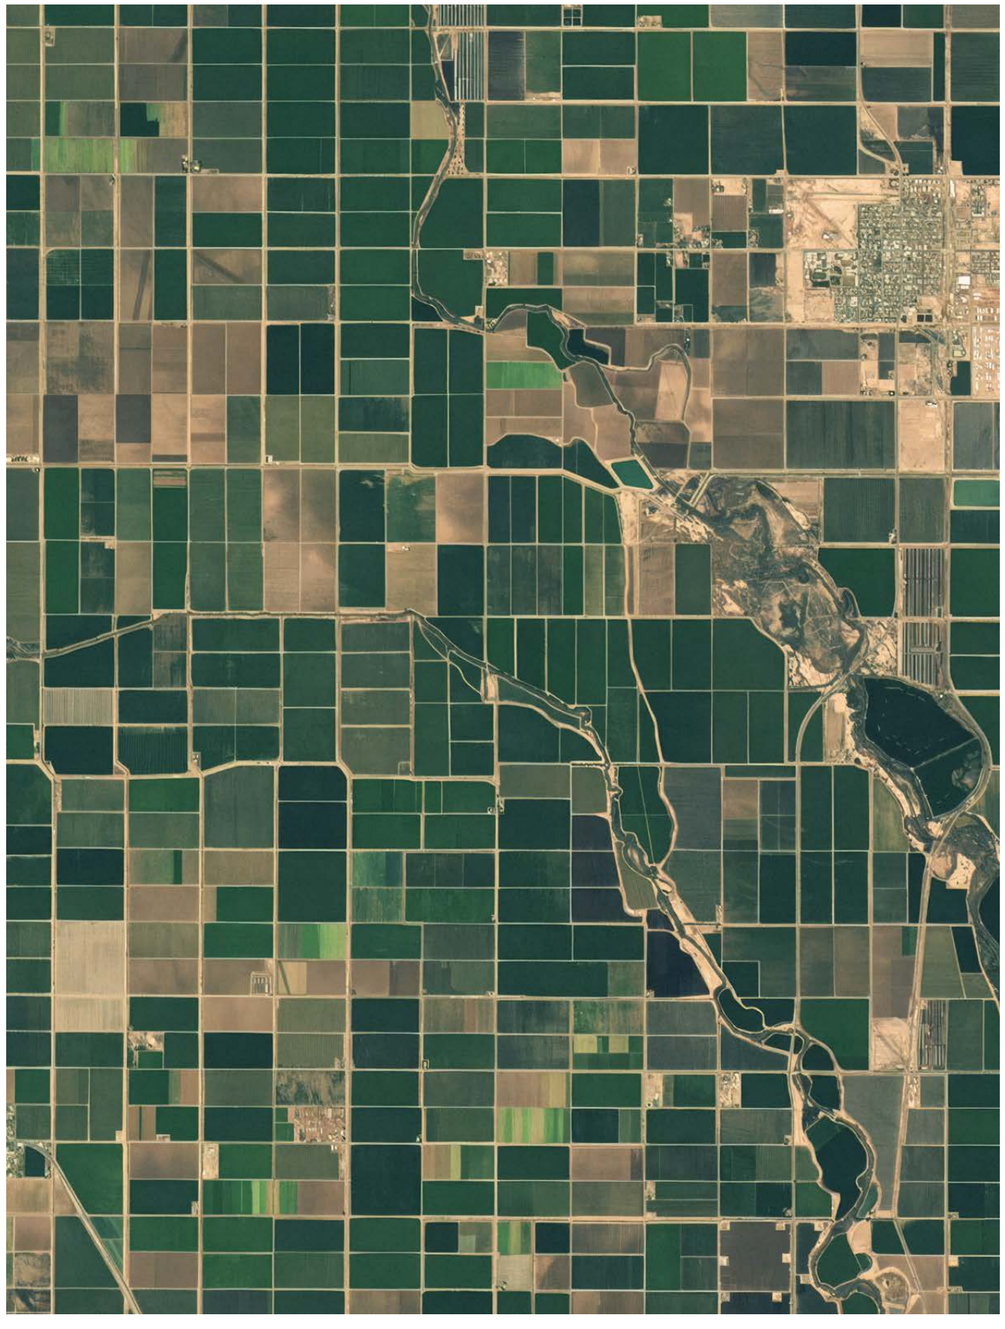

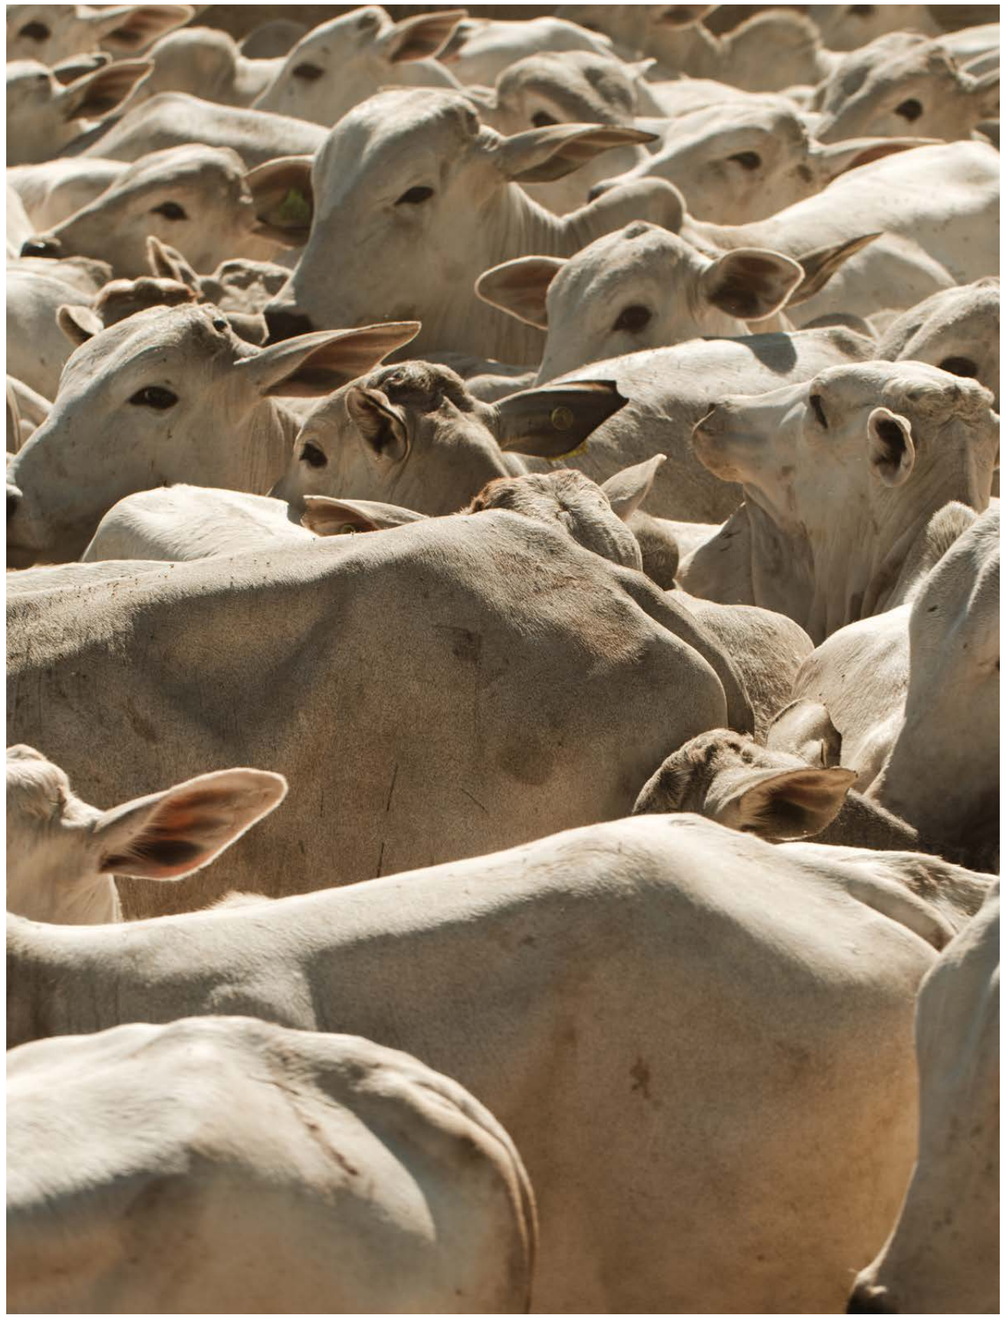

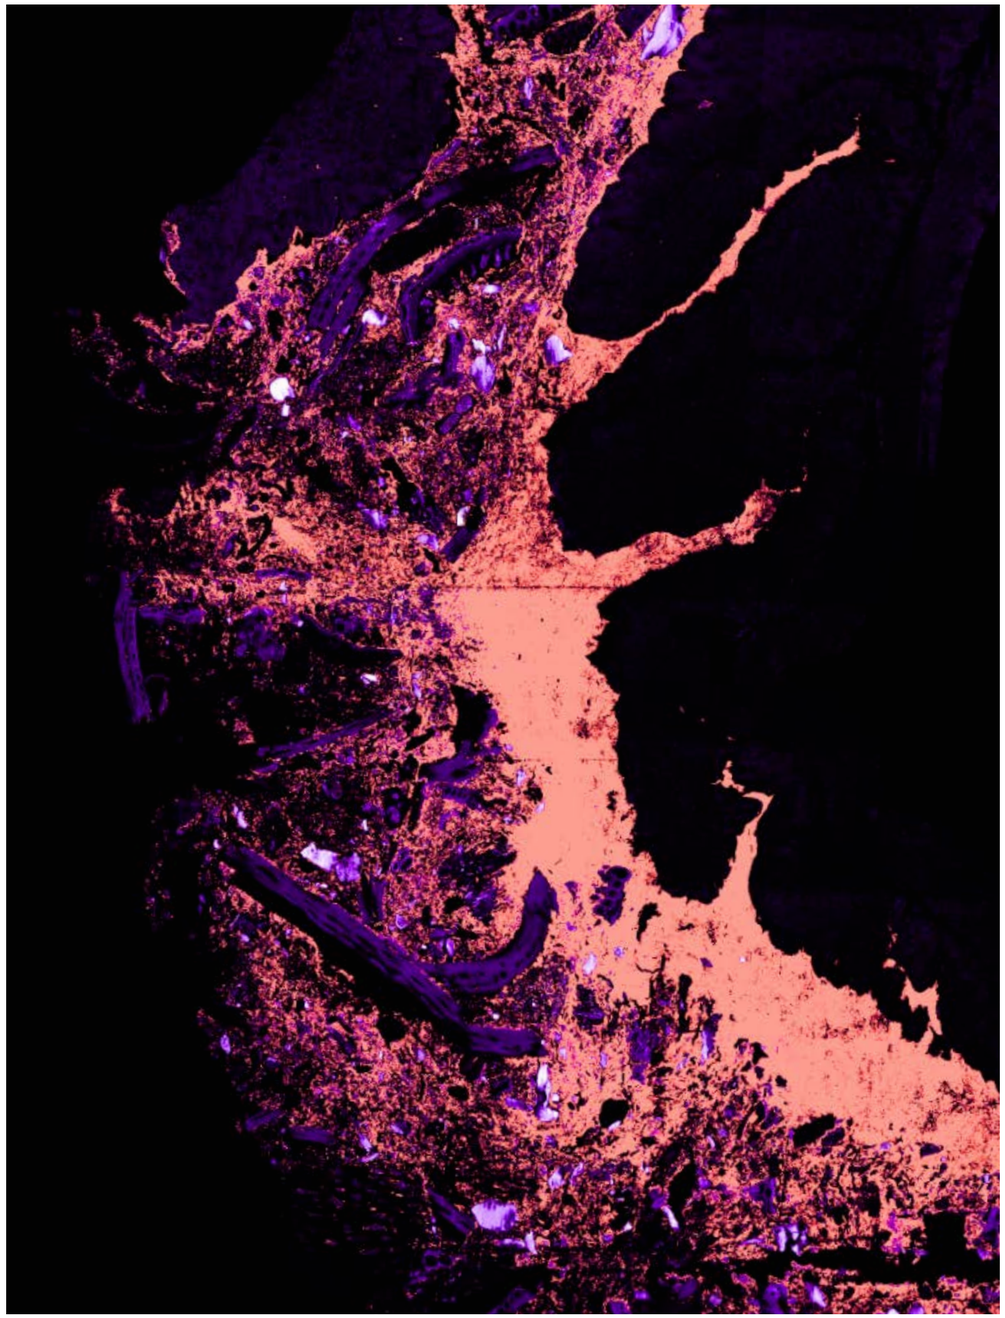

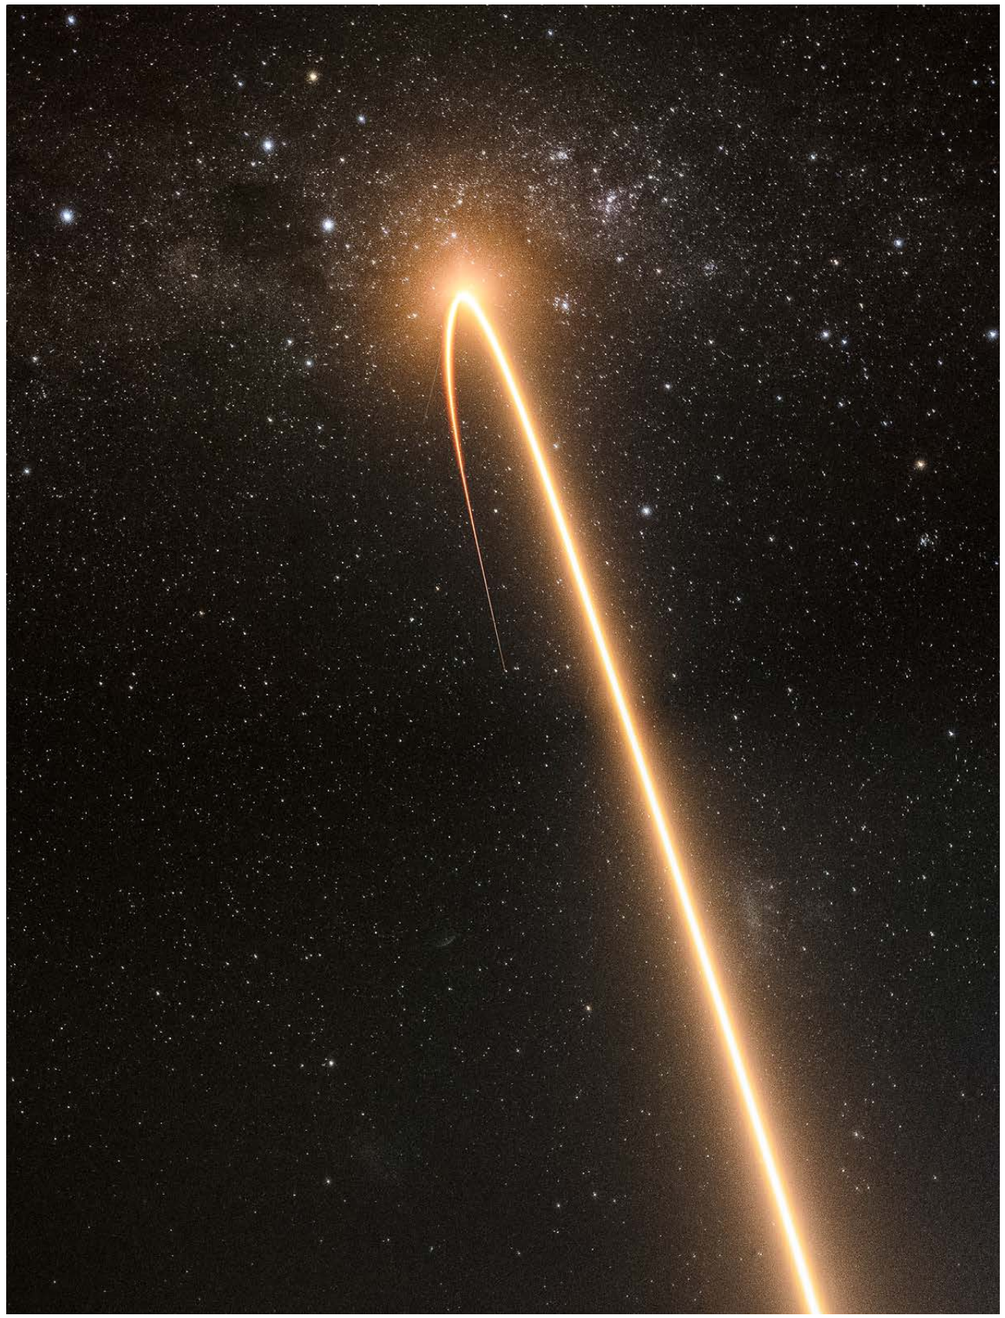

America is the clear world leader in industries such as information technology, finance, and entertainment. What's often forgotten in debates over American competitiveness or the impact of globalization is that the country is also, and has long been, home to a massive amount of high-value manufacturing, whether of chemicals, electronics, machinery, vehicles, or foods. In terms of the value added to global GDP by its manufacturing sector—that is, the value of outputs, minus the cost of inputs—the U.S. ranks second in the world, surpassed only by China. "People don't realize that the United States is, in fact, a leader in manufacturing," says Zachary Bogue, co-founder and managing partner at DCVC.

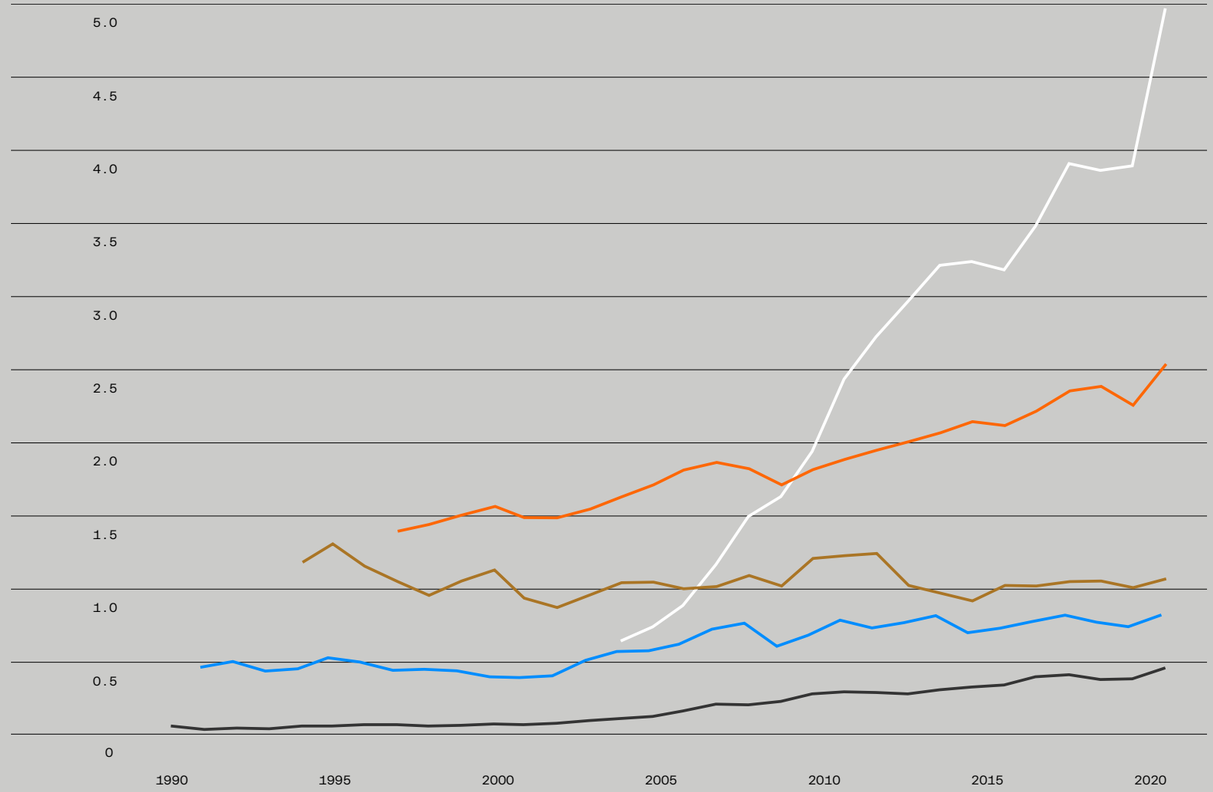

Fig. 0.0.1 **Value added to manufacturing**

Gross output of manufacturers minus the value of intermediate goods and services, in trillions of U.S. dollars per year.

Source: World Bank

We are still a nation of builders.

But if we want to make the competitive products and services that Americans and the world will want to buy *in the future,* while also sparing the environment and creating plenty of high-paying jobs, we can't ease up on efforts to modernize and decarbonize U.S. manufacturing. Nor can we remain dependent on strategic adversaries like China for affordable consumer goods or critical metals and minerals. The challenge now is to expand on our existing industrial base and bring home our most important supply chains, while at the same time shrinking our overall carbon footprint, undoing environmental damage, making our population healthier, and protecting against adversaries.

Engineering this kind of smart expansion will require innovations across multiple technologies:

01

Distributed low-carbon and zero-carbon energy sources, to power data centers and support the electrification of most forms of industry and transportation

02

A healthy workforce supported by modern medical care and cutting-edge treatments, administered at lower cost

03

AI and other forms of large-scale computing

04

Advanced robotics and manufacturing control systems

05

Digital twins and other forms of simulation

06

Clean, low-waste factories built closer to end consumers

07

Reliable supplies of clean water and critical metals and minerals

08

A supply of nutritious food grown on Earth-friendly farms

09

Defense, cybersecurity, and Earth-observation technologies that can detect and deter threats from nation-states and smaller actors

10

An overhauled energy-distribution grid that can compensate for supply fluctuations and environmental disruptions

In other words: ***deep tech*** **of the kind DCVC has long championed using the $4 billion in venture funds we manage.**

↗ Hard at work at Rocket Lab

Each annual edition of this Deep Tech Opportunities Report seeks to describe in broad terms what we see as the most obviously attractive and interesting deep-tech investment areas. This year, that's the prospect of an American industrial renaissance—one that is intertwined with, and indeed accelerated by, efforts to improve the environment. Here at DCVC we believe that private capital and the expertise that the best venture capitalists bring can help spark the required breakthroughs, eventually bringing about the kind of abundance and resilience we've written about in previous editions of this report. In short: *deep-tech innovation has never been more important* for the future of the planet and the security and prosperity of the U.S. and its allies.

As you read this year's report you'll see us emphasizing and elaborating on three major trends:

1

One, obviously, is the role deep-tech companies are playing in the reinvention of American industry. They're rethinking manufacturing methods, transportation and logistics, resource extraction and recovery, and many other areas in order to help American companies build and deliver better, higher-value products at lower expense.

2

The second inescapable trend is the growing importance of computation and artificial intelligence—in particular, foundation models cast in the same mold as the large language models being built by OpenAI, Google, and Anthropic, but trained on highly curated, proprietary data. These models now figure in almost everything deep-tech companies are doing, from building and programming robots to making farms more sustainable to reinventing drug development.

3

Third, we see an inextricable link between deep-tech innovation that bolsters American resilience and innovation that saves the environment. In many ways, they are one and the same. The reality, of course, is that there is no time left to put off the major reductions in greenhouse gas emissions that will be required to slow and eventually stop global temperature increases. Fortunately, smart reindustrialization helps to slow emissions. Reshoring that shortens supply chains and circumvents the wasteful and polluting practices of traditional trading partners is inherently planet-friendly. "Pretty much everything in advanced manufacturing also falls into the bucket of decarbonization," says DCVC general partner Milo Werner. On top of that, renewable and sustainable energy technologies are simply a better investment than fossil fuel–based alternatives. Properly understood, these are areas of common ground between political factions, where entrepreneurs and leaders can work together to create better jobs and a cleaner, more habitable environment.

→ **Alta Resource Technologies** is using biology to source critical minerals from e-waste, rather than faraway mining facilities.

"

You can have prosperity *and* resilience—and you don't have to be a caricature of a nineteenth-century robber baron to do it.

MATT OCKO
MANAGING PARTNER, DCVC

We think these three overarching trends overlap and reinforce one another, and when we invest we often look for companies tapping into them in novel, transformative, defensible ways. Alta Resource Technologies, one of our newest portfolio companies, is a prime example. The company uses computationally optimized proteins to efficiently latch onto and recover rare earth elements from e-waste—with the goal of creating a U.S.-based supply chain for critical minerals that are currently imported from China.

Another example is Relation Therapeutics, a drug discovery company working to treat osteoporosis, fibrosis, and other complex diseases. The company uses genomic and transcriptomic data from human cells to train AI models that help researchers predict which genes and proteins contribute to disease processes. In the lab, it observes how knocking out these genes affects cell function, generating more data that points to biological pathways that could be targeted with drugs. In the largest-ever deal for a seed-stage TechBio company, pharmaceutical giant GSK recently offered hundreds of millions of dollars in milestone payments for the right to develop drugs based on Relation's findings.

A point we hope is clear from this year's report is that computation feeds into smarter product design, which builds new industries that solve challenges in health, manufacturing, and climate, earning high profits that can in turn be reinvested in more computation and a revitalized infrastructure. "You can have your cake and eat it too," says DCVC co-founder and managing partner Matt Ocko. "You can have prosperity *and* resilience—and you don't have to be a caricature of a nineteenth-century robber baron to do it."

Our strategy at DCVC has always been to steer away from the frothy, the faddish, and the fashionable and invest instead in deep-tech companies with new, proprietary, computationally accelerated ways to solve the hardest, most urgent science and engineering problems. Today, more than ever, that strategy is proving profitable, for us, our limited partners, and our portfolio companies. We hope that you'll find inspiration in the stories in this year's report—and that if you haven't yet joined us in these critical efforts, you soon will.

2025

Deep Tech Opportunities Report

This third annual report offers a look at the deep-tech opportunities and companies we see as preeminent in this time of change. We believe the priority for American innovation capital should be rebuilding the nation's manufacturing base—and in such a way that the environment directly benefits, even in cases where that is not the primary motivation. We also see tremendous promise in new, computationally driven medical therapies for better health.

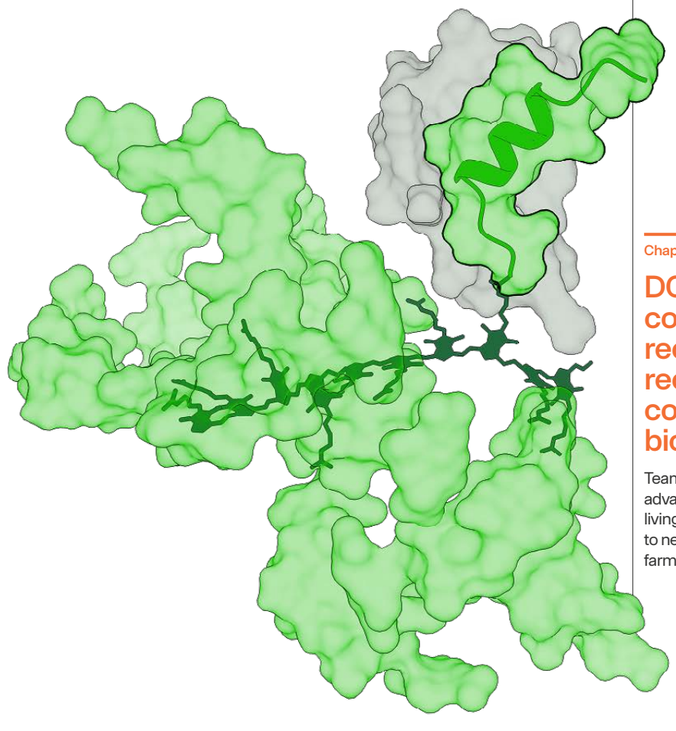

A few words on the investing climate.

Donald Trump's second term alters the competitive landscape for the companies DCVC funds—intriguingly, often for the positive.

AS WE WRITE THIS in May 2025, we are seeing rapid, at times unprecedented changes to the federal government, and to the parts of it that bear most directly on deep tech. These include moves to raise trade barriers, cut funding to grant-giving agencies and research universities, and halt some kinds of energy permitting while accelerating others.

But beneath the headlines surrounding the president's directives, what we also see are forces emerging in Washington, D.C., on Wall Street, and in Silicon Valley that will bolster the domestic industrial renaissance we're describing in these pages—and that are therefore likely to act as tailwinds for the kinds of deep-tech companies we back. As we did in last year's report, we'd like to provide a data-driven summary of those forces.

But first, allow us to point out that **our investment strategy is working.** DCVC is the world's largest pure-play deep-tech venture capital fund. We invest at all stages of company growth, from incubation to pre-seed to seed rounds to Series A and B, and beyond. But because most of our investments begin with early-stage rounds, we need the reserves required to help our portfolio companies all the way across the chasm between early innovation and market dominance (and to provide their customers with the assurance that they'll be around for the long haul). That's why we're grateful to have $4 billion under management, with investments in more than 270 companies over the lifetime of the firm. And we're proud to say that the distributions we paid out to our limited partners in 2024 were among the largest in the fund's history—at a time when distributions across the venture industry, as a share of net asset value, have been hovering near a 10-year low. Meanwhile:

Investors are flocking to deep tech.

The spread of deep learning and other forms of AI and computing across practically every realm of innovation looks to us like a technology supercycle, similar to the internet revolution of 1993-2000, and our companies have long been at its forefront. Today, the rest of the industry is joining us there. Generalist mega-funds such as Andreessen Horowitz, Sequoia Capital, General Catalyst, and Light-speed Venture Partners have dramatically increased the pace of their deep-tech investing—those four firms led funding rounds for 15 early-stage deep-tech companies in 2020 and 46 in 2024. (a16z has even set up an "American Dynamism" fund to invest in aerospace, defense, manufacturing, and related fields, while General Catalyst has unveiled an investing theme it calls "global resilience.") And at new and emerging deep-tech funds, the total assets under management have quintupled over the last five years, from $867 million in 2020 to $5.2 billion in 2024. We welcome this trend—because the more funds focus on deep tech, the more dollars are available for co-investment in the most promising companies.

Figure 0.1.1 U.S. manufacturing capacity growth

In 2024, U.S. manufacturing capacity was up 130 basis points over the previous year, and is growing consistently for the first time in a decade.

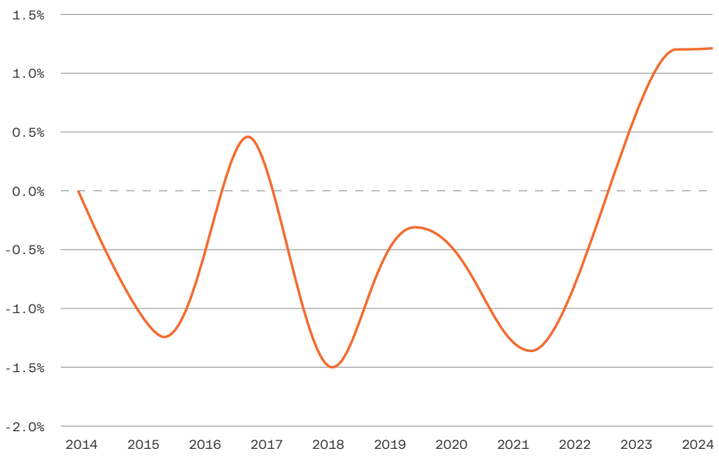

Figure 0.1.2 **Trump bump:** Protectionism during the first Trump administration drove a 550-basis-point increase in the share of foreign direct investment captured by the United States.

