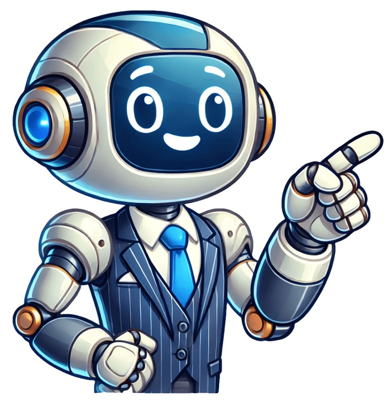

 Author manuscript; available in PMC: 2023 Oct 13. Introduced more than a half-century ago, Granger causality has become a popular tool for analyzing time series data in many application domains, from economics and finance to genomics and neuroscience. Despite this popularity, the validity of this framework for inferring causal relationships among time series has remained the topic of continuous debate. Moreover, while the original definition was general, limitations in computational tools have constrained the applications of Granger causality to primarily simple bivariate vector autoregressive processes. Starting with a review of early developments and debates, this article discusses recent advances that address various shortcomings of the earlier approaches, from models for high-dimensional time series to more recent developments that account for nonlinear and non-Gaussian observations and allow for subsampled and mixed-frequency time series. Keywords: multivariate time series, vector autoregressive model, graphical models, penalized estimation, deep neural networks, mixed-frequency time series There is a range of applications where the interest is in understanding interactions between a set of time series, including in neuroscience, genomics, econometrics, climate science, and social media analysis. For example, in neuroscience, one may seek to understand whether activity in one brain region correlates with later activity in another region, or to decipher instantaneous correlations between regions—both notions of functional connectivity. In genomics, there is an analogous study of gene regulatory networks. In econometrics, one may be interested in how various macroeconomic indicators predict one another. We also have unprecedented levels of data on people’s actions—including social media posts, purchase histories, and political voting records—and want to understand the dependencies between the actions of these individuals. Modern recording modalities and the ability to store and process large amounts of data have escalated the scale at which we seek to do such analyses. In many cases, one may seek notions of causal interactions among the time series but be limited to drawing inferences from observational data without opportunities for experimentation and without known mechanistic models for the observed phenomena. In such cases, Granger (1969) put forth a framework leveraging the temporal ordering inherent to time series in hopes of drawing causal statements restricted to the past causing the future. The framework, in reality, assesses whether one series is predictive of another: A series xi is deemed not to be “causal” of another series xj if leveraging the history of series xi does not reduce the variance of the prediction of series xj. In this review, we distinguish this definition from other standard definitions of causality by referring to it as Granger causality. Although there is a long history of debate about the validity of the Granger causality framework for causal analyses—and justly so—in this review we take the stance that analyzing interactions in time series defined by association has its utility. Granger causality has traditionally relied on assuming a linear vector autoregressive (VAR) model (Lütkepohl 2005) and considering tests on the VAR coefficients in the bivariate setting. However, in real-world systems involving many time series, considering the relationship between just a pair of series can lead to confounded inferences (e.g., Lütkepohl 1982). Network Granger causality aims to adjust for possible confounders or jointly consider multiple series (Eichler 2007, Basu et al. 2015). There are other important limitations of the linear VAR model underlying standard Granger causal analysis that have precluded its broad utility. Some limiting assumptions include assuming (a) real-valued time series with (b) linear dynamics dependent on (c) a known number of past lagged observations, with (d) observations available at a fixed, discrete sampling rate that matches the time scale of the causal structure of interest. In contrast, modern time series are often messy in ways that break a number of these assumptions, including through nonlinear dynamics and irregular sampling. Recent advances have pushed the envelope on where Granger causality can be applied by loosening these restrictions in a variety of ways. We review some of these advances and set the stage for further developments. In Section 2 we review the history of Granger causality, starting with the original definition and assumptions in Section 2.1 and early approaches for testing in Section 2.2. We then turn to network Granger causality and the issues of lag selection and nonstationary VAR models in Section 3. Finally, in Section 4 we review recent advances that move beyond the standard linear VAR model and consider discrete-valued series (Section 4.1), nonlinear dynamics and interactions (Section 4.2), and series observed at different sampling rates (Section 4.3). In his seminal paper, Granger (1969) proposed a notion of causality based on how well past values of a time series yt could predict future values of another series xt. Let $\mathcal{H}$