

# Chapter 10

## Understanding Simulation Results

Andrew Evans, Alison Heppenstall, and Mark Birkin

**Abstract** Simulation modelling is concerned with the abstract representation of entities within systems and their interrelationships; understanding and visualising these results is often a significant challenge for the researcher. Within this chapter we examine particular issues such as finding “important” patterns and interpreting what they mean in terms of causality. We also discuss some of the problems with using model results to enhance our understanding of the underlying social systems which they represent, and we will assert that this is in large degree a problem of isolating causal mechanisms within the model architecture. In particular, we highlight the issues of equifinality and identifiability—that the same behaviour may be induced within a simulation from a variety of different model representations or parameter sets—and present recommendations for dealing with this problem. The chapter ends with a discussion of avenues of future research.

### Why Read This Chapter?

To help you understand the results that a simulation model produces, by suggesting some ways to analyse and visualise them. The chapter concentrates on the internal dynamics of the model rather than its relationship to the outside world.

## 10.1 Introduction

Simulation models may be constructed for a variety of purposes. Classically these purposes tend to centre on either the capture of a set of knowledge or making predictions. Knowledge capture has its own set of issues that are concerned with structuring and verifying knowledge in the presence of contradiction and uncertainty. The problems of prediction, closely associated with calibration and validation, centre around comparisons with real data, for which the methods covered in Chap. 9 (David et al. 2017) are appropriate. In this chapter, however, we look at

---

A. Evans (✉) • A. Heppenstall • M. Birkin  
School of Geography, University of Leeds, Leeds, UK  
e-mail: [a.j.evans@leeds.ac.uk](mailto:a.j.evans@leeds.ac.uk)

what our models tell us through their internal workings and logic, how we might understand/interpret simulation results as results about an attempted simulation of the real world, rather than as results we expect to compare directly with the world. Here then, we tackle the third purpose of modelling: the *exploration* of abstracted systems through simulation. In a sense, this is a purpose predicated only on the limitations of the human mind. By common definition, simulation modelling is concerned with abstract representations of entities within systems and their interrelationships and with the exploration of the ramifications of these abstracted behaviours at different temporal and geographical scales. In a world in which we had larger brains, models would not be required to reveal anything—we would instantly see the ramifications of abstracted behaviours in our heads. To a degree, therefore, models may be seen as replacing the hard joined-up thinking that is required to make statements about the way the world works. This chapter looks at what this simplifying process tells us about the systems we are trying to replicate.

In part, the complications of simulation modelling are a product of the dimensionality of the systems with which we are dealing. Let us imagine that we are tackling a system of some spatio-temporal complexity, for example, the prices in a retail market selling items A, B and C. Neighbouring retailers adjust their prices based on local competition, but the price of raw materials keeps the price surface out of equilibrium. In addition, customers will only buy one of the products at a time, creating a link between the prices of the three items. Here, then, we have three interdependent variables, each of which varies spatio-temporally, with strong auto- and cross-correlations in both time and space. What kinds of techniques can be used to tease apart such complex systems? In Sect. 10.2 of this chapter, we will discuss some of the available methodologies broken down by the dimensionality of the system in question and the demands of the analysis. Since the range of such techniques is extremely sizable, we shall detail a few traditional techniques that we believe might be helpful in simplifying model data that shows the traits of complexity and some of the newer techniques of promise.

Until recently, most social science models represented social systems using mathematical aggregations. We have over 2500 years' worth of techniques to call upon that are founded on the notion that we need to simplify systems as rapidly as we can to the point at which the abstractions can be manipulated within a single human head. As is clear, not least from earlier contributions in this volume, it is becoming increasingly accepted that social scientists might reveal more about systems by representing them in a less aggregate manner. More specifically, the main difference between mathematics and the new modelling paradigm is that we now aspire to work at a scale at which the components under consideration can be represented as having their own discrete histories; mathematics actually works in a very similar fashion to modern models, but at all the other scales. Naturally there are knock-ons from this in terms of the more explicit representation of objects, states and events, but these issues are less important than the additional simulation and analytical power that having a history for each component of a system gives us. Of course, such a "history" may just be the discrete position of an object at a single historical moment, and plainly at this level of complication, the boundary between

such models and, for example, Markov models, is somewhat diffuse; however, as the history of components becomes more involved, so the power of modern modelling paradigms comes to the fore. What is lacking, however, are the techniques that are predicated on these new architectures. Whilst models which are specified at the level of individual entities or “agents” may also be analysed using conventional mathematical techniques, in Sect. 10.3 of the chapter, we will discuss some more novel approaches which are moving the direction of understanding the outputs of these new, unaggregated, models on their own terms.

One of the reasons that simulation models are such a powerful methodology for understanding complex systems is their ability to display aggregate behaviour which goes beyond the simple extrapolation of the behaviour of the individual component parts. In mathematical analysis, such as dynamical systems theory, this behaviour tends to be linked to notions of equilibrium, oscillation and catastrophe or bifurcation. Individual- and agent-based modelling approaches have veered more strongly towards the notion of emergence, which can be defined as “an unforeseen occurrence; a state of things unexpectedly arising” (OED 2010). The concept of emergence is essentially a sign of our ignorance of the causal pathways within a system. Nevertheless, emergence is our clearest hope for developing an understanding of systems using models. We hope that emergence will give us a perceptual shortcut to the most significant elements of a system’s behaviour. When it comes to applications, however, emergence is a rather double-edged blade: emergence happily allows us to see the consequence of behaviours without us having to follow the logic ourselves; however it is problematic in relying upon us to filter out which of the ramifications are important to us. As emergence is essentially a sign of incomplete understanding, and therefore weakly relative, there is no objective definition of what is “important”. one day classification of the kinds of patterns that relate to different types of causal history, but there is no objective manner of recognising a pattern as “important” as such. These two problems, finding “important” patterns (in the absence of any objective way of defining “important”) and then interpreting what they mean in terms of causality, are the issues standing between the researcher and perfect knowledge of a modelled system. In the fourth section of this chapter, we will discuss some of the problems with using model results to enhance our understanding of the underlying social systems which they represent, and we will assert that this is in large degree a problem of isolating causal mechanisms within the model architecture. In particular, we highlight the issues of equifinality and identifiability—that the same behaviour may be induced within a simulation from a variety of different model representations or parameter sets—and present recommendations for dealing with this problem. Since recognising emergence and combating the problems of identifiability and equifinality are amongst the most urgent challenges to effective modelling of complex systems, this leads naturally to a discussion of future directions in the final section of the chapter.

## 10.2 Aggregate Patterns and Conventional Representations of Model Dynamics

Whether a model is based on deductive premises or inferred behaviours, any new understanding of a given modelled system tends to be developed inductively. Modellers examine model outputs, simplify them and then try to work out the cause utilising a combination of hypothesis dismissal, refinement and experimentation. For example, a modeller of a crowd of people might take all the responses of each person over time and generate a single simple mean statistic for each person; these might then be correlated against other model variables. If the correlation represents a real causal connection, then varying the variables should vary the statistic. Proving such causal relationships is not something we often have the ability to do in the real world. During such an analysis, the simplification process is key: it is this that reveals the patterns in our data. The questions are: how do we decide what needs simplifying and, indeed, how simple to make it?

