The Economist

The new era of corporate scandal

Algeria kicks out its dictator

Why Europe's banks are so rotten

Great wheels of China

APRIL 6TH–12TH 2019

Redesigning life

The promise of synthetic biology

TQ

The Economist
Technology Quarterly:
Synthetic biology

→ April 6th 2019

3 Engineering DNA

5 Life, reprogrammed

7 The revolution will be automated

9 Intelligent design

11 Liberation biology

A whole new world

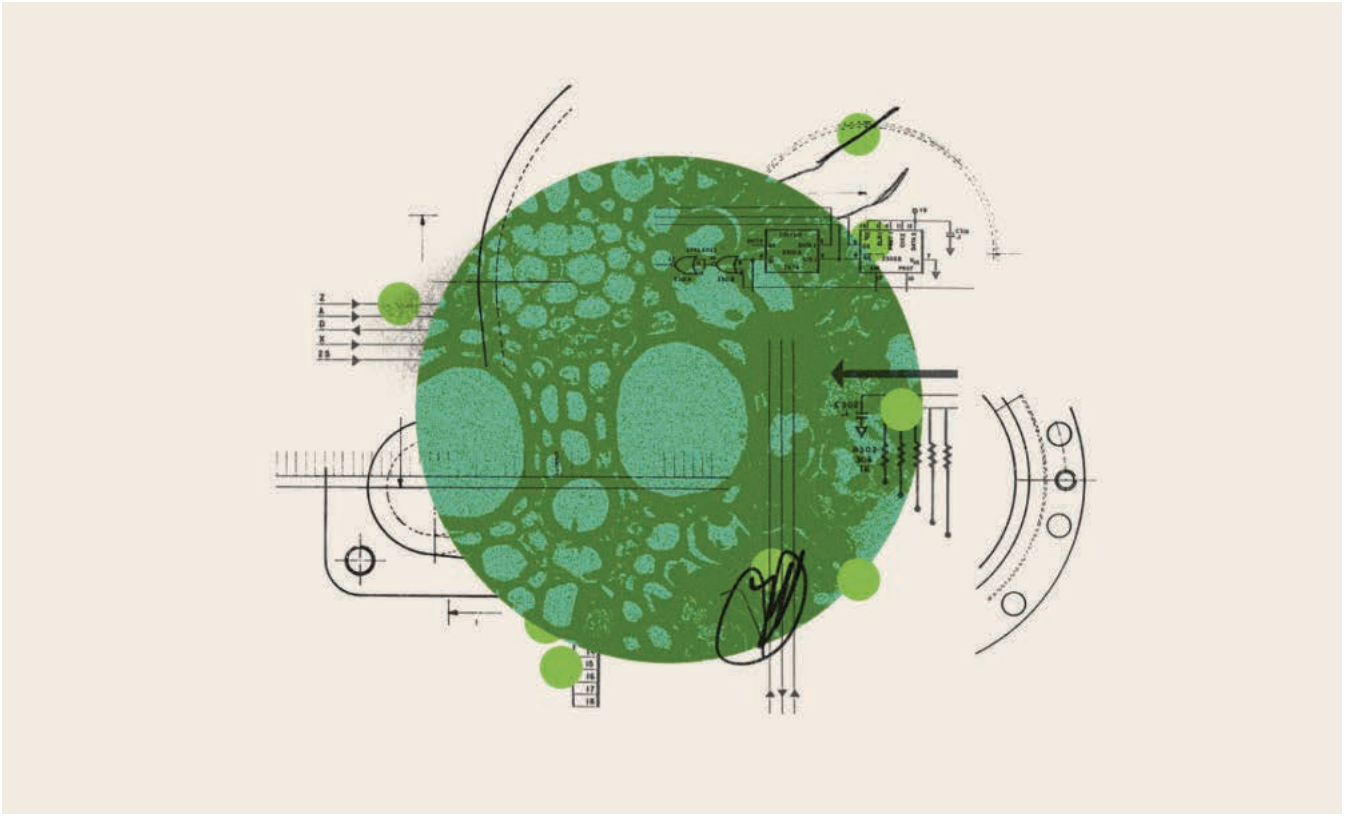

A whole new world

The engineering of living organisms is not yet changing everything. Give it time, says Oliver Morton

Biology is a way of structuring matter at a molecular scale by slotting each atom into its needful place. It is a way of controlling flows of energy on every scale from that of the smallest living cell to that of the whole living planet. It is a way of growing order and surprise in a universe that in all other respects tends towards entropic stagnation. And it is a thicket of limits on how long lives can last and how much life can accomplish.

It is also a way of packing 3,500 excited young people into the Hynes Convention Centre in Boston, Massachusetts. More than 300 teams from 42 countries took part in the annual International Genetically Engineered Machine (iGEM) competition there last October. By encouraging such teams to co-operate and compete in its Grand Jamboree, the iGEM foundation is hoping to create a framework for a synthetic-biology industry which combines molecular biology and engineering to achieve specific goals. Over the summer the young people went from an idea about something biological that might meet a human need, to designing new genes and seeing how well their ideas worked.

The projects presented to judges and peers at iGEM covered a remarkable range. There was an attempt to give bacteria a human sense of smell; there were fungi that could be used to build bases on Mars. The Great Bay team of Chinese high-school students won an award for synthesising the active ingredient of catnip in yeast and bacteria; they think it may help programmes which round up stray cats. Post graduates at the University of Marburg won an award for new genetic tools that will make a very fast-growing bacterium, *Vibrio natriegens*, easier for other engineers to use.

Not all iGEM projects pan out; one of the things teams learn is that, though engineering organisms is now possible, it is still difficult. Life can be very recalcitrant. Even so, there are already 32 startups around the world that began life as iGEM teams. Ginkgo Bioworks, a firm which grew out of teams from MIT that competed in 2004 and 2006, builds new organisms for clients in agriculture and the chemicals industry at an astonishing rate in its labs on the other side of Boston. It has attracted $429m of investment.

However, the Grand Jamboree is not primarily a route to riches. It is a celebration, and an exploration, of technology that will, in time, change the living world far beyond the test tube. Human engineering of the inanimate has produced a range of wonders from cities of towering glass to the fused sand that sits at the heart of computers. It is entirely plausible that engineering the animate could produce wonders as great and as various—and as unimaginable today as skyscrapers and silicon chips were 200 years ago.

Shining, shimmering, splendid

Humans have been turning biology to their own purposes for more than 10,000 years. They have reshaped crops and livestock through selective breeding and changed the structure of ecosystems by moving species around—most notably in the "Columbian exchange" that mixed together the fauna and flora of the New World and the Old. Having learned, in the 1950s, that genes were written on long molecules of DNA like stock prices on tickertape, by the

1970s scientists were able to start to move traits from organisms in which they evolved to organisms in which they could be useful by cutting and pasting pieces of that tickertape. That ability became the basis of the biotechnology industry.

The key enabling technology for synthetic biology is the ability to write new chemical messages on to fresh bits of tickertape, rather than just move nature's old messages from genome to genome. Machines capable of synthesising DNA letter by letter started to appear in the late 1980s. A decade later there were companies offering to write out almost any sequence of DNA letters you asked for and courier them straight to you. No longer limited by the genes they found in nature, biologists were able to get cells to work in whole new ways—to reprogram them.

