Contents

Introduction

Imagine a virtual human, not made of flesh and bone but one made of bits and bytes, and not just any human, but a virtual version of you, accurate at every scale, from the way your heart beats down to the letters of your DNA code.

—*Virtual Humans* movie premiere, Science Museum, London

Within the walls of a nineteenth-century chapel on the outskirts of Barcelona, a heart starts to contract. This is not a real heart but a virtual copy of one that still pounds inside a patient's chest. With its billions of equations, and 100 million patches of simulated cells, the digital twin pumps at a leisurely rate of around one beat per hour as it tests treatments, from drugs to implants.

Though it was deconsecrated many decades ago, the Chapel Torre Girona is still adorned with a cross above its entrance. You can sense a higher power and purpose inside its romantic architecture. There, as sunlight streams through its stained-glass windows, you are confronted by an enormous glass-and-steel room, within which stand three ranks of black cabinets dotted with green lights.

This is MareNostrum (the Roman name for the Mediterranean Sea), a supercomputer on the campus of the Polytechnic University of Catalonia that is used by Peter Coveney along with colleagues across Europe to simulate electrical, chemical, and mechanical processes within the human body. These simulations look just like the real thing, whether a fluttering heart or a lung expanding into the chest. Much more important, however, is that these virtual organs behave like the real thing.

To show the dazzling range and potential of virtual human research, we used MareNostrum to create a movie, with the help of simulations run on other supercomputers, notably SuperMUC-NG in Germany (the suffix *MUC* refers to the code of nearby Munich Airport). Working with an international team, we wanted our *Virtual Humans* movie to showcase where these diverse efforts to create a body *in silico* could take medicine.

FIGURE 1. Still from the *Virtual Humans* movie. (CompBioMed and Barcelona Supercomputing Centre)

In September 2017, we held the premiere in the cavernous IMAX Cinema of the Science Museum in London with Fernando Cucchietti and Guillermo Marin, our colleagues from the Barcelona Supercomputing Centre. Even though we had worked for many months on the movie, gazing up at a pounding virtual heart the size of four double-decker buses still left us a little breathless.

SuperMUC-NG and MareNostrum 4 are among a few hundred or so great computational machines dotted around the world that are being harnessed to model the cosmos, understand the patterns of nature and meet the major challenges facing our society, such as studying how the Earth will cope with climate change, developing low-carbon energy sources and modelling the spread of virtual pandemics.

Just as great medieval cathedrals were raised by architects, masons, geometers and bishops to give humankind a glimpse of the infinite, supercomputers are the cathedrals of the information age, where novel worlds of endless variety, even entire universes, can be simulated within these great engines of logic, algorithms and information.

You can also re-create the inner worlds of the human body, and not just any body, or an average body, but a particular person, from their tissues and organs down to the molecular machines at work within their cells, their component proteins along with their DNA.

FIGURE 2. The MareNostrum supercomputer. (Wikimedia Commons: Gemmaribasmaspoch. CC-BY-SA-4.0)

The eventual aim of this endeavour is to capture life's rhythms, patterns and disorders in a computer, not just of any life or an average life, but of one particular body and one particular life—yours.[1]

At the premiere, we were joined by colleagues who had developed virtual hearts, arteries and veins along with the skeleton and its musculature. On that great IMAX screen in the Science Museum, the packed audience glimpsed a future when drugs can be designed to suit an individual patient, when we can visualise the shimmering movements of a mutated protein in the body, track the turbulent flow of drug particles deep into the lungs, study the surges of blood cells through the brain, and simulate the stresses and strains that play on weakened bones.

