DOWNLOAD PDF INFERRING CAUSAL IMPACT USING BAYESIAN STRUCTURAL TIME-SERIES MODELS

Chapter 1 : An R package for causal inference using Bayesian structural time-series models

Inferring causal impact using Bayesian structural time-series models infer the causal impact that a designed market intervention has exerted on an outcome metric.

Description CausalImpact performs causal inference through counterfactual predictions using a Bayesian structural time-series model. See the package documentation http: In particular, the model assumes that the time series of the treated unit can be explained in terms of a set of covariates which were themselves not affected by the intervention whose causal effect we are interested in. The easiest way of running a causal analysis is to call CausalImpact with data, pre. In this case, a time-series model is automatically constructed and estimated. See Example 1 below. An alternative is to supply a custom model. In this case, the function is called with bsts. See Example 3 below. This can be a zoo object, a vector, a matrix, or a data. In any of these cases, the response variable must be in the first column, and any covariates in subsequent columns. A zoo object is recommended, as its time indices will be used to format the x-axis in plot. This period can be thought of as a training period, used to determine the relationship between the response variable and the covariates. If data is a zoo object with a time attribute, pre. This is the period after the intervention has begun whose effect we are interested in. The relationship between response variable and covariates, as determined during the pre-period, will be used to predict how the response variable should have evolved during the post-period had no intervention taken place. If data is a zoo object with a time attribute, post. One particularly important parameter is prior. For even more control over the model, you can construct your own model using the bsts package and feed the fitted model into CausalImpact , as shown in Example 3. In this case, omit data, pre. Instead only pass in bsts. The model must have been fitted on data where the response variable was set to NA during the post-treatment period. The actual observed data during this period must then be passed to the function in post. This is required if and only if a fitted bsts. Value CausalImpact returns a CausalImpact object containing the original observed response, its counterfactual predictions, as well as pointwise and cumulative impact estimates along with posterior credible intervals. Results can summarised using summary and visualized using plot. The object is a list with the following fields: Time-series object zoo containing the original response response as well as the computed inferences. The added columns are listed in the table below. Summary statistics for the post-intervention period. This includes the posterior expectation of the overall effect, the corresponding posterior credible interval, and the posterior probability that the intervention had any effect, expressed in terms of a one-sided p-value. Note that checking whether the posterior interval includes zero corresponds to a two-sided hypothesis test. In contrast, checking whether the p-value is below alpha corresponds to a one-sided hypothesis test. A suggested verbal interpretation of the results. A list with four elements pre. The field series is a zoo time-series object with the following columns:

DOWNLOAD PDF INFERRING CAUSAL IMPACT USING BAYESIAN STRUCTURAL TIME-SERIES MODELS

Chapter 2 : Inferring Causal Impact Using Bayesian Structural Time-Series Models | the morning paper

Inferring the impact of market interventions is an important and timely problem. Partly because of recent interest in big data, many ï¬•rms have begun to understand that a competitive advantage can be had by systematically using impact measures.

An R package for causal inference using Bayesian structural time-series models What does the package do? This R package implements an approach to estimating the causal effect of a designed intervention on a time series. For example, how many additional daily clicks were generated by an advertising campaign? Answering a question like this can be difficult when a randomized experiment is not available. How does it work? Given a response time series e. This model is then used to try and predict the counterfactual, i. For a quick overview, watch the tutorial video. What assumptions does the model make? As with all non-experimental approaches to causal inference, valid conclusions require strong assumptions. In the case of CausalImpact, we assume that there is a set control time series that were themselves not affected by the intervention. If they were, we might falsely under- or overestimate the true effect. The model also assumes that the relationship between covariates and treated time series, as established during the pre-period, remains stable throughout the post-period see model. How is the package structured? The package is designed to make counterfactual inference as easy as fitting a regression model, but much more powerful, provided the assumptions above are met. The package has a single entry point, the function CausalImpact. Given a response time series and a set of control time series, the function constructs a time-series model, performs posterior inference on the counterfactual, and returns a CausalImpact object. The results can be summarized in terms of a table, a verbal description, or a plot. Creating an example dataset To illustrate how the package works, we create a simple toy dataset. It consists of a response variable y and a predictor x1. The example data has observations. We create an intervention effect by lifting the response variable by 10 units after timepoint Running an analysis To estimate a causal effect, we begin by specifying which period in the data should be used for training the model pre-intervention period and which period for computing a counterfactual prediction post-intervention period. Alternatively, we could specify the periods in terms of dates or time points; see Section 5 for an example. To perform inference, we run the analysis using: The return value is a CausalImpact object. Plotting the results The easiest way of visualizing the results is to use the plot function that is part of the package: The first panel shows the data and a counterfactual prediction for the post-treatment period. The second panel shows the difference between observed data and counterfactual predictions. This is the pointwise causal effect, as estimated by the model. The third panel adds up the pointwise contributions from the second panel, resulting in a plot of the cumulative effect of the intervention. Remember, once again, that all of the above inferences depend critically on the assumption that the covariates were not themselves affected by the intervention. The model also assumes that the relationship between covariates and treated time series, as established during the pre-period, remains stable throughout the post-period. Working with dates and times It is often more natural to feed a time-series object into CausalImpact rather than a data frame. For example, we might create a data variable as follows: Date c "", "" post. Date c "", "" As a result, the x-axis of the plot shows time points instead of indices:

