






ARTICLE

<https://doi.org/10.1057/s41599-019-0371-1>

OPEN

On revolutions

Armand M. Leroi ^{1,2*}, Ben Lambert ^{3,4}, Matthias Mauch ⁵, Marina Papadopoulou ^{1,6}, Sophia Ananiadou⁷, Staffan I. Lindberg ^{8,9} & Patrik Lindenfors^{9,10,11}

ABSTRACT Sometimes the normal course of events is disrupted by a particularly swift and profound change. Historians have often referred to such changes as “revolutions”, and, though they have identified many of them, they have rarely supported their claims with statistical evidence. Here, we present a method to identify revolutions based on a measure of multivariate rate of change called Foote novelty. We define revolutions as those periods of time when the value of this measure is, by a non-parametric test, shown to significantly exceed the background rate. Our method also identifies conservative periods when the rate of change is unusually low. We apply it to several quantitative data sets that capture long-term political, social and cultural changes and, in some of them, identify revolutions — both well known and not. Our method is general and can be applied to any phenomenon captured by multivariate time series data of sufficient quality.

¹Department of Life Sciences, Imperial College London, London SW7 2AZ, UK. ²Data Science Institute, Imperial College London, London SW7 2AZ, UK. ³Department of Zoology, University of Oxford, New Radcliffe House, Woodstock Road, Oxford OX2 6GG, UK. ⁴Department of Infectious Disease Epidemiology, Imperial College London, Medical School Building St Mary's Campus, Norfolk Place, London W2 1PG, UK. ⁵Centre for Digital Musicology, Queen Mary University, Mile End Road, London E1 4NS, UK. ⁶Groningen Institute for Evolutionary Life Sciences, University of Groningen, Nijenborgh 7, 9747 AG Groningen, Netherlands. ⁷National Centre for Text Mining, School of Computer Science, University of Manchester, 131 Princess Street, Manchester M1 7DN, UK. ⁸V-Dem Institute, Gothenburg University, Box 100, Gothenburg SE-405 30, Sweden. ⁹Department of Political Science, Gothenburg University, Box 100, Gothenburg S-405 30, Sweden. ¹⁰Institute for Future Studies, Box 591, SE-101 31 Stockholm, Sweden. ¹¹Centre for the Study of Cultural Evolution and Department of Zoology, Stockholm University, SE-106 91 Stockholm, Sweden. *email: a.leroi@imperial.ac.uk

Introduction

What is a revolution? It seems that the word “revolution” was first applied to sublunary events when parliamentarians, aided by the Dutch, deposed James II from the English throne and so brought about the Glorious Revolution. Since then, it has been applied ever more widely (Cohen, 1986). Responding to the French Revolution of 1789, Friederich Schlegel called for an Aesthetic Revolution in poetry, and so extended the term beyond politics (Heumakers, 2015). In the latter half of the 19th century—an age of revolutions—John Stuart Mill, Karl Marx and Arnold Toynbee, following a French coinage, wrote of the Industrial Revolution (Bezanson, 1922). In the 1950s Alexandre Koyré, Herbert Butterfield, A. R. Hall, and Thomas Kuhn, described the Scientific Revolution (Koyré, 1957; Butterfield, 1950; Hall, 1954; Kuhn, 1957; Cohen, 1994). In his 1972 book, *The Structure of Scientific Revolutions*, Kuhn generalized the idea, arguing that science advanced, if it advanced at all, by revolutions (Kuhn, 1972). The Darwinian Revolution was swiftly identified (Ruse, 1979; Himmelfarb, 1996), as were many others. Indeed, Kuhn’s book prompted something of a revolution in scientific discourse, as scientists themselves took to identifying, or calling for, “paradigm shifts”—Kuhn’s term for a revolution—in their fields. A search of all articles indexed by the Web of Science in 2017 reveals more than two thousand that do so, though many of the purported revolutions seem quite modest in scope (e.g., Seward (2017); Raoult (2017); Lowenstein and Grantham (2017); Lonne (2017)).

For all that, revolutions are hard to pin down. Upon close inspection they often seem to shrink. Pick a revolution, even a famous and well-documented one, and you can be sure to find scholars who have sought to cut it down to size or even deny that it happened at all. “The drastic social changes imputed to the Revolution, seem less clear-cut or not apparent at all.”—thus Simon Schama on how his generation of historians viewed the impact of the French Revolution (Schama, 1989). “There was no such thing as the Scientific Revolution, and this is a book about it.”—so Steve Shapin, in paradoxical mode, on early modern science (Shapin, 1996). Evolutionary biologists may be surprised to learn that the Darwinian Revolution has its skeptics too (Hodge, 2005; Bowler, 1988).

The difficulty of identifying revolutions has plagued the historical natural sciences as well. In the 1980s archeologists labeled the sudden appearance, fifty thousand years ago, of culture as the Human Revolution (Mellars and Stringer, 1989). It was not long before others had dismissed it as the “revolution that wasn’t” (McBrearty and Brooks, 2000). For much of his life Stephen Jay Gould argued that the Darwinian Revolution had run its course and that evolutionary biology needed another (Gould, 2002). (But one not to be confused with the broader Paleobiological Revolution of the 1970s and 80s which he helped shape (Sepkoski, 2012; Sepkoski and Ruse, 2009).) The coping stone of Gould’s new paradigm, an unstable edifice, was the theory of punctuated equilibrium that he proposed with Niles Eldredge in 1972 (Eldredge and Gould, 1972). This theory, shorn of its theoretical structure, postulated that change in fossil lineages is itself best described as stasis interrupted by periods of rapid change rather than gradual evolution. It may seem like a simple matter to decide which, but the ensuing decades-long quarrel among paleontologists about what the fossils show has proved otherwise (Pennel et al., 2013).

Unsurprisingly given its fame, the idea of punctuated evolution has spread to other fields. Natural languages (Atkinson et al., 2008; Dediú and Levinson, 2012; Greenhill, 2017), computer languages (Valverde and Solé, 2015), technology (De Dreu and

van Dijk, 2018), and socio-political structures (Spencer, 1990; Currie and Mace, 2011; Turchin, 2018) have all been claimed to evolve in a punctuated fashion; and theoretical models that explain why they might do so have been developed (Kolodny et al., 2015, 2016). But the idea also appears now in fields as remote from paleontology as management science and policy research (e.g., Flink (2017); Fowler et al. (2017); De Ruiter and Schalk (2017)). In them the term has lost its deeper meaning altogether and is just another way to express the existence, or hope, of revolutions.

The problem is clear. Great revolutions may entail change in many dimensions—ideas, wealth, social roles, political structures, the composition of assemblages of artefacts and species or else their features—but to varying degrees, at varying rates, and with varying starts and ends. A revolution’s visibility, then, depends on where you look. Even when considering the same data, some scholars will see discontinuity where others see continuity—it may be merely a matter of temperament—in the absence of an objective method for distinguishing the two, there is no way to know which of their accounts is more true. It seems desirable, then, to give the detecting of revolutions a statistical foundation.

At a minimum, the idea of a revolution supposes a sudden acceleration in the rate of change. The most direct way to estimate a rate, or its increase, is from time series data, and various methods have been developed to just that. Such methods identify “regime shifts”, “phase shifts” and the like in time series data (Andersen et al., 2009); some of them have even been used to detect punctuated events in evolving lineages (Hunt, 2010, 2012; Hunt et al., 2015). Evolutionary biologists, however, generally lack good time series data: the fossil record is imperfect, and only a few populations of living species have been studied for long periods of time (but see Lambert et al., n.d.). They do, however, have an abundance of phylogenies. Phylogenies provide a window into the past: they allow models of evolutionary rates to be tested from the distribution of phenotypes across extant species (Pennel et al., 2013; Atkinson et al., 2008; Pagel et al., 2006; Bokma, 2008; Harmon, 2010; Venditti and Pagel, 2010; Uyeda et al., 2011; Landis et al., 2012; Duchon, 2017; Landis and Schraiber, 2017). Such studies show that evolutionary rates are not constant, but sometimes accelerate during speciation or else the invasion of new adaptive zones.

