JANUARY 2012

nature physics insight

COMPLEXITY

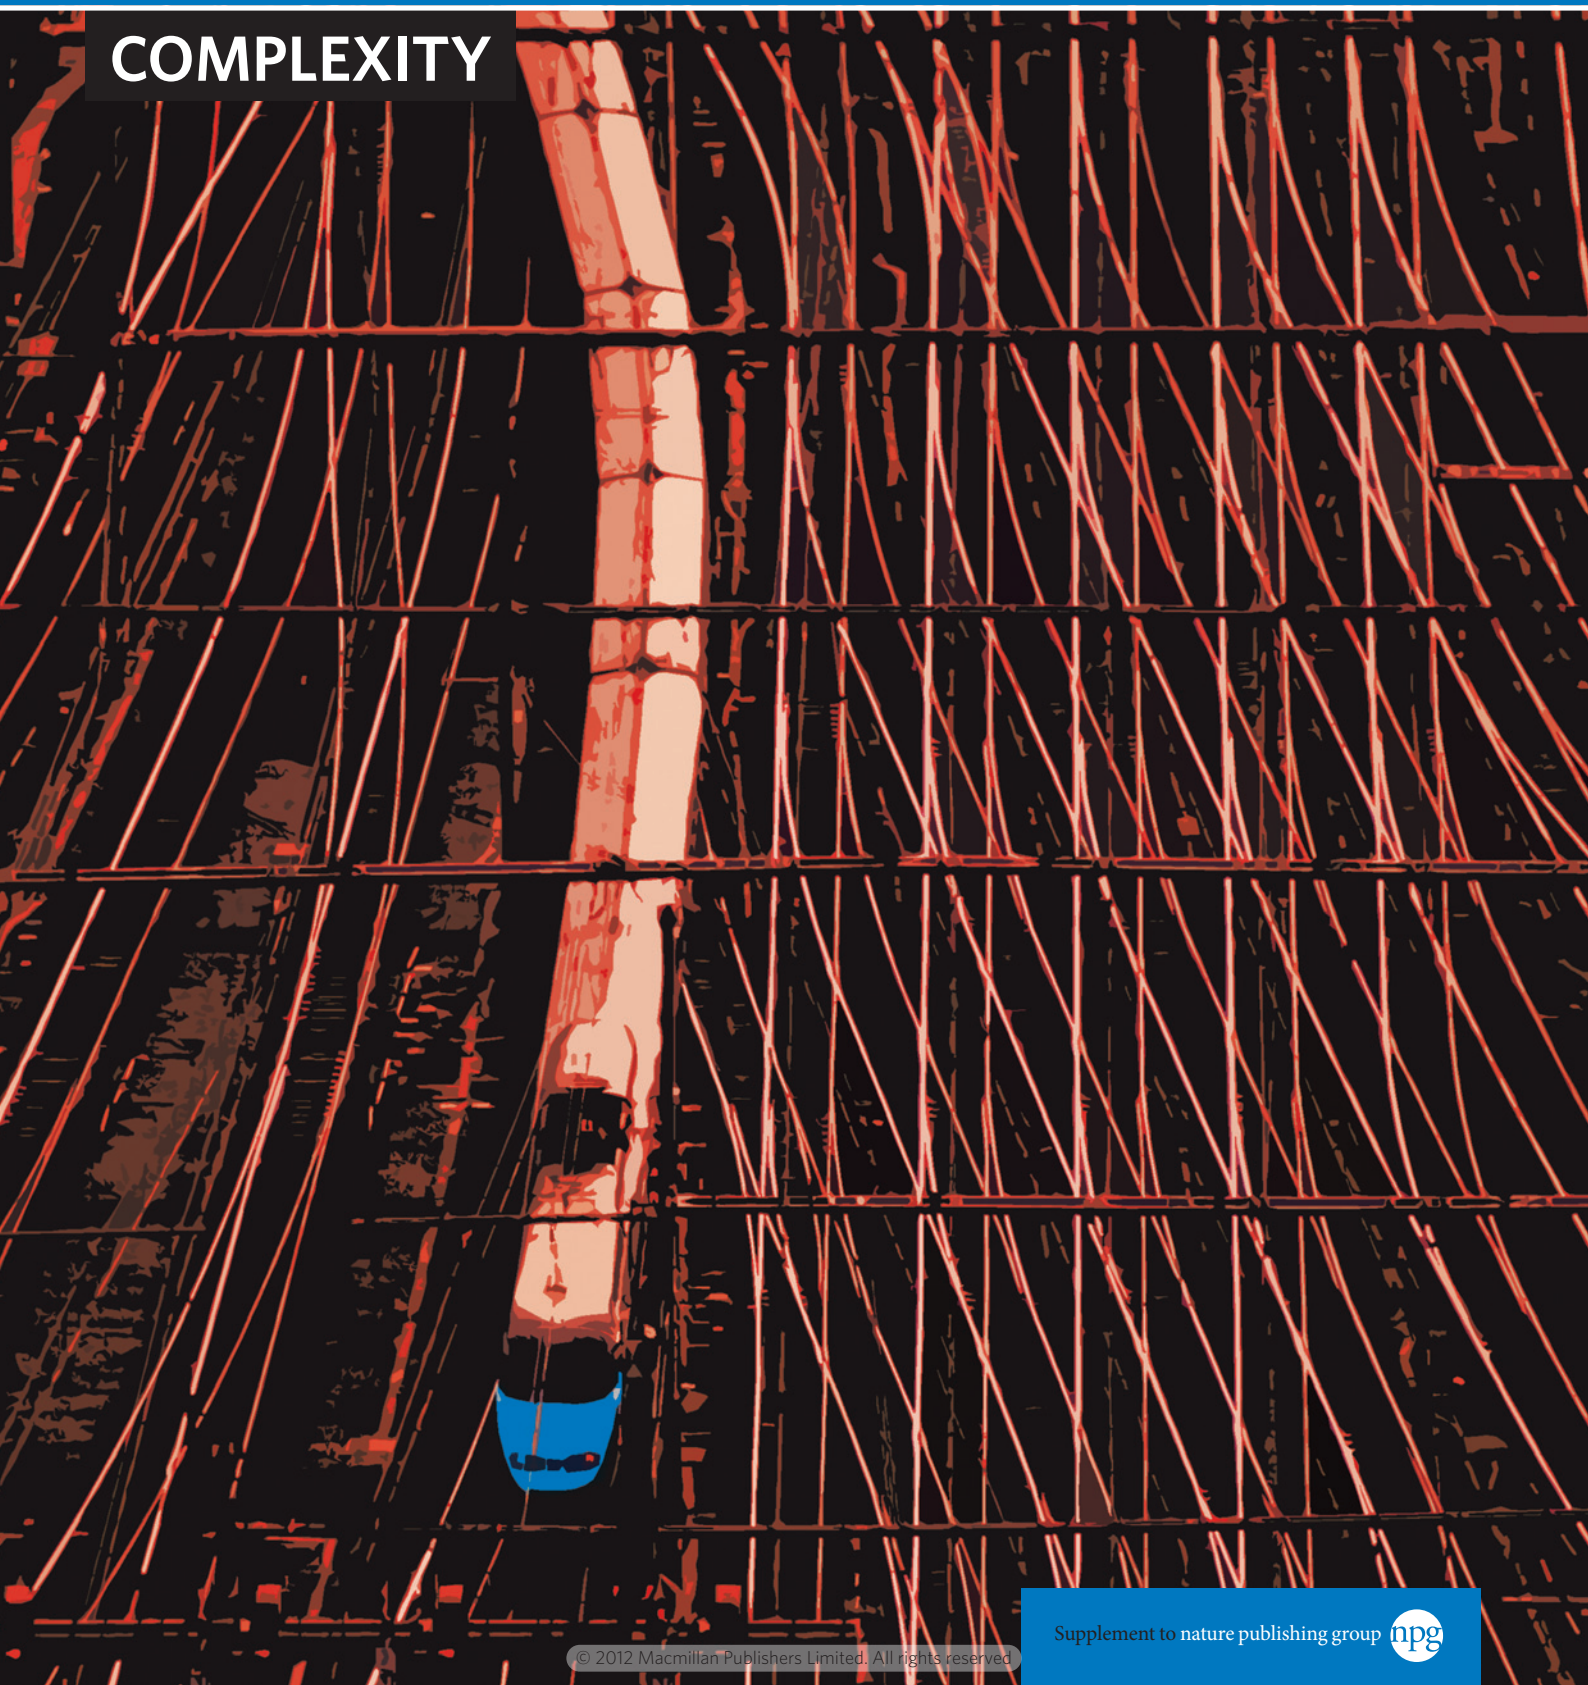

© 2012 Macmillan Publishers Limited. All rights reserved

Supplement to nature publishing group npg

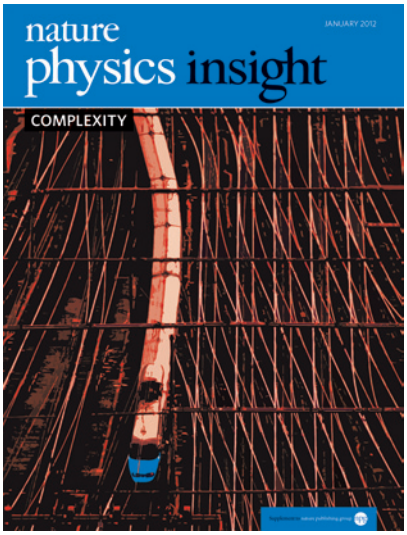

COVER IMAGE

In many large ensembles, the property of the system as a whole cannot be understood by studying the individual entities — neurons in the brain, for example, or transport users in traffic networks. The past decade, however, has seen important progress in our fundamental understanding of what such seemingly disparate 'complex systems' have in common. Image: © Marvin E. Newman/ Getty Images.

NPG LONDON
The Macmillan Building,
4 Crinan Street, London N1 9XW
T: +44 207 833 4000
F: +44 207 843 4563
naturephysics@nature.com

EDITOR
ALISON WRIGHT

INSIGHT EDITOR
ANDREAS TRABESINGER

PRODUCTION EDITOR
JENNY MARSDEN

COPY EDITOR
MELANIE HANVEY

ART EDITOR
KAREN MOORE

EDITORIAL ASSISTANT
AMANDA WYATT

MARKETING
SARAH-JANE BALDOCK

PUBLISHER
RUTH WILSON

EDITOR-IN-CHIEF,
NATURE PUBLICATIONS
PHILIP CAMPBELL

nature publishing group

Complexity

A formal definition of what constitutes a complex system is not easy to devise; equally difficult is the delineation of which fields of study fall within the bounds of 'complexity'. An appealing approach — but only one of several possibilities — is to play on the 'more is different' theme, declaring that the properties of a complex system as a whole cannot be understood from the study of its individual constituents. There are many examples, from neurons in the brain, to transport users in traffic networks, to data packages in the Internet.

Large datasets — collected, for example, in proteomic studies, or captured in records of mobile-phone users and Internet traffic — now provide an unprecedented level of information about these systems. Indeed, the availability of these detailed datasets has led to an explosion of activity in the modelling of complex systems. Data-based models can not only provide an understanding of the properties and behaviours of individual systems, but also, beyond that, might lead to the discovery of common properties between seemingly disparate systems.

Much of the progress made during the past decade or so comes under the banner of 'network science'. The representation of complex systems as networks, or graphs, has proved to be a tremendously useful abstraction, and has led to an understanding of how many real-world systems are structured, what kinds of dynamic processes they support and how they interact with each other. This *Nature Physics* Insight is therefore admittedly inclined towards research in complex networks. As Albert-László Barabási argues in his Commentary, the past decade has indeed witnessed a 'network takeover'. On the other hand, James Crutchfield, in his review of the tools for discovering patterns and quantifying their structural complexity, demonstrates beautifully how fundamental theories of information and computation have led to a deeper understanding of just what 'complex systems' are.

For a topic as broad as complexity, it is impossible to do justice to all of the recent developments. The field has been shaped over decades by advances in physics, engineering, computer science, biology and sociology, and its ramifications are equally diverse. But a selection had to be made, and we hope that this Insight will prove inspiring, and a showcase for the pivotal role that physicists are playing — and are bound to play — in the inherently multidisciplinary endeavour of making sense of complexity.

Andreas Trabesinger, Senior Editor

CONTENTS

© 2012 Macmillan Publishers Limited. All rights reserved

The network takeover

Albert-László Barabási

Reductionism, as a paradigm, is expired, and complexity, as a field, is tired. Data-based mathematical models of complex systems are offering a fresh perspective, rapidly developing into a new discipline: network science.

Reports of the death of reductionism are greatly exaggerated. It is so ingrained in our thinking that if one day some magical force should make us all forget it, we would promptly have to reinvent it. The real worry is not with reductionism, which, as a paradigm and tool, is rather useful. It is necessary, but no longer sufficient. But, weighing up better ideas, it became a burden.

"You never want a serious crisis to go to waste," Ralph Emmanuel, at that time Obama's chief of staff, famously proclaimed in November 2008, at the height of the financial meltdown. Indeed, forced by an imminent need to go beyond reductionism, a new network-based paradigm is emerging that is taking science by storm. It relies on datasets that are inherently incomplete and noisy. It builds on a set of sharp tools, developed during the past decade, that seem to be just as useful in search engines as in cell biology. It is making a real impact from science to industry. Along the way it points to a new way to handle a century-old problem: complexity.

A better understanding of the pieces cannot solve the difficulties that many research fields currently face, from cell biology to software design. There is no 'cancer gene'. A typical cancer patient has mutations in a few dozen of about 300 genes, an elusive combinatorial problem whose complexity is increasingly a worry to the medical community. No single regulation can legislate away the economic malady that is slowly eating at our wealth. It is the web of diverging financial and political interests that makes policy so difficult to implement. Consciousness cannot be reduced to a single neuron. It is an emergent property that engages billions of synapses. In fact, the more we know about the workings of individual genes, banks or neurons, the less we understand the system as a whole. Consequently, an increasing number of the big questions of contemporary science are rooted in the same problem: we hit the limits of reductionism. No need to mount a defence of it. Instead, we need to tackle the real question in front of us: complexity.

The complexity argument is by no means new. It has re-emerged repeatedly during the past decades. The fact that it is still fresh underlines the lack of progress achieved so far. It also stays with us for good reason: complexity research is a thorny undertaking. First, its goals are easily confusing to the outsider. What does it aim to address — the origins of social order, biological complexity or economic interconnectedness? Second, decades of research on complexity were driven by big, sweeping theoretical ideas, inspired by toy models and differential equations that ultimately failed to deliver. Think synergetics and its slave modes; think chaos theory, ultimately telling us more about unpredictability than how to predict nonlinear systems; think self-organized criticality, a sweeping collection of scaling ideas squeezed into a sand pile; think fractals, hailed once as the source of all answers to the problems of pattern formation. We learned a lot, but achieved little: our tools failed to keep up with the shifting challenges that complex systems pose. Third, there is a looming methodological question: what should a theory of complexity deliver? A new Maxwellian formula, condensing into a set of elegant equations every ill that science faces today? Or a new uncertainty principle, encoding what we can and what we can't do in complex systems? Finally, who owns the science of complexity? Physics? Engineering? Biology, mathematics, computer science? All of the above? Anyone?

These questions have resisted answers for decades. Yet something has changed in the past few years. The driving force behind this change can be condensed into a single word: data. Fuelled by cheap sensors and high-throughput technologies,

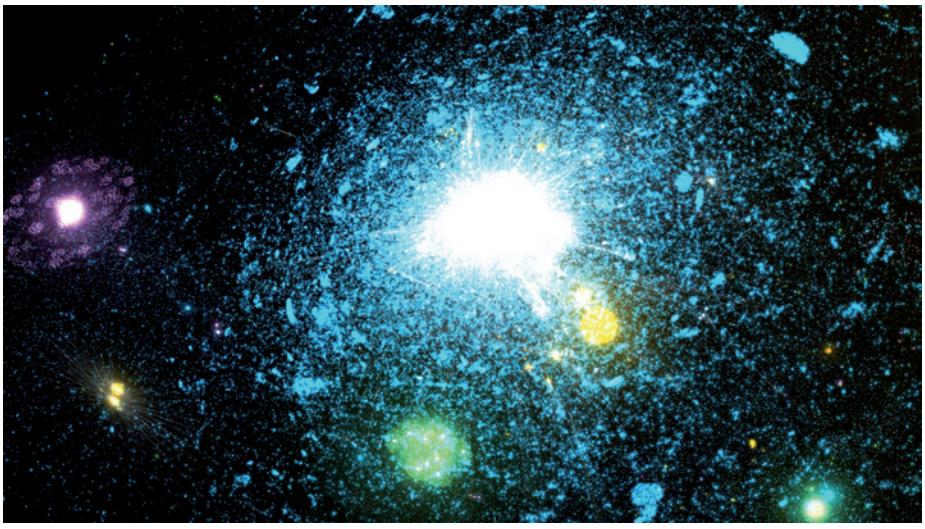

Network universe. A visualization of the first large-scale network explicitly mapped out to explore the large-scale structure of real networks. The map was generated in 1999 and represents a small portion of the World Wide Web[11]; this map has led to the discovery of scale-free networks. Nodes are web documents; links correspond to URLs. Visualization by Mauro Martino, Alec Pawling and Chaoming Song.

© 2012 Macmillan Publishers Limited. All rights reserved

the data explosion that we witness today, from social media to cell biology, is offering unparalleled opportunities to document the inner workings of many complex systems. Microarray and proteomic tools offer us the simultaneous activity of all human genes and proteins; mobile-phone records capture the communication and mobility patterns of whole countries[1]; import–export and stock data condense economic activity into easily accessible databases[2]. As scientists sift through these mountains of data, we are witnessing an increasing awareness that if we are to tackle complexity, the tools to do so are being born right now, in front of our eyes. The field that benefited most from this data windfall is often called network theory, and it is fundamentally reshaping our approach to complexity.

Born at the twilight of the twentieth century, network theory aims to understand the origins and characteristics of networks that hold together the components in various complex systems. By simultaneously looking at the World Wide Web and genetic networks, Internet and social systems, it led to the discovery that despite the many differences in the nature of the nodes and the interactions between them, the networks behind most complex systems are governed by a series of fundamental laws that determine and limit their behaviour.

An increasing number of the big questions of contemporary science are rooted in the same problem: we hit the limits of reductionism.

On the surface, network theory is prone to the failings of its predecessors. It has its own big ideas, from scale-free networks to the theory of network evolution[3], from community formation[4,5] to dynamics on networks[6]. But there is a defining difference. These ideas have not been gleaned from toy models or mathematical anomalies. They are based on data and meticulous observations. The theory of evolving networks was motivated by extensive empirical evidence documenting the scale-free nature of the degree distribution, from the cell to the World Wide Web; the formalism behind degree correlations was preceded by data documenting correlations on the Internet and on cellular maps[7,8]; the extensive theoretical work on spreading processes was preceded by decades of meticulous data collection on the spread of viruses and fads, gaining a proper theoretical footing in the network context[6]. This data-inspired methodology is an important shift compared with earlier takes on complex systems. Indeed, in a survey of the ten most influential papers in complexity, it will be difficult to find one that builds directly on experimental data. In contrast, among the ten most cited papers in network theory, you will be hard pressed to find one that does not directly rely on empirical evidence.

With its deep empirical basis and its host of analytical and algorithmic tools, today network theory is indispensible in the study of complex systems. We will never understand the workings of a cell if we ignore the intricate networks through which its proteins and metabolites interact with each other. We will never foresee economic meltdowns unless we map out the web of indebtedness that characterizes the financial system. These profound changes in complexity research echo major economic and social shifts. The economic giants of our era are no longer carmakers and oil producers, but the companies that build, manage or fuel our networks: Cisco, Google, Facebook, Apple or Twitter. Consequently, during the past decade, question by question and system by system, network science has hijacked complexity research. Reductionism deconstructed complex systems, bringing us a theory of individual nodes and links. Network theory is painstakingly reassembling them, helping us to see the whole again. One thing is increasingly clear: no theory of the cell, of social media or of the Internet can ignore the profound network effects that their interconnectedness cause. Therefore, if we are ever to have a theory of complexity, it will sit on the shoulders of network theory.

The daunting reality of complexity research is that the problems it tackles are so diverse that no single theory can satisfy all needs. The expectations of social scientists for a theory of social complexity are quite different from the questions posed by biologists as they seek to uncover the phenotypic heterogeneity of cardiovascular disease. We may, however, follow in the footsteps of Steve Jobs, who once insisted that it is not the consumer's job to know what they want. It is our job, those of us working on the mathematical theory of complex systems, to define the science of the complex. Although no theory can satisfy all needs, what we can strive for is a broad framework within which most needs can be addressed.

The twentieth century has witnessed the birth of such a sweeping, enabling framework: quantum mechanics. Many advances of the century, from electronics to astrophysics, from nuclear energy to quantum computation, were built on the theoretical foundations that it offered. In the twenty-first century, network theory is emerging as its worthy successor: it is building a theoretical and algorithmic framework that is energizing many research fields, and it is closely followed by many industries. As network theory develops its mathematical and intellectual core, it is becoming an indispensible platform for science, business and security, helping to discover new drug targets, delivering Facebook's latest algorithms and aiding the efforts to halt terrorism.

Who owns the science of complexity?

As physicists, we cannot avoid the elephant in the room: what is the role of physics in this journey? We physicists do not have an excellent track record in investing in our future. For decades, we forced astronomers into separate departments, under the slogan: it is not physics. Now we bestow on them our highest awards, such as last year's Nobel Prize. For decades we resisted biological physics, exiling our brightest colleagues to medical schools. Along the way we missed out on the bio-revolution, bypassing the financial windfall that the National Institutes of Health bestowed on biological complexity, proudly shrinking our physics departments instead. We let materials science be taken over by engineering schools just when the science had matured enough to be truly lucrative. Old reflexes never die, making many now wonder whether network science is truly physics. The answer is obvious: it is much bigger than physics. Yet physics is deeply entangled with it: the Institute for Scientific Information (ISI) highlighted two network papers[3,9] among the ten most cited physics papers of the past decade, and in about a year Chandrashekhar's 1945 tome, which has been the most cited paper in *Review of Modern Physics* for decades, will be dethroned by a decade-old paper on network theory[10]. Physics has as much to offer to this journey as it has to benefit from it.

Although physics has owned complexity research for many decades, it is not without competition any longer. Computer science, fuelled by its poster progenies,

© 2012 Macmillan Publishers Limited. All rights reserved

such as Google or Facebook, is mounting a successful attack on complexity, fuelled by the conviction that a sufficiently fast algorithm can tackle any problem, no matter how complex. This confidence has prompted the US Directorate for Computer and Information Science and Engineering to establish the first network-science programme within the US National Science Foundation. Bioinformatics, with its rich resources backed by the National Institutes of Health, is pushing from a different direction, aiming to quantify biological complexity. Complexity and network science need both the intellectual and financial resources that different communities can muster. But as the field enters the spotlight, physics must assert its engagement if it wants to continue to be present at the table.

As I follow the debate surrounding the faster-than-light neutrinos, I wish deep down for it to be true. Physics needs the shot in the arm that such a development could deliver. Our children no longer want to become physicists and astronauts. They want to invent the next Facebook instead. Short of that, they are happy to land a job at Google. They don't talk quanta — they dream bits. They don't see entanglement but recognize with ease nodes and links. As complexity takes a driving seat in science, engineering and business, we physicists cannot afford to sit on the sidelines. We helped to create it. We owned it for decades. We must learn to take pride in it. And this means, as our forerunners did a century ago with quantum mechanics, that we must invest in it and take it to its conclusion. ❐

Albert-László Barabási is at the Center for Complex Network Research and Departments of Physics, Computer Science and Biology, Northeastern University, Boston, Massachusetts 02115, USA; the Center for Cancer Systems Biology, Dana-Farber Cancer Institute, Boston, Massachusetts 02115, USA; and the Department of Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, Massachusetts 02115, USA.
e-mail: alb@neu.edu

References

1. Onnela, J. P. *et al. Proc. Natl Acad. Sci. USA* **104,** 7332–7336 (2007).
2. Hidalgo, C. A., Klinger, B., Barabási, A. L. & Hausmann, R. *Science* **317,** 482–487 (2007).
3. Barabási, A. L. & Albert, R. *Science* **286,** 509–512 (1999).
4. Newman, M. E. J. *Networks: An Introduction* (Oxford Univ. Press, 2010).
5. Palla, G., Farkas, I. J., Derényi, I. & Vicsek, T. *Nature* **435,** 814–818 (2005).
6. Pastor-Satorras, R. & Vespignani, A. *Phys. Rev. Lett.* **86,** 3200–3203 (2001).
7. Pastor-Satorras, R., Vázquez, A. & Vespignani, A. *Phys. Rev. Lett.* **87,** 258701 (2001).
8. Maslov, S. & Sneppen, K. *Science* **296,** 910–913 (2002).
9. Watts, D. J. & Strogatz, S. H. *Nature* **393,** 440–442 (1998).
10. Barabási, A. L. & Albert, R. *Rev. Mod. Phys.* **74,** 47–97 (2002).
11. Albert, R., Jeong, H. & Barabási, A-L. *Nature* **401,** 130-131 (1999).

© 2012 Macmillan Publishers Limited. All rights reserved

nature physics

INSIGHT | REVIEW ARTICLES

PUBLISHED ONLINE: 22 DECEMBER 2011 | DOI: 10.1038/NPHYS2190

Between order and chaos

James P. Crutchfield

What is a pattern? How do we come to recognize patterns never seen before? Quantifying the notion of pattern and formalizing the process of pattern discovery go right to the heart of physical science. Over the past few decades physics' view of nature's lack of structure—its unpredictability—underwent a major renovation with the discovery of deterministic chaos, overthrowing two centuries of Laplace's strict determinism in classical physics. Behind the veil of apparent randomness, though, many processes are highly ordered, following simple rules. Tools adapted from the theories of information and computation have brought physical science to the brink of automatically discovering hidden patterns and quantifying their structural complexity.

One designs clocks to be as regular as physically possible. So much so that they are the very instruments of determinism. The coin flip plays a similar role; it expresses our ideal of the utterly unpredictable. Randomness is as necessary to physics as determinism—think of the essential role that 'molecular chaos' plays in establishing the existence of thermodynamic states. The clock and the coin flip, as such, are mathematical ideals to which reality is often unkind. The extreme difficulties of engineering the perfect clock[1] and implementing a source of randomness as pure as the fair coin testify to the fact that determinism and randomness are two inherent aspects of all physical processes.

In 1927, van der Pol, a Dutch engineer, listened to the tones produced by a neon glow lamp coupled to an oscillating electrical circuit. Lacking modern electronic test equipment, he monitored the circuit's behaviour by listening through a telephone ear piece. In what is probably one of the earlier experiments on electronic music, he discovered that, by tuning the circuit as if it were a musical instrument, fractions or subharmonics of a fundamental tone could be produced. This is markedly unlike common musical instruments—such as the flute, which is known for its purity of harmonics, or multiples of a fundamental tone. As van der Pol and a colleague reported in *Nature* that year[2], 'the turning of the condenser in the region of the third to the sixth subharmonic strongly reminds one of the tunes of a bag pipe'.

Presciently, the experimenters noted that when tuning the circuit 'often an irregular noise is heard in the telephone receivers before the frequency jumps to the next lower value'. We now know that van der Pol had listened to deterministic chaos: the noise was produced in an entirely lawful, ordered way by the circuit itself. The *Nature* report stands as one of its first experimental discoveries. Van der Pol and his colleague van der Mark apparently were unaware that the deterministic mechanisms underlying the noises they had heard had been rather keenly analysed three decades earlier by the French mathematician Poincaré in his efforts to establish the orderliness of planetary motion[3–5]. Poincaré failed at this, but went on to establish that determinism and randomness are essential and unavoidable twins[6]. Indeed, this duality is succinctly expressed in the two familiar phrases 'statistical mechanics' and 'deterministic chaos'.

Complicated yes, but is it complex?

As for van der Pol and van der Mark, much of our appreciation of nature depends on whether our minds—or, more typically these days, our computers—are prepared to discern its intricacies. When confronted by a phenomenon for which we are ill-prepared, we often simply fail to see it, although we may be looking directly at it. Perception is made all the more problematic when the phenomena of interest arise in systems that spontaneously organize.

Spontaneous organization, as a common phenomenon, reminds us of a more basic, nagging puzzle. If, as Poincaré found, chaos is endemic to dynamics, why is the world not a mass of randomness? The world is, in fact, quite structured, and we now know several of the mechanisms that shape microscopic fluctuations as they are amplified to macroscopic patterns. Critical phenomena in statistical mechanics[7] and pattern formation in dynamics[8,9] are two arenas that explain in predictive detail how spontaneous organization works. Moreover, everyday experience shows us that nature inherently organizes; it generates pattern. Pattern is as much the fabric of life as life's unpredictability.

In contrast to patterns, the outcome of an observation of a random system is unexpected. We are surprised at the next measurement. That surprise gives us information about the system. We must keep observing the system to see how it is evolving. This insight about the connection between randomness and surprise was made operational, and formed the basis of the modern theory of communication, by Shannon in the 1940s (ref. 10). Given a source of random events and their probabilities, Shannon defined a particular event's degree of surprise as the negative logarithm of its probability: the event's self-information is $I_i = -\log_2 p_i$. (The units when using the base-2 logarithm are bits.) In this way, an event, say i, that is certain ($p_i = 1$) is not surprising: $I_i = 0$ bits. Repeated measurements are not informative. Conversely, a flip of a fair coin ($p_{\text{Heads}} = 1/2$) is maximally informative: for example, $I_{\text{Heads}} = 1$ bit. With each observation we learn in which of two orientations the coin is, as it lays on the table.