Chapter 3 : Making Causal Impact Analysis Easy | Stitch Fix Technology â€“ Multithreaded

Inferring Causal Impact Using Bayesian Structural Time-Series Models - Brodersen et al. (Google) Today's paper comes from 'The Annals of Applied Statistics' - not one of my usual sources (!), and a good indication that I'm likely to be well out of my depth again for parts of it.

Nevertheless, it addresses a really interesting and relevant question for companies of all shapes and sizes: Or more precisely, how do I accurately measure the impact that a marketing activity had, so that I can figure out whether or not it had a good ROI and hence guide future actions. This also includes things like assessing the impact of the rollout of a new feature, so you can treat the word marketing fairly broadly in this context. The impact of a marketing event is felt over time, so the measured response is a time series variable e. Hence we are interested in the difference between the observed time series in the real-world, and the time-series that would have been observed if the intervention had not taken place. If we could simply directly measure what would have happened then calculating the difference would be easy. Therefore we need to build a predictive model of what would have happened a synthetic control. This has three components: The time-series behaviour of the response itself, prior to the intervention. The time-series behaviour of other contemporaneous time series that were predictive of the target time-series before the intervention and are assumed to remain so afterwards, on the basis that they also have not received any intervention. In practice, there are often many such series available, and the challenge is to pick the relevant subset to use as contemporaneous controls. This selection is done on the pre-treatment portion of potential controls; but their value for predicting the counterfactual lies in their post-treatment behavior. X1 and X2 are two other markets not subject to the intervention that we can use as control series. The solid blue line in the period before the intervention shows how well the model has been tracking actuals, the dashed blue line the period after the intervention shows what the model predicts would have happened. Now we can subtract the prediction of what would have happened dashed blue line from what actually happened Y to see the true causal impact of the marketing event: The Bayesian structural time-series models that the authors build for this are based on two equations, one is the observation equation linking observed data to a state vector, and the other is a state equation that describes how the state vector evolves over time. For the state itself they use a local linear trend, which is good for short term predictions, combined with a seasonality model where appropriate. Control time series are included via linear regression, with a choice of either static or dynamic coefficients. Structural time-series models allow us to examine the time series at hand and flexibly choose appropriate components for trend, seasonality, and either static or dynamic regression for the controls. The presence or absence of seasonality, for example, will usually be obvious by inspection. A more subtle question is whether to choose static or dynamic regression coefficients. When faced with many potential controls, we prefer letting the model choose an appropriate set. This can be achieved by placing a spike-and-slab prior over coefficients. Put all the pieces together and you get something that looks a bit like this. For the detailed description, see the full paper! The ad campaign ran for six consecutive weeks and were geo-targeted to a randomised set of 95 out of designated market areas DMAs. Since paid clicks were zero before the campaign, one might wonder why we could not simply count the number of paid clicks after the campaign had started. The reason is that paid clicks tend to cannibalise some organic clicks. Since we were interested in the net effect, we worked with the total number of clicks. The overall effect estimated by the model was an additional 88, clicks. The original experiment was actually conducted with a conventional control, which enabled the team to assess how well their model stacked up against it. The conventional control gave an estimated uplift of 84, clicks. Statisticians will probably hate me here for eliding all the details about confidence intervals etc. An important characteristic of counterfactual-forecasting approaches is that they do not require a setting in which a set of controls, selected at random, was exempt from the campaign. We therefore repeated the preceding analysis in the following way: In the absence of a dedicated set of control regions, such industry related time series can be very powerful controls, as they capture not only seasonal variations but also market-specific trends and events though not necessarily advertiser-specific trends. A major strength of the controls chosen here is that time series on web searches are publicly available through Google Trends http: This makes the approach applicable to virtually any kind of intervention. Overall, we expect inferences on the causal impact of designed market interventions to play an increasingly prominent role in providing quantitative accounts of return on investment [Danaher and Rust , Leeflang et al. This is because marketing resources, specifically, can only be allocated to whichever campaign elements jointly provide the greatest return on ad spend ROAS if we understand the causal effects of spend on sales, product adoption or user engagement. At the same time, our approach could be used for many other applications involving causal inference. Examples include problems found in economics, epidemiology, biology or the political and social sciences. With the release of the CausalImpact R package we hope to provide a simple framework serving all of these areas. Structural time series models are being used in an increasing number of applications at Google, and we anticipate that they will prove equally useful in many analysis efforts elsewhere.