Rise of Digital Twins

In engineering, virtual copies are known as digital twins. The concept is usually attributed to a paper by John Vickers and Michael Grieves at the University of Michigan in 2002,[2] which talked of a "Mirrored Spaces Model." NASA coined the term *digital twin* in 2010,[3] and applied this way of thinking to spacecraft.[4] However, the origins of this approach can be glimpsed much earlier. Many cite the Apollo moon programme as one notable example, when simulators on the ground were used as analogue twins of spacecraft. This approach was famously employed in 1970 to help return three astronauts safely to Earth in the aftermath of an explosion 200,000 miles out in space on board the ill-fated Apollo 13 mission.[5]

Today, digital twins are well established. Many industrial processes and machines are too complex for one brain to grasp, so experimenting with their digital twins makes their behaviour easier to explore and understand.[6] Lessons learned this way are transforming the future of manufacturing and, by accelerating automation, altering the future of work. Digital copies of machines, even entire factories, are helping to anticipate hurdles, perfect designs and prevent mistakes before they occur.

Digital twins are used to optimise supply chains and store layouts; General Electric used a twin to boost efficiency at an aluminium smelter in India; a twin of the route of a proposed railway line in north west England—in the form of 18 billion data points harvested by drones—was created to help manage this vast transport project; a "factory of the future" in Australia honed a virtual copy of a robotic workstation before building the real thing; engineers use digital twins to estimate the lifetime of a jet engine and how to maintain it efficiently. Digital twins have been used to help create wind turbines, oil rigs, cars, jet engines, aircraft, spacecraft and more besides. Some believe that digital twin cities hold the key to future urban planning.

Digital twins are emerging in medicine too, thanks to the data revolution in biology. One of the legions of people sifting through health data is Leroy Hood of the Institute for Systems Biology, Seattle. Among the most influential of today's biotechnologists, Hood has worked at the leading edge of medicine, engineering and genetics for decades, dating back to the first human genome programme meeting in 1985. In 2015, he launched a venture that gathered a plethora of data on 5000 patients for five years. All their data were stored in what Hood calls "personal health clouds."* Analysis of a patient's cloud can reveal telltale signals of what Hood calls "pre-pre-disease" that doctors could use to anticipate problems, then intervene to maintain their health.

Hood talks of "scientific wellness," which "leverages personal, dense, dynamic data clouds to quantify and define wellness and identify deviations from well states toward disease." A living embodiment of his approach, the 82-year-old was on sprightly form ("I plan never to retire") when we talked to him about his vision of a "P4" future, where treatments are predictive, preventive, personalised, and participatory. Simulations of the body will help usher in that future by making sense of what patterns in a patient's data hold in store for them.

In reality, of course, we make do with incomplete understanding and incomplete data. But, as advances in weather forecasting have shown, these shortcomings can be overcome to make useful predictions. We have come a long way since 1922 when, in his remarkable book *Weather Prediction by Numerical Process*, the British mathematician Lewis Fry Richardson (1881–1953) outlined the idea of a fantastic forecast factory, where thousands of human "computers," using slide rules and calculators, are coordinated by a "conductor." Richardson mused on whether "some day in the dim future it will be possible to advance the computations faster than the weather advances." But even he went on to admit that his forecast factory was only a dream.

A century later, his extraordinary vision has become a reality. Supercomputers can make predictions a few days into the future with reasonable accuracy by constantly updating sophisticated computer models with data from orbiting satellites, buoys, aircraft, ships and weather stations.

A typical forecasting model relies on a system of equations to simulate whether it is going to rain or shine. There is an equation for momentum, density, and temperature in each of water's three phases (vapour, liquid and solid), and potentially for other chemical variables too, such as the ozone that absorbs harmful ultraviolet radiation. In

* Leroy Hood, interview with Peter Coveney and Roger Highfield, August 12, 2021.

Chapter Two, we spell out why these nonlinear differential equations, notably partial differential equations, rule the climate system. In all, it takes *billions* of equations to model the planet down to a resolution of, currently, around 60 kilometres.* Overall, the model has to take account of ever-changing thermodynamic, radiative and chemical processes working on scales from hundreds of metres to thousands of kilometres, and from seconds to weeks.[7] That represents a tour de force of simulation, one that some claim already approaches the complexity required to model the human brain.