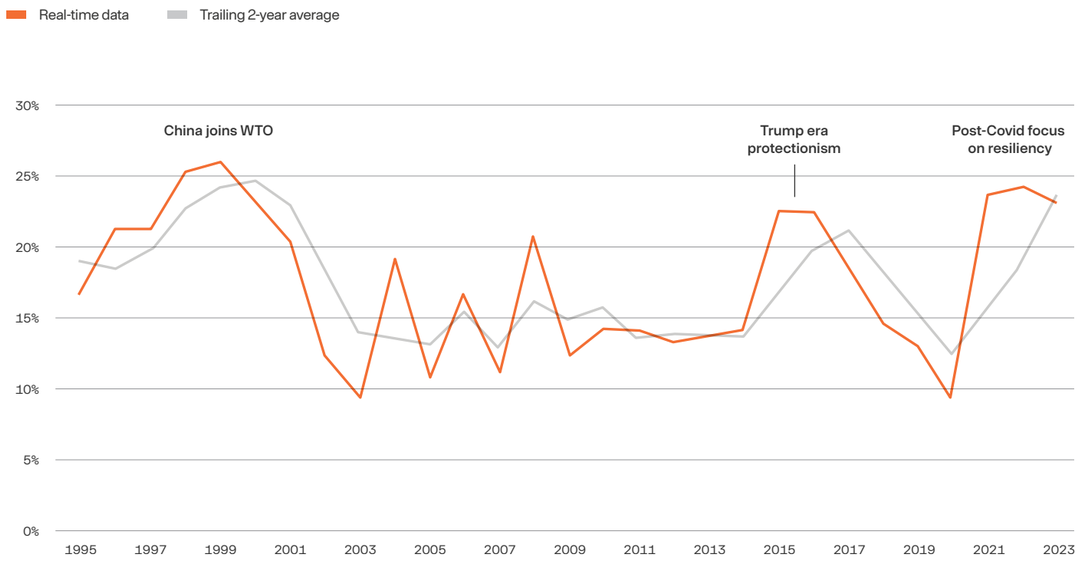

IPOs will be back in the picture if tariff-based market upheaval ends.

In 2022, by one tally, there were only 19 tech IPOs, raising $7.5 billion, but in 2024 there were 65, raising $32.1 billion. This is still far below the pre-Covid norm, but newly public companies (including DCVC portfolio companies such as Oklo and Rocket Lab) were trading extremely well before markets lost value due to tariff unrest. If investor sentiment improves, we expect the IPO market to be healthy. Overall end-of-year performance for 2023 IPOs was 6.1 percent, but for 2024 IPOs it was 83.7 percent.

Consolidation in the venture industry continues.

Even as the venture industry's attention turns to deep tech, the absolute number of active venture firms is shrinking. In 2023, the number of firms that have raised funds within the last eight years went down for the first time since the financial crisis of 2008. At the same time, though, the average fund size hit its highest level since 2008—indicating that more LP dollars are going to fewer top-tier firms. This is concerning in that the firms with the deepest pockets are able to buy into hot companies at inflated deal valuations, and companies with unjustifiably high valuations will inevitably have a harder time raising follow-on rounds and exiting profitably. On the other hand, the fragility of some of these overpriced companies opens up a competitive niche for first-in-category companies commercializing breakthrough technologies with more reasonable valuations—which is where we prefer to play. (↳Continued on page 20)

A drop to drink The All-American Canal snakes toward California's Imperial Valley, carrying water diverted from the Colorado River. Deep-tech entrepreneurs are developing ways to supplement dwindling surface water supplies **(see Opportunity 4.1).**

Protectionism puts wind in the sails of deep-tech companies.

Increased tariffs benefit many of our portfolio companies, most directly by encouraging foreign direct investment in their U.S.-based factories (a phenomenon we saw during the first Trump administration), and by increasing demand for domestically sourced materials, products, and energy. Meanwhile, we've always prioritized investing in companies that have subsidy-free paths to market. The classic DCVC company offers solutions that are better-performing or cheaper, or that cleanly solve a massive technical problem or market headache. Those advantages tend to accrue irrespective of who holds the White House. So we don't believe that a slowdown or reversal of Biden-era tax credits and other programs designed to fight climate change will undermine our portfolio.

We'll need deep tech to build a manufacturing economy powered by zero-carbon energy and blunt the worst effects of global warming.

The Trump administration has defunded some areas of climate science, in reaction to what it views as politicized overregulation of American industry and commerce. At the same time, it has signaled it will make changes, such as streamlining permitting and embracing nuclear and geothermal energy, that could boost clean tech. We also see strong market signals of continued demand for deep-tech environmental solutions, such as RE100, a coalition of more than 400 large corporations committed to switch to 100 percent renewable electricity. Clean capitalism is smart capitalism.

Biopharma hungers to replace loss of revenues.

The patents on many of the biopharma industry's most lucrative drugs and therapies are due to expire over the next few years, and the loss of exclusivity means they'll have to replace enormous amounts of revenue just to maintain current cash inflows—Deloitte puts the tally at $236 billion by 2030. This creates a robust market for mergers, acquisitions, and integrations for the kinds of biotech companies backed by DCVC Bio. Such companies are historically responsible for creating 70 to 80 percent of new drugs.

There's still a craving for innovation in many corners of the U.S. government.

Under the new administration, we expect to see the Pentagon continue to reform its acquisition process to funnel more resources to small companies and speed up the process of fielding new defense technologies (see Chapter 8 of this report). This overall re-engagement between our military leadership and most compelling new tech companies is good for our portfolio and vital to the nation's interests. It's our hope that as part of this doubling-down on America's world-leading strength in innovation, the Trump administration will ensure that basic university research gets the support it requires, and from which the American people have so greatly prospered. Silicon Valley–born deep-tech entrepreneurship depends on fundamental discovery, full stop.

Figure 0.1.3 The tech IPO market is gaining momentum

Both the number of technology IPOs and the post-IPO performance of newly public companies improved in 2024.

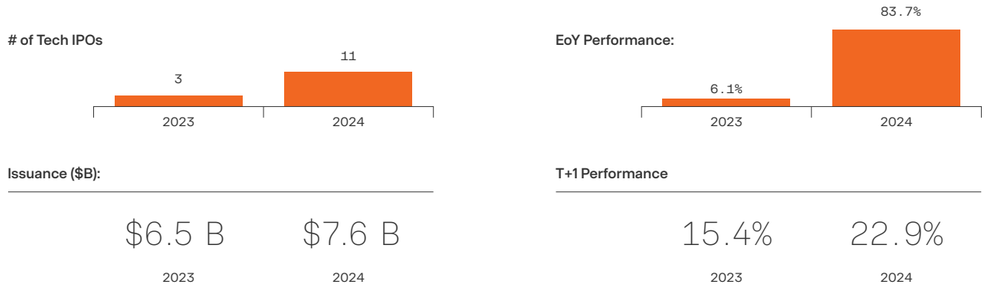

↳ DCVC's real strength is in our team of investors.

Zachary Bogue and Matt Ocko
Founders and managing partners, DCVC

Whether conditions are stormy or smooth, we trust our seasoned crew of general partners and operating partners. They have the expertise to recognize the glimmer of greatness in concepts that are truly novel—you'll read about many of them in the coming pages—and also the experience to help our portfolio companies navigate a changing world and get those ideas to market. Working with deep-tech innovators continues to be our privilege and our passion.

Deep tech will solidify the United States' advantage as a global manufacturing powerhouse while sparing the environment.

As we've already noted, the U.S. has long been a leader in manufacturing. The U.S. manufacturing sector's gross output has been going up more or less steadily for the last two decades, and U.S. manufacturing capacity increased markedly in 2023 and 2024, after a 10-year lull. But to make this linear growth exponential, we must massively increase the variety of the goods we manufacture and the efficiency with which we make and distribute them.

OPPORTUNITIES

COMPANIES

Alta Resource Technologies, Fulfil, Tidal Metals, Unspun

VOICES

James Hardiman, Earl Jones, Milo Werner

Unspun aims to reduce waste and carbon emissions in the garment industry by building 3D weaving robots that start from yarn to create clothing on demand.

Deep tech will solidify the United States' advantage as a global manufacturing powerhouse while sparing the environment.

Imagine that it's time to replace your old kitchen blender. If you visit Amazon you'll see dozens of different blender models from various brands, each with a unique capacity, power level, and features. At your nearest Amazon distribution center, there are probably 100 units of each blender in stock. Each of those units was probably manufactured nine to 12 months ago, likely in Asia, and spent weeks in transit on a container ship.

In aggregate, an enormous amount of energy and raw materials—in the form of plastic, metal, motors, circuit boards, and the like—is locked up in these devices. That's not even counting the fuel expended to get them to the Amazon warehouse, or from Amazon to you.

This stockpiling of goods is what enables the cornucopia of instant choice to which we've become accustomed in the internet era. But what if there were a lower-impact way to provide the same flexibility? What if you could simply tell Amazon which features you wanted, and the necessary parts could be printed and automatically assembled at the warehouse, to be shipped to you the next day? You'd still get the blender you want, but with a far smaller environmental impact and expenditure of vital resources.

The idea of on-demand manufacturing plants positioned close to the end consumer, a concept decades old, is no longer a fantasy. It's becoming real at companies like DCVC-backed Unspun—and it's just one of the ways entrepreneurs are rethinking how industrial production and supply chains should work.

The mission at Unspun is to reduce material and transportation waste in the fashion industry, using automated 3D circular weaving technology that can create woven garments on demand in microfactories. At its R&D facility in Emeryville, Calif., the company's prototype weaving machine pulls in thousands of individual yarns to create a pair of pants in 20 minutes—four times the speed of traditional methods. Now the company is focused on scaling its technology to meet growing demand from customers—including major retailers like Walmart and Decathlon—as well its garment manufacturing partners.

One of Unspun's goals is to make fashionable clothes while sidestepping the textile waste endemic to the fashion industry. Traditionally, long production and transportation timelines have meant that retailers must order items a year or more before they reach stores. Customers' tastes change faster than that, resulting in a constant mismatch between supply and demand—and leaving huge amounts of inventory unsold.

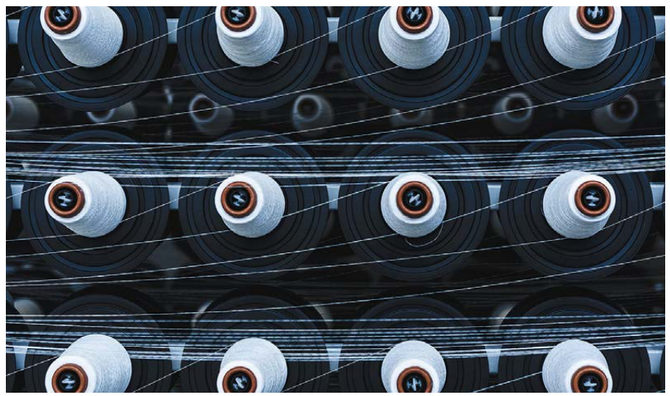

↑ **Unspun's** prototype Vega spinning machine, working faster than the eye can see, unspools thousands of individual yarns and weaves them seamlessly into 3D textiles.

Consumer trends forecaster WGSN estimates that 10 to 30 percent of all garments manufactured are never sold to consumers. Most of this deadstock ends up in landfills or incinerators. (Burberry formerly bought back and burned their overstock simply to prevent discounting.) Researchers at the World Resources Institute and the Apparel Impact Institute estimate that total carbon emissions from the fashion business, counting production, transportation, and disposal of overstock, exceeds a gigaton of CO_2 equivalent per year, or about 2 percent of all global emissions.

Part of the industry's challenge is that yarns are woven into fabric in one factory, prepared and finished at another, and cut and sewn into garments at yet another. Then the finished clothing flows through a labyrinthine global chain of distributors and warehouses, with time lost and carbon burned at every step. By contrast, Unspun's spinning machines allow brands and retailers to leap straight from yarn to the finished product and deliver items to stores in days. That enables stores to order only what they need, since they can easily restock styles and sizes that sell out.

"By delaying the configuration of materials—making products closer to the end customer, growing food closer to the market, weaving clothes closer to the retailer, manufacturing appliances in the region where they'll be purchased—you are only moving the molecules that people are actually using," says DCVC general partner Milo Werner. "You're eliminating waste throughout the whole supply chain." And even as Unspun helps the apparel industry to avoid overstock and reduce emissions, their regional microfactories will create skilled jobs, bringing home an industry that has been largely outsourced to far-away regions with lower labor costs.

↑ **Clothing graveyard** in the Atacama desert of Chile. Thousands of pounds arrive at the landfill each month from the nearby free trade zone of Iquique.

10-30%

of all garments manufactured are never sold or worn. Most of this deadstock ends up in landfills or incinerators.

→ **Fulfil Robotics** | In Mountain View, Calif., the company built the first fully automated micro-fulfillment grocery store.

Fulfil is another exciting example of a DCVC-backed company that's tackling waste by rethinking the supply chain. The United Nations says that in 2022, 19 percent of food produced for consumers was wasted at the retail, food service, and household level, and another 13 percent was lost in the supply chain. That waste not only costs a billion meals a day, but also generates 10 percent of global greenhouse gas emissions. Fulfil builds small fulfillment warehouses for online grocery orders where R2-D2–like shop bots pick and pack items for delivery by services such as DoorDash. Its inventory management software lets store owners track which perishable items are approaching their expiration dates and use dynamic pricing to ensure they're sold before they must be discarded.

If Fulfil automated 5 percent of U.S. grocery shopping, it would reduce food waste by 800 million pounds per year, says CEO Mir Aamir. "Fulfil's AI and robotic approach to the e-commerce grocery segment serves three key roles," says DCVC general partner Alan Cohen. "It reduces the punishing labor costs that make the segment unprofitable today. It can shift operations to 24/7. And it can potentially remove food deserts from underserved areas."

Another business in need of innovation and reshoring is the production of rare-earth elements and the other minerals we need to build everything from electric vehicles to consumer electronics, jet engines, and defense-critical technologies. When China imposed export curbs on gallium in 2023, prices for this silvery metal—used in semiconductors and photovoltaic panels—roughly doubled. Antimony, germanium, and other strategic resources could be subject to similar swings in price and supply, at China's whim.

"The prospect of crippling shortages looms large" as the world races to electrify transportation, digitize industry, and fortify defense capabilities, writes Nathan Ratledge, founder and CEO of Alta Resource Technologies. With seed funding from DCVC, Alta is tackling this problem at the very beginning: the extraction stage. The company uses computationally optimized proteins to pull high-purity rare-earth elements from low-grade sources available domestically, such as e-waste. The process bolsters the domestic supply chain for critical minerals while at the same time sidestepping traditional mining and refining, and its high environmental toll.

"We can't just take the mines and machines in China and start running them here—that's not going to work from a price perspective," says DCVC general partner James Hardiman. "The technology needs to be different in order for us to compete. And so that looks like Alta, where there's a line of sight to being the global low-cost separator for rare earths. Even if it's done in the U.S. at U.S. wages, it will still be cheaper than what China does."

China's control over critical minerals mining extends to magnesium, a strong and light structural metal that could be critical to building new fleets of lightweight vehicles. To help with this problem, DCVC led a seed funding round for Tidal Metals, which is developing a way to extract the metal from its most abundant source: ocean water. The company's patented temperature-swing adsorption vapor pump uses a process called fractional crystallization to pull magnesium out of seawater or brines in the form of anhydrous magnesium chloride, which can then be separated using electrolysis into chlorine and magnesium metal.

"How do we reindustrialize in a way that's environmentally harmless?" asks DCVC operating partner Earl Jones. "Wouldn't it be great if we could get resources out of wastewater or seawater and not have to use big earth-moving devices to dig out ore and then refine that ore using very energy-intensive thermal processes? Tidal Metals leans right into all that."

Reshoring an industrial process that's wasteful and polluting does no one any good. The manufacturing and logistics breakthroughs that will create tomorrow's high-growth companies, we believe, are those that that are cheaper, cleaner, and more energy-efficient than traditional methods.

"How do we reindustrialize in a way that's environmentally harmless?... Wouldn't it be great if we could get resources out of wastewater or seawater?

Earl Jones
Operating partner, DCVC

America needs exponentially greater computing power, in more locations, using less electricity.

DCVC backs companies that use new kinds of computing and robotics to solve hard, trillion-dollar science and engineering problems. And sometimes, we back companies that are reinventing computing itself.

Slip Robotics' automated loading robots can increase worker safety, reduce costs, and boost throughput at loading docks by an order of magnitude.

OPPORTUNITIES

COMPANIES

Agility Robotics, Mythic, Q-CTRL, Slip Robotics

VOICES

James Hardiman, Matt Ocko, Ali Tamaseb

New hardware architectures will make AI more energy-efficient.

The makers of today's frontier AI models say that to make their systems smarter and achieve artificial general intelligence (AGI) and eventually superintelligence, they may need to build computing facilities that are 50 to 1,000 times the size of today's data centers. Such facilities would have correspondingly greater power needs—which is why these companies are already looking to new energy sources such as advanced geothermal and nuclear power. (Microsoft, for example, has contracted to buy *all* of the power from a recommissioned Unit 1 reactor at Three Mile Island.) But before AI providers commit to such a massive scale-up, the industry must grapple with the interlinked problems of energy efficiency and algorithmic efficiency.

Unit 1 of the Three Mile Island nuclear plant shut down in 2019, but it will be restarted in 2028 under a deal between plant owner Constellation Energy and Microsoft, which needs the power for regional data centers.

It's become an article of faith among technology policy experts that massive new investments in computing hardware are essential for U.S. competitiveness. We agree—but we think it's time to get smart and creative about where to find the electricity and water needed to power and cool all this hardware, and how to moderate the industry's overall energy needs.

The leading voice in the push for greater spending on AI infrastructure is OpenAI CEO Sam Altman, who has said that with sufficient investment "it is possible that we'll have superintelligence in a few thousand days." In a January 2025 white paper, OpenAI argued that "chips, data, energy, and talent are the keys to winning on AI—and this is a race America can and must win." The company argued for dramatically increased federal spending on power and data transmission, and warned that if the U.S. doesn't attract the private investment funds available for such projects, this money will flow instead to China-backed projects. A coalition including OpenAI, SoftBank, Oracle, and MGX primed the pump by announcing a joint $100 billion investment in Stargate, a new company tasked with expanding OpenAI's infrastructure, starting with a facility in Texas.

But before companies spend billions or trillions of dollars on new data centers and power facilities, they should be clear about their reasons for choosing to invest in AI over other valid priorities, and they should think in advance about ways to minimize these technologies' environmental impact. The theory that seems to drive OpenAI and its competitors is that the first company or country to achieve AGI will permanently dominate the global high-tech economy. But we don't see this as a safe assumption, given that the core ideas in artificial intelligence are so easily copied, stolen, or reinvented.

Indeed, one of the big news stories in the AI world as we were writing this in the spring of 2025 was the emergence of a Chinese AI startup, DeepSeek, whose open-source R1 reasoning model performed as well as or better than ChatGPT-4o and other frontier models available at the time, but was reportedly trained for a tiny fraction of the money and energy poured into the leading U.S. models. It is too early to tell whether DeepSeek's innovations are original or whether its cost claims are accurate. But the surprise from China did open up an opportunity to "take stock of unrealized efficiency gains in AI compute before we end up building a bunch of 30-year gas plants," Arvind Ravikumar, co-director of the Center for Energy and Environmental Systems Analysis at the University of Texas at Austin, commented on Bluesky. We believe the best reason to invest more in AI infrastructure and R&D here in the United States is not to secure a stranglehold on the technology, but to ensure that AI develops in ways that are safe, equitable, and energy-efficient.

Given the cost of building and powering hyperscale AI hubs, the sector may continue to be dominated by a handful of companies such as Microsoft, Google, Oracle, OpenAI, and Anthropic. But there will always be room for startup-based innovation and venture-scale investment in the field—especially, we believe, in technologies that could supply more computing capability per megawatt. That's what DCVC-backed MosaicML, which is now part of another portfolio company called Databricks, enabled with a platform that allows organizations to train their own generative AI models two to seven times faster than conventional methods.

Figure 2.1.1

Key sources of electricity demand growth
Data centers and electric vehicles will put massive new stresses on the electrical grid.

Source: Center for Strategic and International Studies

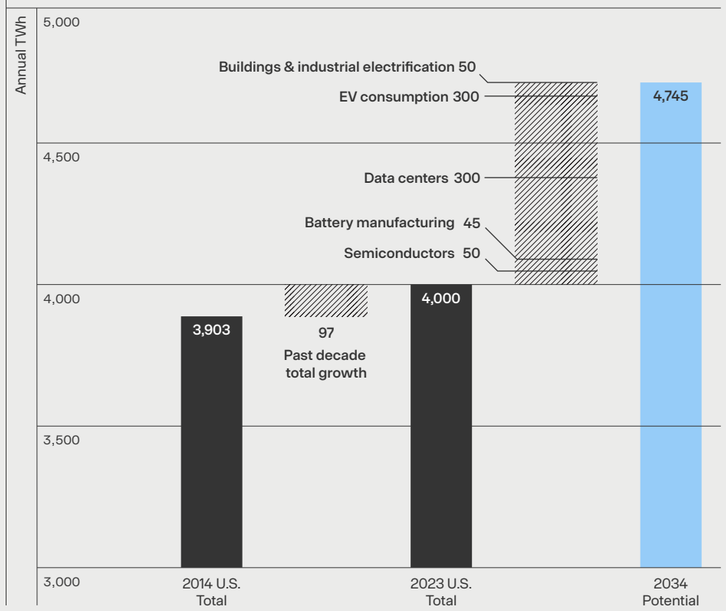

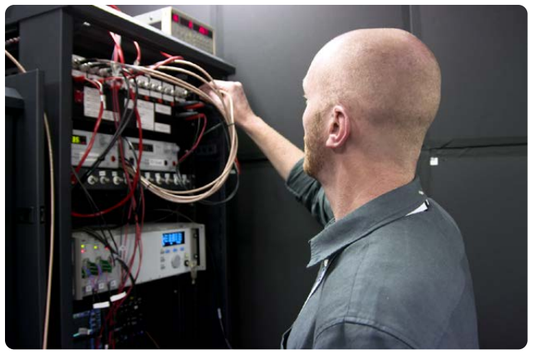

← Q-CTRL
The company builds error suppression software that improves the performance of quantum computers, and has pioneered a new generation of quantum-enabled navigation technology that offers an alternative to GPS.

It's also what Q-CTRL is delivering with its error-suppression software for quantum computers. Q-CTRL, which we've backed since 2018, is helping quantum computing leaders like IBM deal with qubits' stubborn vulnerability to noise, which can cause them to lose the quantum superposition that's key to their huge computing capacity. The company's software detects which qubits in a computation are producing errors and helps operators apply corrections by adjusting the control pulses that manipulate each qubit. Similar error-mitigation methods allowed IBM to prove in 2023 that a 127-qubit circuit could run atomic simulations faster and more precisely than a conventional supercomputer. Error correction opens a potential era of "quantum advantage" in which even noisy quantum computers can provide real value in the near future. (Q-CTRL's technology may have even more immediate applications in the area of quantum sensing, which could be useful for defense, navigation, and energy exploration.)

And already, quantum computers can tackle specialized problems that would be exponentially more expensive, in both time and energy, for classical computers. In late 2024, Google reported that in just five minutes, its Willow quantum chip performed a so-called Random Circuit Sampling (RCS) test that would have taken one of the U.S. government's fastest supercomputers 10 septillion years. (In RCS, a computer runs a series of random quantum operations and then checks for the correct pattern of results; it measures how well a quantum computer can handle complex calculations that would overwhelm traditional computers.) Hartmut Neven, the founder of Google's Quantum AI Lab, predicts that the same scaling advantage that makes quantum computers good at RCS will make them well suited for training certain AI models. "Quantum computation will be indispensable for collecting training data that's inaccessible to classical machines, training and optimizing certain learning architectures, and modeling systems where quantum effects are important," Neven writes. "This includes helping us discover new medicines, designing more efficient batteries for electric cars, and accelerating progress in fusion and new energy alternatives."

But well before we've worked out all the kinks in quantum computers, we'll need to reduce the power demands of classical microchips. That's long been the focus at Mythic, a maker of analog matrix processors that we've backed since 2016. In a traditional digital computer, the parameters used in the training and inference phases of a generative AI program such as a large language model are stored in memory, and must be bused back and forth to a CPU or GPU for processing. In Mythic's chips, parameters are stored directly in the processor, which makes AI computations 18,000 times faster and 1,500 times more energy-efficient. In fact, the more parameters in an AI model, the greater the advantage of Mythic's chips over conventional GPUs, according to Mythic CTO and co-founder Dave Fick. "Our early analysis has indicated order-of-magnitude improvements in each metric of cost, performance, and power over Nvidia and leading LLM startups," Fick writes.

Figure 2.1.2

Hyperscaler U.S./North America Electricity Use
Meta: U.S. data center electricity only. Alphabet: U.S., all electricity. Microsoft: North America, all electricity.

Source: Company reports, Barclays Research

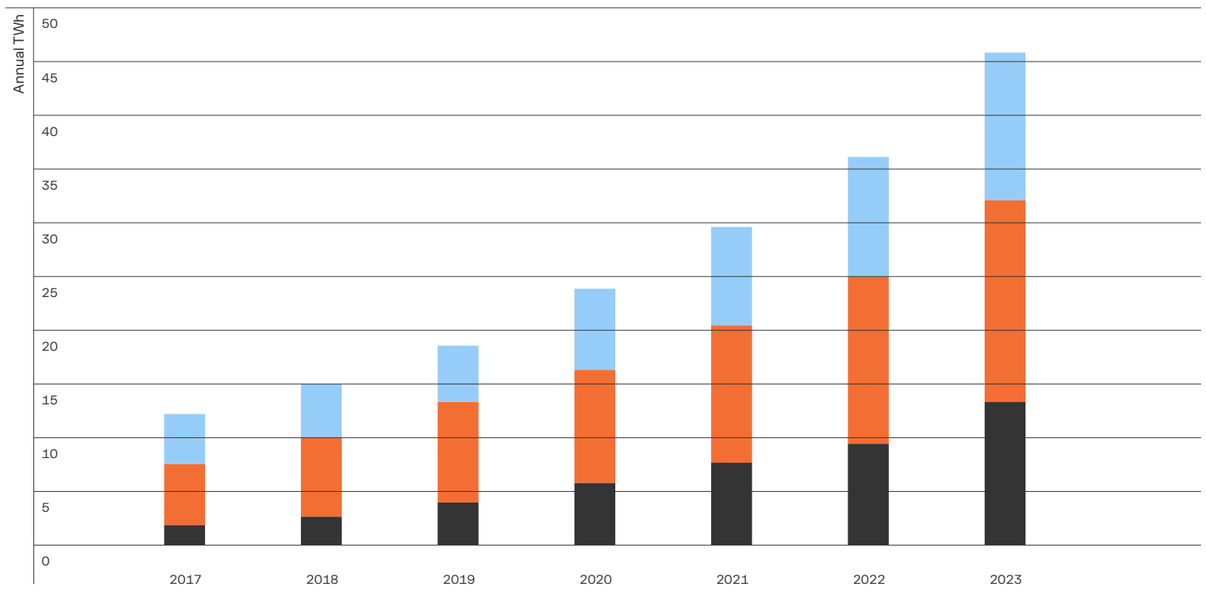

Of course, to measure whether new computing architectures can actually help to lower AI energy demands, we'll need better data from all of the hyperscale computing companies, and we agree with calls from AI scholars for greater transparency and for research into standardized metrics for quantifying AI sustainability. We're also encouraged by progress at companies like Crusoe to build a sustainably powered cloud AI infrastructure, and by xLight's work to make semiconductor manufacturing itself more environmentally friendly. In today's semiconductor fabs, the light source used to etch features onto a silicon surface is laser-produced plasma, or LPP. xLight aims to replace LPP with particle accelerator–driven free electron lasers that reduce a fab's power and water consumption by a factor of 10.

"If the urgency here is for America and the free world to have more compute in more places that works faster and for less power, then we have a whole buffet of options," says DCVC's Ocko. "It could be next-generation chip architectures like Mythic. It could be algorithmic innovations like MosaicML. Or it could be quantum, which is intrinsically a billion to a trillion times more power-efficient per element than conventional chips." AGI itself may or may not be achievable in the near future—indeed, we've argued in the past that it's a shiny object distracting from more transformative work (see p. 24 of our 2024 Deep Tech Opportunities Report). But either way, the U.S. will need more, faster, and greener computing hardware, and that's the conviction guiding our investment in this area.

> "If the urgency here is for America and the free world to have more compute in more places that works faster and for less power, then we have a whole buffet of options."

Matt Ocko
Managing partner, DCVC

Automation is almost always a productivity booster, but the extreme solution—building new "lights-out" factories or warehouses where robots run the show, unimpeded by humans—is affordable only for the largest, richest companies. We look instead for companies developing robots that can work alongside people in existing facilities, handling the slowest and most boring or back-breaking work while freeing humans to do what they do best (namely, thinking and acting flexibly).

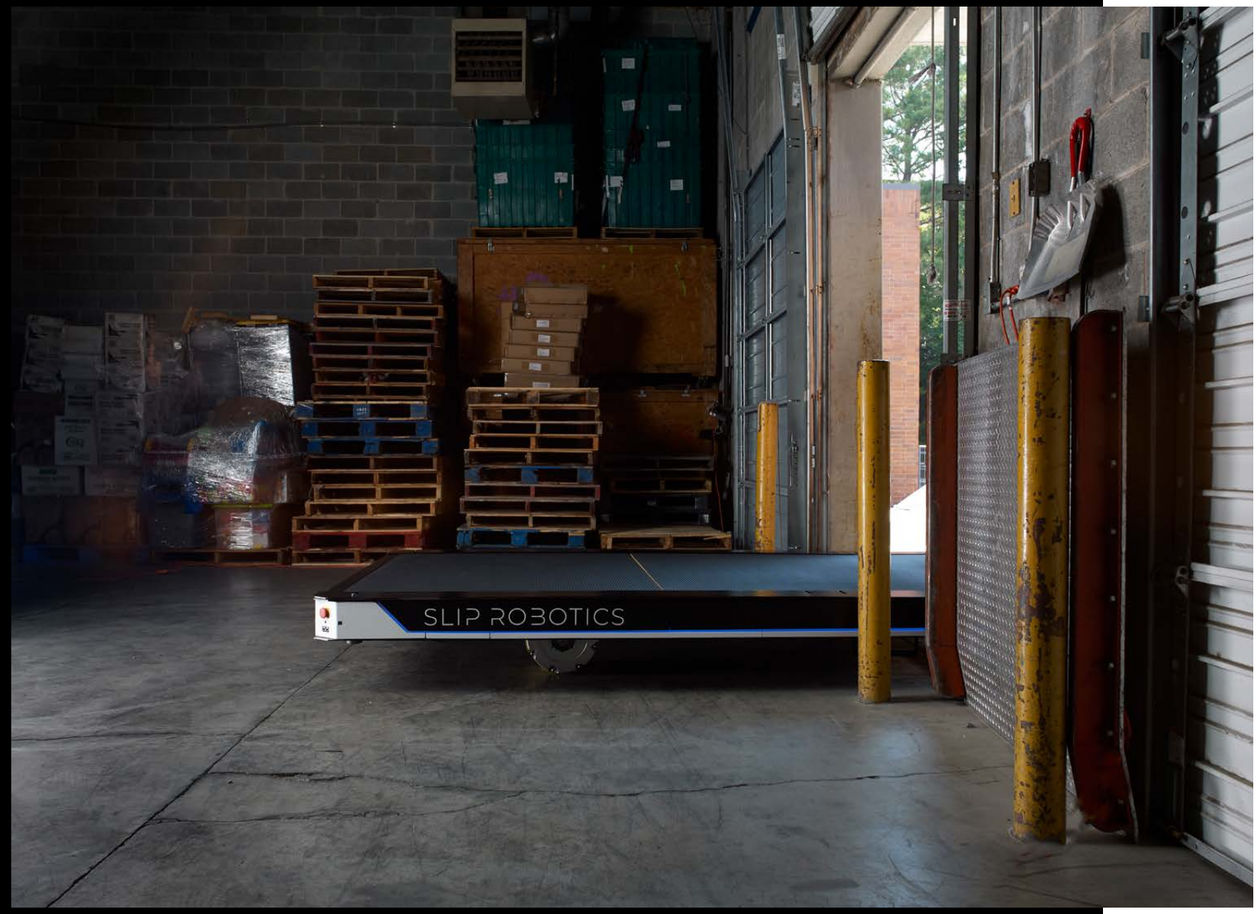

Every day, across the world, millions of trailer trucks back up to loading docks at factories and warehouses. These docks amount to a universal standard interface for the global supply chain, and they aren't likely to be redesigned anytime soon. Unfortunately, they're inefficient and dangerous. Forklift operators traditionally load and empty the trailers one box, crate, or pallet at a time—a job that takes 30 to 60 minutes per truckload. That keeps drivers idle, while also exposing operators to high rates of injury.