We can classify model results by the dimensionality of the outputs. A general classification for social systems would be:

- Single statistical aggregations (1D)
- Time series of variables (2D)
- The spatial distributions of invariants (2D) or variables (3D)
- Spatio-temporal locations of invariants (3D) or variables (4D)
- Other behaviours in multidimensional variable space ( $nD$ )

For simplicity, this assumes that geographical spaces are essentially two-dimensional (while recognising that physical space might also be represented along linear features such as a high street, across networks or within a three-dimensional topographical space for landforms or buildings). It should also be plain that in the time dimension, models do not necessarily produce just a stream of data, but that the data can have complex patterning. By their very nature, individual-level models, predicated as they are on a life cycle, will never stabilise in the way a mathematical model might (Uchmanski and Grimm 1996); instead models may run away or oscillate, either periodically or chaotically.

Methods for aiding pattern recognition in data break down, again, by the dimensionality of the data, but also by the dimensionality of their outputs. It is quite possible to generate a one-number statistic for a 4D spatio-temporal distribution. In some cases, the reduction of dimensionality is the explicit purpose of the technique, and the aim is that patterns in one set of dimensions should be represented as closely as possible in a smaller set of dimensions so they are easier to understand. Table 10.1 below presents a suite of techniques that cross this range (this is not meant to be an exhaustive list; after all, pattern recognition is a research discipline of its own with a whole body of literature including several dedicated journals).

To begin with, let us consider some examples which produce outputs in a single dimension. In other words, techniques for generating global and regional statistics describing the distribution of variables across space, either a physical space or a

**Table 10.1** Pattern recognition techniques for different input and output data dimensions

	1D output	2D output	3D output	4D output	ND
1D input					
2D input	Exploratory statistics	Cluster locating Fourier/wavelet transforms			
3D input	Entropy statistics	Phase diagrams Fourier/wavelet transforms			
4D input	Diffusion statistics	Time slices	Recurrence plots		
<i>n</i> D	Network statistics	Eigenvector analysis	Sammon mapping	Animations	Heuristic techniques

variable space. Such statistics generally tend to be single time slice, but can be generated for multiple time slices to gauge overall changes in the system dynamics.

Plainly, standard aggregating statistics used to compare two distributions, such as the variable variance, will lose much of interest, both spatially and temporally. If we wish to capture the distribution of invariants, basic statistics like nearest-neighbour (Clark and Evans 1954) or the more complex patch shape, fragmentation and connectivity indices of modern ecology (for a review and software, see McGarigal 2002) provide a good starting point. Networks can be described using a wide variety of statistics covering everything from shortest paths across a network to the quantity of connections at nodes (for a review of the various statistics and techniques associated with networks, see Boccaletti et al. 2006; Evans 2010). However, we normally wish to assess the distribution of a variable across a surface—for example, a price surface or a surface of predicted retail profitability. One good set of global measures for such distributions are entropy statistics. Suppose we have a situation in which a model is trying to predict the number of individuals that buy product A in one of four regions. The model is driven by a parameter, beta. In two simulations we get the following results: *simulation one* (low beta), 480, 550, 520 and 450 and *simulation two*, (high beta) 300, 700, 500 and 400. Intuitively the first simulation has less dispersal or variability than the second simulation. An appropriate way to measure this variability would be through the use of entropy statistics. The concept of entropy originates in thermodynamics, where gases in a high-entropy state contain dispersed molecules. Thus high entropy equates to high levels of variability. Entropy statistics are closely related to information statistics where a -entropy state corresponds to a high information state. In the example above, *simulation two* is said to contain more “information” than *simulation one*, because if we approximate the outcome using no information, we would have a flat average—500, 500, 500 and 500—and this is closer to *simulation one* than *simulation two*. Examples of entropy and information statistics include Kolmogorov-Chaitin, mutual information statistics and the Shannon information statistic. Most applications in the literature use customised code for the computation of entropy statistics, although the computation of a limited range of generalised entropy indices is possible within

Stata.<sup>1</sup> Entropy statistics can also be used to describe flows across networks. In this sense they provide a valuable addition to network statistics: most network statistics concentrate on structure rather than the variable values across them. Unless they are looking specifically at the formation of networks over time, or the relationship between some other variable and network structure, modellers are relatively bereft of techniques to look at variation *on* a network.

In the case where variability is caused and constrained by neighbourhood effects, we would expect the variation to be smoother across a region. We generally expect objects in space under neighbourhood effects to obey Tobler's first law of geography (Tobler 1970) that everything is related, but closer things are related more. This leads to spatial auto- or cross-correlation, in which the values of variables at a point reflect those of their neighbours. Statistics for quantifying such spatial auto- or cross-correlation at the global level, or for smaller regions, such as Moran's I and Geary's C, are well established in the geography literature (e.g. Haining 1990); a useful summary can be found in Getis (2007).

Such global statistics can be improved on by giving some notion of the direction of change of the auto- or cross-correlation. Classically this is achieved through semi-variograms, which map out the intensity of correlation in each direction traversed across a surface (for details, see Isaaks and Srivastava 1990). In the case where it is believed that local relationships hold between variables, local linear correlations can be determined, for example, using geographically weighted regression (GWR; for details, see Fotheringham et al. 2002). GWR is a technique which allows the mapping of  $R^2$ s calculated within moving windows across a multivariate surface and, indeed, mapping of the regression parameter weights. For example, it would be possible in our retail results to produce a map of the varying relationship between the amount of A purchased by customers and the population density, if we believed these were related. GWR would not just allow a global relationship to be determined, but also how this relationship changed across a country. One important but somewhat overlooked capability of GWR is its ability to assess how the strength of correlations varies with scale by varying the window size. This can be used to calculate the key scales at which there is sufficient overlap between the geography of variables to generate strong relationships (though some care is needed in interpreting such correlations, as correlation strength generally increases with scale: Robinson 1950; Gehlke and Biehl 1934). Plainly, identifying the key scale at which the correlations between variables improve gives us some ability to recognise key distance scales at which causality plays out. In our example, we may be able to see that the scale at which there is a strong relationship between sales of A and the local population density increases as the population density decreases, suggesting rural consumers have to travel further and a concomitant non-linearity in the model components directing competition.

---

<sup>1</sup>Confusingly, "generalised entropy" methods are also widely used in econometrics for the estimation of missing data. Routines which provide this capability, e.g. in SAS, are not helpful in the description of simulation model outputs!

If, on the other hand, we believe the relationships do not vary smoothly across a modelled surface, we instead need to find unusual clusters of activity. The ability to represent spatial clustering is of fundamental importance, for example, within Schelling's well-known model of segregation in the housing market (Schelling 1969). However clustering is often not so easy to demonstrate within both real data and complex simulation outputs. The most recent techniques use, for example, wavelets to represent the regional surfaces, and these can then be interpreted for cluster-like properties. However, for socio-economic work amongst the best software for cluster detection is the geographical analysis machine (GAM), which not only assesses clustering across multiple scales but also allows assessment of clustering in the face of variations in the density of the population at risk. For example, it could tell us where transport network nodes were causing an increase in sales of A, by removing regions with high sales caused by high population density (the population "at risk" of buying A). Clusters can be mapped and their significance assessed (Openshaw et al. 1988).