That new ability underlay the three turn-of-the-century academic trends which came together to form synthetic biology. One of these was centred on engineers at MIT who had, from the 1960s to 1980s, pioneered the computer and internet revolutions. The plummeting price of DNA-sequencing technology—machines that have only to read, as opposed to write, the tickertape of life, and thus work a lot faster—made it obvious to them that biology, like computing, was based on digital code and capable of making progress at exponential speeds. From this they concluded that cells could, in principle, be engineered in the same way that circuits and software are.

Programming in nature is extremely convoluted, having evolved with no intention or guidance. And there is no helpful manual. But if you could synthesise genes that provided new, simpler ways of doing things, you would be liberated from having to understand the old ones. Life could be transformed into something more amenable to an engineering approach, with well-defined standardised parts. Tom Knight, one of the pioneers at MIT, and his colleagues saw in this sort of biological engineering something similarly world-changing to their work on the early internet and pre-PC computer workstations. And they found a generation of eager students whose first great "wow" moment in the cinema had been the re-engineered dinosaurs of "Jurassic Park".

The second ingredient that went into synthetic biology came from academics who were thinking along similar lines in the opposite direction; instead of trying to work round natural mechanisms they wanted to work towards recreating them. They were particularly interested in the systems by which cells turn genes on and off. Only when a gene is on, or "expressed", will a cell make the protein described by that gene's tickertape sequence. When it is turned off, or "repressed", the protein's production stops. Because proteins are the molecules that carry out almost all the tasks that go on in a cell, which genes are expressed when is fundamental to how cells work—and to how a brain cell, say, differs from a muscle cell, or a cancer cell from a healthy one.

Amid all this revolutionary talk, young companies in the field made a fateful decision to plunge into biofuels

In 2000 two teams published designs for novel genetic "circuits" with which they could control the expression of one gene with a protein made by another. In one of the gene circuits the carefully fashioned genetic switches flicked each other on and off over time. Genetic circuitry like this "repressilator" was child's play compared with the co-ordinated gene expression that evolution has programmed into leaves and eyes. But as one of the creators of the repressilator suggested, perhaps with Richard Dawkins's metaphor of evolution as a blind watchmaker in mind, "at this stage one can learn more by putting together a simple if inaccurate pendulum clock than one can by disassembling the finest Swiss timepiece."

The third ingredient was more practical: metabolic engineering. Life uses proteins called enzymes, which catalyse chemical reactions, to build all the other molecules it needs, with a different enzyme for each step of the construction. Sometimes the end product of such a metabolic pathway is something humans have a use for, such as a hormone, an antibiotic or a pesticide.

Being able to write DNA from scratch allowed metabolic engineers to bring together genes from a number of different organisms to build new pathways, thus offering the prospect of making molecules beyond the reach of chemistry for less than the cost of harvesting them from plants. The most striking project, led by Jay Keasling, a professor at the University of California, Berkeley, was a pathway which created a precursor to artemisinin, a molecule made by a plant called Chinese sweet wormwood that had been discovered to be a very good malaria drug. It was impossible to make the molecule by other means.

Unbelievable sights

As DNA synthesis became more widely available in the early 2000s, the various ways it could engineer new capabilities into organisms came together. By 2002 engineering undergraduates at MIT were using genes bought online to transform bacteria. In 2003 Dr Keasling and colleagues founded a company, Amyris, with an eye to making artemisinin and other useful stuff. The first international conference on synthetic biology took place at MIT the following year, a few months before the first iGEM Jamboree.

The media got wind of the excitement. It was not just that synthetic biology seemed like the sort of thing from which weapons could be made. Scientists playing God is always good copy, even if the creations were, as yet, mostly microbes. Rarely has science offered such a heady scent of Goddishness—with forbidden-fruit undertones of Frankenstein and Faust—and raised so many ethical dilemmas (see leader). Drew Endy, a charismatic young leader in the MIT group, talked of "reimplementing life in a manner of our choosing". George Church of Harvard talked of synthesising not just genes but whole genomes, including, perhaps, those of creatures now extinct. The mammoth might return.

Amid all this revolutionary talk, young companies in the field made a fateful decision to plunge into biofuels. It seemed a noble undertaking: biofuels could usher in the new technology of life while making good the damage done by the old technology of industrial machines. And governments were keen to subsidise them. But scaling up the pathways that produced hydrocarbons by the gram in the lab to the scale of millions of litres proved even more difficult than expected. The capital expense was huge. Worse, the oil price fell steeply. The projects failed.

That made investors very cautious about synthetic biology. But the field attracted a bit of support from some governments, such as those of Britain and Singapore. In America the Pentagon's far-out-ideas department, DARPA, which had taken an early interest,

Ratcheting up

United States, number of investment deals in synthetic biology, by type

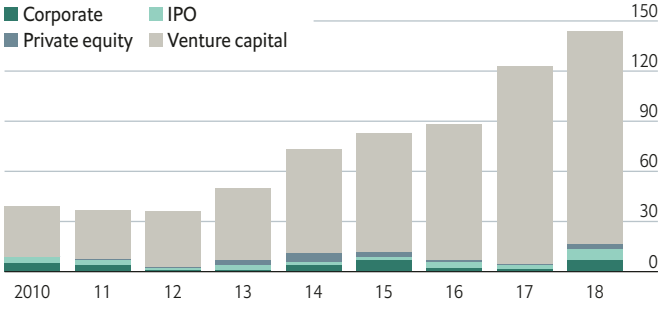

Sources: DCVC; Pitchbook; Silicon Valley Bank; SynBioBeta

created a new office of biology in 2013. Two years later it launched a programme that paid for leading laboratories in the field to put together pathways which could produce 1,000 molecules never created biologically before.

In January 2019 the 1,000th of those molecules was made. It seems an auspicious omen. In the past few years synthetic biology has shown signs of starting to live up to its promise. In part this is because of sustained academic effort and its cumulative gains; in part it is a matter of startup companies in the field finding their feet. But other factors are at play, too.

One was new gene-editing technologies—ways of doctoring existing tickertape a letter at a time. In 2000 there was none; now there is a whole range, and those based on a molecule called CRISPR have proved particularly powerful and easy to use (it was a big part of the Marburg team's victorious iGEM project). This has breathed new life into the idea of making precise changes to genomes, which is what synthetic biology is all about. It has opened new fields for biological research and new investors' wallets (see chart on previous page).

No one saw CRISPR coming. The falling cost of DNA synthesis, on the other hand, was widely foreseen. But it has still been a dramatic enabler. The price of a gene synthesised to order is about a thousandth of what it was in 2000; if you buy in bulk or have the technology in-house it can cost a lot less.

And then there is machine learning. Synthetic biology gets even greater benefits than most other industries from the recent growth in the capabilities of pattern-recognition programs. It is not just that laboratories produce reams of data with which to train such programs. In a paper in 2005 Dr Endy noted that "the designs of natural biological systems are not optimised by evolution for the purposes of human understanding." That is a problem for humans interpreting data and asking questions. For machines, though, understanding is as unnecessary as it is impossible. They just find patterns and uncover rules. This is not science as scientists understand it. But, if rigorously tested, such rules can still be a basis for engineering. There were perfectly good rules for building bridges long before there was a theory of gravity.

While synthetic biology has grown more capable, the promise of two older approaches to the improvement of life has diminished. One, the pharmaceutical industry, seems bound by Eroom's law (Moore's law backwards); the number of new drugs developed per billion dollars of R&D spending falls remorselessly. It was ten in 1970. It is well under one today, and still dropping.