The engineers at Slip Robotics understood that the simplest solution to a problem is often the best one. They designed automated pallets to take over the loading dock. These self-guided floors on wheels slide freight into and out of trailers under the control of software and sensors. Three bots, each carrying up to 12,000 pounds, are enough to fill a trailer, cutting an hour-long job to just five minutes and reducing labor and waiting costs from $100 per load to about $15 per load. The SlipBots themselves are still loaded up by forklift operators, but when these workers don't have to drive their forklifts directly into trailers, accident rates go down by 25 percent, freight damage goes down by 40 percent, and equipment damage goes down by 62 percent, according to the company.

"SlipBots are elegantly simple, fast, and easy to deploy, requiring no integration with IT and no complex environmental mapping," says DCVC general partner James Hardiman. In fact, Slip already has two dozen customers in its pipeline, and its goal is to deploy 450 bots by the end of 2025. It's an innovation that could help the logistics industry meet growing pressure for fast delivery without redesigning warehouses or expanding truck fleets. Slip has succeeded by "focusing on a very narrow problem, taking things on and off a truck," says Hardiman. "It's a constrained, well-defined environment, and you can make a very specialized robot for that."

But versatility, too, can be a virtue. In settings built for humans, nothing less than a humanoid robot with hands, legs, and eyes is needed to take over the tedious and strain-inducing jobs humans don't want, such as carrying totes. That's what Agility Robotics' Digit is designed for.

Digit is 5'9", weighs 200 pounds, and can carry loads up to 35 pounds. Since mid-2024, the world's biggest contract logistics provider, GXO, has been deploying Digit robots alongside human employees at its SPANX fulfillment center in Atlanta, where they move totes from retrieval robots onto conveyors. It's the logistics industry's first formal commercial deployment of humanoid robots, and the key to its success is Digit's ability to adapt to a complex warehouse environment with changing needs, says Adrian Stoch, GXO's chief automation officer. "Whereas previous technologies have been developed for a very specific, very discrete solution, [Digit] can be unpacking boxes today and unloading a trailer tomorrow and assisting our associates with unloading heavy totes the next day," Stoch says.

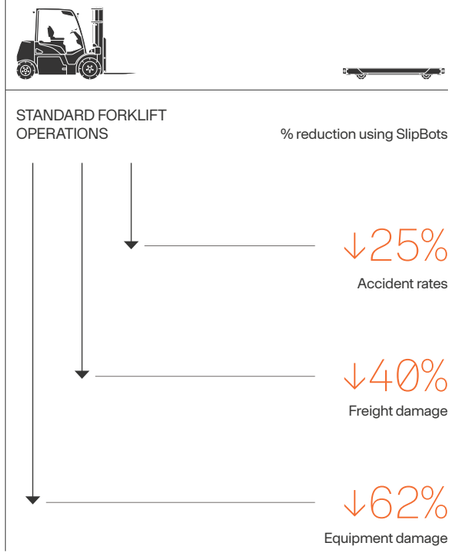

> **Whereas previous technologies have been developed for a very specific, very discrete solution, [Digit] can be unpacking boxes today and unloading a trailer tomorrow and assisting our associates with unloading heavy totes the next.**

Adrian Stoch
Chief automation officer, GXO

That kind of versatility requires a lot more planning. A cloud automation platform called Agility Arc monitors the Digit robots and integrates their work into existing warehouse management systems. But in the future, any organization that needs to move lots of material—think of military logistics flights, for instance—will likely want tireless robot workers like Agility's as part of its hybrid workforce. We believe Agility, with its deep understanding of the physics of motion and the requirements of real workplaces, is the best-in-class company in this market.

We also think that as the people engineering robotic control systems begin to incorporate ideas from generative AI, robots will take up even more roles in human society. The maturation of Windows in the early 1990s marked a "Microsoft moment" when much of the value in the personal computing market suddenly shifted from hardware to operating systems. It seems inevitable that robotics will go through a similar shift, as robots begin to learn from experience and that experience gets encoded in software.

For one thing, the physics-based models that control actuation and movement in most current robots will likely be supplemented or replaced by nimbler predictive models. In this area, we're carefully following the work of the San Francisco startup Physical Intelligence, which is solving challenges such as tokenizing robot actions in a way that allows these actions to be processed by transformers like those at the heart of today's largest AI models. The company believes it can create generalist models for robot control simply by training dexterous robots on everyday tasks such as folding laundry, busing tables, or filling grocery bags.

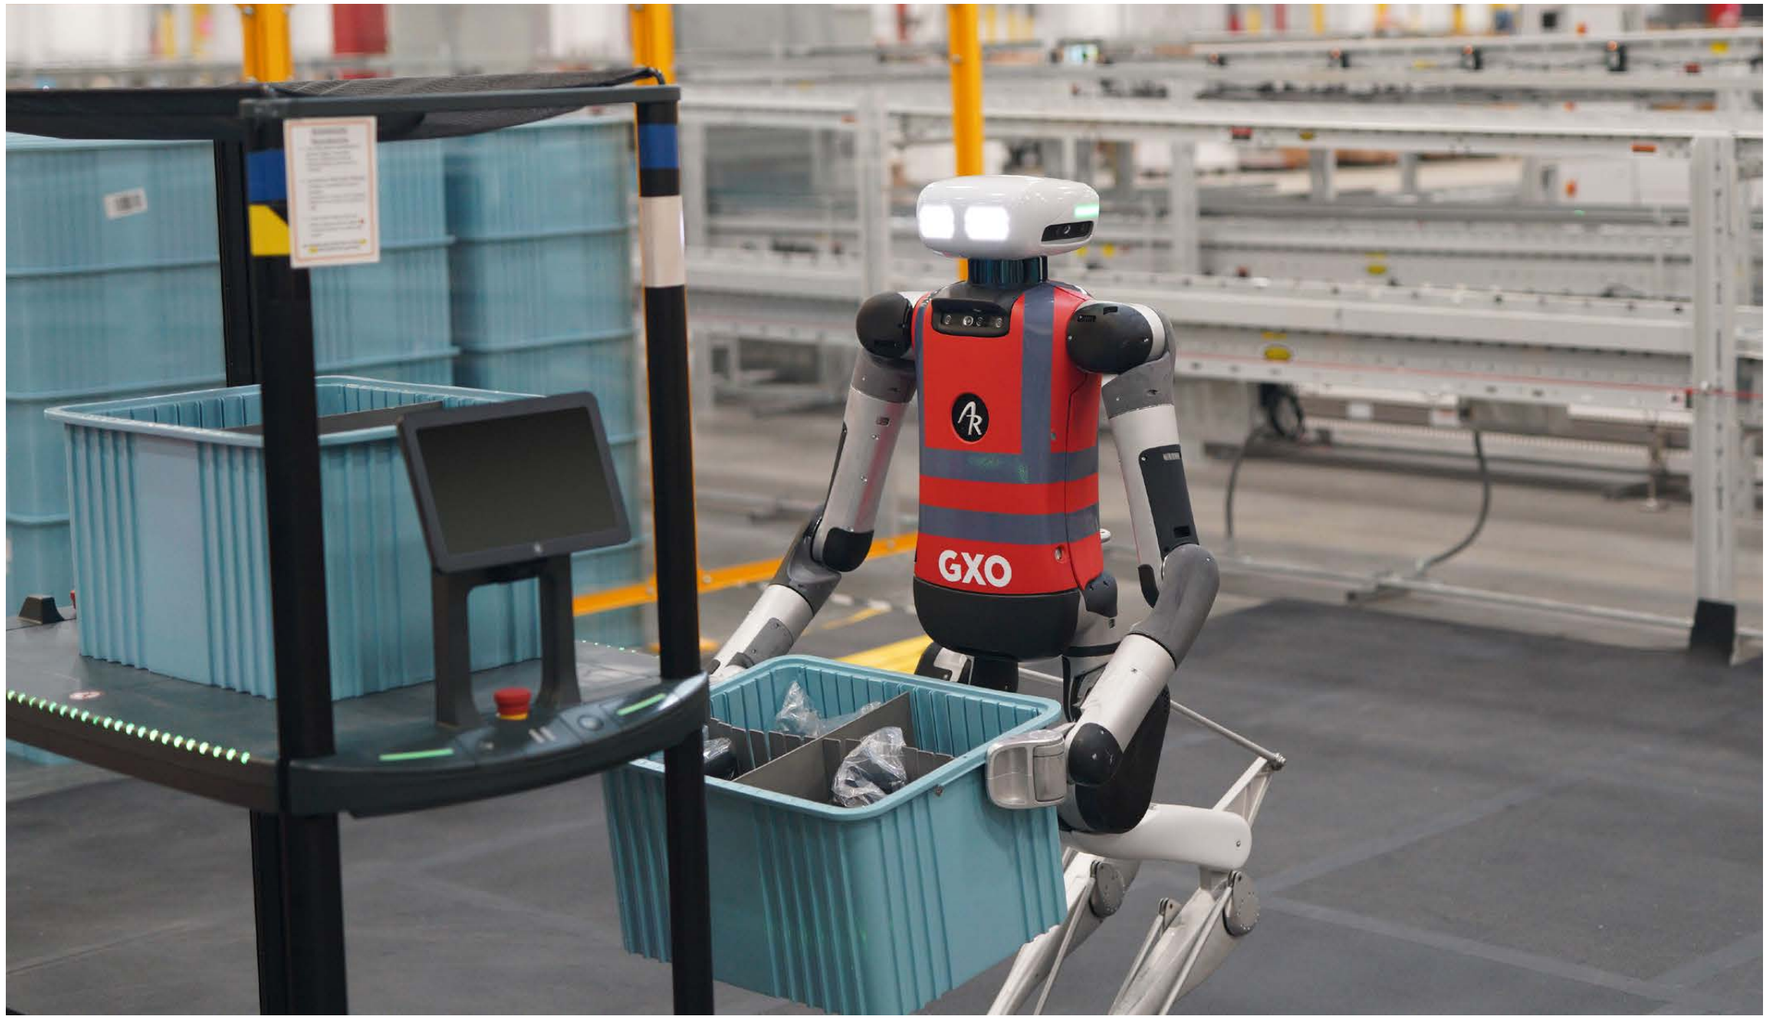

↗ **Agility Robotics'** bipedal Digit robots are hard at work hauling totes in fulfillment centers managed by contract logistics provider GXO.

Similarly, Skild AI is building a foundation AI model as a general-purpose brain for diverse types of robots, from mobile manipulators to quadrupeds to security robots. "With our model, any kind of robot can be agile, dexterous, and safe to interact with people," the company says. Skild's long-term goal is to develop artificial general intelligence that's "rooted in the physical world, challenging the popular notion that AGI can solely arise from digital knowledge."

In the field of intelligent robotics, the advantage will go to companies with a built-in way to gather large amounts of proprietary training data—an area where Agility happens to shine (see the 2024 Deep Tech Opportunities Report, p. 22). In the future, robots will thrive not in depopulated factories or warehouses but alongside us in our homes and workplaces, doing the tasks we teach them and helping us to stay safe and productive.

↓ **Skild AI** makes machine learning algorithms that could help security and inspection robots and other types of bots learn their way around the world.

Chat-based generative AI models such as OpenAI's GPT-4o, Google's Gemini, and Anthropic's Claude are changing the way we search for knowledge. And customized AI models based on similar machine-learning technology are giving researchers new ways to find long-hidden patterns in biology and many other fields. But...

there are certain things chatbots can't and shouldn't be used for. We're not fans of—or investors in—companies touting generative models as replacements for humans in the creative arts, counseling, therapy, or clinical care.

↳ 2024 was the year large AI models went multimodal. In December OpenAI released Sora, a paid ChatGPT feature that generates lifelike short videos based on text prompts like "In a pastel bathroom with a rubber ducky, an adorable dragon made entirely of shampoo bubbles." Days later Google's DeepMind rolled out a similar video-generation tool called Veo 2, and smaller companies like Pika and Luma AI offered their own variations. In essence, these tools make high-quality computer-generated graphics available to average users, in part by taking over many of the tasks traditionally performed by human visual effects artists.

Make no mistake, generative videos can be arrestingly realistic. Who doesn't enjoy seeing Rockefeller Center overrun by golden retrievers? But what these eye-popping videos are usually missing is the spark of human emotion or inspiration. "AI can be an enabler," says DCVC general partner Ali Tamaseb. "It can help a moviemaker create visual effects. But ultimately you still need the human to write the story and to direct the movie." (After the 2023 Writers Guild of America strike, the big Hollywood studios agreed, signing a contract that says AI models aren't writers and that nothing they can create can be considered literary material.)

A similar point applies to AI models meant to function as stand-ins for psychologists, therapists, or physicians, such as Pi, Woebot, Wysa, Youper, talk2us, or even ChatGPT. Some of these models are trained on validated therapeutic techniques such as cognitive behavioral therapy or mindfulness meditation. And the chat screen itself can be a useful space for self-reflection, almost like a paper journal that talks back to you. But these models can be alarmingly sycophantic, and they obviously lack true empathy; the insights of a stochastic parrot are, by definition, canned. It should be no surprise that chatting with a bundle of weighted features and equations, even one trained on trillions of parameters, is less nourishing than connecting with a single human soul.

"I think what a patient wants in a therapeutic or medical encounter is the actual human being, not necessarily the conversation," Tamaseb says. "When you go see the doctor, you get 20 percent better just thanks to the placebo effect. An LLM can offer much more content than your therapist or doctor might, but I don't think it will ever replace the impact of having a real human talking back to you."

That's not to say that we don't see a need for better treatments for depression and anxiety, conditions that, beyond their inherent misery, cost the world's economies 12 billion workdays and $1 trillion per year in lost productivity (see Opportunity 7.4). AI-based therapy offers a seemingly affordable and scalable way to help people learn better coping skills—but it's a technology that's still in its nascent stages, lacking clinical validation or even consistent reporting standards.

In short, AI-based creative tools and AI-based therapy are fields that will no doubt attract venture-scale investments, but for now we plan to steer away, especially when new technologies purport to replace what's deeply and essentially human. AI is extremely useful in high-touch professions as an *adjunct* to human experience and skill—but we hope it will never come to be seen as a valid substitute.

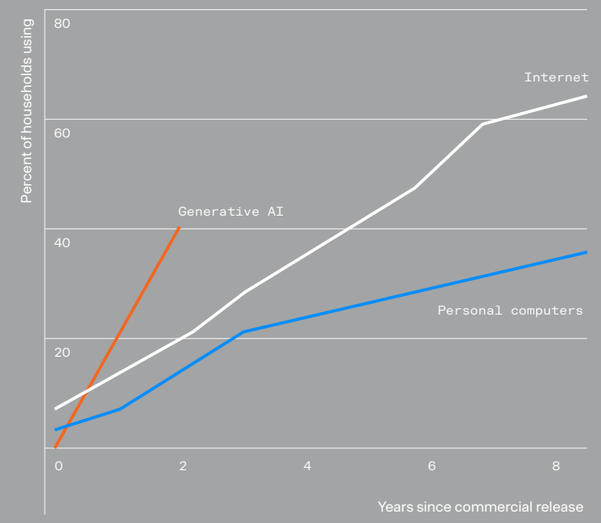

Figure 2.3.1 **Chatbots' explosive growth**

U.S. households have adopted generative AI at a startling pace. Source: IEA.

Halter's smart collars train grassfed cattle to stay inside a movable virtual fence line, allowing "regenerative grazing" that promotes carbon sequestration.

CHAPTER 3.0

Clean technologies are at the center of industrial revitalization.

In 2023 we began making investments from our first dedicated climate fund, DCVC Climate. The fund is designed to help climate-tech companies at an often overlooked, underfunded stage in their life cycles: *industrialization*, when they must make the leap from the demo or prototype stage to commercial scale. Our goal is to help these companies achieve the unit economics needed to prevail over incumbents in their sectors.

Usually, these incumbents employ older industrial processes that depend on fossil inputs. The companies we back at DCVC Climate provide alternatives that are cleaner, cheaper, and more efficient, and that create new resources and stronger supply chains closer to home. In that sense, our climate investing work and our focus on industrial revitalization fully overlap and reinforce one another (see Chapter 1). "The industrial renaissance is about decarbonization," says DCVC general partner Milo Werner. "Advanced manufacturing is solving climate change. They are all the same thing."

OPPORTUNITIES

COMPANIES

Blue River Technology, CH4 Global, ElectronX, Equilibrium, Fervo Energy, Fourth Power, Halter, Oklo, Pivot Bio, Radiant, Sabanto, Twelve, Verdant Robotics

VOICES

Zachary Bogue, Matt Ocko, Dr. Rachel Slaybaugh, Ali Tamaseb, Milo Werner

A two-decade stretch of near-zero growth in electrical demand in the United States is definitively over. In Opportunity 2.1 we noted that the data centers running today's frontier AI models seem set to consume ever-increasing amounts of electricity. Indeed, Barclays Research projects that the energy needs of U.S. data centers will grow by 14 to 21 percent per year through 2030, tripling their overall consumption from 150-175 terawatt-hours (TWh) in 2023 to 560 TWh in 2030, equivalent to 13% of current U.S. electricity usage. But AI is only one of the many emerging sectors that will need more power—a lot more—as the nation builds new industrial capabilities and the economy transitions away from fossil energy. To name a few others:

Steelmakers are under pressure to reduce carbon emissions by switching to new chemistries and energy sources. A single green steel plant under construction in Ohio, if run exclusively on clean hydrogen, would require 160 kilotons of the gas per year. Making that much hydrogen would consume 8.3 TWh of electricity, more than the annual electrical usage of the state of Vermont.

Generative AI models aren't the only computational workloads adding to electricity demand—there's also cryptocurrency mining, which the U.S. Energy Information Administration estimates already represents 0.6 percent to 2.3 percent of all U.S. electricity consumption. This figure seems likely to grow in an era of relaxed crypto regulation.

Consumer demand for electric vehicles may soften if the Trump administration revokes tax credits for new EV purchases, but analysts expect growth to persist as costs come down and the range and utility of the vehicles improve. EV charging will, in turn, increase the burden on electrical grids. One study focusing on Texas' grid projected that after 2030, in the absence of strategies for increasing electricity supply or manage peak loads, EV charging will lead to frequent blackouts and brownouts in the state.

The production of clean hydrogen for energy storage and industrial processes depends on electrolysis. The International Energy Agency (IEA) forecasts that by 2027 this process will claim about 2 percent of the world's renewable energy capacity.

Cement manufacturers are also exploring electrification of their kilns as a way to reduce carbon emissions. DCVC-backed Brimstone is a leader in the reinvention of cement. Its calcium silicate-derived Portland cement actually absorbs more carbon than it releases, and "the cleaner the kiln firing, the more carbon-negative the cement," says Dr. Rachel Slaybaugh, a partner at DCVC.

To offset dwindling supplies of fresh groundwater and surface water, more countries are turning to desalination of seawater or brackish water, which requires at least 1 kilowatt-hour of electricity per cubic meter of water treated. In 2024 the IEA projected that energy demand for desalination will double globally by 2030.

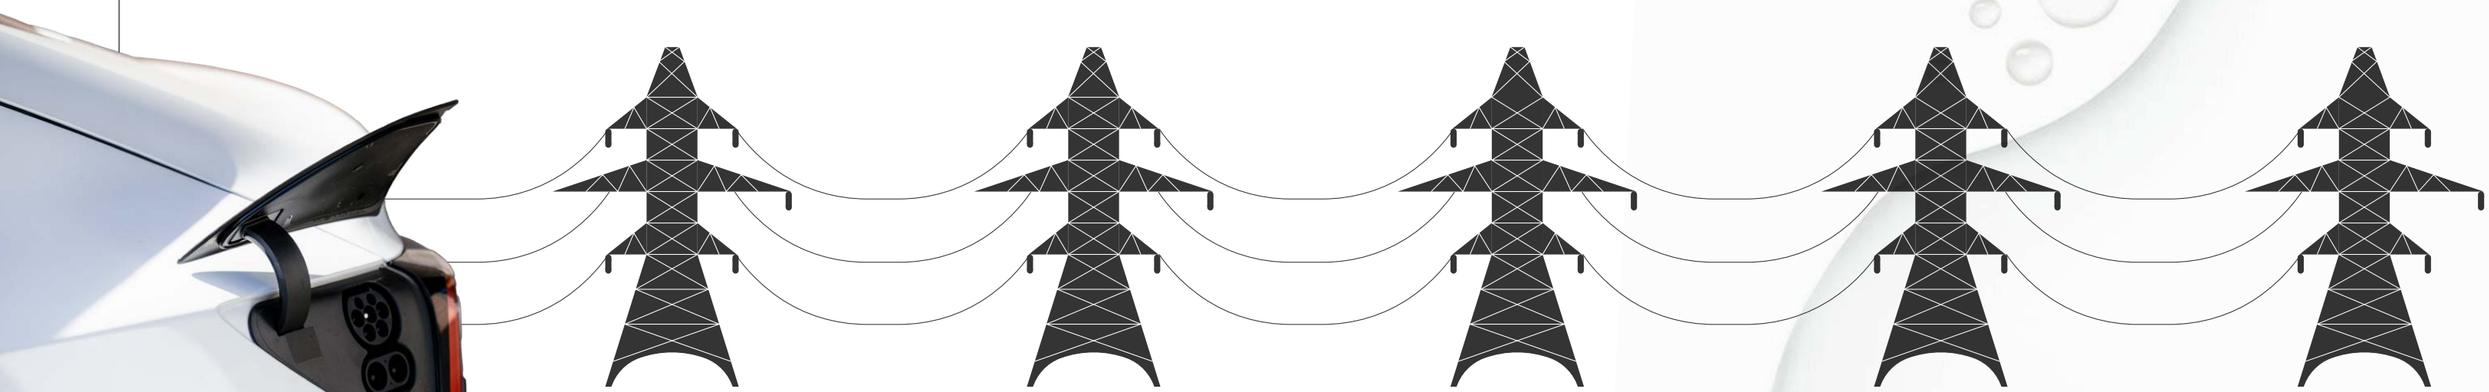

Meeting this rising electrical demand will require a range of generating solutions tailored to the specific needs of each region and industry. Building new coal- or gas-fired plants should be a strategy of last resort, due to their high CO_2 emissions. And thanks to years of innovation and hard work, the economics are beginning to shift against fossil plant construction. It's now far cheaper to build and operate a new solar, wind, or geothermal facility than to build a coal plant. Most renewable sources can't yet match the cost-effectiveness of combined-cycle gas-fired plants, but onshore wind farms and solar photovoltaic facilities are catching up fast, according to a 2024 report commissioned by the U.S. Energy Information Administration.

DCVC Climate backs companies with bold, proven ideas for continuing to bring down the costs of other types of zero-carbon energy, including predictable baseload power. One of our most exciting investments is Fervo Energy, a leader in enhanced geothermal energy. There's an astonishing amount of heat energy trapped below our feet—enough to power all human activity for millennia, according to one MIT report. Yet this heat is difficult to tap, and so geothermal energy currently accounts for only 0.4 percent of electrical generation in the U.S. Fervo uses techniques adapted from the oil and gas sector, including horizontal drilling, hydraulic fracturing, and fiber optic sensing, to reach pockets of geothermal heat previously considered inaccessible.

↗ **Fervo Energy** is the leader in enhanced geothermal energy, using techniques borrowed from oil and gas development to reach hotter, deeper rock.

Fervo activated its first commercial well in 2023, achieving record flow rates and supplying 3.5 megawatts of electricity to a Google data center in Nevada. Now the company is in its scale-up phase, building a 500-megawatt power plant called Cape Station in southwest Utah. Fervo has already lined up hundreds of millions of dollars in new equity financing, credit, and loans—not to mention purchase agreements with big customers who are eager for sources of clean firm power.

Figure 3.1.1

Levelized Cost of Energy Renewable and conventional energy sources are cost-competitive, after accounting for factors such as capital and operating costs.

Source: Lazard.

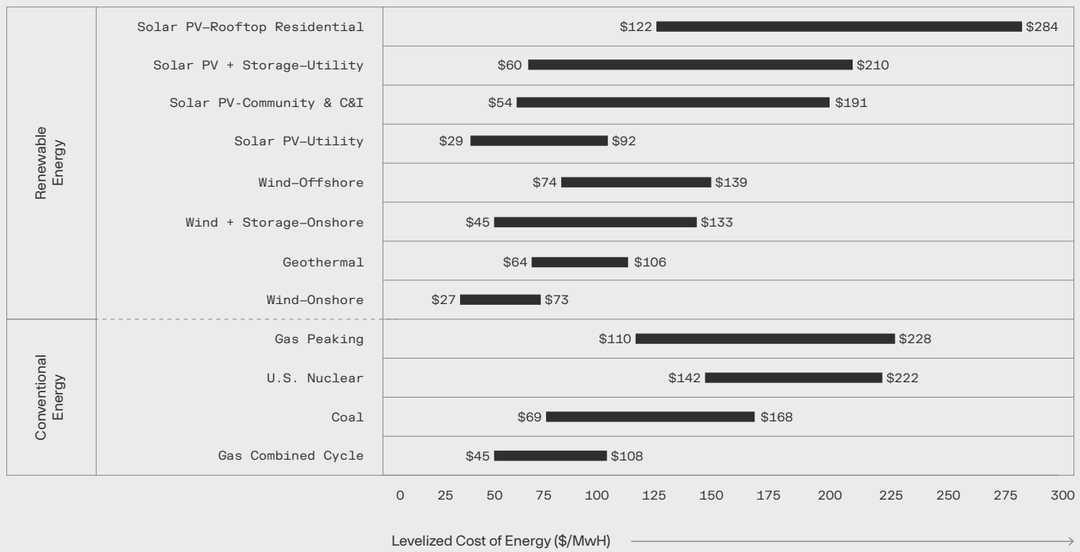

> [Fervo is] proving that they can capitalize these projects in a sustainable way. In 2026 they'll be able to put megawatts on the grid. They have a long project pipeline, with a big jump on lease holding—which is great, because now a lot of people are trying to get geothermal leases. Piece after piece, they're doing the things they said they were going to do.

Rachel Slaybaugh
Partner, DCVC

"Their first project has operated for over a year now, and they've proven it's steady and doesn't have thermal declines outside of what they predicted," Slaybaugh says. "They're proving that they can capitalize these projects in a sustainable way. In 2026 they'll be able to put megawatts on the grid. They have a long project pipeline, with a big jump on lease holding—which is great, because now a lot of people are trying to get geothermal leases. Piece after piece, they're doing the things they said they were going to do."

Geothermal power is attractive in part because it's persistent, regardless of the weather or the time of day. The average geothermal plant can operate at 90 percent of its peak production potential, compared to 37 percent for wind and 11-13 percent for solar. But there are also other technologies that promise to provide predictable power and smooth out the bumps in the supply of wind and solar energy. One is thermal storage—and in this area we're backing a company called Fourth Power, built on the work of MIT mechanical engineering professor Asegun Henry.

The idea of thermal storage is to siphon electricity from the grid during periods of excess production and low prices, store it for hours or days in the form of heat captured in large graphite blocks, then turn it back into electric power when demand and prices are higher. To make all of that practical at scale, Henry developed an innovative liquid-metal circulation system that transfers thermal energy from electric heaters into the graphite blocks, and then carries the energy back out when needed to thermophotovoltaic towers, where it gets converted back into electricity. The Massachusetts-based company is currently testing each of these components in the lab, while getting ready to build its first commercial pilot plant.

The materials that go into Fourth Power's design, like tin and graphite, are far cheaper than those needed for competing storage technologies such as lithium-ion batteries. That's why the company expects it will be able to deliver electricity for $25 per kilowatt-hour (when paired with a renewable power source), less than a tenth the cost of power from grid-scale chemical batteries—and on par with the cost of power from gas-fired generating plants.

→ **Oklo** The Aurora fission reactor, fueled by recycled waste, will produce up to 75MWe of power.

Since Fourth Power can source most of its materials domestically, it's also in a good position to ride out any looming trade war with China or other countries. "If there are really high tariffs, Fourth Power's value proposition gets even stronger, because lithium-ion batteries will become more expensive," Slaybaugh explains.

Soaring electricity demand has also revived conversations about one more source of clean firm power: nuclear. In September 2024, Microsoft struck a 20-year deal with Constellation to buy power for its data centers from a revived Unit 1 reactor at Three Mile Island. And in January 2025 NextEra Energy said it may restart its Duane Arnold nuclear plant in Iowa to feed power to data centers in the Midwest, while TerraPower, founded by Bill Gates, announced it had inked an agreement to sell power from its advanced nuclear project in Wyoming to Sabey Data Centers. Analysts at Morgan Stanley say there's a "nuclear renaissance" on the way that will double global nuclear capacity and soak up $1.5 trillion in new capital investment through 2050.

But not all of that money can or should go into 100-megawatt-plus reactors, which are still notorious for their high capital expenses and operating expenses. Another way to power data centers or other key facilities is to build smaller, more efficient reactors on site.

Oklo—which went public last year, six years after DCVC's first investment in 2018—says its first Aurora liquid metal–cooled fast reactor will come online before 2030. These reactors, which use nuclear waste as fuel, are designed to provide 15 to 50 megawatts of power, perfect for a data center. Oklo said in December that it has signed an agreement with Switch, an operator of large data centers in Georgia, Michigan, and Nevada, to deploy enough Aurora reactors to generate up to 12 gigawatts of power through 2044. And it said in January that it has inked a deal with natural gas generator supplier RPower to roll out a "phased power model" for data center customers even sooner. Under the deal, gas generators deployed to data center sites by RPower would be replaced by Aurora powerhouses as they become commercially available.