The theory describes an information source: a random variable X consisting of a set $\{i = 0, 1, \ldots, k\}$ of events and their probabilities $\{p_i\}$. Shannon showed that the averaged uncertainty $H[X] = \sum_i p_i I_i$—the source entropy rate—is a fundamental property that determines how compressible an information source's outcomes are.

With information defined, Shannon laid out the basic principles of communication[11]. He defined a communication channel that accepts messages from an information source X and transmits them, perhaps corrupting them, to a receiver who observes the channel output Y. To monitor the accuracy of the transmission, he introduced the mutual information $I[X;Y] = H[X] - H[X|Y]$ between the input and output variables. The first term is the information available at the channel's input. The second term, subtracted, is the uncertainty in the incoming message, if the receiver knows the output. If the channel completely corrupts, so

Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, California 95616, USA.
*e-mail: chaos@ucdavis.edu.

that none of the source messages accurately appears at the channel's output, then knowing the output Y tells you nothing about the input and $H[X|Y] = H[X]$. In other words, the variables are statistically independent and so the mutual information vanishes. If the channel has perfect fidelity, then the input and output variables are identical; what goes in, comes out. The mutual information is the largest possible: $I[X;Y] = H[X]$, because $H[X|Y] = 0$. The maximum input–output mutual information, over all possible input sources, characterizes the channel itself and is called the channel capacity:

$$C = \max_{\{P(X)\}} I[X;Y]$$

Shannon's most famous and enduring discovery though—one that launched much of the information revolution—is that as long as a (potentially noisy) channel's capacity C is larger than the information source's entropy rate $H[X]$, there is way to encode the incoming messages such that they can be transmitted error free[11]. Thus, information and how it is communicated were given firm foundation.

How does information theory apply to physical systems? Let us set the stage. The system to which we refer is simply the entity we seek to understand by way of making observations. The collection of the system's temporal behaviours is the process it generates. We denote a particular realization by a time series of measurements: $\ldots x_{-2}x_{-1}x_0x_1\ldots$. The values x_t taken at each time can be continuous or discrete. The associated bi-infinite chain of random variables is similarly denoted, except using uppercase: $\ldots X_{-2}X_{-1}X_0X_1\ldots$. At each time t the chain has a past $X_{:t} = \ldots X_{t-2}X_{t-1}$ and a future $X_{:} = X_tX_{t+1}\ldots$. We will also refer to blocks $X_{t':} = X_tX_{t+1}\ldots X_{t'-1}, t < t'$. The upper index is exclusive.

To apply information theory to general stationary processes, one uses Kolmogorov's extension of the source entropy rate[12,13]. This is the growth rate h_μ:

$$h_\mu = \lim_{\ell\to\infty} \frac{H(\ell)}{\ell}$$

where $H(\ell) = -\sum_{\{x_{:\ell}\}} \Pr(x_{:\ell}) \log_2 \Pr(x_{:\ell})$ is the block entropy—the Shannon entropy of the length-ℓ word distribution $\Pr(x_{:\ell})$. h_μ gives the source's intrinsic randomness, discounting correlations that occur over any length scale. Its units are bits per symbol, and it partly elucidates one aspect of complexity—the randomness generated by physical systems.

We now think of randomness as surprise and measure its degree using Shannon's entropy rate. By the same token, can we say what 'pattern' is? This is more challenging, although we know organization when we see it.

Perhaps one of the more compelling cases of organization is the hierarchy of distinctly structured matter that separates the sciences—quarks, nucleons, atoms, molecules, materials and so on. This puzzle interested Philip Anderson, who in his early essay 'More is different'[14], notes that new levels of organization are built out of the elements at a lower level and that the new 'emergent' properties are distinct. They are not directly determined by the physics of the lower level. They have their own 'physics'.

This suggestion too raises questions, what is a 'level' and how different do two levels need to be? Anderson suggested that organization at a given level is related to the history or the amount of effort required to produce it from the lower level. As we will see, this can be made operational.

Complexities

To arrive at that destination we make two main assumptions. First, we borrow heavily from Shannon: every process is a communication channel. In particular, we posit that any system is a channel that communicates its past to its future through its present. Second, we take into account the context of interpretation. We view building models as akin to decrypting nature's secrets. How do we come to understand a system's randomness and organization, given only the available, indirect measurements that an instrument provides? To answer this, we borrow again from Shannon, viewing model building also in terms of a channel: one experimentalist attempts to explain her results to another.

The following first reviews an approach to complexity that models system behaviours using exact deterministic representations. This leads to the deterministic complexity and we will see how it allows us to measure degrees of randomness. After describing its features and pointing out several limitations, these ideas are extended to measuring the complexity of ensembles of behaviours—to what we now call statistical complexity. As we will see, it measures degrees of structural organization. Despite their different goals, the deterministic and statistical complexities are related and we will see how they are essentially complementary in physical systems.

Solving Hilbert's famous *Entscheidungsproblem* challenge to automate testing the truth of mathematical statements, Turing introduced a mechanistic approach to an effective procedure that could decide their validity[15]. The model of computation he introduced, now called the Turing machine, consists of an infinite tape that stores symbols and a finite-state controller that sequentially reads symbols from the tape and writes symbols to it. Turing's machine is deterministic in the particular sense that the tape contents exactly determine the machine's behaviour. Given the present state of the controller and the next symbol read off the tape, the controller goes to a unique next state, writing at most one symbol to the tape. The input determines the next step of the machine and, in fact, the tape input determines the entire sequence of steps the Turing machine goes through.

Turing's surprising result was that there existed a Turing machine that could compute any input–output function—it was universal. The deterministic universal Turing machine (UTM) thus became a benchmark for computational processes.

Perhaps not surprisingly, this raised a new puzzle for the origins of randomness. Operating from a fixed input, could a UTM generate randomness, or would its deterministic nature always show through, leading to outputs that were probabilistically deficient? More ambitiously, could probability theory itself be framed in terms of this new constructive theory of computation? In the early 1960s these and related questions led a number of mathematicians—Solomonoff[16,17] (an early presentation of his ideas appears in ref. 18), Chaitin[19], Kolmogorov[20] and Martin-Löf[21]—to develop the algorithmic foundations of randomness.

The central question was how to define the probability of a single object. More formally, could a UTM generate a string of symbols that satisfied the statistical properties of randomness? The approach declares that models M should be expressed in the language of UTM programs. This led to the Kolmogorov–Chaitin complexity $KC(\mathbf{x})$ of a string $\mathbf{x}$. The Kolmogorov–Chaitin complexity is the size of the minimal program P that generates $\mathbf{x}$ running on a UTM (refs 19,20):

$$KC(\mathbf{x}) = \text{argmin}\{|P| : UTM \circ P = \mathbf{x}\}$$

One consequence of this should sound quite familiar by now. It means that a string is random when it cannot be compressed: a random string is its own minimal program. The Turing machine simply prints it out. A string that repeats a fixed block of letters, in contrast, has small Kolmogorov–Chaitin complexity. The Turing machine program consists of the block and the number of times it is to be printed. Its Kolmogorov–Chaitin complexity is logarithmic

in the desired string length, because there is only one variable part of P and it stores $\log \ell$ digits of the repetition count ℓ.

Unfortunately, there are a number of deep problems with deploying this theory in a way that is useful to describing the complexity of physical systems.

First, Kolmogorov–Chaitin complexity is not a measure of structure. It requires exact replication of the target string. Therefore, $KC(\mathbf{x})$ inherits the property of being dominated by the randomness in $\mathbf{x}$. Specifically, many of the UTM instructions that get executed in generating $\mathbf{x}$ are devoted to producing the 'random' bits of $\mathbf{x}$. The conclusion is that Kolmogorov–Chaitin complexity is a measure of randomness, not a measure of structure. One solution, familiar in the physical sciences, is to discount for randomness by describing the complexity in ensembles of behaviours.

Furthermore, focusing on single objects was a feature, not a bug, of Kolmogorov–Chaitin complexity. In the physical sciences, however, this is a prescription for confusion. We often have access only to a system's typical properties, and even if we had access to microscopic, detailed observations, listing the positions and momenta of molecules is simply too huge, and so useless, a description of a box of gas. In most cases, it is better to know the temperature, pressure and volume.

The issue is more fundamental than sheer system size, arising even with a few degrees of freedom. Concretely, the unpredictability of deterministic chaos forces the ensemble approach on us.

The solution to the Kolmogorov–Chaitin complexity's focus on single objects is to define the complexity of a system's process—the ensemble of its behaviours[22]. Consider an information source that produces collections of strings of arbitrary length. Given a realization $x_{:\ell}$ of length ℓ, we have its Kolmogorov–Chaitin complexity $KC(x_{:\ell})$, of course, but what can we say about the Kolmogorov–Chaitin complexity of the ensemble $\{x_{:\ell}\}$? First, define its average in terms of samples $\{x^i_{:\ell} : i = 1, \ldots, M\}$:

$$KC(\ell) = \langle KC(x_{:\ell}) \rangle = \lim_{M \to \infty} \frac{1}{M} \sum_{i=1}^{M} KC(x^i_{:\ell})$$

How does the Kolmogorov–Chaitin complexity grow as a function of increasing string length? For almost all infinite sequences produced by a stationary process the growth rate of the Kolmogorov–Chaitin complexity is the Shannon entropy rate[23]:

$$h_\mu = \lim_{\ell \to \infty} \frac{KC(\ell)}{\ell}$$

As a measure—that is, a number used to quantify a system property—Kolmogorov–Chaitin complexity is uncomputable[24,25]. There is no algorithm that, taking in the string, computes its Kolmogorov–Chaitin complexity. Fortunately, this problem is easily diagnosed. The essential uncomputability of Kolmogorov–Chaitin complexity derives directly from the theory's clever choice of a UTM as the model class, which is so powerful that it can express undecidable statements.

One approach to making a complexity measure constructive is to select a less capable (specifically, non-universal) class of computational models. We can declare the representations to be, for example, the class of stochastic finite-state automata[26,27]. The result is a measure of randomness that is calibrated relative to this choice. Thus, what one gains in constructiveness, one looses in generality.

Beyond uncomputability, there is the more vexing issue of how well that choice matches a physical system of interest. Even if, as just described, one removes uncomputability by choosing a less capable representational class, one still must validate that these, now rather specific, choices are appropriate to the physical system one is analysing.

At the most basic level, the Turing machine uses discrete symbols and advances in discrete time steps. Are these representational choices appropriate to the complexity of physical systems? What about systems that are inherently noisy, those whose variables are continuous or are quantum mechanical? Appropriate theories of computation have been developed for each of these cases[28,29], although the original model goes back to Shannon[30]. More to the point, though, do the elementary components of the chosen representational scheme match those out of which the system itself is built? If not, then the resulting measure of complexity will be misleading.

Is there a way to extract the appropriate representation from the system's behaviour, rather than having to impose it? The answer comes, not from computation and information theories, as above, but from dynamical systems theory.

Dynamical systems theory—Poincaré's qualitative dynamics—emerged from the patent uselessness of offering up an explicit list of an ensemble of trajectories, as a description of a chaotic system. It led to the invention of methods to extract the system's 'geometry from a time series'. One goal was to test the strange-attractor hypothesis put forward by Ruelle and Takens to explain the complex motions of turbulent fluids[31].

How does one find the chaotic attractor given a measurement time series from only a single observable? Packard and others proposed developing the reconstructed state space from successive time derivatives of the signal[32]. Given a scalar time series $x(t)$, the reconstructed state space uses coordinates $y_1(t) = x(t)$, $y_2(t) = \mathrm{d}x(t)/\mathrm{d}t, \ldots, y_m(t) = \mathrm{d}^m x(t)/\mathrm{d}t^m$. Here, $m+1$ is the embedding dimension, chosen large enough that the dynamic in the reconstructed state space is deterministic. An alternative is to take successive time delays in $x(t)$ (ref. 33). Using these methods the strange attractor hypothesis was eventually verified[34].

It is a short step, once one has reconstructed the state space underlying a chaotic signal, to determine whether you can also extract the equations of motion themselves. That is, does the signal tell you which differential equations it obeys? The answer is yes[35]. This sound works quite well if, and this will be familiar, one has made the right choice of representation for the 'right-hand side' of the differential equations. Should one use polynomial, Fourier or wavelet basis functions; or an artificial neural net? Guess the right representation and estimating the equations of motion reduces to statistical quadrature: parameter estimation and a search to find the lowest embedding dimension. Guess wrong, though, and there is little or no clue about how to update your choice.

The answer to this conundrum became the starting point for an alternative approach to complexity—one more suitable for physical systems. The answer is articulated in computational mechanics[36], an extension of statistical mechanics that describes not only a system's statistical properties but also how it stores and processes information—how it computes.

The theory begins simply by focusing on predicting a time series $\ldots X_{-2}X_{-1}X_0X_1\ldots$. In the most general setting, a prediction is a distribution $\Pr(X_{t:}|x_{:t})$ of futures $X_{t:} = X_tX_{t+1}X_{t+2}\ldots$ conditioned on a particular past $x_{:t} = \ldots x_{t-3}x_{t-2}x_{t-1}$. Given these conditional distributions one can predict everything that is predictable about the system.

At root, extracting a process's representation is a very straightforward notion: do not distinguish histories that make the same predictions. Once we group histories in this way, the groups themselves capture the relevant information for predicting the future. This leads directly to the central definition of a process's effective states. They are determined by the equivalence relation:

$$x_{:t} \sim x_{:t'} \Leftrightarrow \Pr(X_{t:}|x_{:t}) = \Pr(X_{t:}|x_{:t'})$$

The equivalence classes of the relation $\sim$ are the process's causal states $\mathcal{S}$—literally, its reconstructed state space, and the induced state-to-state transitions are the process's dynamic $\mathcal{T}$—its equations of motion. Together, the states $\mathcal{S}$ and dynamic $\mathcal{T}$ give the process's so-called ϵ-machine.

Why should one use the ϵ-machine representation of a process? First, there are three optimality theorems that say it captures all of the process's properties[36–38]: prediction: a process's ϵ-machine is its optimal predictor; minimality: compared with all other optimal predictors, a process's ϵ-machine is its minimal representation; uniqueness: any minimal optimal predictor is equivalent to the ϵ-machine.

Second, we can immediately (and accurately) calculate the system's degree of randomness. That is, the Shannon entropy rate is given directly in terms of the ϵ-machine:

$$h_\mu = -\sum_{\sigma\in\mathcal{S}} \Pr(\sigma) \sum_{\{x\}} \Pr(x|\sigma) \log_2 \Pr(x|\sigma)$$

where $\Pr(\sigma)$ is the distribution over causal states and $\Pr(x|\sigma)$ is the probability of transitioning from state σ on measurement x.

Third, the ϵ-machine gives us a new property—the statistical complexity—and it, too, is directly calculated from the ϵ-machine:

$$C_\mu = -\sum_{\sigma\in\mathcal{S}} \Pr(\sigma) \log_2 \Pr(\sigma)$$

The units are bits. This is the amount of information the process stores in its causal states.

Fourth, perhaps the most important property is that the ϵ-machine gives all of a process's patterns. The ϵ-machine itself—states plus dynamic—gives the symmetries and regularities of the system. Mathematically, it forms a semi-group[39]. Just as groups characterize the exact symmetries in a system, the ϵ-machine captures those and also 'partial' or noisy symmetries.

Finally, there is one more unique improvement the statistical complexity makes over Kolmogorov–Chaitin complexity theory. The statistical complexity has an essential kind of representational independence. The causal equivalence relation, in effect, extracts the representation from a process's behaviour. Causal equivalence can be applied to any class of system—continuous, quantum, stochastic or discrete.

Independence from selecting a representation achieves the intuitive goal of using UTMs in algorithmic information theory—the choice that, in the end, was the latter's undoing. The ϵ-machine does not suffer from the latter's problems. In this sense, computational mechanics is less subjective than any 'complexity' theory that per force chooses a particular representational scheme.

To summarize, the statistical complexity defined in terms of the ϵ-machine solves the main problems of the Kolmogorov–Chaitin complexity by being representation independent, constructive, the complexity of an ensemble, and a measure of structure.

In these ways, the ϵ-machine gives a baseline against which any measures of complexity or modelling, in general, should be compared. It is a minimal sufficient statistic[38].

To address one remaining question, let us make explicit the connection between the deterministic complexity framework and that of computational mechanics and its statistical complexity. Consider realizations $\{x_{:\ell}\}$ from a given information source. Break the minimal UTM program P for each into two components: one that does not change, call it the 'model' M; and one that does change from input to input, E, the 'random' bits not generated by M. Then, an object's 'sophistication' is the length of M (refs 40,41):

$$\mathrm{SOPH}(x_{:\ell}) = \mathrm{argmin}\{|M| : P = M + E, x_{:\ell} = UTM \circ P\}$$

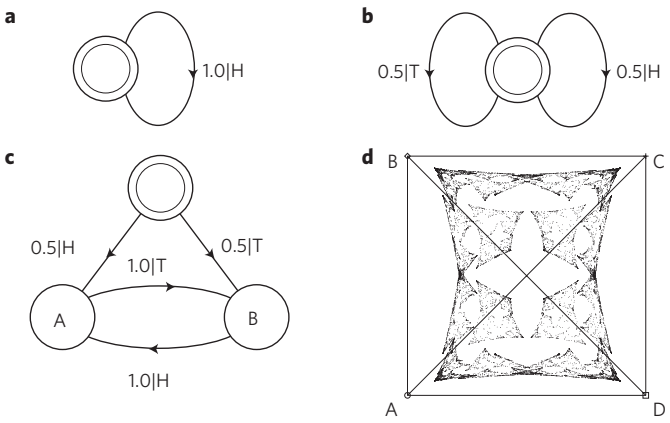

Figure 1 | ϵ-machines for four information sources. **a**, The all-heads process is modelled with a single state and a single transition. The transition is labelled $p|x$, where $p \in [0;1]$ is the probability of the transition and x is the symbol emitted. **b**, The fair-coin process is also modelled by a single state, but with two transitions each chosen with equal probability. **c**, The period-2 process is perhaps surprisingly more involved. It has three states and several transitions. **d**, The uncountable set of causal states for a generic four-state HMM. The causal states here are distributions $\Pr(A;B;C;D)$ over the HMM's internal states and so are plotted as points in a 4-simplex spanned by the vectors that give each state unit probability. Panel **d** reproduced with permission from ref. 44, © 1994 Elsevier.

As done with the Kolmogorov–Chaitin complexity, we can define the ensemble-averaged sophistication $\langle\mathrm{SOPH}\rangle$ of 'typical' realizations generated by the source. The result is that the average sophistication of an information source is proportional to its process's statistical complexity[42]:

$$KC(\ell) \propto C_\mu + h_\mu \ell$$

That is, $\langle\mathrm{SOPH}\rangle \propto C_\mu$.

Notice how far we come in computational mechanics by positing only the causal equivalence relation. From it alone, we derive many of the desired, sometimes assumed, features of other complexity frameworks. We have a canonical representational scheme. It is minimal and so Ockham's razor[43] is a consequence, not an assumption. We capture a system's pattern in the algebraic structure of the ϵ-machine. We define randomness as a process's ϵ-machine Shannon-entropy rate. We define the amount of organization in a process with its ϵ-machine's statistical complexity. In addition, we also see how the framework of deterministic complexity relates to computational mechanics.

Applications

Let us address the question of usefulness of the foregoing by way of examples.

Let's start with the Prediction Game, an interactive pedagogical tool that intuitively introduces the basic ideas of statistical complexity and how it differs from randomness. The first step presents a data sample, usually a binary times series. The second asks someone to predict the future, on the basis of that data. The final step asks someone to posit a state-based model of the mechanism that generated the data.

The first data set to consider is $x_{:0} = \ldots$HHHHHHH—the all-heads process. The answer to the prediction question comes to mind immediately: the future will be all Hs, $x_{:} =$ HHHHH…. Similarly, a guess at a state-based model of the generating mechanism is also easy. It is a single state with a transition labelled with the output symbol H (Fig. 1a). A simple model for a simple process. The process is exactly predictable: $h_\mu = 0$

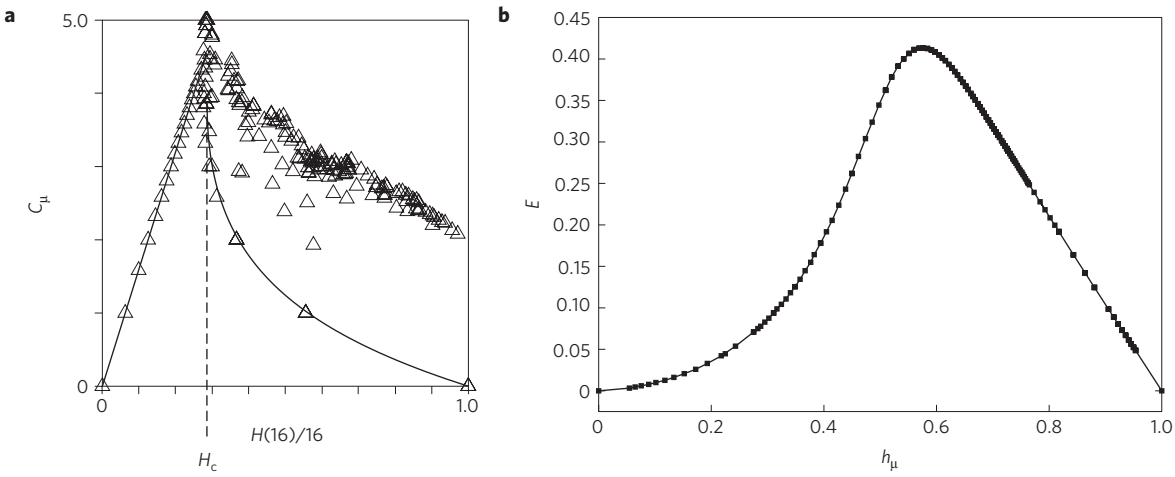

Figure 2 | Structure versus randomness. a, In the period-doubling route to chaos. **b**, In the two-dimensional Ising-spinsystem. Reproduced with permission from: **a**, ref. 36, © 1989 APS; **b**, ref. 61, © 2008 AIP.

bits per symbol. Furthermore, it is not complex; it has vanishing complexity: $C_\mu = 0$ bits.

The second data set is, for example, $x_{:0} = \ldots$THTHTTHTHH. What I have done here is simply flip a coin several times and report the results. Shifting from being confident and perhaps slightly bored with the previous example, people take notice and spend a good deal more time pondering the data than in the first case.

The prediction question now brings up a number of issues. One cannot exactly predict the future. At best, one will be right only half of the time. Therefore, a legitimate prediction is simply to give another series of flips from a fair coin. In terms of monitoring only errors in prediction, one could also respond with a series of all Hs. Trivially right half the time, too. However, this answer gets other properties wrong, such as the simple facts that Ts occur and occur in equal number.

The answer to the modelling question helps articulate these issues with predicting (Fig. 1b). The model has a single state, now with two transitions: one labelled with a T and one with an H. They are taken with equal probability. There are several points to emphasize. Unlike the all-heads process, this one is maximally unpredictable: $h_\mu = 1$ bit/symbol. Like the all-heads process, though, it is simple: $C_\mu = 0$ bits again. Note that the model is minimal. One cannot remove a single 'component', state or transition, and still do prediction. The fair coin is an example of an independent, identically distributed process. For all independent, identically distributed processes, $C_\mu = 0$ bits.

In the third example, the past data are $x_{:0} = \ldots$HTHTHTHTH. This is the period-2 process. Prediction is relatively easy, once one has discerned the repeated template word $w =$ TH. The prediction is $x_{:} =$ THTHTHTH$\ldots$. The subtlety now comes in answering the modelling question (Fig. 1c).

There are three causal states. This requires some explanation. The state at the top has a double circle. This indicates that it is a start state—the state in which the process starts or, from an observer's point of view, the state in which the observer is before it begins measuring. We see that its outgoing transitions are chosen with equal probability and so, on the first step, a T or an H is produced with equal likelihood. An observer has no ability to predict which. That is, initially it looks like the fair-coin process. The observer receives 1 bit of information. In this case, once this start state is left, it is never visited again. It is a transient causal state.

Beyond the first measurement, though, the 'phase' of the period-2 oscillation is determined, and the process has moved into its two recurrent causal states. If an H occurred, then it is in state A and a T will be produced next with probability 1. Conversely, if a T was generated, it is in state B and then an H will be generated. From this point forward, the process is exactly predictable: $h_\mu = 0$ bits per symbol. In contrast to the first two cases, it is a structurally complex process: $C_\mu = 1$ bit. Conditioning on histories of increasing length gives the distinct future conditional distributions corresponding to these three states. Generally, for p-periodic processes $h_\mu = 0$ bits symbol^{-1} and $C_\mu = \log_2 p$ bits.

Finally, Fig. 1d gives the ϵ-machine for a process generated by a generic hidden-Markov model (HMM). This example helps dispel the impression given by the Prediction Game examples that ϵ-machines are merely stochastic finite-state machines. This example shows that there can be a fractional dimension set of causal states. It also illustrates the general case for HMMs. The statistical complexity diverges and so we measure its rate of divergence—the causal states' information dimension[44].

As a second example, let us consider a concrete experimental application of computational mechanics to one of the venerable fields of twentieth-century physics—crystallography: how to find structure in disordered materials. The possibility of turbulent crystals had been proposed a number of years ago by Ruelle[53]. Using the ϵ-machine we recently reduced this idea to practice by developing a crystallography for complex materials[54–57].

Describing the structure of solids—simply meaning the placement of atoms in (say) a crystal—is essential to a detailed understanding of material properties. Crystallography has long used the sharp Bragg peaks in X-ray diffraction spectra to infer crystal structure. For those cases where there is diffuse scattering, however, finding—let alone describing—the structure of a solid has been more difficult[58]. Indeed, it is known that without the assumption of crystallinity, the inference problem has no unique solution[59]. Moreover, diffuse scattering implies that a solid's structure deviates from strict crystallinity. Such deviations can come in many forms—Schottky defects, substitution impurities, line dislocations and planar disorder, to name a few.

The application of computational mechanics solved the longstanding problem—determining structural information for disordered materials from their diffraction spectra—for the special case of planar disorder in close-packed structures in polytypes[60]. The solution provides the most complete statistical description of the disorder and, from it, one could estimate the minimum effective memory length for stacking sequences in close-packed structures. This approach was contrasted with the so-called fault

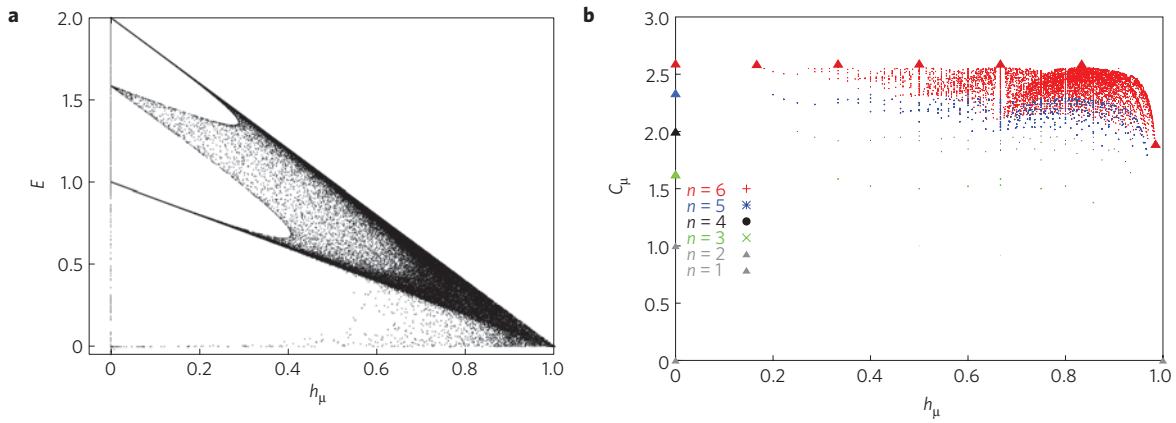

Figure 3 | Complexity–entropy diagrams. a, The one-dimensional, spin-1/2 antiferromagnetic Ising model with nearest- and next-nearest-neighbour interactions. Reproduced with permission from ref. 61, © 2008 AIP. **b**, Complexity–entropy pairs (h_μ, C_μ) for all topological binary-alphabet ϵ-machines with $n = 1, \ldots, 6$ states. For details, see refs 61 and 63.

model by comparing the structures inferred using both approaches on two previously published zinc sulphide diffraction spectra. The net result was that having an operational concept of pattern led to a predictive theory of structure in disordered materials.

As a further example, let us explore the nature of the interplay between randomness and structure across a range of processes. As a direct way to address this, let us examine two families of controlled system—systems that exhibit phase transitions. Consider the randomness and structure in two now-familiar systems: one from nonlinear dynamics—the period-doubling route to chaos; and the other from statistical mechanics—the two-dimensional Ising-spin model. The results are shown in the complexity–entropy diagrams of Fig. 2. They plot a measure of complexity (C_μ and **E**) versus the randomness ($H(16)/16$ and h_μ, respectively).

One conclusion is that, in these two families at least, the intrinsic computational capacity is maximized at a phase transition: the onset of chaos and the critical temperature. The occurrence of this behaviour in such prototype systems led a number of researchers to conjecture that this was a universal interdependence between randomness and structure. For quite some time, in fact, there was hope that there was a single, universal complexity–entropy function—coined the 'edge of chaos' (but consider the issues raised in ref. 62). We now know that although this may occur in particular classes of system, it is not universal.