This has excited interest in fundamentally new approaches to medicine. One is reprogramming cells to do helpful therapeutic things. Immune-system cells are the most obvious candidates. The cells of the microbiome—the interlinked bacterial ecosystems that thrive on skin and in guts—are another possibility.

A dazzling place I never knew

The second ailing improver of life is the petrochemical industry. Synthetic biology's push into biofuels was not fundamentally misguided; fossil hydrocarbons do have to be replaced. The mistake was rushing into a bulk market with low margins: petrol. Some companies are now using synthetic biology to replace more upmarket molecules from the same crude oil which end up in fragrances and food additives with far more added value. Others are looking at making plastics environmentally friendly. As their technologies prove themselves at increasing scales, and as their technical prowess allows them to expand their repertoire to cheaper bulk products, these efforts could eat the petrochemical industry from within like some world-saving parasitic wasp.

Synthetic-biology executives say their worry is not money, but focus and time. Every firm has more revolutionary-looking projects than it can pursue. And no one knows how long it will take the projects to pay off. As the gnomic aphorism at the end of Mr Endy's emails has it: "Our victory inevitable. Our timing uncertain". ■

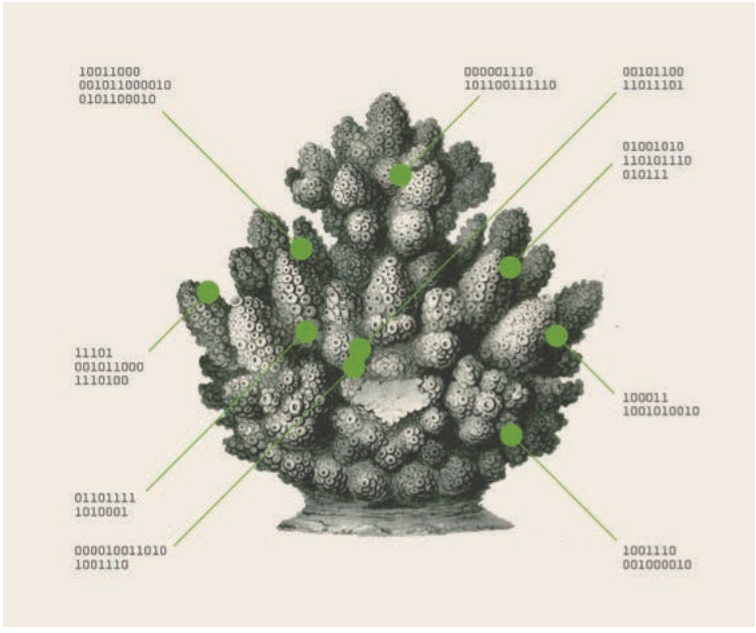

How it works

Reprogramming life

Synthetic biology differs from everything that has come before

TO UNDERSTAND BIOLOGY, synthetic or otherwise, you have to understand how proteins are made. Proteins, which carry out almost all the basic functions of life, from respiration to reproduction, are all made of 20 smaller molecules strung together into a chain. The shapes those chains fold up into in order to fulfil different functions are complex and incredibly hard to predict. But they are almost all entirely determined by the order of these smaller molecules, which are called amino acids.

The gene for a given protein is simply the list, in order, of the amino acids needed to make it. This information is written down in the genome as a sequence of DNA bases—A, C, T and G, the letters on the ticker-tape—in the same way data in a computer are stored as a string of 1s and 0s. The program that turns these DNA sequences into sequences of amino acids is the genetic code. It assigns a fixed amino-acid meaning to each of the 64 different three-letter words (such as ACT or CTG), known as "codons", that can be made using DNA's four bases.

Imagine a codebook with codons in one column and the names of the amino acids in another. To decode a gene, look up its codons one by one and write down their amino-acid meanings. It is a simple, rule-based undertaking—an algorithm. The cell carries out the same algorithm. But instead of a code book which matches codons to amino acids, it uses codon-recognising, amino-acid-carrying molecules called tRNAs and a mechanism called a ribosome which provides a place for those tRNAs to interact with a copy of the gene. The act of reading the gene, codon by codon, is the act of creating the protein, amino acid by amino acid (see diagram overleaf).

When it executes an algorithm this way, biology looks like computer science. But it is important to appreciate that biology does not deal with information the way humans do. In human programs, the logic and the machinery that acts on it are kept separate. Computer users can change a program in blithe ignorance of the

From code to protein

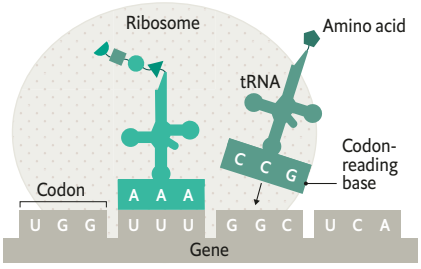

Codons on the gene in the ribosome are recognised by tRNAs carrying amino acids

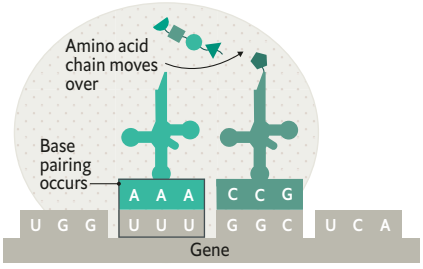

The growing amino acid chain is stuck to the latest amino-acid, specified by the codon next in line

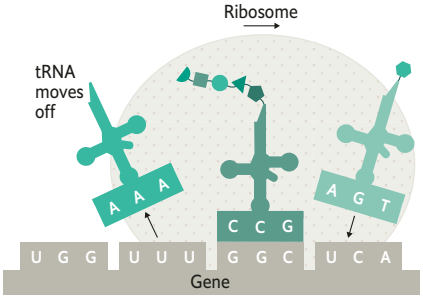

The ribosome moves along the gene, so that the next codon's amino acid can be put in place

physical principles and peculiarities built into the hardware that it runs on. But evolution cares nothing for such distinctions. All its processing is just a matter of molecules interacting—the way that tRNAs stick to codons as if to velcro, the way the shape of the ribosome forces amino-acids together, and so on. From the simple gene-to-protein translation of the ribosome to the extraordinary synchronised symphony which turns a fertilised egg into a whole human, biological information and its implementation are all but inseparable. Life runs not on software and hardware, but in allware. That makes it highly resistant to human reprogramming.

It can, though, be hacked. From the 51 amino acids of human insulin, which in 1978 became the first product made by the first biotech company, Genentech, to artificial antibodies containing more than a thousand of the things, biotechnology consists almost entirely of getting cells to produce proteins they would not normally make by cutting a gene out of one organism and dumping it, often unceremoniously, into another.

Most of these proteins have been natural ones. Nature is well stocked with proteins that do useful things—regulate blood sugar, kill pests or break down grime on laundry. Putting the genes for such proteins into the genomes of bacteria that will then secrete insulin, or of crops that need pest resistance, or fermentation tanks churning out supplements for detergent, was an obvious moneymaker. But the preference for the natural was, until recently, also driven by necessity. Designing a protein from scratch was impossibly hard. So was getting suites of proteins from different sources to work together.

That is no longer true. Protein design and DNA synthesis now make it possible to produce proteins that, separately or together, do things nature does not. They remain imperfect. But because DNA sequences are cheap it is possible to try out lots of variations to see which actually work.