Often, simulations will be concerned with variations in the behaviour of systems, or their constituent agents, over time. In common with physical systems, social and economic systems are often characterised by periodic behaviour, in which similar states recur, although typically this recurrence is much less regular than in many physical systems. For example, economic markets appear to be characterised by irregular cycles of prosperity and depression. Teasing apart a model can provide nonintuitive insights into such cycles. For example, Heppenstall et al. (2006) considered a regional network of petrol stations and showed within an agent simulation how asymmetric cyclical variations in pricing (fast rises and slow falls), previously thought to be entirely due to a desire across the industry to maintain artificially high profits, could in fact be generated from more competitive profit maximisation in combination with local monitoring of network activity. While it is, of course, not certain these simpler processes cause the pattern in real life, the model exploration does give researchers a new explanation for the cycles and one that can be investigated in real petrol stations.

In trying to detect periodic behaviour, wavelets are rapidly growing in popularity (Graps 2004). In general, one would assume that the state of the simulation can be represented as a single variable which varies over time (let's say the average price of A). A wavelet analysis of either observed or model data would decompose this trend into chunks of time at varying intervals, and in each interval the technique identifies both a long-term trend and a short-term fluctuation. Wavelets are therefore particularly suitable for identifying cycles within data. They are also useful as filters for the removal of noise from data and so may be particularly helpful in trying to compare the results from a stylised simulation model with observed data which would typically be messy, incomplete or subject to random bias. It has been argued that such decompositions are fundamentally helpful in establishing a basis for forecasting (Ramsey 2002).

Wavelets are equally applicable in both two and three dimensions. For example, they may be useful in determining the diffusion of waves across a two-dimensional space and over time and can be used to analyse, for example, the relationship

between wave amplitude and propagation distance. Viboud et al. (2006) provide a particularly nice example of such a use, looking at the strength of the propagation of influenza epidemics as influenced by city size and average human travel distances in the USA. Other more traditional statistics, such as the Rayleigh statistic (Fisher et al. 1987; Korie et al. 1998), can also be used to assess the significance of diffusion from point sources.

In addition to global and regional aggregate statistics of single variables or cross-correlations, it may be that there is simply too great a dimensionality to recognise patterns in outputs and relate them to model inputs. At this point it is necessary to engage in multidimensional scaling. If an individual has more than four characteristics, then multidimensional scaling methods can be used to represent the individuals in two or three dimensions. In essence, the problem is to represent the relation between individuals such that those which are most similar in  $n$ -dimensions still appear to be closest in a lower-dimensional space which can be visualised more easily. The most popular technique is Sammon mapping. This method relies on the ability to optimise an error function which relates original values in high-dimensional space to the transformed values. This can be achieved using standard optimisation methods within packages such as MATLAB or using a number of bespoke R packages. Multidimensional scaling can be useful in visualising the relative position of different individuals within a search space, for exploring variations in a multi-criteria objective function within a parameter space or for comparing individual search paths within different simulations (Pohlheim 2006).

Eigenvector methods are another form of multidimensional scaling. Any multidimensional representation of data in  $n$ -dimensional space can be transformed into an equivalent space governed by  $n$  orthogonal eigenvectors. The main significance of this observation is that the principal eigenvector constitutes the most efficient way to represent a multidimensional space within a single value. For example, Moon, Schneider and Carley (Moon et al. 2006) use the concept of “eigenvector centrality” within a social network to compute a univariate measure of relative position based on a number of constituent factors.

Eigenvector analyses, however, can be nonintuitive to those not used to them. Somewhat simpler presentations of multidimensional data can be made using clustering techniques. These collapse multidimensional data so that individual cases are members of a single group or cluster, classified on the basis of a similarity metric. The method may therefore be appropriate if the modeller wishes to understand the distribution of an output variable in relation to the combination of several input variables. Cluster analysis is easy to implement in all the major statistics packages (R, SAS, SPSS). The technique is likely to be most useful in empirical applications with a relatively large number of agent characteristics (i.e. six or more) rather than in idealised simulations with simple agent rules. One advantage of this technique over others is that it is possible to represent statistical variation within the cluster space, for example, by displaying the interquartile variation in the attribute variable within clusters.



### 10.3 Individual Patterns, Novel Approaches and Visualisation

Plainly aggregate statistics like those above are a useful way of simplifying individual-level data, both in terms of complexity and dimensionality. However, they are the result of over 2500 years of mathematical development in a research environment unsuited to the mass of detail associated with individual-level data. Now, computers place us in the position of being able to cope with populations of individual-level data at a much smaller scale. We still tend to place our own understanding at the end of an analytical trail, constraining the trail to pass through some kind of simplification and higher level of aggregation for the purposes of model analysis. Despite this, it is increasingly true that individual-level data is dealt with at the individual level for the body of the analysis, and this is especially true in the case of individual-level modelling, in which experimentation is almost always enacted at the individual level. Whether it is really necessary to simplify for human understanding at the end of an analysis is not especially clear. It may well be that better techniques might be developed to do this than those built on an assumption of the necessity of aggregation.

At the individual level, we are interested in recognising patterns in space and time, seeing how patterns at different scales affect each other, and then using this to say something about the behaviour of the system/individuals. Patterns are often indicators of the attractors to which individuals are drawn in any given system and present a shortcut to understanding the mass of system interactions. However, it is almost as problematic to go through this process to understand a model as it is, for example, to derive individual-level behaviours from real large-size spatio-temporal datasets of socio-economic attributes. The one advantage we have in understanding a model is that we do have some grip on the foundation rules at the individual scale. Nonetheless, understanding a rule and determining how it plays out in a system of multiple interactions are very different things. Table 10.2 outlines some of the problems.

Despite the above, our chief tool for individual-level understanding without aggregation is, and always has been, the human ability to recognise patterns in

**Table 10.2** Issues related to understanding a model at different levels of complexity

Complexity	Issues
Spatial	What is the impact of space (with whom do individuals initiate transactions and to what degree)?
Temporal	How does the system evolve?
Individuals	How do we recognise which individual behaviours are playing out in the morass of interactions?
Relationships	How do we recognise and track relationships?
Scale	How can we reveal the manner in which individual actions affect the large-scale system and vice versa?

masses of data. Visualisation, for all its subjectivity and faults, remains a key element of the research process. The standard process is to present one or more attributes of the individuals in a map in physical or variable space. Such spaces can then be evolved in movies or sliced in either time or space (Table 10.3 shows some examples). In general, we cannot test the significance of a pattern without first recognising it exists, and to that extent significance testing is tainted by the requirement that it tests our competency in recognising the correct pattern as much as that the proposed pattern represents a real feature of the distribution of our data. Visualisation is also a vital tool in communicating results within the scientific community and to the wider public. The former is not just important for the transmission of knowledge, but because it allows others to validate the work. Indeed, the encapsulation of good visualisation techniques within a model framework allows others to gain deeper understanding of one's model, and to experiment at the limits of the model—what Grimm (2002) calls “visual debugging”. Good model design starts like the design of any good application, with an outline of what can be done to make it easy to use, trustworthy and simple to understand. Traditionally, user interface design and visualisation have been low on the academic agenda, to the considerable detriment of both the science and the engagement of taxpayers. Fortunately, in the years since the turn of the millennium, there has been an increasing realisation that good design engages the public and that there is a good deal of social science research that can be built on that engagement. Orford et al. (1999) identify computer graphics, multimedia, the World Wide Web and virtual reality as four visualisation technologies that have recently seen a considerable evolution within the social sciences. There is an ever-increasing array of visualisation techniques at our disposal: Table 10.3 presents a classification scheme of commonly used and more novel visualisation methods based on the dimensionality and type of data that is being explored.

