Industry-Driven Visual Analytics for Understanding Financial Timeseries Models

David Jonker
Uncharted Software
Toronto, Canada
0000-0002-8454-5096

Richard Brath
Uncharted Software
Toronto, Canada
0000-0001-6006-2092

Scott Langevin
Uncharted Software
Victoria, Canada
0000-0002-3309-8935

***Abstract*—Timeseries models are used extensively in financial services, for example, to quantify risk and predict economics. However, analysts also need to comprehend the structure and behavior of these models to better understand and explain results. We present a methodology, derived from extensive industry experience, to aid explanation through integrated interactive visualizations that reveal model structure and behavior of constituent timeseries factors, thereby increasing understanding of the model, the domain and the sensitivities. Expert feedback indicates alignment with mental models.**

Keywords—timeseries model, timeseries visualization, factor model.

I. Beyond Prediction: Understanding Timeseries models

Modeling timeseries is a common task in financial services broadly applicable for diverse analytical tasks. For example: financial ratios are used for valuation; regression models may be used for pricing; risk models may be used for estimating credit quality and market exposures; attribution models for estimating which factors contribute to returns; and nowcasting models for predicting economic measures before official release. Over time these models have become automated (e.g. automatic selection of features and choice of model and parameters) and become more complex (e.g. large number of time series features).

In our experience a key challenge in model deployment and usage is the communication and understanding of these models by downstream consumers. Users are reluctant to use a model that they do not understand. If they were not directly involved in the design of the model, then explanation is required. Such an explanation often takes the form of a document outlining the technical approach and the equations. However, this approach 1) assumes analysts are familiar with statistical formulae; 2) requires analysts to cross-reference between model results and documentation rather than direct observation into the model *in situ*; 3) requires the reader to map abstracted examples described in the document to specific concrete scenarios and data in the interface.

Furthermore, the understanding of a model goes beyond confidence in the forecast variables. Using a model only for an immediate predictive point value is highly useful for answering *what* (e.g. what is my exposure, what is the value of GDP going to be, what is the probability of default). Beyond the *what*, there is value in conceptualization, validation and internalization of the model. This aids understanding *why* and *what-if* (e.g. why did this outcome occur, what-if this scenario occurs).

Visualizing causality, visual inference and predictions is a top unsolved visualization problem [1]. What is needed is increased transparency in a form that facilitates ease of deep understanding of models (e.g. the Explainable AI challenge articulated by [2]. This challenge includes aspects such as model structure, behavior, quality, and change over time in a holistic approach to facilitate rapid internalization of the model. Our primary contribution is a visual analytics approach to aid this deep understanding of complex timeseries models.

II. Background on Timeseries Models and Visual Analytics

We have implemented many visual analytic systems utilizing timeseries models. Visual analytics is “the science of analytical reasoning facilitated by interactive visual interfaces” [3]. There are many different aspects of a model that are relevant to visualize, including the model output, the input data, the model structure, and changes in the model structure over time. There are also various aspects of the model that a user might want to manipulate, such as what-if modification of the input values to create scenarios; defining output targets and solving for inputs; adjustment to model structure and inputs; and so on. Harnessed together, these can create a deep understanding of a model leading to improved decision making [4].

A. Model Structure

In most models there is an intrinsic structure between the inputs and the output. This may be explicit, such as a hierarchy in an attribution model or a directed acyclic graph in a causal model or neural network. This structure may be implicit: for example, a flat multi-regression model may be better explained by grouping more strongly correlated input variables (thereby forming a hierarchical grouping).

Viewing the structure of the model can be insightful to understand how the inputs are assembled, their relationships, and their relative weights. Prior data visualizations have focused on depicting the structure of causal models as acyclic directed graphs (i.e. node-link diagrams), for example in biology [5]; neuroscience [6]; network anomalies [7]; and visualization [8].