It turned out, though, that the general situation is much more interesting[61]. Complexity–entropy diagrams for two other process families are given in Fig. 3. These are rather less universal looking. The diversity of complexity–entropy behaviours might seem to indicate an unhelpful level of complication. However, we now see that this is quite useful. The conclusion is that there is a wide range of intrinsic computation available to nature to exploit and available to us to engineer.

Finally, let us return to address Anderson's proposal for nature's organizational hierarchy. The idea was that a new, 'higher' level is built out of properties that emerge from a relatively 'lower' level's behaviour. He was particularly interested to emphasize that the new level had a new 'physics' not present at lower levels. However, what is a 'level' and how different should a higher level be from a lower one to be seen as new?

We can address these questions now having a concrete notion of structure, captured by the ϵ-machine, and a way to measure it, the statistical complexity C_μ. In line with the theme so far, let us answer these seemingly abstract questions by example. In turns out that we already saw an example of hierarchy, when discussing intrinsic computational at phase transitions.

Specifically, higher-level computation emerges at the onset of chaos through period-doubling—a countably infinite state ϵ-machine[42]—at the peak of C_μ in Fig. 2a.

How is this hierarchical? We answer this using a generalization of the causal equivalence relation. The lowest level of description is the raw behaviour of the system at the onset of chaos. Appealing to symbolic dynamics[64], this is completely described by an infinitely long binary string. We move to a new level when we attempt to determine its ϵ-machine. We find, at this 'state' level, a countably infinite number of causal states. Although faithful representations, models with an infinite number of components are not only cumbersome, but not insightful. The solution is to apply causal equivalence yet again—to the ϵ-machine's causal states themselves. This produces a new model, consisting of 'meta-causal states', that predicts the behaviour of the causal states themselves. This procedure is called hierarchical ϵ-machine reconstruction[45], and it leads to a finite representation—a nested-stack automaton[42]. From this representation we can directly calculate many properties that appear at the onset of chaos.

Notice, though, that in this prescription the statistical complexity at the 'state' level diverges. Careful reflection shows that this also occurred in going from the raw symbol data, which were an infinite non-repeating string (of binary 'measurement states'), to the causal states. Conversely, in the case of an infinitely repeated block, there is no need to move up to the level of causal states. At the period-doubling onset of chaos the behaviour is aperiodic, although not chaotic. The descriptional complexity (the ϵ-machine) diverged in size and that forced us to move up to the meta-ϵ-machine level.

This supports a general principle that makes Anderson's notion of hierarchy operational: the different scales in the natural world are delineated by a succession of divergences in statistical complexity of lower levels. On the mathematical side, this is reflected in the fact that hierarchical ϵ-machine reconstruction induces its own hierarchy of intrinsic computation[45], the direct analogue of the Chomsky hierarchy in discrete computation theory[65].

Closing remarks

Stepping back, one sees that many domains face the confounding problems of detecting randomness and pattern. I argued that these tasks translate into measuring intrinsic computation in processes and that the answers give us insights into how nature computes.

Causal equivalence can be adapted to process classes from many domains. These include discrete and continuous-output HMMs (refs 45,66,67), symbolic dynamics of chaotic systems[45],

molecular dynamics[68], single-molecule spectroscopy[67,69], quantum dynamics[70], dripping taps[71], geomagnetic dynamics[72] and spatiotemporal complexity found in cellular automata[73–75] and in one- and two-dimensional spin systems[76,77]. Even then, there are many remaining areas of application.

Specialists in the areas of complex systems and measures of complexity will miss a number of topics above: more advanced analyses of stored information, intrinsic semantics, irreversibility and emergence[46–52]; the role of complexity in a wide range of application fields, including biological evolution[78–83] and neural information-processing systems[84–86], to mention only two of the very interesting, active application areas; the emergence of information flow in spatially extended and network systems[74,87–89]; the close relationship to the theory of statistical inference[85,90–95]; and the role of algorithms from modern machine learning for nonlinear modelling and estimating complexity measures. Each topic is worthy of its own review. Indeed, the ideas discussed here have engaged many minds for centuries. A short and necessarily focused review such as this cannot comprehensively cite the literature that has arisen even recently; not so much for its size, as for its diversity.

I argued that the contemporary fascination with complexity continues a long-lived research programme that goes back to the origins of dynamical systems and the foundations of mathematics over a century ago. It also finds its roots in the first days of cybernetics, a half century ago. I also showed that, at its core, the questions its study entails bear on some of the most basic issues in the sciences and in engineering: spontaneous organization, origins of randomness, and emergence.

The lessons are clear. We now know that complexity arises in a middle ground—often at the order–disorder border. Natural systems that evolve with and learn from interaction with their immediate environment exhibit both structural order and dynamical chaos. Order is the foundation of communication between elements at any level of organization, whether that refers to a population of neurons, bees or humans. For an organism order is the distillation of regularities abstracted from observations. An organism's very form is a functional manifestation of its ancestor's evolutionary and its own developmental memories.

A completely ordered universe, however, would be dead. Chaos is necessary for life. Behavioural diversity, to take an example, is fundamental to an organism's survival. No organism can model the environment in its entirety. Approximation becomes essential to any system with finite resources. Chaos, as we now understand it, is the dynamical mechanism by which nature develops constrained and useful randomness. From it follow diversity and the ability to anticipate the uncertain future.

There is a tendency, whose laws we are beginning to comprehend, for natural systems to balance order and chaos, to move to the interface between predictability and uncertainty. The result is increased structural complexity. This often appears as a change in a system's intrinsic computational capability. The present state of evolutionary progress indicates that one needs to go even further and postulate a force that drives in time towards successively more sophisticated and qualitatively different intrinsic computation. We can look back to times in which there were no systems that attempted to model themselves, as we do now. This is certainly one of the outstanding puzzles[96]: how can lifeless and disorganized matter exhibit such a drive? The question goes to the heart of many disciplines, ranging from philosophy and cognitive science to evolutionary and developmental biology and particle astrophysics[96]. The dynamics of chaos, the appearance of pattern and organization, and the complexity quantified by computation will be inseparable components in its resolution.

Received 28 October 2011; accepted 30 November 2011; published online 22 December 2011

References

1. Press, W. H. Flicker noises in astronomy and elsewhere. *Comment. Astrophys.* **7,** 103–119 (1978).
2. van der Pol, B. & van der Mark, J. Frequency demultiplication. *Nature* **120,** 363–364 (1927).
3. Goroff, D. (ed.) in *H. Poincaré New Methods of Celestial Mechanics, 1: Periodic And Asymptotic Solutions* (American Institute of Physics, 1991).
4. Goroff, D. (ed.) *H. Poincaré New Methods Of Celestial Mechanics, 2: Approximations by Series* (American Institute of Physics, 1993).
5. Goroff, D. (ed.) in *H. Poincaré New Methods Of Celestial Mechanics, 3: Integral Invariants and Asymptotic Properties of Certain Solutions* (American Institute of Physics, 1993).
6. Crutchfield, J. P., Packard, N. H., Farmer, J. D. & Shaw, R. S. Chaos. *Sci. Am.* **255,** 46–57 (1986).
7. Binney, J. J., Dowrick, N. J., Fisher, A. J. & Newman, M. E. J. *The Theory of Critical Phenomena* (Oxford Univ. Press, 1992).
8. Cross, M. C. & Hohenberg, P. C. Pattern formation outside of equilibrium. *Rev. Mod. Phys.* **65,** 851–1112 (1993).
9. Manneville, P. *Dissipative Structures and Weak Turbulence* (Academic, 1990).
10. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27** 379–423; 623–656 (1948).
11. Cover, T. M. & Thomas, J. A. *Elements of Information Theory* 2nd edn (Wiley–Interscience, 2006).
12. Kolmogorov, A. N. Entropy per unit time as a metric invariant of automorphisms. *Dokl. Akad. Nauk. SSSR* **124,** 754–755 (1959).
13. Sinai, Ja G. On the notion of entropy of a dynamical system. *Dokl. Akad. Nauk. SSSR* **124,** 768–771 (1959).
14. Anderson, P. W. More is different. *Science* **177,** 393–396 (1972).
15. Turing, A. M. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc. 2* **42,** 230–265 (1936).
16. Solomonoff, R. J. A formal theory of inductive inference: Part I. *Inform. Control* **7,** 1–24 (1964).
17. Solomonoff, R. J. A formal theory of inductive inference: Part II. *Inform. Control* **7,** 224–254 (1964).
18. Minsky, M. L. in *Problems in the Biological Sciences* Vol. XIV (ed. Bellman, R. E.) (Proceedings of Symposia in Applied Mathematics, American Mathematical Society, 1962).
19. Chaitin, G. On the length of programs for computing finite binary sequences. *J. ACM* **13,** 145–159 (1966).
20. Kolmogorov, A. N. Three approaches to the concept of the amount of information. *Probab. Inform. Trans.* **1,** 1–7 (1965).
21. Martin-Löf, P. The definition of random sequences. *Inform. Control* **9,** 602–619 (1966).
22. Brudno, A. A. Entropy and the complexity of the trajectories of a dynamical system. *Trans. Moscow Math. Soc.* **44,** 127–151 (1983).
23. Zvonkin, A. K. & Levin, L. A. The complexity of finite objects and the development of the concepts of information and randomness by means of the theory of algorithms. *Russ. Math. Survey* **25,** 83–124 (1970).
24. Chaitin, G. *Algorithmic Information Theory* (Cambridge Univ. Press, 1987).
25. Li, M. & Vitanyi, P. M. B. *An Introduction to Kolmogorov Complexity and its Applications* (Springer, 1993).
26. Rissanen, J. Universal coding, information, prediction, and estimation. *IEEE Trans. Inform. Theory* **IT-30,** 629–636 (1984).
27. Rissanen, J. Complexity of strings in the class of Markov sources. *IEEE Trans. Inform. Theory* **IT-32,** 526–532 (1986).
28. Blum, L., Shub, M. & Smale, S. On a theory of computation over the real numbers: NP-completeness, Recursive Functions and Universal Machines. *Bull. Am. Math. Soc.* **21,** 1–46 (1989).
29. Moore, C. Recursion theory on the reals and continuous-time computation. *Theor. Comput. Sci.* **162,** 23–44 (1996).
30. Shannon, C. E. Communication theory of secrecy systems. *Bell Syst. Tech. J.* **28,** 656–715 (1949).
31. Ruelle, D. & Takens, F. On the nature of turbulence. *Comm. Math. Phys.* **20,** 167–192 (1974).
32. Packard, N. H., Crutchfield, J. P., Farmer, J. D. & Shaw, R. S. Geometry from a time series. *Phys. Rev. Lett.* **45,** 712–716 (1980).
33. Takens, F. in *Symposium on Dynamical Systems and Turbulence,* Vol. 898 (eds Rand, D. A. & Young, L. S.) 366–381 (Springer, 1981).
34. Brandstater, A. *et al.* Low-dimensional chaos in a hydrodynamic system. *Phys. Rev. Lett.* **51,** 1442–1445 (1983).
35. Crutchfield, J. P. & McNamara, B. S. Equations of motion from a data series. *Complex Syst.* **1,** 417–452 (1987).
36. Crutchfield, J. P. & Young, K. Inferring statistical complexity. *Phys. Rev. Lett.* **63,** 105–108 (1989).

37. Crutchfield, J. P. & Shalizi, C. R. Thermodynamic depth of causal states: Objective complexity via minimal representations. *Phys. Rev. E* **59,** 275–283 (1999).
38. Shalizi, C. R. & Crutchfield, J. P. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.* **104,** 817–879 (2001).
39. Young, K. *The Grammar and Statistical Mechanics of Complex Physical Systems*. PhD thesis, Univ. California (1991).
40. Koppel, M. Complexity, depth, and sophistication. *Complexity* **1,** 1087–1091 (1987).
41. Koppel, M. & Atlan, H. An almost machine-independent theory of program-length complexity, sophistication, and induction. *Information Sciences* **56,** 23–33 (1991).
42. Crutchfield, J. P. & Young, K. in *Entropy, Complexity, and the Physics of Information* Vol. VIII (ed. Zurek, W.) 223–269 (SFI Studies in the Sciences of Complexity, Addison-Wesley, 1990).
43. William of Ockham *Philosophical Writings: A Selection, Translated, with an Introduction* (ed. Philotheus Boehner, O. F. M.) (Bobbs-Merrill, 1964).
44. Farmer, J. D. Information dimension and the probabilistic structure of chaos. *Z. Naturf.* **37a,** 1304–1325 (1982).
45. Crutchfield, J. P. The calculi of emergence: Computation, dynamics, and induction. *Physica D* **75,** 11–54 (1994).
46. Crutchfield, J. P. in *Complexity: Metaphors, Models, and Reality* Vol. XIX (eds Cowan, G., Pines, D. & Melzner, D.) 479–497 (Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, 1994).
47. Crutchfield, J. P. & Feldman, D. P. Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos* **13,** 25–54 (2003).
48. Mahoney, J. R., Ellison, C. J., James, R. G. & Crutchfield, J. P. How hidden are hidden processes? A primer on crypticity and entropy convergence. *Chaos* **21,** 037112 (2011).
49. Ellison, C. J., Mahoney, J. R., James, R. G., Crutchfield, J. P. & Reichardt, J. Information symmetries in irreversible processes. *Chaos* **21,** 037107 (2011).
50. Crutchfield, J. P. in *Nonlinear Modeling and Forecasting* Vol. XII (eds Casdagli, M. & Eubank, S.) 317–359 (Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, 1992).
51. Crutchfield, J. P., Ellison, C. J. & Mahoney, J. R. Time's barbed arrow: Irreversibility, crypticity, and stored information. *Phys. Rev. Lett.* **103,** 094101 (2009).
52. Ellison, C. J., Mahoney, J. R. & Crutchfield, J. P. Prediction, retrodiction, and the amount of information stored in the present. *J. Stat. Phys.* **136,** 1005–1034 (2009).
53. Ruelle, D. Do turbulent crystals exist? *Physica A* **113,** 619–623 (1982).
54. Varn, D. P., Canright, G. S. & Crutchfield, J. P. Discovering planar disorder in close-packed structures from X-ray diffraction: Beyond the fault model. *Phys. Rev. B* **66,** 174110 (2002).
55. Varn, D. P. & Crutchfield, J. P. From finite to infinite range order via annealing: The causal architecture of deformation faulting in annealed close-packed crystals. *Phys. Lett. A* **234,** 299–307 (2004).
56. Varn, D. P., Canright, G. S. & Crutchfield, J. P. *Inferring Pattern and Disorder in Close-Packed Structures from X-ray Diffraction Studies, Part I: ϵ-machine Spectral Reconstruction Theory* Santa Fe Institute Working Paper 03-03-021 (2002).
57. Varn, D. P., Canright, G. S. & Crutchfield, J. P. Inferring pattern and disorder in close-packed structures via ϵ-machine reconstruction theory: Structure and intrinsic computation in Zinc Sulphide. *Acta Cryst. B* **63,** 169–182 (2002).
58. Welberry, T. R. Diffuse x-ray scattering and models of disorder. *Rep. Prog. Phys.* **48,** 1543–1593 (1985).
59. Guinier, A. *X-Ray Diffraction in Crystals, Imperfect Crystals and Amorphous Bodies* (W. H. Freeman, 1963).
60. Sebastian, M. T. & Krishna, P. *Random, Non-Random and Periodic Faulting in Crystals* (Gordon and Breach Science Publishers, 1994).
61. Feldman, D. P., McTague, C. S. & Crutchfield, J. P. The organization of intrinsic computation: Complexity-entropy diagrams and the diversity of natural information processing. *Chaos* **18,** 043106 (2008).
62. Mitchell, M., Hraber, P. & Crutchfield, J. P. Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Syst.* **7,** 89–130 (1993).
63. Johnson, B. D., Crutchfield, J. P., Ellison, C. J. & McTague, C. S. *Enumerating Finitary Processes* Santa Fe Institute Working Paper 10-11-027 (2010).
64. Lind, D. & Marcus, B. *An Introduction to Symbolic Dynamics and Coding* (Cambridge Univ. Press, 1995).
65. Hopcroft, J. E. & Ullman, J. D. *Introduction to Automata Theory, Languages, and Computation* (Addison-Wesley, 1979).
66. Upper, D. R. *Theory and Algorithms for Hidden Markov Models and Generalized Hidden Markov Models*. Ph.D. thesis, Univ. California (1997).
67. Kelly, D., Dillingham, M., Hudson, A. & Wiesner, K. Inferring hidden Markov models from noisy time sequences: A method to alleviate degeneracy in molecular dynamics. Preprint at http://arxiv.org/abs/1011.2969 (2010).
68. Ryabov, V. & Nerukh, D. Computational mechanics of molecular systems: Quantifying high-dimensional dynamics by distribution of Poincaré recurrence times. *Chaos* **21,** 037113 (2011).
69. Li, C-B., Yang, H. & Komatsuzaki, T. Multiscale complex network of protein conformational fluctuations in single-molecule time series. *Proc. Natl Acad. Sci. USA* **105,** 536–541 (2008).
70. Crutchfield, J. P. & Wiesner, K. Intrinsic quantum computation. *Phys. Lett. A* **372,** 375–380 (2006).
71. Goncalves, W. M., Pinto, R. D., Sartorelli, J. C. & de Oliveira, M. J. Inferring statistical complexity in the dripping faucet experiment. *Physica A* **257,** 385–389 (1998).
72. Clarke, R. W., Freeman, M. P. & Watkins, N. W. The application of computational mechanics to the analysis of geomagnetic data. *Phys. Rev. E* **67,** 160–203 (2003).
73. Crutchfield, J. P. & Hanson, J. E. Turbulent pattern bases for cellular automata. *Physica D* **69,** 279–301 (1993).
74. Hanson, J. E. & Crutchfield, J. P. Computational mechanics of cellular automata: An example. *Physica D* **103,** 169–189 (1997).
75. Shalizi, C. R., Shalizi, K. L. & Haslinger, R. Quantifying self-organization with optimal predictors. *Phys. Rev. Lett.* **93,** 118701 (2004).
76. Crutchfield, J. P. & Feldman, D. P. Statistical complexity of simple one-dimensional spin systems. *Phys. Rev. E* **55,** 239R–1243R (1997).
77. Feldman, D. P. & Crutchfield, J. P. Structural information in two-dimensional patterns: Entropy convergence and excess entropy. *Phys. Rev. E* **67,** 051103 (2003).
78. Bonner, J. T. *The Evolution of Complexity by Means of Natural Selection* (Princeton Univ. Press, 1988).
79. Eigen, M. Natural selection: A phase transition? *Biophys. Chem.* **85,** 101–123 (2000).
80. Adami, C. What is complexity? *BioEssays* **24,** 1085–1094 (2002).
81. Frenken, K. *Innovation, Evolution and Complexity Theory* (Edward Elgar Publishing, 2005).
82. McShea, D. W. The evolution of complexity without natural selection—A possible large-scale trend of the fourth kind. *Paleobiology* **31,** 146–156 (2005).
83. Krakauer, D. Darwinian demons, evolutionary complexity, and information maximization. *Chaos* **21,** 037111 (2011).
84. Tononi, G., Edelman, G. M. & Sporns, O. Complexity and coherency: Integrating information in the brain. *Trends Cogn. Sci.* **2,** 474–484 (1998).
85. Bialek, W., Nemenman, I. & Tishby, N. Predictability, complexity, and learning. *Neural Comput.* **13,** 2409–2463 (2001).
86. Sporns, O., Chialvo, D. R., Kaiser, M. & Hilgetag, C. C. Organization, development, and function of complex brain networks. *Trends Cogn. Sci.* **8,** 418–425 (2004).
87. Crutchfield, J. P. & Mitchell, M. The evolution of emergent computation. *Proc. Natl Acad. Sci. USA* **92,** 10742–10746 (1995).
88. Lizier, J., Prokopenko, M. & Zomaya, A. Information modification and particle collisions in distributed computation. *Chaos* **20,** 037109 (2010).
89. Flecker, B., Alford, W., Beggs, J. M., Williams, P. L. & Beer, R. D. Partial information decomposition as a spatiotemporal filter. *Chaos* **21,** 037104 (2011).
90. Rissanen, J. *Stochastic Complexity in Statistical Inquiry* (World Scientific, 1989).
91. Balasubramanian, V. Statistical inference, Occam's razor, and statistical mechanics on the space of probability distributions. *Neural Comput.* **9,** 349–368 (1997).
92. Glymour, C. & Cooper, G. F. (eds) in *Computation, Causation, and Discovery* (AAAI Press, 1999).
93. Shalizi, C. R., Shalizi, K. L. & Crutchfield, J. P. *Pattern Discovery in Time Series, Part I: Theory, Algorithm, Analysis, and Convergence* Santa Fe Institute Working Paper 02-10-060 (2002).
94. MacKay, D. J. C. *Information Theory, Inference and Learning Algorithms* (Cambridge Univ. Press, 2003).
95. Still, S., Crutchfield, J. P. & Ellison, C. J. Optimal causal inference. *Chaos* **20,** 037111 (2007).
96. Wheeler, J. A. in *Entropy, Complexity, and the Physics of Information* volume VIII (ed. Zurek, W.) (SFI Studies in the Sciences of Complexity, Addison-Wesley, 1990).

Acknowledgements

I thank the Santa Fe Institute and the Redwood Center for Theoretical Neuroscience, University of California Berkeley, for their hospitality during a sabbatical visit.

Additional information

The author declares no competing financial interests. Reprints and permissions information is available online at http://www.nature.com/reprints.

nature physics

INSIGHT | REVIEW ARTICLES

PUBLISHED ONLINE: 22 DECEMBER 2011 | DOI: 10.1038/NPHYS2162

Communities, modules and large-scale structure in networks

M. E. J. Newman

Networks, also called graphs by mathematicians, provide a useful abstraction of the structure of many complex systems, ranging from social systems and computer networks to biological networks and the state spaces of physical systems. In the past decade there have been significant advances in experiments to determine the topological structure of networked systems, but there remain substantial challenges in extracting scientific understanding from the large quantities of data produced by the experiments. A variety of basic measures and metrics are available that can tell us about small-scale structure in networks, such as correlations, connections and recurrent patterns, but it is considerably more difficult to quantify structure on medium and large scales, to understand the 'big picture'. Important progress has been made, however, within the past few years, a selection of which is reviewed here.

A network is, in its simplest form, a collection of dots joined together in pairs by lines (Fig. 1). In the jargon of the field, a dot is called a 'node' or 'vertex' (plural 'vertices') and a line is called an 'edge'. Networks are used in many branches of science as a way to represent the patterns of connections between the components of complex systems[1–6]. Examples include the Internet[7,8], in which the nodes are computers and the edges are data connections such as optical-fibre cables, food webs in biology[9,10], in which the nodes are species in an ecosystem and the edges represent predator–prey interactions, and social networks[11,12], in which the nodes are people and the edges represent any of a variety of different types of social interaction including friendship, collaboration, business relationships or others.

In the past decade there has been a surge of interest in both empirical studies of networks[13] and development of mathematical and computational tools for extracting insight from network data[1–6]. One common approach to the study of networks is to focus on the properties of individual nodes or small groups of nodes, asking questions such as, 'Which is the most important node in this network?' or 'Which are the strongest connections?' Such approaches, however, tell us little about large-scale network structure. It is this large-scale structure that is the topic of this paper.

The best-studied form of large-scale structure in networks is modular or community structure[14,15]. A community, in this context, is a dense subnetwork within a larger network, such as a close-knit group of friends in a social network or a group of interlinked web pages on the World Wide Web (Fig. 1). Although communities are not the only interesting form of large-scale structure—there are others that we will come to—they serve as a good illustration of the nature and scope of present research in this area and will be our primary focus.

Communities are of interest for a number of reasons. They have intrinsic interest because they may correspond to functional units within a networked system, an example of the kind of link between structure and function that drives much of the present excitement about networks. In a metabolic network[16], for instance—the network of chemical reactions within a cell—a community might correspond to a circuit, pathway or motif that carries out a certain function, such as synthesizing or regulating a vital chemical product[17]. In a social network, a community might correspond to an actual community in the conventional sense of the word, a group of people brought together by a common interest, a common location or workplace or family ties[18].

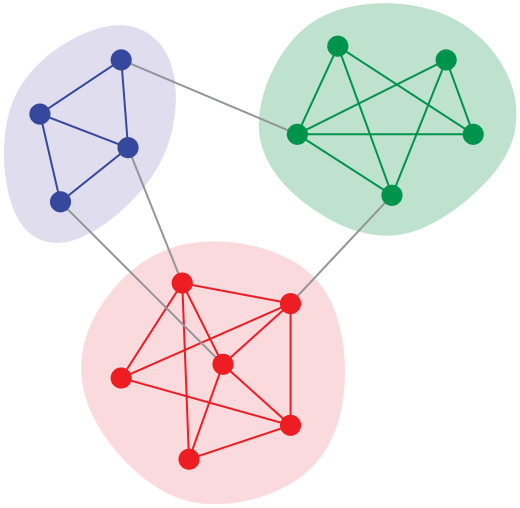

Figure 1 | Example network showing community structure. The nodes of this network are divided into three groups, with most connections falling within groups and only a few between groups.

However, there is another reason, less often emphasized, why a knowledge of community structure can be useful. In many networks it is found that the properties of individual communities can be quite different. Consider, for example, Fig. 2, which shows a network of collaborations among a group of scientists at a research institute. The network divides into distinct communities as indicated by the colours of the nodes. (We will see shortly how this division is accomplished.) In this case, the communities correspond closely to the acknowledged research groups within the institute, a demonstration that indeed the discovery of communities can point to functional divisions in a system. However, notice also that the structural features of the different communities are widely varying. The communities highlighted in red and light blue, for instance, appear to be loose-knit groups of collaborators working together in various combinations, whereas the groups in yellow and dark blue are both organized around a central hub, perhaps a group

Department of Physics and Center for the Study of Complex Systems, University of Michigan, Ann Arbor, Michigan 48109, USA. e-mail: mejn@umich.edu.

NATURE PHYSICS DOI:10.1038/NPHYS2162

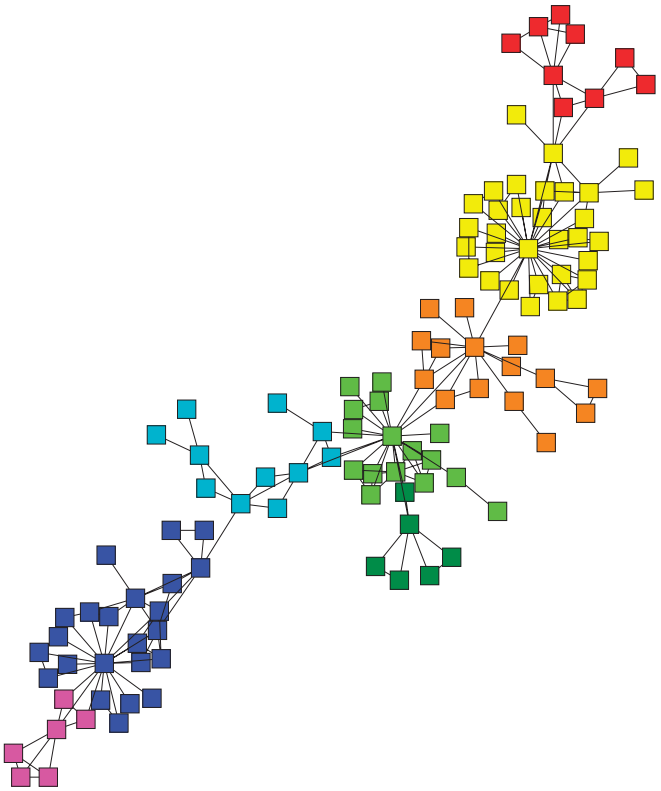

Figure 2 | A network of collaborations among scientists at a research institute. Nodes in this network represent the scientists and there is an edge between any pair of scientists who co-authored a published paper during the years of the study. Colours represent communities, as discovered using a modularity-maximization technique.

leader or principal investigator of some kind. Distinctions such as these, which may be crucial for understanding the behaviour of the system, become apparent only when one looks at structure on the community level.

The network in this particular example has the nice property that it is small enough and sparse enough to be drawn clearly on the page. One does not need any calculations to pick out the communities in this case: a good eye will do the job. However, when we are working with larger or denser networks, networks that can have thousands or even millions of nodes (or a smaller number of nodes but very many edges), clear visualization becomes impossible and we must turn instead to algorithmic methods for community detection and the development of such methods has been a highly active area of research in the past few years[15].

The community-detection problem is challenging in part because it is not very well posed. It is agreed that the basic problem is to find locally dense regions in a network, but this is not a precise formulation. If one is to create a method for detecting communities in a mechanical way, one must first define exactly what one means by a community. Researchers have been aware of this issue from the outset and have proposed a wide variety of definitions, based on counts of edges within and between communities, counts of paths across networks, spectral properties of network matrices, information-theoretic measures, random walks and many other quantities. With this array of definitions comes a corresponding array of algorithms that seek to find the communities so defined[14,15,19–31]. Unfortunately, it is no easy matter to determine which of these algorithms are the best, because the perception of good performance itself depends on how one defines a community and each algorithm is necessarily good at finding communities according to its own definition. To get around this circularity, we typically take one of two approaches. In the first, algorithms are tested against real-world networks for which there is an accepted division into communities, often based on additional measurements that are independent of the network itself, such as interviews with participants in a social network or analysis of the text of web pages. If an algorithm can reliably find the accepted structure then it is considered successful. In the second approach, algorithms are tested against computer-generated networks that have some form of community structure artificially embedded within them. A number of standard benchmark networks have been proposed for this purpose, such as the 'four groups' networks[14] or so-called the LFR benchmark networks[32]. A number of studies have been published that compare the performance of proposed algorithms in these benchmark tests[33,34]. Although these approaches do set concrete targets for performance of community-detection methods, there is room for debate over whether those targets necessarily align with good performance in broader real-world situations. If we tune our algorithms to solve specific benchmark problems we run the risk of creating algorithms that solve those problems well but other (perhaps more realistic) problems poorly.