Another classification scheme of these techniques that is potentially very useful comes from Andrienko et al. (2003). This classification categorises techniques based on their applicability to different types of data:

- “Universal” techniques that can be applied whatever the data, e.g. querying and animation
- Techniques revealing existential change, e.g. time labels, colouring by age, event lists and space-time cubes
- Techniques about moving objects, e.g. trajectories, space-time cubes and snapshots in time
- Techniques centred on thematic/numeric change, e.g. change maps, time series and aggregations of attribute values

For information on other visualisation schemes, see Cleveland (1983), Hinneburg et al. (1999) and Gahegan (2001).

In each case, the techniques aim to exploit the ease with which humans recognise patterns (Muller & Schumann Müller and Schumann 2003). Pattern recognition is, at its heart, a human attribute, and one which we utilise to understand models, no matter how we process the data. The fact that most model understanding is founded

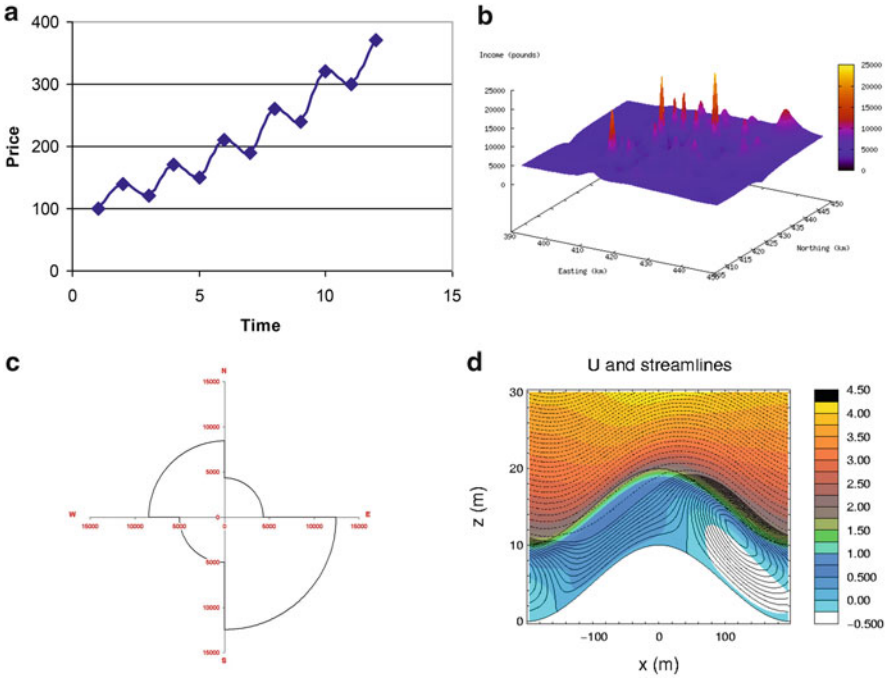
**Table 10.3** Classification of visualisation methods according to dimensionality and type of data

	Method	Pro	Con
Spatial 1D/2D	Map: overlay; animated trajectory representation (e.g. arrows); snapshots	View of whole trajectory of an object	Cannot analyse trajectory of movement. If several objects cross paths, cannot tell whether objects met at crossing point or visited points at different times
	Spatial distribution, e.g. choropleth maps	Gives a snapshot of an area.	Cannot see how a system evolves through time. Aggregate view of area. Only represents one variable; hard to distinguish relationships
Temporal 1D	Time-series graphs/linear and cyclical graphs	Show how the system (or parameters) change over time	No spatial element. Hard to correlate relationships between multivariate variables
	Rank clocks (e.g. Batty 2006)	Good for visualising change over time in ranked order of any set of objects	No spatial element
	Rose diagrams	Good for representation of circular data, e.g. wind speed and direction	No spatial element
	Phase diagram	Excellent for examining system behaviour over time for one or two variables	No spatial element. Gets confusing quickly with more than two variables
Spatio-temporal 3D/4D	Map animation (e.g. Patel and Hudson-Smith 2012)	Can see system evolving spatially and temporally	Hard to quantify or see impacts of individual behaviour, i.e. isolated effects
	Space-time cube (Andrienko et al. 2003)	Can contain space-time paths for individuals	Potentially difficult to interpret
	Recurrence plot	Reveals hidden structures over time and in space	Computationally intensive. Methods difficult to apply. Have to generate multiple snapshots and run as an animation
	Vector plotting/contour slicing (Ross and Vosper 2003)	Ability to visualise 2D or 3D data and multiple dimensional dataset	Hard to quantify individual effects

on a human recognition of a “significant” pattern is somewhat unfortunate, as we will bring our own biases to the process. At worst we only pay attention to those patterns that confirm our current prejudices: what Wyszomirski et al. (1999) call the *WYWIWYG*—What You Want is What You Get—fallacy. At best, we will only recognise those patterns that match the wiring of the human visual system and our cultural experiences. The existence of visualisation techniques generally points up the fact that humans are better at perceiving some patterns than others, and in some media than others—it is easier to see an event as a movie and not a binary representation of the movie file displayed as text. However, in addition to standard physiological and psychological restrictions on pattern recognition consistent to all people, it is also increasing apparent there are cultural differences in perceptions. Whether there is some underlying biological norm for the perception of time and space is still moot (Nisbett and Masuda 2003; Boroditsky 2001), but it is clear that some elements of pattern recognition vary by either culture or genetics (Nisbett and Masuda 2003; Chua et al. 2005). Even when one looks at the representation of patterns and elements like colour, there are clear arguments for a social influence on the interpretation of even very basic stimuli into perceptions (Roberson et al. 2004). Indeed, while there is a clear and early ability of humans to perceive moving objects in a scene as associated in a pattern (e.g. Baird et al. 2002), there are cultural traits associated with the age at which even relatively universal patterns are appreciated (Clement et al. 1970). The more we can objectify the process, therefore, the less our biases will impinge on our understanding. In many respects it is easier to remove human agents from data comparison and knowledge development than pattern hunting, as patterns are not something machines deal with easily. The unsupervised recognition of even static patterns repeated in different contexts is far from computationally solved (Bouvrie and Sinha 2007), though significant advances have been made in recent years (Druzhkov and Kustikova 2016). Most pattern-hunting algorithms try to replicate the process found in humans, and in that sense one suspects we would do better to skip the pattern hunting and concentrate on data consistency and the comparison of full datasets directly. At best we might say that an automated “pattern” hunter that wasn’t trying to reproduce the human ability would instead seek to identify attractors within the data.

Figure 10.1 presents several visualisation methods that are commonly found in the literature, ranging from 1D time-series representation (a) to contour plots (d) that could be potentially used for 4D representation.

Visualisations are plainly extremely useful. Here we’ll look at a couple of techniques that are of use in deciphering individual-level data: phase maps and recurrence plots. Both techniques focus on the representation of individual-level states and the relationships between stated individuals.