In visualizing graphs as node-link structures one typically uses circles for nodes, lines or directed arrows for links, and short textual labels for annotating nodes or links. This approach is used in popular graph visualization tools (e.g. D3.js, Gephi, Cytoscape, YeD, etc). However, this paradigm is limited in the amount of information each model element can convey, and relies on cumbersome drill-downs to see more detailed sets of attributes associated with each element. For example, a node typically only visualizes three data attributes using color, size and a textual label. It is difficult to show a long timeseries – [9] use animation, however a transient display relies on visual memory making it cognitively difficult to compare across periods [10]. The approach does not provide affordances for manipulation of input/output values. Finally, these tools do not provide a good means for embedding commentary directly into the visual representation, thereby requiring cross-referencing to other blocks of prose.

B. Model Behavior

Timeseries analysis is highly prevalent in financial services. A capital markets expert at one of the largest firms providing data and analytical software estimates that as much as 90% of analysis involves timeseries. In the simplest use cases, analysts are interested in understanding the trends and events in timeseries, which can be explicitly depicted as timeseries charts. The behavior of the timeseries is readily visible as anomalies, trends, changes, gaps, reversals and other phenomena are readily perceived.

Explicit timeseries depictions allow users to understand the past (*what happened*) and future prediction (*what will happen*) e.g. model prediction, upcoming events, trend). However we have found that users also want to uncover *why* (e.g. why a particular anomaly, gap, reversal, etc., occurred). This enables them to assess, generalize, predict and act on similar patterns in the future. This requires interactive visualization techniques to explore underlying causes, for example, in stock market data this may include timeseries of indexes, peers, economics, fundamentals, earnings, news, social media, business activity and so on [11].

A simple approach is to visualize only model inputs and outputs, [12] bypassing the representation of model structure: however for larger models with hundreds or thousands of inputs this approach is not scalable and limits model understanding. Others have focused on explicitly visualizing internal or hidden nodes e.g. [13].

In the case of timeseries models, this analysis involves exploring underlying timeseries features in the model to understand which features are contributing to predicted model changes [14]. Understanding the relations and contributions across the timeseries features aids in assessing potential causes and facilitates understanding why a change occurred. This provides transparency and insight into the model, and potential increased understanding of a domain of interest.

Without this transparency, the analyst is asked to place wholesale trust in the totality of the models. In reality to many analysts the *why* may be more important than the accuracy of a predicted outcome. The goal is to understand how things are related, and it's that understanding that enables the analyst to act.

C. Model Sensitivity

Neither graphs nor timeseries express aspects such as amplification, dampening, prevention, or transience [15]. Causality can be deterministic, or probabilistic with varying levels of likelihood. Aspects such as uncertainty, soft evidence, and non-linear behavior also impact model behavior and are important to convey [16]. What-if scenarios and stress tests can be used to assess a model's response to different conditions.

D. Model Quality and Changes

Quantitative measures of model quality, such as forecast error, mean absolute error, and root mean squared error can also provide value by validating or comparing models using known ground truth. Models also change over time: some features become more important, some irrelevant, relationships shift, some new features are added, and so on. Comparing models quantitatively helps rank them against known ground truth. However, it is important for analysts to qualitatively explore, contrast, and understand varied model characteristics.

E. Annotation, Commentary and Narrative

In addition to explicit representations of the model, the addition of prose to supplement all the above is important, particularly for novice users and for casual consumers who mainly interact with the model infrequently. Explanation is becoming an important feature in machine learning and timeseries models, Narrative explanations have a history in visualization wherein the explanations are tightly coupled with visualization using a number of different techniques as first discussed in [17], and are now a broad area of research [18]. Note that this text can be metadata (such as references to data sources, calculations and so on), user authored notes, on-the-fly collaborative dialogue, semi-automated or fully automated natural language generation (NLG) to highlight or explain insights, or annotate resultant models.

III. Visual Analytics of Timeseries Models

Based on the above discussion, it is a significant challenge to capture these many aspects of model explanation. Working with expert users, we have prioritized the following goals for visual analytic timeseries models: 1) explicitly represent the underlying timeseries behavior; 2) represent the model structure;

3) allow for model manipulation; and 4) facilitate annotation and commentary.