→ **Radiant Nuclear** The Kaieidos portable microreactor fits on a truck.

An even smaller and more portable fission reactor called Kaleidos is under development at DCVC-backed Radiant. Kaleidos will be small enough to fit on the back of a semitrailer and will generate about 1.2 megawatts of power, making it a suitable power source for a small hospital, a military facility, or an emergency command center, replacing dirty and noisy diesel generators. Since we last discussed Radiant in our 2024 Deep Tech Opportunities Report, the company has secured $166 million in new funding in a Series C financing round led by DCVC; successfully completed a passive cooldown test demonstrating that Kaleidos can safely shut down and cool off without power; and finalized plans to test a Kaleidos prototype at U.S. Department of Energy's Idaho National Laboratory in 2026.

"Radiant is exciting because they're showing that you can get something done in a reasonable amount of time with a reasonable amount of money," Slaybaugh says. "Their team is on fire to go build stuff. 'Let's get building, let's get moving' is a new narrative in an industry that has maybe struggled to do that."

Now that the AI revolution and the electrification of industry and transportation are creating unprecedented power needs, the demand for zero-carbon electricity is finally catching up to the vision of innovators in areas like next-generation nuclear power, geothermal power, and thermal storage. We've been proud to back these innovators for many years, and we can't wait to see what kinds of further advances their ideas unlock.

$1.5 T

Expected new capital investment in nuclear power through 2050, according to Morgan Stanley

Chains of carbon atoms form the backbone of every molecule of hydrocarbon fuel, from ethane to methane to kerosene. Taking these atoms out of the ground by drilling for oil and gas ultimately adds to the net burden of greenhouse gases in the atmosphere. But carbon-neutral fuel is becoming a reality, as deep-tech innovators figure out how to pull the carbon we need for jet travel and other hard-to-decarbonize sectors from industrial sources and from the atmosphere itself.

Nature has an efficient way to use sunlight, air, and water to create hydrocarbons: it's called photosynthesis. Twelve, which we've backed since 2018, invented an industrial equivalent of photosynthesis in the form of a "black leaf" electrolyzing membrane embedded with a novel CO_2-reducing catalyst. Pumping CO_2 and H_2O through an electrified stack of these membranes yields CO and H_2. This combination is known as syngas, and it can be used as a feedstock to make many other hydrocarbons, including jet fuel. If the power for the electrolyzer comes from renewable sources, and if the CO_2 comes from industrial facilities where it otherwise would have been vented into the air, then the overall process is highly carbon-negative.

→ **Twelve**
E-Jet sustainable aviation fuel from Twelve is an identical drop-in replacement for conventional fuel, at 90% lower emissions.

At a plant under construction in Moses Lake, Wash., Twelve plans to produce up to a million gallons per year of its E-Jet sustainable aviation fuel. Last year the company signed a 14-year agreement to supply 260 million gallons of E-Jet to International Airlines Group (IAG), the holding company for British Airways, Iberia, and Aer Lingus. DCVC's Zachary Bogue calls this agreement "a significant vote of confidence" in Twelve's technology, but adds that it also signals a shift in the way we perceive and use CO_2. "By converting a destructive byproduct of industrial activities into something as valuable as aviation fuel, Twelve is leading a paradigm shift in redefining the possibilities of carbon transformation," Bogue says.

Circularity Fuels is also contributing to that redefinition. The company, which was incubated here at DCVC and is led by entrepreneur-in-residence Stephen Beaton, has developed a reactor called Ouro that can take CO_2 out of the atmosphere or an industrial waste stream and turn it directly into almost any hydrocarbon fuel or feedstock. The reactor combines sorbents and proprietary catalysts to reduce the number of steps involved in carbon conversion, leapfrogging the complex and expensive refining processes used by most other electrofuel makers.

Circularity's first product is high-purity methane. This is the most expensive form of methane to make if you're starting from fossil sources, and Circularity can already produce it at lower than the prevailing cost. It's selling this methane to companies that use it to grow diamond, graphene, and other advanced carbon materials.

Beaton wants to scale up manufacturing of the Circularity reactors so that the company can offer electrofuel methane as a carbon-neutral, drop-in replacement for natural gas, which we already know how to store and transport efficiently. Says Bogue: "This type of drop-in replacement is crucial for transitioning off fossil fuels, because the world doesn't have a century to build out transportation and storage networks for other energy-carrying molecules like hydrogen, ammonia, or methanol."

There's carbon all around us, in the air and in our factories' waste streams. As soon as companies like Twelve and Circularity Fuels have perfected cheap ways to extract it and make it into useful chemicals, we'll be able to leave fossil carbon where, given the option, it belongs—in geological formations deep underground.

> By converting a destructive byproduct of industrial activities into something as valuable as aviation fuel, Twelve is leading a paradigm shift in redefining the possibilities of carbon transformation.

Zachary Bogue
Managing partner, DCVC

The increasing volatility of electrical supply and demand means it's time to invest more in grid-stabilizing technologies.

As extreme weather events become more commonplace due to human-induced atmospheric warming, the fragile and aging U.S. electrical grid will come under increasing stress, even as power providers struggle to meet rising demand from data center growth and the electrification of industry (SEE OPPORTUNITY 3.1). That will put a premium on technologies that help smooth out supply and demand, whether through energy storage or digital arbitrage.

In each of the last four decades, the frequency, severity, and cost of floods, cyclones, droughts, heat waves, wildfires, freezes, and compound weather disasters have steadily risen (see Figure 3.3.1). Between 1980 and 2024, the United States suffered an average of nine weather events per year with losses exceeding $1 billion; but in 2024 alone there were 27 such disasters. The U.S. energy grid, much of which was built in the 1960s and 1970s, wasn't designed to withstand these kinds of challenges, which can cause catastrophic damage to pipelines and transmission lines. One harbinger was the February 2021 deep freeze in Texas, which shut down power to the natural gas supply chain and cut off fuel to power plants, leaving millions without electricity for days and causing 246 deaths, mostly from hypothermia. The Department of Energy warns that these cascading failures will grow more frequent unless the U.S. invests in making the grid smarter and more resilient.

One way to do that is to build energy storage facilities that can charge up when electricity is cheap and abundant and feed electricity back into the grid when it's needed—whether to supply emergency power during outages, shore up supply at times of peak demand, or simply fill in for renewable power when the sun isn't shining or the wind isn't blowing. Thanks to innovation and economies of scale in the electric vehicle business, the price of lithium-ion batteries has been dropping rapidly, making it cheaper for companies to lash hundreds or thousands of batteries together to build such utility-scale storage facilities.

Figure 3.3.1

Billion-dollar weather disasters, 1980-2024.
The total cost of the 403 tracked weather events exceeds $2.9 trillion—with $183 billion in damage occurring in 2024 alone.

Source: NOAA.

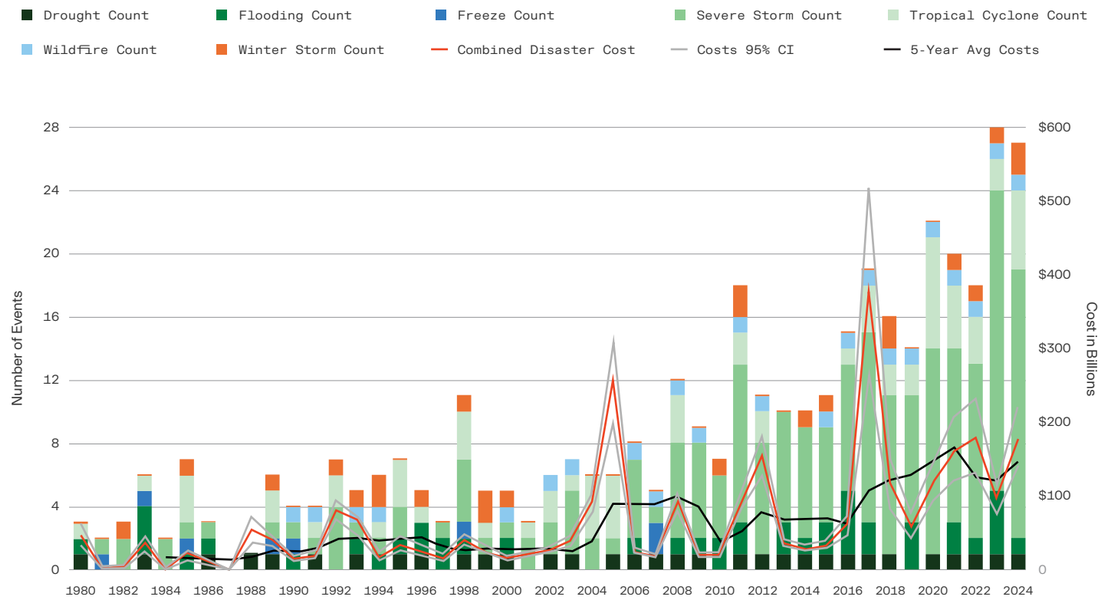

→ In February 2021, a deep freeze in Texas crippled natural gas power plants and caused hundreds of hypothermia deaths.

> Unless you have the real-time data, it's hard to understand how the grid is actually performing … With grid transparency you increase the capability of simulating it and managing it more effectively.

Milo Werner
General partner, DCVC

Even so, there's a need for more financing for these large battery projects. That's where Equilibrium Energy comes in. The Texas-based company was founded in 2021, the same year as the state's grid failure. CEO Ryan Hanley, a veteran of Tesla's grid battery division, saw a need for a new type of contract to accelerate battery deployment, similar to the power-purchase agreements that have boosted progress in the wind, solar, and next-generation nuclear industries. Under Equilibrium's so-called tolling agreements, the company pays battery owners for the right to manage their facilities and bid the power into wholesale electrical markets. Sophisticated predictive modeling of supply and demand allows the company to buy and sell this power at the most profitable times.

A facility under Equilibrium's management in Texas was the state's top-performing battery by revenue in 2024, and now the company is offering more types of energy portfolio management and expanding into California, where it expects that its combination of power-system modeling and machine-learning techniques will offer battery owners predictable revenue and make grid storage look like a safer investment. "Unless you have the real-time data, it's hard to understand how the grid is actually performing," says DCVC general partner Milo Werner. "With grid transparency you increase the capability of simulating it and managing it more effectively."

(We should note that grid-scale batteries like those Equilibrium manages have one significant limitation: their storage duration is only a few hours per charge. In Opportunity 3.1 we described a promising alternative, namely graphite-based thermal storage systems like Fourth Power's, which can hold energy efficiently for up to 100 hours. Thermal storage is perfect for offsetting the variable nature of sun and wind energy, and also for blunting the effects of extreme weather events. "We see Fourth Power's solution as critical because it enables the continued growth of variable power, and complements our existing investments," says DCVC's Slaybaugh.)

Another DCVC-backed company working to counter energy market volatility is ElectronX. In regions like California and Texas where there are lots of wind and solar facilities, variable factors such as bad weather and the number of hours of daylight can cause spot prices for electricity to swing wildly. "In the same day, it can go from a couple hundred dollars [per kilowatt-hour] on the negative side to a couple hundred dollars on the positive side," notes DCVC's Ali Tamaseb. "It's even more volatile than crypto. And one of the reasons is that there aren't enough financial products to smooth out the markets." ElectronX is building a digital financial exchange, currently awaiting approval by the Commodity Futures Trading Commission, that will offer short-duration derivatives designed to help investors hedge against electricity price risks on an hour-by-hour basis. "It enables people to put more money to work building renewable assets—battery storage, solar, everything—if they know they're not going to get screwed," Tamaseb explains. "It makes it a normal, predictable business. Anything that helps the grid to transition to this new world of renewables is good in our minds."

Grid management and optimization is an example of a DCVC investing area that makes sense even in an era of shifting government policy on renewables and clean energy. "The core thesis at Equilibrium and ElectronX is that the electrical grid is volatile and will keep getting more volatile, and that is still true," Slaybaugh says. "We don't see a real risk for them in a future without the Inflation Reduction Act or with tariffs. In fact, if there's less financial support for clean energy, it becomes even more important for those existing assets to operate effectively—and so the need for grid management goes up."

Methane (CH_4) is both an important carrier of energy and a powerful greenhouse gas—so we need smarter, more efficient ways to obtain and contain it. We invest in technologies for making methane using carbon extracted from the air, and in innovators working to reduce agricultural emissions of methane and help the oil and gas industry spot and reduce methane leaks.

In Opportunity 3.2 we introduced DCVC-incubated Circularity Fuels, whose mission is to find efficient ways to use the carbon in the atmosphere and in industrial waste streams to make hydrocarbon fuels and feedstocks—starting with methane. This molecule's advantages are that its carbon-hydrogen bonds carry so much energy, and that it can be stored and transported so easily and cheaply. Its huge downside is that when released into the atmosphere it absorbs infrared radiation—trapping 120 times as much heat as CO_2 does—meaning that we can't afford to keep freeing more of it from fossil deposits and letting it leak away.

Circularity's idea is that if we're going to keep using methane and other hydrocarbon fuels, we should make them from the carbon that's already all around us. The company's Ouro Reactor uses renewable electricity to make CO_2 from the air or from industrial sources into CH_4 and other hydrocarbons. In a way, it's the inverse of an internal combustion engine. "Instead of burning hydrocarbons to make energy, we're taking renewable energy from solar farms or wind farms and turning that energy into hydrocarbon fuels," says Stephen Beaton, Circularity's CEO. At industrial scale, this carbon-neutral process would leverage the large existing network of natural gas pipelines and storage tanks, and the company "would be able to take advantage of the best parts of the fossil fuel infrastructure while avoiding the most costly processes," Beaton says. "That would be the quickest way to get off of fossil fuels."

But there's one important thing to note about that existing infrastructure: it's leaky and needs upgrades. A Stanford-led study published last year in *Nature* estimated that U.S. natural gas wells, pipelines, storage, and transmission facilities emit at least 6 million tons of methane annually, three times greater than the official government estimate. The scale of the waste and pollution is startling. In some areas, such as the New Mexico portion of the Permian Basin, more than nine percent of all methane produced escapes into the air.

Most of the data for the Stanford study was provided by Insight M, a DCVC portfolio company that conducts high-frequency aerial surveys of oil and gas basins and infrastructure. The company's proprietary spectrometers are mounted on small planes and detect methane by measuring the absorption of reflected sunlight. The goal is not to shame big emitters but simply to provide reliable data they can use to identify leaks and "keep gas in the pipe," to quote Insight M. The company says its reports have saved customers $500 million in gas value and helped to keep over 125 billion cubic feet (2.5 million metric tons) of methane out of the atmosphere—which, in CO_2-equivalent terms, is like taking 50 million cars off the road for a year. "The free world needs more energy, but more cleanly, and that's one of the throughlines" both at DCVC and at Insight M, says DCVC's Ocko.

Believe it or not, however, the single biggest source of human-induced methane emissions isn't the oil and gas industry. It's cattle.

→ **Circularity Fuels** is using wind power to make methane from CO_2 in the air.

↓ **Insight M** flies over oil and gas basins to detect methane leaks.

↳ **CH4 Global** grows seaweed to make a belch-reducing cattle feed additive.

The single biggest source of human-induced methane emissions isn't the oil and gas industry. It's cattle.

As grass and other foods ferment in the stomachs of cows and steers, as well as goats, sheep, and buffaloes, methane is a natural byproduct. Globally, these animals belch 80 to 95 million metric tons of methane per year. Fortunately, there's a cattle feed additive, derived from a red seaweed called *Asparagopsis*, that can reduce enteric fermentation by up to 90 percent. CH4 Global, which both DCVC and DCVC Bio have backed since 2021, launched the world's first commercial-scale *Asparagopsis* production facility this year, and it has already won approval to administer its Methane Tamer additive at cattle feed lots in Australia. The company expects that it will help avert a billion tons of methane emissions by 2030. "That's a massive agricultural benefit," Ocko says. "They let people continue to deliver large volumes of healthy food with less impact."

Innovators are finding better ways to lean on nature to remove carbon from the atmosphere, and to measure and finance such projects.

To arrest global warming, it's not enough to stop pouring CO_2 into the atmosphere; we also need to take huge amounts of it *out*. The most ambitious technology-based methods for doing that, such as direct air capture, remain uneconomical. But nature-based carbon removal projects are proliferating. It's a sector that seems ripe for venture capital investment.

The clearing and burning of forests, mostly to create new cropland and grazing land, releases about 8 gigatons of CO_2 into the atmosphere each year. That's more than the total annual greenhouse gas emissions of the United States (about 6.3 gigatons), making deforestation one of the biggest contributors to climate change. At the same time, trees and other vegetation store huge amounts of CO_2 as they grow—about 15 gigatons per year—which means the world's forests are still a net carbon sink. It follows that nations should invest more in efforts to slow deforestation and accelerate reforestation and afforestation (establishing tree cover on formerly non-forested land).

One tree can absorb more than 48 pounds of CO_2 per year.

Governments, NGOs, and private companies have made important moves in that direction over the last two decades, aided by the creation of a voluntary carbon market—originally laid out in the 1997 Kyoto Protocol—that allows entities investing in reforestation and other emissions reductions projects to receive carbon credits or offsets in return. Unfortunately, the carbon credit system has developed serious pathologies. Organizations responsible for setting carbon standards and issuing credits have, in some cases, grossly overstated the actual emissions reductions of approved projects, investigations have found. In other cases, credits have been shown to lack "additionality," meaning the projects and/or carbon savings the credits helped fund would have occurred anyway. Double counting is also a problem—that's when both an entity selling a carbon credit and the entity buying the credit claim the same emissions reduction, creating the illusion that carbon is being eliminated far faster than it really is. And carbon savings aren't always permanent; forests planted with carbon credits burn down, too, putting the carbon back in the air.

To restore trust in carbon credits and get investment flowing faster, organizations such as Verra and the Integrity Council for the Voluntary Carbon Market (ICVCM) are rolling out more rigorous standards for removal projects, built around principles such as quantification and tracking, transparency, and third-party validation. To win ICVCM's approval, a project's promised carbon removals must be additional ("i.e., they would not have occurred in the absence of the incentive created by carbon credit revenues"), permanent, and not double-counted. Verra's new ABACUS label for carbon credits comes with similar requirements.

Reforestation and afforestation, including agroforestry, are on a short list of project types that both ICVCM and Verra consider effective at sequestering carbon. Recently, a group of American technology companies called the Symbiosis Coalition—including Google, Meta, Microsoft, and Salesforce—promised to buy enough carbon credits from developers of nature-based projects to remove up to 20 million tons of carbon from the atmosphere by 2030. The coalition will start by investing in reforestation and agroforestry, using ABACUS as its standard of quality. "Science tells us we cannot wait to invest in nature restoration," writes the coalition's executive director, Julia Strong. "Symbiosis aims to supercharge that necessary climate action now."

And this is where deep-tech innovators come into the picture. It's not easy to measure exactly how much CO_2 is being locked away in a forest as it grows. To help with the quantification and tracking of reforestation and afforestation projects—and, ultimately, to bolster investor confidence and keep the dollars flowing—there's a new crop of companies specializing in measurement, reporting, and verification (MRV). A company called Kanop, for example, applies deep-learning models to interferometric synthetic-aperture radar (SAR) data from satellites to estimate canopy heights and overall forest biomass. Another startup called Nadar fuses multiple types of satellite data, including optical, SAR, and lidar, to monitor changes in forest stock. Upstream Tech offers a geospatial platform called Lens that lets project managers access public and commercial satellite imagery and measure landscape changes over time. All of these companies help project managers minimize field time and bypass older, more error-prone methods for estimating forest biomass, such as manually measuring the diameters of tree trunks.

According to the Nature Tech Collective, venture investors put $1.23 billion into MRV companies between 2018 and 2023. We expect to see a virtuous circle set in, as more reliable measurement technology helps unlock more investment dollars for nature-based carbon removal, in turn encouraging more landowners to undertake removal projects, in turn expanding demand for more and better MRV.

"The only method that's ever been shown to remove atmospheric CO_2 at scale is nature-based solutions, like reforestation done right," says DCVC's Bogue. "We're seeing people trying to bring out new types of data products that will actually quantify that whole market. So there's some activity here, and it's an area I'm actively watching."

> The only method that's ever been shown to remove atmospheric CO_2 at scale is nature-based solutions.

Zachary Bogue
Managing partner, DCVC

Humanity has already reinvented agriculture several times over to keep up with population growth. Now we must do so again, to grow food that's healthier, cheaper, and less destructive to the natural environment. DCVC and DCVC Bio have invested in a range of companies creating more efficient methods for growing crops and raising livestock.

We mentioned above that our portfolio company CH4 Global makes a seaweed-based additive for cattle feed that alters their digestion, drastically reducing the amount of methane they belch up. That's just one example of the kind of innovative agtech company we look for. Sometimes these firms are using the tools of biotech; sometimes they're advancing the state of the art in hardware or robotics.

One of the biggest environmental stressors from modern agriculture is the production of synthetic nitrogen fertilizer. The traditional Haber-Bosch process for pulling nitrogen from the atmosphere and storing it as ammonia dates back to 1913 and has been an enormous boon, allowing farmers to turn unproductive land into fertile fields that have helped to feed billions of people. Unfortunately, the process requires immense supplies of electricity, as well as fossil feedstocks such as natural gas; overall, fertilizer production is responsible for about 2 percent of global CO_2 emissions. On top of that, excessive use of synthetic fertilizers pollutes waterways with nitrogen runoff and releases vast amounts of the greenhouse gas nitrous oxide (N_2O)—an even bigger contributor to climate change, in CO_2-equivalent terms, than methane from cattle.

Pivot Bio, which we've backed since 2014, offers an alternative. The company identified strains of bacteria that live in the rhizomes of crops such as soybeans and turned off genetic brakes on nitrogen fixation, enabling them to make natural ammonia fertilizer from the air. Peer-reviewed studies published in 2024 and 2025 showed that the microbes in Pivot's PROVEN 40 additive deliver an extra 35 to 40 pounds of nitrogen per acre. Farmers who inoculate their soil with PROVEN 40 can buy that much less synthetic fertilizer from traditional suppliers, while also lowering N_2O emissions from their fields. In fact, a nitrogen credit system created by Pivot—which allows companies that use crops like corn to pay farmers for verifiable reductions in synthetic fertilizer use—was named by Time as one of the best inventions of 2024.

↓ **Pivot Bio's** microbes convert atmospheric nitrogen into a form available to corn plants.

Another way farmers can use less synthetic fertilizer and pesticides is to apply them more precisely. We backed Blue River Technology, which created a sprayer attachment for tractors that uses computer vision and machine learning to tell crops from weeds and apply herbicides only when needed. John Deere bought the company in 2017 and sells the system under its See & Spray brand. In 2019 DCVC Bio followed on with an investment in Verdant Robotics, a kind of big brother to Blue River, and we have also backed Sabanto, whose automation packages turn the tractor already in the farmer's barn into an autonomous vehicle. With robot tractors, farmers can mow, till, aerate, or seed their fields with higher precision and a fraction of the labor. Our most recent investment in this area is in AgZen, a DCVC Bio company that uses AI and computer vision to monitor whether crop protection and fertilizers are reaching their intended targets; the system helps farmers improve yields, safety, and compliance, while using 30 to 50 percent less chemistry.

2%

Share of global greenhouse gas emissions from synthetic nitrogen fertilizer production

4%

Share of global greenhouse emissions from agricultural N_2O

Pasture-based livestock farming or "regenerative grazing" is yet another burgeoning form of sustainable, planet-friendly agriculture. It's a humane alternative to confined feeding operations that allows grassfed animals to graze freely on pasture, where they naturally aerate and fertilize the soil and promote carbon sequestration. A DCVC company called Halter makes wireless, solar-powered smart collars that use sounds, vibrations, and low-energy pulses to train grassfed dairy and beef cattle to stay inside a movable virtual fence line. When it's time to rotate to a new pasture, farmers simply redraw the lines. The system also allows farmers to shift cows off poor terrain or eroded areas, monitor individual animals for health changes, and allocate the right amount of grass per cow, maximizing milk production. Just as with CH4, Pivot, Blue River, Verdant, and Sabanto, Halter's technology helps farmers work more intelligently and profitably while making their operations more environmentally friendly and sustainable.

Reindustrialization halts without water.

Water supplies limit growth of all kinds—agricultural, residential, and especially industrial, given the large water requirements of facilities such as data centers and semiconductor fabrication plants. One key to industrial revitalization in the U.S., therefore, will be investing in technologies that provide reliable sources of clean, fresh water.

OPPORTUNITIES

COMPANIES

Aquafortus,
Tidal Metals, ZwitterCo

VOICES

Zachary Bogue,
Earl Jones

In many regions, human activity is draining surface waters and aquifers faster than nature can recharge them. Weather disruptions exacerbate the problem: changing rainfall patterns due to global warming mean that one in three people around the world live in water-deprived areas, up from one in five in 1990, according to the United Nations. To support population growth, reindustrialization, and the clean energy transition, we must be creative about where we find water.

Water has never been free. The aqueducts of the ancient Persian, Roman, Incan, and Aztec empires were enormously expensive public-works projects designed to keep cities and farms growing. But Rome wasn't full of hyperscale data centers, which use 400 million gallons of water per day in the U.S. alone. Today, with demand for fresh water reaching new extremes, it's more important than ever to find new sources of H_2O and to lower the cost of obtaining and treating it. As we've said in past editions of the Deep Tech Opportunities Report, these are natural tasks for deep-tech innovators.

One type of water that's currently viewed by energy producers as a headache—but could be turned into a rich resource, using new technology—is the briny "produced water" that comes out of the ground alongside oil and gas, at a rate of 4 to 6 barrels of produced water for every barrel of oil. In the fossil-rich Permian Basin region of Texas and New Mexico, well operators generate about 840 million gallons of produced water each day. The current practice is to inject most of it back underground, but that loosens geologic formations and causes earthquakes, including a magnitude 5.0 quake that struck near Toyah, Texas, on February 14, 2025. On top of that, well operators are "literally running out of underground pore space to dispose of this water," says DCVC operating partner Earl Jones. "If you can't get rid of the water, there's no oil."

Aquafortus, which we've backed since 2021, is scaling up a solvent extraction technology that recovers clean water—as well as valuable minerals—from produced brine. The company cycles brine through a series of towers full of absorbent and regenerant materials that trap water and isolate salts (see infographic, page 65). The clean water can be reused for industrial purposes, while the salts can be processed to extract industry-critical minerals such as lithium, magnesium, cobalt, strontium, bromine, and iodine. Aquafortus' field research facility in Colorado has already produced more than 20,000 barrels of clean water from briny wastewater, and the company has agreements to scale up the process for a number of oil and gas companies, including Occidental Petroleum. "What we're trying to do is reframe this from a discussion of how you dispose of a waste product to how do we extract value from a resource and create a vision of water abundance," Jones says.

That could be especially important for regions that don't have enough conventional water from rivers and aquifers but do have a lot of brine or brackish water. On the New Mexico side of the Permian Basin alone, produced-water extraction amounts to 160 million gallons per day, which is roughly equal to the state's total daily municipal water consumption. A means of treating and reusing this water, perhaps for industrial or agricultural purposes, would be a big win for the state.

Figure 4.1.1

AI's water footprint—about 13 billion liters per year from Google's U.S. data centers alone—includes water evaporating from the cooling towers of steam-generating power plants and chiller systems for server racks.

Source: UC Riverside.

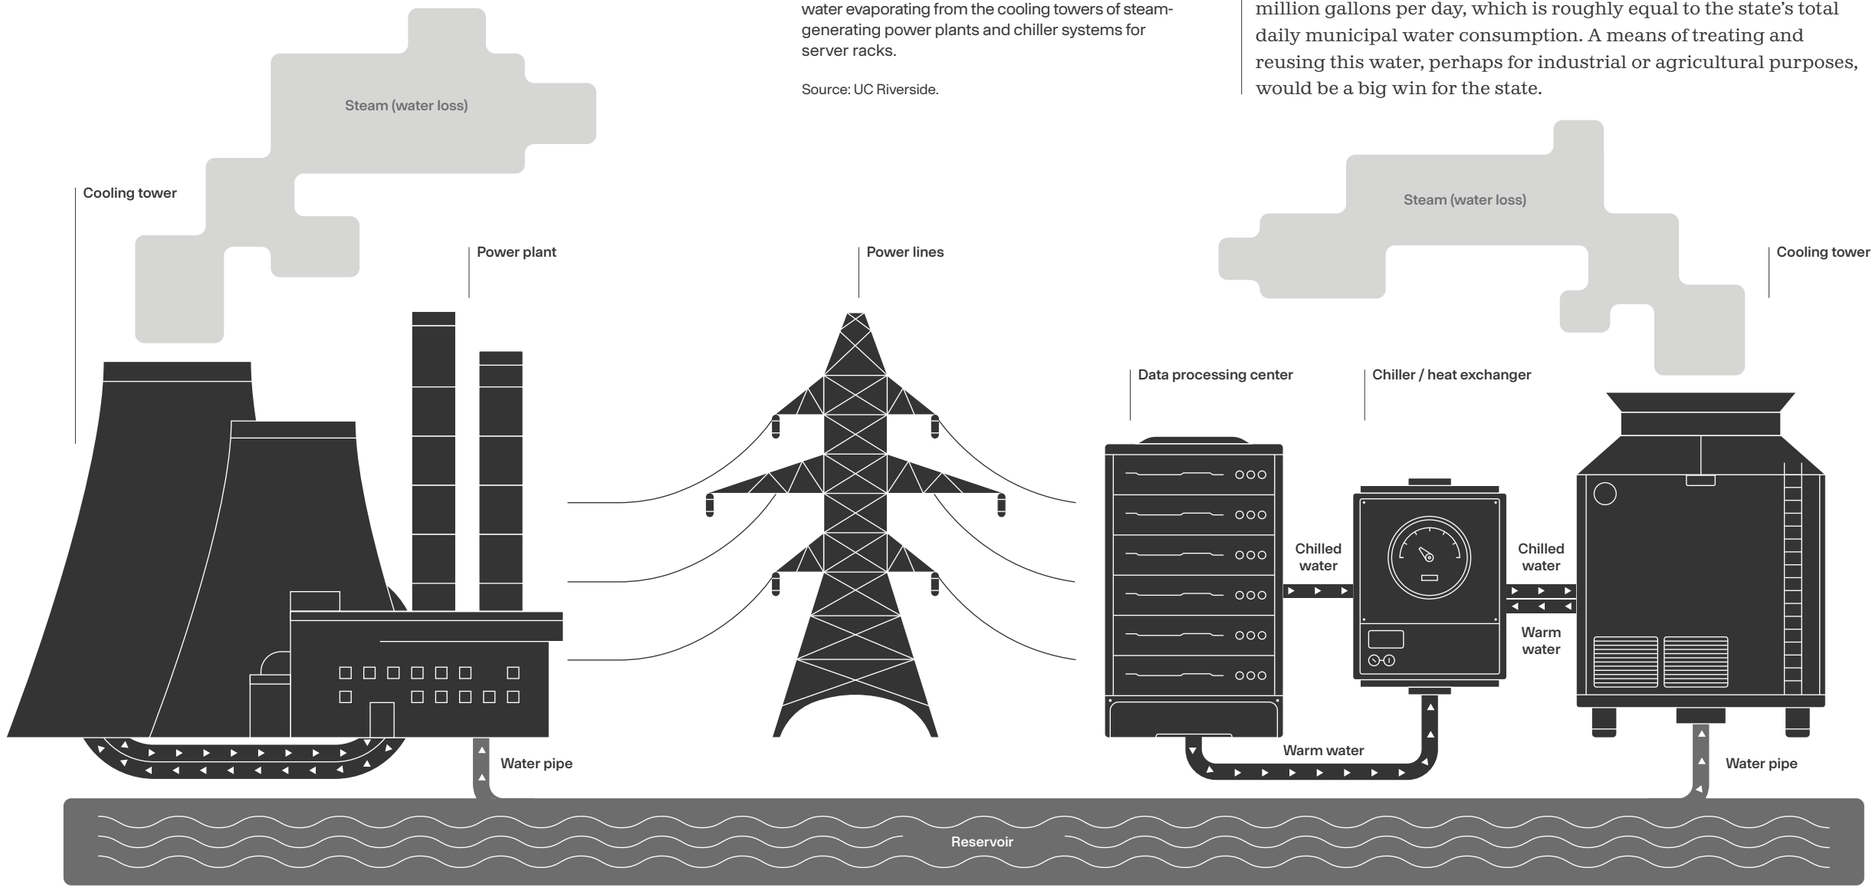

Another DCVC portfolio company that's also in the business of helping customers reclaim, rather than discharge, dirty water is ZwitterCo. It's a developer of superfiltration and reverse-osmosis membranes that stick to water rather than the compounds they filter out. They do that by incorporating zwitterions—molecules that have equal numbers of positively and negatively charged branches and that can therefore attract H_2O and repel the proteins, fats, oils, and other organic compounds that normally foul or clog water filters and reduce their throughput. ZwitterCo membranes still clog up, but far more slowly—and they can be cleaned much more easily, without caustic chemicals, which means they don't need to be taken out of service or replaced nearly as often as traditional filters.