This is a crucial issue and one that is worth bearing in mind as we take a look in the following sections at the present state of research on community detection. As we will see, however, researchers have, in spite of the difficulties, come up with a range of approaches that return real, useful information about the large-scale structure of networks, and in the process have learned much, both about individual networks that have been analysed and about mathematical methods for representing and understanding network structure.

Hierarchical clustering

Studies of communities in networks go back at least to the 1970s, when a number of techniques were developed for their detection, particularly in computer science and sociology. In computer science the problem of graph partitioning[35], which is similar but not identical to the problem of community detection, has received attention for its engineering applications, but the methods developed, such as spectral partitioning[36] and the Kernighan–Lin algorithm[37], have also been fruitfully applied in other areas. However, it is the work of sociologists that is perhaps the most direct ancestor of modern techniques of community detection.

An early, and still widely used, technique for detecting communities in social networks is hierarchical clustering[5,11]. Hierarchical clustering is in fact not a single technique but an entire family of techniques, with a single central principle: if we can derive a measure of how strongly nodes in a network are connected together, then by grouping the most strongly connected we can divide the network into communities. Specific hierarchical clustering methods differ on the particular measure of strength used and on the rules by which we group strongly connected nodes. Most common among the measures used are the so-called structural equivalence measures, which focus on the number n_{ij} of common network neighbours that two nodes i, j have. In a social network of friendships, for example, two people with many mutual friends are more likely to be close than two people with few and thus a count of mutual friends can be used as a measure of connection strength. Rather than using the raw count n_{ij}, however, one typically normalizes it in some way, leading to measures such as the Jaccard coefficient and cosine similarity. For example, the cosine similarity σ_{ij} between nodes i and j is defined by

$$\sigma_{ij} = \frac{n_{ij}}{\sqrt{k_i k_j}}$$

where k_i is the degree of node i (that is, the number of connections it has). This measure has the nice property that its

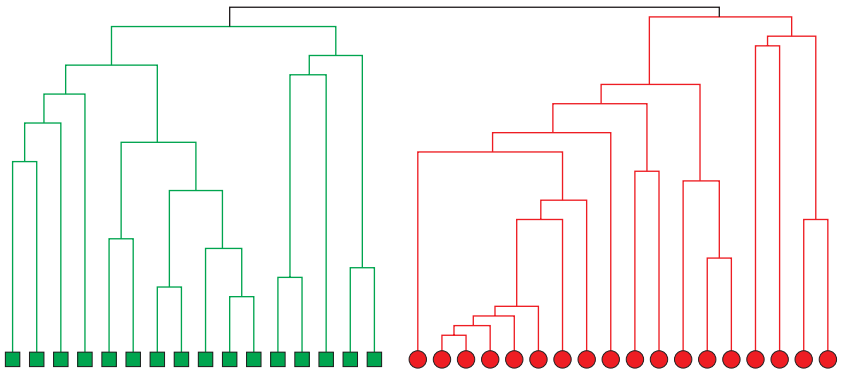

Figure 3 | Average-linkage clustering of a small social network. This tree or 'dendrogram' shows the results of the application of average-linkage hierarchical clustering using cosine similarity to the well-known karate-club network of Zachary[38], which represents friendship between members of a university sports club. The calculation finds two principal communities in this case (the left and right subtrees of the dendrogram), which correspond exactly to known factions within the club (represented by the colours).

value falls always between zero and one—zero if the nodes have no common neighbours and one if they have all their neighbours in common.

Once one has defined a measure of connection strength, one can begin to group nodes together, which is done in hierarchical fashion, first grouping single nodes into small groups, then grouping those groups into larger groups and so forth. There are a number of methods by which this grouping can be carried out, the three common ones being the methods known as single-linkage, complete-linkage and average-linkage clustering. Single-linkage clustering is the most widely used by far, primarily because it is simple to implement, but in fact average-linkage clustering generally gives superior results and is not much harder to implement.

Figure 3 shows the result of applying average-linkage hierarchical clustering based on cosine similarity to a famous network from the social networks literature, Zachary's karate-club network[38]. This network represents patterns of friendship between members of a karate club at a US university, compiled from observations and interviews of the club's 34 members. The network is of particular interest because during the study a dispute arose among the club's members over whether to raise club fees. Unable to reconcile their differences, the members of the club split into two factions, with one faction departing to start a separate club. It has been claimed repeatedly that by examining the pattern of friendships depicted in the network (which was compiled before the split happened) one can predict the membership of the two factions[14,20,26,27,38–40].

Figure 3 shows the output of the hierarchical clustering procedure in the form of a tree or 'dendrogram' representing the order in which nodes are grouped together into communities. It should be read from the bottom up: at the bottom we have individual nodes that are grouped first into pairs, and then into larger groups as we move up the tree, until we reach the top, where all nodes have been gathered into one group. In a single image, this dendrogram captures the entire hierarchical clustering process. Horizontal cuts through the figure represent the groups at intermediate stages.

As we can see, the method in this case joins the nodes together into two large groups, consisting of roughly half the network each, before finally joining those two into one group at the top of the dendrogram. It turns out that these two groups correspond precisely to the groups into which the club split in real life, which are indicated by the colours in the figure. Thus, in this case the method works well. It has effectively predicted a future social phenomenon, the split of the club, from quantitative data measured before the split occurred. It is the promise of outcomes such as this that drives much of the present interest in networks.

Hierarchical clustering is straightforward to understand and to implement, but it does not always give satisfactory results. As it exists in many variants (different strength measures and different linkage rules) and different variants give different results, it is not clear which results are the 'correct' ones. Moreover, the method has a tendency to group together those nodes with the strongest connections but leave out those with weaker connections, so that the divisions it generates may not be clean divisions into groups, but rather consist of a few dense cores surrounded by a periphery of unattached nodes. Ideally, we would like a more reliable method.

Optimization methods

Over the past decade or so, researchers in physics and applied mathematics have taken an active interest in the community-detection problem and introduced a number of fruitful approaches. Among the first proposals were approaches based on a measure known as betweenness[14,21,41], in which one calculates one of several measures of the flow of (imaginary) traffic across the edges of a network and then removes from the network those edges with the most traffic. Two other related approaches are the use of fluid-flow[19] and current-flow analogies[42] to identify edges for removal; the latter idea has been revived recently to study structure in the very largest networks[30]. A different class of methods are those based on information-theoretic ideas, such as the minimum-description-length methods of Rosvall and Bergstrom[26,43] and related methods based on statistical inference, such as the message-passing method of Hastings[25]. Another large class exploits links between community structure and processes taking place on networks, such as random walks[44,45], Potts models[46] or oscillator synchronization[47]. A contrasting set of approaches focuses on the detection of 'local communities'[23,24] and seeks to answer the question of whether we can, given a single node, identify the community to which it belongs, without first finding all communities in the network. In addition to being useful for studying limited portions of larger networks, this approach can give rise to overlapping communities, in which a node can belong to more than one community. (The generalized community-detection problem in which overlaps are allowed in this way has been an area of increasing interest within the field in recent years[22,31].)

However, the methods most heavily studied by physicists, perhaps unsurprisingly, are those that view the community-detection problem by analogy with equilibrium physical processes and treat it as an optimization task. The basic idea is to define a quantity that is high for 'good' divisions of a network and low for 'bad' ones, and then to search through possible divisions for the one with the highest score. This approach is similar to the minimization

of energy when finding the ground state or stable state of a physical system, and the connection has been widely exploited. A variety of different measures for assigning scores have been proposed, such as the so-called E/I ratio[48], likelihood-based measures[49] and others[50], but the most widely used is the measure known as the modularity[18,51].

Suppose you are given a network and a candidate division into communities. A simple measure of the quality of that division is the fraction of edges that fall within (rather than between) communities. If this fraction is high then you have a good division (Fig. 1). However, this measure is not ideal. It is maximized by putting all nodes in a single group together, which is a correct but trivial form of community structure and not of particular interest. A better measure is the so-called modularity, which is defined to be the fraction of edges within communities minus the expected value of that fraction if the positions of the edges are randomized[51]. If there are more edges within communities than one would find in a randomized network then the modularity will be positive and large positive values indicate good community divisions.

Let A_{ij} be equal to the number of edges between nodes i and j (normally zero or one); A_{ij} is an element of the 'adjacency matrix' of the network. It can be shown that for a network with m edges in total, the expected number that fall between nodes i and j if the positions of the edges are randomized is given by $k_i k_j/2m$, where k_i is again the degree of node i. Thus, the actual number of edges between i and j minus the expected number is $A_{ij} - k_i k_j/2m$ and the modularity Q is the sum of this quantity over all pairs of nodes that fall in the same community. If we label the communities and define s_i to be the label of the community to which node i belongs, then we can write

$$Q = \frac{1}{2m}\sum_{ij}\left[A_{ij} - \frac{k_i k_j}{2m}\right]\delta_{s_i,s_j}$$

where δ_{ij} is the Kronecker delta and the leading constant $1/2m$ is included only by convention—it normalizes Q to measure fractions of edges rather than total numbers but its presence has no effect on the position of the modularity maximum.

The modularity takes precisely the form $H = -\sum_{ij} J_{ij}\delta_{s_i,s_j}$ of the Hamiltonian of a (disordered) Potts model, apart from a minus sign, and hence its maximization is equivalent to finding the ground state of the Potts model—the community assignments s_i act similarly to spins on the nodes of the network. Unfortunately, direct optimization of the modularity by an exhaustive search through the possible spin states is intractable for any but the smallest of networks, and faster indirect (but exact) algorithms have been proved rigorously not to exist[52]. A variety of approximate techniques from physics and elsewhere, however, are applicable to the problem and seem to give good, but not perfect, solutions with relatively modest computational effort. These include simulated annealing[17,53], greedy algorithms[54,55], semidefinite programming[28], spectral methods[56] and several others[40,57]. Modularity maximization forms the basis for other more complex approaches as well, such as the method of Blondel *et al.*[27], a multiscale method in which modularity is first optimized using a greedy local algorithm, then a 'supernetwork' is formed whose nodes represent the communities so discovered and the greedy algorithm is repeated on this supernetwork. The process iterates until no further improvements in modularity are possible. This method has become widely used by virtue of its relative computational efficiency and the high quality of the results it returns. In a recent comparative study it was found to be one of the best available algorithms when tested against computer-generated benchmark problems of the type described in the introduction[34].

Figure 2, showing collaboration patterns among scientists, is an example of community detection using modularity maximization. One of the nice features of the modularity method is that one does not need to know in advance the number of communities contained in the network: a free maximization of the modularity, in which the number of communities is allowed to vary, will tell us the most advantageous number, as well as finding the exact division of the nodes among communities.

Although modularity maximization is efficient, widely used and gives informative results, it—like hierarchical clustering—has deficiencies. In particular, it has a known bias in the size of the communities it finds—it has a preference for communities of size roughly equal to the square root of the size of the network[58]. Modifications of the method have been proposed that allow one to vary this preferred size[59,60], but not to eliminate the preference altogether. The modularity method also ignores any information stored in the positions of edges that run between communities: as modularity is calculated by counting only within-group edges, one could move the between-group edges around in any way one pleased and the value of the modularity would not change at all. One might imagine that one could do a better job of detecting communities if one were to make use of the information represented by these edges.

In the past few years, therefore, researchers have started to look for a more principled approach to community detection, and have gravitated towards the method of block modelling, a method that traces its roots back to the 1970s (refs 61,62), but which has recently enjoyed renewed popularity, with some powerful new methods and results emerging.

Block models

Block modelling[63–67] is in effect a form of statistical inference for networks. In the same way that we can gain some understanding from conventional numerical data by fitting, say, a straight line through data points, so we can gain understanding of the structure of networks by fitting them to a statistical network model. In particular, if we are interested in community structure then we can create a model of networks that contain such structure, then fit it to an observed network and in the process learn about community structure in that observed network, if it exists.

A simple example of a block model is a model network in which one has a certain number n of nodes and each node is assigned to one of several labelled groups or communities. In addition, one specifies a set of probabilities p_{rs}, which represent the probability that there will be an edge between a node in group r and a node in group s. This model can be used, for instance, in a generative process to create a random network with community structure. By making the edge probabilities higher for pairs of nodes in the same group and lower for pairs in different groups, then generating a set of edges independently with exactly those probabilities, one can produce an artificial network that has many edges within groups and few between them—the classic community structure.

However, we can also turn the experiment around and ask, 'If we observe a real network and we suppose that it was generated by this model, what would the values of the model's parameters have to be?' More precisely, what values of the parameters are most likely to have generated the network we see in real life? This leads us to a 'maximum likelihood' formulation of the community-detection problem. The probability, or likelihood, that an observed network was generated by this block model is given by

$$L = \prod_{i<j} p_{s_i s_j}^{A_{ij}} (1 - p_{s_i s_j})^{1 - A_{ij}}$$

where A_{ij} is an element of the adjacency matrix, as before, and s_i is again the community to which node i belongs. Now

we simply maximize this quantity over the probabilities p_{rs} and the communities s_i. Again we have turned the detection of communities into an optimization problem, albeit a harder one than the modularity-maximization problem. The values of the probabilities p_{rs} are usually of lesser interest to us, but if we can find the community parameters s_i that maximize the likelihood then we have solved our community-detection problem.

Although it seems elegant and well-founded in principle, the surprising thing about this approach, at least as we have described it here, is that it does not work well. Figure 4a shows an example application of (a slight variant of) the method to a network of weblogs, or 'blogs'—personal web pages maintained by individuals or groups, on which they publish their thoughts on topics of their choosing. This particular network, which was assembled by Adamic and Glance[68], is composed of blogs about US politics that were active around the time of the US presidential election in 2004, and the edges in the network represent web hyperlinks between blogs. Adamic and Glance showed that this network was strongly divided into two communities, one of left-leaning (that is, liberal) blogs, which commonly link to one another, and the other of right-leaning (conservative) ones, which also link to one another, but that there were few links between left and right. The communities appear as roughly the left and right halves of the network as it is drawn in Fig. 4a. The colours in the figure show the division of the network into two communities found with the maximum likelihood method above, and it is clear that the method has failed to find the known division in this case. What has gone wrong?

On closer inspection, we find that the method fails in this case because it does not take into account the wide variation among the degrees of nodes in the network. In this network (and many others) degrees vary over a great range, whereas degrees in the block model are Poisson distributed and narrowly peaked about their mean. This means, in effect, that there is no choice of parameters for the model that gives a good fit to the data. Fitting this block model is similar to fitting a straight line through an inherently curved set of data points—you can do it, but it is unlikely to give you a meaningful answer.

It turns out, however, that one can fix such problems by suitably modifying the model. Figure 4b shows a different fit to the same network using now a 'degree-corrected' block model that allows for widely varying degrees[49]. As the figure shows, the model now finds a division that corresponds closely to the known division between left- and right-leaning blogs. The moral of the story is that it is not hard to come up with models so unrealistic that they will not fit the observed network for any parameter values and one must guard against this possibility if the method is to work.

Once we deal with this issue, however, the block-model method has some promising features. If we have found the parameter values for the best fit of the model to an observed network, we can then plug those values back into the model and use the model to generate further networks that are similar to the original network, but not identical. This ability to generate similar networks can be used, for instance, to guess at the locations of possible missing edges in a network. For many networks our data are incomplete or unreliable, and there may be edges missing from the recorded structure. Looking at a large selection of generated networks that are similar to the original, one can find edges that appear often in the generated networks but not in the original; such edges turn out to be reliable candidates for missing data. Guimerá and Sales-Pardo[69] have shown that this approach is at least as accurate as, and often better than, previous methods for predicting missing edges.

Another nice feature of the block-model method is that it lends itself to many variants that are suitable for particular types of problem. For instance, in some problems we can, with some effort, carry out experiments to determine the community membership of a few nodes, and the goal is to determine the rest. In recent work, Yan *et al.*[70] have devised a variant of the block-model method in which one can use the model to determine on which nodes these experiments should be done, by looking for the nodes whose membership information will be most useful, in the sense that it will tell us as much as possible not only about the measured nodes but also about the membership of other nodes in the network. They show that the accuracy of community detection can be enormously improved by carrying out just a few experiments on nodes carefully chosen using this technique.

However, perhaps the most promising feature of the block-model method is that it is not limited to detecting traditional community structure in networks. In principle, any type of structure that can be formulated as a probabilistic model can be detected, including overlapping communities, bipartite or k-partite

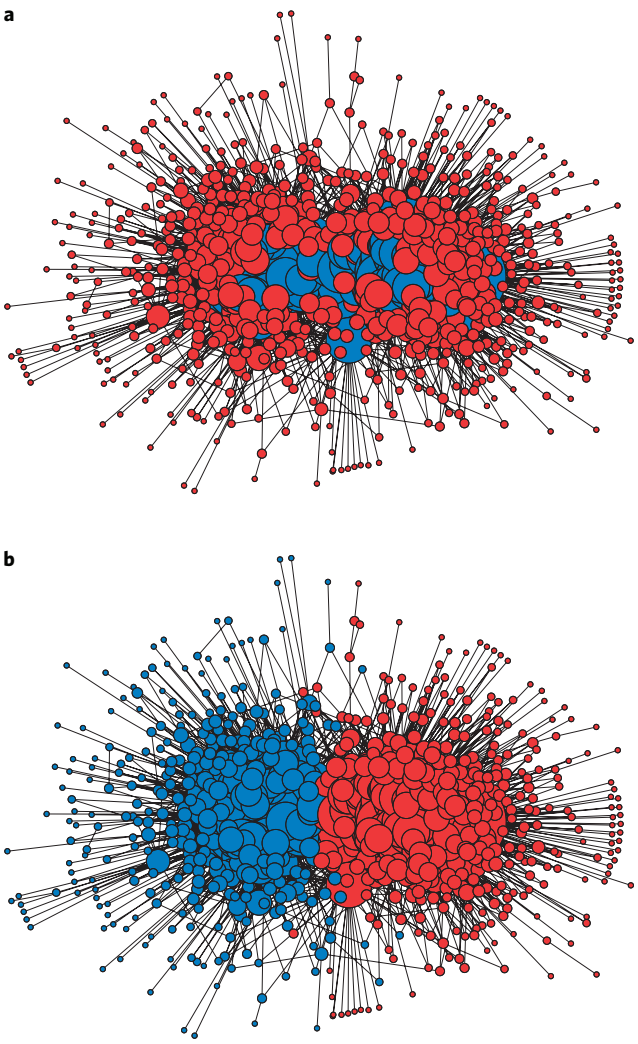

Figure 4 | Analysis of a network of links between web sites about US politics. The two panels represent the divisions found in a network of political weblogs using two different versions of the block model method. **a**, Division into two communities discovered using a fit to the basic block model described in the text, which fails to find the acknowledged division of the network into politically left- and right-leaning communities. **b**, Division using a block model that corrects for the broad distribution of node degrees in the network. This division corresponds closely to the acknowledged one. Figure reproduced with permission from ref. 49, © 2011 APS. Network data taken from ref. 68.

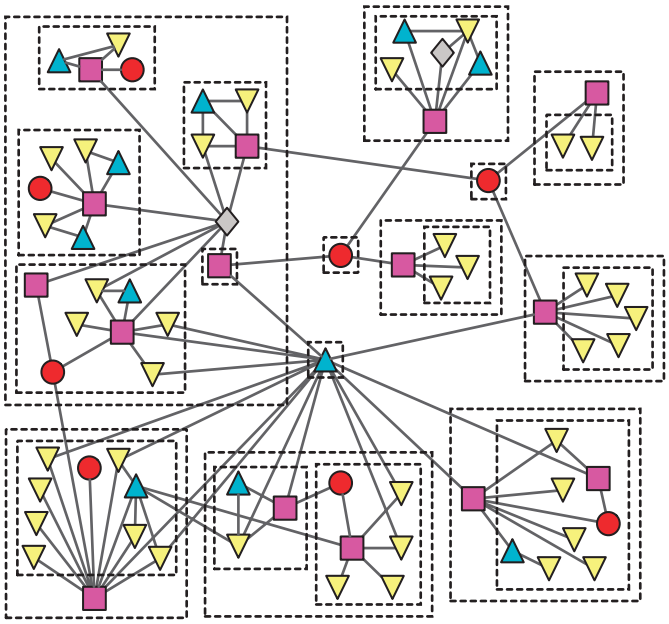

Figure 5 | Hierarchical divisions in a food web of grassland species. Outlined sets of nodes represent groups of species at different levels in the hierarchy. For clarity only two levels in the hierarchy are shown, although five levels were found in some parts of the network. Reproduced from ref. 71.

structures, communities within communities and many others. The field is only just beginning to explore the wide range of possibilities that this approach offers, but Fig. 5 shows one example, drawn from my own work[71]. In this study we examined the food web of a grassland ecosystem—the network of predator–prey interactions between species—and searched for a generalized form of hierarchical community structure in which groups divide into subgroups and subsubgroups and so on. Using a model that employs a tree structure reminiscent of the dendrogram of Fig. 3 to represent the hierarchy of groups, and edge probabilities that depend on shortest paths through the tree, we were able to discover an entire spectrum of structure within the network, spanning the range from small motifs of a few nodes to the size of the entire network. Of particular note in this example is the way in which the method groups host species (squares) with their parasites (yellow triangles), but at the next level in the hierarchy also gathers the parasites separately into their own groups. In some sense, the parasites have more in common with each other than with their host, and hence can be thought of as belonging to a separate group, even though they have no direct interactions with one another through the food web. The calculation realizes this and divides the network accordingly.

Conclusion

The study of network structure and its links with the function and behaviour of complex systems is a large and active field of endeavor, with new results appearing daily and an energetic community of researchers working on both methods and applications. Some of the ideas discussed here are now well established and widely used, whereas others, such as the block-model methods, are being actively researched and developed, and there are many others still that there is not room to describe in this article. The pace of developments is, if anything, accelerating, and the field offers substantial promise for those in physics, biology, the social sciences and elsewhere, for whom the ability to make sense of the structures, large and small, found in networks can open a new window on the behaviour of systems of many kinds.

References

1. Albert, R. & Barabási, A-L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74,** 47–97 (2002).
2. Dorogovtsev, S. N. & Mendes, J. F. F. Evolution of networks. *Adv. Phys.* **51,** 1079–1187 (2002).
3. Newman, M. E. J. The structure and function of complex networks. *SIAM Rev.* **45,** 167–256 (2003).
4. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Hwang, D-U. Complex networks: Structure and dynamics. *Phys. Rep.* **424,** 175–308 (2006).
5. Newman, M. E. J. *Networks: An Introduction* (Oxford Univ. Press, 2010).
6. Cohen, R. & Havlin, S. *Complex Networks: Structure, Stability and Function* (Cambridge Univ. Press, 2010).
7. Faloutsos, M., Faloutsos, P. & Faloutsos, C. On power-law relationships of the internet topology. *Comput. Commun. Rev.* **29,** 251–262 (1999).
8. Pastor-Satorras, R. & Vespignani, A. *Evolution and Structure of the Internet* (Cambridge Univ. Press, 2004).
9. Pimm, S. L. *Food Webs* 2nd edn (Univ. Chicago Press, 2002).
10. Pascual, M. & Dunne, J. A. (eds) *Ecological Networks: Linking Structure to Dynamics in Food Webs* (Oxford Univ. Press, 2006).
11. Wasserman, S. & Faust, K. *Social Network Analysis* (Cambridge Univ. Press, 1994).
12. Scott, J. *Social Network Analysis: A Handbook* 2nd edn (Sage, 2000).
13. Costa, L. da F., Rodrigues, F. A., Travieso, G. & Boas, P. R. V. Characterization of complex networks: A survey of measurements. *Adv. Phys.* **56,** 167–242 (2007).
14. Girvan, M. & Newman, M. E. J. Community structure in social and biological networks. *Proc. Natl Acad. Sci. USA* **99,** 7821–7826 (2002).
15. Fortunato, S. Community detection in graphs. *Phys. Rep.* **486,** 75–174 (2010).
16. Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N. & Barabási, A-L. The large-scale organization of metabolic networks. *Nature* **407,** 651–654 (2000).
17. Guimerà, R. & Amaral, L. A. N. Functional cartography of complex metabolic networks. *Nature* **433,** 895–900 (2005).
18. Newman, M. E. J. & Girvan, M. Finding and evaluating community structure in networks. *Phys. Rev. E* **69,** 026113 (2004).
19. Flake, G. W., Lawrence, S. R., Giles, C. L. & Coetzee, F. M. Self-organization and identification of Web communities. *IEEE Comput.* **35,** 66–71 (2002).
20. Zhou, H. Distance, dissimilarity index, and network community structure. *Phys. Rev. E* **67,** 061901 (2003).
21. Radicchi, F., Castellano, C., Cecconi, F., Loreto, V. & Parisi, D. Defining and identifying communities in networks. *Proc. Natl Acad. Sci. USA* **101,** 2658–2663 (2004).
22. Palla, G., Derényi, I., Farkas, I. & Vicsek, T. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* **435,** 814–818 (2005).
23. Bagrow, J. P. & Bollt, E. M. Local method for detecting communities. *Phys. Rev. E* **72,** 046108 (2005).
24. Clauset, A. Finding local community structure in networks. *Phys. Rev. E* **72,** 026132 (2005).
25. Hastings, M. B. Community detection as an inference problem. *Phys. Rev. E* **74,** 035102 (2006).
26. Rosvall, M. & Bergstrom, C. T. An information-theoretic framework for resolving community structure in complex networks. *Proc. Natl Acad. Sci. USA* **104,** 7327–7331 (2007).
27. Blondel, V. D., Guillaume, J-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *J. Stat. Mech.* **2008,** P10008 (2008).
28. Agrawal, G. & Kempe, D. Modularity-maximizing network communities via mathematical programming. *Eur. Phys. J. B* **66,** 409–418 (2008).
29. Hofman, J. M. & Wiggins, C. H. Bayesian approach to network modularity. *Phys. Rev. Lett.* **100,** 258701 (2008).
30. Leskovec, J., Lang, K., Dasgupta, A. & Mahoney, M. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Math.* **6,** 29–123 (2009).
31. Ahn, Y-Y., Bagrow, J. P. & Lehmann, S. Link communities reveal multiscale complexity in networks. *Nature* **466,** 761–764 (2010).
32. Lancichinetti, A., Fortunato, S. & Radicchi, F. Benchmark graphs for testing community detection algorithms. *Phys. Rev. E* **78,** 046110 (2008).
33. Danon, L., Duch, J., Diaz-Guilera, A. & Arenas, A. Comparing community structure identification. *J. Stat. Mech.* P09008 (2005).
34. Lancichinetti, A. & Fortunato, S. Community detection algorithms: A comparative analysis. *Phys. Rev. E* **80,** 056117 (2009).
35. Schaeffer, S. E. Graph clustering. *Comput. Sci. Rev.* **1,** 27–64 (2007).
36. Pothen, A., Simon, H. & Liou, K-P. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.* **11,** 430–452 (1990).
37. Kernighan, B. W. & Lin, S. An efficient heuristic procedure for partitioning graphs. *Bell Syst. Tech. J.* **49,** 291–307 (1970).
38. Zachary, W. W. An information flow model for conflict and fission in small groups. *J. Anthropol. Res.* **33,** 452–473 (1977).