These studies all concern the evolution of isolated characteristics such as body size. But revolutions are usually thought of as times when *many* things change at once. For this reason, here we define revolutions as statistically significant local increases in the *multivariate* rate of change relative to the background rate. Following this definition, we outline a revolution-detector that identifies such rate changes in time series by means of a non-parametric permutation test. Existing multivariate time series segmentation methods work by dividing series into self-contained windows that assume a time-invariant (typically parametric) relationship (Omranian et al., 2015; Preuss et al., 2015). Our method, by contrast, uses the local multivariate rate of change of series across neighboring periods to classify time points into “revolutionary” or “conservative” periods. As such, it most resembles some methods that have been used to study spatial change in multivariate ecological data sets (Kent et al., 2006). We have previously introduced our method while applying to the evolution of American popular music (Mauch et al., 2015). Here we refine its statistical basis, apply it to several large data sets that capture changes in political, social, and cultural systems over time, and identify a variety of revolutions that are well known, as well as some that are not.

Detecting revolutions

A sketch of the method. We assume that a method for detecting revolutions should consider many characteristics of a population simultaneously, that is, be underpinned by a multivariate metric of change. The idea behind our method is as follows. We begin with a collection of time series of summary statistics that capture the evolving properties of a population. These summary statistics might be frequencies or means or raw observations. For example, such summary statistics might be gene frequencies in an evolving population of fruit flies, topic frequencies in an evolving population of novels, or socio-economic variables in an evolving population of nations. The power of our method to detect revolutions is greatest when series contain ample time steps, and when sampling variances of the summary statistics are small compared to overall variability (i.e., across data pooled over all time periods). These summary statistics will inevitably fluctuate over time. We take a revolution, however, to occur only if many variables change in concert, with exceptional speed.

To locate these revolutionary periods, we calculate a multivariate temporal distance matrix, which computes a univariate measure of discrepancy metering the difference between observations across all possible pairs of time periods. Any distance metric can be used to measure this discrepancy and can be chosen to best suit the data at hand. The temporal distance matrix captures the multivariate rate of change and can be usefully visualized in a heat map displaying pairwise distances between pairs of time periods (rows being the temporally-ordered first time period of a pair and columns being the second). Often this heat map will be without structure and appear as if differences between pairs occur at random. Revolutions, however, induce a quite distinct checkerboard pattern in this matrix. Within the checks, which are blocks of years, the rate of multivariate change is relatively low, but among them it is high. Where the checks meet is where a revolution may have occurred.

To search for the footprints of revolution, we use a quantitative method called Foote novelty. We explain how it works in detail below, but for the moment it is enough to say that Foote novelty mimics our visual perception of checkerboards by applying straightforward mathematical operations to the distance matrix in an attempt to locate checkerboards, should they exist. For a given period, Foote novelty contrasts the variation among diagonal elements (which represent differences between consecutive time periods) of the temporal distance matrix with that among elements from the anti-diagonals (which represent differences between non-consecutive periods) over a given window of time (e.g., five years). Applying Foote novelty to the temporal distance matrix thus returns series representing rates of change over time in which the rate represents a sort of running average, and in which the run length is dictated by the chosen time window. By altering the window width, we can capture the relative rates of change at different temporal scales ranging from one time-step to as many as the data permit. Relatively high Foote novelty scores indicate a potential revolution.

Of course, the Foote novelty scores wobble about over time for any given time series, even when no revolution is present. So we need some way of distinguishing normal, unrevolutionary, variation from extraordinary, revolutionary, variation. We do so by means of a non-parametric statistical test, which compares Foote novelty series calculated on the actual distance matrix with those series obtained from calculation on matrices that have been randomly permuted along their diagonals. These random permutations produce bootstrap samples of the original distance matrices. These, in turn, provide a baseline measurement of Foote Novelty used to define a threshold level, above which a revolution is deemed to have occurred. Like any statistical test, ours is vulnerable to both Type I and II errors. So, using simulations, we have investigated the conditions under which it holds and those

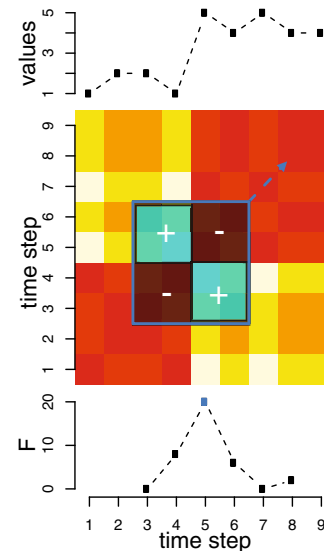


Fig. 1 How Foote novelty works: a simple example. Top. A univariate time series. Middle. A distance matrix of the time series overlain by the Foote novelty kernel. In the distance matrix, higher values are coloured lighter. The 2×2 matrix marked with a blue box delimits the FN kernel and shows the natural checkerboard structure of distances at points of rapid change. Bottom. The Foote novelty values, showing a peak in the middle at the point of maximum change in the series, F_4^2 , also colored blue.

under which it might fail, and suggest some procedures to mitigate these errors.

Calculating Foote novelty. Foote novelty is a multivariate metric of change that has its origin in signal processing (Foote, 2000). Although particularly suited to estimating change in a multivariate time series, here we explain its workings using a single time series.

Consider the time series (1, 2, 2, 1, 5, 4, 5, 4, 4) and its distance matrix (Fig. 1). It has an obvious change point between the fourth and fifth element: the data before (1, 2, 2, 1) and after (5, 4, 5, 4, 4) being both quite homogenous. Calculate a distance matrix, plot them as a heatmap, and these periods appear as distinct blocks of low local variation along the main diagonal (the darker shaded blocks in Fig. 1). By contrast, pairwise distances between data points before and after such a change are larger, resulting in two off-diagonal blocks of high-cross variability (the lighter shaded blocks in Fig. 1) in the distance matrix. The result is a checkerboard. And where the blocks meet pinpoints a large change.

In order to capture this checkerboard structure, Foote devised a kernel that itself looks like a checkerboard, being composed of two pairs of blocks of size k —the half-width—with the diagonal blocks equal to -1 and the off-diagonal components equal to $+1$. For example, the kernel for $k = 2$ is given by,

$$C^2 = \begin{pmatrix} 1 & 1 & -1 & -1 \\ 1 & 1 & -1 & -1 \\ -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix}}_{\text{cross variability}} - \underbrace{\begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{pmatrix}}_{\text{local variability}}$$

This kernel is moved along the main diagonal from one target time step to the next and the elementwise product is taken of the matrix and the values it overlaps (Fig. 1). The diagonal components of the Foote kernel (“local variability”) estimates the differences within the series before and after the target time step, while its off-diagonal components (“cross variability”) estimates the differences among them. The Foote novelty score, then, depends on their relative magnitudes. More formally, Foote novelty at target time step t is:

$$F_t^k = C^k \odot D_t^k,$$

where C^k is the Foote kernel with half-width k , D_t^k represents the submatrix of the temporally-ordered distance matrix D centered about diagonal element (t, t) and of the same dimensions as the Foote kernel, and \odot is the Hadamard (i.e. elementwise) product.

Figure 1 shows how Foote novelty can be calculated for our synthetic one-dimensional time series. In practice, we use a kernel with two small modifications. First, we follow Foote in imposing a radially symmetric Gaussian taper with a standard deviation of $2 \times 0.4k$, to remove edge effects. This amounts to multiplying each grid point by a Gaussian kernel $\frac{1}{4\pi \times 0.4k} e^{-(x^2+y^2)/2(2 \times 0.4k)^2}$, where x and y are the horizontal and vertical distances of the grid point from the center of C . This gives distances closer to the target time point more weight than those further away. Second, in order to have a central point of reference, we add a “cross” of zeros between the blocks of the kernel C . As a result, the size of the whole kernel is $2k + 1$, and the value F_t^k corresponds precisely to the kernel centered at t . The resulting $k = 2$ kernel is given by (to within numerical error),

$$C^2 = \begin{pmatrix} 0.013 & 0.023 & 0 & -0.023 & -0.013 \\ 0.023 & 0.042 & 0 & -0.042 & -0.023 \\ 0 & 0 & 0 & 0 & 0 \\ -0.023 & -0.042 & 0 & 0.042 & 0.023 \\ -0.013 & -0.023 & 0 & 0.023 & 0.013 \end{pmatrix}$$

Assuming that D is large (e.g., 100×100 time points) relative to C (e.g., 4×4), F_t^k can be calculated for different points in time. The calculation of the Foote novelty series $\{F_1^k, F_2^k, \dots, F_T^k\}$ hence amounts to sliding C along the central diagonal of distance matrix D , calculating F_t^k as we go.