A. Interactive Financial Ratio Decomposition

In the process of exploring preliminary concepts, we created a rapid prototype of financial ratios in Microsoft Excel. Model structure was explicitly represented left to right, with the timeseries data for each element depicted in the vertical dimension (Figure 1). The analyst can visually trace which variables have a similar correlated trend to the output variable. In this example, decreasing ROE is impacted by ratios such as financial leverage, book value and liabilities.

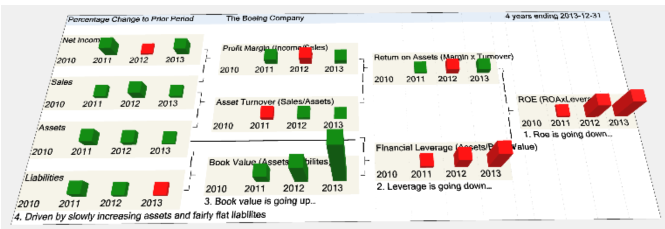

Fig. 1. Quick prototype showing decomposition of financial ratios.

The use of the model structure and the timeseries data together in this prototype was promising, but it was challenging to perceive and correlate behaviors over time with discrete bars and few data points. The colored 3D bars limited the ability to perceive up/down trends, and the textual annotations were difficult to read in perspective or obscured by 3D bars.

B. Full Scale Industry Application: Nowcasting

Incorporating lessons learned from several such prototypes and domains, we have since designed and implemented more advanced and comprehensive solutions for a variety of needs such as *nowcasting*. These solutions are scalable to larger datasets and more complex models, tailored for interactive, collaborative, multi-user, always-on, web-based environments.

1) Nowcasting Introduction

Nowcasting has become an important technique in capital markets. Nowcasting seeks to predict near-term financially-significant values such as GDP, company sales or merger activity. Accurate prediction of these values in advance of their official release provides opportunity to capture a profit ahead of the price movement of stocks upon their release. The Federal Reserve Bank of Atlanta publishes a United States GDP nowcasting model [19], with regular updates on their website www.frbatlanta.org/cqer/research/gdpnow.aspx. Various hedge funds use alternative datasets to nowcast estimates of earnings and other indicators before official filings, such as using satellite imagery of night illumination as a factor to estimate GDP (e.g. go.spaceknow.com/africa-lights-index/), social media to estimate company sentiment compared to stock prices, or tracking movement of corporate jets for potential insights into sales and mergers (e.g. www.quandl.com/).