Chapter 4 : Bayesian structural time series - Wikipedia

data: Time series of response variable and any covariates. This can be a zoo object, a vector, a matrix, or a calendrierdelascience.com any of these cases, the response variable must be in the first column, and any covariates in subsequent columns.

Standard constraints for this optimization problem include: The MarketMatching package specifies the well known Sakoe-Chiba band when calling dtw which allows the user to specify a maximum allowed time difference between two matched data points. Dynamic Time Warping Example To see how this works, consider the following example. Note that the code in this example is not needed to run the MarketMatching package; the package will set it up for you. This is simply to demonstrate the details behind the scene. The following code shows the two time series as well as how data points are connected: It also helps to look at the actual cost matrix and the optimal alignment path that leads to the minimal distance. This can be illustrated by the cumulative cost matrix: This is a large number considering that the distance between Zurich and Copenhagen is The weights that are applied to the local distances are determined by the step pattern used. When the default step pattern is chosen, diagonal steps are weighted by a factor of 2 while other steps receive a weight of 1. We can visualize this by using the plot function provided by the dtw package: As mentioned, the MarketMatching package utilizes the CausalImpact package written by Kay Brodersen at Google see [1] to do the post period inference. Fit a Bayesian structural time series model using data prior to the pre-intervention period. The model can include the control markets as linear regression components with spike-and-slab priors. Based on this model, generate counterfactual predictions for the post-intervention period assuming that the intervention did not take place. In a pure Bayesian fashion, leverage the counterfactual predictions to quantify the causal impact of the intervention. Second, estimating control markets effects with Bayesian priors captures the uncertainty of the relationship between the test market and the control markets. This is critical as it ensures that the counterfactual predictions are not rigidly relying on historical relationships between the test and control markets that may be carrying large standard errors. Moreover the spike-and-slab priors help avoid overfitting by promoting a sparseness during market variable selection. As a result, this approach produces robust counterfactual expectations for the post period that factors in uncertainties in historical market relationships as well as unobserved trends. Moreover, we can calculate posterior intervals through sampling to gauge confidence in the magnitude of causal impact and estimate the posterior probability that the causal impact is non-existent. Some Technical Details When MarketMatching calls CausalImpact the following structural time series model state space model is created for the pre-intervention period: The local level term defines how the latent state evolves over time and is often referred to as the unobserved trend. See [1] and [3] for more details. Once this model is in place, we can create a synthetic control series by predicting the values for the post period and then compare to the actual values to estimate the impact of the event. In order to gauge the believability of the estimated impact, posterior intervals can be created through sampling in a pure Bayesian fashion. We can also compute the tail probability of a non-zero impact. The posterior inference is conveniently provided by CausalImpact package. Note that the CausalImpact package can fit much more complicated structural time series models with seasonal terms as well as dynamic coefficients for the linear regression component. However, MarketMatching package requests the more conservative model based on the assumptions that the control markets will handle seasonality and that static control market coefficients is sufficient. About Spike-and-Slab Priors As mentioned in the overview, we can select the final control markets while fitting structural time series model. The dynamic time warping step pre-screens the markets, while picking the final markets can be treated as a variable selection problem. In the CausalImpact package this is done by applying spike-and-slab priors to the coefficients of the linear regression terms. Spike-and-slab prior consist of two parts: This is typically a product of independent Bernoulli distributions one for each variable , where the parameters probability of getting chosen can be set according to the expected model size. The slab part is a wide-variance Gaussian prior that shrinks the non-zero coefficients toward some value usually zero. This helps combat multicollinearity, which is rampant since markets tend to be highly correlated. This approach is a powerful way of reducing a large set of correlated markets into a parsimonious model that averages over a smaller set of markets. Moreover, since this is a Bayesian model, the market coefficients follow random distributions and we can incorporate the uncertainties of the historical relationships when forecasting as opposed to relying on a rigid encoding based on fixed coefficients. For more details in spike-and-slab priors, see [1]. Using a pure goodness-of-fit-based measures is not meaningful as a higher SE will always guarantee a better historical fit. In addition, the goal of the model is not to fit the data in the post-intervention period which means that we cannot use the post intervention period as a hold-out sample. Here are some tips to deciding the size of the standard error: In general, larger values of the standard error leads to wider posterior forecast intervals and hence results are more likely to be inconclusive. If you know a priori that the test market series is volatile due to unexplained noise, choose 0. The MAPE measures the historical fit and the Durbin-Watson statistic measures the level of autocorrelation in the residuals. This analysis will show the tradeoff between a larger standard error versus fit and ill-behaved residuals. We want to choose a standard error that is as small as possible in order to rely more on the predictive value coming from the control markets, but not at any cost. The MarketMatching package produces charts that help make this tradeoff. Note that the Durbin-Watson statistic should be as close to 2 as possible. When you cannot make a decision, choose 0. Note on Structural Time Series Models The class of structural time series models deployed by the CausalImpact package provides the most flexible and transparent approach to modeling time series data. This model can be recast as the structural model with a local level term the model described above. References [1] CausalImpact version 1. Scott and Hal Varian, http: Do you want to build amazing products with amazing peers?