Statistical inference. In real data, the distance matrix, D , is nearly always based on a multivariate time series. Foote novelty, F_t^k , is generally positive and varies as the underlying variables fluctuate in value. We therefore define revolutions as periods when its value is statistically significantly higher than in the rest of the series.

To determine this we compare the observed F_t^k values to the distribution of F_t^k values obtained from randomly permuting the distance matrix. In our original test, we permuted the distance matrix on its axes (Mauch et al., 2015); here, following a suggestion by T. Underwood, H. Long, R. J. So, and Y. Zhu (pers. comm.), we permute on the diagonals.

By this we mean the following. A given distance matrix of dimension $p \times p$ has $p - 1$ diagonals that can be independently permuted (not $2(p - 1)$ because distance matrices must be symmetric). The longest non-degenerate diagonal (since the central diagonal is composed only of zeros) is the vector of elements, $v_1 = (D_{1,2}, D_{2,3}, D_{p-1,p})$ and so on. To permute the matrix along its diagonals, we visit each v_j and randomly sort its elements. The same changes are made on both sides of the central diagonal so that the matrix remains symmetric. Permuting the diagonals of the distance matrix retains more of its structure than

permuting its axes does, and so provides a more robust null model. The motivation for this test is based on our empirical results and because it is comparable to carrying out a permutation of the time series in blocks of varying lengths greater than or equal to one time period.

The kernel half-width, k , can be as small as 1 or as large as the data allow, but different k show different aspects of change. Foote novelty acts rather like a microscope. Small k values zoom in on short-term heterogeneities that large k values may obliterate, and large k values may reveal long-term variation invisible at smaller fields of view. A sustained period of rapid change will tend to produce revolutionary signals at many different k , but more complex patterns of rate variation will result in conflicting signals. For example, a large k may well identify a single, long, revolution where a smaller one identifies two or more. The shifting picture of the rate landscape that emerges as we adjust the focus of our Foote Novelty microscope is not a weakness of the method, but a consequence of making the scale of analysis explicit. In practice, we examine all half-widths that the data allow and identify revolutions by their consistency in a given region.

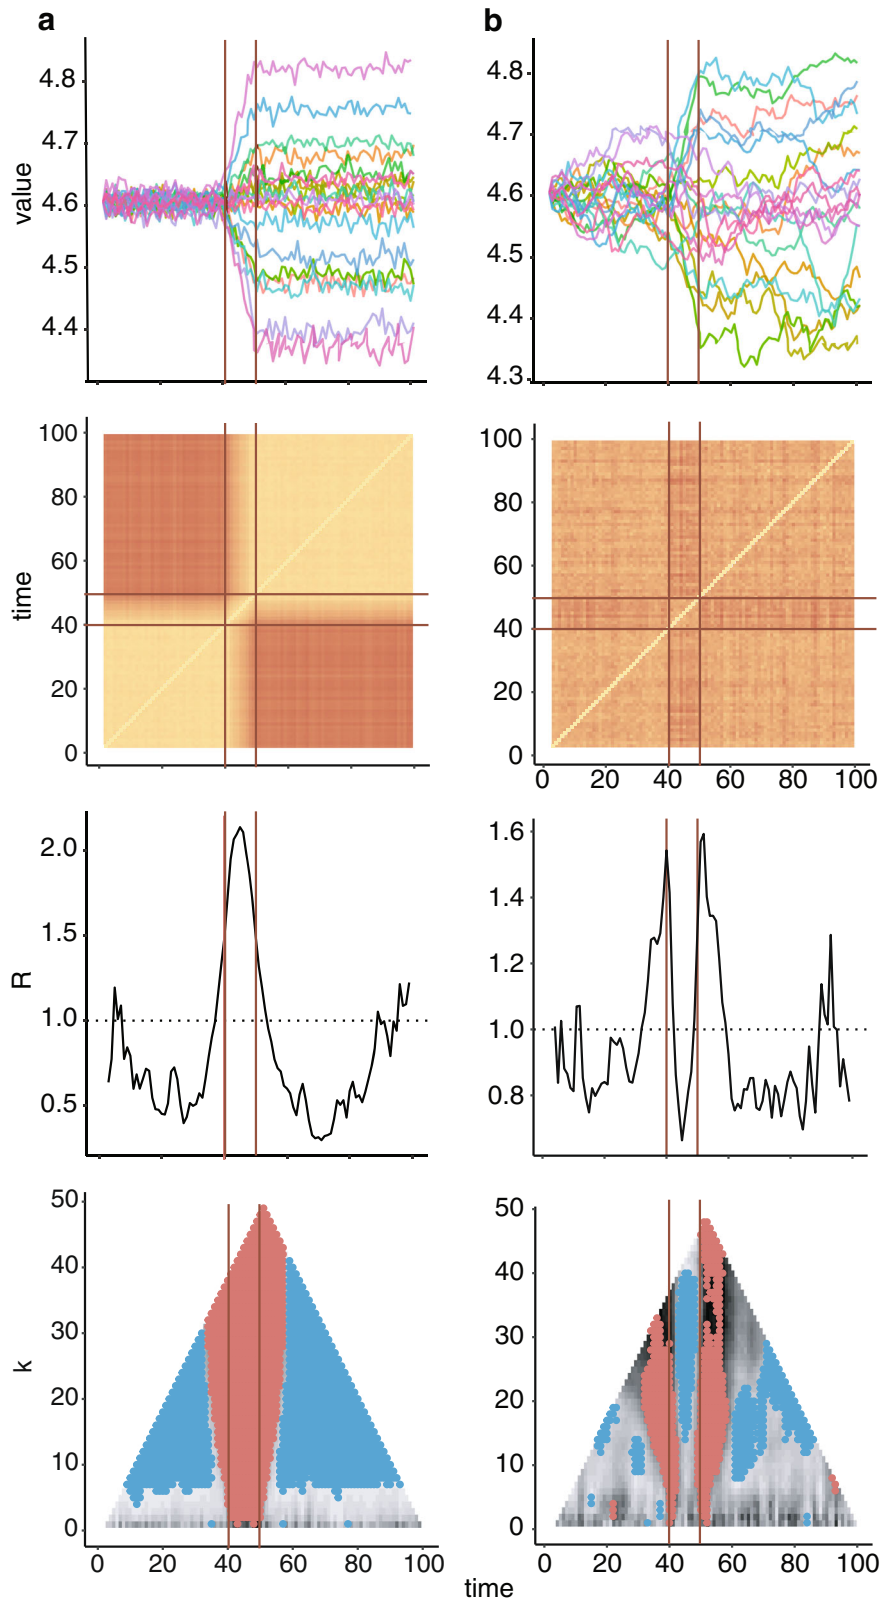
Any significance value is, of course, arbitrary, and we would also like a general picture of fluctuations in the rate of change regardless of whether or not they are statistically significant. To this end we propose an index, R_t , which captures the relative rate of change at a given time point, t . Assuming a set K of desired kernel half-widths, this index is constructed by first standardizing every F_t^k estimate by the average over all valid time points for its half-width, \bar{F}^k , and then averaging the standardized values over all $k \in K$ estimated for that time point to give a single value:

$$R_t = \frac{1}{|K|} \sum_{k \in K} \frac{F_t^k}{\bar{F}^k}$$

If $R_t > 1$, then the rate of change at a given time point is greater than the average rate of change in the entire series; if $R_t < 1$ then it is smaller. Note that statistical inference does not depend directly on the value of R_t . This means that R_t may sometimes have a relatively high value even when no revolution has been detected. This is particularly true at the start and finish of time series where the statistical power to detect revolutions diminishes.

Having identified a revolution, we would also like to know which variables contribute to it. One simple way to find out involves removing variables from the data set one at a time and re-running the analysis. Variables which, when removed, yield fewer statistically significant F_t^k in a given revolution contribute to it; those which yield more obscure it. Joint effects can be tested by removing combinations of variables.