"A lot of these industrial waters are so biologically contaminated that they've been challenging to treat with membranes," Jones explains. "Digestates, leachates, textiles, fermentation broths, refinery wastes, ag waste, rendering plants—all are big sources of wastewaters that are really expensive to treat." The company already has more than 50 customers across 15 countries, including farms, which can now filter the water out of manure and reuse it for crop irrigation. Power plants are another potential customer: they already use reverse-osmosis membranes to purify dirty or brackish surface waters for use in steam turbines or cooling towers, and swapping out their old elements for ZwitterCo's would reduce fouling, maintenance, and operating expenses. An added bonus is that simply by cleaning ZwitterCo's filters, customers can recover and reuse the trapped nutrients and other materials.

> What we're trying to do is reframe this from a discussion of how you dispose of a waste product to how do we extract value from a resource and create a vision of water abundance.
>
> **Earl Jones**
> Operating partner, DCVC

Recovery is also the goal at Tidal Metals, a water-tech company we described in Chapter 1. Tidal is building pumps that extract magnesium from seawater. The upshot could be a cheap, clean, domestic supply of this extremely useful metal. "Magnesium is the strongest, lightest, structural metal, and if we can make it domestically, we can lightweight all of our vehicles," says DCVC's Bogue. American manufacturers will also need magnesium for the next generation of aircraft, and for magnesium-based batteries, which could achieve higher energy densities than lithium-ion batteries.

Tidal, Aquafortus, and ZwitterCo each show in their own ways how advanced water treatment is a win-win proposition. The technology doesn't simply provide the clean water we need to run factories and power plants and irrigate farms, but it turns waste products into assets and creates new sources of critical minerals and materials. In the end, we believe, American reindustrialization and the transition to zero-carbon energy won't happen without cheap, energy-efficient water reuse as part of the package.

Figure 4.1.2

The Aquafortus process removes water from brine through solvent extraction, then produces clean water through membrane filtration.

01

In the Absorption Tower, a dry absorbent material is added to the bottom, and briny wastewater is added to the top. Water transfers into the absorbent, leaving all but 2 percent of the salt behind in a concentrated brine.

02

The wet absorbent is piped to the bottom of the Regeneration Tower, where dry regenerant material floats at the top. Low-salinity water transfers into the regenerant, and the dry absorbent is recycled.

03

Wet regenerant is pumped through a system of tube-shaped nanofiltration and reverse osmosis membranes at high pressure. These membranes, like ZwitterCo's, attract water and reject regenerant (which spares the membranes from fouling or scaling from the residual salt in the regenerant). The filtration tubes produce clean water, and the dry regenerant is recycled.

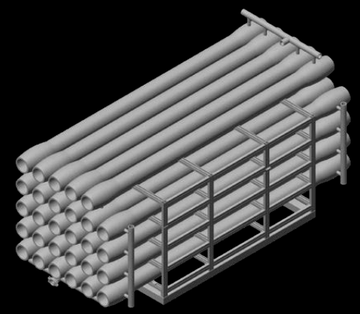

Water offsets, patterned after familiar carbon offsets, could be a big win-win for industry, allowing operators of water-intensive facilities such as data centers to expand without increasing net water consumption.

Fresh water is a finite yet indispensable resource that's increasingly costly to obtain and/or remediate. That's the fundamental realization driving the deep-tech water innovation we talked about in the previous section; it's also the fundamental notion behind the cluster of ideas that goes by labels such as water resilience, water neutrality, water replenishment, water balancing, water stewardship, or water offsets.

In a water-neutral development, builders offset 100% of their projected water demand by covering the cost of water-efficiency projects elsewhere. It's an approach that's been tested or studied by a handful of cities in Brazil, Mexico, the U.K., and California. And it's also taking hold in the corporate world. In 2007, Coca-Cola committed to replenish 100 percent of the water used in its finished products globally, and it says it's been meeting that goal every year since 2015. In 2021, Google joined the water stewardship movement, announcing that by 2030 it aims to replenish 120 percent of the water it consumes at its offices and data centers.

In practice, "replenishment" often translates to support for water conservation projects that will balance or offset corporate activities. Google's recent collaboration with agtech startup N-Drip, which aims to replace wasteful flood irrigation with a gravity-powered micro irrigation system, is an example. Google is paying farmers in the Platte River basin of Nebraska to convert 1,000 acres of corn and soybean fields from flood irrigation to N-Drip's system, offsetting some of the water use of its data center in Papillion, Neb.

At the same time, Google is paying farmers in southeast Nebraska to adopt water conservation technology from a San Francisco startup called Arable. The company makes sensors that monitor weather, crop growth, soil moisture, and other parameters, feeding the data into software that advises farmers on when and how much to irrigate. In demonstrations, the company says, the Arable system has helped corn and soybean farmers reduce irrigation by 22 percent.

Microsoft, too, has pledged to replenish more water than it consumes by 2030, and said in 2023 that it had invested more than $16 million in 49 projects around the world, saving more than 16 billion gallons of water.

"When we think about water being a constraint on industrial revitalization, companies like Microsoft and Google stepping up for water offsets is a great thing," says DCVC's Jones. "Just as with carbon offsets, there has to be validation. We're still in the early stages. But if it matures into something stable, we could see a lot of industrial companies funding water use reduction. These farm projects can be a way to prime the pump."

→ **No-pump, no-flood irrigation** Corn farmers in Nebraska are using a drip irrigation system funded by Google's water replenishment program.

TechBio companies are using data to rethink how we develop new medical treatments.

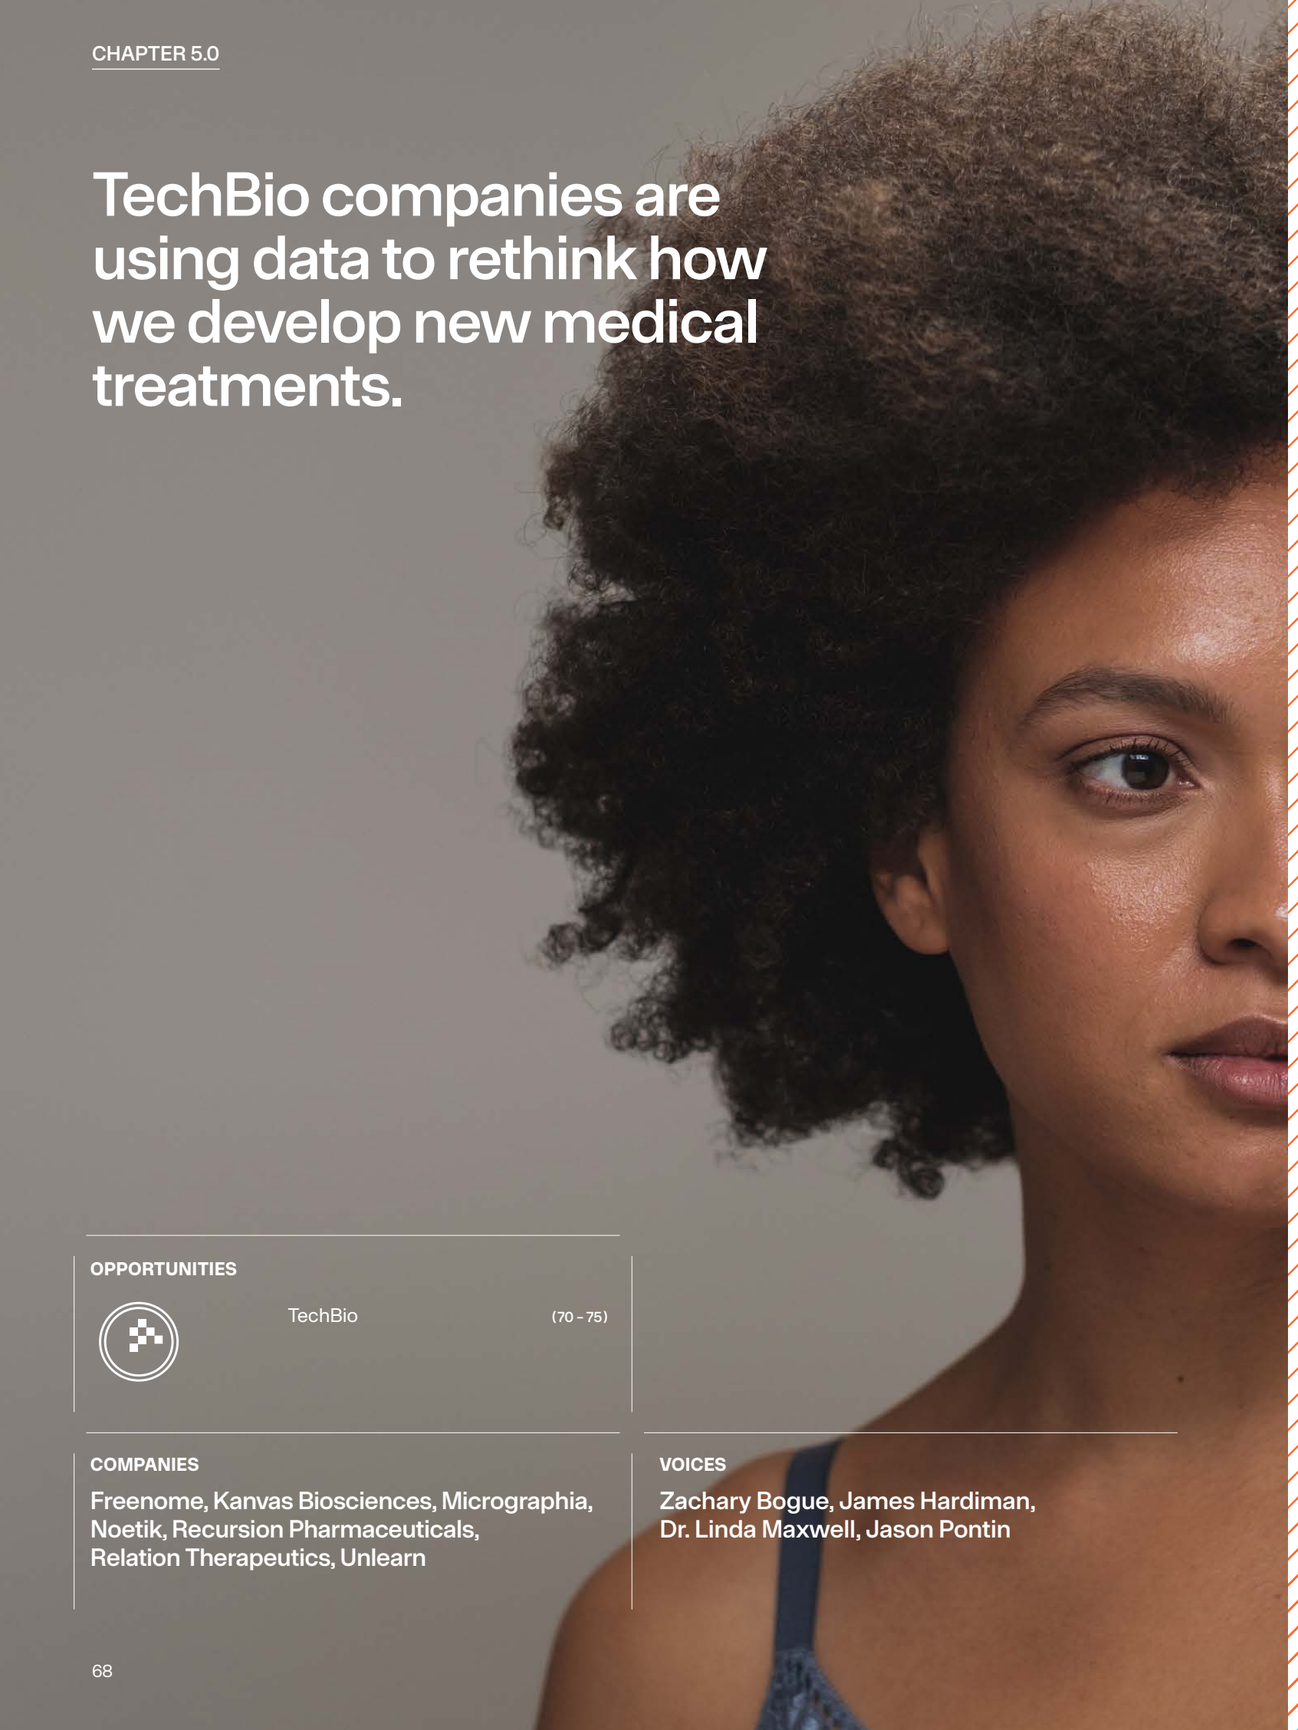

OPPORTUNITIES

TechBio (70 – 75)

COMPANIES

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

VOICES

Zachary Bogue, James Hardiman, Dr. Linda Maxwell, Jason Pontin

Unlearn creates digital twins of clinical subjects to help companies design smaller, faster drug trials.

We think there's room in the life science industry for a new kind of company that relies on a mix of deep lab expertise and data-driven computing—including machine vision, digital twinning and simulation, and predictive analytics—to get new treatments to market much faster and more efficiently than older approaches. We call that combination TechBio, and we're seeing the idea take shape in promising ways across multiple portfolio companies.

"TechBio companies, first and foremost, set out to create highly structured, proprietary data sets, then run machine learning or other AI models on that data to get to either therapeutic or diagnostic targets," explains DCVC's Bogue. "Last year, TechBio was still emerging, but now it's hitting its stride, with major proof points."

We've mentioned many of our TechBio portfolio companies in previous reports, but all of them have progressed substantially. Relation Therapeutics, for example, is proving that a hybrid system combining single-cell genomics, AI modeling of gene and protein networks, and large-scale lab analytics on effector cells can help identify promising new drug targets for undertreated diseases in less time than conventional drug discovery and development.

An effector cell is a human cell type chosen by Relation for its involvement in a complex disease—say, osteoblasts, a kind of bone cell, in the case of the company's osteoporosis research—which is then modified using CRISPR and other tools to remove genes the company's AI model predicts may be contributing to pathology. In the lab, the company studies how knocking out these genes one or two at a time in different cells affects protein expression and other markers (such as mineralization rates, in osteoblasts). Then it deploys more machine-learning algorithms to identify and prioritize the potential disease pathways that could best be modulated with drug molecules.

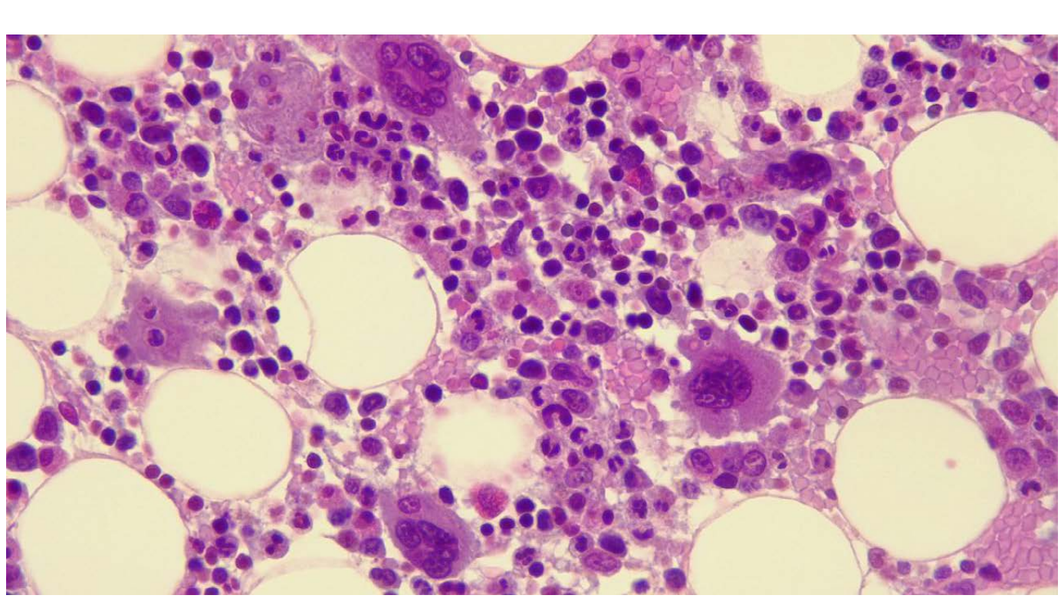

↖ **Relation Therapeutics** builds proprietary datasets for AI model training by studying gene-edited human cells in the lab.

"Last year, TechBio was still emerging, but now it's hitting its stride, with major proof points."

Zachary Bogue
Managing partner, DCVC

"Relation has shown that they can massively increase the number of targets, and reduce the failure rates and time involved in drug development.

Jason Pontin
General partner, DCVC

The work is helping Relation zoom in on possible targets for new osteoporosis drugs—and it's also attracting the interest of pharmaceutical giants like GSK. In a multifaceted deal announced in December 2024, GSK committed to $108 million in upfront and collaboration-based payments to develop drugs for fibrotic disease and osteoarthritis. Relation is also in line for potential preclinical, development, commercial, and sales milestone payments averaging $200 million per target, along with tiered royalties on net sales of products. We believe it's the largest deal ever for a seed-stage biotech or TechBio company.

"If we want life-changing medicines for people who are at risk for conditions like osteoporosis, we need to understand complex diseases. What's unique about Relation is they have had a 'lab in the loop' from the beginning, using effector cells to test their hypotheses," says DCVC general partner Jason Pontin. "With these GSK deals, and with the identification of 20 novel targets in osteoporosis in only two years, Relation has shown that they can massively increase the number of targets, and reduce the failure rates and time involved in drug development."

Other DCVC- and DCVC Bio–backed TechBio companies are advancing just as rapidly, as we explain on the following pages.

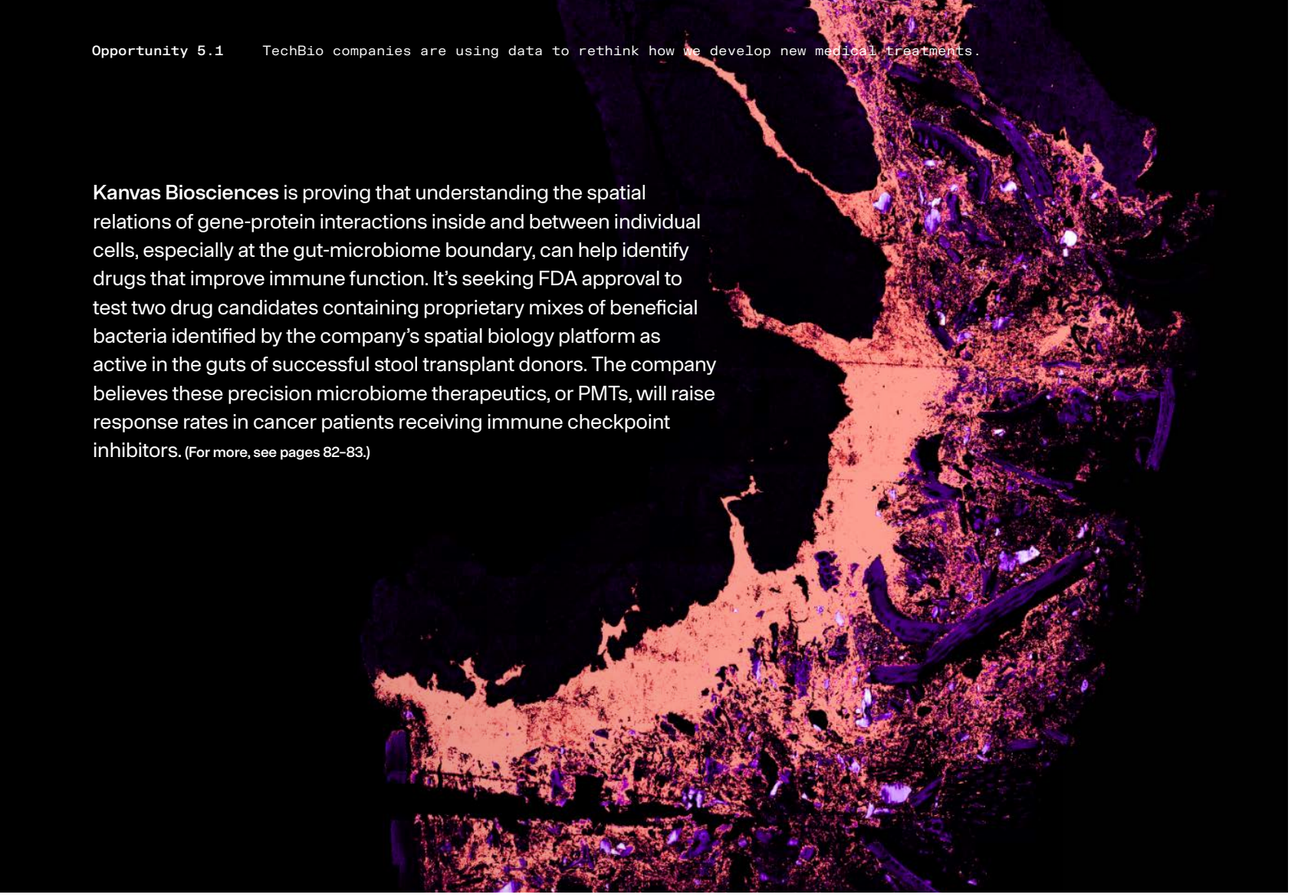

Kanvas Biosciences is proving that understanding the spatial relations of gene-protein interactions inside and between individual cells, especially at the gut-microbiome boundary, can help identify drugs that improve immune function. It's seeking FDA approval to test two drug candidates containing proprietary mixes of beneficial bacteria identified by the company's spatial biology platform as active in the guts of successful stool transplant donors. The company believes these precision microbiome therapeutics, or PMTs, will raise response rates in cancer patients receiving immune checkpoint inhibitors. (For more, see pages 82–83.)

Unlearn is showing how AI can speed drugs to market by helping companies rethink the design of clinical trials. The company's AI models create and analyze digital twins of potential trial participants to predict their health outcomes under placebo, which allows researchers to design smaller or shorter trials. In a collaboration with Johnson & Johnson reported in 2024, Unlearn showed that trials of Alzheimer's drugs using digital twins could shrink their control arms by 33 percent without losing any statistical power—a finding that could save pharmaceutical companies tens of millions of dollars in large drug trials.

Noetik is demonstrating that new foundation models trained on huge amounts of imaging data can speed drug development. The company used images showing spatial patterns of gene expression in almost 40 million cells from over 1,000 patient tumor samples to train a transformer model called OCTO-vc, which can now simulate how virtual cells would behave in different tissue environments. "For example, a virtual T cell can be introduced into various parts of a tumor to infer whether it will be activated and capable of killing cancer cells," the company explains. That's data that could help the company design new, precision immunotherapies for cancer. "The team's speed of execution in building one of the most sophisticated AI-enabled oncology discovery engines in two years is unprecedented," says DCVC's Hardiman.

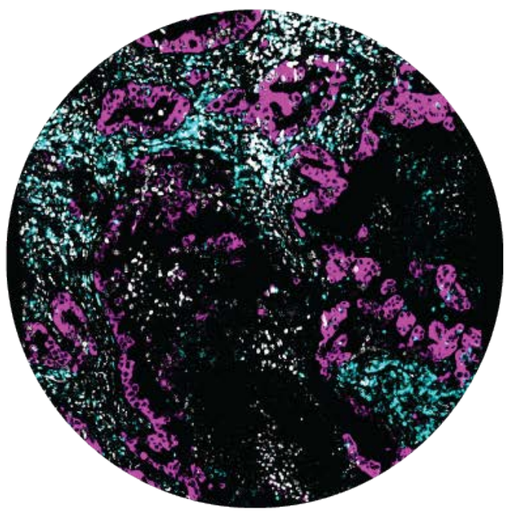

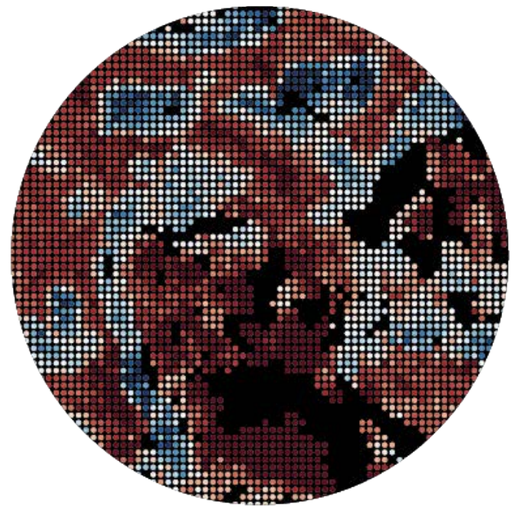

Freenome, too, uses machine learning to study human samples, but with a focus on the bloodstream and the cell-free DNA and proteins that circulate there. Patterns in these biomarkers can indicate the presence of cancer before it's detectable by other means. In a study of its blood-based test for colorectal cancer on nearly 50,000 patients reported last year, Freenome found that the test correctly detected 57 percent of Stage I cancers, 100 percent of Stage II, 82 percent of Stage III, and 100 percent of Stage IV, for a 79 percent sensitivity rate overall. That's in the same ballpark as detection rates for traditional screening methods such as colonoscopies—but Freenome's test is non-invasive, meaning it could make early detection of colorectal cancer far more accessible.

Recursion Pharmaceuticals, a DCVC-backed company that went public in 2021, has advanced multiple drugs through early-stage clinical trials. The company trains AI models on lab data showing how genetic changes, toxins, and other insults affect cell morphology. It then runs those models on its BioHive-2 supercomputer, made up of a few thousand Nvidia GPUs, to predict how different drugs will affect the human body. In 2024 the company dosed its first patients in trials of drugs against solid tumors, lymphoma, and *Clostridioides difficile* infection. And in January Recursion said it had received the go-ahead to test a drug against B-cell malignancies in the U.K. and a drug against small-cell lung cancer in the U.S.

Micrographia Bio, which we've backed since 2022, is similar to Kanvas Bio and Noetik in that they're combining microscopy and machine learning to measure the spatial distribution of protein-protein interactions in cells, allowing them to infer the state of the cells from imagery alone. Their nanometer-scale spatial proteomics platform, boosted by cutting-edge optics and AI-enabled image analysis, allows them to measure how perturbations such as drug molecules affect cells over time.

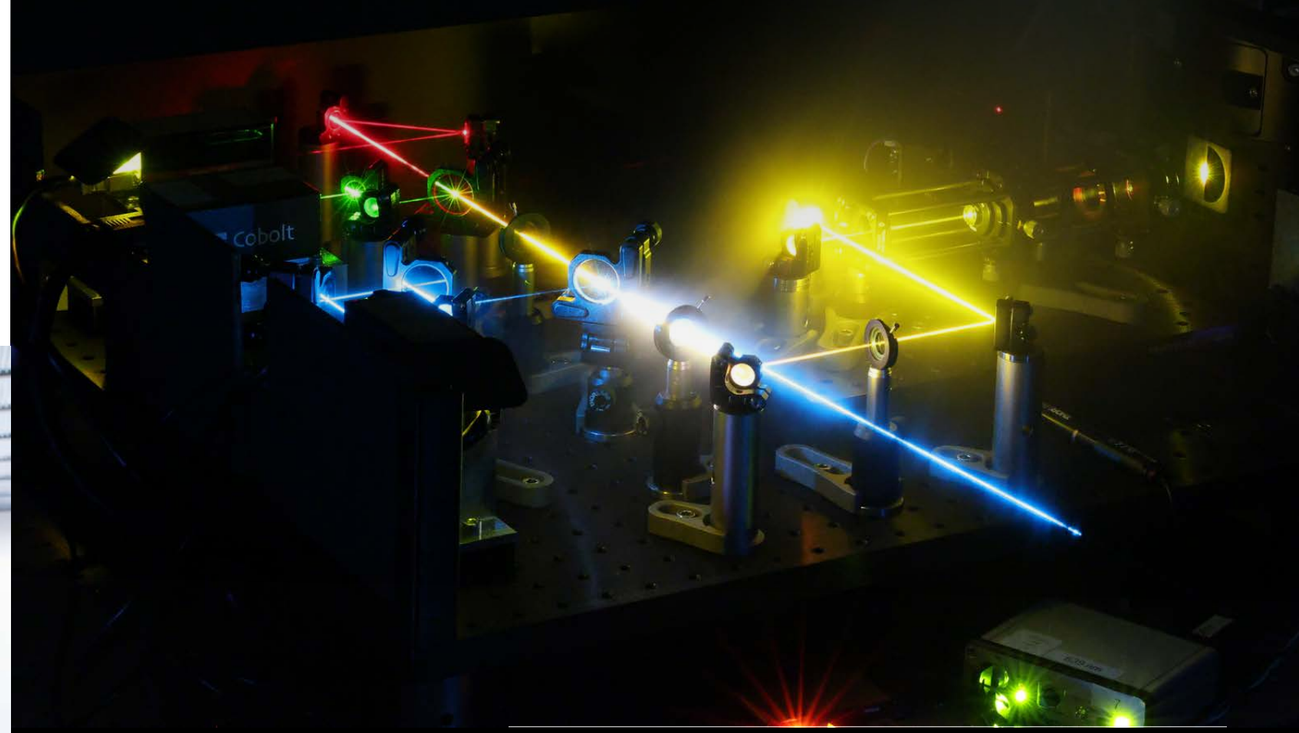

DCVC's Ocko defines TechBio this way: "Deep machine learning and AI, leading to an understanding of foundational biological systems, with a lab in the loop or similar acquisition of real-world data to formulate and vet those systems, and actionable outcomes from the use of those systems such as a molecule or a recommendation." By that definition, all of DCVC's TechBio portfolio companies are advancing the state of the art, and we expect to see them creating drugs and tests that benefit millions.