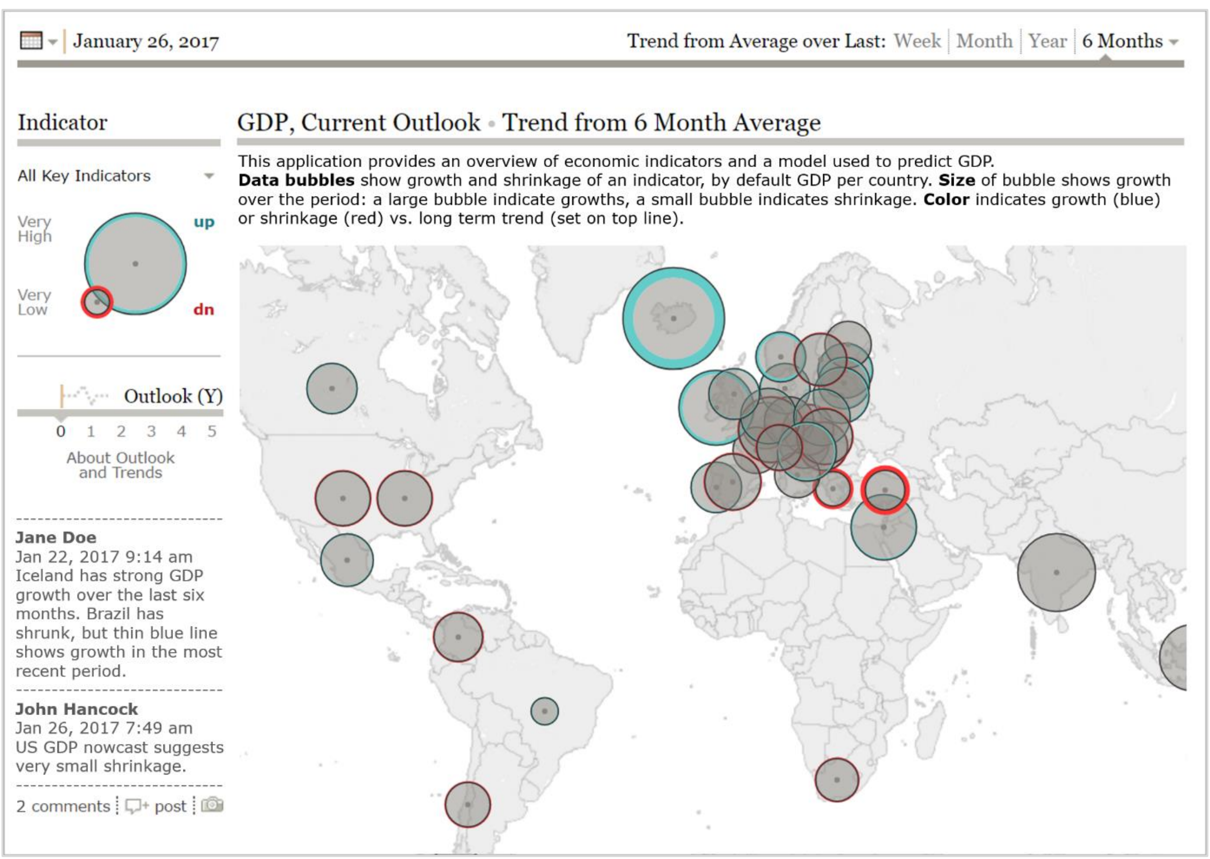

Fig. 2. Nowcast overview with zoomable map, legend, introductory text and analyst commentary.

2) GDP Nowcasting Visualization

Our GDP nowcast visualization example shown in Figure 2 starts with a global view. This view allows for comparison across countries, each country represented by a bubble, with key metrics of size (growth) and change (to prior period, indicated by color, red indicating shrinking GDP, green indicating increasing GDP). Note that the application includes a various configurations (such as time periods), a legend (top left), instructional commentary (top) and analyst authored commentary (bottom left side panel).

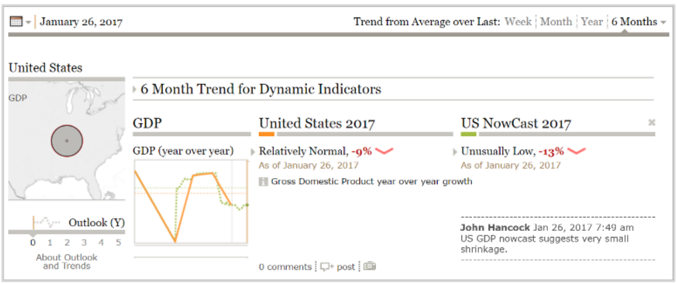

Fig. 3. Nowcast focused on USA.

The analyst can pick a particular country, in this example, United States (Fig. 3). The map is reduced (top left) and the analyst can view a timeseries chart with both nowcast (green dotted line) and official values (solid amber line). In the timeseries chart, the yellow line (actual GDP growth) starts high on the left of the chart, significantly drops to a low value in middle, then rebounds. The green line (the nowcast value) captures most of this rebound almost immediately after the low, thereby providing significant advance indication that the drop was a one-time anomaly.