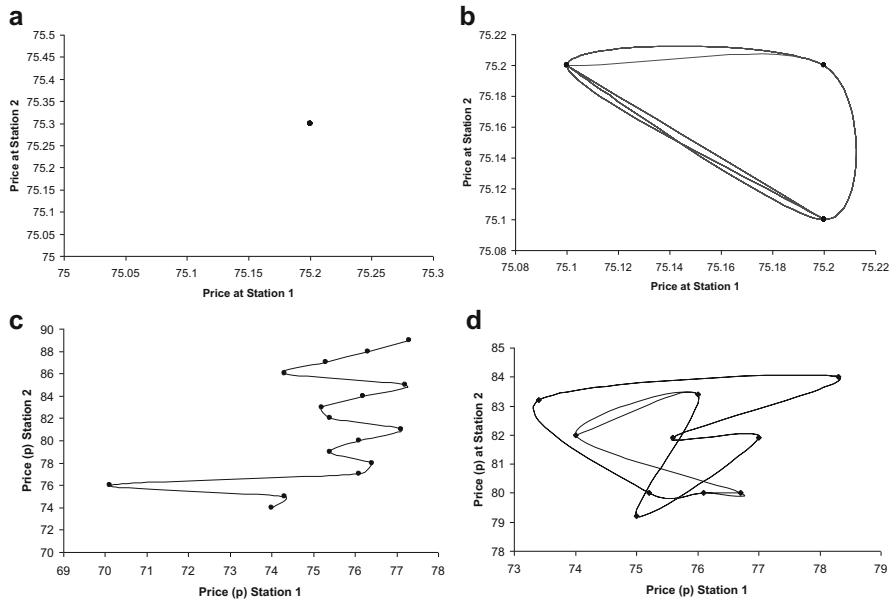


**Fig. 10.1** Examples of different visualisation methods. (a) 1D Time-series graph (idealised data). (b) 3D interpolated map (idealised data). (c) Rose diagram. (d) Contour plot

### 10.3.1 Phase Maps

Phase-space maps are commonly used by physicists to study the behaviour of physical systems. In any graphical representation, a phase-space map represents an abstract view of the behaviour of one or more of the system components. These can be particularly useful to us as we can plot the behaviour of our system over time. This allows us to understand how the system is evolving and whether it is chaotic, random, cyclical or stable (Fig. 10.2).

Each of the graphs produced in Fig. 10.2 is a representation of the coincident developments in two real neighbouring city centre petrol stations in Leeds (UK) over a 30-day period (sampled every other day). Figure 10.2a represents a stable system. Here, neither of the stations is changing in price and, thus, a fixed point is produced. However, this behaviour could easily change if one or both of the stations alter its price. This behaviour is seen in Fig. 10.2b. Both stations are changing their prices each day (from 75.1p to 75.2p to 75.1p); this creates a looping effect; the stations are cycling through a pattern of behaviour before returning to their starting point. Note that the graph appears to reveal a causative link between the two stations as they are never simultaneously low. Figure 10.2c, d shows a more varied pattern of behaviour between the stations. In Fig. 10.2c, one point is rising in price, whilst

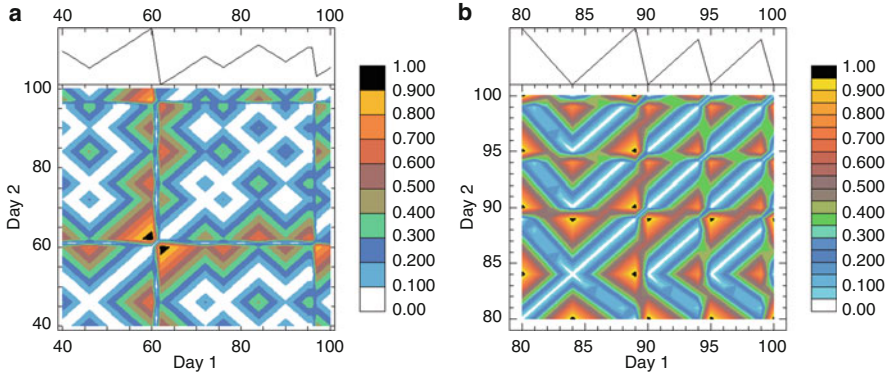


**Fig. 10.2** Examples of different types of behaviour found in urban petrol stations (Leeds). (a) Stable. (b) Looping. (c) Two types of behaviour. (d) Chaotic

the other is oscillating. In Fig. 10.2(d), there is no apparent pattern in the displayed behaviour. Simply knowing about these relationships is valuable information and allows us a greater understanding of this system, its behaviour and its structure. For example, it may be that the only difference between the graphs is one of distance between the stations, but we would never see this unless the graphs allowed us to compare at a detailed level the behaviours of stations that potentially influence each other.

### 10.3.2 Recurrence Plots

Recurrence plots (RPs) are a relatively new technique for the analysis of time-series data that allows both visualisation and quantification of structures hidden within data or exploration of the trajectory of a dynamical system in phase space (Eckmann et al. 1987). They are particularly useful for graphically detecting hidden patterns and structural changes in data as well as examining similarities in patterns across a time-series dataset (where there are multiple readings at one point). RPs can be also used to study the nonstationarity of a time series as well as to indicate its degree of aperiodicity (Casdagli 1997; Kantz and Schreiber 1997). These features make RPs a very valuable technique for characterising complex dynamics in the time domain



**Fig. 10.3** Example of Recurrence Plots. (a) RP of the change in price at a retail outlet over 100 days. (b) illustrates how oscillations in the change in the price data are represented in the RP

(Vasconcelos et al. 2006), a factor reflected in the variety of applications that RPs can now be found in ranging from climate variation (Marwan and Kruths 2002) and music (Foote and Cooper 2001) to heart rate variability (Marwan et al. 2002).

Essentially a RP is constructed via a matrix where values at a pair of time steps are compared against each other. If the system at the two snapshots is completely different, the result is 1.0 (black), while completely similar periods are attributed the value 0.0 (represented as white). Through this, a picture of the structure of the data is built up. Figure 10.3a shows the RP of the change in price at a retail outlet over 100 days. Above the RP is a time-series graph diagrammatically representing the change in price. Changes in price, either increases, decreases or oscillations, can be clearly seen in the RP. Figure 10.3b illustrates how oscillations in the change in the price data are represented in the RP.

Early work on this area has shown that there is considerable potential in the development and adaptation of this technique. Current research is focused on the development of cross-reference RPs (consideration of the phase-space trajectories of two different systems in the same phase space) and spatial recurrence plots.

## 10.4 Explanation, Understanding and Causality

Once patterns are recognised, “understanding” our models involves finding explanations highlighting the mechanisms within the models which give rise to these patterns. The process of explanation may be driven with reference to current theory or developing new theory. This is usually achieved through:

1. Correlating patterns visually or statistically with other parts of the model, such as different geographical locations, or with simulations with different starting values.
2. Experimentally adjusting the model inputs to see what happens to the outputs.

### 3. Tracking the causal processes through the model.

It may seem obvious, and yet it is worth pointing out, that model outputs can only causally relate to model inputs, not additional data in the real world. Plainly insights into the *system* can come from comparison with external data that is correlated or miscorrelated with model outputs, but this is not the same as understanding your model and the way it currently represents the system. One would imagine that this means that understanding of a model cannot be facilitated by comparing it with other, external, data, and yet it can often be worth:

4. Comparing model results with real-world data, because the relationships between real data and both model inputs and model outputs may be clearer than the relationships between these two things within the model.

Let's imagine, for example, a model that predicts the location of burglaries across a day in a city region where police corruption is rife. The model inputs are known offenders' homes, potential target locations and attractiveness, the position of the owners' homes and the police, who prefer to serve the wealthy. We may be able to recognise a pattern of burglaries that moves, over the course of the day, from the suburbs to the city centre. Although we have built into our model the fact that police respond faster to richer people, we may find, using (1), that our model doesn't show less burglaries in rich areas, because the rich areas are so spatially distributed that the police response times are stretched between them. We can then alter the weighting of the bias away from the wealthy (2) to see if it actually reduces the burglary rate in the rich areas by placing police nearer these neighbourhoods as an ancillary effect of responding to poor people more. We may be able to fully understand this aspect of the model and how it arises (3), but still have a higher than expected burglary rate in wealthy areas. Finally, it may turn out (4) that there is a strong relationship between these burglaries and real data on petrol sales, for no other reason than both are high at transition times in this social system, when the police would be most stretched between regions—suggesting in turn that the change in police locations over time is as important as their positions at any one time.