Thus, for example, there are companies working on new metabolic pathways which combine enzymes freshly discovered through the sequencing of ever more genomes, enzymes long understood and enzymes significantly re-engineered. It is an exacting craft, or art; it requires not just finding the right enzymes but also bringing about the carefully balanced levels of gene expression needed if a dozen or more of them are to work together, not to mention tweaking the underlying metabolism to prevent things produced by the new pathway from disrupting those already there. But if the work is done well, it seems now to be the case that more or less any small molecule found in nature can be made by yeast or bacteria in a fermentation tank.

Two particularly interesting possibilities are the cannabinoids made by marijuana and the variations on opium and morphine made by poppies. Cannabinoids come in a remarkably wide number of forms, some psychoactive, some not, some therapeutic, some not, many legal for some purposes in some jurisdictions, many illegal for all purposes elsewhere. A set of cannabinoid-synthesising pathways described by Dr Keasling and colleagues this February offers therapeutic and recreational possibilities along these lines which will be explored by a new company called Demetrix. A hugely ambitious 20-protein pathway capable of producing morphine and its relatives, developed by a former student of Dr Keasling's, Christina Smolke, offers perhaps more profound possibilities. Dr Smolke has founded a company, Antheia, which aims to use her new know-how to make opiates that are cheaper and so more accessible to the tens of millions around the world unable to get pain relief, and also to make opiates that are less addictive.

Protein design and DNA synthesis make it possible to produce proteins that do things nature does not

Breaking the code

A more radical possibility, at least in terms of chemistry, than remaking and improving natural compounds is to create enzymes to catalyse chemical reactions nature never carries out. Take the task of sticking a carbon atom to a silicon atom. Human chemists are pretty good at this, and the organo-silicon compounds they thus create are used in electronics, pharmaceuticals, building materials, breast implants and more. Nature, though, does not use carbon-silicon bonds, and so no natural enzymes make them.

In 2016 Frances Arnold, of Caltech, corrected nature's deficit, using evolution to create an enzyme which stuck silicon to carbon and opened up a whole new realm of chemistry to biology. She now guides her directed-evolution technique, which won her a Nobel prize in 2018, with machine learning, the better to alleviate the watchmaker's blindness. She believes that synthetic biology can in principle create enzymes for most of the reactions today's chemists bring about with rare catalysts, high temperatures and pressures, or environmentally unfriendly solvents.

As well as making new proteins, it is also possible to make new RNAs. This is how CRISPR gene-editing works. A molecule of RNA is created that velcroes itself to a specific sequence in the genome; a companion protein then slices through the bit of DNA thus highlighted. Once the DNA is broken, a new gene, or gene fragment, can be inserted into the gap. If you put a gene describing the CRISPR RNA and its protein into a cell in such a way that it gets expressed only under certain conditions, you have a cell whose genome can be reprogammed by remote control.

If you write an organism's genome from scratch you can make it easier to mess around with in a number of ways. A coalition of ten laboratories around the world is currently rewriting the entire genome of *Saccharomyces cerevisiae*, brewer's yeast, in order to make it an even better test bed for genetic research than it already is. To this end they are carefully stitching together the most appropriate versions of over 6,000 genes as well as most of the sometimes vital

gubbins found between them—over 12m bases of DNA in all. One of the things the project is writing into the genome is a system that will make it cut itself up and reshuffle its genes when told to. This technology should provide a powerful new tool for the study of evolution, says Tom Ellis of Imperial College, London.

A deeper way in which what is known as "Sc2.0" differs from *Saccharomyces cerevisiae* proper is that it operates with a slightly different genetic code. Three of DNA's 64 codons describe not an amino acid but an action: specifically, "stop". These three codons—TAG, TAA and TGA—tell the ribosome and its tRNAs: "This is the end of the gene. Add no more amino acids, we're done with this one." In the re-engineered yeast, though, only two of these three stop codons are used. Wherever the natural, baseline yeast genome marks the end of a protein-coding sequence with a TAG codon, the scientists writing Sc2.0 use one of the other stop codons, TAA or TGA. This means that in Sc2.0 TAG means nothing—and so can be made to mean something new.

Nature uses 20 amino acids in its proteins. But there are hundreds of others that could be used, some of which would confer interesting new properties. In Sc2.0 it will be possible to make the TAG codon "mean" one of these other amino acids by designing a new tRNA molecule that recognises the codon and new enzymes to stick an amino acid to that molecule. Cells thus equipped will be able to use an amino acid no natural cell has ever used before.

Nor does the process have to stop there. The genetic code uses 61 codons to code for just 20 amino acids; in some cases there are six codon "synonyms" for a given amino acid. Writing an organism's DNA in a form missing particular synonyms is a compositional task similar to choosing to avoid using a common linguistic symbol, such as "e", in a short bit of writing; the upshot may look slightly ungainly, but you can do it. Rewrite the code with fewer synonyms, and you have more codons to devote to non-canonical amino acids. One therapeutic option this might open up is drugs that bacterial defences cannot cope with. Bacteria have evolved to counter everyday proteins; put in amino acids they have never seen before and some of those defences no longer work.

Bespoke genetic codes have attractions beyond a larger vocabulary. It is the universality of today's genetic code that allows viruses to force the cells which they attack to do their bidding, making their viral proteins from their viral genes. A genome that uses a different genetic code would be impregnable to such attack; the virus's genes would no longer describe the proteins it needs. Recoding could thus make cells immune to any viral attack; indeed, there is already work on achieving this in bacteria.

If it works, this sort of recoding could be very helpful to existing biotechnology. Fermentation tanks that never get wiped out by infections and antibody-producing cell lines that could not harbour viruses would be a great boon. It is possible to imagine changes in the way codons code for amino acids so radical that parts of synthetic biology become a separate creation, parallel biospheres based on the original but no longer in contact with it, populated by creatures which neither infect nor are infected, that are linked to the rest of life only through the intentionality of design.

A hint of such strangeness could be seen in a paper published in *Science*, a journal, this January by Stephen Benner of the Foundation for Applied Molecular Evolution in Florida and his colleagues. They have created double helices in which the existing bases, A, T, C and G, are supplemented by Z, P, S and B. This *hachimoji* ("eight letters") DNA offers much denser data storage than evolution has had at its disposal for the past 4bn years. With eight letters to play with, for example, you could recode the genome to use doublets, rather than triplets, as codons, if you redesigned the ribosome, the tRNAs and a bunch of other stuff, too.

Would anyone want to? The potential of the existing code is enormous, the range of proteins it can, in principle, describe is barely yet explored; there might seem to be no need for such showing off. At the same time, engineers do like to tinker. ■

Automation

An industrial revolution

Remaking life means automating biology

Zach Serber worked at Amyris, a synthetic-biology pioneer, when the company was trying to crack the biofuel market. Seeing brilliant metabolic engineering fail to make a business led him and his co-founders at Zymergen, a company based in Emeryville, California, to take their new company in a different direction. They would not try to manufacture or sell things. They would offer their synthetic biology as a way of making businesses already using biotechnology more profitable. This is, at the moment, the model used by a number of leading synthetic-biology companies. At its heart is the automation of experiment.

Biotechnology is already a bigger business than many people realise. Rob Carlson of Bioeconomy Capital, an investment company, calculates that money made from creatures which have been genetically engineered accounted for about 2% of American GDP in 2017. The contribution was split between three industries. Pharmaceuticals and crops, contributing $137bn and $104bn respectively, are the ones that the public knows about. The third sector, industrial biotechnology, is much less visible but even more lucrative, worth $147bn or more (see chart). Chemicals used for many purposes—raw materials for plastics, food additives, some fragrances and biofuels—are already being churned out at scale by altered micro-organisms in fermentation tanks.

