Measuring emergence via nonlinear Granger causality

Anil K. Seth

Department of Informatics, University of Sussex, Brighton, BN1 9QJ
a.k.seth@sussex.ac.uk, www.anilseth.com

Abstract

The concept of emergence is central to artificial life and complexity science, yet quantitative, intuitive, and easy-to-apply measures of emergence are surprisingly lacking. Here, I introduce a just such a measure, G-emergence, which operationalizes the notion that an emergent process is both *dependent upon* and *autonomous from* its underlying causal factors. G-emergence is based on a nonlinear time series analysis adapted from 'Granger causality' and it provides a measure not only of emergence but also of apparent 'downward causation'. I illustrate the measure by application to a canonical example of emergence, an agent-based simulation of bird flocking, and I discuss its potential impact on perhaps the most challenging of all scientific problems involving emergence: consciousness.

The maturation of artificial life and complexity science over recent years has given rise to renewed interest in emergence. Although the concept of emergence has a long philosophical history (Broad, 1925; Kim, 1999), its essence is simple enough: An emergent property is somehow 'more than the sum' of its component parts. Emergent properties appear rife within complex systems of all kinds: biological, cognitive, social, and technological. Broadly speaking, artificial life and complexity science focus on explaining phenomena that seem to involve emergence, and models constructed under these auspices are often described as emergent (Bedau, 2003). It is therefore surprising and significant that quantitative and easy-to-apply *measures* of emergence are mostly lacking. This is unfortunate because the ability to measure a phenomenon is an essential step towards its effective scientific description (Chang, 2004).

In this paper I will first differentiate several notions of emergence and by doing so briefly illustrate some relevant conceptual challenges. I will then introduce 'G-emergence', a new measure which operationalizes the intuition that an emergent process is simultaneously *autonomous from* and *dependent upon* its underlying causal factors. G-emergence is easy to apply, and I illustrate it by application to a canonical example of emergence: bird flocking. I end by discussing related measures, how it can defuse the metaphysically awkward notion of 'downward causation', and how it may shed new light one of the most recalcitrant problems in science: the relation between neural mechanism and phenomenal experience.

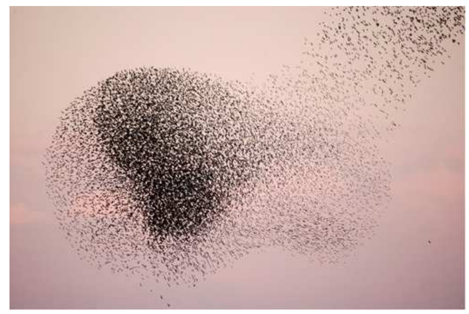

Figure 1: A flock of starlings about to roost.

Varieties of emergence

Intuitively, emergence refers either to a macro-level property that is 'more than the sum of' the micro-level parts ('property' or 'synchronic' emergence) or to the appearance of a qualitatively distinctive new phenomenon over time ('temporal' or 'diachronic' emergence). A striking example of property emergence is a flock of starlings wheeling in the sky before they roost: the flock seems to have a shape and trajectory of its own, which appears to exceed those of the individual birds (Figure 1). Temporal emergence is well illustrated by the appearance of new morphological features during embryogenesis and development. This paper focuses on measuring property emergence, though new opportunities for measuring temporal emergence are also identified.

Following Bedau (1997, 2003), both property emergence and temporal emergence can be differentiated into three categories: strong, weak, and nominal [similar decompositions can be found in (van Gulick, 2001; Bar-Yam, 2004)]. The least controversial of these is nominal emergence, which is simply the notion of a kind of property that can be possessed by macro-level objects or processes but not by their micro-level constituents. For example, a circle is nominally emergent from the set of points from which it is constructed. Because nominally emergent properties can be derived trivially I will not discuss them any further.

Most challenging and controversial is the notion of strong emergence, which involves two closely related claims. First, a macro-level property is in principle not identifiable from micro-level observations. Second, macro-level properties have irreducible causal powers. The first claim rejects mechanistic explanations altogether, apparently calling a halt to scientific advance in the absence of new fundamental principles of nature (Chalmers, 2006). The second raises the difficult notion of 'downward causation'. Downward causation is problematic firstly because it contravenes the plausible doctrine that 'the macro is the way it is in virtue of the how things are at the micro', an idea that has been expressed variously as 'causal fundamentalism' (Jackson and Pettit, 1992) or 'supervenience' (Kim, 1999). A second challenge raised by downward causation is that of resolving conflicts between micro-level and macro-level causes (Bedau, 2003). Even so, the main problem with strong emergence may lie in its scientific irrelevance (Bedau, 2003). The only recurrent example of strong emergence in the scientific literature is that of the emergence of conscious states (e.g., qualia) from neurobiological processes (Sperry, 1969; Chalmers, 2006), which may speak more to our lack of understanding of consciousness than to our grasp of deep principles of emergence. I will return to this possibility later on.