Thanks to the torrent of biomedical data available today, along with ever more powerful theory and computation, we believe simulations will revolutionise biology just as much as they have transformed meteorology. The American meteorologist Cleveland Abbe (1838–1916) once declared how progress in his field depended on "the consecration of the physicist and mathematician to this science."[8] To echo his 1895 vision of forecasting, we look forward to the day when it is not enough to know someone is unwell—we want to be able to understand if they will fall sick and why, so that we can make them better and for longer.

Optimism about the potential of digital twins in medicine is bolstered by our current ability to forecast weather, which would amaze Abbe. We take the daily forecasts for granted, but this feat of prediction is truly extraordinary. Markus Covert of Stanford University, who has developed virtual cells, remarked that "prediction of storms such as Hurricane Sandy ten days in advance of landfall—with the corresponding evacuation of hundreds of residents, saving both lives and property—could arguably be ranked as among the great technical triumphs in human history."[9]

When it comes to climate forecasts, plans are under way to create a "digital twin" of Earth that would simulate the atmosphere, ocean, ice, and land down to a resolution of one kilometre, providing forecasts of the risks of floods, droughts, and fires, along with the swirling ocean eddies that shift heat and carbon around the planet. This European model, Destination Earth, will fold in other data, such as energy use, traffic patterns and human movements (traced by mobile

* Tim Palmer, email to Peter Coveney, June 2, 2021.

phones), to reveal how climate change will affect society—and how society could alter the trajectory of climate change in what some already call the Anthropocene, a geological epoch where human activity is having a significant impact on our planet.[10]

The details of creating a digital twin of our own planet Earth are staggering. Take clouds, as one example. They are made of water, which is also the main ingredient of the human body (around 68%[11]). Unlike us, however, clouds seem simple—great plumes of water droplets or ice crystals floating in the sky. Their formation is critical to our ability to predict weather, important for our understanding of the effects of global heating and central to controversial schemes to curb climate change through geoengineering.[12]

From cumulus tufts with beguiling shapes to great sheets of grey, clouds are a beautiful example of how complexity can result from simplicity, as droplets of water are borne on air currents of convection. As these droplets condense inside clouds, a little heat is released, making the clouds buoyant. At great heights, where temperatures fall well below freezing, the droplets turn into ice crystals, giving the resulting cirrus clouds a wispy, feathery look.

Within a cloud, processes at the smallest scales govern the formation of droplets. But, though microscopic, these features and interactions have large-scale, macroscopic, effects. The smaller and more numerous the droplets, the more that light is scattered. At the scale of micrometres, turbulence accelerates cloud formation and triggers rain showers.[13] Large-scale air motions can create vast cloud systems that can span a continent. By reflecting light into space, clouds can cool the Earth's surface, which is why some believe they should be nurtured to help curb runaway global warming.[14]

Essentially all the laws that underpin cloud formation are known, so we should be able to represent how they evolve in terms of known mathematical equations. The hope is to achieve the same for virtual humans, even down to the last water molecule. This may sound fantastical, but optimism that mathematics can describe the warm, complex, dynamic world of the body dates back centuries. The English physician William Harvey (1578–1657) relied on calculations in his demonstration of the circulation of the blood,[15] while in 1865 the French physiologist Claude Bernard (1813–1878) stated that "the application of mathematics to natural phenomena is the aim of all science."[16]

Our ability to create a virtual copy of a person depends on describing the body with the language of mathematics. Although a work in progress, equations written using calculus, which express rates of change, can already depict complex processes uncovered by molecular biologists, cell biologists and many others in the biosciences. These mathematical expressions—ordinary and partial differential equations—can describe at every instant how blood pressure varies depending on where you make a measurement in the body or track an electrical impulse as it speeds along a neuron in the brain, or how quickly a virus steals into a person's airway.

To put these equations to work, all that is needed to start calculating are the boundary conditions for the problem at hand. This could mean the state of a neuron or an infected cell at a given time or at various time intervals, their rates of change at various instants or the upper and lower limits of a given quantity. These conditions tether the mathematics to reality so we can make forecasts about the body, or "healthcasts," by analogy with the weather.