39. White, D. R. & Harary, F. The cohesiveness of blocks in social networks: Connectivity and conditional density. *Sociol. Methodol.* **31,** 305–359 (2001).
40. Duch, J. & Arenas, A. Community detection in complex networks using extremal optimization. *Phys. Rev. E* **72,** 027104 (2005).
41. Wilkinson, D. M. & Huberman, B. A. A method for finding communities of related genes. *Proc. Natl Acad. Sci. USA* **101,** 5241–5248 (2004).
42. Wu, F. & Huberman, B. A. Finding communities in linear time: A physics approach. *Eur. Phys. J. B* **38,** 331–338 (2004).
43. Rosvall, M. & Bergstrom, C. T. Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems. *PLoS One* **6,** e18209 (2011).
44. Zhou, H. & Lipowsky, R. *Network Brownian Motion: A New Method to Measure Vertex–Vertex Proximity and to Identify Communities and Subcommunities* 1062–1069 (Lecture Notes in Computer Science, Vol. 3038, Springer, 2004).
45. Pons, P. & Latapy, M. *Proc. 20th International Symposium on Computer and Information Sciences* 284–293 (Lecture Notes in Computer Science, Vol. 3733, Springer, 2005).
46. Reichardt, J. & Bornholdt, S. Detecting fuzzy community structures in complex networks with a Potts model. *Phys. Rev. Lett.* **93,** 218701 (2004).
47. Boccaletti, S., Ivanchenko, M., Latora, V., Pluchino, A. & Rapisarda, A. Detection of complex networks modularity by dynamical clustering. *Phys. Rev. E* **75,** 045102 (2007).
48. Karckhardt, D. & Stern, R. Informal networks and organizational crises: An experimental simulation. *Soc. Psychol. Q.* **51,** 123–140 (1988).
49. Karrer, B. & Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83,** 016107 (2011).
50. Li, Z., Zhang, S., Wang, R-S., Zhang, X-S. & Chen, L. Quantitative function for community detection. *Phys. Rev. E* **77,** 036109 (2008).
51. Newman, M. E. J. Mixing patterns in networks. *Phys. Rev. E* **67,** 026126 (2003).
52. Brandes, U. *et al. Proc. 33rd International Workshop on Graph-Theoretic Concepts in Computer Science* (Lecture Notes in Computer Science,Vol. 4769, Springer, 2007).
53. Medus, A., Acuña, G. & Dorso, C. O. Detection of community structures in networks via global optimization. *Physica A* **358,** 593–604 (2005).
54. Clauset, A., Newman, M. E. J. & Moore, C. Finding community structure in very large networks. *Phys. Rev. E* **70,** 066111 (2004).
55. Wakita, K. & Tsurumi, T. in *Proc. IADIS International Conference, WWW/Internet 2007* (eds Isaías, P., Nunes, M. B. & Barroso, J.) 153–162 (IADIS Press, 2007).
56. Newman, M. E. J. Modularity and community structure in networks. *Proc. Natl Acad. Sci. USA* **103,** 8577–8582 (2006).
57. Shuzhuo, L., Yinghui, C., Haifeng, D. & Feldman, M. W. A genetic algorithm with local search strategy for improved detection of community structure. *Complexity* **15,** 53–60 (2010).
58. Fortunato, S. & Barthélémy, M. Resolution limit in community detection. *Proc. Natl Acad. Sci. USA* **104,** 36–41 (2007).
59. Reichardt, J. & Bornholdt, S. Statistical mechanics of community detection. *Phys. Rev. E* **74,** 016110 (2006).
60. Arenas, A., Fernandez, A. & Gomez, S. Analysis of the structure of complex networks at different resolution levels. *New J. Phys.* **10,** 053039 (2008).
61. Breiger, R. L., Boorman, S. A. & Arabie, P. An algorithm for clustering relations data with applications to social network analysis and comparison with multidimensional scaling. *J. Math. Psychol.* **12,** 328–383 (1975).
62. Holland, P. W., Laskey, K. B. & Leinhardt, S. Stochastic blockmodels: Some first steps. *Soc. Networks* **5,** 109–137 (1983).
63. Snijders, T. A. B. & Nowicki, K. Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *J. Classification* **14,** 75–100 (1997).
64. Nowicki, K. & Snijders, T. A. B. Estimation and prediction for stochastic blockstructures. *J. Am. Stat. Assoc.* **96,** 1077–1087 (2001).
65. Airoldi, E. M., Blei, D. M., Fienberg, S. E. & Xing, E. P. Mixed membership stochastic blockmodels. *J. Mach. Learning Res.* **9,** 1981–2014 (2008).
66. Goldenberg, A., Zheng, A. X., Feinberg, S. E. & Airoldi, E. M. A survey of statistical network structures. *Found. Trends Mach. Learning* **2,** 1–117 (2009).
67. Bickel, P. J. & Chen, A. A nonparametric view of network models and Newman–Girvan and other modularities. *Proc. Natl Acad. Sci. USA* **106,** 21068–21073 (2009).
68. Adamic, L. A. & Glance, N. *Proc. WWW-2005 Workshop on the Weblogging Ecosystem* (2005).
69. Guimerà, R. & Sales-Pardo, M. Missing and spurious interactions and the reconstruction of complex networks. *Proc. Natl Acad. Sci. USA* **106,** 22073–22078 (2009).
70. Yan, X., Zhu, Y., Rouquier, J-B. & Moore, C. in *Proc. 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association of Computing Machinery, 2011).
71. Clauset, A., Moore, C. & Newman, M. E. J. Hierarchical structure and the prediction of missing links in networks. *Nature* **453,** 98–101 (2008).

Acknowledgements

Some of the work described here was financially supported by the US National Science Foundation under grants DMS–0405348 and DMS–0804778.

Additional information

The author declares no competing financial interests. Reprints and permissions information is available online at http://www.nature.com/reprints.

REVIEW ARTICLES | INSIGHT
PUBLISHED ONLINE: 22 DECEMBER 2011 | DOI: 10.1038/NPHYS2160
nature physics

Modelling dynamical processes in complex socio-technical systems

Alessandro Vespignani[1,2]

In recent years the increasing availability of computer power and informatics tools has enabled the gathering of reliable data quantifying the complexity of socio-technical systems. Data-driven computational models have emerged as appropriate tools to tackle the study of dynamical phenomena as diverse as epidemic outbreaks, information spreading and Internet packet routing. These models aim at providing a rationale for understanding the emerging tipping points and nonlinear properties that often underpin the most interesting characteristics of socio-technical systems. Here, using diffusion and contagion phenomena as prototypical examples, we review some of the recent progress in modelling dynamical processes that integrates the complex features and heterogeneities of real-world systems.

Questions concerning how pathogens spread in population networks, how blackouts can spread on a nationwide scale, or how efficiently we can search and retrieve data on large information structures are generally related to the dynamics of spreading and diffusion processes. Social behaviour, the spread of cultural norms, or the emergence of consensus may often be modelled as the dynamical interaction of a set of connected agents. Phenomena as diverse as ecosystems or animal and insect behaviour can all be described as the dynamic behaviour of collections of coupled oscillators. Although all these phenomena refer to very different systems, their mathematical description relies on very similar models that depend on the definition and characterization of a large number of individuals and their interactions in spatially extended systems.

The modelling of dynamical processes is a research field that crosses different disciplines and has developed an impressive array of methods and approaches, ranging from simple explanatory models to realistic approaches capable of providing quantitative insight into real-world systems. Initially these models used simplistic assumptions for the micro-processes of interaction and were mostly concerned with the study of the emerging macro-level behaviour. This interest has favoured the use of techniques akin to statistical physics and the analysis of nonlinear, equilibrium and non-equilibrium physical systems in the study of collective behaviour in social and population systems. In recent years, however, the increase in interdisciplinary work and the availability of system-level high-quality data has opened the way to data-driven models aimed at a realistic description of complex socio-technical systems. Modelling approaches to dynamical processes in complex systems have been expanded into schemes that explicitly include spatial structures and have thus grown into a multiscale framework in which the various possible granularities of the system are considered through different approximations. These models offer a number of interesting and sometimes unexpected behaviours whose theoretical understanding represents a new challenge that has considerably transformed the mathematical and conceptual framework for the study of dynamical processes in complex systems.

Dynamical processes and phase transitions

The study of dynamical processes and the emergence of macro-level collective behaviour in complex systems follows a conceptual route essentially similar to the statistical physics approach to non-equilibrium phase transitions. A prototypical example is that of contagion processes. Epidemiologists, computer scientists and social scientists share a common interest in studying contagion phenomena and rely on very similar spreading models for the description of the diffusion of viruses, knowledge and innovations[1–5]. All these processes define a contagion dynamics that can be seen as an actual biological pathogen that spreads from host to host, or a piece of information or knowledge that is transmitted during social interactions. Let us consider the simple susceptible–infected–recovered (SIR) epidemic model. In this model, infected individuals (labelled with the state I) can propagate the contagion to susceptible neighbours (labelled with the state S) with rate λ, while infected individuals recover with rate μ and become removed from the population. This is the prototypical model for the spread of infectious diseases where individuals recover and are immune to disease after a typical time that, on average, can be expressed as the inverse of the recovery rate. A classic variation of this model is the susceptible–infected–susceptible (SIS) model, in which individuals revert to the susceptible state with rate μ, modelling the possibility of re-infection of individuals. The mapping between epidemic models and non-equilibrium phase transitions was pointed out in physics long ago, making those models of very broad relevance also outside the area of information and disease spreading. The static properties of the SIR model can indeed be mapped to an edge-percolation process[6]. Analogously, the SIS model can be regarded as a generalization of the contact-process model[7], widely studied as the paradigmatic example of an absorbing-state phase transition with a unique absorbing state[8].

A cornerstone feature of epidemic processes is the presence of the so-called epidemic threshold[1]. In a fully homogeneous population, the behaviour of the SIR model is controlled by the reproductive number $R_0 = \beta/\mu$, where $\beta = \lambda\langle k\rangle$ is the per-capita spreading rate, which takes into account the average number of contacts $\langle k\rangle$ of each individual. The reproductive number simply identifies the average number of secondary cases generated by a primary case in an entirely susceptible population and defines an epidemic threshold such that only if $R_0 \geq 1$ ($\beta \geq \mu$) can epidemics reach an endemic state and spread into a closed population. The SIS and SIR models are indeed characterized by a threshold defining the transition between two very different regimes. These regimes are determined by the values of the disease parameters, and characterized by

[1]Department of Physics, College of Computer and Information Sciences, Bouvé College of Health Sciences, Northeastern University, Boston, Massachusetts 02115, USA, [2]Institute for Scientific Interchange (ISI), Torino, 10133, Italy. e-mail: a.vespignani@neu.edu.

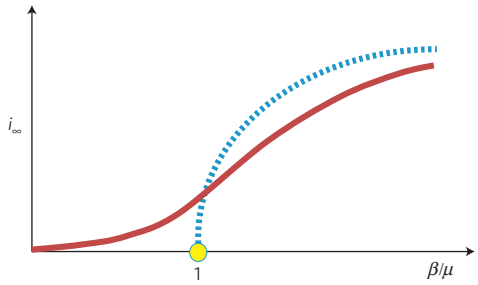

Figure 1 | Phase diagram of epidemic models. Illustration of the behaviour of the prevalence i_∞ for the SIS and SIR model in a heterogeneous network (solid line) as a function of the spreading rate β/μ, compared with the theoretical prediction for a homogeneous network (dashed line). The figure clearly shows the difference between homogeneous and heterogeneous networks, where the epidemic threshold is shifted to very small values. For scale-free networks with degree distribution exponent $\gamma \leq 3$, however, the associated prevalence i_∞ is extremely small over a large range of values of β/μ. In other words, as noted since the first work on epidemic spreading in complex networks, the bad news about the suppression (or very small value) of the epidemic threshold is balanced by the very low prevalence attained by the epidemic[46].

the global parameter i_∞, which identifies the density of infected individuals (or nodes in a network) in the infinite-time limit. In the limit of an infinitely large population, this density is zero below the threshold and assumes a finite value above the threshold. From this perspective we can consider the epidemic threshold as the critical point of the system and i_∞ as representing the order parameter characterizing the transition. Below the critical point the system relaxes in a frozen state with null dynamics—the healthy phase. Above this point, a dynamical state characterized by a macroscopic number of infected individuals sets in, defining an infected phase (Fig. 1).

Many other pioneering works in the area of social sciences use simple dynamical models to explore the emergence of macro-level collective behaviour as a function of the micro-level processes acting among the agents of a large population[9–11], and the incursions by statistical physicists in the area of social sciences have become very frequent (see, for example, the recent review by Castellano *et al.*[12]). A first class of models is represented by behavioural models where the attributes of agents are binary variables similar to Ising spins, as in the case of the voter model[13], the majority-rule model[14,15] and the Sznajd model[16]. In other instances, further realism has been introduced by the use of continuous opinion variables[17–19]. Along the path opened by Axelrod[11], models in which opinions or cultures are represented by vectors of cultural traits have introduced the notion of bounded confidence: an agent will not interact with any other agent, independently of their opinions, unless the opinions are close enough.

Finally, there is a vast class of models that focus on the analysis of diffusion processes as a tool to study phase transitions and emergent phenomena in simple models mimicking the routing of information packets in technological systems and networks. In this case the focus is on what lies behind the appearance of congestion and traffic self-similarity[20–26]. In traffic problems, one of the main issues is that the diffusion process is not random but determined by recurrent patterns, reinforcing mechanisms and routing strategies that represent formidable challenges to the modelling of systems[27]. Interestingly, it is the study of traffic dynamics in the Internet and the World Wide Web that has made clear the central role of networks and their structural properties in the understanding and characterization of dynamical processes in real-world systems.

Box 1 | The heterogeneous mean-field approach.

The heterogeneous mean-field approach generalizes, for the case of networks with arbitrary degree distribution, the equations describing the dynamical process, by considering degree-block variables grouping nodes within the same degree class k. If we consider the SIS model, the variables describing the system are i_k and s_k, which respectively represent the fraction of nodes with degree k in the infected and susceptible class. The evolution equation for the infected individual is

$$\frac{\mathrm{d}i_k(t)}{\mathrm{d}t} = -\mu i_k + \lambda[1-i_k(t)]k\Theta_k(t)$$

The first term just expresses the fact that any node in the infected state may recover with rate μ. The second term, which generates new infected individuals, is proportional to the probability of transmission λ, the degree k, the probability $1-i_k$ that a vertex with degree k is not infected, and the density Θ_k of infected neighbours of vertices of degree k, which is the probability of contacting an infected individuals. As we are still assuming a mean-field description of the system, the latter term is the average probability that any given neighbour of a vertex of degree k is infected. This quantity can be expressed as $\Theta_k(t) = \sum_{k'} P(k'|k) i_{k'}(t)$, which is the average over all possible degrees k' of the probability $P(k'|k)$ that any edge of a node of degree k is pointing to a node of degree k' times the probability $i_{k'}$ that the node is infected. This expression can be further simplified by considering a random network in which the conditional probability does not depend on the originating node. In this case we have that $P(k'|k) = k'P(k')/\langle k\rangle$, following simply from the fact that any edge has a probability proportional to the degree itself of pointing to a node with degree k' (see ref. 38). On substituting the expression for Θ in the main equation and adopting the early-epidemic assumption (that is, assuming that all second-order terms of i_k and r_k can be neglected), we readily recover the topology-dependent epidemic threshold result, $\lambda/\mu = \langle k\rangle/\langle k^2\rangle$.

Following the results obtained with the HMF assumption, a number of rigorous results that link the network topology to the epidemic threshold have been derived[53,57,58]. These results relate the epidemic threshold to the largest eigenvalue of the adjacency matrix of the network, showing that the HMF does not recover the correct behaviour for the SIS model when the degree distribution of the graph $P(k) \sim k^{-\gamma}$ has $\gamma > 3$. The rigorous results refer to quenched networks where the adjacency matrix is fixed in time. The HMF assumption instead, in its mean-field perspective, is equivalent to a system in which edges are continuously reshuffled so that the elements of the adjacency matrix are defined by the effective probabilities $k_i k_j / \sum_i k_i$ that two nodes i and j with degree k_i and k_j, respectively, are connected. This consideration clearly shows the shortcomings of the HMF assumption in the case of systems where the timescale of the transmission or infection is very short with respect to the duration of the contact and the adjacency matrix can be considered as quenched. The HMF can be considered, however, as a description of the system closer to reality in situations where the transmission occurs on rapidly varying networks; this is for instance the case for many influenza-like illnesses, where the infectious period is much longer than the duration of contacts responsible for the transmission[57].

Complex networks and dynamical processes

We live in an increasingly interconnected world, where infrastructures composed of different technological layers inter-operate

NATURE PHYSICS DOI:10.1038/NPHYS2160

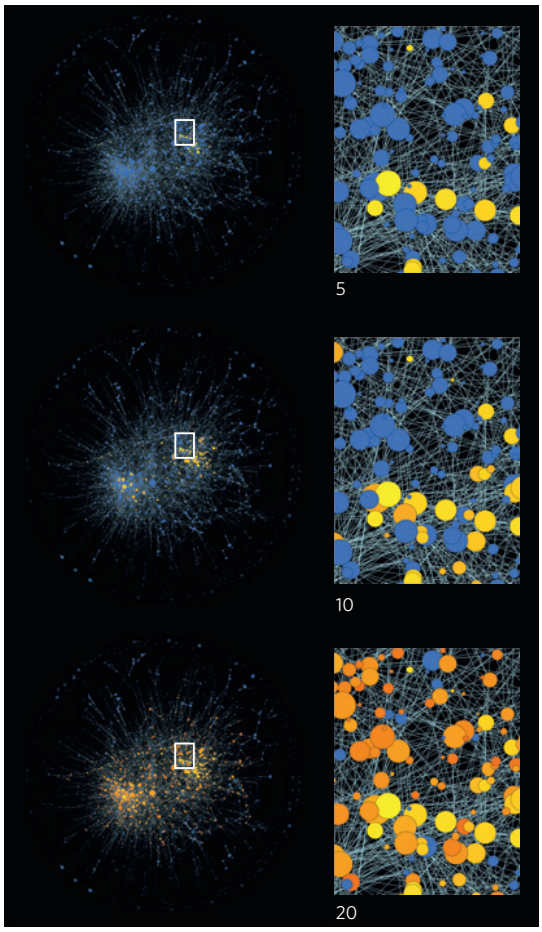

Figure 2 | Progression of an epidemic process. The progression of a susceptible–infected (SI) epidemic in a heavy-tailed network at three snapshots of the process, corresponding to time $t = 5$, 10 and 20, measured in unitary time integration steps of the model. The SI model assumes that infected nodes will spread the infection indefinitely to neighbours with rate α. In this case we know that the system is eventually completely infected, whatever the spreading rate of the infection. However, we can highlight the effect of topological fluctuations on the spreading hierarchy. Susceptible nodes are coloured blue and infected nodes are coloured from yellow to red according to the time of infection (red corresponding to later times). The size of a node is proportional to the node degree. In general, the first nodes to be infected are the large hubs with high degree, then the epidemic progresses in time by a dynamical cascade through degree classes, finally affecting low-degree nodes.

within the social component that drives their use and development. Examples are the Internet, the World Wide Web, mobile technologies, and transportation and mobility infrastructures[28–34]. The multiscale nature and complexity of these networks are crucial features in understanding and managing socio-technical systems and the dynamical processes occurring on top of them. For this reason, in the past decade, the study of models unfolding on complex networks has generated a body of work that includes results of conceptual and practical relevance[35–40]. The resilience of networks, their vulnerability to attacks, and their synchronization properties are all drastically affected by topological heterogeneities. Consensus formation, disease spreading and the accessibility of information can benefit or be impaired by the connectivity pattern of the population or infrastructure we are looking at. Network science has thus become pervasive in the study of complex systems and presented us with a number of surprising discoveries that have steered our way of thinking on dynamical processes in socio-technical systems.

One of the most important features affecting dynamical processes in real-world networks is the presence of dynamic self-organization and the lack of characteristic scales—typical hallmarks of complex systems[40–44]. Although those characteristics have long been acknowledged as a relevant factor in determining the properties of dynamical processes, many real-world networks exhibit levels of heterogeneity that were not anticipated until a few years ago. In particular, the various statistical distributions characterizing these networks are generally heavy-tailed, skewed, and varying over several orders of magnitude. This is a very peculiar feature, typical of many natural and artificial complex networks, characterized by virtually infinite degree fluctuations, where the degree k of a given node represents its number of connections to other nodes. In contrast to regular lattices and homogeneous graphs, characterized by nodes having a typical degree k close to the average $\langle k \rangle$, such networks are structured in a hierarchy where a few nodes (the hubs) have very high degree whereas the vast majority of nodes have lower degrees. This feature is usually manifest in a heavy-tailed degree distribution, often approximated by a power-law behaviour of the form $P(k) \sim k^{-\gamma}$, which implies a non-negligible probability of finding vertices with very high degree[40,42–44]. Furthermore, the presence of large-scale fluctuations associated with heavy-tail distributions is also observed for the intensity carried by the connecting links, transport flows, and other basic quantities that go beyond the connectivity description of the network[45].

The presence of large-scale fluctuations virtually acting at all scales of the network connectivity pattern calls for a mathematical analysis where the variables characterizing each node of the network explicitly enter the description of the system. Unfortunately, the general solution, handling the master equation of the system, is hardly, if ever, achievable—even for very simple dynamical processes. For this reason, a viable theoretical approach has to be based on techniques such as mean-field and deterministic continuum approximations, which usually provide the understanding of the basic phenomenology and phase diagram of the process under study. In both cases, the heterogeneous nature of the network-connectivity pattern is introduced by aggregating variables according to a degree-block formalism that assumes that all nodes with the same degree k are statistically equivalent[38,46,47]. This assumption allows the grouping of nodes in degree classes, yielding a convenient representation of the system. For instance, if for each node i we associate a corresponding state σ_i characterizing its dynamical state, a convenient representation of the system is provided by the quantity S_k, which indicates the number of nodes of degree k in the dynamical state $\sigma = s$, and the corresponding degree-block density of nodes of degree k in the state s

$$s_k = \frac{S_k}{V_k}$$

where V_k is the number of nodes of degree k. Finally, the global averages on the network are given by the expression

$$\rho_s = \sum_k P(k) s_k$$

where ρ_s is the probability that any given node is in the state s. This formalism defines a mean-field approximation within each degree class, relaxing, however, the overall homogeneity assumption on the degree distribution[38]. This framework, first introduced for the description of epidemic processes, is at the basis of the heterogeneous mean-field (HMF) approach that allows the analytical study of dynamical processes in complex networks by writing mean-field dynamical equations for each degree class variable. An example

Box 2 | The particle–network framework.

The particle–network framework extends the HMF approach to the case of a reaction–diffusion system in which particles (or individuals) diffuse on a network with arbitrary topology. A convenient representation of the system is therefore provided by quantities defined in terms of the degree k

$$N_k = \frac{1}{V_k}\sum_{i|k_i=k} N_i$$

where V_k is the number of nodes with degree k and the sums run over all nodes i having degree k_i equal to k. The degree-block variable N_k represents the average number of particles in nodes with degree k. The use of the HMF approach amounts to the assumption that nodes with degree k, and thus the particles in those nodes, are statistically equivalent. In this approximation the dynamics of particles randomly diffusing on the network is given by a mean-field dynamical equation expressing the variation in time of the particle subpopulations $N_k(t)$ in each degree block k. This can simply be written as:

$$\frac{\partial N_k}{\partial t} = -d_k N_k(t) + k\sum_{k'} P(k'|k) d_{k'k} N_{k'}(t)$$

The first term of the equation just considers that only a fraction of particles d_k moves out of the node per unit time. The second term accounts for particles diffusing from its neighbours into the node of degree k. This term is proportional to the number of links k times the average number of particles coming from each neighbour. The number of particles arriving from each neighbour is thus equal to that of particles $d_{k'k}N_{k'}(t)$ diffusing on any edge connecting a node of degree k' with a node of degree k, averaged over the conditional probability $P(k'|k)$ that an edge belonging to a node of degree k is pointing to a node of degree k'. Here the term $d_{k'k}$ is the diffusion rate along the edges connecting nodes of degree k and k'. The rate at which individuals leave a subpopulation with degree k is then given by $d_k = k\sum_{k'} P(k'|k) d_{kk'}$. The function $P(k'|k)$ encodes the topological connectivity properties of the network and allows the study of different topologies and mixing patterns. The above equation explicitly introduces the diffusion of particles into the description of the system. The equation can easily be generalized to particles with different states, and reacting among themselves, by adding a reaction term to the above equations. For instance, the generalization of the SIR model described in the main text would consider three types of particle, denoting infected, susceptible and recovered individuals. The reaction taking place among individuals in the same node would be the usual contagion process among susceptibles and infected individuals, and the spontaneous recovery of infected individuals.

The analysis of a simple diffusion process immediately indicates the importance of network topology. In a random network with arbitrary degree distribution, the stationary state reached by a swarm of particles diffusing with the same diffusive rate yields $N_k \sim k$ and the probability to find a single diffusing walker in a node of degree k is

$$p_k = \frac{k}{\langle k\rangle}\frac{1}{V}$$

where V is the total number of nodes in the network. This expression implies that the higher the degree of the nodes, the greater the probability to be visited by the walker. This observation has profound consequences for the way we can discover, retrieve and rank information in complex networks. The PageRank algorithm[117] is in this respect a major breakthrough, based on the idea that a viable ranking depends on the topological structure of the network, and is defined by essentially simulating the random surfing process on the web graph. The most important pages are simply those with the highest probability of being discovered if the web-surfer had infinite time to explore the web. Analogously, search processes can take advantage of this property using degree-biased searching algorithms that bias the routing of messages towards nodes with high degree[115,116].

of the HMF approach is given in Box 1 for the case of the SIS model. The HMF technique is often the first line of attack towards understanding the effects of complex connectivity patterns on dynamical processes and it has been used widely in a broad range of phenomena, although with different names and specific assumptions, depending on the problem at hand. Although it contains several approximations, the HMF approach readily shows that the heterogeneity found in the connectivity pattern of many networks may drastically affect the unfolding of the dynamical process.

The classic example for the effect of degree heterogeneity on dynamical processes in complex networks is epidemic spreading. The previously discussed result of the presence of an epidemic threshold in the SIR and SIS models is obtained under the assumption that each individual in the system has, to a first approximation, the same number of connections $k \simeq \langle k\rangle$. However, social heterogeneity and the existence of 'super-spreaders' have long been known in the epidemics literature[48]. Generally, it is possible to show that the reproductive rate R_0 is renormalized by fluctuations in the transmissibility or contact pattern as $R_0 \to R_0(1+f(\nu))$, where $f(\nu)$ is a positive and increasing function of the standard deviation ν of the individual transmissibility or connectivity pattern[49]. In particular, by generalizing the dynamical equations of the SIS model, the HMF approach yields that the disease will affect a finite fraction of the population only if $\beta/\mu \geq \langle k\rangle^2/\langle k^2\rangle$, that is the ratio between the first and second moments of the degree distribution[38,46,47]. This readily suggests that the topology of the network enters the very definition of the epidemic threshold. Furthermore, this implies that in heavy-tailed networks such that $\langle k^2\rangle \to \infty$, in the limit of infinite network size, we have a null epidemic threshold. Although this is not the case in any finite-size real-world network[50,51], larger heterogeneity levels lead to smaller epidemic thresholds (Fig. 1). This is an important result, which indicates that heterogeneous networks behave very differently from homogeneous networks with respect to physical and dynamical processes. Indeed, the heterogeneous connectivity pattern of networks affects also the dynamical progression of the epidemic process, which results in a striking hierarchical dynamics in which the infection propagates from higher-degree to lower-degree classes. The infection first takes control of the high-degree vertices in the network, then rapidly invades the network via a cascade through progressively lower-degree classes (Fig. 2). It also turns out that the time behaviour of epidemic outbreaks and the growth of the number of infected individuals are governed by a timescale τ proportional to the ratio between the first and second moment of the network's degree distribution, thus suggesting a velocity of progression that increases with the heterogeneity of the network[52].

The change of framework suggested by the network heterogeneity in the case of epidemic processes has triggered many studies

NATURE PHYSICS DOI:10.1038/NPHYS2160

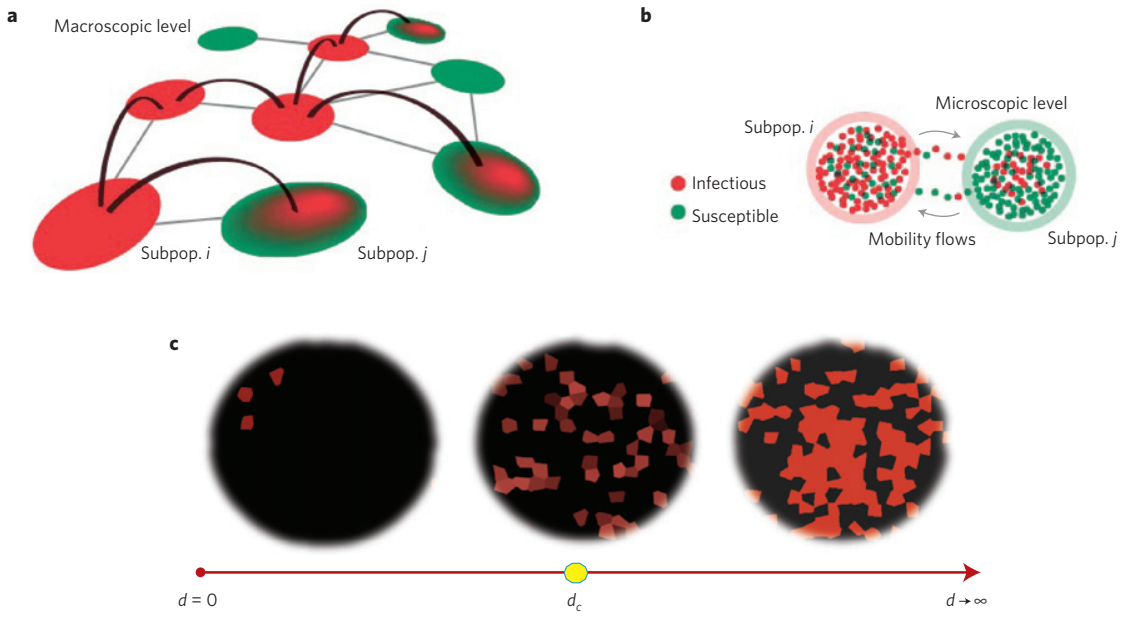

Figure 3 | Illustration of the global threshold in reaction–diffusion processes. a, Schematic of the simplified modelling framework based on the particle–network scheme. At the macroscopic level the system is composed of a heterogeneous network of subpopulations. The contagion process in one subpopulation (marked in red) can spread to other subpopulations as particles diffuse across subpopulations. **b**, At the microscopic level, each subpopulation contains a population of individuals. The dynamical process, for instance a contagion phenomena, is described by a simple compartmentalization (compartments are indicated by different coloured dots). Within each subpopulation, individuals can mix homogeneously, or according to a subnetwork, and can diffuse with rate d from one subpopulation to another, following the edges of the network. **c**, A critical value d_c of the diffusion strength for individuals or particles identifies a phase transition between a regime in which the contagion affects a large fraction of the system and one in which only a small fraction is affected (see the discussion in the text). Panels **a** and **b** reproduced from ref. 118.

aimed at providing a more rigorous analytical basis for the results obtained with the HMF and other approximate methods exploring different spreading models[53–58]. Equally important is the research activity concerned with developing dynamical ad hoc strategies for network protection; targeted immunization strategies and targeted prophylaxis that evolve with time might be particularly effective in the control of epidemics on heterogeneous patterns, compared with massive uniform vaccinations or stationary interventions[59–62]. Following the results on epidemic processes, an avalanche of studies addressed the study of the effect of the network's structure on the behaviour of the most widely used classes of dynamical processes. For instance, in the area of synchronization it has been shown that networks with heavy-tailed degree distributions, and therefore a large number of hubs, are more difficult to synchronize than homogeneous networks, a counterintuitive insight dubbed the paradox of heterogeneity[63–66]. In the case of packet-traffic routing, homogeneous networks have typically much larger congestion thresholds than heterogeneous graphs[67–69]. Finally, a wealth of surprising results, often overturning the common wisdom obtained by studies on regular networks, have been harvested on the voter and the Axelrod models[70–73], and many other models for the emergence of cooperation[38,74].