In between strong emergence and nominal emergence lies the useful notion of weak emergence (Bedau, 1997, 2003), according to which a macro-level property is derived from the interaction of micro-level components but in complicated ways such that the macro-level property has no simple micro-level explanation. In contrast to strong emergence, weakly emergent properties are in principle identifiable from micro-level components, and in contrast to nominal emergence, the micro-to-macro inferential pathways must be non-trivial. According to Bedau, weakly emergent macro-level properties are *ontologically dependent* on and reducible to micro-level causal factors, but at the same time they are *epistemologically irreducible* due to the complexity of the micro-to-macro inferential pathways.

What exactly does it mean for a macro-level property to be epistemologically irreducible? Bedau's answer is that a weakly emergent (epistemologically irreducible) property is underivable from its micro-level parts *except by simulation*. This is an all-or-none classification. Either a macro-level property can be derived by some explanatory short-cut, in which case weak emergence does not apply, or it cannot, in which case the micro-level causal factors need to be simulated explicitly in order to derive the macro-level property.

In this paper I consider a continuous version of weak emergence, in which a macro property is weakly emergent *to the extent that* it is not identifiable from micro-level observations. This variation is valuable firstly because for many systems it may not be possible to prove 'underivability except by simulation', and secondly because from the perspective of measurement, a continuous value is much more useful than a binary classification.

Measuring weak emergence

To derive a continuous measure of weak emergence, I take as a starting point the idea that a weakly emergent macro-level property is simultaneously (i) *autonomous from* and (ii) *dependent upon* its underlying causal factors (Bedau, 1997). To operationalize this notion statistically, I propose that a macro-variable M can be measured as weakly emergent from a set of micro-variables $\mathbf{m}$ ($\mathbf{m} = m_1 \ldots m_N$) to the extent that: (i) past observations of M help predict future observations of M with greater accuracy than predictions based on past observations of $\mathbf{m}$ alone, and (ii) past observations of $\mathbf{m}$ help predict future observations of M with greater accuracy than predictions based on past observations of M alone.

The first condition provides an objective measure of the non-triviality of micro-to-macro inferential pathways, and the second checks for micro-to-macro causal dependence. This definition is relative to a choice of macro and micro description levels and is also relative to a choice of prediction method. As described below, an appropriate framework for prediction is provided by a statistical definition of causality first introduced by Granger (1969), in recognition of which the present measure is called *G-emergence*.

G-causality

In 1969 Granger introduced the idea of 'Granger causality' (G-causality) as a formalization, in terms of linear regression modelling, of Wiener's intuition that Y 'causes' X if knowing Y helps predict the future of X (Granger, 1969; Seth, 2007a). According to G-causality, Y causes X if the inclusion of past observations of Y reduces the prediction error of X in a linear regression model of X and Y, as compared to a model which includes only previous observations of X. Since its introduction, G-causality has found wide application in economics and many other fields including neuroscience and climatology (Ding et al., 2006; Seth, 2008a). To illustrate G-causality, suppose that the temporal dynamics of two time series, $X_1(t)$ and $X_2(t)$ (both of length T), can be described by a bivariate autoregressive model:

$$X_1(t) = \sum_{j=1}^{p} A_{11,j} X_1(t-j) + \sum_{j=1}^{p} A_{12,j} X_2(t-j) + \xi_1(t)$$

$$X_2(t) = \sum_{j=1}^{p} A_{21,j} X_1(t-j) + \sum_{j=1}^{p} A_{22,j} X_2(t-j) + \xi_2(t)$$

where p is the maximum number of lagged observations included in the model (the model order, $p < T$), A contains the coefficients of the model, and ξ_1, ξ_2 are the residuals (prediction errors) for each time series. If the variance of ξ_1 (or ξ_2) is reduced by the inclusion of the X_2 (or X_1) terms in the first (or second) equation, then it is said that X_2 (or X_1) *G-causes* X_1 (or X_2). Assuming that X_1 and X_2 are covariance stationarity (i.e., unchanging mean and variance), the magnitude of this interaction can be measured by the log ratio of the prediction error variances for the restricted (R) and unrestricted (U) models:

$$gc_{2\rightarrow 1} = log\frac{var(\xi_{1R(12)})}{var(\xi_{1U})},$$

where $\xi_{1R(12)}$ is derived from the model omitting the $A_{12,j}$ (for all j) coefficients in the first equation and ξ_{1U} is derived from the full model. Importantly, G-causality is easy to generalize to the multivariate case in which the G-causality of X_1 is tested in the context of multiple variables $X_2 \dots X_N$ ($X_i \neq X_j$ for all $X_{i,j}$). In this case, X_2 G-causes X_1 if knowing X_2 reduces the variance in X_1's prediction error when the activities of all other variables $X_3 \dots X_n$ are also included in the regression model (see below). For a tutorial introduction to Granger causality, see Seth (2007a).