TechMed—the merger of human medical expertise and computing—is remaking healthcare.

Using computing and data, we can reinvent and drastically improve existing medical diagnostics, devices, and healthcare delivery systems. We've been investing in this sector for a while, and its promise is starting to pay off. Thanks to AI, the cost of developing and testing new machines and procedures is going down drastically, and we think we'll soon see unicorn and decacorn companies emerging.

Kanvas Biosciences uses its new spectral light-sheet microscope to capture images of GI tissue at unmatched speed and resolution.

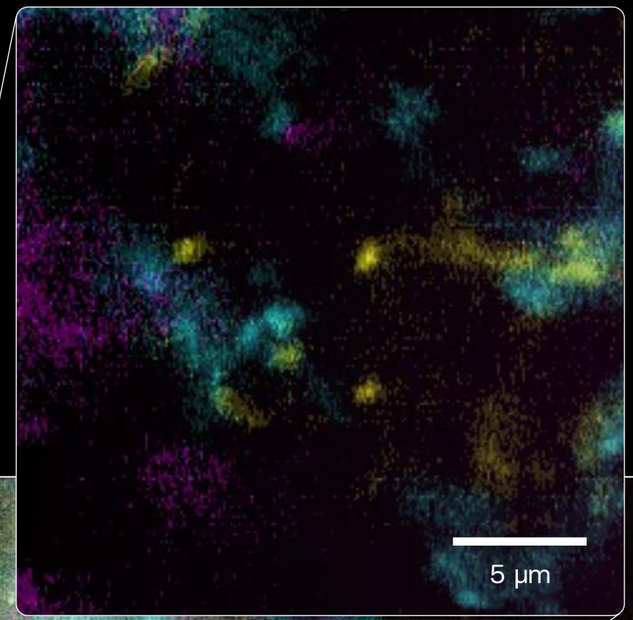

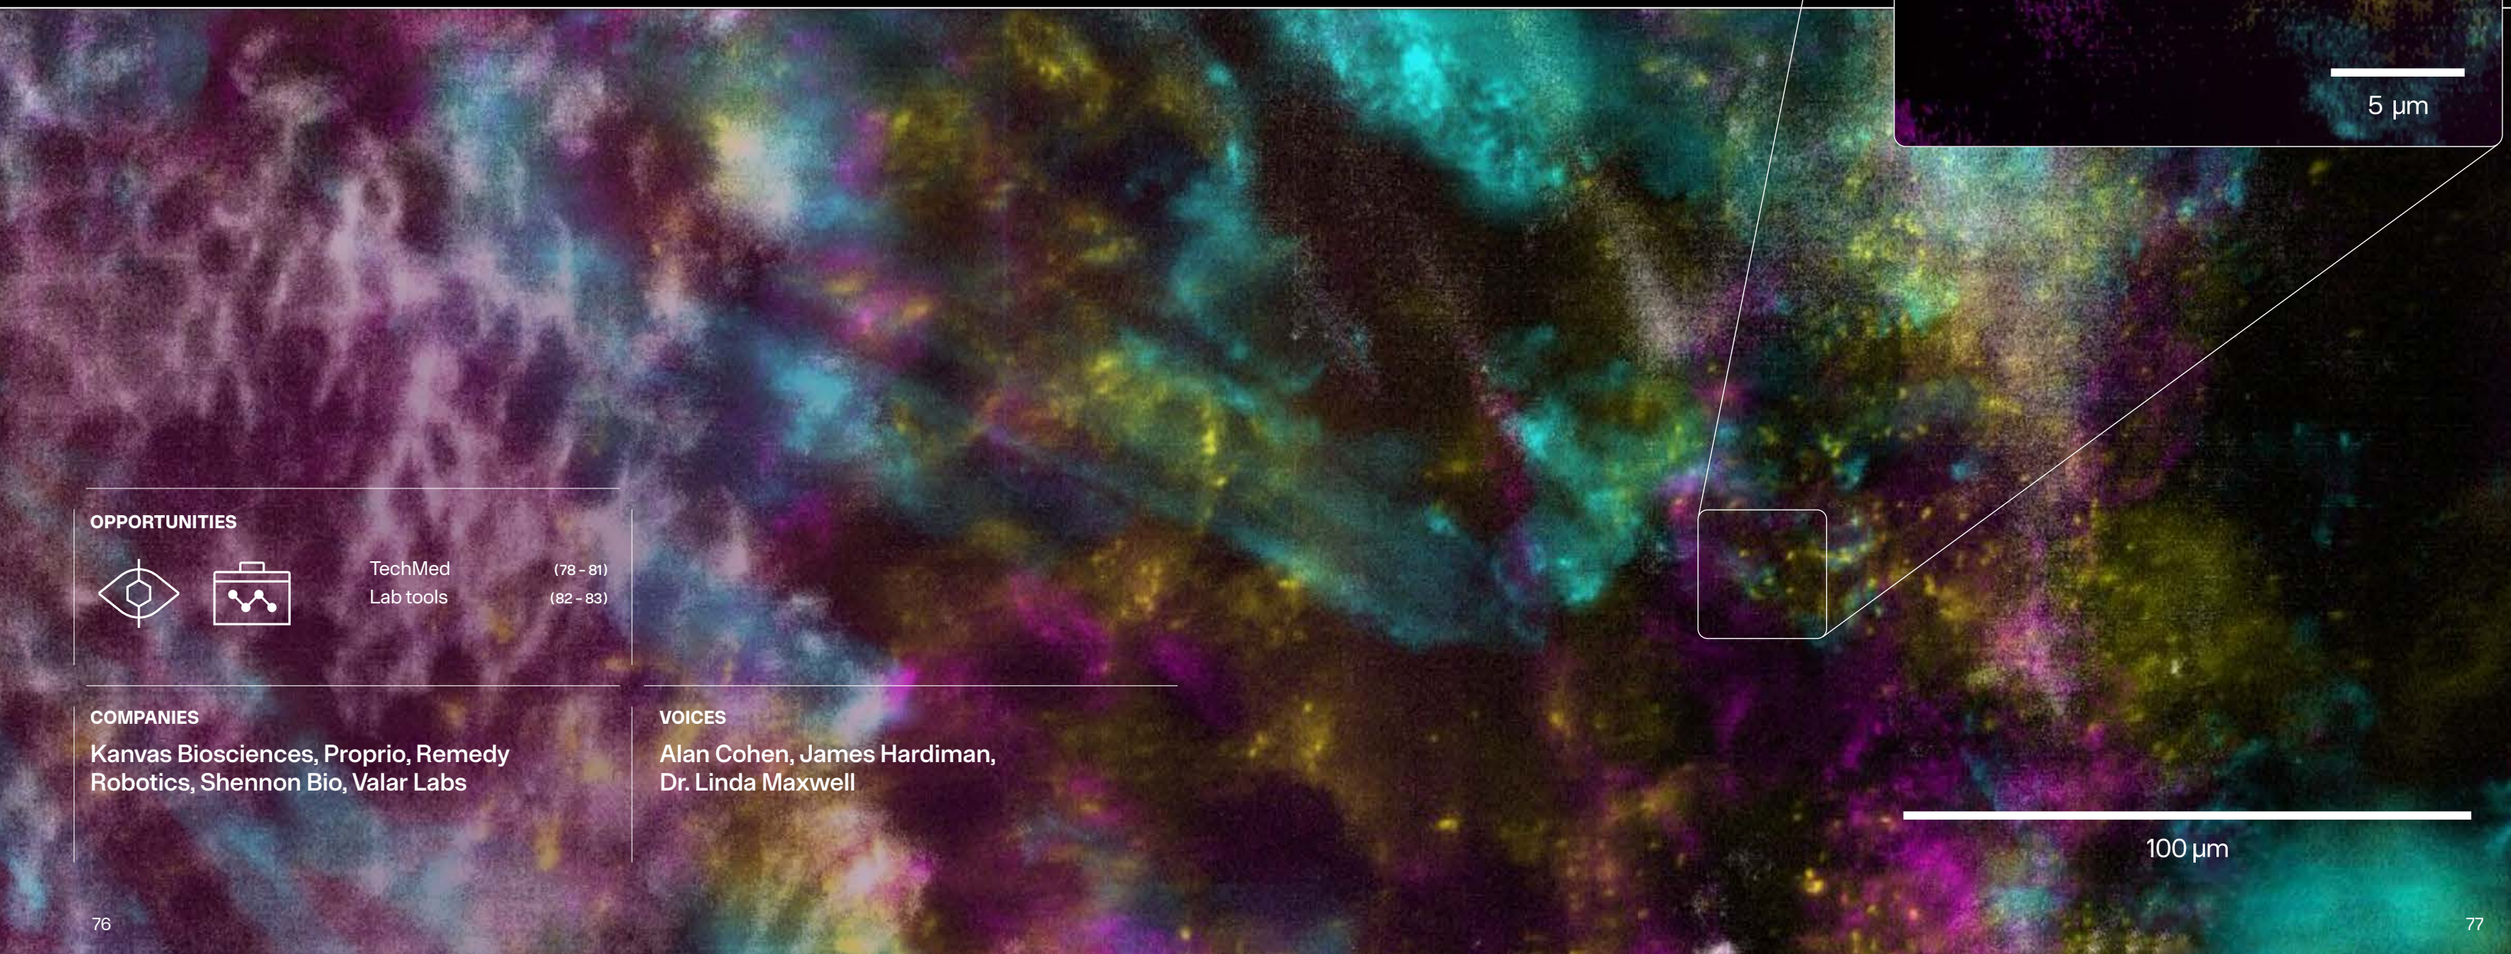

OPPORTUNITIES

COMPANIES

Kanvas Biosciences, Proprio, Remedy Robotics, Shennon Bio, Valar Labs

VOICES

Alan Cohen, James Hardiman, Dr. Linda Maxwell

A robot such as the da Vinci Surgical System can allow a surgeon to perform a minimally invasive procedure with greater precision. A large language model such as Google's Med-PaLM 2 or MedLM can help a physician explore possible diagnoses they might have missed or create more thorough medical notes for a patient's health record. The human providers in these scenarios are performing at a higher level than they could have alone—and we think those are just small portents of the level of co-piloting between human and machine we will soon see in healthcare.

After his defeat by IBM's Deep Blue computer in 1997, chess grandmaster Garry Kasparov proposed a new form of the game called "cyborg chess" or "centaur chess" with two human opponents, each consulting a computer chess engine, each playing at a superhuman level thanks to the human-machine collaboration. The concept and the label quickly spread to other fields such as financial advising, legal research, intelligence analysis, and software development. And it pops up frequently in science fiction and futurism; Adrian Hon, in his book *A New History of the Future in 100 Objects*, saw a near future with "amplified teams" of three to seven people supported by communications and AI software, with "the ability to outperform any individual human or any lesser-integrated team at any intellectual task."

Radiology is one field already being transformed by AI in this way. In a recent trial involving 80,000 women in Sweden, teams consisting of an AI system and a radiologist identified 20 percent more breast cancers than teams pairing two human radiologists, at half the workload per physician. We're convinced that this centaur model will continue to spread through healthcare. And we can already see it playing out in several companies in our TechMed portfolio.

Remedy Robotics, to take one example, is building a remotely operated endovascular surgical robot that could provide immediate care to patients suffering cardiovascular emergencies, even at hospitals where there's no qualified surgeon. Machine-learning software built into Remedy's tools enables enhanced visualization and precise control, allowing for remote treatment and bringing down complication rates.

Then there's Valar Labs, where researchers have built machine-learning models that can serve as what CEO Anirudh Joshi has called "a co-pilot for the oncologist." The company takes tissue samples from solid tumors and analyzes them using deep-learning models trained on hundreds of thousands of tissue images labeled by trained pathologists. The models predict how fast the patient's cancer will progress and whether it will respond to various forms of therapy.

The company's first product, Vesta, is designed for early-stage bladder cancer patients; it predicts whether a given patient needs aggressive treatment such as a cystectomy or whether a form of immunotherapy called Bacillus Calmette-Guérin (BCG) treatment might slow tumor growth. Joshi and the company's other founders believe that a similar approach will help doctors treat other types of tumors, such as pancreatic cancer. "The thing with foundation models in biology is that there are going to be many different problems you can point them at," says DCVC general partner James Hardiman.

Valar is, in effect, "giving pathologists superpowers," adds DCVC general partner Alan Cohen. "TechMed companies are focused on creating intelligent software systems that can augment and optimize the human technique involved in the practice of medicine itself."

Figure 6.1.1

Valar's AI algorithm differentiates between different types of tumor cells, and even predicts how patients will respond to various treatments, by learning to recognize features in thousands of pathology images labeled by experts.

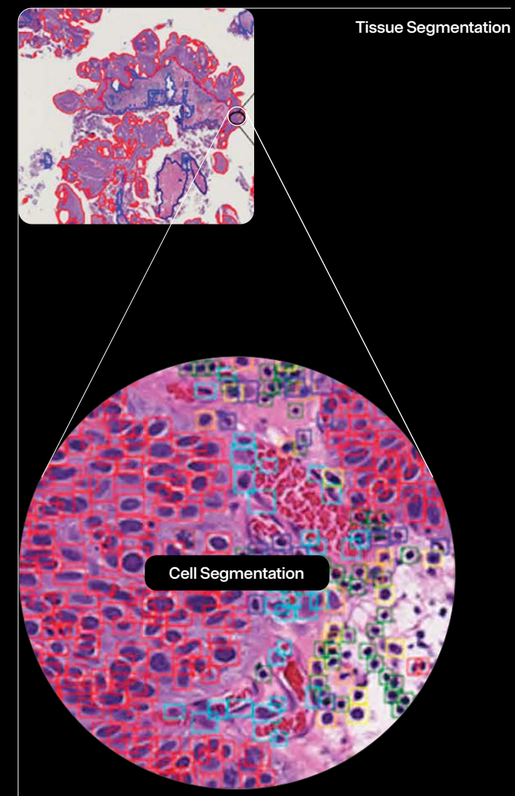

> The thing with foundation models in biology is that there are going to be many different problems you can point them at.

James Hardiman
General partner, DCVC

Proprio is doing exactly this for specialists in spine and cranial surgery. The rate of complications after lumbar spine surgeries or spine fusion procedures is shockingly high, with over 30 percent of patients needing revision surgery within five years, according to one study. When surgeons can track instruments and align implants precisely, minimizing deviations from their preoperative plan, complication rates drop dramatically. That's where Proprio comes in. The company places depth-sensitive lightfield cameras over the surgical area, feeding data into deep-learning models that construct magnified 3D representations of patient anatomy. This helps surgeons "see" through blood and tissue and align tools precisely.

The FDA approved the system in 2023 and it's been used in over 80 surgeries at multiple medical centers, with zero revisions required so far. "The equivalent might be JARVIS in the Iron Man suit, or the heads-up display in an F-35," says Cohen. "Pilots used to fly based on sight and maybe radar, but increasingly they have all these other systems that are informing their decision path. I think we'll make other investments like that, where you are effectively mastering an enormous array of data and information and building a digital twin that informs a medical professional in what they are doing."

Importantly, innovations like those at Proprio and Valar promise to bring huge cost savings to the healthcare industry. Better intraoperative data means faster surgeries and higher OR utilization. Fewer revision surgeries means lower costs for insurers. AI assistance can reduce the cost of cancer screening. Picking the right cancer treatment first helps avoid expensive and debilitating tests and treatments.

But ultimately, the best reason to make doctors into centaurs is to help patients. "There's a lot of technology that doesn't actually move the needle on outcomes. They are nice-to-haves," says Dr. Linda Maxwell, an operating partner at DCVC. "Proprio is an exemplar of really smart technology that does all the great things like improving the bottom line and efficiency. But where the rubber hits the road—the patient outcomes—they are demonstrating meaningful impact. The patient is going to be out of the hospital sooner, and they're less likely to need revision work."

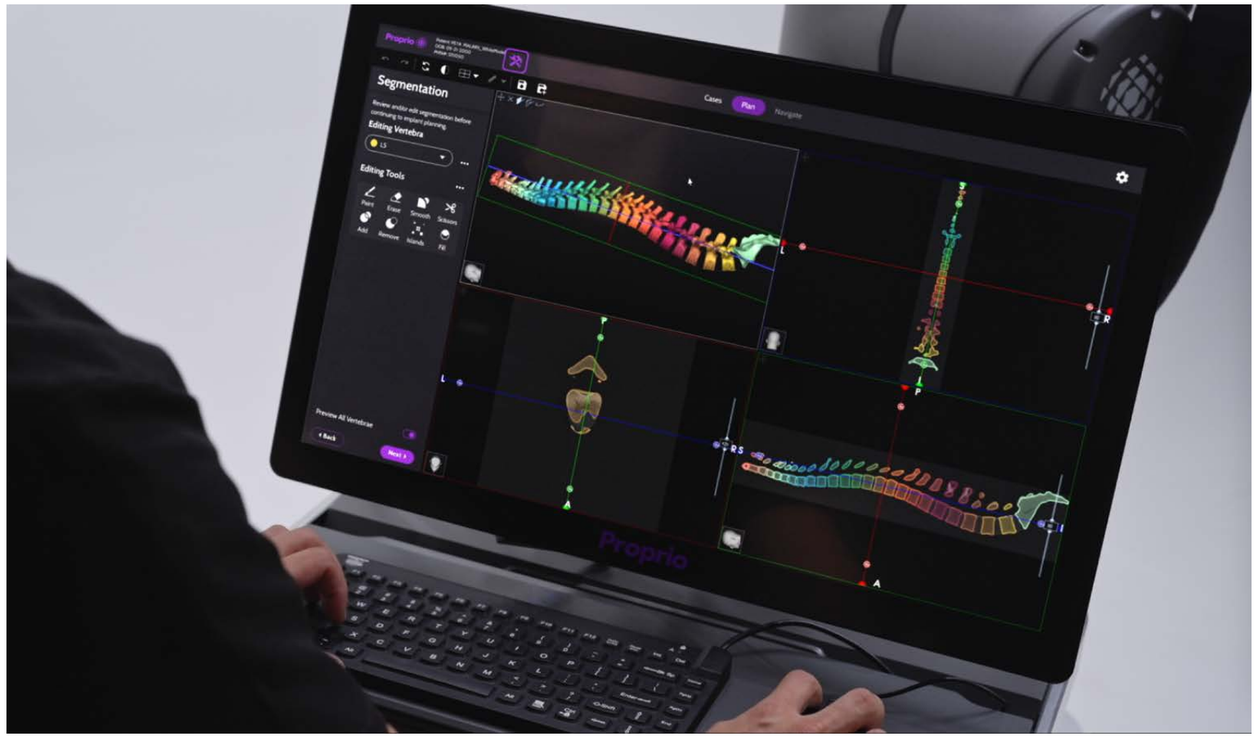

> TechMed companies are focused on creating intelligent software systems that can augment and optimize the human technique involved in the practice of medicine itself.

Alan Cohen
General partner, DCVC

> Where the rubber hits the road—the patient outcomes—[Proprio is] demonstrating meaningful impact. The patient is going to be out of the hospital sooner, and they're less likely to need revision work.

Dr. Linda Maxwell
Operating partner, DCVC

New lab techniques in areas like spectrometry, microfluidics, and microscopy are speeding up life science R&D in fundamental ways and getting treatments to patients faster. These enabling technologies—spanning the characterization, modeling, and manipulation of biology—will lead to new discoveries and clinical advances in the same way that bigger, higher-energy particle accelerators extend our understanding of physics.

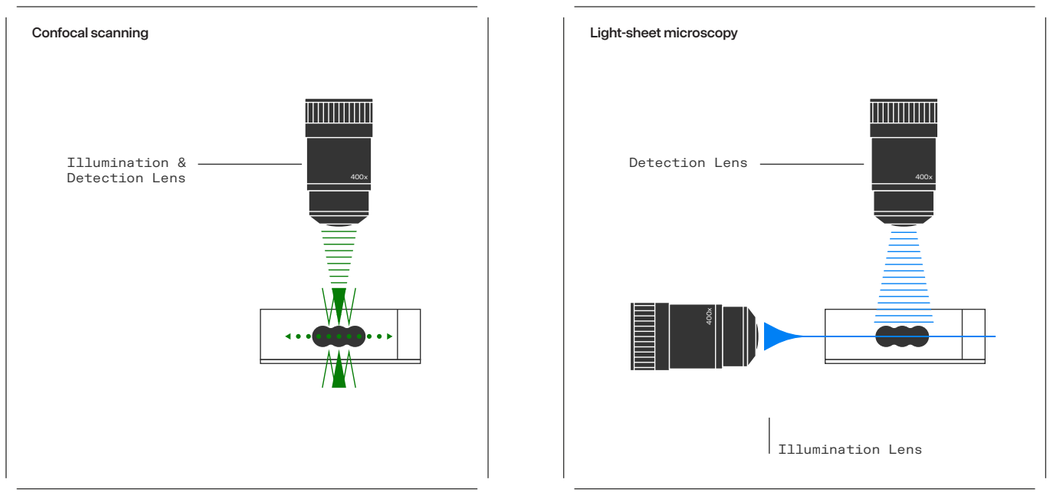

Figure 6.2.1

Kanvas Biosciences is accelerating its effort to understand how gut bacteria and host cells interact by building a custom spectral light-sheet microscope. It speeds up imaging by illuminating an entire sheet of tissue at a time.

Consider cancer immunotherapy, mentioned in Chapter 5 of this report. Its goal is to harness the body's immune system by directing the cytotoxic powers of T cells or B cells against tumor cells. That means, first, identifying the functionally activated T or B cells that recognize antigens on tumor cells, so that larger populations of those cells can be created outside or inside the body. But given the human bloodstream contains hundreds of billions of T or B cells, this is a classic needle-in-a-haystack problem, and isolating the right cells normally requires weeks or months. That's "a heartbreaking problem in a domain where the last thing a patient has is time," says DCVC's Hardiman.

Shennon Biotechnologies, a DCVC portfolio company, has developed a proprietary, ultra-high-throughput platform that can identify the needed cells in just hours. Using microfluidics, the company can isolate single cells, expose them to a pool of antigens, and assess their functional responses—and do it on a massive scale, screening millions of cells per hour. The company tests the selected T cells in the lab and in animals to ensure they'll hit the targeted tumors, then works with pharmaceutical partners to develop them into therapies. "They can fish out these rare T cells in a patient sample a hundred times better than anybody else can, and therefore they can discover new drugs, and that's very exciting," Hardiman says.

Another increasingly important tool for life-science innovators is liquid chromatography-mass spectrometry, or LC-MS. The technology, which can separate a chemical sample into its molecular components and then identify them, dates back to the 1970s. We wouldn't normally get moon-eyed over machines that have long been workhorses in the biotechnology, pharmaceutical, agrochemical, and food processing industries, but in recent years LC-MS manufacturers such as Agilent, Thermo Fisher, SCIEX, Waters Corporation, and Shimadzu have built instruments with greatly improved sensitivity, throughput, automation, and analytical software. That means, for instance, that developers of cutting-edge therapies such as antisense oligonucleotides (ASOs) and other RNA- or DNA-based medicines have a ready way to measure whether these molecules are active or stable in the body, helping them move more quickly from preclinical studies to clinical trials.

And another of our portfolio companies, Kanvas Biosciences, is advancing the science of microscopy in a way that will likely lead to better cancer treatments. Since 2021, when we first invested, the company has been using multiplexed spectral imaging of human tissue to measure the way the microorganisms in our GI tracts interact with the human cells that surround them. That data fuels AI models that help the company identify combinations of microorganisms that can, for example, boost the effects of immune checkpoint inhibitors for cancer. Up to now, Kanvas has gathered images using commercial confocal microscopes—a slow process, since these instruments can scan only a single point at a time. But now Kanvas has built a light-sheet microscope that illuminates a whole slice of tissue with a planar sheet of laser light (see Figure 6.2.1), with all the spectral variety and spatial resolution the company needs to identify bacteria and their RNA transcripts. Tissues that would have taken five minutes to image on a confocal microscope can now be scanned in just seconds—giving Kanvas a mountain of new data to understand host-microbiome interactions and design new drugs.

> "Every time you get a new measurement device, you're able to probe deeper, observe things you haven't observed before, and make new discoveries."
>
> 
> **James Hardiman**
> General partner, DCVC

Many of the TechMed, TechBio, and traditional biotech companies DCVC backs are benefiting from such tools, and it's part of the reason we're optimistic that a new wave of powerful, healthspan-increasing therapies is coming. "There's something really exciting about the emergence of new tools and new capabilities," Hardiman says. "Every time you get a new measurement device, you're able to probe deeper, observe things you haven't observed before, and make new discoveries."

DCVC Bio companies are recognizing and reordering the components of biology.

Without exception, companies in the DCVC Bio portfolio consist of top science teams using data and AI to help doctors cure undertreated diseases, help society advance agriculture into a new era, and help designers create objects from new materials.
That strategy is working:

- DCVC Bio–backed **Creyon Bio** struck a deal with Lilly that will bring up to $1 billion in milestone-based payments for the company's assistance building safe, effective RNA-targeting drugs.
- **Radionetics Oncology** earned $140 million in funding, also from Lilly, to boost its research on targeting radiopharmaceuticals to cell-surface receptors on tumor cells.
- **Orca Bio** has created a robotic, machine-learning-driven cell sorting platform that makes bone marrow transplants safer and more effective. In March, Phase III clinical trial results showed that Orca's lead immunotherapy doubled the number of transplant patients who survive free of graft-versus-host disease, and the company is preparing to launch its first product, Orca-T.
- **Elo Life** is running field trials in Latin America with Dole to validate its TR4-fungus-resistant bananas. The fungus threatens to wipe out banana farms around the world.

We think teams using proprietary data and advanced AI, informed by scientific experience, can untangle complex biology and show the way to new medicines, materials, and farming methods, and that's where we've focused our investing. In the following pages, we explore a few fields where this magic combination of people and technologies will soon make a big difference.

OPPORTUNITIES

COMPANIES

Creyon Bio, Elo Life, Grove Biopharma, Latus Bio, Newleos Therapeutics, Orca Bio, Radionetics Oncology

VOICES

Dr. John Hamer, Dr. Kiersten Stead

The first attempts to advance gene therapy into the clinic in the late 1990s and early 2000s ended with dramatic failures that set back progress for a decade or more. Researchers learned much in the process, but while a second wave of gene therapy companies has begun to find ways around the earlier problems, they've still struggled to create therapies with the desired specificity. Today we're seeing a third generation of companies developing exquisitely customized viral vectors and other advanced techniques that can deliver therapeutic genes to the body with far more specificity and at much lower doses and cost.

Gene therapy—the quest to cure diseases caused by single-gene defects by replacing or mitigating the missing or mutated genes—is a beautiful idea that's proved fiendishly difficult to execute. After one clinical trial volunteer, 18-year-old Jesse Gelsinger, died in 1999 from a massive immune reaction to the adenovirus vector used to carry the OTC gene into his body, and after half the children in a subsequent European trial of a gene therapy for immune deficiency developed leukemia, funding for gene therapy disappeared and most researchers left the field.

Much later, researchers and companies ventured back into gene therapy using new delivery vehicles for getting DNA into patient's cells, such as modified adeno-associated viruses (AAV—a natural resident virus in humans) and lentiviruses. Those methods have proved far more specific and therefore safer. But researchers and commercial drug developers have run into new challenges such as vector efficiency—getting the new genes to the right cells, in sufficient numbers to make significant amounts of whatever payload is desired—and difficulties scaling up viral vector production while maintaining quality.

Even when new gene therapies have seemed promising, biotech companies and their pharmaceutical industry partners have run into challenges getting them all the way to the clinic, or convincing insurance companies to cover them once they're approved. In late 2024 Pfizer ended a seven-year collaboration with Sangamo Therapeutics to produce a gene therapy for hemophilia A, sending Sangamo's stock value plunging. The therapy is a one-time treatment using an AAV to deliver a functioning version of the gene encoding a clotting protein that hemophilia A patients lack, and it was a winner in Phase III clinical trials. But it would have come with a high price tag and would have had to compete with cheaper (albeit inconvenient and non-curative) treatments like infusions or chronic injectable drugs. So Pfizer removed the treatment from the market after poor sales.

Similarly, Bluebird Bio withdrew its gene therapy for beta thalassemia from the European market in 2021, saying it couldn't get government payers to cover the initial $1.8 million price tag for the one-shot therapy. The company also ran into delays manufacturing its gene therapy for sickle cell disease. To date, Bluebird, despite having launched three separate gene therapies, has never been profitable. (See Opportunity 7.2 for a discussion of the challenges companies and patients face paying for one-and-done curative treatments such as gene therapy.)

"A number of companies in the space have seen their market capitalization tank, because this is all turning out to be a lot more commercially difficult than people imagined," says Dr. John Hamer, managing partner at DCVC Bio. "Everybody had a different twist or trick, but what they didn't have was a deep understanding of the therapeutic opportunities that were out there."

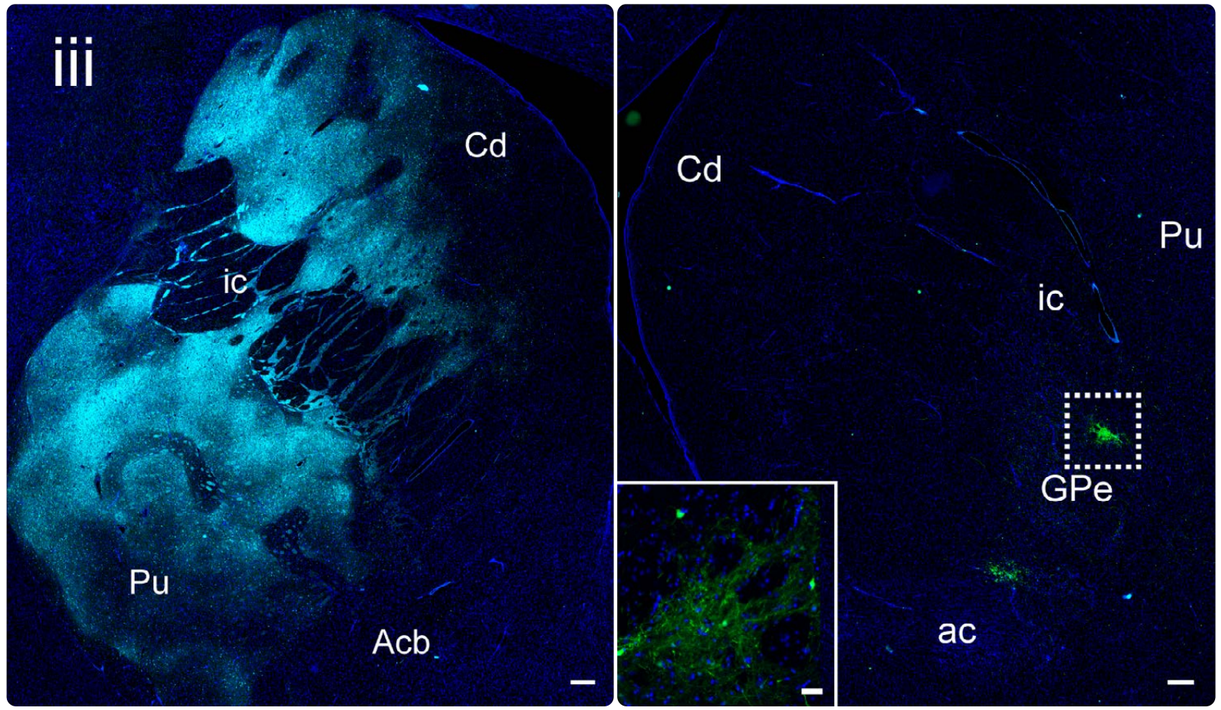

↓ **Latus Bio**
In a rhesus macaque monkey brain, far more neurons took up and expressed a fluorescent reporter gene when delivered by Latus' novel AAV-DB-3 vector (left).