As well as being the biggest biotechnology market, this is also the one best suited to companies seeking to offer innovation as a service. Testing drugs and genetically modified crops is a long and costly business. Replacing one strain of industrial yeast with a better one can be done in a week. Industrial customers tend to know what they want and synthetic biology promises a lot of value. Tim Fell, the boss of Synthace, a synthetic-biology software company in London, says that in one project the company engineered a 200-fold increase in the rate at which bacteria produced something useful (he cannot say what) in just four weeks.

About 75% of Zymergen's business, according to Dr Serber, is helping companies re-engineer, for industrial purposes, microbes they are already using, in order to increase production, reduce costs or both. The company is built around machine-learning programs that suggest changes to the genome which could produce an organism and setting—temperature, nutrient balance, and so on—that improves on the status quo. In this fiercely empirical process

Growth industry

United States, estimated biotechnology revenues, $bn

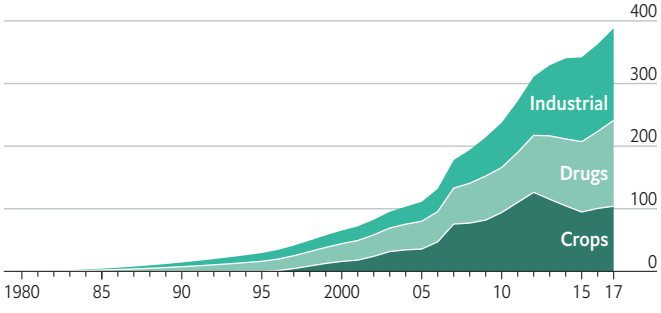

Source: Bioeconomy Capital

Zymergen makes DNA tweaks of all sorts, most of them to sequences that regulate gene expression. These tweaks, says Dr Serber, have helped customers for its "molecular technology" make better margins on hundreds of thousands of tonnes of product.

Arzeda, based in the Interbay district of Seattle, has a similar business model and is based on similar technologies. But where Zymergen concentrates on empirically derived ways to improve productivity, the expertise of Arzeda's machine-learning systems and scientists is in applying a theoretical understanding of how the shape into which proteins fold determines their function, thus making them better at what they do, or able to do something new. It brands itself "the protein design company".

Ginkgo, the iGEM-born startup in Boston, is another variation on the business-to-business theme. Its focus is not on the specifics of genome-based machine-learning or protein design, though it does both, so much as on developing a broader expertise in the remaking of microbes. It calls itself "the organism company".

Their use of massive amounts of synthesised DNA is producing a new way of doing biology on an industrial scale

Means of production

The three companies may differ in details of their approaches, but the big picture unites them. All of them see their current business-to-business approach as a stepping stone, a way of honing their techniques, teaching their machine-learning programs and bringing in cash as they develop products of their own. Arzeda talks of making tulipin, which among other things can greatly improve the qualities of perspex. That improvement is not so great as to justify harvesting it from its native tulips, but Arveda's proteins mean you do not need to. Ginkgo is spinning out joint ventures with clients to work in specific areas. In 2018 it created a business with Bayer, a chemical company, to develop microbes which would make fertiliser inside a plant's root system. It has another spin-out working on cannabis, and has just announced a third one developing plant proteins for use in vegetarian foods, including meat substitutes. Zymergen is looking at materials for electronics.

They are also united in their zeal for high-throughput experiments. Their use of massive amounts of synthesised DNA is producing a new way of doing biology on an industrial scale

During the five years that Jason Kelly, Ginkgo's chief executive, spent in Dr Endy's lab at MIT in the 2000s he reckons he may have ordered 50,000 bases of commercially synthesised DNA—a pretty profligate amount at the time. Today Ginkgo orders synthetic DNA sequences at 50,000 times that rate, using them to change the genomes of thousands of organisms a day. In 2017 it bought a DNA-synthesis company, Gen9, bringing all its production capacity in-house. That has not sated its appetite. It has a contract with Twist Bioscience, the world's biggest DNA-synthesis company, for a billion base pairs over the coming years.

Arzeda is smaller, but Alexandre Zanghellini, its boss, says it still manages to order around 10,000 new DNA sequences a week, each of which is then put into a particular microbe so that its computers' assumptions about how changes in the sequence change the function of proteins can be tested. Often these DNA sequences are not even looked at by humans before they arrive by courier.

Drinking from such a firehose of DNA increasingly requires experiments designed and managed by computers. Ginkgo spent years programming computers to supervise experiments and robots to carry them out and then finding and removing the innumerable bugs with which those programs were afflicted. For ten years, according to Dr Kelly, doing lab work using the partially automated foundries thus created was considerably slower for the company's designers than doing it themselves would have been. But having to use the automated systems meant having to improve them. A couple of years ago, Dr Kelly says, Ginkgo reached a point where its foundries were as productive, in terms of person hours for work done, as an expert researcher. Now he pegs them as ten times more productive, and says the margin is growing.

Automation increases not just the amount of research that can be done, but also its complexity. Much biological experimentation takes place in trays of 96 "microwells", or miniature test tubes. Humans tend to design experiments using these wells quite simply; do A to one subset, B to another, and so on. A computer can design experimental strategies that are much more complex, picking a wider range of hypotheses to test, and then testing many more hypotheses per tray. For properly programmed robot flunkies, the most recondite experimental schemes are a doddle. According to Markus Gershater, the chief scientific officer at Synthace, the gains software and automation offer experimental design can be just as important as gains in speed and throughput.

The role of machine learning in these labs means they have an enormous appetite for data. Most biology labs do without mass spectrometers, analytic tools which rapidly sort through samples molecule by molecule and characterise every one of them. They are expensive and produce more data than most people need. Synthetic-biology companies love them.

More data offer computers a clearer idea of what is going on; they also show what is going wrong. Most biologists at the bench have a sense that the living material they work with is not really to be trusted. Biology, they say, unlike physics, is unreliable. The "noise" in experiments can often swamp the signal you are looking for. Getting an experiment to work pretty regularly is good enough. In part this may be true. But it is hard not to think that much of the unreliability is with the biologists, not the biology. How else to explain why studies repeatedly find that many results reported in research papers cannot be replicated in other labs?

The problem is not just human error. It is also human nescience. There are things going on in a lab that experimenters do

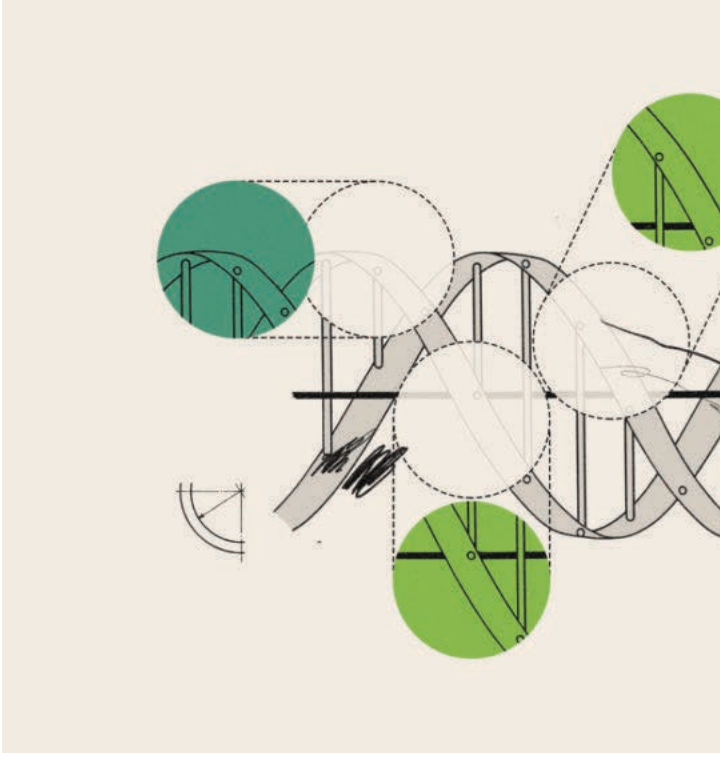

not notice, but the creatures they experiment on do. The more data a system gathers, the more can be discovered, if necessary, about what actually happened, and that will surely help replicability.