G-autonomy

A simple extension of G-causality allows quantification of the 'statistical autonomy' of a variable with respect to a set of other variables (Seth, 2007b). In this case, instead of asking whether the prediction error of X_1 is reduced by including past observations of X_2, we ask whether the prediction of error of X_1 is reduced by inclusion of its *own* past, given a set of external variables. That is, a variable X_1 is *G-autonomous* to the extent that its own past states help predict its future states over and above predictions based on past states of a set of external variables $X_2 \dots X_N$. By analogy with G-causality, the G-autonomy of X_1 with respect to X_2 is given by:

$$ga_{X_1|X_2} = log\frac{var(\xi_{1R(11)})}{var(\xi_{1U})},$$

where $\xi_{1R(11)}$ is derived from the model omitting the $A_{11,j}$ (for all j) coefficients in the Granger equations.

G-autonomy amplifies the notion of autonomy as 'self-determination' in contrast to other more abstract notions such as 'organizational closure' (Varela, 1979). It is consistent with the notion that a (behaviorally) autonomous system should not be fully determined by its environment, and that a random system should not have high autonomy (Bertschinger et al., 2008). Put simply, a variable is G-autonomous to the extent that it is dependent on its own history and that these dependencies are not accounted for by external factors. Previously I have shown that G-autonomy behaves as expected for simple model systems and can increase as a result of evolutionary adaptation (Seth, 2007b).

G-emergence

Having defined G-causality and G-autonomy the extension to G-emergence is straightforward. A macro-variable M is *G-emergent* from a set of micro-variables $\mathbf{m}$ if and only if (i) M is G-autonomous with respect to $\mathbf{m}$ and (ii) M is G-caused by $\mathbf{m}$. A simple measure for the G-emergence of M from $\mathbf{m}$ is therefore given by

$$ge_{M|\mathbf{m}} = ga_{M|\mathbf{m}}\left(\frac{1}{N}\sum_{i=1}^{N} gc_{m_i\rightarrow M}\right)$$

This measure captures three basic intuitions about weak emergence (Bedau, 1997): that it is a subset of nominal emergence, that it involves dependence on underlying processes, and that it involves autonomy from underlying processes. Importantly, $ge_{M|\mathbf{m}}$ will be zero either if M is independent of $\mathbf{m}$ or if M is fully predicted by $\mathbf{m}$.

Under what circumstances could G-emergence be high? A macro variable could be G-emergent from a set of micro variables if there are 'hidden' or 'latent' influences, i.e., relevant micro causal factors not represented in the regression. However, even if all micro causal factors are present G-emergence can still arise because of dependence on the prediction algorithm used. It is plausible, and indeed necessary for G-emergence to be useful in practice, that in some cases the macro variable is more epistemically transparent to the prediction algorithm than is the collection of micro variables. This, too, is consistent with Bedau's weak emergence, where 'underivability except by simulation' is replaced by '(un)predictability by Granger causality'.

An obvious criticism of G-emergence as measured using linear modeling is that a macro-variable may appear to be G-emergent in virtue of being a nonlinear function of its micro-level components. Clearly, a satisfying measure of emergence should not rely on the failure of linear methods to detect nonlinear dependencies. Fortunately, it is easy to extend G-causality (and hence G-autonomy and G-emergence) to nonlinear situations, for instance by a Taylor expansion:

$$\begin{aligned}
X_1(t) = &\sum_{k=1}^{q}\sum_{j=1}^{p} A_{11,j,k}{X_1}^k(t-j) + \\
&\sum_{k=1}^{q}\sum_{j=1}^{p} A_{12,j,k}{X_2}^k(t-j) + \\
&\sum_{k=1}^{q}\sum_{j=1}^{p} A_{13,j,k}{X_3}^k(t-j) + \xi_1(t) \qquad (1)\\
X_2(t) = &\sum_{k=1}^{q}\sum_{j=1}^{p} A_{21,j,k}{X_1}^k(t-j) + \\
&\sum_{k=1}^{q}\sum_{j=1}^{p} A_{22,j,k}{X_2}^k(t-j) +
\end{aligned}$$

$$\sum_{k=1}^{q}\sum_{j=1}^{p} A_{23,j,k} {X_3}^k(t-j) + \xi_2(t) \quad (2)$$

$$X_3(t) = \sum_{k=1}^{q}\sum_{j=1}^{p} A_{31,j,k} {X_1}^k(t-j) + \sum_{k=1}^{q}\sum_{j=1}^{p} A_{32,j,k} {X_2}^k(t-j) + \sum_{k=1}^{q}\sum_{j=1}^{p} A_{33,j,k} {X_3}^k(t-j) + \xi_3(t) \quad (3)$$