Let us look at each of these methodologies for developing understanding in turn.

**Correlation** Most social scientists will be familiar with linear regression as a means for describing data or testing for a relationship between two variables; there is a long scientific tradition of correlating data between models and external variables, and this tradition is equally applicable to intra-model comparisons. Correlating datasets is one of the areas where automation can be applied. As an exploratory tool, regression modelling has its attractions, not least its simplicity in both concept and execution. Simple regressions can be achieved in desktop applications like Microsoft Excel, as well as all the major statistical packages (R, SAS, SPSS, etc.). Standard methodologies are well known for cross-correlation of both continuous normal data and time series. However even for simple analyses with a single input and single output variable, linear regression is not always an appropriate technique. For example, logistic regression models will be more appropriate for binary response data, Poisson models will be superior when values in the dependent



tend to be highly clustered, while binomial models may be the most effective when observations are highly dispersed around the mean. An interesting example is Fleming and Sorenson (2001) in which binomial estimates of technological innovation are compared to the complexity of the invention measured by both the number of components and the interdependence between those components. In behavioural space, methodologies such as association rule making (e.g. Hipp et al. 2002) allow the Bayesian association of behavioural attributes. It is worth noting that where models involve a distribution in physical space, this can introduce problems, in particular where the model includes neighbourhood-based behaviours and therefore the potential to develop spatial auto and cross-correlations. These alter the sampling strategies necessary to prove relationships—a full review of the issues and methodologies to deal with them can be found in Wagner and Fortin (2005).

**Experimentation** In terms of experimentation, we can make the rather artificial distinction between sensitivity testing and “what if?” analyses—the distinction is more one of intent than anything. In sensitivity analysis one perturbs model inputs slightly to determine the stability of the outputs, under the presumption that models should match the real world in being insensitive to minor changes (a presumption not always well founded). In “what if?” analyses, one alters the model inputs to see what would happen under different scenarios. In addition to looking at the output values at a particular time slice, the stability or otherwise of the model, and the conditions under which this varies, also gives information about the system (Grimm 1999).

**Tracking Causality** Since individual-based models are a relatively recent development, there is far less literature dealing with the tracking of causality through models. It helps a little that the causality we deal with in models, which is essentially a mechanistic one, is far more concrete than the causality perceived by humans, which is largely a matter of the repeated coincidence of events. Nevertheless, backtracking through a model to mark a causality path is extremely hard, primarily for two reasons. The first is what we might call the “find the lady problem”—that the sheer number of interactions involved in social processes tends to be so large we don’t have the facilities to do the tracking. The second issue, which we might call the “drop in the ocean problem”, is more fundamental as it relates to a flaw in the mathematical representation of objects, that is, that numbers represent aggregated quantities, not individuals. When transacted objects in a system are represented with numbers greater than one, it is instantly impossible to reliably determine the path taken by a specific object through that system. For objects representing concepts, either numerical (e.g. money) or nonnumerical (e.g. a meme), this isn’t a problem (one dollar is much like any other; there is only one Gangnam style to know). However, for most objects such aggregations place ambiguous nodes between what would otherwise be discrete causal pathways. Fortunately, we tend to use numbers in agent models as a methodology to cope with our ignorance (e.g. in the case of calibrated parameters) or the lack of the computing power we’d need to deal with individual objects and their transactional histories (e.g. in the case of a variable like “number of children”). As it happens, every day brings improvements to both.

In addition, the last 10 years or so has seen considerable theoretical advances in the determination of the probabilities of causation (e.g. Granger 1980; Pearl and Verma 1991; Greenland and Pearl 2006). For now, however, the tracking of causality is much easier if the models build in appropriate structures from the start. While they are in their infancy, techniques like process calculi (Worboys 2005) and Petri nets show the potential of this area.

The inability to track causality leads to the perennial problem of identifiability, that is, that a single model outcome may have more than one history of model parameters that leads to it. Identifiability is part of a larger set of issues with confirming that the model in the computer accurately reflects the system in the real world—the so-called equifinality issue. These are issues that play out strongly during model construction from real data and when validating a model against real data, and a review of techniques to examine these problems, including using model variation to determine the suitability of variables and parameters, can be found in Evans (2012). At the model stage we are interested in, however, we at least have the advantage that there is only one potential model that may have created the output—the one running. Nevertheless, the identifiability of the parameters in a running model still makes it hard to definitively say when model behaviour is reasonable. For those modelling for prediction, this is of little consequence—as long as the model gives consistently good predictions it may as well be a black box. However, if we wish to tease the model apart and look at how results have emerged, these issues become more problematic.

The mechanisms for dealing with these problems are pragmatic:

1. Examine the stability of the calibration process and/or the state of internal variables that weren't inputs or outputs across multiple runs.
2. Validate internal variables that weren't inputs or outputs used in any calibration against real data.
3. Run the model in a predictive mode with as many different datasets as possible—the more the system can replicate reality at output, the more likely it is to replicate reality internally. If necessary engage in inverse modelling: initialize parameters randomly and then adjust them over multiple runs until they match all known outputs.

Of these, by far the easiest, but the least engaged with, is checking the stability of the model in parameter space (see Evans 2012 for a review). Various AI techniques have been applied to the problem of optimising parameters to fit model output distributions to some predetermined pattern (such as a “real-world” distribution). However, the stability of these parameterizations and the paths AIs take to generate them are rarely used to examine the degree to which the model fluctuates between different states, let alone to reflect on the nature of the system. The assumption of identifiability is that the more parameterized a model, the more likely it is a set of parameter values can be derived which fit the data but don't represent the true values. However, in practice the limits on the range of parameter values within any given model allow us an alternative viewpoint: that the more parameterized rules in a model, the more the system is constrained by the potential range of the elements

in its structure and the interaction of these ranges. For example, a simple model  $a = b$  has no constraints, but  $a = b/c$ , where  $c = \text{distance between } a \text{ and } b$ , adds an additional constraint even though there are more parameters. As such rules build up in complex systems, it is possible that parameter values become highly constrained, even though, taken individually, any given element of the model seems reasonably free. This may mean that if a system is well modelled, exploration of the model's parameter space by an AI might reveal the limits of parameters within the constraints of the real complex system. For example, Heppenstall et al. (2007) use a genetic algorithm to explore the parameterisation of a petrol retail model/market and find that while some GA-derived parameters have a wide range, others consistently fall around specific values that match those derived from expert knowledge of the real system.

The same issues as hold for causality hold for data uncertainty and error. We have little in the way of techniques for coping with the propagation of either through models (see Evans 2012 for a review). It is plain that most real systems can be perturbed slightly and maintain the same outcomes, and this gives us some hope that errors at least can be suppressed; however we still remain very ignorant as to how such homeostatic forces work in real systems and how we might recognise or replicate them in our models. Data and model errors can breed patterns in our model outputs. An important component of understanding a model is understanding when this is the case. If we are to use a model to understand the dynamics of a real system and its emergent properties, then we need to be able to recognise novelty in the system. Patterns that result from errors may appear to be novel (if we are lucky), but as yet there is little in the way of toolkits to separate out such patterns from truly interesting and new patterns produced intrinsically.