Curative therapies are obviously best for patients, and today a third wave of gene therapy companies has emerged to provide them, using new technologies to create even safer and more targeted ways of delivering payloads to the cells where they can do their work. One of these is Latus Bio, where DCVC Bio first invested in 2024. The company confronted the reality that AAVs, the most common vehicles for modern gene therapy, often lack the needed potency, specificity, and manufacturability, and can still cause tissue toxicity and immune side effects.

An AAV's tropism—the specific types of target cells it can "infect"—is determined by the structure of its capsid, the icosahedral shell that protects its DNA cargo. If the capsid structure is well-matched to the target cell type, the virus can bind, enter, and release its curative payload more efficiently. Latus, founded by Dr. Beverly Davidson of the Children's Hospital of Philadelphia, has built a massively parallel, high-throughput platform for screening a range of modified AAV capsids, looking for those that will deliver genes exactly where they're needed while limiting toxicity.

> "We're starting to see another generation of gene therapies that use some very, very specific targeting technology to target gene therapies to the right cells."
>
> 
>
> **Dr. John Hamer**
> Managing partner, DCVC Bio

Figure 7.1.1

Delivering new genes deep inside the brain In a Latus study, the novel AAV-DB-3 vector transferred genes to many cortical regions (blue), while the older AAV5 only reached cells around the injection area (green).

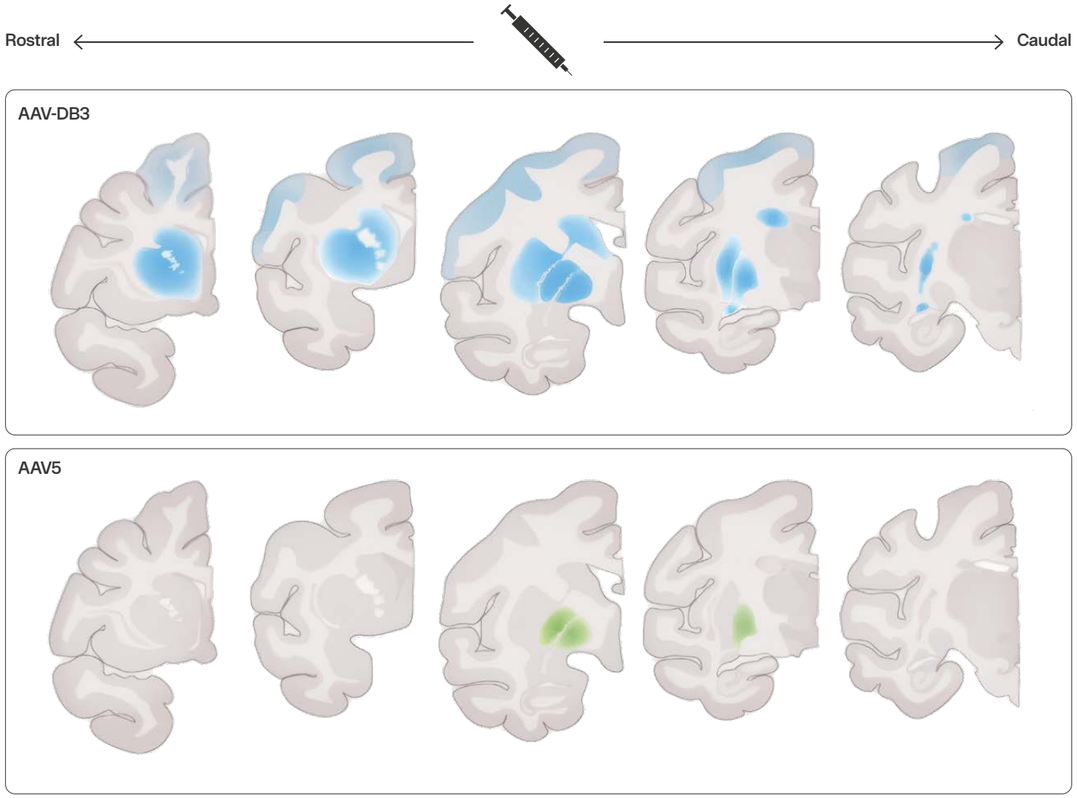

The company has developed capsids that will allow treatment of neurological disorders such as Huntington's disease and Batten disease. It will be able to dose young children who have Batten starting this year, and will follow on soon after with a program to address the devastating life sentence of Huntington's. There, the strategy is not to replace the huntingtin gene itself (the neuronal protein that, when mutated, causes the disease) but to silence the genes for MSH3, which regulates the expansion of the mutated repeat sequence in the disease. Importantly, Latus' capsid-screening platform allows it to find AAVs that are easy to manufacture, and that are effective at delivering large amounts of curative payloads to specific cells in the brain, the kidney, or the heart. This specificity will mean lowering off-tissue toxicity, manufacturing complications, dose, and cost, says DCVC Bio managing partner Dr. Kiersten Stead. "Modified AAVs are going to be the workhorses of the industry for delivering genes and other payloads," she says.

Slowly but surely, then, gene therapy companies are coming to grips with each of the problems that have plagued the researchers since the 1990s. "We're starting to see another generation of gene therapies that use some very, very specific targeting technology to target gene therapies to the right cells," Hamer says.

RNA- or DNA-based treatments such as gene therapy and cell therapy have come with high upfront costs, but are curative, and are more economical than chronic treatments over the course of a patient's life. Regulators should work with insurers to cover these costs, and with insurers and manufacturers to introduce new business models such as gradual payment plans.

In the previous section we put a spotlight on the third-generation cell and gene therapies that are offering prospective cures for previously intractable diseases. Developers and manufacturers, with high R&D and manufacturing costs, have put seemingly eye-popping price tags on many of these drugs. For example, Orchard Therapeutics' drug Lenmeldy, an infusion of genetically modified stem cells approved for the treatment of the rare neurological disease metachromatic leukodystrophy (MLD), was priced at $4.25 million, a world record. Hemgenix, CSL Behring's gene therapy for hemophilia B, costs $3.5 million, while Zynteglo, Bluebird Bio's gene therapy for beta thalassemia, costs $2.8 million.

↗ **Lenmeldy,** a stem cell infusion that stops or slows MLD with a single treatment, is priced in the U.S. at $4.25 million.

But before you spit out your coffee, consider the bigger economic picture. The standard therapy for hemophilia B includes regular prophylactic infusions of clotting factor IX, at an annual cost of $700,000 to $800,000 per patient and a lifetime cost of $21 million to $23 million. Three years after their single-dose Hemgenix treatments, 94 percent of patients no longer needed prophylactic treatment. That makes the CSL Behring therapy look like a real bargain.

Here's the challenge: payers such as insurance companies, self-insured employers, and governments have figured out how to cover the costs of most chronic treatments, but they're having trouble digesting the high upfront expense of one-and-done, curative treatments. More specifically, they know that individuals often change insurance providers, so they're incentivized to reject coverage of curative treatments, in the hope that the patients who need them will simply leave and reduce their insurers' costs. The mismatch creates access barriers for patients and unpredictable economics for manufacturers. "The problem has been the recalcitrance of payers to adapt—of regulators to say to insurers, 'You have to cover these,' or of insurers to put in a business model that makes sense, like a payment plan that is more spread out," says DCVC Bio's Stead.

What kind of business model would make sense? Some manufacturers already offer outcomes-based agreements under which payers and patients receive a rebate if the therapy isn't as effective as hoped. These arrangements are intended to shift some of the risk away from payers, but they require rigorous tracking, and they delay rather than lower high upfront payments for patients who do respond. In short, this model is not working.

We're more attracted to several other innovative proposals. One is a subscription model where payers would pay a fixed annual fee to access a therapy for a defined patient population. Quantile Health, a new data-science company we admire, has introduced a Netflix-like reimbursement platform in which payers pay a fixed cost reflecting the underlying risk of their patient population, and manufacturers agree to provide access to treatments to anyone in these populations. Another model is amortization, which would allow payers to spread the costs of expensive medicines across many years.

All of these models would require closer collaboration among regulators, insurers, and manufacturers to ensure that groundbreaking treatments are accessible. "There's got to be some sort of institutionalized or regulated cooperation to make sure that one-and-dones get covered," Stead says, "because right now we're on a path to incentivize the U.S. healthcare system to prioritize chronic treatments even when curative treatments are available."

Proposed Payment Model	Pro(s)	Con(s)
Status quo	▪ Lets the market set prices ▪ Incentivizes innovation	▪ High upfront costs for one-time curative treatments ▪ Payment is not necessarily tied to efficacy
Efficacy-based	▪ Cost is tied to health outcomes ▪ Incentivizes data collection	▪ Outcomes may not be measurable ▪ Stakeholders may not agree on efficacy
Mortgage	▪ Insurers can pay over time, increasing access to costly one-time treatments ▪ Can increase affordability of high-cost drugs	▪ Costs are spread out rather than reduced ▪ Continued payment even for poor outcomes
Indication-specific	▪ Good for drugs that have multiple uses ▪ Price can also be tied to volume	▪ Can cause confusion for payers and patients ▪ May lead to fragmented care
Netflix	▪ All-you-can-treat, including costly one-time curatives, for one fee ▪ Can reduce overuse of treatments like antibiotics	▪ Determining pricing structure may be difficult ▪ Delinking price to volume may lead to market instability

Figure 7.2.1

Learning to love costly one-time treatments
Today's insurance system is broken, incentivizing chronic treatments even when they have higher lifetime costs.

> There's got to be some sort of institutionalized or regulated cooperation to make sure that one-and-dones get covered.

Dr. Kiersten Stead
Managing partner, DCVC Bio

Autophagy is the cell's mechanism for breaking down and cleaning up accumulated junk, and it's inhibited or absent in some cells in people with Alzheimer's disease and Parkinson's disease. We're excited about the idea of creating autophagy-modulating drugs that deliberately target and degrade the messy macromolecular complexes often found in diseased cells.

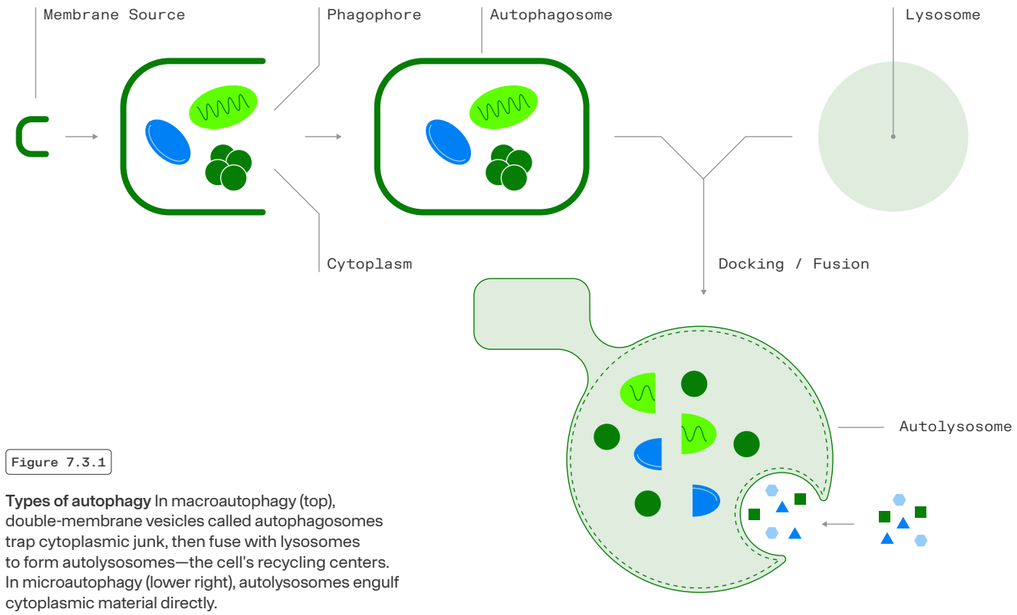

Figure 7.3.1

Types of autophagy In macroautophagy (top), double-membrane vesicles called autophagosomes trap cytoplasmic junk, then fuse with lysosomes to form autolysosomes—the cell's recycling centers. In microautophagy (lower right), autolysosomes engulf cytoplasmic material directly.

The word *autophagy* comes from the Greek for "self-eating." The process is part of regular maintenance inside cells, which must remove and recycle clutter such as damaged organelles and misfolded proteins in order to maintain homeostasis. It works by enveloping cellular material in vesicles called autophagosomes. These fuse with special digestive organelles called lysosomes, which carry dozens of types of enzymes that cut the material up into recyclable components. When this junk-collection process is working correctly, it suppresses aging and prolongs life. When it's not, the result can be cancer and neurodegenerative disease.

Among the causes of Parkinson's disease, for example, are mutations in the genes *PRKN* and *PINK1*, which encode proteins that normally help with mitophagy, the breakdown of impaired mitochondria. As they release energy, mitochondria also release reactive oxygen species (ROS), which can damage DNA, RNA, proteins, and lipids; so when failing mitochondria pile up in the cell, so do ROS. In the brain, this process can damage critical dopamine-producing neurons near the basal ganglia, which regulate voluntary motor control, leading to the tremors and slow movement familiar to Parkinson's patients. Separately, impaired autophagy allows the buildup of α-synuclein and tau proteins, which accumulate into structures called Lewy bodies and neurofibrillary tangles, further damaging the dopaminergic neurons.

> "What we're really interested in is, can we trigger autophagy in certain disease and cell types to get rid of damaged organelles or large protein complexes that are causing problems?
>
> **Dr. John Hamer**
> Managing partner, DCVC Bio

Existing medications for Parkinson's disease manage the symptoms—for example, by increasing dopamine production—without fixing the underlying causes. If researchers could identify a drug that restores normal autophagy in the affected cells, it could lead to an actual cure for Parkinson's. And the same could be true for other hard-to-treat problems where impaired autophagy plays a role, including diabetes, heart disease, inflammatory bowel disease, non-alcoholic fatty liver disease, and cancer. "What we're really interested in is, can we trigger autophagy in certain disease and cell types to get rid of damaged organelles or large protein complexes that are causing problems?" explain DCVC Bio's Hamer.

It's well-known that a master signaling protein called mTOR helps to suppress autophagy, and that mTOR inhibitors such as rapamycin can induce it. But that may be a blunderbuss approach, given that mTOR has so many other important functions. (Starvation, too, may inhibit mTOR and induce autophagy, but few would consider it a cure for cancer or neurological diseases.) Today, more and more researchers are turning their attention to the search for more precise autophagy-modulating agents.

What we'd love to see is a systematic, high-throughput, data-driven search for molecules or compounds that can regulate autophagy precisely—activating it only in specific cell types, or only to destroy specific kinds of proteins, without undesirable off-target effects. That will likely require a far better understanding of the molecular mechanisms of autophagy, better assays to measure and monitor it, and (eventually) rigorous testing of possible interventions in animals and human clinical trials.

Here at DCVC Bio, we are incubating a company in this area, so watch this space. "We have some drugs that trigger autophagy accidentally, but there's nothing where we can deliberately target and degrade a large macromolecular complex," says Stead. "That's what we're pretty interested in."

Anxiety is a massive, costly, and historically undertreated mental health scourge. The best we can say about existing drugs and treatment regimens is that they're better than nothing. Here at DCVC Bio, we see a need for a massive reinvestment in the search for drugs that safely lessen anxiety, including existing compounds that have been shelved by big pharma companies.

There's no mental health problem more widespread than debilitating anxiety, a category encompassing generalized anxiety disorder, social anxiety disorder, panic disorder, and phobias. According to the NIH, about 19 percent of adults had an anxiety disorder in the past year—more females (23.4 percent) than males (14.3 percent)—and about 31 percent of American adults will experience an anxiety disorder at some point in their lifetimes. The World Health Organization says the condition affects more than 300 million people globally, but that only 28 percent of people who need treatment actually get it.

Cognitive behavioral therapy (CBT) teaches important skills for dealing with anxiety. It works even better in combination with drug treatment such as selective serotonin reuptake inhibitors (SSRIs) and serotonin-norepinephrine reuptake inhibitors, also commonly used as antidepressants. Benzodiazepines can also help for a limited period, but they can cause fatigue, cognitive impairment, and addiction. (There's also a groundswell of interest in hallucinogenic compounds such as psilocybin and ketamine for the treatment of anxiety and other conditions, but they're difficult to study in a controlled fashion, and whatever their merits, they're not scalable to millions of people as long as treatment protocols require a physician be present to administer the drug and monitor the patient.)

Figure 7.4.1

The economics of mental health Helping people live happier, less-stressed lives would also reduce a huge drag on the U.S. economy. Source: National Bureau of Economic Research.

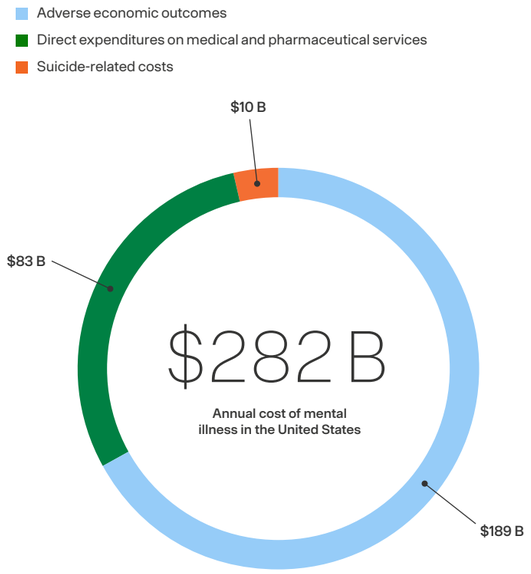

The truth is that no existing therapy is as effective as medical professionals and patients would like. Most people relapse after treatment ends—in one study, anxiety returned for 26 to 45 percent of patients who discontinued antidepressants, and another study showed that 48 percent of patients who received CBT were still symptomatic 2 to 14 years later.

"There's a real need in this space for good anxiety drugs, without the side effect profiles of benzodiazepines and other medicines, that can treat general anxiety, social anxiety, PTSD, and a number of different indications," says DCVC Bio's Hamer.

Our fund placed its first bet in this area in February 2025, participating in a Series A financing round for Boston-based Newleos Therapeutics. The company launched with four drugs in its pipeline, all licensed from Roche. Newleos' lead drug candidate, NTX-1955, is a molecule that enhances signaling in the brain by gamma-aminobutyric acid, or GABA. GABA is an inhibitory neurotransmitter that lets chloride ions into neurons and makes them less likely to fire. By slowing down this firing, the GABA system helps to keep us calm, which is why it's already the target of numerous anti-anxiety drugs.

→ **Structural and functional MRI meta-analysis** Brain regions involved in emotional processing are consistently altered in patients with generalized anxiety disorder. Source: Brazilian Journal of Psychiatry.

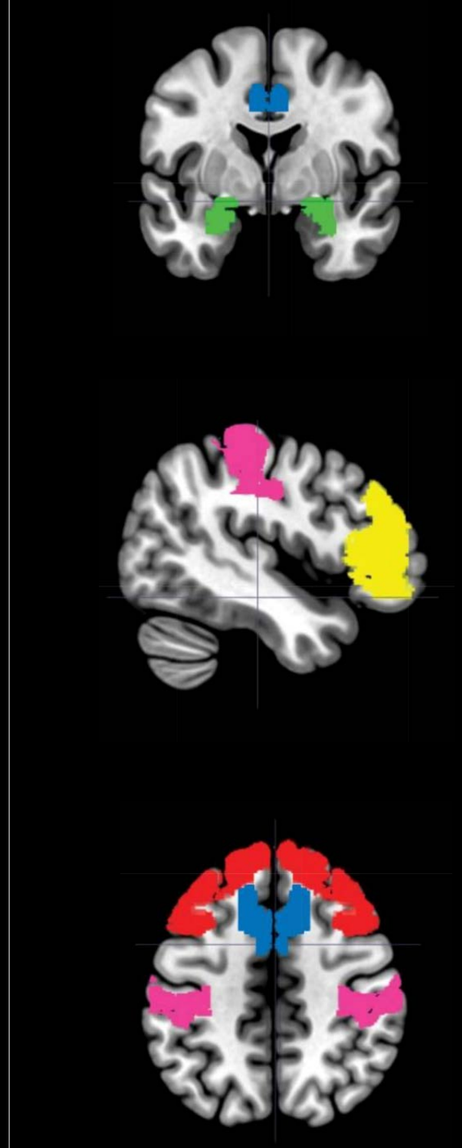

↑ Gray matter volumes and hyperactivation:

- Amygdala
- Anterior cingulate cortex
- Postcentral gyrus

↓ Gray matter volumes and hypoactivation in:

- Dorsolateral prefrontal cortex

↓ Gray matter volumes and hyperactivation in:

- Ventrolateral prefrontal cortex

We're also watching with high hopes as other companies explore new approaches to treating anxiety:

A drug called etifoxine has dual effects on the GABA system. It binds to certain subunits of the $GABA_A$ receptor directly, enhancing its action, while also stimulating production of neurosteroids that themselves up-regulate the receptor. Etifoxine has long been approved for use against anxiety in France and a few smaller markets, and it's considered as effective as benzodiazepines, but with fewer side effects. California-based GABA Therapeutics says it will move into clinical trials of GRX-917, an analog of etifoxine that's easier to metabolize, this year.

Swiss startup Synendos Therapeutics is moving toward clinical trials of a selective endocannabinoid reuptake inhibitor called SYT-510. Endocannabinoids are natural substances in the body that activate the same receptors that are sensitive to phytocannabinoids such as THC, the psychoactive molecule in cannabis. By limiting the reabsorption of endocannabinoids by neurons, the drug would, in theory, help restore dysregulated neurotransmission in the brain, treating anxiety and mood disorders.

Tezampanel, which blocks receptors for glutamate, the most common neurotransmitter in vertebrates, was developed by Lilly as a treatment for migraines but later shelved. Animal studies show it may be an effective anti-anxiety drug.

"

There's a real need in this space for good anxiety drugs, without the side effect profiles of benzodiazepines and other medicines.

Dr. John Hamer
Managing partner, DCVC Bio

In general, we like the idea of refilling the pipeline of potential anxiety treatments by restarting research on shelved, abandoned, or orphan compounds. There's even hope that AI could accelerate this repurposing process. One group of researchers writing last year in *Nature Medicine* showed that a graph neural network model, when allowed to comb through decades of medical literature, was efficient at exploring shared molecular pathways and pathologies and identifying candidate drugs for conditions that lack approved treatments.

We know that the world needs better anxiety treatments as soon as possible—and the advantage of the Newleos-style in-licensing approach is that when a drug has a known safety and efficacy profile, moving it to the clinic can be far faster and cheaper.

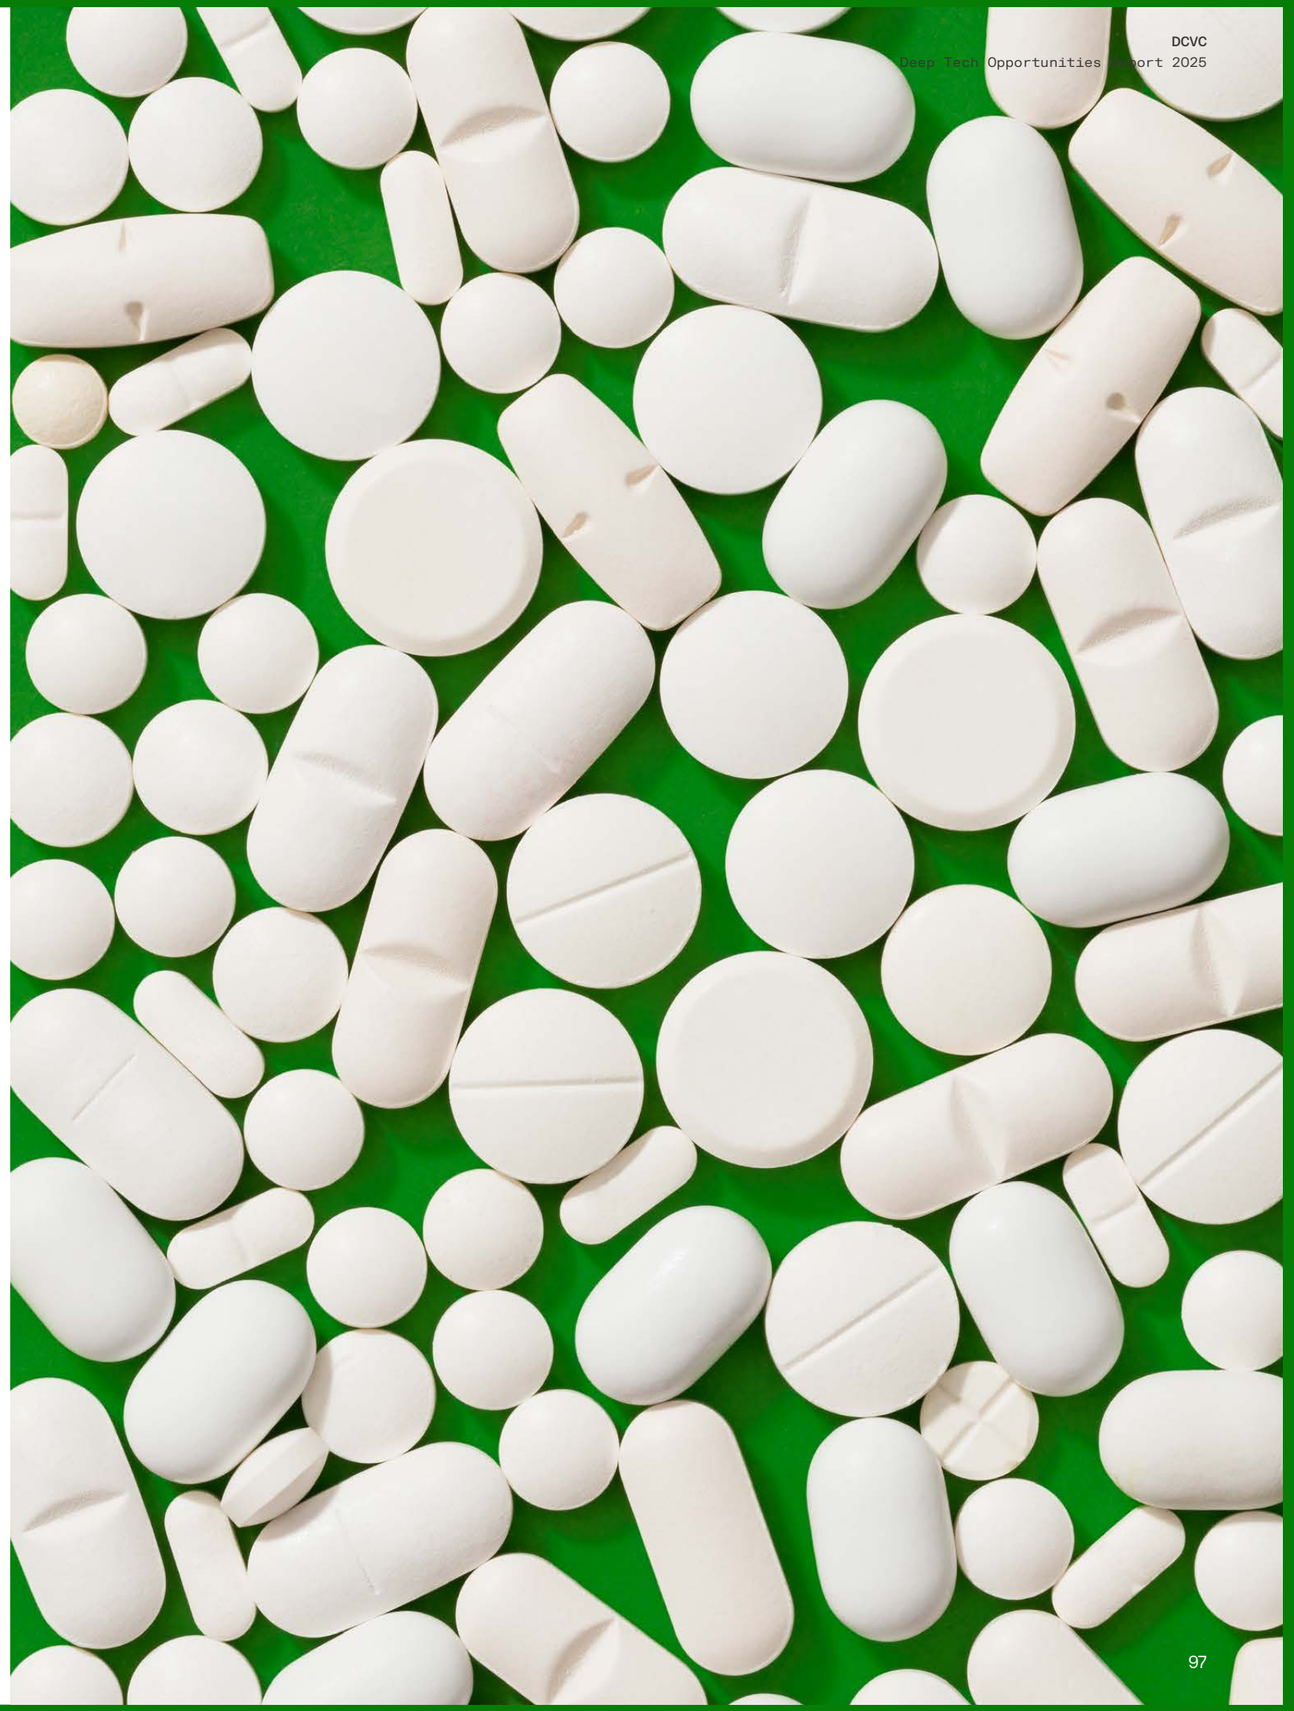

There's an exciting new way to put protein-like drugs inside cells, where they can help against hitherto untreatable cancers and other diseases like neurodegeneration.

The modern pharmacopeia has generally been divided into small-molecule drugs that can easily cross the cell membrane and hit targets in the cytosol or the nucleus, and large-molecule drugs that are too big to slip through and that therefore work outside of cells, usually by modifying signaling *between* cells. It's been nearly impossible to design and ferry large-molecule drugs into the cell interior and keep them in one piece—until now.