where q is the number of polynomial terms to be included in the Taylor expansion; note the set of variables has been expanded to three in order to illustrate extension to the multivariate ($n > 2$) case. In this example:

$$ge_{X_1|X_2,X_3} = \log\frac{var(\xi_{1R(11)})}{var(\xi_{1U})} \times \frac{1}{2}\left(\log\frac{var(\xi_{1R(12)})}{var(\xi_{1U})} + \log\frac{var(\xi_{1R(13)})}{var(\xi_{1U})}\right) \quad (4)$$

where, following the previous convention, $\xi_{1R(ab)}$ is derived from the model omitting the A_{ab} coefficients in (3). A value of linear or nonlinear G-emergence can be considered statistically significant if the corresponding G-autonomy and G-causality measures are themselves statistically significant. This can be assessed by F-tests on the null hypothesis that the coefficients in A_{11} (G-autonomy) and $A_{12} \ldots A_N$ (G-causality) are zero (Granger, 1969; Geweke, 1982).

It is worth noting that the concept of G-emergence does not depend on using a particular method for nonlinear regression. There exist other more sophisticated methods than Taylor expansions which can be less sensitive to noisy observations and which involve fewer parameters. For example, Ancona et al. (2004) have shown that radial basis functions can serve as effective regression kernels for measuring nonlinear Granger causality. However, for present purposes the Taylor method is preferable because (i) it is simple to describe and to implement, (ii) statistical significance can easily be assessed, and (iii) it supplies an explicit formula for G-emergence (4). Finally, note that the value of G-emergence will depend on the set of micro-variables included in $\mathbf{m}$. Therefore, in heterogenous systems it will be possible to identify a *G-emergence set* as that set of micro-variables which maximizes $ge_{M|\mathbf{m}}$.

Example: Flocking

I now show that G-emergence behaves appropriately in a simple computational model of property emergence. As noted above, a canonical example of property emergence is flocking behavior among birds. In a seminal work in artificial life, Reynolds (1987) showed that visually compelling bird flocking can be simulated by combining three simple rules for simulated birds (boids):

- *aggregation*, each boid tends to fly towards the perceived centre-of-mass (CM) of the flock,
- *avoidance*, each boid tends to avoid colliding with other nearby boids,
- *matching*, each boid tends to align its velocity with that of other nearby boids.

Here, a simple boids simulation is used to test whether visually compelling flocking correlates with high G-emergence of the CM of the flock (the macro-variable) with respect to the trajectories of the individual boids (the micro-variables).

$N = 10$ boids were simulated in a toroidal square environment of length 200 (all dimensions and distances are in arbitrary units; speeds are given in units per time-step). Boids were initialized with positions and velocities randomly chosen from the range [0,200] (x, y position), $[0, 2\pi]$ (heading), and [3,9] (speed). At each time-step the heading α_i and speed s_i of each boid i were updated synchronously according to:

$$\alpha_i = \alpha_i + a_1\theta_1 + a_2(\pi + \theta_2) + a_3\theta_3 + r_1,$$
$$s_i = s_i + a_4\overline{ds} + r_2,$$

where θ_1 is the bearing to the perceived CM (i.e., the CM not including boid i), θ_2 is the bearing to the nearest boid, θ_3 is the bearing to the mean heading of all other boids within a 20 unit range, $\overline{ds}$ is the difference between the speed of boid i and the mean speed of all other boids within 20 units, and r_1 and r_2 are random numbers in the range [-0.01,0.01]. The parameter vector a (all $a \in [0, 1]$) determines the relative contribution of each factor. Toroidal distances were calculated in the standard way, according to the minimum distance either across, or not across, the boundary. CM positions were calculated iteratively in order to minimize the toroidal distance to each boid (i.e., not as the average boid position, which leads to boundary artifacts).

Three different conditions were tested. Condition R (random) produced near-random boid behavior ($a_R = [0.01, 0.01, 0.01, 0.01]$). Condition L (low) evoked poor flocking behavior by imposing a strong dependence on velocity matching; boids in this condition tended to move in semi-rigid formations ($a_L = [0.1, 0.1, 0.6, 0.6]$). Condition H (high) evoked compelling flocking behavior; the parameter set ($a_H = [0.1, 0.3, 0.3, 0.3]$) was selected by hand. Examples of boid and CM trajectories from each condition are shown in Figure 2. Although static images do not fully capture the dynamic nature of flocking it is clear that boid trajectories in condition H are more flock-like than those in conditions L and R.