Figure 2a shows the method in action on simulated data. We simulated twenty stationary series, each of which represents a measured variable, for 100 time points (Fig. 2a first row). By a “stationary” series, here we mean one whose mean is constant over time, that is, one that shows stasis before and after (though, obviously, not during) a revolution. Starting at time point 40, we introduced a revolution by allowing the variables to undergo a directional change for ten generations after which the series become stationary again. The rate discontinuity can be clearly seen in the distance matrix (Fig. 2a second row). To identify the revolution, we estimated F_t^k for all kernel half-widths, k , and time points, t , allowed by the data, in this case $1 \leq k \leq 49$, and calculated the rate index, R_t , for all time points. A sharp rate discontinuity is visible between time points 37 and 53, where $R_t > 1$ (Fig. 2a third row). Finally, we determine the statistical significance for each F_t^k estimate (Fig. 2a fourth row).



On the face of it, considering all k , there is strong evidence for a revolution spanning time points 37–57. But we have carried out 2401 significance tests over all k , of which 60 ($\alpha = 0.05/2$) are expected to be significant due to chance alone. In fact, we find that 766 are significant, strongly suggesting that the series

contains at least one real revolution. This final inferential step is equivalent to a Bonferroni correction. Since the test’s resolution decreases as k increases, the most accurate estimate is given by the smallest k at which the revolution appears: in this case, $k = 4$, where it spans time points 42–49—very close to the real values of

Fig. 2 Identifying revolutions using Foote novelty in simulated series. **a** Evolution of 20 simulated stationary time series with a revolution in the middle. **b** Evolution of 20 simulated undirected random walk time series with a revolution in the middle. In both sets of simulations, the standard deviation of series perturbations in non-revolutionary periods is set at $\sigma = 1$. During the revolutions, which start at time point 40, the size of the change in each time point is increased until time point 50, when the revolution ends. The amount by which each variable, i , changes during a revolution is drawn from a normal distribution. First row from top: Evolution of the time series. Second row: Distance matrices among time points: dark reds are increasingly dissimilar. Third row: The rate of change index, R_i , which is the sum of the F_t^k values for any time point i over all k , relative to the sum of the mean F_t^k values over all time points. Fourth row: Identifying revolutions by Foote novelty. Each cell represents the F_t^k estimate for a given half-width, k and time point; the color of the cell gives the relative F_t^k value, light gray being low and dark gray being high. Note that this color scale is only comparable within any given plot. Statistically significant ($\alpha = 0.05/2$) revolutionary periods are overlain in red; conservative periods are blue. In both cases, we identify a revolution in the correct region, but at larger half-widths, the resolution becomes coarser. Statistically significant time points which are not contiguous with the simulated revolution are false positives. Note that, for the random walk series, the undifferenced data are shown but the distance matrices, R_i and F_t^k values are all based on first differences. This means that only revolution boundaries are expected to have high F_t^k values.

40–50. A few statistically significant F_t^k values are seen well outside of the simulated revolution; these are false positives and we discuss their identification below.

Sensitivity and specificity. In the example shown in Fig. 2a, we simulated stationary variables with a revolution characterized by large changes in a variable's mean occurring over relatively long periods of time. To see whether our method works in other kinds of series, we applied it to many sets of simulated time series and then counted the revolutions detected. In these simulations, we varied: (1) the *persistence* of the series, ρ , (2) the magnitude of change in variable values during a revolution, that is, its *strength*, s , and (3) the *length* of the revolution, l and (4) the number of variables measured, n , giving 3300 combinations of parameter values, for which we simulated 20 replicates each, or 66,000 sets of series in all.

The persistence of a series measures the autocorrelation in an autoregressive order-one process. In these simulations, persistence varied between 0 (fully mean stationary) to 1 (random walk); the revolution length between 2 and 18 time units, and variable number between 10 and 190 (See Supplementary Information Materials and Methods for details). For each of these 66,000 sets, we estimated F_t^k for five kernel half-width, k , determined which were statistically significant, and then used these data sets to study the trade-off between sensitivity and specificity by estimating the rate of false positives (Type I errors) and false negatives (Type II errors).

We investigated the rate of false positives in series with no revolutions ($s = 0$). In this subset of the simulations, only three parameters vary: the persistence of the series, ρ , the number of variables, n , and the kernel half-width, k . Here the overall number of (false) revolutions detected should be equal to, or less than, $\alpha = 0.05$. For fully stationary series ($\rho = 0$), we found that this was so, however, as the series became more persistent the rate of false positives increased, so that in random walks ($\rho = 1$), revolutions were detected, on average, in 16% of the series (Fig. 3a). Thus, like many econometric tests, ours requires stationary series.

The risk of false-positives climbs above the set significance threshold at $\rho = 0.25$ (Fig. 3a). Persistent series can be made stationary by taking their first differences, $\bar{x}(t) - \bar{x}(t - 1)$, where \bar{x} is the mean at time steps t and $t - 1$ and, when we do so, we find that the rate of false positives is, once again, equal to or below the set significance threshold regardless of persistence (Fig. 3b). Figure 2b illustrates the effect of differencing on one set of random walk time series with a revolution introduced between time points 40–50. Now the revolution appears as spikes in F_t^k and R_i , marking its start and end and a set of significant F_t^k values between time points 32–41 and 46–57. Smaller k values (e.g., $k = 8$) give the most accurate estimates of the revolution's boundaries as time points 39–41 and 49–50 (Fig. 2b third and fourth rows).

We investigated the rate of false negatives in all series which contained revolutions ($s > 0$). When applied to levels, we found that, regardless of persistence, our test fails to detect about 22% of revolutions (Fig. 3b). Differencing reduces the power of the test considerably when applied to stationary series, but only slightly in highly persistent series (Fig. 3b). Focusing on the two extreme cases, stationary series ($\rho = 0$) and random walk series ($\rho = 1$) made stationary by differencing, we find that our method often fails to identify short and weak revolutions ($l \leq 6$, $s \leq 0.5$) in data sets based on few variables ($n \leq 10$), particularly when analyzed using very small half-widths ($k = 1$) (Fig. 3c).

In order to balance the risk of Type I and II errors when applying our test to real data, we recommend that investigators first estimate the overall persistence, $\bar{\rho}$, of the set of time series. If the series prove to be stationary or weakly persistent ($\bar{\rho} \leq 0.25$), then the test can be safely applied to the original data. But if the series are even moderately persistent, then it should be applied to the first differences.

The prevalence of revolutions. To illustrate our method, we applied it to several real data sets, all of which document changes in the historical record over the course of decades. Some of these use unbounded, continuous, data (Fig. 4); others, frequencies (Fig. 5). The first concerns a familiar subject: the spread and retreat of democracy across the globe in the course of the 20th century. In 1991, the political scientist Samuel Huntington identified three great global “waves” of democratization (Huntington, 1991). The first wave began around 1820; the second is associated with post-War War II de-colonization, and the third began in 1974 and is associated with the collapse of European and Latin American dictatorships, the fall of the Iron Curtain in 1989, and the spread of democracy in Africa. Huntington evidently based his argument on a simple count of “democracies” without either defining what he meant by the term or presenting any data. Here, using much better data, we ask whether our method can identify the second and third of his waves.

To do this we use the V-Dem data set. This data set, the work of many scholars, rates the degree to which the world's nation states were democratic over the course of the 20th century by means of a large number of ordinal variables that capture, in fine detail, the political structure of a given state in a given year (Coppedge, 2016). V-Dem provides indices where these variables have been aggregated to five higher-level quantitative variables that capture the degree to which a state exhibits: (i) freedom of expression; (ii) freedom of association; (iii) clean elections; (iv) an elected executive and, (v) universal suffrage (see Supplementary Information Materials and Methods for details). Fig. 4a (top) shows the yearly means of these variables averaged over the states extant in a given year (≤ 174). Consistent with previous V-Dem studies (Lindberg et al., 2014; Lüthmann, 2018), it shows that global democracy has increased over the course of the 20th century but that the rate at which it has done so has not been