Reaction–diffusion processes and computational thinking

Although most approaches assume systems in which each node of the network corresponds to a single individual, it is of crucial importance for the study of many phenomena to provide a general understanding of processes where the multiple occupancy of nodes is a key feature. Examples of multiple occupancy are provided by chemical reactions, in which different molecules or atoms diffuse in space and may react whenever in close contact. Mechanistic metapopulation epidemic models, where particles represent people moving between different locations, and the routing of information packets in technological networks provides relevant examples in the case of socio-technical systems[75–79]. All those phenomena fall into the category of reaction–diffusion processes, where each node i is allowed to have any non-negative integer number of particles N_i so that the total particle population of the system is $N = \sum N_i$. The particle–network framework extends the heterogeneous mean-field approach to reaction–diffusion systems in networks with arbitrary degree distribution (Box 2). Particles diffuse along the edges connecting nodes, with a diffusion coefficient that depends on the node degree and/or other nodes' attributes. Within each node, particles may react according to different schemes characterizing the interaction dynamic of the system.

The consideration of complex networks in reaction–diffusion systems has broadened our knowledge of non-equilibrium reaction–diffusion systems in heterogeneous systems. For instance, the Turing mechanism represents a classical model for the formation of self-organized spatial structures in non-equilibrium activator–inhibitor systems. By studying the Turing mechanism[80] in systems with heterogeneous connectivity patterns it has been found that the relevant instabilities of the systems are localized in a set of vertices with degree inversely proportional to the characteristic scale of diffusion[81]. Interestingly, and contrary to other models and systems where the hubs are the playmakers, the segregation process takes place mainly in vertices of low degree.

Another interesting example is that of simple epidemic processes, such as the SIR model in a metapopulation context[79,82–90]. In this case, each node of the network is a subpopulation (ideally an urban area) connected by a transportation system (the edges of the network) that allows individuals to move from one subpopulation to another (Fig. 3). If we assume a diffusion rate d for each individual and consider that the single-population reproductive number of the SIR model is $R_0 > 1$, we can easily identify two different limits. If $d = 0$, any epidemic occurring in a given subpopulation

will remain confined; no individual could travel to a different subpopulation and spread the infection across the system. In the limit $d \to \infty$ we have that individuals are constantly wandering from one subpopulation to the other and the system is in practice equivalent to a well-mixed unique population. In this case, as $R_0 > 1$, the epidemic will spread across the entire system. A transition point between these two regimes occurs at a threshold value d_c of the diffusion rate, identifying a global invasion threshold. This threshold cannot be uncovered by continuous models as it is related to the stochastic diffusion rate of single individuals. Furthermore, the global invasion threshold is affected by the connectivity fluctuations of the metapopulation network. In particular, the greater the network heterogeneity, the smaller the value of the diffusion rate above which the epidemic may globally invade the metapopulation system. This result assumes a particular relevance, as it explains why travel restrictions seem to be highly ineffective in containing epidemics: the complexity and heterogeneity of present-day transport networks favour considerably the global spread of infectious diseases. Only infeasibly tight mobility restrictions, reducing global travel fluxes by 90% or more, would be effective[84,91,92].

Reaction–diffusion models lend themselves to the implementation of large-scale computer simulations (Monte-Carlo and individual-based simulations) that allow one to track microscopically the state of each node and the evolution of the dynamical process. At the most detailed level, the introduction of agent-based models has enabled the usual modelling perspective to be extended further by simulating the population and embedding environment on an individual-by-individual basis. An example is epidemic modelling, where spatially structured and agent-based models at various granularities (country, inter-city, intra-city) have been pushed to the computational limits with the integration of huge amount of data describing the flows of people and/or animals[93–97]. These models can generate results at an unprecedented level of detail and have been used successfully in the analysis and anticipation of real epidemics, such as the 2009 H1N1 pandemic[98,99]. Computer simulations thus become valuable in allowing both *in silico* experiments that would be infeasible in real systems and the capability to analyse and forecast scenarios. This computational approach is also helping to guide researchers in identifying typical nonlinear behaviour and tipping points[100] not accessible by analytical means, using the numerical simulations as a novel experimental workbench[101,102].

Co-evolution, timescale and control

Although in recent years our understanding of dynamical processes in complex networks has progressed at an exponential pace, there are still a number of major challenges that keep the research community actively engaged. The first challenge stems from the fact that the analysis of dynamical processes is generally performed in the presence of a timescale separation between the network evolution and the dynamical process unfolding on its structure. In one limit we can consider the network as quenched in its connectivity pattern, thus evolving on a timescale that is much longer that the dynamical process itself. In the other limiting case, the network evolves on a timescale much shorter than the dynamical process, which thus effectively disappears from the definition of the interaction among individuals such that this interaction can be conveniently replaced by effective random coupling. Although the timescale separation is extremely convenient with a view to the numerical and analytical tractability of the models, networks generally evolve on a timescale that might be comparable to that of the dynamical process. Furthermore, the network properties used in defining models generally represent a time-integrated static snapshot of the system. However, in many systems the timing and duration of interactions define processes on a timescale very different from, and often conflicting with, those of the time-integrated view. This highlights the importance of considering the concurrency of network evolution and dynamical processes in realistic models to avoid misleading conclusions[103–106].

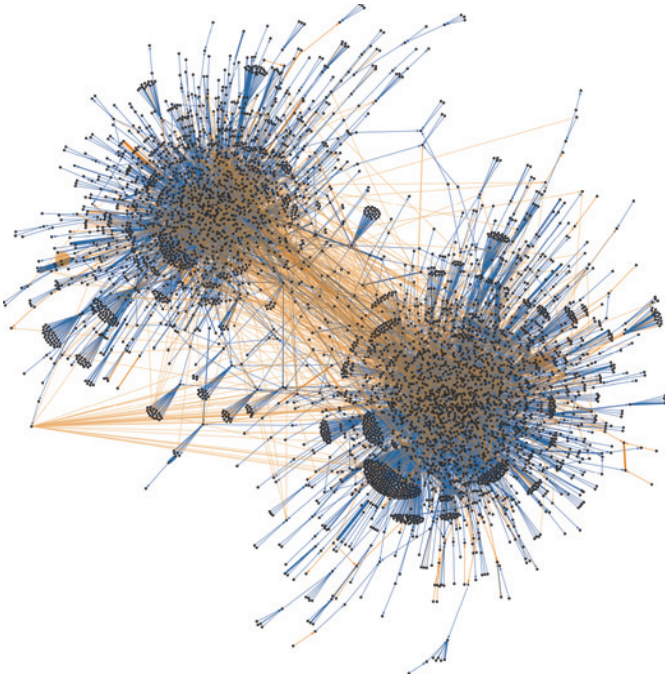

Figure 4 | Visualization of the dynamical network generated by Twitter interactions. Twitter is a microblogging tool that allows users to post and relay ('re-tweet') short messages. The topic of the message is signalled by short identifiers (@mentions, #hash-tags and urls). This feature allows one to trace the spreading of specific discussion topics (also called memes). The figure shows the diffusion network for the tag #gop. Each node corresponds to an individual user. Blue edges represent re-tweets and orange edges represent mentions. Two communities are clearly visible, corresponding to politically left- and right-leaning users[113]. Communications between the two communities take place primarily through the use of mentions, while within a group communication occurs through re-tweets. The figure, obtained using the Truthy infrastructure[114], clearly exemplifies the co-evolution of the communication network with the spreading process.

A second challenge is the co-evolution of networks with the dynamical process. Access to the mathematical and statistical laws that characterize the interplay and feedback mechanisms between the network evolution and the dynamical processes is extremely important, especially in social systems, where the adaptive nature of agents is of paramount importance[106–108]. The spreading of an opinion is affected by the interaction among individuals, but the presence and/or establishment of interaction among individuals is affected by their opinion. This issue is increasingly relevant in the area of the modern social networks populating the information-technology ecosystem, such as those defined by the Facebook and Twitter applications. In this case the network and the spread of information cannot be defined in isolation, because of rapidly changing interactions and modes of communication that depend on the type of information exchanged and the adaptive behaviour of individuals (Fig. 4).

The adaptive behaviour of individuals to the dynamical processes they are involved in represents another modelling challenge, as it calls for the understanding of the feedback among different and competing dynamical processes. For instance, relatively little systematic work has been done to provide coupled behaviour–disease models able to close the feedback loop between

behavioural changes triggered in the population by an individual's perception of the disease spread and the actual disease spread[109,110]. Similar issues arise in many areas where we find competing processes of adaptation and awareness to information or knowledge spreading in a population[111].

Finally, the overall goal is not only to understand complex systems, mathematically describe their structure and dynamics, and predict their behaviour, but also to control their dynamics. Also in this case, although control theory offers a large set of mathematical tools for steering engineered and natural systems, we are just taking the first steps towards a full understanding of how the network heterogeneities influence our ability to control the network dynamics and how the network evolution impacts controllability[112].

Conclusions

There are no doubts that a complete understanding of complex socio-technical systems requires diving into the specifics of each system by adopting a domain-specific perspective. Data-driven models, however, are generating new questions, the answers to which should preferably be analytical and applicable to a wide range of systems. What are the fundamental limits to predictability with computational modelling? How does our understanding depend on the level of accuracy of our description and knowledge of the state of the system? The research community needs, now more than ever, the kind of basic theoretical understanding that would help discriminate between what is relevant and what is superfluous in the description of socio-technical systems. This is a crucial endeavour if we want to complement data-driven approaches with a conceptual understanding that would help guide the management, prediction and control of dynamical processes in complex systems—a conceptual understanding that necessarily descends from the study of the dynamical models and processes presented here.

References

1. Keeling, M. J. & Rohani, P. *Modeling Infectious Diseases in Humans and Animals* (Princeton Univ. Press, 2008).
2. Goffman, W. & Newill, V. A. Generalization of epidemic theory: An application to the transmission of ideas. *Nature* **204,** 225–228 (1964).
3. Rapoport, A. Spread of information through a population with socio-structural bias: I. Assumption of transitivity. *Bull. Math. Biol.* **15,** 523–533 (1953).
4. Tabah, A. N. Literature dynamics: Studies on growth, diffusion, and epidemics. *Annu. Rev. Inform. Sci. Technol.* **34,** 249–286 (1999).
5. Lloyd, A. L. & May, R. M. How viruses spread among computers and people. *Science* **292,** 1316–1317 (2001).
6. Grassberger, P. On the critical behavior of the general epidemic process and dynamical percolation. *Math. Biosci.* **63,** 157–172 (1983).
7. Harris, T. E. Contact interactions on a lattice. *Ann. Prob.* **2,** 969–988 (1974).
8. Marro, J. & Dickman, R. *Nonequilibrium Phase Transitions in Lattice Models* (Cambridge Univ. Press, 1999).
9. Granovetter, M. Threshold models of collective behavior. *Am. J. Sociol.* **83,** 1420–1443 (1978).
10. Nowak, A., Szamrej, J. & Latané, B. From private attitude to public opinion: A dynamic theory of social impact. *Psychol. Rev.* **97,** 362–376 (1990).
11. Axelrod, R. *The Complexity of Cooperation* (Princeton Univ. Press, 1997).
12. Castellano, C., Fortunato, S. & Loreto, V. Statistical physics of social dynamics. *Rev. Mod. Phys.* **81,** 591–646 (2009).
13. Krapivsky, P. L. Kinetics of monomer–monomer surface catalytic reactions. *Phys. Rev. A* **45,** 1067–1072 (1992).
14. Galam, S. Minority opinion spreading in random geometry. *Eur. Phys. J. B* **25,** 403–406 (2002).
15. Krapivsky, P. L. & Redner, S. Dynamics of majority rule in two-state interacting spin systems. *Phys. Rev. Lett.* **90,** 238701 (2003).
16. Sznajd-Weron, K. & Sznajd, J. Opinion evolution in closed community. *Int. J. Mod. Phys. C* **11,** 1157–1165 (2000).
17. Deffuant, G., Neau, D., Amblard, F. & Weisbuch, G. Mixing beliefs among interacting agents. *Adv. Complex Syst.* **3,** 87–98 (2000).
18. Hegselmann, R. & Krause, U. Opinion dynamics and bounded confidence models, analysis and simulation. *J. Art. Soc. Soc. Sim.* **5,** 2 (2002).
19. Ben-Naim, E., Krapivsky, P. L. & Redner, S. Bifurcations and patterns in compromise processes. *Physica D* **183,** 190–204 (2003).
20. Leland, W. E., Taqqu, M. S., Willinger, W. & Wilson, D. V. On the self-similar nature of Ethernet traffic. *IEEE/ACM Trans. Netw.* **2,** 1–15 (1994).
21. Csabai, I. 1/f noise in computer network traffic. *J. Phys. A* **27,** L417–L42 (1994).
22. Solé, R. V. & Valverde, S. Information transfer and phase transitions in a model of internet traffic. *Physica A* **289,** 595–605 (2001).
23. Willinger, W., Govindan, R, Jamin, S., Paxson, V. & Shenker, S. Scaling phenomena in the Internet: Critically examining criticality. *Proc. Natl Acad. Sci. USA* **99,** 2573–2580 (2002).
24. Valverde, S. & Solé, R. V. Internet's critical path horizon. *Eur. Phys. J. B* **38,** 245–252 (2004).
25. Tadić, B., Thurner, S. & Rodgers, G. J. Traffic on complex networks: Towards understanding global statistical properties from microscopic density fluctuations. *Phys. Rev. E* **69,** 036102 (2004).
26. Crovella, M. E. & Krishnamurthy, B. *Internet Measurements: Infrastructure, Traffic and Applications* (John Wiley, 2006).
27. Helbing, D. Traffic and related self-driven many particle systems. *Rev. Mod. Phys.* **73,** 1067–1141 (2001).
28. Albert, R., Jeong, H. & Barabási, A-L. Internet: Diameter of the World-Wide Web. *Nature* **401,** 130–131 (1999).
29. Pastor-Satorras, R. & Vespignani, A. *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge Univ. Press, 2004).
30. Brockmann, D., Hufnagel, L. & Geisel, T. The scaling laws of human travel. *Nature* **439,** 462–465 (2006).
31. Onnela, J-P. *et al.* Structure and tie strengths in mobile communication networks. *Proc. Natl Acad. Sci. USA* **104,** 7332–7337 (2007).
32. González, M. C., Hidalgo, C. A. & Barabási, A-L. Understanding individual human mobility patterns. *Nature* **453,** 779–782 (2008).
33. Lazer, D *et al.* Life in the network: The coming age of computational social science. *Science* **323,** 721–723 (2009).
34. Vespignani, A. Predicting the behavior of tecno-social systems. *Science* **325,** 425–428 (2009).
35. Albert, R. & Barabási, A-L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74,** 47–97 (2002).
36. Boccaletti, S. *et al.* Complex networks: Structure and dynamics. *Phys. Rep.* **424,** 175–308 (2006).
37. Dorogovtsev, S. N., Goltsev, A. V. & Mendes, J. F. F. Critical phenomena in complex networks. *Rev. Mod. Phys.* **80,** 1275–1335 (2008).
38. Barrat, A., Barthelemy, M. & Vespignani, A. *Dynamical Processes on Complex Networks* (Cambridge Univ. Press, 2008).
39. Cohen, R. & Havlin, S. *Complex Networks: Structure, Robustness and Function* (Cambridge Univ. Press, 2010).
40. Newman, M. E. J. *Networks: An Introduction* (Oxford Univ. Press, 2010).
41. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393,** 440–442 (1998).
42. Barabási, A-L. & Albert, R. Emergence of scaling in random networks. *Science* **286,** 509–512 (1999).
43. Dorogovtsev, S. N. & Mendes, J. F. F. *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford Univ. Press, 2003).
44. Amaral, L. A. N., Scala, A., Barthlemy, M. & Stanley, H. E. Classes of small-world networks. *Proc. Natl Acad. Sci. USA* **97,** 11149–11154 (2005).
45. Barrat, A., Barthlemy, M., Pastor-Satorras, R. & Vespignani, A. The architecture of complex weighted networks. *Proc. Natl Acad. Sci. USA* **101,** 3747–3752 (2004).
46. Pastor-Satorras, R. & Vespignani, A. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.* **86,** 3200–3203 (2001).
47. Moreno, Y., Pastor-Satorras, R. & Vespignani, A. Epidemic outbreaks in complex heterogeneous networks. *Eur. Phys. J. B* **26,** 521–529 (2002).
48. Hethcote, H. W. & Yorke, J. A. Gonorrhea: Transmission and control. *Lect. Notes Biomath.* **56,** 1–105 (1984).
49. Anderson, R. M. & May, R. M. *Infectious Diseases in Humans* (Oxford Univ. Press, 1992).
50. May, R. M. & Lloyd, A. L. Infection dynamics on scale-free networks. *Phys. Rev. E* **64,** 066112 (2001).
51. Pastor-Satorras, R. & Vespignani, R. Epidemic dynamics in finite size scale-free networks. *Phys. Rev. E* **65,** 035108(R) (2002).
52. Barthelemy, M., Barrat, A., Pastor-Satorras, R. & Vespignani, A. Velocity and hierarchical spread of epidemic outbreaks in scale-free networks. *Phys. Rev. Lett.* **92,** 178701 (2004).
53. Wang, Y., Chakrabarti, D., Wang, G. & Faloutsos, C, in *Proc. 22nd International Symposium on Reliable Distributed Systems (SRDS'03)* 25–34 (IEEE, 2003).
54. Boguna, M., Pastor-Satorras, R. & Vespignani, A. Absence of epidemic threshold in scale-free networks with degree correlations. *Phys. Rev. Lett.* **90,** 028701 (2003).
55. Castellano, C. & Pastor-Satorras, R. Routes to thermodynamic limit on scale-free networks. *Phys. Rev. Lett.* **100,** 148701 (2008).
56. Chatterjee, S. & Durrett, R. Contact processes on random graphs with power law degree distributions have critical value 0. *Ann. Probab.* **37,** 2332–2356 (2009).

57. Castellano, C. & Pastor-Satorras, R. Thresholds for epidemic spreading in networks. *Phys. Rev. Lett.* **105,** 218701 (2010).
58. Durrett, R. Some features of the spread of epidemics and information on a random graph. *Proc. Natl Acad. Sci. USA* **107,** 4491–4498 (2010).
59. Pastor-Satorras, R. & Vespignani, A. Immunization of complex networks. *Phys. Rev. E* **65,** 036104 (2001).
60. Cohen, R., Havlin, S. & Ben-Avraham, D. Efficient immunization strategies for computer networks and populations. *Phys. Rev. Lett.* **91,** 247901 (2003).
61. Holme, P. Efficient local strategies for vaccination and network attack. *Europhys. Lett.* **68,** 908–914 (2004).
62. Goldenberg, J., Shavitt, Y., Shir, E. & Solomon, S. Distributive immunization of networks against viruses using the 'honey-pot' architecture. *Nature Phys.* **1,** 184–188 (2005).
63. Motter, A. E., Zhou, C. S. & Kurths, J. Enhancing complex-network synchronization. *Europhys. Lett.* **69,** 334–340 (2005).
64. Motter, A. E., Zhou, C. S. & Kurths, J. Network synchronization, diffusion, and the paradox of heterogeneity. *Phys. Rev. E* **71,** 016116 (2005).
65. Gómez-Gardeñes, J., Campillo, M., Floria, L. M. & Moreno, Y. Dynamical organization of cooperation in complex topologies. *Phys. Rev. Lett.* **98,** 108103 (2007).
66. Korniss, G. Synchronization in weighted uncorrelated complex networks in a noisy environment: Optimization and connections with transport efficiency. *Phys. Rev. E* **75,** 051121 (2007).
67. Arenas, A., Díaz-Guilera, A. & Guimerà, R. Communication in networks with hierarchical branching. *Phys. Rev. Lett.* **86,** 3196–3199 (2001).
68. Guimerà, R., Arenas, A., Díaz-Guilera, A. & Giralt, F. Dynamical properties of model communication networks. *Phys. Rev. E* **66,** 026704 (2002).
69. Sreenivasan, S., Cohen, R., López, E., Toroczkai, Z. & Stanley, H. E. Structural bottlenecks for communication in networks. *Phys. Rev. E* **75,** 036105 (2007).
70. Castellano, C., Loreto, V., Barrat, A., Cecconi, F. & Parisi, D. Comparison of voter and Glauber ordering dynamics on networks. *Phys. Rev. E* **71,** 066107 (2005).
71. Sood, V. & Redner, S. Voter model on heterogeneous graphs. *Phys. Rev. Lett.* **94,** 178701 (2005).
72. Suchecki, K., Eguíluz, V. M. & San Miguel, M. Conservation laws for the voter model in complex networks. *Europhys. Lett.* **69,** 228–234 (2005).
73. Klemm, K., Eguíluz, V. M., Toral, R. & San Miguel, M. Nonequilibrium transitions in complex networks: A model of social interaction. *Phys. Rev. E.* **67,** 026120 (2003).
74. Santos, F. C., Pacheco, J. M. & Lenaerts, T. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl Acad. Sci. USA* **103,** 3490–3494 (2006).
75. van Kampen, N. G. *Stochastic Processes in Physics and Chemistry* (North-Holland, 1981).
76. Bolker, B. M. & Grenfell, T. Chaos and biological complexity in measles dynamics. *Proc. Trans. R. Soc. Lond. B* **251,** 75–81 (1993).
77. Keeling, M. J. & Rohani, P. Estimating spatial coupling in epidemiological systems: A mechanistic approach. *Ecol. Lett.* **5,** 20–29 (2002).
78. Sattenspiel, L. & Dietz, K. A structured epidemic model incorporating geographic mobility among regions. *Math. Biosci.* **128,** 71–91 (1995).
79. Watts, D., Muhamad, R., Medina, D. C. & Dodds, P. S. Multiscale resurgent epidemics in a hierarchical metapopulation model. *Proc. Natl Acad. Sci. USA* **102,** 11157–11162 (2005).
80. Turing, A. M. The chemical basis of morphogenesis. *Phil. Trans. R. Soc. Lond.* **B237,** 37–72 (1952).
81. Nakao, H. & Mikhailov, A. S. Turing patterns in network-organized activator-inhibitor systems. *Nature Phys.* **6,** 544–550 (2010).
82. Colizza, V., Pastor-Satorras, R. & Vespignani, A. Reaction–diffusion processes and metapopulation models in heterogeneous networks. *Nature Phys.* **3,** 276–282 (2007).
83. Colizza, V. & Vespignani, A. Invasion threshold in heterogeneous metapopulation networks. *Phys. Rev. Lett.* **99,** 148701 (2007).
84. Colizza, V. & Vespignani, A. Epidemic modeling in metapopulation systems with heterogeneous coupling pattern: Theory and simulations. *J. Theor. Biol.* **251,** 450–467 (2008).
85. Barthélemy, M., Godrèche, C. & Luck, J-M. Fluctuation effects in metapopulation models: Percolation and pandemic threshold. *J. Theor. Biol.* **267,** 554–564 (2010).
86. Saldana, J. Continuous-time formulation of reaction–diffusion processes on heterogeneous metapopulations. *Phys. Rev. E* **78,** 012902 (2008).
87. Ni, S. & Weng, W. Impact of travel patterns on epidemic dynamics in heterogeneous spatial metapopulation networks. *Phys. Rev. E* **79,** 016111 (2009).
88. Ben-Zion, Y., Cohena, Y. & Shnerba, N. M. Modeling epidemics dynamics on heterogenous networks. *J. Theor. Biol.* **264,** 197–204 (2010).
89. Balcan, D. & Vespignani, A. Phase transitions in contagion processes mediated by recurrent mobility patterns. *Nature Phys.* **7,** 581–586 (2011).
90. Belik, V., Geisel, T. & Brockmann, D. Natural human mobility patterns and spatial spread of infectious diseases. *Phys. Rev. X* **1,** 011001 (2011).
91. Cooper, B. S., Pitman, R. J., Edmunds, W. J. & Gay, N. J. Delaying the international spread of pandemic influenza. *PLoS Med.* **3,** e12 (2006).
92. Hollingsworth, T. D., Ferguson, N. M. & Anderson, R. M. Will travel restrictions control the international spread of pandemic influenza? *Nature Med.* **12,** 497–499 (2006).
93. Hufnagel, L., Brockmann, D. & Geisel, T. Forecast and control of epidemics in a globalized world. *Proc. Natl Acad. Sci. USA* **101,** 15124–15129 (2004).
94. Eubank, S. *et al.* Modelling disease outbreaks in realistic urban social networks. *Nature* **429,** 180–184 (2004).
95. Longini, I. M. *et al.* Containing pandemic infleunza at the source. *Science* **309,** 1083–1087 (2005).
96. Ferguson, N. M. *et al.* Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* **437,** 209–211 (2005).
97. Colizza, V., Barrat, A., Barthlemy, M., Valleron, M. A. J. & Vespignani, A. Modeling the worldwide spread of pandemic influenza: Baseline case and containment interventions. *PLoS Med.* **4,** e13 (2007).
98. Balcan, D. *et al.* Seasonal transmission potential and activity peaks of the new influenza A(H1N1): A Monte Carlo likelihood analysis based on human mobility. *BMC Med.* **7,** 45 (2009).
99. Merler, S., Ajelli, M., Pugliese, A. & Ferguson, N. M. Determinants of the spatiotemporal dynamics of the 2009 H1N1 pandemic in Europe: Implications for real-time modelling. *PLoS Comput. Biol.* **7,** e1002205 (2011).
100. Gladwell, M. *The Tipping Point: How Little Things Can Make a Big Difference* (Little, Brown and Company, 2002).
101. Helbing, D. & Yu, W. The outbreak of cooperation among success-driven individuals under noisy condition. *Proc. Natl Acad. Sci. USA* **106,** 3680–3685 (2009).
102. Xie, J. *et al.* Social consensus through the influence of commited minorities. *Phys. Rev. E* **84,** 011130 (2011).
103. Morris, M. & Kretzschmar, M. Concurrent partnerships and the spread of HIV. *AIDS* **11,** 641–648 (1997).
104. Moody, J. The importance of relationship timing for diffusion: Indirect connectivity and STD infection risk. *Soc. Forces* **81,** 25–56 (2002).
105. Isella, L. *et al.* What's in a crowd? Analysis of face-to-face behavioral networks. *J. Theor. Biol.* **271,** 166–180 (2011).
106. Volz, E. & Meyers, L. A. Epidemic thresholds in dynamic contact networks. *J. R. Soc. Interface* **6,** 233–241 (2009).
107. Holme, P. & Newman, M. E. J. Nonequilibrium phase transition in the coevolution of networks and opinions. *Phys. Rev. E* **74,** 056108 (2006).
108. Centola, D., Gonzalez-Avella, J. C., Eguiluz, V. M. & San Miguel, M. Homophily, cultural drift, and the co-evolution of cultural groups. *J. Conflict Resolution* **51,** 905–929 (2007).
109. Funk, S., Salathé, M. & Jansen, V. A. A. Modelling the inuence of human behaviour on the spread of infectious diseases: A review. *J. R. Soc. Interface* **7,** 1247–1256 (2010).
110. Perra, N., Balcan, D., Goncalves, B. & Vespignani, A. Towards a characterization of behavior–disease models. *PLoS ONE* **6,** e23084 (2011).
111. Bauch, C. T. & Earn, D. J. Vaccination and the theory of games. *Proc. Natl Acad. Sci. USA* **101,** 13391–13394 (2004).
112. Liu, Y-Y., Slotine, J-J. & Barabasi, A-L. Controllability of complex networks. *Nature* **473,** 167–173 (2011).
113. Conover, M. *et al. Proc. 5th International Conference on Weblogs and Social Media (ICWSM)* 89–96 (2011).
114. Ratkiewicz, J. *et al. Proc. 20th International Conference Companion on World Wide Web (WWW '11)* 249–252 (ACM, 2001).
115. Kim, B. J., Yoon, C. N., Han, S. K. & Jeong, H. Path finding strategies in scale-free networks. *Phys. Rev. E* **65,** 027103 (2002).
116. Adamic, L. A., Lukose, R. M., Puniyani, A. R. & Huberman, B. A. Search in power-law networks. *Phys. Rev. E* **64,** 046135 (2001).
117. Brin, S. & Page, L. The anatomy of a large-scale hypertextual Web search engine. *Comput. Netw. ISDN Syst.* **30,** 107–117 (1998).
118. Bajardi, P. *et al.* Human mobility networks, travel restrictions, and the global spread of 2009 H1N1 pandemic. *PLoS ONE* **6,** e16591 (2011).

Acknowledgements

I thank B. Goncalves and N. Perra for their help with the figures and a critical reading of the manuscript. This work has been partially funded by the NIH R21-DA024259, DTRA-1-0910039 and NSF CCF-1101743 and NSF CMMI-1125095 awards. The work has been also partly sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the US Government.