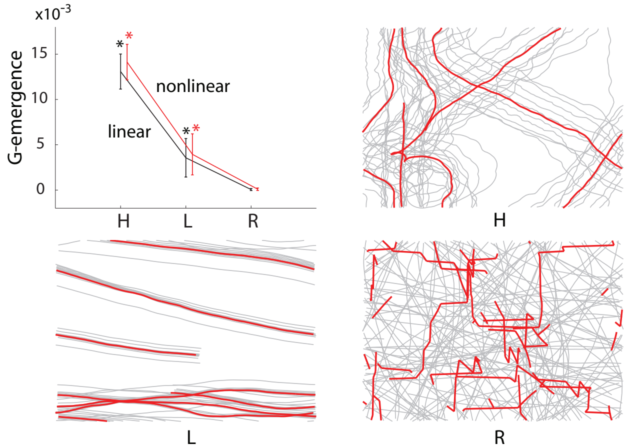

Figure 2: G-emergence of the centre-of-mass (CM) of a boid flock. Top left: Mean and standard deviation linear and nonlinear G-emergence by condition (asterisks show statistical significance). Other panels: Example trajectories (500 time-step segment) of the boids (grey) and CM (red) in condition H (high G-emergence), L (low G-emergence), and R (random).

For each condition the boid simulation was run 25 times with each run lasting 5000 time-steps; for each run the x, y coordinates of each boid and the global CM were recorded. Several preprocessing steps were carried out prior to calculation of G-emergence. In order to reduce the dimensionality of the dataset and to provide further robustness against boundary effects, each x, y coordinate pair was transformed into a single variable reflecting distance from the centre of the environment. The first 500 data points were removed to eliminate initial transients and each resulting time series was transformed into its zero-mean equivalent. Finally, each time series was first-order differenced in order to ensure covariance stationarity (Seth, 2005). Following preprocessing, for each run in each condition both linear and nonlinear G-emergence of the CM were computed using ordinary least squares regression. I chose a model order $p = 5$ and (for the nonlinear analysis) a polynomial order $q = 3$. The model order was selected based on the average Akaike information criterion (Seth, 2007a) across all 75 runs.

Figure 2 shows the mean linear and nonlinear G-emergence of the CM in each condition. Confirming the prediction that high G-emergence tracks compelling flocking, both linear and nonlinear measures show significantly higher values of G-emergence in condition H than in conditions L and R. All values of G-emergence in conditions in H and L were significant ($P < 10^{-5}$ for G-autonomy and G-causality, two-tailed t-test); those in condition R were not.

To test the behavior of G-emergence across different parameter combinations in the boids model, I computed linear and nonlinear G-emergence for each parameter vector in the

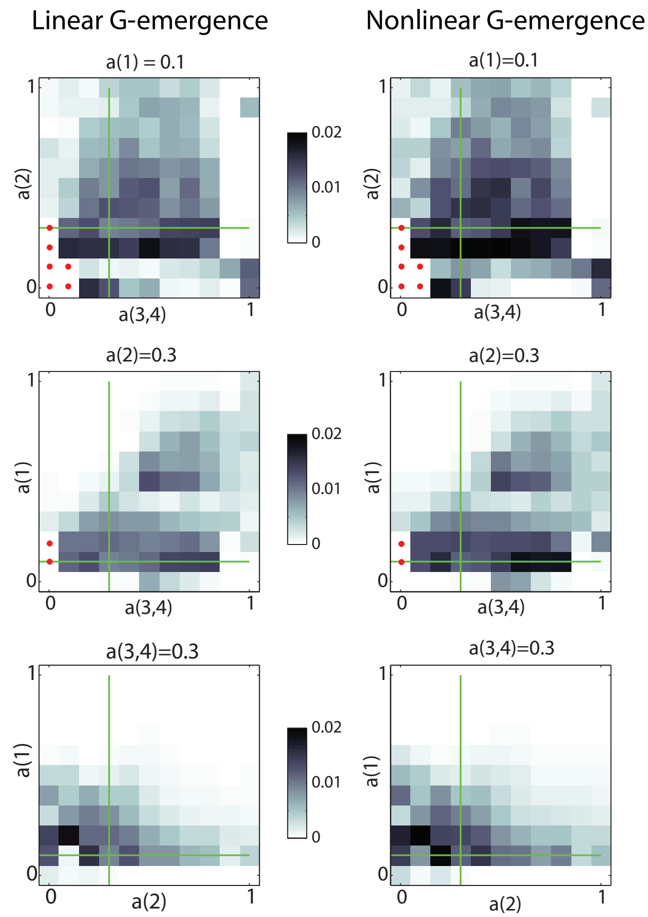

Figure 3: Parameter space of the boids model. The parameter vector a_H is indicated by the intersection of the green lines. Grey scale shows average linear and nonlinear G-emergence of the global CM. Each value is the average of three evaluations of 5000 time-steps each. Red dots indicate parameter combinations which lead to reliably non-stationary time-series.

space $a_{(1,2,3)} \in [0.0, 0.1, \ldots, 1.0]$. Parameters a_3 and a_4 were yoked together because they both influence the same rule (velocity matching) and three evaluations were carried out for each vector, requiring a total of $11 \times 11 \times 11 \times 3 = 3993$ evaluations. Figure 3 shows G-emergence for three orthogonal cross-sections through the three-dimensional parameter space; in each cross-section the vector corresponding to a_H (condition H) is marked by the intersection of the green lines.