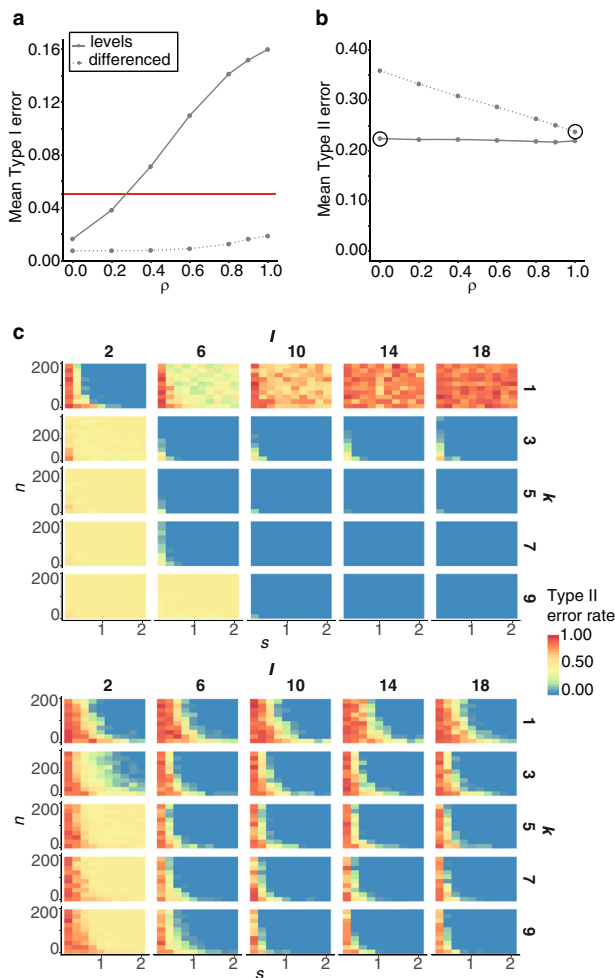


Fig. 3 Performance of revolution detection by Foote novelty based on simulated series. a Mean Type I error (false positive) rate as function of persistence, ρ , in series without revolutions ($s = 0$). When estimated on levels (solid gray line), the observed Type I error rate quickly increases above the significance level, $\alpha = 0.05$ (solid red line), but when differenced it does not (dotted gray line). The risk of false-positives climbs above the set significance threshold at $\rho = 0.25$. **b** Mean Type II error (false negative) rate as function of persistence in series with revolutions ($s > 0$). When estimated on levels (solid line), the observed Type II error rate is around 22%, but is higher when estimated on differences, decreasing as persistence increases. **c** A closer look at Type II error rates in stationary series ($\rho = 0$) (Top) and first-differenced random walk series ($\rho = 1$) (Bottom) as a function of the kernel half-width, k , number of variables in the simulation, n , the strength of the revolution, s , and its length, l . These plots are an expansion of the data in **b** marked with a circle. In both cases, our method tends to fail to identify short and weak revolution ($l \leq 6, s \leq 0.5$), in data sets based on few variables ($n \leq 10$), particularly when analyzed using very small half-widths ($k = 1$). Although the distribution of false negatives differs somewhat between the two sets of series, the overall mean false negative rates are very similar, 22% and 23%, respectively.

constant. We first estimated the persistence, $\bar{\rho}$, of the series and, finding that it was ≥ 0.25 , took the first difference (Supplementary Information Table 1). Our index, R , shows that the relative rate of change was elevated in the 1940s, early 1960s and between 1974–1999 (Fig. 4a). We then carried out 3192 significance tests over all k of which 208 were significant ($\alpha = 0.05/2$), many more than the 78 expected by chance alone, suggesting that the

series contains at least one real revolution. The years in which the rate of change is significantly higher than the background rate fall into four nearly contiguous groups: 1944–1949, 1962, 1975–1985, and 1989–1996 which we then identify as distinct “revolutions” (Supplementary Information Table 2).

Even when differenced, the entire series proved to be more persistent than desirable if we wish to avoid a high rate of Type I errors ($\bar{\rho} = 0.437$), but visual examination of the data suggested that, outside of the inferred revolutions, the series was close to stationary. To test this idea we estimated the persistence of periods before, between and after our inferred revolutions, and found that they were indeed acceptably non-persistent ($\bar{\rho} = 0.255$). We also took the second differences of the entire series, which made it overall stationary ($\bar{\rho} = -0.275$), and even so found revolutions in 1947–1948 and 1990–1992, albeit reduced in size. Thus, we are confident that the revolutions we identified are not due to the general persistence of the series.

These revolutions are very consistent with Huntington’s “waves”, if we allow that his “third wave” is composed of two distinct sub-waves (c.f., Kurzman, 1998; McFaul, 2002; Way, 2005). Interestingly, the 1962 revolution—by far the most weakly supported of the four—is an anti-democratic one caused by military coups in Indonesia, Pakistan, Greece, Nigeria, Turkey, and many Latin American countries. Huntington identified this phenomenon too and labeled it a “reverse wave” as have previous V-Dem studies (Mechkova et al., 2017). But we can add some detail to this picture. Analysis of the contributions of individual variables shows that, where the revolution of the late 1940s was due to changes in political structures, the 1977–1984 and 1989–1996 revolutions were due to an increase of personal liberty (Supplementary Information Table 3). Revolutions, unsurprisingly, differ in their natures and causes. Thus, our method can identify times of rapid political change of the sort that political scientists and historians have spotted using less formal methods.

We now turn to another familiar phenomenon: American pop music (Fig. 5a). Pop music is also said to undergo revolutionary change as new genres rise and fall, but unlike the spread of democracy there is little consensus as to when those revolutions occurred and what, exactly, changed in them (Frith, 1988; Tschmuck, 2006). We have previously studied the evolution of the US Billboard Hot 100, 1960–2010 (Mauch et al., 2015). In that study, we assayed 17,094 songs for 16 harmonic and timbral features and, using an earlier version of our method, claimed the existence of three revolutions: in the mid-1960s, early 1980s, and late 1980s–early 1990s. We re-analyzed these data using our improved testing procedure and, finding that the series is highly persistent, took the first differences. We find that $R_t > 1$ during 1967–1969, 1971, 1978, 1982–1983, 1986–1989, 1994–1995, 1998–2000, and 2005. We carried out 552 tests over all k of which 58 show a significantly elevated rate of change, more than the 14 expected by chance alone ($\alpha = 0.05/2$); these fall into three revolutions: 1968–1969, 1982–1983, 1986–1988. These are very close to the revolutions that we previously identified and that are due, respectively, to the rise of rock-related chords and timbres (aggressive percussion) in the 1960s, the revival of guitar-heavy rock and the arrival of drum-machine percussion in the early 1980s and, in the late 1980s, the rise of hip hop at the expense of rock and pop-related timbres (Supplementary Information Table 3). Note that since here we used differenced data, rather than levels, these are the boundaries of revolutions and not, as previously, their entire span. This accounts for the small discrepancy of dates between this analysis and the earlier one.

Besides these data sets we also applied our test to five others: the car models sold in the USA, 1950–2010 (Fig. 4b); a data set on the crimes committed in England and Wales 1900–2000 (Fig. 4c);