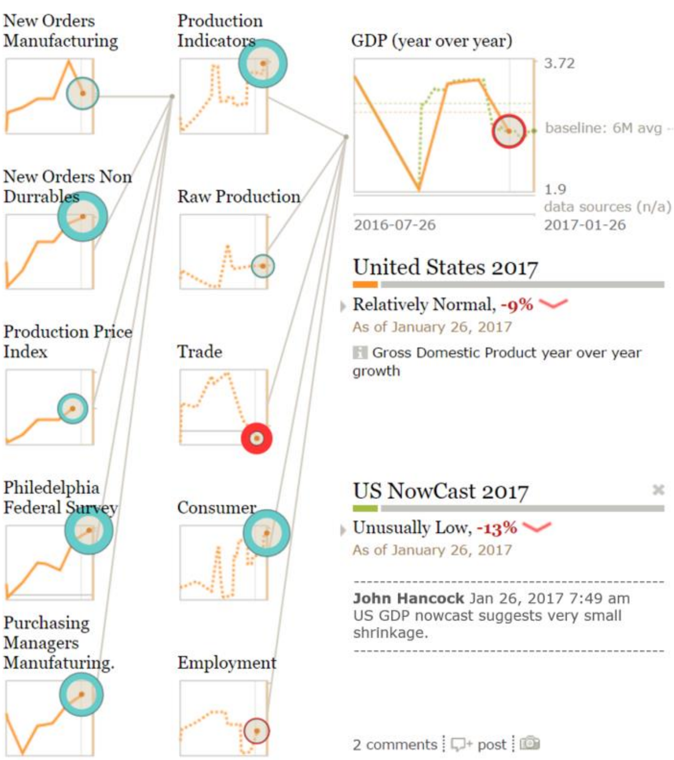

Fig. 4. Nowcast showing output (right), factors (center), and inputs (left).

The analyst can then drill-down to show the detailed structure and component timeseries which comprise the model. The nowcast is shown top right, with connections to the underlying factors (center column) and inputs (left column); as shown in Figure 4. Note that underlying inputs, such as Philadelphia Fed Survey or Production Price Index are solid lines, whereas the factors Production, Trade, Consumer are synthetic modeled values and are differentiated by dotted lines. The analyst can trace from the nowcast, back to the factors, and then back to the factor inputs. In this example, there is a recent drop in GDP, and a corresponding recent drop in trade.

Note that there can be many inputs, far more than fit on the screen at once. The factors (center column) provide a view on intermediate aggregates. Any one of these can be clicked to further drill-down to see the underlying input variables related to that factor. In Figure 5, the factor “Production Indicators” has been selected and promoted to the large visual in the upper right corner with the underlying timeseries on the left.

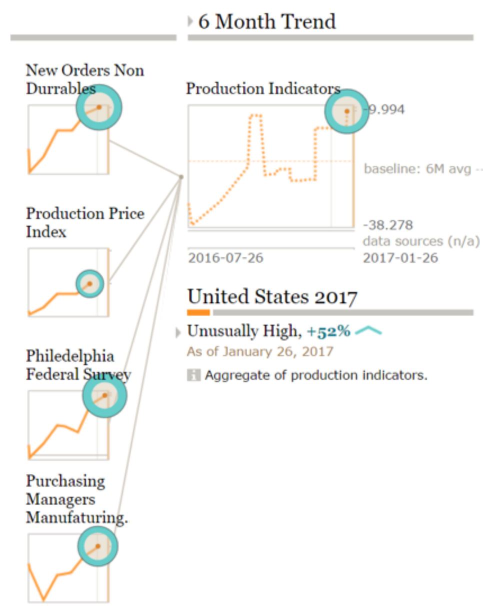

Fig. 5. Focus on a particular factor.

The analyst can also create *what-if* scenarios. When selected, the data points for the input variables can be dragged up/down and the model recomputed. A number of variables can be changed to create hypothetical scenarios to assess the potential impact of changes in the world according to the model (e.g. explore the potential impact to GDP of trade impact due to tariffs; or the impact to GDP of low unemployment and full factory utilization). The analyst benefits from seeing the change in the model results; can compare that to their real-world knowledge to potentially question the model; and can uncover new potential scenarios and outcomes previously unconsidered.

3) Discussion and Expert Evaluation

We have more than twenty years visual analytics experience working with expert users in the financial services domain, including data providers, banks, mutual funds and regulators.