But while we accept that the laws of nature are universal, in one critical and practical sense the life sciences—by which we mean biology and medicine—are quite different from the physical sciences—physics and chemistry—that we use to describe clouds. They are more empirical, more dependent on making measurements and doing experiments and, until now, less dependent on theoretical understanding.

Theory, that is, the mathematical representation of the laws of nature, plays a relatively diminished role in medicine and biology. Even the Darwin-Wallace theory of evolution, regarded by some as the greatest scientific theory of all, does not admit a mathematical description. This might sound shocking, but the reality is that, while basic predictions about the patterns of inheritance have been made since Gregor Mendel studied peas in the nineteenth century, the course of evolution is not possible to predict in any quantitative manner.[17]

Some influential figures are only too aware of this shortcoming. Paul Nurse, director of the Francis Crick Institute in London and former assistant editor of the *Journal of Theoretical Biology*, told us how he was weary of reading papers that use clever technology to make measurements that come to “barely any significant conclusions.”* In an opinion article for the journal *Nature*, he cited Sydney Brenner (1927–2019), his old friend and fellow Nobelist: “We are drowning in a sea of data and starving for knowledge.”[18] He complained to us that the importance of theory and the principles of life are relatively neglected in favour of cramming facts, knowledge and information. Biology “does have ideas, so why aren’t we talking about them?”

Yet biology, like the rest of science, is undoubtedly governed by the laws of nature. To be sure, there are no-go areas for moral and ethical reasons based on human arguments, but there is absolutely every reason to believe that we should be able to understand a particular scientific aspect of how an organism works and capture that insight in the form of mathematics. To create Virtual You, we need to go beyond the current use of theory in making post hoc rationalisations in biology, after studies are carried out, to using theory to guide experiments and make predictions.

Uniting Science

Science is balkanised. The notion of dividing academic inquirers into tribes dates back to ancient Greece with Socrates (c. 469–399 BCE), his student Plato (c. 428–347 BCE) and, in turn, Plato’s student Aristotle (384–322 BCE).[19] Within a few decades, however, Timon of Philius (c. 320–230 BCE) moaned about the squabbling of “bookish cloisterlings” at the Museum of Alexandria. By the sixteenth century, Francis Bacon (1561–1626) and other philosophers were mourning the splintering of human knowledge.

By the mid-nineteenth century, the disciplinary boundaries of the modern university had taken root, each with its own customs, language, funding streams, establishments and practices. In *Virtual You*, we intend to show that today’s research is more than a baggy collection of fragmented efforts—it is a grand and complementary mosaic of data, models, mechanisms and technology. The big picture of how the human body works is beginning to heave into view.

* Paul Nurse, interview with Peter Coveney and Roger Highfield, September 25, 2021.

Just as there is no privileged point of view of the human body, so each perspective from each discipline is equally important. Each is complementary and, if united and consistent, remarkable new insights can emerge. If we look, for example, at the great molecular biology revolution that dates from the 1950s, when physicists and chemists tackled biology, and biologists used techniques developed by physicists, we can see that this vital atomic view of proteins, enzymes and other molecules of living things perfectly complements existing insights into heredity and evolution, marking a powerful unification of knowledge known as consilience.

The simple idea at the heart of this book is that the convergence of many branches of science—patient data, theory, algorithms, AI and powerful computers—is taking medicine in a new direction, one that is quantitative and predictive. We will show how mathematics can capture an extraordinary range of processes at work in living things, weigh up developments in computer hardware and software and then show how the human body can be portrayed *in silico*, holding up a digital mirror to reflect our possible futures.

This is a story that builds on multidisciplinary ideas we set out in our earlier books, *The Arrow of Time*[20] and *Frontiers of Complexity*.[21] In the first, we discussed how to reconcile a deep problem at the heart of science: that time is represented in different ways by different theories and at different length scales, ranging from the microscopic to the macroscopic. In the latter, we showed how complexity in mathematics, physics, biology, chemistry and even the social sciences is transforming not only the way we think about the universe, but also the very assumptions that underlie conventional science, and how computers are essential if we are to explore and understand this complexity. Nowhere is this more relevant than in the efforts to create the virtual human. In *Virtual You*, we draw these threads together within a broad tapestry of research, both historical and contemporary.