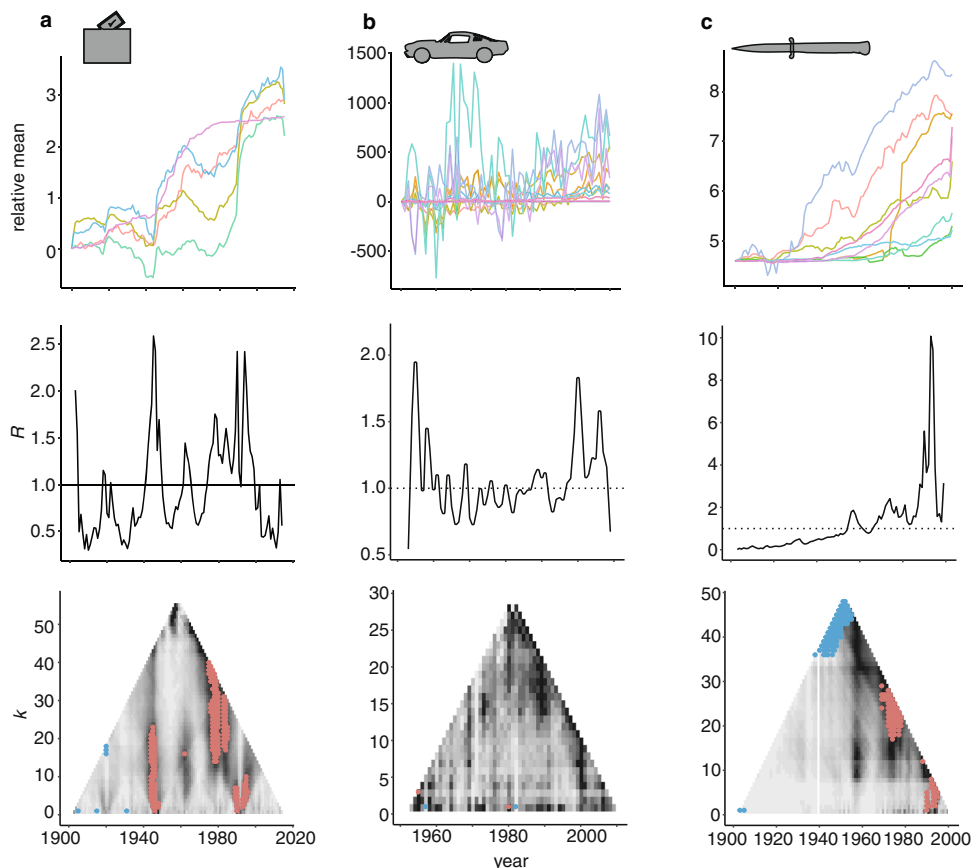


Fig. 4 Cultural Revolutions: unbounded variables. **a** global democracy; **b** car models; **c** crime rates per hundred thousand, UK. Top row of each series. Trends shown as values normalized to the first time point. Middle row. The rate of change index, R_t . Bottom row. Identifying revolutions by Foote novelty. Each cell represents the F_k^t estimate for a given half-width, k and time point; the color of the cell gives the relative F_k^t value, light gray being low and dark gray being high. Note that this color scale is only comparable within any given plot. Statistically significant ($\alpha = 0.05/2$) revolutionary periods are overlain in red; conservative periods in blue; when necessary Foote novelty tests were done on differences.

the common names given to newborn girls in the USA, 1945–2010 (Fig. 5b); the articles published in the *British Medical Journal*, 1960–2008 (Fig. 5c), and American, Irish, and English novels published between 1840 and 1890 (Fig. 5d) (See Supplementary Information Materials and Methods for details). Of these series, two: the girls names and English and Welsh crime rates, showed strong evidence of revolutions.

The girls names showed revolutions particularly in the years 1974–1975 and 1988–1991 (Fig. 5b; Supplementary Information Table 3). These dates mark when a set of names—Jessica, Ashley, Lauren, Amanda, and Amber among others—become swiftly and immensely fashionable and then, about 15 years later, passé and replaced by names such as Emma, Isabella, Olivia and Hannah (Supplementary Information Table 3, Supplementary Information Fig. 1). Of course, baby names change in frequency all the time (Lieberson, 2000); it is the fact that several of them rose and fell in tandem that makes their dynamics revolutionary.

For the crime data set (Fig. 4c), we carried out 2304 tests, of which 138 show a significantly elevated rate of change, many more than the 58 expected by chance alone ($\alpha = 0.05/2$). We detected two periods of revolutionary change: general increase in crime between 1965 and 1978 and, then, a general decrease from 1989–1995. The former is the increase of crime rates—and, for decades, their accelerating rate of increase—that occurred in Western democracies after 1960, part of what Francis Fukuyama called “The Great Disruption” (Fukuyama, 1991); the latter is the well known sharp decline in crime rates in the 1990s (Pinker,

2011). The rise was mostly due to an increase in criminal damage and robbery; the decline mostly due to a decline in burglary.

These examples show that our method can be applied to quite different data sets: some are count data (e.g., baby names), while others are continuous traits (e.g., measure of democracy); some aggregate many individual entities that exist only in a single time interval (e.g., pop songs), while others track the evolution of a collection of entities over time (e.g., the democratic qualities of nations): all it requires is that we can estimate a distance in feature-space between intervals in a time series. Using it we have convincingly identified revolutions—some well known, others not—in several data sets, but not in all of them. This is as expected. After all, revolutions are, by definition, rare.

Conclusion

We began this paper by defining a revolution as a period of time in which the multivariate rate of change is demonstrably higher than at other times. This is most likely to occur when several variables show simultaneous increases in the rate of change. Thus, our definition captures the classical idea of a revolution as a rapid, correlated, change in many properties of a system. The magnitude of change in a revolution—what we have called its strength—may be large or small in absolute terms: what matters is its size relative to the variance of change across the entire series. The period over which it occurs—what we have called its span—may be short or long.

A revolution cannot, however, span an entire time series. This is true even when all variables are changing constantly. To see this

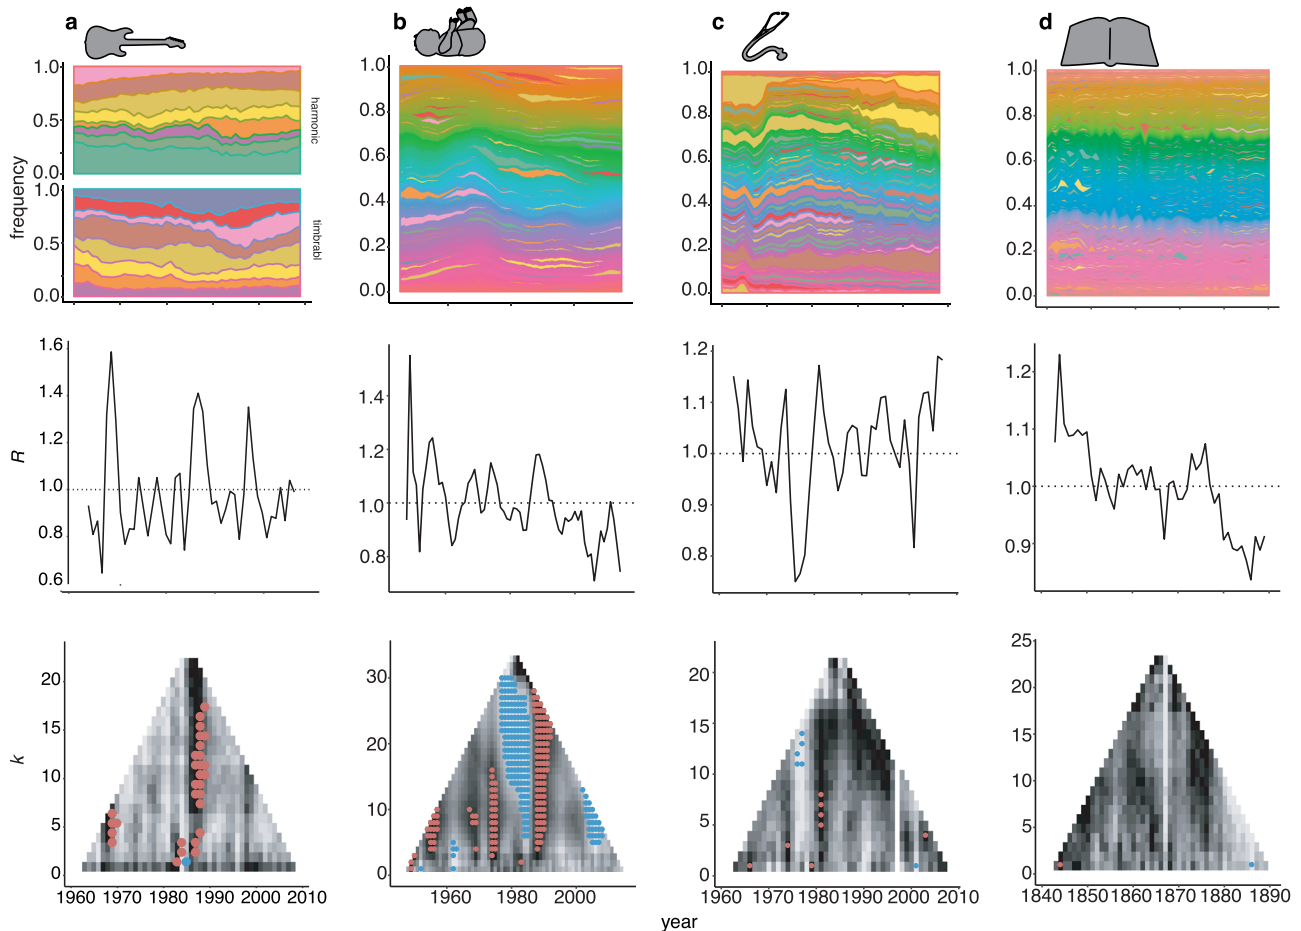


Fig. 5 Cultural revolutions: frequency variables. **a** pop music: Billboard Hot 100, USA; **b** newborn girls’ names, USA; **c** BMJ articles; **d** English, Irish and American novels. Top row of each series. Trends of frequencies shown as stacked area plots. Middle row. The rate of change index, R_t . Bottom row. Identifying revolutions by Foote novelty. Each cell represents the F_t^k estimate for a given half-width, k and time point; the color of the cell gives the relative F_t^k value, light gray being low and dark gray being high. Note that this color scale is only comparable within any given plot. Statistically significant ($\alpha = 0.05/2$) revolutionary periods are overlain in red; conservative periods in blue; when necessary Foote novelty tests were done on differences.