Several aspects of the above cross-sections are notable. First, linear and nonlinear G-emergence are strongly correlated suggesting that even linear measures can provide insight into emergent properties in some complex systems. Second, in most regions of parameter space G-emergence changes smoothly, suggesting it is a robust measure. However, some regions show sharp transitions, for example be-

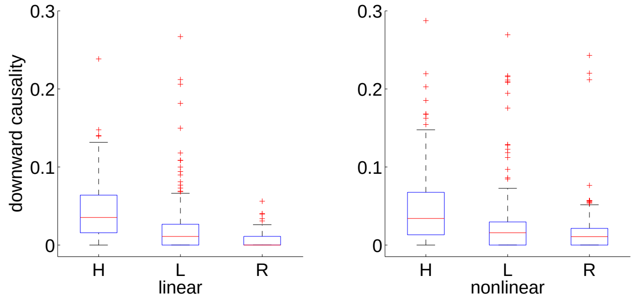

Figure 4: Downward causation is significantly higher in condition H than in conditions L or R. Boxplots show linear and nonlinear Granger causality from the global CM to individual boids, calculated separately for each boid for all 25 runs in each condition (i.e., 250 values per boxplot). Non-significant causalities were set to zero (nominal threshold of 0.01, Bonferroni corrected to 10^{-5}). Resulting distributions are non-normal and differences between conditions were tested using the Wilcoxon rank sum test. For both linear and nonlinear analyses all pairwise comparisons among medians were significantly different ($P < 10^{-3}$). Each boxplot shows lower quartile, median, and upper quartile values; whiskers show range of remaining data and '+' denote outliers.

tween some vectors with $a_1 = 0$ and neighboring vectors. The sensitivity of G-emergence to these transitions indicates that it can usefully identify parameter regions of complex models in which non-trivial weak emergence is present.

Downward causation

A common intuition regarding emergence is that it involves 'downward' causation from macro-levels to micro-levels. For proponents of strong emergence, downward causation is in fact an essential aspect of what it means to be emergent (Kim, 1999). However, physical interpretations of downward causality pose tricky metaphysical problems, for example, how to resolve competing micro- and macro- causes (Bedau, 2003). G-emergence, being statistically defined, presents a metaphysically innocent alternative according to which downward causality is reflected by G-causality from the macro-variable(s) to the micro-variable(s).

Figure 4 shows downward (Granger) causation from the global CM to the individual boid trajectories, for both linear and nonlinear G-causality measures. Averages are taken across all boids and across all 25 runs in each condition. Consistent with an association between emergence and downward causation, both measures of downward causation are significantly higher in condition H than in conditions R or L. Despite this result, it seems possible in principle for weak emergence to occur without downward causation (of course strong emergence requires downward causality by definition). Having separately applicable measures of weak emergence and downward causation makes it possible to explore conditions (if any) in which emergence and downward causation do not occur together, potentially refining and deepening the concept of emergence.

Discussion

In this paper I have introduced a method for detecting the degree of weak emergence in a system by performing physical measurements on it. Because this measure it based on a statistical interpretation of causality it sidesteps conceptual pitfalls such as competition among micro- and macro-causes, and it provides an objective and graded assessment of the non-triviality of micro-to-macro inferential pathways. MATLAB (Mathworks, Natick, MA) code for calculating G-emergence from arbitrary time-series data is provided on the author's website, www.anilseth.com.

Diachronic emergence

Diachronic (temporal) emergence refers to the appearance of new properties over time, as exemplified by evolution and development. A diachronically emergent process is by definition statistically non-stationary and therefore not amenable to direct measurement by G-emergence. Nonetheless, it is plausible that a diachronically emergent process is bracketed by statistically stationary periods with different G-emergence properties. In this way, G-emergence could be used to indirectly to infer diachronic emergence.

Relation to other measures

The intuition that differences in predictive ability may be important in defining macro-level properties is shared by (Shalizi and Moore, 2006). However these authors focus on clarifying the concept of a macro-state and they do not explicitly combine measures of autonomy and causal dependence. Rather, one process is called emergent from another if it has a higher 'predictive efficiency' than the process it derives from. Their measure of predictive efficiency is based on information-theoretic model reconstruction [the epsilon-machine concept, see (Crutchfield, 1994)], which is powerful but less easy to apply in practice than the time series metrics described here. A related approach is taken by Polani (2006) in which an 'emergent description' involves a further step of decomposing systems into independent informational sub-components.