Industrialisation helps in other ways, too. One piece of kit popular in labs that can afford it is the Echo 655, built by Labcyte. Like a pipetting system it moves drops of fluid from one set of wells to another. But by creating those drops with tightly focused ultrasound rather than suction it does so in much smaller amounts, much more accurately and with no contaminating contact. Smaller wells—up to 1,536 on a tray—mean more throughput and less spent on the chemical reagents the experiments use up. At the far end of this trend towards the tiny and precise is a system made by a startup called Berkeley Lights which has wells which contain but a single cell, manipulated entirely with laser beams.

So powerful is this new automation that it has drawn Synthace away from its original intention of making organisms to providing software as a service instead. The company has developed a computing environment called Antha, where researchers can say what they want done in relatively high-level terms, confident that machines will optimise the experiment's design for the client's instruments and tell the instruments what to do.

A startup called Transcriptic wants to go even further, operating "labs in the cloud" where an experimenter at a terminal anywhere in the world can get a set of experiments done in an automated facility they never even see. Mr Kelly thinks that, at least for the sort of work Ginkgo does, the time is not yet right for such radical approaches—having the people designing the organisms and the foundries that make them under one roof matters a lot. But it is clear that the trend to automation is not yet played out.

From the ratcheting of the ribosome on up, there is something mechanical about life. In foundries like Ginkgo's it is hard to avoid the sense of that mechanistic model moving out from the cells embodying it and into the sparsely inhabited systems studying, manipulating and redesigning them. There is an uncanny feedback loop between the machineries of cell and laboratory which is eating away at the gap in between them. ■

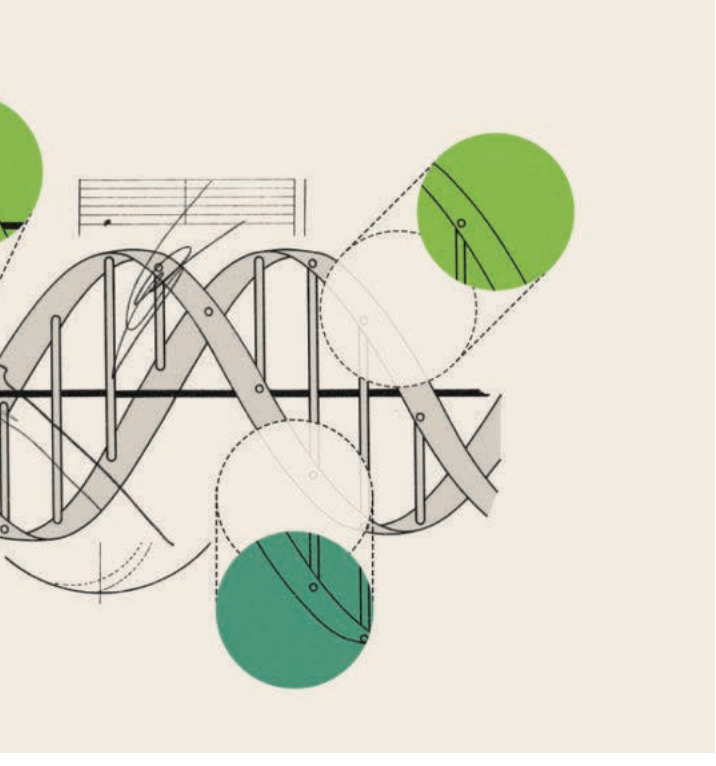

Applications

The new stuff

The uses of intelligent design

Life on earth uses perhaps 5m different proteins. It is by no means clear what each of them does. In even the simplest bacterium there are proteins with jobs that scientists cannot identify—but which the bacterium clearly considers vital, since without the genes for those proteins it dies.

But if many specifics are still hazy, the cumulative capability of all the things that natural proteins can do is well known: it is the living world. All the chemical and physical cleverness that life is capable of, from dandelion seeds to coral reefs, jellyfish to brains, is there because proteins did stuff.

Proteins create the materials of wood and leaf, flesh and bone. A couple of blocks away from Zymergen a startup called Bolt Threads supplies the rag trade, and its own clothing subsidiary, with threads made of proteins from spider silk, and leather from fungal mycelia. As well as being able to explore new physical properties for such materials and, in principle, make them more cheaply, it is also able to offer them to people who object to having silk worms boiled and cattle skinned for their finery. Stella McCartney, a designer, is working with the firm to turn its materials into vegan-friendly fashions.

Many companies are developing products that seek to mimic the taste or texture of meat. Impossible Foods, based across San Francisco Bay in Redwood City, relies on engineered microbes for bulk supplies of the leghemoglobin protein, normally found in the roots of some plants, that makes its completely plant-based "impossible burgers" bloody without the blood. Others, including Ginkgo, are working on similar products. Meat without livestock could, in principle, be a very climate-friendly technology. But for consumers who prize the environment and yet distrust genetic engineering, the technology may raise concerns.

Some investors worry, too. Vijay Pande, who runs a biological-engineering fund at Andreesen Horowitz, a venture-capital firm, sees in some plant-based foods echoes of the rush into biofuels that blighted synthetic biology's early days. He hears of company founders who are basically trying to make mince, he says, which is a cheap product with which to compete. To earn money they will have to make a lot of hamburgers. That said, at least one company is working on synthetic *foie gras*, which might reduce the suffering of geese and sell at a luxurious price.

With food and fabric already covered, and many of the existing 5m proteins still unexplored, it may be hard to imagine why anyone should want more proteins. But there is much more to investigate. Consider the class of all proteins containing 66 amino acids. Because there are 20 different amino-acid possibilities for each of the links in the chain, there are in principle 20^{66} such proteins. That is roughly the same as the number of subatomic particles in the visible universe. And a 66-amino-acid protein is a tiddler.

It is into the cosmically vast sea of proteins nature has never made that David Baker of the University of Washington has set his course. In the 2000s Dr Baker was a world leader in the field of predicting what the structure of a natural protein would be on the basis of the order of its amino acids. This is a fiendishly difficult problem; the way the chains fold up is subject to incredibly subtle chemical forces that have very large effects. But the Baker lab got pretty good at it—good enough to spin out Arzeda, the protein-design company. Five years ago Dr Baker decided it was time to use what he had learned not to understand old proteins, but to design

absolutely new ones in shapes of his choosing which nature has never explored.