consider a collection of variables changing as directed random walks. Since each variable diverges from its original value linearly over time, its rate of change at any time, hence D over any interval, will be, within the limits of stochastic variation, constant as will F_t^k . Thus, viewed retrospectively, although there can be perpetually high rates of change, there are no perpetual revolutions. We can, however, find ourselves perpetually embroiled in revolution. When evolution is super-linear—we are thinking here of patterns such as that expressed in Moore’s law of the evolution of semi-conductor density (Moore, 1965)—the rate of change, D over any interval, and F_t^k , all increase monotonically. In such a series, a revolution will shift as the series grows so that it always defines the cutting edge. Thus, there is a sense in which permanent revolutions can, and probably do, exist.

We have focused on identifying revolutions simply because times of great change capture the imagination and are invariably the subject of scholarly debate. But significance levels are, of course, arbitrary, and the number of revolutions identified will change as they do. They may even be dispensed with. In their absence F_t^k , and its summary index, R_t , provides a simple way of measuring, and visualizing, local variation in rates of change. We note that evolutionary biologists commonly compare rates of evolution using measures such as the darwin and the haldane. Although both can be applied to any kind of time series data, both are univariate and generally estimated over an entire series

(Lambert et al., n.d.), and so are not well suited to estimating temporal variation in rates of multivariate evolution.

In all our data sets, all variables had non-zero values. However, it is possible to imagine revolutions in which some variables become irrelevant even as others arise. To give a concrete example, consider car design. Over fifty years of car evolution we detected much change, but no revolutions. Now, however, electric cars are upon us. Some of their features are much like those of their fossil-fueled ancestors (e.g., door number), but some (e.g., cylinder number, gear number) are not applicable, others can still be measured but are radically different (e.g., the relationship between maximum torque and RPM), while yet others are altogether new (e.g., power train battery capacity). Such changes in the salience of variables can be handled by our method and, if they have a sufficiently swift and strong effect on the multivariate distribution, will appear as a revolution. However, the revolution they will surely bring about seems to be of a different kind than any involving merely quantitative changes, however rapid, in mean horsepower or chassis length. The fundamental distinction is between revolutions that entail changes in the relationships among variables or, more formally, their variance-covariance structure, and those that do not. We think of the former as “structural” revolutions (c.f., Snodgrass, 1980) and the latter “non-structural”— note that they are subsets of the revolutions that our method detects, but leave the problem of telling them apart for future research.

Our method can be applied to identifying dramatic changes in any multivariate time series data of sufficient length and quality. In biology, it might be applied to the study of gene expression profiles, the evolution of gene frequencies or morphology (e.g., Hunt et al., 2015; Tu et al., 2005; Bergland et al., 2014). But the idea of revolution has its origin in historiography and so we have focused on political, social and cultural phenomena. As large data sets capturing their evolution become available (Michel, 2011; Hughes et al., 2012; Rodriguez Zivica et al., 2013; Perc, 2013; Klingenstein et al., 2014; Rule et al., 2015; Bearman, 2015), it will be increasingly possible to infer the quantitative patterns of history and so test general explanations for their causes (Kolodny et al., 2015).

Data and code availability

The code and data are available at: <https://github.com/Armand1/A-revolution-detector>.

Received: 16 March 2019; Accepted: 14 November 2019;