According to the 'contextual emergence' of Atmanspacher (2007), derivation of macro-level properties requires knowledge of micro-level properties and of contingent contextual conditions, the latter defined in terms of stability criteria according to a dynamical systems analysis. This concept diverges from the doctrine of causal fundamentalism (or supervenience) by proposing that micro-level properties offer necessary but not sufficient conditions for deriving macro-level properties, which is suggestive of strong emergence. An explicit measure of strong emergence is offered by Bar-Yam (2004) which is based on measuring the entropy of a system at multiple scales. Oscillations in 'multiscale variety' are suggested to reveal constraints on the values of multiple variables which are not present among subsets of these variables, and the existence of such constraints is taken to indicate strong emergence. However, since on the present account strong emergence rejects mechanistic explanations altogether, a full analysis of Bar-Yam's measure is beyond the present scope.

Phase transitions

Physicists have recently become interested in the onset of collective behaviors among boid-like self-propelling particles (Vicsek et al., 1995; Gregoire et al., 2003). In such systems, phase transitions can be observed among 'gaseous' phases (each particle moves independently), 'liquid' phases (particles move collectively but still diffuse with respect to each other) and 'solid' phases (particles move collectively and remain fixed with respect to each other). Plausibly, these phases correspond respectively to conditions R, H, and L of the present model and the sharp boundaries noted in figure 3 may correspond to phase transitions. However, phase transition analyses tend to focus on the dynamics of transition and assume emergent behavior is phenomenologically obvious in some phases and is absent in others. In contrast, the present focus is on detecting the *degree* of emergence by making physical measurements on a system.

Strong emergence and consciousness

As already noted, strong emergence differs fundamentally from weak emergence in that (strongly) emergent properties are suggested to be causally irreducible to their micro-level components and to exert downwardly causal influences on these components (Kim, 2006). Strong emergence thus poses a radical challenge to science because it implies that there exist real properties in the world that do not 'bottom out' in known sorts of physical interactions.

David Chalmers has made explicit a recurring idea, which is that there is exactly one clear case of a strongly emergent phenomenon, and that is the phenomenon of consciousness (Chalmers, 2006). It seems that two commonly held intuitions about consciousness drive this suspicion. First, the idea that even complete knowledge of the physical interactions sustained by brains will not provide an understanding what is like to have a conscious experience: this is the infamous 'hard problem' of consciousness. Second, the intuition that conscious states have causal efficacy in the world, as exemplified by the notion of free will but which runs through all aspects of consciousness; after all, why have experiences at all if they don't do anything? These intuitions map cleanly onto the defining features of strong emergence, namely that macro-level properties in principle cannot be identified from micro-level observations, and that macro-level properties have irreducible causal powers.

These intuitions can however be challenged. First, to expect a scientific resolution to the 'hard problem' as it is presently conceived may be to misunderstand the role of science in explaining nature. A scientific theory cannot presume to replicate the experience it describes or explains; a theory of a hurricane is not a hurricane (Seth and Edelman, 2008). If the phenomenal aspect of experience is irreducible, so is the fact that physics has not explained why there is something rather than nothing, and this has not prevented physicists from laying bare many mysteries. Second, consciousness can be functionally efficacious without recourse to downward causation. It is entirely plausible that certain neural mechanisms support useful functions in virtue of the fact that they entail conscious experiences (Seth, 2008b). For example, the neural mechanisms underlying consciousness may serve to integrate large amounts of information over short time periods, leading to functionally effective high-dimensional discriminations among a large repertoire of sensorimotor scenes (Tononi and Edelman, 1998). Such information integration may entail conscious qualia in just the same way that the molecular structure of hemoglobin entails a particular spectroscopic response: it simply could not be otherwise (Edelman, 2003). Moreover, experiences of 'free will' and 'volition' are just experiences like any other, and there is a wealth of experimental evidence showing, unsurprisingly, that awareness of a voluntary action is preceded by recognizable signatures in neural activity (Libet, 1985). Together, these points suggest that the association of consciousness with strong emergence does not rest on solid ground.

In contrast, it is very likely that the connection between neural mechanism and conscious experience involves *weak* emergence in many ways. A striking feature of conscious experience is that it seems more than the sum of its parts (each conscious experience is a unity) and that it has a vivid temporality (William James' 'stream of consciousness'). Models of consciousness that can be analyzed in terms of weak emergence therefore have the potential to *explain* features of phenomenology in terms of dynamical processes at the level of neural mechanism. The development and experimental testing of such 'explanatory correlates' (Seth and Edelman, 2008) is a highly promising avenue towards a scientific description of consciousness. It is exciting to consider that measures of weak emergence may eventually find utility in accounting for apparent free will and in crossing the explanatory gap between neural mechanism and phenomenal experience.