One use for such shapes might be to encourage particular types of crystal lattice to grow. Some bacteria which live on plants have evolved proteins which produce a lattice similar enough to that in ice crystals to "seed" the growth of ice, making frost more likely. (Ski resorts now use such proteins to help make snow.) Proteins designed to seed the growth of other crystals could help in the formation of things more interesting than frost, like the atomic lattices of semiconductors used for computers. Another use might be to build molecular motors. Alexis Courbet in Dr Baker's lab has created a protein wheel which can spin on an axle. There are already markets for such tiny pieces of machinery; $20bn a year is spent on the micro-electro-mechanical devices used in things such as mobile-phone motion sensors, car components and switches for optical circuits. Protein-based mechanisms could allow far higher levels of complexity, though probably not soon.

Perhaps the most striking recent design in Dr Baker's lab, though, is a set of proteins created by Zibo Chen. Like DNA molecules in their double helices, these proteins can stick to each other by means of molecular velcro (which is actually, as it is in DNA and RNA, a process called hydrogen bonding). But, also like DNA molecules, they will stick together only if the velcro on the two molecules is complementary. In DNA complementarity is about the sequence of bases. In proteins it is there in the shapes Dr Chen has given his proteins. He has made a family of 64 proteins each of which stick only to one other, making 32 pairs.

Such designs could be used to make a protein's function conditional—for example by designing a protein which cannot do its job unless an extra module is velcroed to it. Conditionality like this might provide a way to reprogram cells that does not need to engage with the mechanisms nature uses and so can be designed with the clear certainty of human software and hardware design. Reprogrammed cells are already influencing cancer therapy.

The most striking recent development in cancer treatment re-engineers the immune system's T-cells—cells which prowl through the body looking for proteins they don't like on the surfaces of other cells. What is called CAR-T therapy starts with the gene for a chimeric antigen receptor (CAR), a protein which sits on the surface of a cell. It is possible to tailor this gene to decide what the protein it describes will recognise; for therapy the chosen target will be a protein specific to the patient's cancer cells. Doctors take T-cells from the body, equip them with the gene for the cancer-recognising receptor, and then put them back in. When the CAR protein recognises a cancer cell from its telltale protein, it orders the T-cell to kill it.

One of the problems that has dogged the development of CAR-T therapies is that the CARs' recognition of cancers is not perfect, which means the T-cells can attack innocent bystander cells, causing severe, sometimes lethal, side-effects in some patients. This is where a system developed by Wendell Lim and his colleagues at the University of California, San Francisco, comes in. They improved the reliability of CAR-T cells using a synthetic version of a protein called Notch, which, like the CARs, generates a signal inside the cell when it recognises a protein outside it.

Dr Lim and his team built a simple two-gene circuit (see diagram overleaf). One gene causes the cells to produce a Notch protein that recognises a specific molecule on the surface of cancer cells. The other produces a CAR which recognises a second such telltale. But the CAR gene will produce its protein only if it is switched on—and the signal that switches it on comes via the Notch protein. This means that the cell goes on its programmed rampage only if both the telltale signs are seen; the first activates the Notch which produces the CAR that recognises the second. Electrical engineers will recognise this as an AND gate: you need both of two inputs to be present if you want the output.

Getting this seemingly straightforward system to work reliably in cells was hard. But once it was done, the value was obvious. In late 2017 the company Dr Lim had founded two years earlier to work on the problem, Cell Design Labs, was bought by Gilead, a therapeutics company, in a deal worth $567m—more than 16 times what had been invested in Cell Design Labs up to that point.

Reprogrammed cells are already influencing cancer therapy

Dr Lim is now investigating the possibility of slightly more complex circuits, for example one that requires three simultaneous stimuli (in circuit speak, two AND gates) or one stimulus and either one of two others, but not both (an AND and an OR). Even very simple circuits might make cells into much more discriminating therapists. Cancer is not the only application. Diseases of the immune system might be treatable by cells taken out of the body, reprogrammed and put back in. "Regenerative therapies" which use stem cells—cells that can give rise to a number of different types of cell—to repair damaged tissue and organs might also benefit from programming which would tell the cells when and where to do their stuff.

Michael Elowitz of Caltech, one of the inventors of the "repressilator" that was one of the first ever artificial genetic circuits, imagines a more radical system that needs no T-cells. Imagine putting the short-lived copies of the genes for a small genetic circuit into every cell in the body. One protein described would be lethal to the cell. A second protein might velcro itself to this first protein, perhaps using some of Dr Chen's highly specific hydrogen bonding, in a way that usually stops the first protein from doing anything awful. But under some particular circumstances—if, for example, the cell were making a protein typical of cancer—the velcro would not stick, and the lethal protein would become active, killing the cell.

A circuit that simple would be ludicrously dangerous; it would be like having a gun pointed at every cell in the body. But safeguards could be added, making the lethality conditional on more

than one factor—just as the Notch system provides an extra level of control for CAR-T therapy for cancer. Dr Elowitz thinks it likely there is a whole new field of medicine to be built from such systems, one that cures not through small drugs that get everywhere, or more sophisticated proteins, such as antibodies, that target specific cells, but through cells that become medicines, or surgeons, themselves.

New molecules developed by synthetic biology can in principle be turned to all the purposes—food, fabric, medicine, recreation, even, if applied to wood or coral, shelter—for which humans use the non-synthetic kind. But as Dr Courbet's little nanomotors suggest, they might do even more. One of the most impressive possibilities is to use them to deal with the global glut of information.

The world currently produces many exabytes (billions of gigabytes) of data a day, and it could produce a great deal more. One estimate suggests that driverless cars may produce 4,000 gigabytes each every day. Those are data that could be learned from, or used forensically after mishaps.

But storage is an issue. Storing a day's worth of the world's data using the most high-density storage medium in current use would require enough very expensive magnetic tape to cover dozens of basketball courts. Alternatively, you could store all those trillions of ones and zeros in just 20 grams of DNA. You could put everything that happened in a century into a single warehouse and expect it to last fundamentally uncorrupted for thousands of years. No other memory system comes close. And that is before you start to factor in adding *hachimoji* bases to increase the density.

Get it all onto one disc

Last year Dr Carlson of BioEconomy worked with Microsoft on a project that showed how data could be coded into DNA and retrieved from it using a scheme like that employed by memory chips. Other researchers have suggested that some simple forms of data processing could even be carried out on data while they are stored in DNA form.

The problem, Dr Carlson points out, is that DNA synthesis currently costs a million times too much for this to be an affordable way of storing data. But being a million times too expensive is not necessarily the hurdle it used to be. Machines that sequence DNA got much more than a million times cheaper between 1995 and 2015; sequencing is now so cheap that in 2020, two decades after the first announcement of a human genome being sequenced, people at the J. Craig Venter Institute in San Diego talk of sequencing a million of them. And big computer companies have a record of getting technologies to improve exponentially for as long as that improvement is physically possible: that is how they managed to live by Moore's law for so long. It is not a coincidence that, after Ginkgo and their billion-base contract, the second-biggest customer at Twist, a leading DNA-synthesis firm, is Microsoft. ■

The on switch

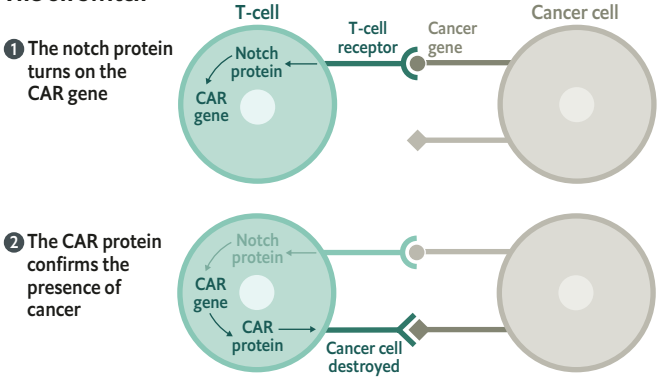

The future

Liberation biology

What is life for?