Additional information

The authors declare no competing financial interests. Reprints and permissions information is available online at http://www.nature.com/reprints.

PROGRESS ARTICLE | INSIGHT
PUBLISHED ONLINE: 22 DECEMBER 2011 | DOI: 10.1038/NPHYS2180
nature physics

Networks formed from interdependent networks

Jianxi Gao[1,2], Sergey V. Buldyrev[3], H. Eugene Stanley[1] and Shlomo Havlin[4]*

Complex networks appear in almost every aspect of science and technology. Although most results in the field have been obtained by analysing isolated networks, many real-world networks do in fact interact with and depend on other networks. The set of extensive results for the limiting case of non-interacting networks holds only to the extent that ignoring the presence of other networks can be justified. Recently, an analytical framework for studying the percolation properties of interacting networks has been developed. Here we review this framework and the results obtained so far for connectivity properties of 'networks of networks' formed by interdependent random networks.

The interdisciplinary field of network science has attracted a great deal of attention in recent years[1–30]. This development is based on the enormous number of data that are now routinely being collected, modelled and analysed, concerning social[31–39], economic[14,36,40,41], technological[40,42–48] and biological[9,13,49,50] systems. The investigation and growing understanding of this extraordinary volume of data will enable us to make the infrastructures we use in everyday life more efficient and more robust.

The original model of networks, random graph theory, was developed in the 1960s by Erdős and Rényi, and is based on the assumption that every pair of nodes is randomly connected with the same probability, leading to a Poisson degree distribution. In parallel, in physics, lattice networks, where each node has exactly the same number of links, have been studied to model physical systems. Although graph theory is a well-established tool in the mathematics and computer science literature, it cannot describe well modern, real-life networks. Indeed, the pioneering 1999 observation by Barabasi[2], that many real networks do not follow the Erdős–Rényi model but that organizational principles naturally arise in most systems, led to an overwhelming accumulation of supporting data, new models and computational and analytical results, and to the emergence of a new science, that of complex networks.

Complex networks are usually non-homogeneous structures that in many cases obey a power-law form in their degree (that is, number of links per node) distribution. These systems are called scale-free networks. Real networks that can be approximated as scale-free networks include the Internet[3], the World Wide Web[4], social networks[31–39] representing the relations between individuals, infrastructure networks such as those of airlines[51], networks in biology[9,13,49,50], in particular networks of protein–protein interactions[10], gene regulation and biochemical pathways, and networks in physics, such as polymer networks or the potential-energy-landscape network. The discovery of scale-free networks led to a re-evaluation of the basic properties of networks, such as their robustness, which exhibit a drastically different character than those of Erdős–Rényi networks. For example, whereas homogeneous Erdős–Rényi networks are extremely vulnerable to random failures, heterogeneous scale-free networks are remarkably robust[4,5]. A great part of our current knowledge on networks is based on ideas borrowed from statistical physics, such as percolation theory, fractals and scaling analysis. An important property of these infrastructures is their stability, and it is thus important that we understand and quantify their robustness in terms of node and link failures. Percolation theory was introduced to study network stability and predicted the critical percolation threshold[5]. The robustness of a network is usually either characterized by the value of the critical threshold analysed using percolation theory[52] or defined as the integrated size of the largest connected cluster during the entire attack process[53]. The percolation approach was also proved to be extremely useful in addressing other scenarios, such as efficient attacks or immunization[6,7,54,55], and for obtaining optimal paths[56] as well as for designing robust networks[53]. Network concepts have also proven to be useful for the analysis and understanding of the spread of epidemics[57,58], and the organizational laws of social interactions, such as friendships[59,60] or scientific collaborations[61,62]. Ref. 63 investigated topologically biased failure in scale-free networks network and control of the robustness or fragility through fine-tuning of the topological bias in the failure process.

A large number of new measures and methods have been developed to characterize network properties, including measures of node clustering, network modularity, correlation between degrees of neighbouring nodes, measures of node importance and methods for the identification and extraction of community structures. These measures demonstrated that many real networks, and in particular biological networks, contain network motifs—small specific subnetworks—that occur repeatedly and provide information about functionality[9]. Dynamical processes, such as flow and electrical transport in heterogeneous networks, were shown to be significantly more efficient when compared with Erdős–Rényi networks[64,65]. Furthermore, it was shown that networks can also possess self-similar properties, so that under proper coarse graining (or, renormalization) of the nodes the network properties remain invariant[19].

However, these complex systems were mainly modelled and analysed as single networks that do not interact with or depend on other networks. In interacting networks, the failure of nodes in one network generally leads to the failure of dependent nodes in other networks, which in turn may cause further damage to the first network, leading to cascading failures and catastrophic consequences. It is known, for example, that blackouts in various countries have been the result of cascading failures between interdependent systems such as communication and power grid systems[67,68]. Furthermore, different kinds of critical infrastructure are also coupled together, such as systems of water and food supply, communications, fuel, financial transactions and power generation and transmission. Modern technology has

[1]Center for Polymer Studies and Department of Physics, Boston University, Boston, Massachusetts 02215, USA, [2]Department of Automation, Shanghai Jiao Tong University, 800 Dongchuan Road, Shanghai 200240, China, [3]Department of Physics, Yeshiva University, New York, New York 10033, USA, [4]Department of Physics, Bar-Ilan University, 52900 Ramat-Gan, Israel. *e-mail: havlin@ophir.ph.biu.ac.il.

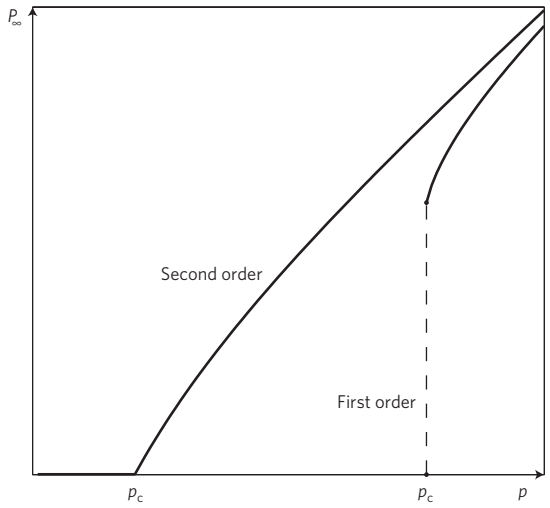

Figure 1 | Schematic demonstration of first- and second-order percolation transitions. In the second-order case, the giant component is continuously approaching zero at the percolation threshold $p = p_c$. In the first-order case, the giant component approaches zero discontinuously.

produced infrastructures that are becoming more and more interdependent, and understanding how robustness is affected by these interdependences is one of the main challenges faced when designing resilient infrastructures[67,69–72]. In recent decades, research was carried out in applied science on cataloguing, analysing and modelling the interdependences in critical infrastructure as well as modelling cascading failures in coupled critical infrastructure networks[40,42–48]. However, no systematic mathematical framework, such as percolation theory, is currently available for adequately addressing the consequences of disruptions and failures occurring simultaneously in interdependent critical infrastructures.

Recently, motivated by the fact that modern, crucially important infrastructures significantly interact, a mathematical framework was developed[73] to study percolation in a system of two interdependent networks subject to cascading failure. The analytical framework is based on a generating-function formalism widely used for studies of percolation and structure within a single network[73–75]. The framework for interdependent networks enables us to follow the dynamics of the cascading failures as well as to derive the analytic solutions for the final steady state. It was found[73] that certain types of interdependent network were significantly more vulnerable than their non-interacting counterparts. The failure of even a small number of elements within a single network may trigger a catastrophic cascade of events that destroys the global connectivity. For a fully interdependent case in which each node in one network depends on a functioning node in other networks and vice versa, a first-order discontinuous phase transition, which is dramatically different from the second-order continuous phase transition found in isolated networks (Fig. 1), was found[73]. This phenomenon is caused by the presence of two types of link: connectivity links within each network; and dependence links between networks. Connectivity links enable the network to carry out its function and dependence links represent the fact that the function of a given node in one network depends crucially on nodes in other networks. The case of connectivity links between the different networks was studied in ref. 66. It was shown[76] that, when the dependence coupling between the networks is reduced, at a critical coupling strength the percolation transition becomes second order.

More recently, two important generalizations of the basic model of ref. 73 have been developed:

One generalization takes into account that in real-world scenarios the initial failure of important nodes (or hubs) may be not random but targeted. A mathematical framework for understanding the robustness of interdependent networks under an initial targeted attack has been studied in ref. 77. The authors of that work developed a general technique that uses the random-attack problem to map the targeted-attack problem in interdependent networks.

The other generalization takes into account that, in real-world scenarios, the assumption that each node in network A depends on one and only one node in network B and vice versa may not be valid. To correct this shortcoming, a theoretical framework for understanding the robustness of interdependent networks with a random number of support and dependence relationships has been developed and studied[78].

In all of the above studies[73,76–78], the dependent pairs of nodes in both networks were chosen randomly. Thus when high-degree nodes in one network depend with a high probability on low-degree nodes of another network the configuration becomes vulnerable. To quantify and better understand this phenomenon, we proposed two 'intersimilarity' measures between the interdependent networks[79]. On the one hand, intersimilarity occurs in interdependent networks when nodes with similar degrees tend to be interdependent. On the other hand, it occurs if the neighbours of interdependent nodes in each network also tend to be interdependent. Refs 79–81 found that as the interdependent networks become more intersimilar the system becomes more robust. A system composed of an interdependent world-wide seaport and airport networks and the world-wide airport network was studied in ref. 79, where it was found that well-connected seaports tend to couple with well-connected airports, and two ways of measuring the intersimilarity of interdependent networks were developed. The case in which all pairs of interdependent nodes in both networks have the same degree was solved analytically in ref. 82.

The robustness of a two-coupled-networks system has been studied for dependence coupling[73] and for connectivity coupling[66]. Very recently a more realistic coupled network system with both dependence and connectivity links between the coupled networks was studied[83]. Using a percolation approach, rich and unusual phase transition phenomena were found, including a mixed first-order and second-order hybrid transition. This hybrid transition shows that a discontinuous jump in the size of the giant component (as in a first-order transition) is followed by a continuous decrease to zero (as in a second-order transition).

Previous studies of isolated networks in which dependence links cause cascading failure fall into two categories:

The first studies failures due to network overload when the network flow is a physical quantity, for example, in power transmission systems, transportation networks, or Internet traffic[84–87]. The models produced by these studies demonstrate that when an overloaded node stops traffic flow, the choosing of alternative paths can overload other nodes, and a cascading failure that disables the entire network can result.

The second is studies that produce models based on local dependences, such as the decision-making of interacting agents[11]. In these models the state of a node depends on the state of its neighbours, that is, a failing node will cause its neighbours to also fail.

The rich phenomena found in interdependent networks and the insights obtained from the percolation framework developed in refs 73,76 have led to a better understanding of the effect of dependence links within single isolated networks. A percolation approach for a single network in the presence of random dependence links was developed recently[88–90]. The results show that cascading failures occur, yielding a first-order transition, and that

the percolation threshold of the network significantly increases with an increase in the number of dependence links.

Generating functions for a single network

We begin by describing the generating-function formalism[74] for a single network that will also be useful in studying interdependent networks. We assume that all N_i nodes in network i are randomly assigned a degree k from a probability distribution $P_i(k)$, and are randomly connected with the only constraint that the node with degree k has exactly k links[91]. We define the generating function of the degree distribution

$$G_i(x) \equiv \sum_{k=0}^{\infty} P_i(k) x^k \tag{1}$$

where x is an arbitrary complex variable. Using equation (1), the average degree of network i is

$$\langle k \rangle_i = \sum_{k=0}^{\infty} k P_i(k) = \left. \frac{\partial G_i}{\partial x} \right|_{x \to 1} = G_i'(1) \tag{2}$$

In the limit of infinitely large networks $N_i \to \infty$, the random connection process can be modelled as a branching process in which an outgoing link of any node has a probability $kP_i(k)/\langle k \rangle_i$ of being connected to a node with degree k, which in turn has $k-1$ outgoing links. Using equations (1) and (2), the generating function of this branching process is defined as

$$H_i(x) \equiv \frac{\sum_{k=0}^{\infty} P_i(k) k x^{k-1}}{\langle k \rangle_i} = \frac{G_i'(x)}{G_i'(1)} \tag{3}$$

Let f_i be the probability that a randomly selected link does not lead to the giant component. If a link leads to a node with $k-1$ outgoing links this probability is f_i^{k-1}. Thus $H_i(f_i)$ also has the meaning that a randomly selected link does not lead to the giant component and hence f_i satisfies the recursive relation equation $f_i = H_i(f_i)$. The probability that a node with degree k does not belong to the giant component is f_i^k and hence the probability that a randomly selected node belongs to the giant component is $g_i = 1 - G_i(f_i)$.

Once a fraction $1-p$ of nodes is randomly removed from a network, the generating function remains the same, but with a new argument $z_i \equiv px + 1 - p$ (ref. 75). Accordingly, owing to the definition of f_i and g_i the probability that a randomly chosen surviving node belongs to a giant component is given by

$$g_i(p) = 1 - G_i[pf_i(p) + 1 - p] \tag{4}$$

where $f_i(p)$ satisfies

$$f_i(p) = H_i[pf_i(p) + 1 - p] \tag{5}$$

Thus $P_{\infty,i}$, the fraction of nodes that belongs to the giant component, is given by the product[75]

$$P_{\infty,i} = pg_i(p) \tag{6}$$

As p decreases, the non-trivial solution $f_i < 1$ of equation (5) gradually approaches the trivial solution $f_i = 1$. Accordingly, $P_{\infty,i}$ gradually approaches zero as in a second-order phase transition and becomes zero when two solutions of equation (5) coincide at $p = p_c$. At this point the straight line corresponding to the left-hand side of equation (5) becomes tangent to the curve corresponding to its right-hand side, yielding

$$p_c = 1/H'_i(1) \tag{7}$$

For example, for Erdős–Rényi networks[92–94], characterized by a Poisson degree distribution, using equations (1), (3) and (7) we obtain

$$G_i(x) = H_i(x) = \exp[\langle k \rangle_i (x-1)] \tag{8}$$

$$g_i(p) = 1 - f_i(p) \tag{9}$$

$$f_i(p) = \exp\{p\langle k \rangle_i [f_i(p) - 1]\} \tag{10}$$

and using equations (7) and (8)

$$p_c = \frac{1}{\langle k \rangle_i} \tag{11}$$

Finally, using equations (6), (9) and (10), we obtain a direct equation for $P_{\infty,i}$

$$P_{\infty,i} = p[1 - \exp(-\langle k \rangle_i P_{\infty,i})] \tag{12}$$

Framework of two partially interdependent networks

A generalization of the percolation theory of two fully interdependent networks[73] has been developed by Parshani *et al.*[76], where a more realistic case of a pair of partially interdependent networks has been studied. In this case, both interacting networks have a certain fraction of completely autonomous nodes whose function does not directly depend on the nodes of the other network. It has been found that, once the fraction of autonomous nodes increases above a certain threshold, the abrupt collapse of the interdependent networks characterized by a first-order transition observed in ref. 73 changes, at a critical coupling strength, to a continuous second-order transition as in classical percolation theory[52].

In the following we describe in more detail the framework developed in ref. 76. This framework consists of two networks A and B with the numbers of nodes N_A and N_B, respectively. Within network A, the nodes are randomly connected by A edges with degree distribution $P_A(k)$, whereas the nodes in network B are randomly connected by B edges with degree distribution $P_B(k)$. The average degrees of the networks A and B are a and b respectively. In addition, a fraction q_A of network A nodes depends on the nodes in network B and a fraction q_B of network B nodes depends on the nodes in network A. We assume that a node from one network depends on no more than one node from the other network, and if node A_i depends on node B_j, and B_j depends on A_k, then $k = i$. The latter condition, which we call a no-feedback condition (Fig. 2), excludes configurations that completely collapse even for fully interdependent networks once a single node is removed[78]. We assume that the initial removal of nodes from network A is a fraction $1-p$.

Next we present the formalism for the cascade process step by step (Fig. 3). After an initial removal of nodes, the remaining fraction of nodes in network A is $\psi_1' \equiv p$. The initial removal of nodes will disconnect some nodes from the giant component. The remaining functional part of network A therefore constitutes a fraction $\psi_1 = \psi_1' g_A(\psi_1')$ of the network nodes, where $g_A(\psi_1')$ is defined by equations (4) and (5). As a fraction q_B of nodes from network B depends on nodes from network A, the number of nodes in network B that become non-functional is

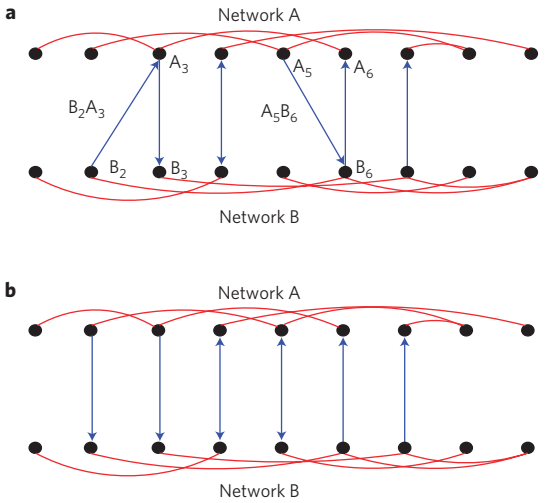

Figure 2 | Differences between the feedback condition and no-feedback condition. a,b In the case of feedback (**a**), node A_3 depends on node B_2, and node $B_3 \neq B_2$ depends on node A_3, whereas if there is no feedback (**b**) this is forbidden. The blue links between the two networks show the dependence links and the red links in each network show the connectivity links, which enable each network to function.

$(1-\psi_1)q_B = q_B[1-\psi'_1 g_A(\psi'_1)]$. Accordingly, the remaining fraction of network B nodes is $\phi'_1 = 1-q_B[1-\psi'_1 g_A(\psi'_1)]$, and the fraction of nodes in the giant component of network B is $\phi_1 = \phi'_1 g_B(\phi'_1)$.

Following this approach we can construct the sequence, ψ'_t and ϕ'_t, of the remaining fraction of nodes at each stage of the cascade of failures. The general form is given by

$$\begin{aligned} \psi'_1 &\equiv p \\ \phi'_1 &= 1-q_B[1-pg_A(\psi'_1)] \\ \psi'_t &= p[1-q_A(1-g_B(\phi'_{t-1}))] \\ \phi'_t &= 1-q_B[1-pg_A(\psi'_{t-1})] \end{aligned} \tag{13}$$

To determine the state of the system at the end of the cascade process we look at ψ'_τ and ϕ'_τ at the limit of $\tau \to \infty$. This limit must satisfy the equations $\psi'_\tau = \psi'_{\tau+1}$ and $\phi'_\tau = \phi'_{\tau+1}$ because eventually the clusters stop fragmenting and the fractions of randomly removed nodes at steps τ and $\tau+1$ are equal. Denoting $\psi'_\tau = x$ and $\phi'_\tau = y$, we arrive in the stationary state at a system of two equations with two unknowns,

$$\begin{aligned} x &= p\{1-q_A[1-g_B(y)]\} \\ y &= 1-q_B[1-g_A(x)p] \end{aligned} \tag{14}$$

The giant components of networks A and B at the end of the cascade of failures are, respectively, $P_{\infty,A} = \psi_\infty = xg_A(x)$ and $P_{\infty,B} = \phi_\infty = yg_B(y)$. Figure 4 shows the excellent agreement for the cascading failures in the giant component between computer simulations and the analytical results. The analytical results were obtained by recursive relations (13), where $g_A(\psi'_t)$ and $g_B(\phi'_t)$ are computed using equations (9) and (10).

Equation (14) can be illustrated graphically by two curves crossing in the (x,y) plane. For sufficiently large q_A and q_B the curves intersect at two points $(0<x_0, 0<y_0)$ and $(x_0<x_1<1, y_0<y_1<1)$. Only the second solution (x_1, y_1) has a physical meaning. As p decreases, the two solutions become closer to each other, remaining inside the unit square $(0<x<1; 0<y<1)$, and at a certain threshold $p=p_c$ they coincide: $0<x_0=x_1=x_c<1$, $0<y_0=y_1=y_c<1$.

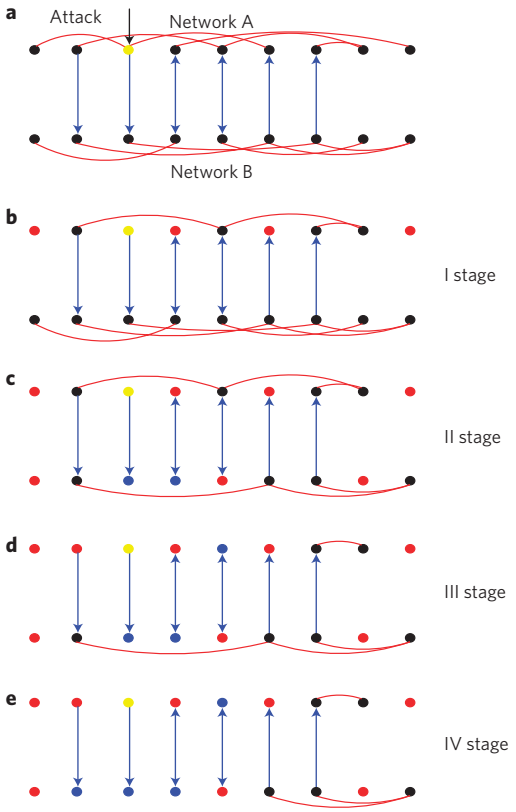

Figure 3 | Description of the dynamic process of cascading failures on two partially interdependent networks, which can be generalized to *n* partially interdependent networks. The black nodes represent the survival nodes, the yellow node represents the initially attacked node, the red nodes represent the nodes removed because they do not belong to the largest cluster and the blue nodes represent the nodes removed because they depend on the failed nodes in the other network. In each stage, for one network, we first remove the nodes that depend on the failed nodes in the other network or on the initially attacked nodes. Next we remove the nodes that do not belong to the largest cluster of the network.

For $p<p_c$ the non-trivial solution corresponding to the intersection abruptly disappears. Thus for sufficiently large q_A and q_B, $P_{\infty,A}$ and $P_{\infty,B}$ as a function of p show a first-order phase transition. As q_B decreases, the intersection of the curves moves out of the unit square; therefore, for small enough q_B, $P_{\infty,A}$ as a function of p shows a second-order phase transition. For the graphical representation of equation (14) and all possible solutions see Fig. 3 in ref. 76.

In a recent study[95], it was shown that a pair of interdependent networks can be designed to be more robust by choosing the autonomous nodes to be high-degree nodes. This choice mitigates the probability of catastrophic cascading failure.

Framework for a network of interdependent networks

In many real systems there are more than two interdependent networks, and diverse infrastructures—water and food supply networks, communication networks, fuel networks, financial transaction networks or power-station networks—can be coupled together[69,70]. Understanding the way system robustness is affected by such interdependences is one of the main challenges when designing resilient infrastructures.

Here we review the generalization of the theory of a pair of interdependent networks[73,76] to a system of n interacting networks[96], which can be graphically represented (Fig. 5) as a network of networks (NON). We develop an exact analytical

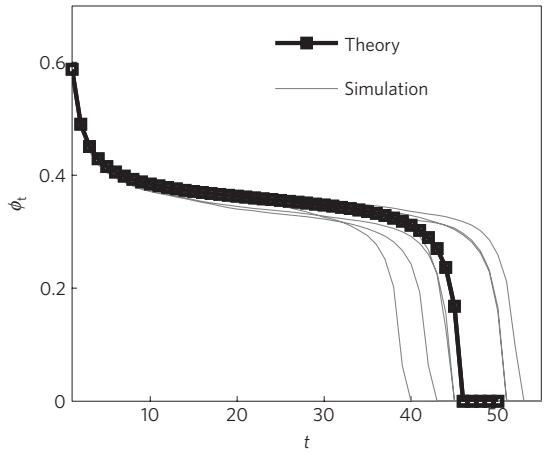

Figure 4 | Cascade of failures in two partially interdependent Erdős–Rényi networks. The giant component ϕ_t for every iteration of the cascading failures is shown for the case of a first-order phase transition with the initial parameters $p = 0.8505$, $a = b = 2.5$, $q_A = 0.7$ and $q_B = 0.8$. In the simulations, $N = 2 \times 10^5$ with over 20 realizations. The grey lines represent different realizations. The squares represent the average over all realizations and the black line is obtained from equation (13).

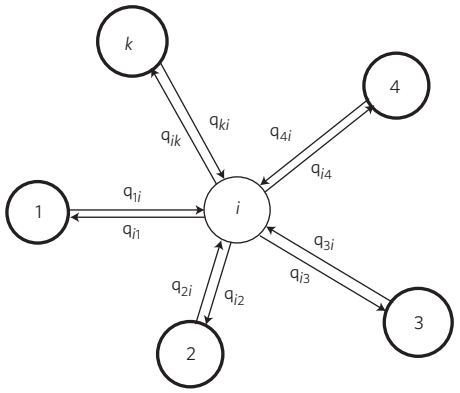

Figure 5 | Schematic representation of a NON. Circles represent interdependent networks, and the arrows connect the partially interdependent pairs. For example, a fraction of q_{3i} of nodes in network i depend on the nodes in network 3. The networks that are not connected by the dependence links do not have nodes that directly depend on one another.

approach for percolation of an NON system composed of n fully or partially interdependent randomly connected networks. The approach is based on analysing the dynamical process of the cascading failures. The results generalize the known results for percolation of a single network ($n = 1$) and the $n = 2$ result found in refs 73,76, and show that, whereas for $n = 1$ the percolation transition is a second-order transition, for $n > 1$ cascading failures occur and the transition becomes first order. Our results for n interdependent networks suggest that the classical percolation theory extensively studied in physics and mathematics is a limiting case of $n = 1$ of a general theory of percolation in NON. As we shall discuss here this general theory has many features that are not present in the classical percolation theory.

In our generalization, each node in the NON is a network itself and each link represents a fully or partially dependent pair of networks. We assume that each network i ($i = 1, 2, \ldots, n$) of the NON consists of N_i nodes linked together by connectivity links. Two networks i and j form a partially dependent pair if a certain fraction $q_{ji} > 0$ of nodes of network i directly depends on nodes of network j, that is, they cannot function if the nodes in network j on which they depend do not function. Dependent pairs are connected by unidirectional dependence links pointing from network j to network i. This convention symbolizes the fact that nodes in network i receive supply from nodes in network j of a crucial commodity, for example electric power if network j is a power grid.

We assume that after an attack or failure only a fraction of nodes p_i in each network i will remain. We also assume that only nodes that belong to a giant connected component of each network i will remain functional. This assumption helps explain the cascade of failures: nodes in network i that do not belong to its giant component fail, causing failures of nodes in other networks that depend on the failing nodes of network i. The failure of these nodes causes the direct failure of the dependent nodes in other networks, failures of isolated nodes in them and further failure of nodes in network i, and so on. Our goal is to find the fraction of nodes $P_{\infty,i}$ of each network that remain functional at the end of the cascade of failures as a function of all fractions p_i and all fractions q_{ij}. We assume that all networks in the NON are randomly connected networks characterized by a degree distribution of links $P_i(k)$, where k is a degree of a node in network i. We further assume that each node a in network i may depend with probability q_{ji} on only one node b in network j.

We can study different models of cascading failures in which we vary the survival time of the dependent nodes after the failure of the nodes in other networks on which they depend and the survival time of the disconnected nodes. We conclude that the final state of the networks does not depend on these details but can be described by a system of equations somewhat analogous to the Kirchhoff equations for a resistor network. This system of equations has n unknowns x_i. These represent the fractions of nodes that survive in network i after the nodes that fail in the initial attack are removed, and also the nodes depending on the failed nodes in other networks at the end of cascading failure are removed, but without considering yet the further failing of nodes due to the internal connectivity of the network. The final giant component of each network can be found from the equation $P_{\infty,i} = x_i g_i(x_i)$, where $g_i(x_i)$ is the fraction of the remaining nodes of network i that belong to its giant component given by equation (4).

First we shall discuss the more complex case of the no-feedback condition. The unknowns x_i satisfy the system of n equations,

$$x_i = p_i \prod_{j=1}^{K} [q_{ji} y_{ji} g_j(x_j) - q_{ji} + 1] \tag{15}$$

where the product is taken over the K networks interlinked with network i by the partial dependence links (Fig. 3) and

$$y_{ij} = \frac{x_i}{q_{ji} y_{ji} g_j(x_j) - q_{ji} + 1} \tag{16}$$

has the meaning of the fraction of nodes in network j that survive after the damage from all the networks connected to network j except network i is taken into account. The damage from network i must be excluded owing to the no-feedback condition. In the absence of the no-feedback condition, equation (15) becomes much simpler as $y_{ji} = x_j$. Equation (15) is valid for any case of interdependent NON, whereas equation (16) represents the no-feedback condition.

Four examples of a NON solvable analytically

In this section we present four examples that can be explicitly solved analytically: (1) a tree-like Erdős–Rényi fully dependent

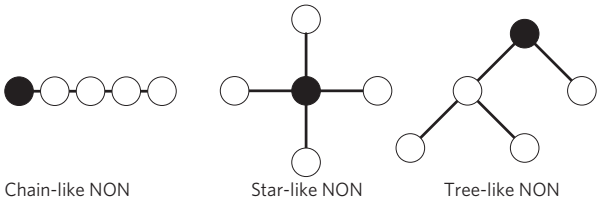

Figure 6 | Three types of loopless NON composed of five coupled networks. All have the same percolation threshold and the same giant component. The dark node represents the origin network on which failures initially occur.