Conclusions

Scientific progress in understanding a phenomenon relies on the ability to measure that phenomenon. 'Emergence' has thus far resisted the development of useful measures, perhaps because of a suspicion that it necessarily involves violation of mechanistic/reductionistic explanations. But this suspicion is only valid for 'strong' emergence and proposed measures of strong emergence are correspondingly difficult to apply and interpret (Bar-Yam, 2004). In this paper I have developed and illustrated a quantitative, intuitive, and practically straightforward measure of weak emergence. G-emergence is based on the intuition that emergent properties are both *dependent on* and *autonomous from* their components (Bedau, 1997) and is operationalized using linear and nonlinear time series analysis.

In a simulation of bird flocking, visually compelling flocking behavior is accompanied by high G-emergence as compared to random movement or flight in rigid formations. High G-emergence is also accompanied by downward (Granger) causation from the flock to each boid, though this may not be the case for all systems. Finally, G-emergence provides a platform for measuring other sorts of emergence; for instance 'temporal emergence' and/or 'self-organization' could be measured as the change in G-emergence between two different time periods.

References

Ancona, N., Marinazzo, D., and Stramaglia, S. (2004). Radial basis function approaches to nonlinear granger causality of time series. *Physical Review E*, 70:056221.

Atmanspacher, H. (2007). Contextual emergence from physics to cognitive neuroscience. *Journal of consciousness studies*, 14:18–36.

Bar-Yam, Y. (2004). A mathematical theory of strong emergence using multiscale variety. *Complexity*, 9(6):15–24.

Bedau, M. (1997). Weak emergence. *Philosophical Perspectives*, 11:375–399.

Bedau, M. (2003). Downward causation and the autonomy of weak emergence. *Principia*, 6:5–50.

Bertschinger, N., Olbrich, E., Ay, N., and Jost, J. (2008). Autonomy: an information theoretic perspective. *Biosystems*, page in press.

Broad, C., editor (1925). *The mind and its place in nature*. Routledge and Kegan Paul, London.

Chalmers, D. (2006). Strong and weak emergence. In Clayton, P. and Davies, P., editors, *The re-emergence of emergence*. Oxford University Press, Oxford.

Chang, H., editor (2004). *Inventing temperature: Measurement and scientific progress*. Oxford University Press, New York, NY.

Crutchfield, J. (1994). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54.

Ding, M., Chen, Y., and Bressler, S. (2006). Granger causality: Basic theory and application to neuroscience. In Schelter, S., Winterhalder, M., and Timmer, J., editors, *Handbook of Time Series Analysis*, pages 438–460. Wiley, Wienheim.

Edelman, G. (2003). Naturalizing consciousness: A theoretical framework. *Proceedings of the National Academy of Sciences, USA*, 100(9):5520–5524.

Geweke, J. (1982). Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association*, 77:304–13.

Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37:424–438.

Gregoire, G., Chate, H., and Tu, Y. (2003). Moving and staying together without a leader. *Physica D*, 181:151–170.

Jackson, F. and Pettit, P. (1992). In defence of explanatory ecumenism. *Economics and Philosophy*, 8:1–21.

Kim, J. (1999). Making sense of emergence. *Philosophical Studies*, 95:3–36.

Kim, J. (2006). Emergence: core ideas and issues. *Synthese*, 151:547–559.

Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and brain sciences*, 8:529–566.

Polani, D. (2006). Emergence, intrinsic structure of information, and agenthood. *Interjournal Complex Systems*, page 1973.

Reynolds, C. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34.

Seth, A. (2005). Causal connectivity of evolved neural networks during behavior. *Network: Computation in Neural Systems*, 16:35–54.

Seth, A. (2007a). Granger causality. *Scholarpedia*, 2(7):1667.

Seth, A. (2007b). Measuring autonomy via multivariate autoregressive modelling. In Almeida e Costa, F., editor, *Proceedings of the Ninth European Conference on Artificial Life*, pages 475–485. Springer-Verlag.

Seth, A. (2008a). Causal networks in simulated neural systems. *Cognitive Neurodynamics*, 2:49–64.

Seth, A. (2008b). Functions of consciousness. In Banks, W., editor, *Springer Encyclopedia of Consciousness*, page in press. Springer Verlag.

Seth, A. and Edelman, G. (2008). Consciousness and complexity. In Meyers, R., editor, *Springer Encyclopedia of Complexity and Systems Science*, page in press. Springer Verlag.

Shalizi, C. and Moore, C. (2006). What is a macrostate: Subjective observations and objective dynamics. http://arxiv.org/abs/cond-mat/0303625.

Sperry, R. (1969). A modified concept of consciousness. *Psychological Review*, 76:532–536.

Tononi, G. and Edelman, G. (1998). Consciousness and complexity. *Science*, 282:1846–1851.

van Gulick, R. (2001). Reduction, emergence and other recent options on the mind-body problem. *Journal of Consciousness Studies*, 8:1–34.

Varela, F. (1979). *Principles of biological autonomy*. Elsevier, North Holland, NY.

Vicsek, T., Czirok, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75:1226.