Virtual You

This is the first account of the global enterprise to create a virtual human aimed at the general reader. Hundreds of millions of dollars have been spent in the past two decades on the effort that has been organised through initiatives such as the International Physiome Project,[22] America's Cancer Patient Digital Twin,[23] the European Virtual Physiological Human,[24] the Human Brain Project[25] and another Europe-wide effort led by University College London to which we both contribute, Computational Biomedicine, or CompBioMed for short.

All are united by a single objective. As one workshop held in Tokyo declared: "The time is now ripe to initiate a grand challenge project to create over the next 30 years a comprehensive, molecules-based, multi-scale, computational model of the human ('the virtual human'), capable of simulating and predicting, with a reasonable degree of accuracy, the consequences of most of the perturbations that are relevant to healthcare."[26] That virtual vision was unveiled more than a decade ago—in February 2008—and its future is fast approaching.

In the following pages, we will take you on a fantastic voyage through the body, its organ systems, cells and tissues along with the deformable protein machines that run them. We hope to convince you that, in coming decades, virtual twins of cells, organs, and populations of virtual humans will increasingly shape healthcare. This organising principle for twenty-first-century medicine will enable doctors for the first time to look forward to—and predict—what is in store for you, including the effects of proposed therapies. This marks a stark contrast with today's approach where doctors, in effect, look back at what happened to similar (though nonidentical) patients in similar (though nonidentical) circumstances.

In the long term, virtual cells, organs and humans—along with populations of virtual humans—will help to evolve the current generation of one-size-fits-all medicine into truly personalised medicine. Your digital twin will help you understand what forms of diet, exercise and lifestyle will offer you the healthiest future. Ultimately, the rise of these digital twins could pave the way for methods to enhance your body and your future. As we discuss in our concluding chapter, virtual humans will hold up a mirror to reflect on the very best that you can be.

The following four chapters focus on the fundamental steps that are required to create a digital twin: harvest diverse data about the body (Chapter One); craft theory to make sense of all these data

(Chapter Two) and use mathematics to understand the fundamental limits of simulations; harness computers to put the spark of life into mathematical understanding of the human body (Chapter Three); blend the insights of natural and artificial intelligence to interpret data and to shape our understanding (Chapter Four).

In Chapters Five to Eight, we show the consequences of taking these steps and begin to build a digital twin, from virtual infections (Chapter Five) to cells, organs, metabolism and bodies. Along the way, in Chapter Six, we encounter the fifth step necessary for the creation of Virtual You. Can we stitch together different mathematical models of different physical processes that operate across different domains of space and time within the body? We can, and the ability to customise a virtual heart to match that of a patient marks one extraordinary example (Chapter Seven), along with modelling the body and its organ systems (Chapter Eight). In Chapter Nine, we discuss "Virtual You 2.0," when the next generation of computers will overcome shortcomings of the current generation of "classical" digital computers.

In our last chapter, we examine the many opportunities, along with ethical and moral issues, that virtual humans will present. Digital twins will challenge what we mean by simple terms such as "healthy." Are you really healthy if your digital twin predicts that—without a treatment or a change in lifestyle—you will not live out your potential life span? You may feel "well," but are you really well if simulations suggest that you are destined to spend a decade longer in a care home than necessary? If a virtual human can become the substrate for human thought, how will we come to regard our digital copy? Finally, in an appendix, we examine a provocative question raised by using computers to simulate the world: Is it possible to re-create the fundamental physics of the cosmos from simple algorithms?

So, to the first of our foundational chapters. This poses the most basic question of all. If we are to create digital twins, how well do we have to know ourselves? To create Virtual You, we need to understand how much data and what kinds are sufficient for a digital twin to be animated by a computer.

As Aristotle once remarked, knowing yourself is the beginning of all wisdom.

FIGURE 3. Virtual anatomical twin. One of the detailed high-resolution anatomical models created from magnetic resonance image data of volunteers. (IT'IS Foundation)

Index

A page number in *italics* refers to a figure.