These analysts include model builders, portfolio managers, researchers, traders and economists. These experts provide feedback through the model and visualization development lifecycle ranging from specific requests for functionality to supporting observations. By engaging experts, we have made a number of enhancements to better facilitate the goal of understanding. For example, detailed tick labels are not required on every x and y axis. Removing them reduces clutter, allows more thumbnail timeseries to fit into the display, and still conveys behavior. Through this ongoing engagement with experts, we have been able to extend the initial application to uses in different domains including economics, commodities and social data. Figure 6 shows an example of the application applied to economics data, with summary metric on the right (consumption), economic factors (center, such as CPI, unemployment, exports, imports, etc), and underlying input timeseries far left (e.g. Manufacturing PMI).

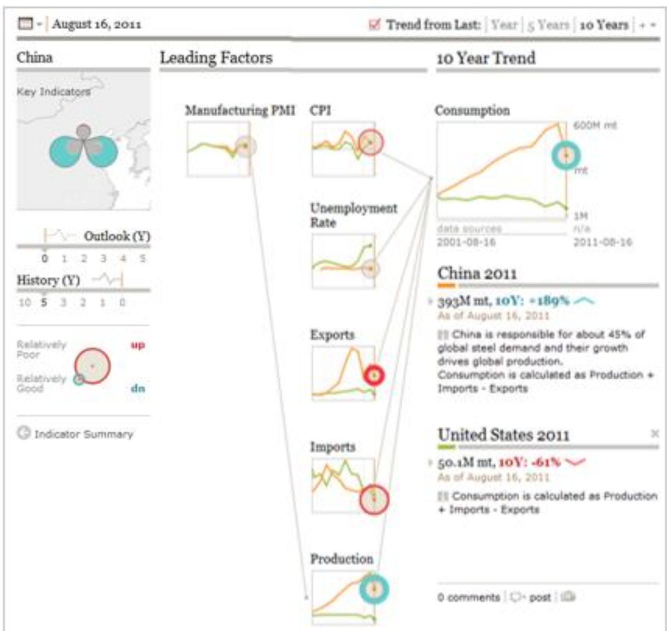

Fig. 6. Similar version of the application applied to economics data.

Explicit visual representation of the timeseries allows analysts to leverage their real-world knowledge. For example, a shift in a timeseries that corresponds to a real-world event (e.g. politics, news) helps the analyst build confidence that the visible model and the data aligns with their mental model. In user sessions, we have observed analysts visually highlight a timeseries node and state confirmatory facts, such as “Yes, China production skyrocketed through that period.”

Explicit interaction with model inputs or model structure allow analysts to explore characteristics of the model, for example, sensitivity to a particular variable or find conditions such as a model step or discontinuity. This is difficult to infer in any other way and again can be confirmed by prior knowledge. Analyst observations reveal their ability to re-create specific what-if scenarios to assess the comprehensiveness of the model, for example: “This condition is unlikely but did happen during an unusual weather event. Some market participants had an enormous liability.”

Embedded model commentary is user-driven and largely focused on metadata documentation (i.e. explaining the variable or a calculation). Additional user-generated commentary typically are insights of the model builder acting as an author and sharing those insights with downstream consumers, potentially constructed in a narrative sequence. For larger models with many data points, many variable and many potential interactions, these insights help guide analyst attention and explain the key points.

An aspect in our approach that we did not address is the visual depiction of relative influence and change in influence of inputs over time. Analysts want to compare these influences with their prior knowledge and to uncover new emerging relations or strengthening relations. One analyst commented “When I was in the trading pit, we’d all go quiet to wait for the report of one number, like CPI Energy for example. But then after a couple of months it would be a different number that we’d wait for, like reserves.”

One aspect commented on by analysts was the need to represent quality metrics such as forecast error and mean forecast error. Analysts also required automatic grouping of inputs in large scale automated models with hundreds to thousands of features. We are currently exploring deeper integration of quality metrics into representations; as well as automated hierarchical grouping of features (which creates more intermediate synthetic timeseries useful for detailed visual inspection), by leveraging our work in large scale graph decomposition and visualization [20].

These analyst insights linking understanding of model structure and behavior to real-world knowledge confirms our belief that a key user goal is the internalization of the model to facilitate and enrich their understanding of the model and the world.

IV. Conclusion

Our approach to model explanation to date is a combination of 1) explicit representation of model structure; 2) explicit representation of timeseries behavior; 3) direct interaction with input data for scenario building; and 4) integrated commentary. The alignment between this explanation and the analyst's mental model facilitates learning and internalization.