Currently our best option for understanding model artefacts is model-to-model comparisons. These can be achieved by varying one of the following contexts while holding the others the same: the model code (the model, libraries and platform), the computer the model runs on or the data it runs with (including internal random number sequences). Varying the model code (for instance, from Java to C++ or from an object-orientated architecture to a procedural one) is a useful step in that it ensures the underlying theory is not erroneously dependent on its representation. Varying the computer indicates the level of errors associated with issues like rounding and number storage mechanisms, while varying the data shows the degree to which model and theory are robust to changes in the input conditions. In each case, a version of the model that can be transferred between users, translated onto other platforms and run on different data warehouses would be useful. Unfortunately, however, there is no universally recognised mechanism for representing models abstracted from programming languages. Mathematics, UML and natural languages can obviously fill this gap to a degree, but not in a manner that allows for complete automatic translation. Even the automatic translation of computer languages is far from satisfactory when there is a requirement that the results be understood by humans so errors in knowledge representation can be checked. In addition, many such translations work by producing the same binary executable. We also need standard ways of comparing the results of models, and

these are no more forthcoming. Practitioners are only really at the stage where we can start to talk about model results in the same way (see, e.g. Grimm et al. 2006). Consistency in comparison is still a long way off, in part because statistics for model outputs and validity are still evolving and in part because we still don't have much idea which statistics are best applied and when (for one example bucking this trend, see Knudsen and Fotheringham 1986).

## 10.5 Future Directions

Recognising patterns in our modelled data allows us to:

1. Compare it with reality for validation.
2. Discover new information about the emergent properties of the system.
3. Make predictions.

Of these, discovering new information about the system is undoubtedly the hardest, as it is much easier to spot patterns you are expecting. Despite the above advances, there are key areas where current techniques do not match our requirements. In particular, these include:

1. Mechanisms to determine when we do not have all the variables we need to model a system and which variables to use.
2. Mechanisms to determine which minor variables may be important in making emergent patterns through non-linearities.
3. The tracking of emergent properties through models.
4. The ability to recognise all but the most basic patterns in space over time.
5. The ability to recognise action across distant spaces over space and time.
6. The tracking of errors, error acceleration and homeostatic forces in models.

While we have components of some of these areas, what we have is but a drop in the ocean of techniques we need. In addition, the vast majority of our techniques are built on the 2500 years of mathematics that resolved to simplify systems that were collections of individuals because we lacked the ability (either processing power or memory) to cope with the individuals *as* individuals. Modern computers have given us this power for the first time, and, as of yet, the ways we describe such systems have not caught up, even if we accept that some reduction in dimensionality and detail is necessary for a human to understand our models. Indeed in the long run, it might be questioned whether the whole process of model understanding and interpretation might be divorced from humans and delegated instead to an artificially intelligent computational agency that can better cope with the complexities directly.

## Further Reading

Statistical techniques for spatial data are reviewed by McGarigal (2002) while for network statistics good starting points are Newman (2003) and Boccaletti et al. (2006), with more recent work reviewed by Evans (2010). For information on coping with auto-/cross-correlation in spatial data, see Wagner and Fortin (2005). Patel and Hudson-Smith (2012) provide an overview of the types of simulation tool (virtual worlds and virtual reality) available for visualising the outputs of spatially explicit agent-based models. Evans (2012) provides a review of techniques for analysing error and uncertainty in models, including both environmental/climate models and what they can bring to the agent-based field. He also reviews techniques for identifying the appropriate model form and parameter sets.

## References

- Andrienko, N., Andrienko, G., & Gatalaky, P. (2003). Exploratory spatio-temporal visualisation: An analytical review. *Journal of Visual Languages and Computing*, 14(6), 503–541.
- Baird, A. A., et al. (2002). Frontal lobe activation during object permanence: Data from near-infrared spectroscopy. *NeuroImage*, 16, 1120–1126.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., & Hwang, D.-U. (2006). Complex networks: Structure and dynamics. *Physics Reports*, 424(4–5), 175–308.
- Boroditsky, L. (2001). Does language shape thought? Mandarin and English speakers' conceptions of time. *Cognitive Psychology*, 43, 1–22.
- Batty, M. (2006). Rank clocks. *Nature*, 444, 592–596.
- Bouvier, J. V., & Sinha, P. (2007). Visual object concept discovery: Observations in congenitally blind children, and a computational approach. *Neurocomputing*, 70(13–15), 2218–2233.
- Casdagli, M. (1997). Recurrence plots revisited. *Physica D: Nonlinear Phenomena*, 108(1–2), 12–44.
- Chua, H. F., Boland, J. E., & Nisbett, R. E. (2005). Cultural variation in eye movements during scene perception. *Proceedings of the National Academy of Sciences of the United States of America*, 102(35), 12629–12633.
- Clark, P. J., & Evans, F. C. (1954). Distance to nearest neighbor as a measure of spatial relationships in populations. *Ecology*, 35(4), 445–453.
- Clement, D. E., Sistrunk, F., & Guenther, Z. C. (1970). Pattern perception among Brazilians as a function of pattern uncertainty and age. *Journal of Cross-Cultural Psychology*, 1(4), 305–313.
- Cleveland, W. S. (1983). *Visualising data*. New Jersey: Hobart Press.
- David, N., Fachada, N., & Rosa, A. C. (2017). Verifying and validating simulations. doi:[https://doi.org/10.1007/978-3-319-66948-9\\_9](https://doi.org/10.1007/978-3-319-66948-9_9).
- Druzhkov, P. N., & Kustikova, V. D. (2016). A survey of deep learning methods and software tools for image classification and object detection. *Pattern Recognition and Image Analysis*, 26(1), 9–15.
- Eckmann, J. P., Kamphorst, S. O., & Reulle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9), 973–977.
- Evans, A. J. (2010). Complex spatial networks in application. *Complexity*, 16(2), 11–19.
- Evans, A. J. (2012). Uncertainty and error. In A. J. Heppenstall, A. T. Crooks, L. M. See, & M. Batty (Eds.), *Agent-based models of geographical systems*. Berlin: Springer. Chapter 15.
- Fisher, N., Lewis, T., & Embleton, B. (1987). *Statistical analysis of spherical data*. Cambridge: Cambridge University Press.