THOSE GIVEN to grand statements about the future often proclaim this to be the century of biology in the same way that the 20th century was that of physics and the 19th was the century of chemistry. Synthetic biology's potential provides a basis for such boosterism; life reprogrammed to produce useful new products, take new forms and act in helpful ways.

Honouring that promise will not be easy. Understanding biology's capacity to process information, and thus control itself, is a much more challenging puzzle than mastering the parcelled world of software and hardware. Taking years to create a working AND gate is therapeutically very promising. But it is a very small step on the way to controlling life as a coder controls a computer.

Still, the fact that synthetic biology recapitulates some early aspects of the computer revolution should not lead people to ignore crucial differences. One is that those who created the modern world of computers did not have powerful computers to help them in their task. Today's synthetic biologists do. Their work builds on, and grows out of, the computer revolution, and this may speed it up a lot. Frances Arnold of Caltech compares life's programming to a symphony composed by evolution, and today's biology by design as being roughly at the level of learning to hold the composer's pencil. That is why she likes to harness evolution to remake things, rather than design from scratch. But it is evolution that is guided with machine learning and directed by human creativity to write songs that humans want.

People have imagined such things before. In the early 20th century scientists and writers inspired by the new power of genetics described "biotopias" eerily reminiscent of the dreams of synthetic biology. In H.G. Wells's "Men Like Gods" (1923) plants "had been trained and bred to make new and unprecedented secretions, waxes, gums, essential oils and the like, of the most desirable quality", which could serve as a slightly flowery mission statement for half the companies in this TQ. In Charlotte Perkins Gilman's "Herland" (1915), a race of parthenogenetic women live in a cornucopian Eden they have fashioned through science to meet their every need.

As Jim Endersby, a historian, points out, some aspects of these biotopias shock modern sensibilities. For one thing, their inhabitants engineer themselves, as well as their environments, in a way that eugenicists like Wells saw as entirely unproblematic. It was the treatment of people as means not ends in this way that Aldous Huxley rebelled against in his own contribution to the genre, "Brave New World"(1932). Synthetic biology will certainly get caught up in the post-eugenic discussions of such matters that CRISPR has brought to the fore today.

Other now-alien attitudes in those early biotopias also throw light on today's arguments. Wells and Gilman saw no problem with deliberately exterminating species; it was a reasonable, even natural, imposition of beneficial order. Today such possibilities are being discussed for real, but with a lot less equanimity. "Gene drives"—genetic systems which, seemingly paradoxically, use sexual transmission to spread sterility—offer a way that CRISPR-technologies might be used to try to wipe out disease vectors, such as the species of mosquito that spread malaria.

Some look at a death toll of hundreds of thousands of people a year and see in this an elegant solution; there are plenty of non-malarial species of mosquito around to pick up any ecological slack. Others ask by what right they might bring an extinction

about, what risks they would expose other less dispensable insects to and what sort of informed consent they might possibly get. Steered by this discussion, research on gene drives—funded by, among others, the Bill and Melinda Gates Foundation—is increasingly, and wisely, focused on ways to break chains of transmission by crashing mosquito populations locally and temporarily, rather than globally, for ever. Insect populations so depleted that there are no longer enough to feed the summer swallows—a feature of Wells's biotopia—is the sort of problem synthetic biologists talk of putting right by cutting back on the need for pesticides, not something they would seek to engineer.

Really, auks

It is not just that today's biotopians are averse to extinctions. Some talk of reversing them—of using the tools of their trade to bring back the passenger pigeon, the great auk, the woolly mammoth, the American chestnut and others. Genomes preserved in museums or permafrost can be sequenced, and the genomes found in related species reprogrammed to produce something similar. In a small gesture in this direction, Ginkgo has made a scent that smells of a type of hibiscus that is now extinct.

This idea, too, meets with scepticism, even repugnance. Some feel that the results would be a Potemkin creation—new creatures that preserve a mere semblance of what has been lost, rather than restoring its essence. Some environmentalists also see it as a grotesque caricature of a problem that their movement has long suffered from: concentrating on a few high-profile species while ignoring the wholesale destruction of others that are less glamorous. But a recent report commissioned by the International Union for the Conservation of Nature found that conservation does need new tools, and synthetic biology offers opportunities in that respect—while also bringing with it risks, both direct and indirect, that need to be assessed on a case-by-case basis and in a precautionary way. Some conservationists are keen to see how well gene drives can wipe out invasive species on islands.

That tools so radical might be used to conserve or preserve, not disrupt, might seem a bit of a contradiction—even perverse. But it is worth considering that the changes wrought by synthetic biology could refashion humankind's relationship with the natural world at a technical and conceptual level and at the same time bring little dislocation to everyday life. As Kelly and Zach Weinersmith put it in their book "Soonish", synthetic biology may be "like Frankenstein, except the monster spends the whole book dutifully making medicine and industrial inputs".

At a "Build-a-Cell" workshop in San Diego this February the assembled researchers noted how hard it was to communicate to the public the remarkable scope of their ambitions: creating genomes and the cells to house them from almost first principles. If you appreciate the conceptual bravura of an organism with no ancestors, or that even discussing such a thing would have seemed insane just 25 years ago, this is staggering. If you do not, such synthetic life seems just to be, well, more life. And life is both already a miracle and the most everyday one. Cell is a cell is a cell.

It may be that the public is on to something. The application of genetic technology to human health and enhancement will be hotly debated. So will worries about how such technologies can be kept out of dangerous hands. But the fundamental change in the relation of the human and the natural may not seem so dramatic.

Consider the Colombian exchange. Shuffling together the ecosystems of the New World and the Old was a huge event in terms of both biology and human history. It wiped out populations and overhauled ways of life. But today an Indian cooking with chilli, or a German smoking a pipe of tobacco, or a Mexican admiring a mustang running free give little thought to the alien biology they are using and appreciating. Synthetic biology's innovations may be similarly woven into the background of the world, but without the concomitant suffering.

Even the most gentle transformation in the relationship between people and nature, though, may bring harm to some. Interacting with nature is often something the poorest depend on most. If synthetic biology replaces prized natural products with cheaper artificial ones, it will break ties to nature that are both meaningful and economically important, particularly in the developing world.

Take vanilla pods. The key flavour molecule in vanilla is already synthesised by chemists. Vanilla pods, though, contain a range of related molecules which provides something richer and consumers value them for that reason. Now synthetic biology might match or surpass nature's subtlety. So Dr Endy, now at Stanford, expected remonstration when he met a Mexican vanilla farmer at a synthetic-biology discussion. Instead, he found a man enthused. Vanilla was not just his livelihood but also a cause for ceremony, a provider of solace, a source of stories. And synthetic biology used properly, he thought, would not replace it as such, but could instead enhance it. It could bring new subtleties to valued scents, welcome strangeness to a well-loved story. The connection between people and plant might be deepened, not displaced.

It is easy to assume that reprogrammed life is a lesser life, innately commercial and desacrilised—that as the machineries of cell and laboratory become ever more tightly bound, they will squeeze out something that is human, or natural, or both, which ought to sit in the space between them. But it is also possible that a new appreciation can grow out of that space, a sense of what life is and could be, extended and enriched by new understanding. Think of it as a tune not yet composed, or the catnip scent of a lost flower reimagined and smelled, as if blooming for the first time. ■