Future challenges include introducing model quality, feature weights and changes over time into these explanations while retaining the ease of use and ease of exploration to achieve better deep understanding.

References

[1] Chaomei Chen. Top 10 Unsolved Information Visualization Problems. *IEEE Computer Graphics And Applications*, 25(4):12–16, 2005.

[2] David Gunning. Explainable Artificial Intelligence (XAI). https://www.darpa.mil/attachments/DARPA-BAA-16-53.pdf. accessed: 10/16/2018.

[3] Kristin A Cook and James J Thomas. *Illuminating the path: The research and development agenda for visual analytics*. Technical report, Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2005.

[4] Mica R Endsley and Erik S Connors. Situation awareness: State of the art. In *Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century*, 2008 IEEE, pages 1–4. IEEE, 2008.

[5] Rainer Opgen-Rhein and Korbinian Strimmer. From Correlation To Causation Networks: A Simple Approximate Learning Algorithm And Its Application To High-Dimensional Plant Gene Expression Data. *BMC Systems Biology*, 1(1):37, 2007.

[6] Anil K Seth. A Matlab Toolbox For Granger Causal Connectivity Analysis. *Journal Of Neuroscience Methods*, 186(2):262–273, 2010.

[7] Hao Zhang, Maoyuan Sun, Danfeng Daphne Yao, and Chris North. Visualizing Traffic Causality For Analyzing Network Anomalies. In *Proceedings of the 2015 ACM International Workshop on International Workshop on Security and Privacy Analytics*, pages 37–42. ACM, 2015.

[8] Nivedita Kadaba, Pourang Irani, and Jason Leboe. Visualizing Causal Semantics Using Animations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1254–1261, 2007.

[9] Niklas Elmqvist and Philippas Tsigas. Growing Squares: Animated Visualization Of Causal Relations. In *Proceedings Of The 2003 ACM Symposium On Software Visualization*, pages 17–ff. ACM, 2003.

[10] Edward Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 1983.

[11] Eugene Sorenson and Richard Brath. Financial Visualization Case Study: Correlating Financial Timeseries And Discrete Events To Support Investment Decisions. In *Information Visualisation (IV), 2013 17th International Conference*, pages 232–238. IEEE, 2013.

[12] Josua Krause, Adam Perer, and Enrico Bertini. Using Visual Analytics To Interpret Predictive Machine Learning Models. arXiv preprint arXiv:1606.05685, 2016.

[13] Hendrik Strobelt, Sebastian Gehrmann, Bernd Huber, Hanspeter Pfister, Alexander M Rush, et al. Visual Analysis Of Hidden State Dynamics In Recurrent Neural Networks. *CoRR*, abs/1606.07461, 2016.

[14] David Jonker. Linked Visible Behaviors: A System For Exploring Causal Influence. *AHFE Cross-Cultural-Decision-Making (CCDM) conference*, 2012.

[15] Miao Yao. Visualizing Causality In Context Using Animation. Simon Fraser University, 2007.

[16] William Wright and Thomas Kapler. Challenges In Visualizing Complex Causality Characteristics. In *IEEE Pacific Visualization*. 2018. Proceedings. IEEE, 2018.

[17] Edward Segel and Jeffrey Heer. Narrative Visualization: Telling Stories With Data. *IEEE Transactions On Visualization And Computer Graphics*, 16(6):1139–1148, 2010.

[18] Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Carpendale. *Data Driven Storytelling*. A K Peters and CRC Press., 2018.

[19] Patrick Higgins, GDPNow: A Model for GDP “Nowcasting”. *Federal Reserve Bank of Atlanta.* 2014. https://www.frbatlanta.org/-/media/documents/research/publications/wp/2014/wp1407.pdf

[20] David Jonker, Scott Langevin, David Giesbrecht, Michael Crouch, and Nathan Kronenfeld. Graph Mapping: Multi-Scale Community Visualization Of Massive Graph Data. *Information Visualization*, 16(3):190–204, 2017.