Published online: 07 January 2020

References

- Andersen T, Carstensen J, Hernández-García E, Duarte CM (2009) Ecological thresholds and regime shifts: approaches to identification. *Trends Ecol Evol* 24:49–57
- Atkinson QD, Meade A, Venditti C, Greenhill SJ, Pagel M (2008) Languages evolve in punctuated bursts. *Sci* 319:588–588
- Bearman P (2015) Big data and historical social science. *Big Data Soc* 2:1–5
- Bergland A, Behrman E, O'Brien K, Schmidt P, Petrov D (2014) Genomic evidence of rapid and stable adaptive oscillations over seasonal time scales in *Drosophila*. *PLoS Genet* 10:e1004775
- Bezanson A (1922) The early use of the term industrial revolution. *Quart J Econ* 36:343–349
- Bokma F (2008) Detection of punctuated equilibrium by bayesian estimation of speciation and extinction rates, ancestral character states and rates of anagenetic and cladogenetic evolution on a molecular phylogeny. *Evolution* 62:2718–2726
- Bowler P (1988) *The non-Darwinian revolution. reinterpreting a historical myth.* Johns Hopkins University Press
- Butterfield H (1950) *The origins of modern science.* G. Bell and Son
- Cohen HF (1994) *The scientific revolution: a historiographical inquiry.* University of Chicago Press
- Cohen IB (1986) *Revolution in science.* Harvard, Harvard University Press
- Coppedge M et al. (2016) V-Dem Country-Year Dataset v6. Varieties of Democracy (V-Dem) Project. V-Dem Institute, Gothenburg
- Currie TE, Mace R (2011) Mode and tempo in the evolution of socio-political organization: reconciling darwinian and spencerian evolutionary approaches in anthropology. *Phil Trans R Soc Biol Sci* 366:1108–1117
- Dediu D, Levinson SC (2012) Abstract profiles of structural stability point to universal tendencies, family-specific factors, and ancient connections between languages. *PLoS ONE* 7:0045198
- De Dreu CKW, van Dijk MA (2018) Climatic shocks associate with innovation in science and technology. *PLoS ONE* 13:1–16
- Duchen P et al. (2017) Inference of evolutionary jumps in large phylogenies using Lévy processes. *Syst Biol* 66:950–963
- Eldredge N, Gould SJ (1972) Punctuated equilibria: an alternative to phyletic gradualism. In Schopf, T. (ed) *Models in Paleobiology.* Freeman Cooper, New York
- Flink CM (2017) Rethinking punctuated equilibrium theory: a public administration approach to budgetary changes. *Policy Stud J* 45:101–120
- Foote J (2000) Automatic audio segmentation using a measure of audio novelty. In: *IEEE International Conference on Multimedia and Expo.* Institute of Electrical and Electronic Engineers. Vol. 1, pp 452–455
- Fowler L, Neaves TT, Terman JN, Cosby AG (2017) Cultural penetration and punctuated policy change: explaining the evolution of US energy. *Policy Rev Policy Res* 34:559–577
- Frith S (1988) *Music for pleasure.* Polity Press, Cambridge
- Fukuyama F (1991) *The great disruption: human nature and the reconstitution of social order.* Profile
- Gould SJ (2002) *The structure of evolutionary theory.* Harvard University Press
- Greenhill SJ et al. (2017) Evolutionary dynamics of language systems. *Proc Natl Acad Sci USA* 114:E8822–E8829
- Hall AR (1954) *The scientific revolution, 1500–1800: The formation of the modern scientific attitude.* Beacon Press
- Harmon LJ et al. (2010) Early bursts of body size and shape evolution are rare in comparative data. *Evolution* 64:2385–2396
- Heumakers A (2015) *De Esthetische Revolutie. Boom*
- Himmelfarb G (1996) *Darwin and the Darwinian revolution.* Ivan R. Dee
- Hodge M (2005) Against revolution and evolution. *J Hist Biol* 38:101–121
- Hughes JM, Foti NJ, Krakauer DC, Rockmore DN (2012) Quantitative patterns of stylistic influence in the evolution of literature. *Proc Natl Acad Sci USA* 109:7682–7686
- Hunt G (2010) Evolution in fossil lineages: paleontology and the origin of species. *Am Natural* 176:S61–S76
- Hunt G (2012) Measuring rates of phenotypic evolution and the inseparability of tempo and mode. *Paleobiology* 38:351–373
- Hunt G, Hopkins MJ, Lidgard S (2015) Simple versus complex models of trait evolution and stasis as a response to environmental change. *Proc Natl Acad Sci USA* 112:4885–4890
- Huntington S (1991) Democracy's third wave. *J Democracy* 2:12–34
- Kent M, Moyeed RA, Reid CL, Pakeman R, Weaver R (2006) Geostatistics, spatial rate of change analysis and boundary detection in plant ecology and biogeography. *Progress Phys Geogr* 30:201–231
- Klingenstein S, Hitchcock T, DeDeo S (2014) The civilizing process in London's Old Bailey. *Proc Natl Acad Sci USA* 111:9419–9424
- Kolodny O, Creanza N, Feldman MW (2015) Evolution in leaps: the punctuated accumulation and loss of cultural innovations. *Proc Natl Acad Sci USA* 112: E6762–E6769
- Kolodny O, Creanza N, Feldman MW (2016) Game-changing innovations: how culture can change the parameters of its own evolution and induce abrupt cultural shifts. *PLoS Comput Biol* 12:1–15
- Koyré A (ed) (1957) *From the closed world to the infinite universe.* Johns Hopkins Press
- Kuhn T (1957) *The Copernican revolution: planetary astronomy in the development of western thought.* Harvard University Press
- Kuhn T (1972) *The structure of scientific revolutions.* University of Chicago Press
- Kurzman C (1998) Waves of democratization. *Stud Comparat Int Dev* 33:42–64
- Lambert B et al. (n.d.) The pace of modern culture. *Nat Hum Behav* (in press)
- Landis MJ, Schraiber JG (2017) Pulsed evolution shaped modern vertebrate body sizes. *Proc Natl Acad Sci USA* 114:13224–13229
- Landis MJ, Schraiber JG, Liang M (2012) Phylogenetic analysis using Lévy processes: finding jumps in the evolution of continuous traits. *Syst Biol* 62:193–204
- Liebersohn S (2000) A matter of taste: how names, fashions and culture change. *Yale*
- Lindberg SI, Coppedge M, Gerring J, Teorell J (2014) A new way to measure democracy. *J Democracy* 25:159–169
- Lonne I (2017) Terrestrial slopes in northern high latitudes: a paradigm shift regarding sediment origin, composition, and dynamic evolution. *Geomorphology* 276:180–202
- Lowenstein J, Grantham JJ (2017) Residual renal function: a paradigm shift. *Kidney Int* 91:561–565
- Lührmann A et al. (2018) Democracy at dusk? v-dem annual report 2017. Technical Report, V-Dem Institute, University of Gothenburg
- Mauch M, MacCallum RM, Levy M, Leroi AM (2015) The evolution of popular music: USA 1960–2010. *R Soc Open Sci* 2:150081
- McBrearty S, Brooks AS (2000) The revolution that wasn't: a new interpretation of the origin of modern human behavior. *J Hum Evol* 39:453–563
- McFaul M (2002) The fourth wave of democracy and dictatorship: noncooperative transitions in the postcommunist world. *World Politics* 54:212–244
- Mechkova N, Lührmann A, Lindberg SI (2017) How much democratic backsliding? *J Democracy* 28:162–169
- Mellars P, Stringer C (1989) *The human revolution: behavioural and biological perspectives on the origins of modern humans.* Edinburgh University Press
- Michel JB et al. (2011) Quantitative analysis of culture using millions of digitized books. *Science* 331:176–182
- Moore GE (1965) Cramming more components onto integrated circuits. *Electronics* 38:114–117
- Omrani N, Mueller-Roebber B, Nikoloski Z (2015) Segmentation of biological multivariate time-series data. *Sci Rep* 5:8937
- Pagel M, Venditti C, Meade A (2006) Large punctuational contribution of speciation to evolutionary divergence at the molecular level. *Science* 314:119–121
- Pennel MW, Harmon LJ, Uyeda JC (2013) Is there room for punctuated equilibrium in macroevolution? *Trends Ecol Evol* 29:23–30
- Perc M (2013) Self-organization of progress across the century of physics. *Sci Rep* 3:1720
- Pinker S (2011) *The better angels of our nature: a history of violence and humanity.* Viking
- Preuss P, Puchstein R, Dette H (2015) Detection of multiple structural breaks in multivariate time series. *J Am Stat Assoc* 110:654–668
- Raoult D (2017) Double paradigm shift for the antibiotics' activity on viruses: Zika's lesson. *Proc Natl Acad Sci USA* 114:E1045

- RodriguezZivica PH, Shifres F, Cecchic GA (2013) Perceptual basis of evolving Western musical styles. *Proc Natl Acad Sci USA* 110:10034–10038
- De Ruiter R, Schalk J (2017) Explaining cross-national policy diffusion in national parliaments: a longitudinal case study of plenary debates in the Dutch parliament. *Acta Politica* 52:133–155
- Rule A, Cointet J-P, Bearman PS (2015) Lexical shifts, substantive changes, and continuity in state of the union discourse, 1790–2014. *Proc Natl Acad Sci* 112:10837–10844
- Ruse M (1979) *The Darwinian revolution: science red in tooth and claw*. The University of Chicago Press
- Schama S (1989) *Citizens: a chronicle of the French revolution*. Viking
- Sepkoski D (2012) Rereading the fossil record: the growth of paleobiology as an evolutionary discipline. University of Chicago Press
- Sepkoski D, Ruse M (eds) (2009) *The paleobiological revolution: essays on the growth of modern paleontology*. University of Chicago Press
- Seward JB (2017) Paradigm shift in medical data management: Big data and small data. *JACC: Cardiovasc Imaging* 10:1304–1306
- Shapin S (1996) *The scientific revolution*. University of Chicago Press
- Snodgrass A (1980) *Archaic Greece: the age of experiment*. University of California Press
- Spencer C (1990) On the tempo and mode of state formation: neoevolutionism reconsidered. *J Anthropol Archaeol* 9:1–30
- Tschmuck P (2006) *Creativity and innovation in the music industry*. Springer, Dordrecht
- Tu BP, Kudlicki A, Rowicka M, McKnight SL (2005) Logic of the yeast metabolic cycle: temporal compartmentalization of cellular processes. *Science* 310:1152–1158
- Turchin P et al. (2018) Quantitative historical analysis uncovers a single dimension of complexity that structures global variation in human social organization. *Proc Natl Acad Sci* 115:E144–E151
- Uyeda JC, Hansen TF, Arnold SJ, Pienaar J (2011) The million-year wait for macroevolutionary bursts. *Proc Natl Acad Sci* 108:15908–15913
- Valverde S, Solé RV (2015) Punctuated equilibrium in the large-scale evolution of programming languages. *J R Soc Interface* 12:20150249
- Venditti C, Pagel M (2010) Speciation as an active force in promoting genetic evolution. *Trends Ecol Evol* 25:14–20
- Way L (2005) Authoritarian state building and the sources of regime competitiveness in the fourth wave: the cases of Belarus, Moldova, Russia, and Ukraine. *World Politics* 57:231–261

Acknowledgements

We thank Simon deDeo, Tim Hitchcock, and Matthew Jockers for supplying data. Arnold Heumakers, Tim Hitchcock, James Rosindell, and Ted Underwood for comments on the paper. B.L. was supported by EPSRC grant code: EP/F500394/1. V-Dem data collection was supported by European Research Council, Grant 724191; Riksbankens

Jubileumsfond, Grant M13-0559:1; Knut and Alice Wallenberg Foundation to Wallenberg Academy Fellow, Grant 2013.0166; as well as by internal grants from the Vice-Chancellor's office, the Dean of the College of Social Sciences, and the Department of Political Science at University of Gothenburg and the Marianne and Marcus Wallenbergs Foundation Grant 2017.0049.

Author contributions

A.M.L., B.L., and M.M. designed the study, contributed new reagents/analytic tools, collected data, carried out analysis, and wrote the paper. M.P. contributed new reagents/analytic tools; P.L., S.L., and S.A. contributed data. All authors gave final approval for publication.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1057/s41599-019-0371-1>.

Correspondence and requests for materials should be addressed to A.M.L.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020