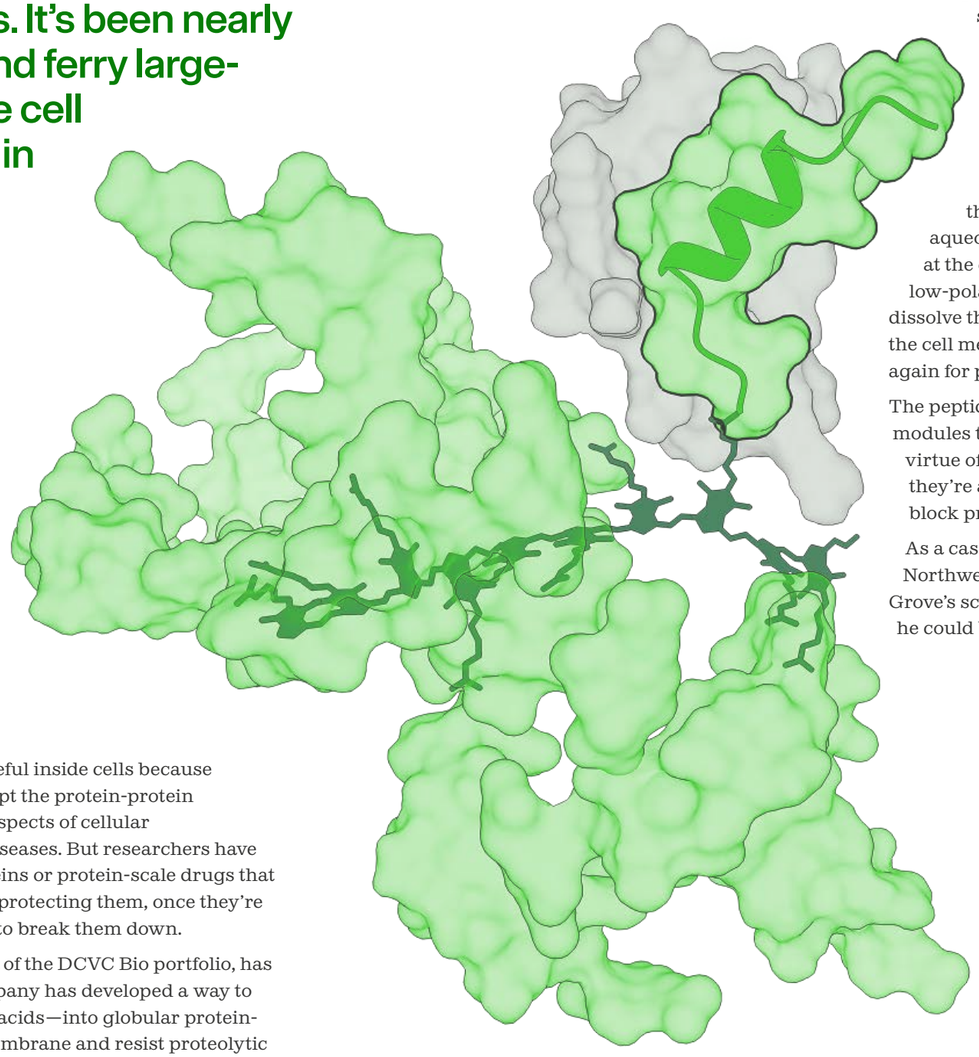

→ **Packaging complex peptides**
Grove Biopharma can tune the properties of its Precision-Linked Proteomimetics by attaching different peptide branches to the molecules' spines (dark green).

Big, protein-sized drugs could be useful inside cells because smaller ones typically can't interrupt the protein-protein interactions that control so many aspects of cellular functioning—and that go wrong in myriad diseases. But researchers have had great difficulty *(a)* designing actual proteins or protein-scale drugs that can pass through the cell membrane, and *(b)* protecting them, once they're inside, from the proteolytic enzymes that try to break them down.

Chicago-based Grove Biopharma, a new part of the DCVC Bio portfolio, has come up with an ingenious answer. The company has developed a way to package complex peptides—chains of amino acids—into globular protein-like structures that can travel through the membrane and resist proteolytic degradation. Once inside, these molecules can interact with the disordered protein targets involved in cancer and other conditions.

Grove calls these new molecules Precision-Linked Proteomimetics, or PLPs. Each PLP is a customizable synthetic molecule consisting of a central chain with multiple side chains. If laid out flat, a PLP would resemble an array of bristles protruding from the head of a brush. Each bristle is a separate peptide, and the peptides are attached via linker units to backbone precursors, which, in turn, link up to form the head of the brush—the molecule's spine (see figure). These versatile molecules represent "a fantastic evolution of biotechnology," says DCVC Bio's Stead.

One talent of PLPs is that they can respond dynamically to different targets and environments. When the peptide bristles are folded up, they create a kind of shell around the spine. The bristles' highly polar electrical properties make the PLPs water-soluble in this state, meaning they can travel well in the aqueous environments inside and outside cells. But at the cell wall, the PLPs can unfold, exposing the low-polarity backbone, which allows the molecules to dissolve through the phospholipid bilayer that makes up the cell membrane. Once inside, the molecules fold up again for protection from proteolysis.

The peptides themselves, along with any Grove-designed modules they may carry, are the important cargo. By virtue of their custom-designed amino acid sequences, they're able to bind to specific proteins inside cells and block processes that can lead to disease.

As a case study, Dr. Nathan Gianneschi—the Northwestern University chemical biologist who is Grove's scientific founder—showed in a mouse model that he could build a targeted PLP that blocks an important neuronal protein, VCP, from binding with mutant huntingtin protein. When natural peptides block this binding, it's been shown to prevent the mitochondrial breakdown that characterizes Huntington's disease.

And that's just one example of the many kinds of protein-protein interactions that could be modulated using PLPs. The company is already building PLPs that can interfere with the action of Myc, a family of transcription factor proteins that activate genes driving cell proliferation and are known to work overtime in most tumor cells. Scientists think that inhibiting Myc directly would slow or stop tumor growth, but despite massive efforts, they've been unable to find drugs that target Myc itself—it's considered "undruggable." The hope at Grove is that disrupting the interaction between Myc and WDR5, a separate protein recruited by Myc to activate its target genes, would slow many cancers where Myc is overexpressed, such as mixed-lineage leukemia (an aggressive form of acute myeloid leukemia). Grove also plans to develop a treatment for advanced prostate cancer by targeting AR-V7, an androgen receptor variant that continues to promote cell growth even in men who've had androgen-blocking therapy.

Stead says DCVC Bio has long wanted to invest in a company with a technology for binding to and inhibiting complex, disordered targets inside cells. "And when we came across Grove, we realized that we'd stumbled onto exactly that," she says. "Grove's peptides have known structures that bind to known targets. By themselves these peptides can't get into cells. But when they're configured into the Grove PLP format, all of a sudden they can become very specific and potent drugs."

> By themselves these peptides can't get into cells. But when they're configured into the Grove PLP format, all of a sudden they can become very specific and potent drugs.
>
> **Dr. Kiersten Stead**
> Managing partner, DCVC Bio

The industrial renaissance sweeping through the U.S. economy will yield faster, cheaper, more effective defense and intelligence systems.

The reinvention of U.S.-based manufacturing described in Chapter 1.0 will be a boon for aerospace and defense companies, which help the nation outpace , contain, and deter strategic opponents. At the same time, a revitalized defense and security infrastructure, operating at lower cost, will help protect our new investments in industrial technology. This flywheel will be powered, in part, by startup innovation.

Founding members and operating partners at DCVC have spent the better part of 40 years in pursuit of national defense and intelligence advantage. We're grateful that the U.S. has a strong base of large, experienced, deep-pocketed defense contractors—the so-called primes. But as we observe asymmetric threats such as drones shaping the conflicts in Ukraine and the Middle East, and as we watch the U.S. enter into a new space race with China, we think the Defense Department's new willingness to source technology from younger, nimbler companies such as Anduril, Palantir, and Hadrian is a welcome change.

Anduril, for example, recently beat Boeing, Lockheed Martin, and Northrop Grumman for a contract to prototype an autonomous fighter jet for the Air Force and Navy, and it's also supplying Microsoft with software for an Army augmented-reality headset—projects that could make U.S. armed forces more capable and efficient. Ultimately, collaborations between innovative startups and primes can leverage the strengths of both kinds of companies.

OPPORTUNITIES

COMPANIES

Capella Space, Fortem Technologies, Impulse Space, Planet, Reality Defender, Rocket Lab, SentinelOne

VOICES

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

Rocket Lab designed the fairing and first stage of its Neutron "mega constellation launcher" to be fully reusable.

Startup-built technologies are bringing us new ways to dominate in the air and in space.

Building a community of smaller defense providers in the United States is a critical part of the larger reindustrialization effort we talked about at the beginning of this report. We're proud to be part of that emerging ecosystem through our portfolio companies, which are making space more accessible and navigable, defending our secure facilities and armed forces against asymmetric threats, and helping us understand what's happening here on Earth.

One of several investments that established DCVC as a leading space systems investor was Rocket Lab. The company went public in 2021 and is making great strides as a provider of small satellite launch services using its Electron rocket and its unique 3D-printed Rutherford engine. The company's goal is to make space access more affordable, and it has already become the second-busiest U.S.-based launch provider, deploying 204 satellites to low Earth orbit (LEO) in 58 launches. It's now moving into the medium-lift market with its reusable Neutron rocket, which will see its first test flight this year.

↗ **Rocket Lab**
The Electron delivers small satellites (up to 300 kg) to low-earth orbit.

↑ **Impulse Space**
The company is filling an unmet need for quick, affordable orbital transfer.

"Since the first U.S. satellite was launched into orbit in 1958, space has been the domain of wealthy governments and multinational contractors that could afford the multibillion-dollar cost of every mission," says DCVC's Ocko. "Rocket Lab changes that math." So does Ursa Major, another space startup we admire; they're also using additive manufacturing to build a variety of rocket motors suitable for everything from hypersonic missiles to heavy launch vehicles.

Impulse Space, a more recent investment for us, aims to reduce the costs in a different part of the launch business: moving payloads from LEO to medium Earth orbit (MEO) and geosynchronous orbit (GEO). The company's Mira spacecraft, which first flew in 2023, is an orbital transfer vehicle that can precisely maneuver payloads of up to 300 kilograms. It's built for rendezvous and proximity operations, such as the planned refueling of the U.S. Space Force's Tetra-5 spacecraft in 2026. And its Helios booster, scheduled for its first flight in 2026, is designed to ride a larger rocket such as a SpaceX Falcon 9 to LEO, then rapidly lift 5+ tons of payload to MEO or GEO.

"

SpaceX is like a bus that drops you off in LEO, and then Impulse takes you up to your proper orbit, much less expensively than SpaceX can.

Matt O'Connell
Operating partner, DCVC

"SpaceX is like a bus that drops you off in LEO, and then Impulse takes you up to your proper orbit, much less expensively than SpaceX can," explains DCVC operating partner Matt O'Connell. "There's a lot of business to be done there, including a new generation of communications satellites. And there's significant government interest in what they're building. If a government satellite were to stop functioning, being able to get a replacement satellite up quickly and inexpensively has a lot of appeal." Just as SpaceX has helped the U.S. space business maintain its advantage in low-cost access to LEO, in other words, Impulse Space will give the nation greater flexibility and dominance in medium and high orbits.

Two DCVC-backed Earth observation companies continue to demonstrate why commercial access to space is so important. Capella Space owns a constellation of synthetic-aperture radar (SAR) satellites that can see the surface at night and through clouds, at very high resolution and with frequent repeat visits. That combination of capabilities is useful to the U.S. government, which relied on Capella's images to predict the Russian invasion of Ukraine. Capella announced in May that it had won a contract to develop new SAR image acquisition modes to support the DoD's Hybrid Space Architecture, which is designed to provide secure communications for military operations. That same month, quantum computing leader IonQ [NASDAQ: IONQ], which is establishing a space-to-space and space-to-ground quantum internet, said it had signed a definitive agreement to acquire the company.

Similarly, Planet's Dove, SkySat, and next-generation Pelican constellations are built to gather medium- and high-resolution optical images of every spot on Earth's landmass, helping with environmental monitoring, security, and intelligence gathering. While the company has long focused on data services, it's also begun to build dedicated Pelican satellite constellations for specific customers, including one that recently signed a seven-year, $230 million contract, Planet's biggest ever. (Both Capella and Planet have partnered with Rocket Lab to deploy their constellations.)

Closer to the planet's surface, Fortem Technologies has been showing how venture-backed companies can help against the growing threat of unmanned aerial systems, or UAS. Whether in the form of store-bought, rotor-driven drones or fixed-wing reconnaissance or munitions drones, these systems are considered by the DoD to be "an inexpensive, accessible, flexible, expendable, and plausibly deniable way to carry out armed attacks and project outsized power." Fortem's DroneHunter F700 interceptor is designed to counter all types of UAS with measures such as tether nets or parachute nets. It's in use in Ukraine to detect and capture Iranian Shahed-136 "kamikaze" drones deployed by Russia. With the aid of Fortem's real-time risk assessment software, it can also be used to secure the airspace around sports stadiums and other event venues.

An American defense-tech startup we respect (but haven't invested in) is Shield AI, whose innovative Hivemind system works as an autonomous AI pilot for small drones, UAS, and fighter jets. Brandon Tseng, Shield's president and co-founder, has warned that China's military is moving at "blazingly fast speeds" to build disruptive weapons such as hypersonic missiles, while the U.S. largely fails to invest in transformational defense products. We might not put it that harshly, but we think this idea is directionally correct. To compete with peer, near-peer, and asymmetric adversaries, the nation needs radically new ideas—and many of those ideas are bubbling up at small companies.

↑ **Fortem Technologies**
The DroneHunter F700 can capture or disable an offending drone using its net guns and parachute. Captured drones can be towed to a specified location.

↓ **Planet**
Under a contract with NASA, the company provides satellite data such as this image of the Athabasca Sand Dunes in Saskatchewan, Canada.

↑ **Capella Space**
Technicians prepare an Acadia satellite for launch at Rocket Lab's Launch Complex 1 in New Zealand.

America, its allies, and their biggest companies operate in an environment of constant cyberattacks and threats of theft and sabotage against critical systems. Venture-funded companies can help defend against this kind of threat.

SentinelOne, an early DCVC portfolio company, grew into one of the preeminent public companies defending against cyberattacks. The company's AI autonomously collects and analyzes data from email, cloud environments, networks, and all the devices connected to them in order to detect and respond to cyber threats in real time. We invested in SentinelOne at its founding in 2013, it went public in 2021, and it's now used by four of the Fortune 10 companies and hundreds of the Fortune Global 2000.

Today, we continue this work with our investment in Reality Defender, which has built a deepfake-detection system that can help governments and corporations flag AI-generated content and manipulation in images, video, audio, and text files. Social engineering scams employing generative AI have rapidly evolved from a theoretical threat into a real one: In 2024 a U.S. senator was tricked into participating in a call with a deepfaked version of Ukrainian Minister of Foreign Affairs Dmytro Kuleba, and an employee at U.K. engineering firm Arup transferred $25 million to criminals after a deepfaked video call. A survey by compliance platform provider ISMS found that 35 percent of U.S. businesses have experienced a deepfake incident in the past year, and that it's now the second most common type of cybersecurity attack.

35%

of U.S. businesses have experienced a deepfake incident in the past year, and it's now the second most common type of cybersecurity attack.

> "[Reality Defender has] been at it for longer, they've been pioneers in this field, and they've been creating awareness that this is a real problem now and will become a larger problem in the future.

Ali Tamaseb
General partner, DCVC

Reality Defender is the first and most respected company to address the problem, winning customers in banking, transportation, defense, and other areas where fraud could cost fortunes, reputations, or even lives. Customers plug into its ensemble of detection models via a web application, or an API integrating directly into enterprise or on-premises software. Reality Defender is currently working with ElevenLabs, a creator of voice-cloning technology, to enhance its algorithms for detecting AI-generated audio. (That deal came after ElevenLabs' software was used to doctor a recording of former president Joe Biden that reached thousands of voters in New Hampshire as part of an illegal voter-suppression campaign.) And it's building a plugin for Zoom that can advise users in real time about whether the person they think they're seeing on their screen is a fake.

"Customers realize this is the best solution out there," says DCVC partner Ali Tamaseb. "They've been at it for longer, they've been pioneers in this field, and they've been creating awareness that this is a real problem now and will become a larger problem in the future. That's why every type of company is interested, from financial services to public relations to banks to Zoom."

DCVC's Ocko points out that effective cybersecurity is an important element of national defense, given that the internet touches so much of our infrastructure. "This gray zone of 360-degree warfare that we face is one of very nasty, very persistent cyberattacks against water, power, communications," Ocko says. "From the very first, from SentinelOne through Reality Defender, we've been anchor investors in this area."

Reality Defender Deep Fake Analysis
Today's generative AI models can create a deepfake video from as little as 10 seconds of footage of a real person. Reality Defender uses its own advanced machine-learning models to detect synthetic faces, manipulated movements, and other telltale signs in applications such as web conferencing and authentication.

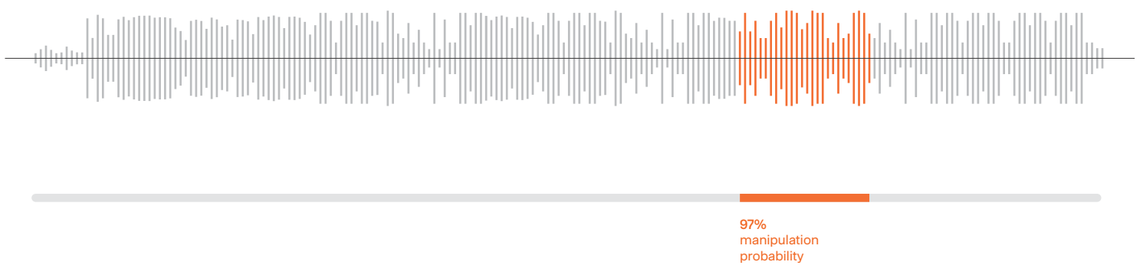

Henry Park
2% manipulation probability

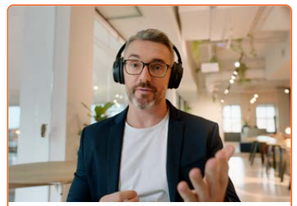

Glenn Rodgers
85% manipulation probability

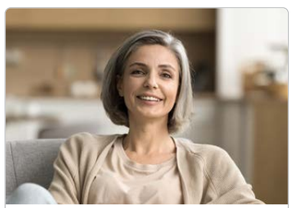

Monica Andersson
4% manipulation probability

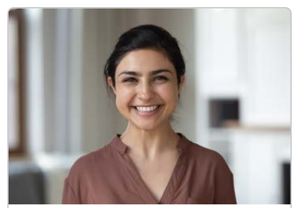

Millie Johnson
3% manipulation probability

Autonomous surface vehicles (ASVs) and autonomous underwater vehicles (AUVs) have a range of defense applications, including reconnaissance, surveilling coastal defenses, and deterring expansion-minded adversaries such as China. We've looked closely at several commercial ASV and AUV makers, and we've already invested in one company that may be able to supply the power sources these vehicles need to travel or loiter for long periods.

Earlier in this report we talked about the need for more autonomous technology in warehouses (see Opportunity 2.2) and in the sky (see Opportunity 6.1). But for industrial, security, defense, and scientific purposes, the nation also needs to build bigger fleets of autonomous vehicles that work underwater.

AUVs can help oil and gas companies map the seafloor before they build drilling infrastructure or lay pipelines. Defense agencies use them to monitor shipping lanes, deploy mine countermeasures, track enemy submarines, surveil and protect underwater cables and pipelines, and deliver payloads to strategic underwater locations. And that's to say nothing of the scientific or forensic uses of AUVs, such as studying life on the ocean floor or assisting in air crash investigations.

The obvious advantage of robots over conventional submarines for such applications is that they don't risk human lives. They also operate at lower cost and require less maintenance and fuel.

We haven't invested in this sector yet, but we're impressed by a number of young companies that are helping to advance the state of the art. One is Ocean Aero, whose innovative solar-powered Triton vehicle can both sail on the surface and dive, carrying payloads such as video and thermal imagers, sonar, radar, acoustic arrays, or other sensors. Saronic Technologies makes three speedboat-like ASVs, from 6 feet to 24 feet in length, that can be used alone or in swarms to deploy sensors or munitions. And ThayerMahan builds a variety of surface and sub-surface vehicles, including Outpost, a wave glider with a towed acoustic array that can loiter in a given location for up to three months, and SeaScout, a towed platform carrying synthetic-aperture sonar—analogous to Capella Space's synthetic-aperture radar—for making high-resolution images of the seafloor and its contents.

"When Admiral Michael Gilday was the chief of naval operations, he made the point that autonomy is the future," says DCVC operating partner Clay Hutmacher. "It's a way to maintain our advantage over the Chinese, as it's one of the few areas where we still enjoy a technological edge." (Gilday told CBS News in 2023 that 40 percent of the U.S. fleet will be unmanned by 2045.)

But one of the factors holding back the development of all undersea technology, including AUVs and ASVs, is power. Most autonomous vehicles rely on lithium-ion batteries, which can power a mid-size AUV for only three days or so. The batteries for remote devices such as deep-ocean listening stations last longer, but not long enough for comfort. "If you have to go replace a battery for a detection net on the bottom of the sea, that's a whole combat operation," Hutmacher observes.

One alternative is radioisotope power systems, or RPS, an area we've come to understand through our investment in Zeno Power Systems. The company makes shoebox-sized "nuclear batteries" in which the heat generated by the decay of radioisotopes from nuclear waste is used to produce electricity. Long used by NASA for interplanetary missions, RPS units can be made small and safe enough to power small satellites and underwater vehicles, Zeno believes; in 2023 the company won a DoD contract to demonstrate an RPS on the seabed, with possible applications including long-endurance sensors and charging stations for AUVs. In 2024 Zeno signed a deal with Atomic Alchemy—which is being acquired by another DCVC portfolio company, Oklo—under which Atomic Alchemy will supply Zeno with strontium-90, americium-241, and other radioisotopes needed for its system.

"If you can put a power source on [an AUV or deep-ocean sensor] that lasts 50 or 100 years, that's a complete game changer," Hutmacher says. "Then you can start to imagine underwater loitering munitions that you have down there for a rainy day, like a sleeper agent waiting to be activated and deployed if necessary." Given tensions with China, says Hutmacher, the Pentagon is likely to put funding behind such ideas, which "might be one of the few ways we can maintain any kind of deterrent presence long-term" in sensitive areas like the South China Sea and the Taiwan Strait.

"Autonomy is the future…It's a way to maintain our advantage over the Chinese, as it's one of the few areas where we still enjoy a technological edge."

Clay Hutmacher
Operating partner, DCVC

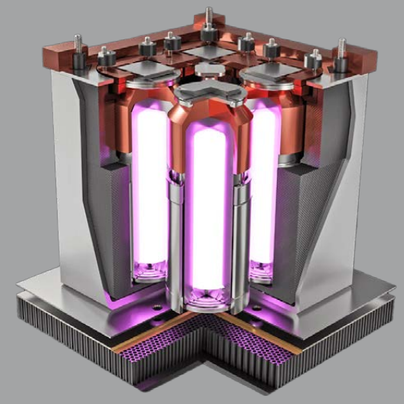

Figure 8.3.1

Zeno's nuclear batteries convert the natural decay of radioactive material into reliable thermal and electrical power.

We ended the previous section by noting that electric power is a nettlesome constraint on the range and capabilities of autonomous craft in the oceans. The same is true in the sky and on the battlefield. One solution could come in the form of wireless transfer of electrical energy via radio frequency waves, lasers, or electromagnetic fields.

When it comes to powering their drones, ground-based robots, remote sensors, or soldier-worn equipment, defense manufacturers have long had the same two choices as carmakers: internal combustion engines or batteries. Engines run longer but require fuel and maintenance. Batteries are quieter but have limited endurance, often resulting in drone flight times of less than an hour; and in combat conditions, recharging or swapping batteries isn't always an option.

Wireless power transmission (WPT) has the potential to supply these devices with continuous power or in-flight recharging. In this sector, we're a proud backer of Redwood City, Calif.-based Reach Power.

Like other WPT companies, Reach takes advantage of the simple fact that radio waves carry energy. In the case of conventional AM or FM radio, it's microwatts or milliwatts of energy—just enough to drive the electronics inside a radio—and the radio waves are modulated to carry information such as music or voice signals. Reach Power's system uses high-frequency RF waves to transmit hundreds of watts of pure energy. In a Reach WPT network, a sophisticated transmitter looks for devices equipped with receiving antennas or rectennas, and then a beam-forming router finds the safest and most efficient pathways to those antennas, avoiding people and utilizing reflections where possible. Reach can also transmit power across meshed networks of modular transmitters and receivers.

Over the last two years, Reach has partnered with In-Q-Tel, the U.S. Air Force Research Laboratory, and the Defense Advanced Research Projects Agency to demonstrate WPT transmitters, receivers, and relay nodes that could someday power sensors, drones, robots, automated guided vehicles, and other devices in the field. In 2024 the company showed that it could beam 256 watts of power across a mesh network to a 5-pound quadcopter drone in flight at NASA Ames Research Center, though at a distance of only 6 meters. The company's challenge now is to transmit higher amounts of power across greater distances with less energy loss, while operating within regulatory guidelines and minimizing the risks to humans as well as interference with other electronic systems.

Those are obstacles the company is hugely incentivized to overcome. The ability to deliver high power to multiple moving devices "is a holy grail for many hard-to-solve defense and intelligence community use cases," in the words of In-Q-Tel executive vice president George Hoyem. DCVC's Hutmacher predicts that "Reach's wireless power, especially if they're able to project it over longer distances, will be very well received in the U.S. government and eventually the commercial market."

And of course, everything Reach learns from working with the military and intelligence communities will help it advance products for wireless power transmission in factories, container terminals, warehouses, or retail locations. Working with the military, says Reach founder and CEO Chris Davlantes, "represents a significant step forward in the development of a true wireless energy web network."

> "Reach's wireless power, especially if they're able to project it over longer distances, will be very well received in the U.S. government and eventually the commercial market.
>
> **Clay Hutmacher**
> Operating partner, DCVC

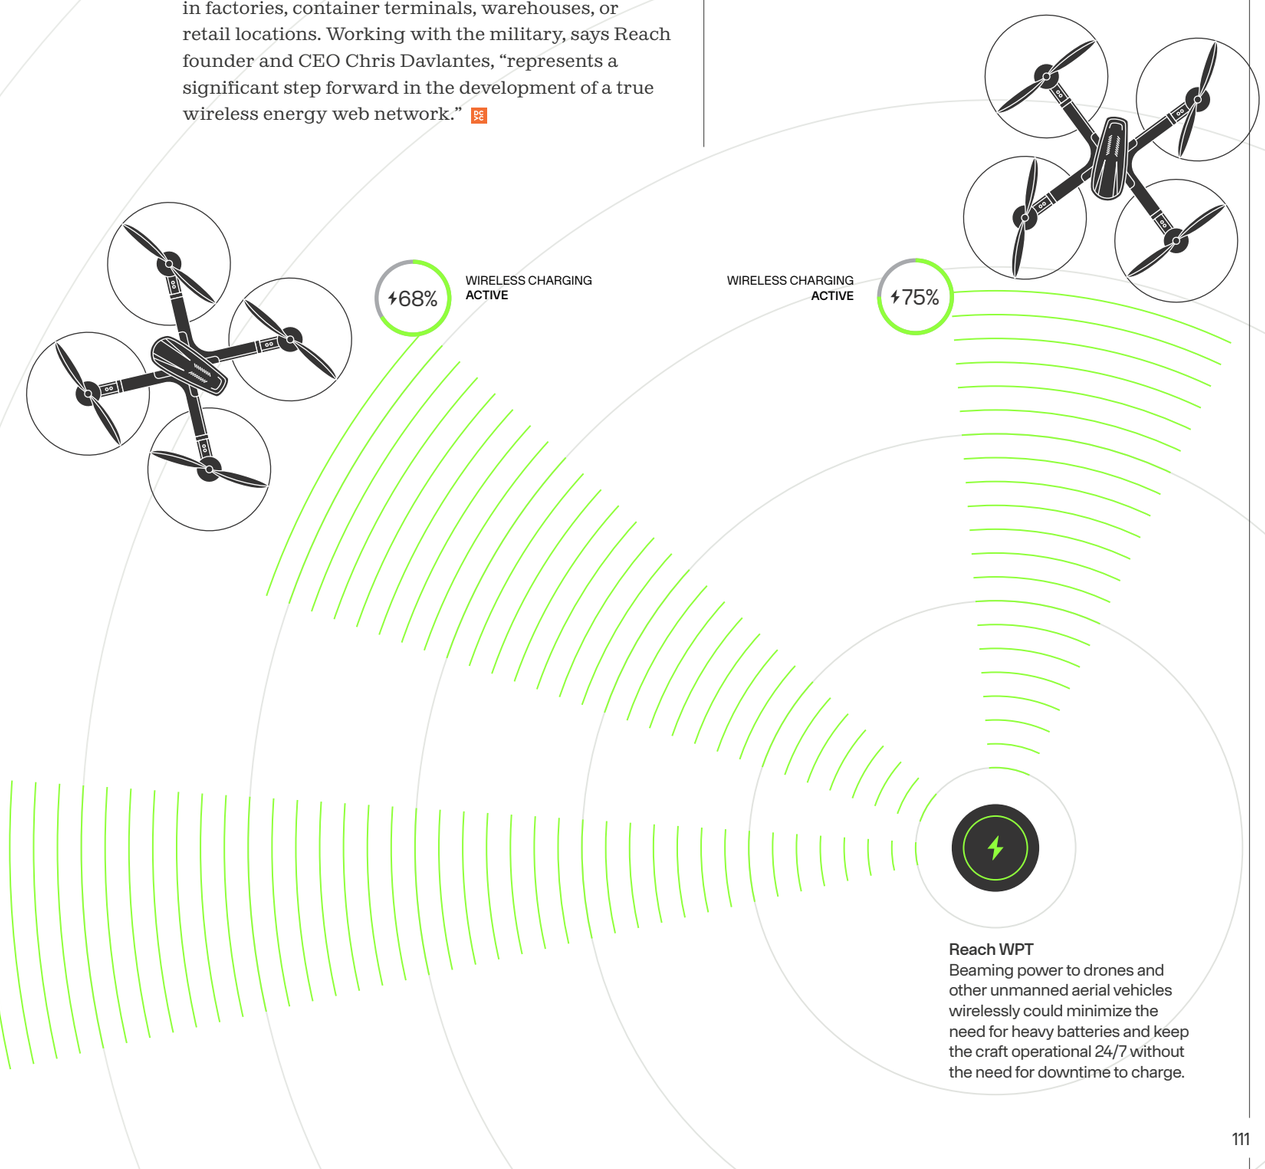

Reach WPT
Beaming power to drones and other unmanned aerial vehicles wirelessly could minimize the need for heavy batteries and keep the craft operational 24/7 without the need for downtime to charge.

→ **Fervo Energy** carries out a cross flow test, resulting in a rare release of steam from its geothermal wells. There's virtually unlimited heat energy below the Earth's surface, but geothermal projects have historically been limited to specific hotspots. Fervo is using advanced drilling techniques to change that.

Here at DCVC, we've been investing in deep tech since 2010. With every passing year, the point of deep tech has become clearer in our minds.

America and its allies need inexpensive, energy-efficient, environmentally friendly ways to build the things we need, power our economies, protect ourselves from security threats, keep everyone nourished, and extend healthy lifespans.

The power of computing, when brilliant entrepreneurs apply it to hard science and engineering problems, can help us realize all these goals. And at the same time, it can help bring factories and their supply chains back to the United States, where they can operate more cleanly and create high-paying jobs.

DCVC is deep tech venture capital.
We welcome your interest in our work.

Thank you to the following DCVC partners, principals, and associates

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

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 June 2025

©DCVC 2025

Deep Tech Opportunities Report

2025

270 University Avenue
Palo Alto, CA 94301

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

dcvc.com