- Fleming, L., & Sorenson, O. (2001). Technology as a complex adaptive system: Evidence from patent data. *Research Policy*, 30, 1019–1039.
- Foote, J., & Cooper, M. (2001). Visualising music structure and rhythm via self-similarity. In *Proceedings of the international computer music conference, ICMC'01, Havana, Cuba* (pp. 419–422). San Francisco: ICMA.
- Fotheringham, A. S., Brunson, C., & Charlton, M. (2002). *Geographically weighted regression: The analysis of spatially varying relationships*. Chichester: Wiley.
- Gahegan, M. (2001). Visual exploration in geography: Analysis with light. In H. J. Miller & J. Han (Eds.), *Geographic data mining and knowledge discovery* (pp. 260–287). London: Taylor & Francis.
- Gehlke, C. E., & Biehl, H. (1934). Certain effects of grouping upon the size of correlation coefficients in census tract material. *Journal of the American Statistical Association*, 29(Supplement), 169–170.
- Getis, A. (2007). Reflections on spatial autocorrelation. *Regional Science and Urban Economics*, 37(4), 491–496.
- Granger, C. W. J. (1980). Testing for causality: A personal viewpoint. *Journal of Economic Dynamics and Control*, 2, 329–352.
- Graps, A. (2004). *Amara's wavelet page*. <http://www.amara.com/current/wavelet.html>
- Greenland, S., & Pearl, J. (2006). *Causal diagrams* (Technical report, R-332). Los Angeles: UCLA Cognitive Systems Laboratory. [http://ftp.cs.ucla.edu/pub/stat\\_ser/r332.pdf](http://ftp.cs.ucla.edu/pub/stat_ser/r332.pdf)
- Grimm, V. (1999). Ten years of individual-based modelling in ecology: What have we learned and what could we learn in the future? *Ecological Modelling*, 115(2), 129–148.
- Grimm, V. (2002). Visual debugging: A way of analyzing, understanding, and communicating bottom-up simulation models in ecology. *Natural Resource Modelling*, 15, 23–38.
- Grimm, V., et al. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2), 115–126.
- Haining, R. (1990). *Spatial data analysis in the social and environmental sciences*. Cambridge: Cambridge University Press.
- Heppenstall, A. J., Evans, A. J., & Birkin, M. H. (2006). Using hybrid agent-based systems to model spatially-influenced retail markets. *Journal of Artificial Societies and Social Simulation*, 9(3). <http://jasss.soc.surrey.ac.uk/9/3/2.html>
- Heppenstall, A. J., Evans, A. J., & Birkin, M. H. (2007). Genetic algorithm optimisation of a multi-agent system for simulating a retail market. *Environment and Planning B: Urban Analytics and City Science*, 34(6), 1051–1070.
- Hinneburg, A., Keim, D. A., & Wawryniuk, M. (1999). HD-eye: Visual mining of high-dimensional data. *IEEE Computer Graphics and Applications*, 19(5), 22–31.
- Hipp, J., Güntzer, U., & Nakhaeizadeh, G. (2002). Data mining of association rules and the process of knowledge discovery in databases. In P. Perner (Ed.), *Advances in data mining. (Lecture Notes in Computer Science, 2394)* (pp. 207–226). Berlin: Springer.
- Isaaks, E. H., & Srivastava, R. M. (1990). *Applied geostatistics*. North Carolina: Oxford University Press USA.
- Kantz, H., & Schreiber, T. (1997). *Non-linear time series analysis*. Cambridge: Cambridge University Press.
- Knudsen, D. C., & Fotheringham, A. S. (1986). Matrix comparison, goodness-of-fit, and spatial interaction modelling. *International Regional Science Review*, 10, 127–147.
- Korie, S., et al. (1998). Analysing maps of dispersal around a single focus. *Environmental and Ecological Statistics*, 5(4), 317–344.
- Marwan, N., & Kruths, J. (2002). Nonlinear analysis of bivariate data with cross recurrence plots. *Physics Letters A*, 302(5–6), 299–307.
- Marwan, N., Wessel, N., Meyerfeldt, U., Schirdewan, A., & Kurths, J. (2002). Recurrence-plot-based measures of complexity and their application to heart-rate-variability data. *Physical Review E*, 66(2), 026702.
- McGarigal, K. (2002). Landscape pattern metrics. In A. H. El-Shaarawi & W. W. Piegorsch (Eds.), *Encyclopedia of environmentrics* (Vol. 2, pp. 1135–1142). Chichester: Wiley.

- Moon, I.-C., Schneider, M., & Carley, K. (2006). Evolution of player skill in the America's Army game. *SIMULATION*, 82(11), 703–718.
- Müller, W., & Schumann, H. S. (2003). Visualisation methods for time-dependent data: An overview. In: S. Chick, P. J. Sánchez, D. Ferrin, & D. J. Morrice (Eds.), *Proceedings of winter simulation 2003*, New Orleans, LA, 7–10 December 2003. <http://informs-sim.org/wsc03papers/090.pdf>
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45, 167–256.
- Nisbett, R. E., & Masuda, T. (2003). Culture and point of view. *Proceedings of the National Academy of Science USA*, 100(19), 11163–11170.
- OED. (2010). *Oxford English dictionary*. <http://www.oed.com/>
- Openshaw, S., Craft, A. W., Charlton, M., & Birch, J. M. (1988). Investigation of leukaemia clusters by use of a geographical analysis machine. *Lancet*, 331(8580), 272–273.
- Orford, S., Harris, R., & Dorling, D. (1999). Geography: Information visualisation in the social sciences. *Social Science Computer Review*, 17(3), 289–304.
- Patel, A., & Hudson-Smith, A. (2012). Agent tools, techniques and methods for macro and microscopic simulation. In A. J. Heppenstall, A. T. Crooks, L. M. See, & M. Batty (Eds.), *Agent-based models of geographical systems*. Berlin: Springer. chapter 18.
- Pearl, J., & Verma, T. S. (1991). A theory of inferred causation. In J. A. Allen, R. Fikes, & E. Sandewall (Eds.), *Proceedings of the 2nd international conference on principles of knowledge representation and reasoning (KR'91)*, Cambridge, MA, USA, April 22–25, 1991 (pp. 441–452). San Mateo: Morgan Kaufmann.
- Pohlheim, H. (2006). Multidimensional scaling for evolutionary algorithms: Visualisation of the path through search space and solution space using Sammon mapping. *Artificial Life*, 12(2), 203–209.
- Ramsey, J. B. (2002). Wavelets in economics and finance: Past and future. *Studies in Nonlinear Dynamics & Econometrics*, 6(3), 1–27.
- Roberson, D., Davidoff, J., Davies, I. R. L., & Shapiro, L. R. (2004). The development of color categories in two languages: A longitudinal study. *Journal of Experimental Psychology: General*, 133(4), 554–571.
- Robinson, W. S. (1950). Ecological correlations and the behaviour of individuals. *American Sociological Review*, 15, 351–357.
- Ross, A. N., & Vosper, S. B. (2003). Numerical simulations of stably stratified flow through a mountain pass. *Quarterly Journal of Royal Meteorological Society*, 129, 97–115.
- Schelling, T. C. (1969). Models of segregation. *The American Economic Review*, 59(2), 88–493.
- Tobler, W. R. (1970). A computer model simulation of urban growth in the Detroit region. *Economic Geography*, 46(2), 234–240.
- Uchmanski, J., & Grimm, V. (1996). Individual-based modelling in ecology: What makes the difference? *Trends in Ecology & Evolution*, 11, 437–441.
- Vasconcelos, D. B., Lopes, S. R., Kurths, J., & Viana, R. L. (2006). Spatial recurrence plots. *Physical Review E*, 73(5), 056207.
- Viboud, C., et al. (2006). Synchrony, waves, and spatial hierarchies in the spread of influenza. *Science*, 312(5772), 447–451.
- Wagner, H. H., & Fortin, M.-J. (2005). Spatial analysis of landscapes: Concepts and statistics. *Ecology*, 86, 1975–1987.
- Worboys, M. F. (2005). Event-oriented approaches to geographic phenomena. *International Journal of Geographic Information Science*, 19(1), 1–28.
- Wyszomirski, T., Wyszomirska, I., & Jarzyna, I. (1999). Simple mechanisms of size distribution dynamics in crowded and uncrowded virtual monocultures. *Ecological Modelling*, 115(2–3), 253–273.