NON, (2) a tree-like random regular fully dependent NON, (3) a loop-like Erdős–Rényi partially dependent NON and (4) a random regular network of partially dependent Erdős–Rényi networks. All cases represent different generalizations of percolation theory for a single network. In all examples except (3) we apply the no-feedback condition.

(1) We solve explicitly[96] the case of a tree-like NON (Fig. 6) formed by n Erdős–Rényi networks[92–94] with the same average degrees k, $p_1 = p$, $p_i = 1$ for $i \neq 1$ and $q_{ij} = 1$ (fully interdependent). From equations (15) and (16) we obtain an exact expression for the order parameter, the size of the mutual giant component for all p, k and n values,

$$P_\infty = p[1 - \exp(-kP_\infty)]^n \tag{17}$$

Equation (17) generalizes known results for $n = 1, 2$. For $n = 1$, we obtain the known result $p_c = 1/k$, equation (11), of an Erdős–Rényi network and $P_\infty(p_c) = 0$, which corresponds to a continuous second-order phase transition. Substituting $n = 2$ in equation (17) yields the exact results of ref. 73.

Solutions of equation (17) are shown in Fig. 7a for several values of n. The special case $n = 1$ is the known Erdős–Rényi second-order percolation law, equation (12), for a single network. In contrast, for any $n > 1$, the solution of (17) yields a first-order percolation transition, that is, a discontinuity of P_∞ at p_c.

Our results show (Fig. 7a) that the NON becomes more vulnerable with increasing n or decreasing k (p_c increases when n increases or k decreases). Furthermore, for a fixed n, when k is smaller than a critical number $k_{\min}(n)$, $p_c \geq 1$, meaning that for $k < k_{\min}(n)$ the NON will collapse even if a single node fails[96].

(2) In the case of a tree-like network of interdependent random regular networks[97], where the degree k of each node in each network is assumed to be the same, we obtain an exact expression for the order parameter, the size of the mutual giant component for all p, k and n values,

$$P_\infty = p\left\{1 - \left\{p^{\frac{1}{n}} P_\infty^{\frac{n-1}{n}}\left[\left(1 - \left(\frac{P_\infty}{p}\right)^{\frac{1}{n}}\right)^{\frac{k-1}{k}} - 1\right] + 1\right\}^k\right\}^n \tag{18}$$

Numerical solutions of equation (18) are in excellent agreement with simulations. Comparing with the results of the tree-like Erdős–Rényi NON, we find that the robustness of n interdependent random regular networks of degree k is significantly higher than that of the n interdependent Erdős–Rényi networks of average degree k. Moreover, whereas for an Erdős–Rényi NON there exists a critical minimum average degree $k = k_{\min}$ that increases with n (below which the system collapses), there is no such analogous $k_{\min}$ for the random regular NON system. For any $k > 2$, the random regular NON is stable, that is, $p_c < 1$. In general, this is correct for any network with any degree distribution, $P_i(k)$, such that $P_i(0) = P_i(1) = 0$, that is, for a network without disconnected or singly connected nodes[97].

(3) In the case of a loop-like NON (for dependences in one direction) of n Erdős–Rényi networks[96], all the links are unidirectional, and the no-feedback condition is irrelevant. If the initial attack on each network is the same, $1 - p$, $q_{i-1i} = q_{n1} = q$ and $k_i = k$, using equations (15) and (16) we obtain that P_∞ satisfies

$$P_\infty = p(1 - e^{-kP_\infty})(qP_\infty - q + 1) \tag{19}$$

Note that if $q = 1$ equation (19) has only a trivial solution $P_\infty = 0$, whereas for $q = 0$ it yields the known giant component of a single network, equation (12), as expected. We present numerical solutions of equation (19) for two values of q in Fig. 7b. Interestingly, whereas for $q = 1$ and tree-like structures equations (17) and (18) depend on n, for loop-like NON structures equation (19) is independent of n.

(4) For NONs where each ER network is dependent on exactly m other Erdős–Rényi networks (the case of a random regular network of Erdős–Rényi networks), we assume that the initial attack on each network is $1 - p$, and each partially dependent pair has the same q in both directions. The n equations of equation (15) are exactly the same owing to symmetries, and hence P_∞ can be obtained analytically,

$$P_\infty = \frac{p}{2^m}(1 - e^{-kP_\infty})[1 - q + \sqrt{(1-q)^2 + 4qP_\infty}]^m \tag{20}$$

from which we obtain

$$p_c = \frac{1}{k(1-q)^m} \tag{21}$$

Again, as in case (3), it is surprising that both the critical threshold and the giant component are independent of the number of networks n, in contrast to tree-like NON (equations (17) and (18)), but depend on the coupling q and on both degrees k and m. Numerical solutions of equation (20) are shown in Fig. 7c, and the critical thresholds p_c in Fig. 7c coincide with the theory, equation (21).

Remark on scale-free networks

The above examples regarding Erdős–Rényi and random regular networks have been selected because they can be explicitly solved analytically. In principle, the generating function formalism presented here can be applied to randomly connected networks with any degree distribution. The analysis of the scale-free networks with a power-law degree distribution $P(k) \sim k^{-\lambda}$ is extremely important, because many real networks can be approximated by a power-law degree distribution, such as the Internet, the airline network and social-contact networks, such as networks of scientific collaboration[2,10,51]. Analysis of fully interdependent scale-free networks[73] shows that, for interdependent scale-free networks, $p_c > 0$ even in the case $\lambda \leq 3$ for which in a single network $p_c = 0$. In general, for fully interdependent networks, the broader the degree distribution the greater p_c for networks with the same average degree[73]. This means that networks with a broad degree distribution become less robust than networks with a narrow degree distribution. This trend is the opposite of the trend found in non-interacting isolated networks. The explanation of this phenomenon is related to the fact that in randomly interdependent networks the hubs in one network may depend on poorly connected nodes in another. Thus the removal of a randomly selected node in one network may cause a failure of a hub in a second network, which in turn renders many singly connected

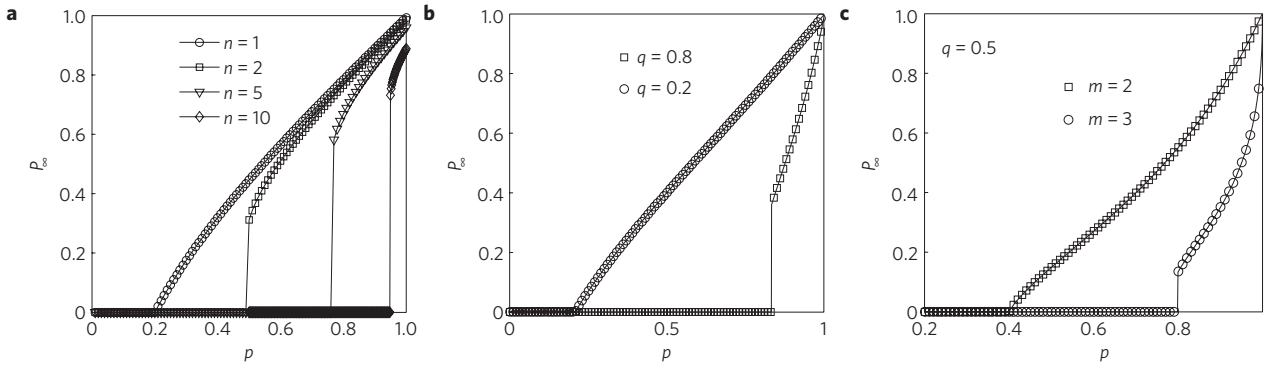

Figure 7 | The fraction of nodes in the giant component P_∞ as a function of p for three different examples. a, A tree-like fully ($q = 1$) interdependent NON; P_∞ is shown as a function of p for $k = 5$ and several values of n. The results are obtained using equation (17). Note that increasing n from $n = 2$ yields a first-order transition. **b**, A loop-like NON; P_∞ is shown as a function of p for $k = 6$ and two values of q. The results are obtained using equation (19). Note that increasing q yields a first-order transition. **c**, A random regular network of Erdős–Rényi networks; P_∞ is shown as a function of p, for two different values of m when $q = 0.5$. The results are obtained using equation (20), and the number of networks, n, can be any number with the condition that any network in the NON connects exactly to m other networks. Note that changing m from 2 to $m > 2$ changes the transition from second order to first order (for $q = 0.5$).

nodes non-functional, and the multiplying damage travels back to the first network. This explanation is corroborated by the analytical proof in ref. 82, which shows that if the degrees of the interdependent nodes coincide, then a network with a broader degree distribution will become more robust than a network with a narrower degree distribution, that is, the behaviour characteristic of non-interacting networks is restored. Ref. 82 also reports that, for fully interdependent scale-free networks with equal degrees of interdependent pairs, $p_c = 0$ for $\lambda < 3$. Moreover, the percolation transition is a discontinuous first-order phase transition if and only if $H_i'(1) < \infty$, that is, if the degree distribution has a finite second moment. For fully interdependent networks with uncorrelated degrees of interdependent nodes, the percolation transition is always a discontinuous phase transition[73,76]. These results, as well as the results of ref. 79, show the need to study more realistic situations in which the interdependent networks have various correlations in the dependences and connectivities. A recent study of partially interdependent scale-free networks shows that, although the giant component decreases significantly owing to cascading failures, p_c is always zero as long as $q < 1$ (D. Zhou *et al.*, unpublished).

Remaining challenges

We have reviewed recent studies of the robustness of a system of interdependent networks. In interacting networks, when a node in one network fails it usually causes dependent nodes in other networks to fail, which, in turn, may cause further damage in the first network and results in a cascade of failures with catastrophic consequences. Our analytical framework enables us to follow the dynamic process of the cascading failures step by step and to derive steady-state solutions. Interdependent networks appear in all aspects of life, nature and technology. Transportation systems include railway networks, airline networks and other transportation systems. Some properties of interacting transportation systems have been studied recently[79,80]. In the field of physiology, the human body can be regarded as a system of interdependent networks. Examples of such interdependent NON systems include the cardiovascular system, the respiratory system, the brain neuron system and the nervous system. In biology, the function of each protein is determined by its interacting proteins, which can be described by a network. As many proteins are involved in a number of different functions, the protein-interaction system can be regarded as a system of interacting networks. In the field of economics, networks of banks, insurance companies and business firms are interdependent.

Thus far, only a very few real-world interdependent systems have been analysed using the percolation approach[71,79,80]. We expect our present work to provide insights leading to a further analysis of real data on interdependent networks. The benchmark models we present here can be used to study the structural, functional and robustness properties of interdependent networks. Because, in real NONs, individual networks are not randomly connected and their interdependent nodes are not selected at random, it is crucial that we understand the many types of correlation that exist in real-world systems and that we further develop the theoretical tools to include such correlations. Further studies of interdependent networks should focus on an analysis of real data from many different interdependent systems and on the development of mathematical tools for studying real-world interdependent systems.

Many real-world networks are embedded in space, and the spatial constraints strongly affect their properties[30]. We need to understand how these spatial constraints influence the robustness properties of interdependent networks[79,80]. Other properties that influence the robustness of single networks, such as the dynamic nature of the configuration in which links or nodes appear and disappear and the directed nature of some links, as well as problems associated with degree–degree correlations and clustering, should be also addressed in future studies of coupled network systems. It is also important to investigate the case when a node in one network is supplied by multiple nodes in an interdependent network. In realistic interdependent pairs of networks i and j, a node in network i may depend on s supply nodes in network j and the total supply of a commodity received by this node from network j must be greater than a certain threshold s_c. In the case of $s_c = 0$ and random selection of the supply nodes, this problem was solved in ref. 78 for two interdependent networks, and this solution can be straightforwardly generalized for an arbitrary NON by replacing equation (15) with

$$x_i = p_i \prod_{j=1}^{K} \{1 - q_{ji} G^{ji}[1 - x_j g_j(x_j)]\} \qquad (22)$$

where $G^{ji}(x)$ is the generating function of the distribution of the supply degree s of nodes in network i that depend on the supply from nodes in network j. When $s = 1$ for all such nodes, $G^{ji}(x) = x$ and equation (22) reduces to equation (15) with $y_{ji} = x_j$, that is, in the absence of the no-feedback condition. More complex cases of multiple supply nodes await further investigation.

It is very important to find a way of improving the robustness of interdependent infrastructures. Our studies thus far show that

there are three methods to achieve this goal: increase the fraction of autonomous nodes[76], particularly nodes with high degree[95]; design the dependence links such that they connect the nodes with similar degrees[79,82] and protect the high-degree nodes against attack[95].

A coupled network in which the interlinks, that is, the links between different networks, are connectivity links was studied in ref. 66. The robustness of this system is greatly improved when compared with a system in which the interlinks are dependence links. A systematic study of the competing effects of a NON in which the interlinks are both dependence and connectivity interlinks is needed. Interesting results on a model containing both dependence and connectivity interlinks have been obtained[83]. Finally, we mention an early study of the Ising model on coupled networks[98]. Also, interacting networks with respect to climate systems were studied in ref. 99.

References

1. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. *Nature* **393,** 440–442 (1998).
2. Barabási, A. L. & Albert, R. Emergence of scaling in random networks. *Science* **286,** 509–512 (1999).
3. Faloutsos, M., Faloutsos, P. & Faloutsos, C. On power-law relationships of the internet topology. *Comput. Commun. Rev.* **29,** 378–382 (2000).
4. Albert, R., Jeong, H. & Barabási, A. L. Error and attack tolerance of complex networks. *Nature* **406,** 378–382 (2000).
5. Cohen, R., Erez, K., Ben-Avraham, D. & Havlin, S. Resilience of the Internet to random breakdown. *Phys. Rev. Lett.* **85,** 4626–4628 (2000).
6. Callaway, D. S., Newman, M. E. J., Strogatz, S. H. & Watts, D. J. Network robustness and fragility: Percolation on random graphs. *Phys. Rev. Lett.* **85,** 5468–5471 (2000).
7. Cohen, R., Erez, K., Ben-Avraham, D. & Havlin, S. Breakdown of the Internet under intentional attack. *Phys. Rev. Lett.* **86,** 3682–3685 (2001).
8. Strogatz, S. H. Exploring complex networks. *Nature* **410,** 268–276 (2001).
9. Milo, R. *et al.* Network motifs: Simple building blocks of complex networks. *Science* **298,** 824–827 (2002).
10. Albert, R. & Barabási, A. L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74,** 47–97 (2002).
11. Watts, D. J. A simple model of global cascades on random networks. *Proc. Natl Acad. Sci. USA* **99,** 5766–5771 (2002).
12. Newman, M. E. J. The structure and function of complex networks. *SIAM Rev.* **45,** 167–256 (2003).
13. Dorogovtsev, S. N. & Mendes, J. F. F. *Evolution of Networks: From Biological Nets to the Internet and WWW (Physics)* (Oxford Univ. Press, 2003).
14. Bonanno, G., Caldarelli, G., Lillo, F. & Mantegna, R. N. Topology of correlation-based minimal spanning trees in real and model markets. *Phys. Rev. E* **68,** 046130 (2003).
15. Barrat, A., Barthelemy, M., Pastor-Satorras, R. & Vespignani, A. The architecture of complex weighted networks. *Proc. Natl Acad. Sci. USA* **101,** 3747–3752 (2004).
16. Newman, M. E. J. & Girvan, M. Finding and evaluating community structure in networks. *Phys. Rev. E* **69,** 026113 (2004).
17. Satorras, R. P. & Vespignani, A. *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge Univ. Press, 2004).
18. Gallos, L. K., Cohen, R. & Argyrakis, P. *et al.* Stability and topology of scale-free networks under attack and defense strategies. *Phys. Rev. Lett.* **94,** 188701 (2005).
19. Song, C., Havlin, S. & Makse, H. A. Self-similarity of complex networks. *Nature* **433,** 392–395 (2005).
20. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Hwang, D. U. Complex networks: Structure and dynamics. *Phys. Rep.* **424,** 175–308 (2006).
21. Newman, M. E. J., Barabási, A-L. & Watts, D. J. *The Structure and Dynamics of Networks* (Princeton Univ. Press, 2006).
22. Caldarelli, G. & Vespignani, A. *Large Scale Structure and Dynamics of Complex Webs* (World Scientific, 2007).
23. Barrát, A., Barthélemy, M. & Vespignani, A. *Dynamical Processes on Complex Networks* (Cambridge Univ. Press, 2008).
24. Cohen, R. & Havlin, S. *Complex Networks: Structure, Robustness and Function* (Cambridge Univ. Press, 2010).
25. Kitsak, M. *et al.* Identification of influential spreaders in complex networks. *Nature Phys.* **6,** 888–893 (2010).
26. Newman, M. E. J. *Networks: An Introduction* (Oxford Univ. Press, 2010).
27. Cohen, R. & Havlin, S. *Complex Networks: Structure, Robustness, and Function* (Cambridge Univ. Press, 2010).
28. West, B. J. & Grigolini, P. *Complex Webs: Anticipating the Improbable* (Cambridge Univ. Press, 2011).
29. Barthélemy, M. Spatial networks. *Phys. Rep.* **499,** 1–101 (2011).
30. Li, D., Kosmidis, K., Bunde, A. & Havlin, S. Dimension of spatially embedded networks. *Nature Phys.* **7,** 481–484 (2011).
31. Snijders, T. A. B., Pattison, P. E., Robins, G. L. & Handcock, M. S. New specifications for exponential random graph models. *Sociol. Methodol.* **36,** 99–153 (2006).
32. Borgatti, S. P. Identifying sets of key players in a network. *Comput. Math. Org. Theor.* **12,** 21–34 (2006).
33. Onnela, J-P. *et al.* Structure and tie strengths in mobile communication networks. *Proc. Natl Acad. Sci. USA* **104,** 7332–7336 (2007).
34. Faust, K. & Zvezki, M. Comparing social networks: Size, density and local structure. *Linear Algebr. Appl.* **3,** 185–216 (2006).
35. Handcock, M. S., Raftery, A. E. & Tantrum, J. M. Model-based clustering for social networks. *J. R. Stat. Soc. A* **170,** 301–354 (2007).
36. Jackson, M. O. & Rogers, B. W. Meeting strangers and friends of friends: How random are social networks? *Am. Econom. Rev.* **97,** 890–915 (2007).
37. Kleinberg, J. The convergence of social and technological networks. *Commun. ACM* **51,** 66–72 (2008).
38. Liben-Nowell, D. & Kleinberg, J. Tracing information flow on a global scale using internet chain-letter data. *Proc. Natl Acad. Sci. USA* **105,** 4633–4638 (2008).
39. Borgatti, S. P., Mehra, A., Brass, D. & Labianca, G. Network analysis in the social sciences. *Science* **323,** 892–895 (2009).
40. Joost, R. Inoperability input–output modeling of disruptions to interdependent economic systems. *Syst. Eng.* **9,** 20–34 (2006).
41. Jackson, M. O. *Social and Economic Networks (Economics, Physics, Sociology)* (Princeton Univ. Press, 2008).
42. Zimmerman, R. Decision-making and the vulnerability of interdependent critical infrastructure. *2004 IEEE Int. Conf. Syst. Man Cybern.* **5,** 4059–4063 (2005).
43. Mendonca, D. & Wallace, W. A. Impacts of the 2001 World Trade Center attack on New York City critical infrastructures. *J. Infrast. Syst.* **12,** 260–270 (2006).
44. Robert, B., Morabito, L. & Christie, R. D. The operational tools for managing physical interdependencies among critical infrastructures. *Int. J. Crit. Infrastruct.* **4,** 353–367 (2008).
45. Reed, D. A., Kapur, K. C. & Christie, R. D. Methodology for assessing the resilience of networked infrastructure. *IEEE Syst. J.* **3,** 174–180 (2009).
46. Bagheri, E. & Ghorbani, A. A. UML-CI: A reference model for profiling critical infrastructure systems. *Inform. Syst. Front.* **12,** 115–139 (2009).
47. Mansson, D., Thottappillil, R., Backstrom, M. & Ludvika, H. V. V. Methodology for classifying facilities with respect to intentional EMI. *IEEE Trans. Electromagn. Compat.* **95,** 46–52 (2009).
48. Johansson, J. & Hassel, H. An approach for modelling interdependent infrastructures in the context of vulnerability analysis. *Reliab. Eng. Syst. Saf.* **95,** 1335–1344 (2010).
49. Alon, U. Biological networks: The tinkerer as an engineer. *Science* **301,** 1866–1867 (2003).
50. Khanin, R. & Wit, E. How scale-free are biological networks. *J. Comput. Biol.* **13,** 810–818 (2006).
51. Colizza, V., Barrat, A., Barthelemy, M. & Vespignani, A. Prediction and predictability of global epidemics: The role of the airline transportation network. *Proc. Natl Acad. Sci. USA* **103,** 2015–2020 (2006).
52. Bunde, A. & Havlin, S. *Fractals and Disordered Systems* (Springer, 1996).
53. Schneider, C. M., Araújo, N. A. M., Moreira, A. A., Havlin, S. & Herrmann, H. J. Mitigation of malicious attacks on networks. *Proc. Natl Acad. Sci. USA* **108,** 3838–3841 (2011).
54. Cohen, R., Havlin, S. & Ben-Avraham, D. Efficient immunization strategies for computer networks and populations. *Phys. Rev. Lett.* **91,** 247901 (2003).
55. Chen, Y., Paul, G., Havlin, S., Liljeros, F. & Stanley, H. E. Finding a better immunization strategy. *Phys. Rev. Lett.* **101,** 058701 (2008).
56. Braunstein, L. A., Buldyrev, S. V., Cohen, Havlin, S. & Stanley, H. E. Optimal paths in disordered complex networks. *Phys. Rev. Lett.* **91,** 168701 (2003).
57. Pastor-Satorras, R. & Vespignani, A. Epidemic spreading in scale-free network. *Phys. Rev. Lett.* **86,** 3200–3203 (2001).
58. Balcan, D. *et al.* Multiscale mobility networks and the large scale spreading of infectious diseases. *Proc. Natl Acad. Sci. USA* **106,** 21484–21489 (2009).
59. Palla, G., Derenyi, I., Farkas, I. & Vicsek, T. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* **435,** 814–818 (2005).
60. Kossinets, G. & Watts, D. Empirical analysis of an evolving social network. *Science* **311,** 88–90 (2006).
61. Newman, M. E. J. The structure of scientific collaboration networks. *Proc. Natl Acad. Sci. USA* **98,** 404–409 (2001).
62. Girvan, M. & Newman, M. E. J. Community structure in social and biological networks. *Proc. Natl Acad. Sci. USA* **99,** 7821–7826 (2002).
63. Moreira, A. A., Andrade, J. S. Jr, Herrmann, H. J. & Indekeu, J. O. How to make a fragile network robust and vice versa. *Phys. Rev. Lett.* **102,** 019701 (2009).

64. Lopez, E., Buldyrev, S. V., Havlin, S. & Stanley, H. E. Anomalous transport in scale-free networks. *Phys. Rev. Lett.* **94,** 248701 (2005).
65. Boguñá, M. & Krioukov, D. Navigating ultrasmall worlds in ultrashort time. *Phys. Rev. Lett.* **102,** 058701 (2009).
66. Leicht, E. A. & D'Souza, R. M. Percolation on interacting networks. Preprint at http://arxiv.org/abs/0907.0894. (2009).
67. Rosato, V. Modeling interdependent infrastructures using interacting dynamical models. *Int. J. Crit. Infrastruct.* **4,** 63–79 (2008).
68. US–Canada Power System Outage Task Force *Final Report on the August 14th 2003 Blackout in the United States and Canada: Causes and Recommendations* (The Task Force, 2004).
69. Peerenboom, J., Fischer, R. & Whitfield, R. in *Proc. CRIS/DRM/IIIT/NSF Workshop Mitigating the Vulnerability of Critical Infrastructures to Catastrophic Failures* (2001).
70. Rinaldi, S., Peerenboom, J. & Kelly, T. Identifying, understanding, and analyzing critical infrastructure interdepedencies. *IEEE Control. Syst. Magn.* **21,** 11–25 (2001).
71. Yağan, O., Qian, D., Zhang, J. & Cochran, D. Optimal allocation of interconnecting links in cyber-physical systems: Interdependence, cascading failures and robustness. http://www.ece.umd.edu/~oyagan/Journals/Interdependent_Journal.pdf (2011).
72. Vespignani, A. The fragility of interdependency. *Nature* **464,** 984–985 (2010).
73. Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E. & Havlin, S. Catastrophic cascade of failures in interdependent networks. *Nature* **464,** 1025–1028 (2010).
74. Newman, M. E. J., Strogatz, S. H. & Watts, D. J. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E* **64,** 026118 (2001).
75. Shao, J., Buldyrev, S. V., Braunstein, L. A., Havlin, S. & Stanley, H. E. Structure of shells in complex networks. *Phys. Rev. E* **80,** 036105 (2009).
76. Parshani, R., Buldyrev, S. V. & Havlin, S. Interdependent networks: Reducing the coupling strength leads to a change from a first to second order percolation transition. *Phys. Rev. Lett.* **105,** 048701 (2010).
77. Huang, X., Gao, J., Buldyrev, S. V., Havlin, S. & Stanley, H. E. Robustness of interdependent networks under targeted attack. *Phys. Rev. E (R)* **83,** 065101 (2011).
78. Shao, J., Buldyrev, S. V., Havlin, S. & Stanley, H. E. Cascade of failures in coupled network systems with multiple support-dependence relations. *Phys. Rev. E* **83,** 036116 (2011).
79. Parshani, R., Rozenblat, C., Ietri, D., Ducruet, C. & Havlin, S. Inter-similarity between coupled networks. *Europhys. Lett.* **92,** 68002–68006 (2010).
80. Gu, C. *et al.* Onset of cooperation between layered networks. *Phys. Rev. E* **84,** 026101 (2011).
81. Cho, W., Coh, K. & Kim, I. Correlated couplings and robustness of coupled networks. Preprint at http://arxiv.org/abs/1010.4971. (2010).
82. Buldyrev, S. V., Shere, N. W. & Cwilich, G. A. Interdependent networks with identical degrees of mutually dependent nodes. *Phys. Rev. E* **83,** 016112 (2011).
83. Hu, Y., Ksherim, B., Cohen, R. & Havlin, S. Percolation in interdependent and interconnected networks: Abrupt change from second to first order transition. *Phys. Rev. E* (in the press). Preprint at http://arxiv.org/abs/1106.4128 (2011).
84. Sachtjen, M. L., Carreras, B. A. & Lynch, V. E. Disturbances in a power transmission system. *Phys. Rev. E* **61,** 4877–4882 (2000).
85. Motter, A. E. & Lai, Y. C. Cascade-based attacks on complex networks. *Phys. Rev. E* **66,** 065102 (2002).
86. Moreno, Y., Pastor, S. R., Vázquez, A. & Vespignani, A. Critical load and congestion instabilities in scale-free networks. *Europhys. Lett.* **62,** 292–298 (2003).
87. Motter, A. E. Cascade control and defense in complex networks. *Phys. Rev. Lett.* **93,** 098701 (2004).
88. Parshani, R., Buldyrev, S. V. & Havlin, S. Critical effect of dependency groups on the function of networks. *Proc. Natl Acad. Sci. USA* **108,** 1007–1010 (2011).
89. Bashan, A., Parshani, R. & Havlin, S. Percolation in networks composed of connectivity and dependency links. *Phys. Rev. E* **83,** 051127 (2011).
90. Bashan, A. & Havlin, S. The combined effect of connectivity and dependency links on percolation of networks. *J. Stat. Phys.* **145,** 686–695 (2011).
91. Molloy, M. & Reed, B. The size of the giant component of a random graph with a given degree sequence. *Combin. Probab. Comput.* **7,** 295–305 (1998).
92. Erdős, P. & Rényi, A. On random graphs I. *Publ. Math.* **6,** 290–297 (1959).
93. Erdős, P. & Rényi, A. On the evolution of random graphs. *Inst. Hung. Acad. Sci.* **5,** 17–61 (1960).
94. Bollobás, B. *Random Graphs* (Academic, 1985).
95. Schneider, C. M., Araújo, N. A. M., Havlin, S. & Herrmann, H. J. Towards designing robust coupled networks. Preprint at http://arxiv.org/abs/1106.3234 (2011).
96. Gao, J., Buldyrev, S. V., Havlin, S. & Stanley, H. E. Robustness of a network of networks. *Phys. Rev. Lett.* **107,** 195701 (2011).
97. Gao, J, Buldyrev, S. V., Havlin, S. & Stanley, H. E. Robustness of a tree-like network of interdependent networks. Preprint at http://arxiv.org/abs/1108.5515 (2011).
98. Suchecki, K. & Holyst, J. A. Ising model on two connected Barabasi–Albert networks. *Phys. Rev. E* **74,** 011122 (2006).
99. Donges, J. F., Schultz, H. C. H., Marwan, N., Zou, Y. & Kurths, J. Investigating the topology of interacting networks. *Eur. Phys. J. B* (2011, in the press).

Acknowledgements

We thank R. Parshani for helpful discussions. We thank the DTRA (Defense Threat Reduction Agency) and the Office of Naval Research for support. J.G. also thanks the Shanghai Key Basic Research Project (grant no 09JC1408000) and the National Natural Science Foundation of China (grant no 61004088) for support. S.V.B. acknowledges the partial support of this research through the B. W. Gamson Computational Science Center at Yeshiva College. S.H. thanks the European EPIWORK project, Deutsche Forschungsgemeinschaft (DFG) and the Israel Science Foundation for financial support.

Additional information

The authors declare no competing financial interests. Reprints and permissions information is available online at http://www.nature.com/reprints. Correspondence and requests for materials should be addressed to H.E.S.