Chapter 5 : RPubs - Inferring Causal Impact Using Bayesian Structural Time-Series Models

An important problem in econometrics and marketing is to infer the causal impact that a designed market intervention has exerted on an outcome metric over time. This paper proposes to infer causal impact on the basis of a diffusion-regression state-space model that predicts the counterfactual market.

MR Digital Object Identifier: Synthetic control methods for comparative case studies: The economic costs of conflict: A case study of the basque country. Empirical strategies in labor economics. Handbook of Labor Economics 3 â€” On making causal claims: A review and recommendations. Using the longitudinal structure of earnings to estimate the effect of training programs. Identification and inference in nonlinear difference-in-differences models. Modeling Internet firm survival using Bayesian dynamic models with time-varying coefficients. Program evaluation with high-dimensional data. The Practice of Econometrics: How much should we trust differences-in-differences estimates? Models of causal inference: Going beyond the Neyman-Rubin-Holland theory. Variational Bayesian mixed-effects inference for classification studies. A new data mining approach to estimate causal effects of policy interventions. Experimental and Quasi-Experimental Designs for Research. Minimum wages and employment: A case study of the fast food industry in New Jersey and Pennsylvania. On Gibbs sampling for state space models. Evaluating online ad campaigns in a pipeline: Causal models at scale. The practical implementation of Bayesian model selection. The Russoâ€"Williamson Theses in the social sciences: Causal inference drawing on two types of evidence. Causal inference and statistical fallacies. Determining the optimal return on investment for an advertising campaign. The simulation smoother for time series models. Inference with difference-in-differences and other panel data. A simple and efficient simulation smoother for state space time series analysis. Data augmentation and dynamic linear models. Variable selection via Gibbs sampling. Approaches for Bayesian variable selection. Rao-Blackwellization for Bayesian variable selection and model averaging in linear and binary regression: A novel data augmentation approach. Generalized least squares inference in panel and multilevel models with serial correlation and fixed effects. Econometric evaluation of social programs, Part I: Causal models, structural models and econometric policy evaluation. Do all and only causes raise the probabilities of effects? In Causation and Counterfactuals. Economic theory and causal inference. In Philosophy of Economics 13 U. A review of causal inference for biomedical informatics. Creating lift versus building the base: Current trends in marketing dynamics. Shortcomings of marginal analysis for wage-employment problems. Here, there, and everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising. Does retail advertising work? A Bayesian foundation for individual learning under uncertainty. Natural and quasi-experiments in economics. Counterfactuals and Causal Inference: Methods and Principles for Social Research. Bayesian analysis of latent threshold dynamic models. Data augmentation for support vector machines. The search for political experiments in nature. Estimating causal effects of treatments in randomized and nonrandomized studies. Statistical inference for causal effects, with emphasis on applications in epidemiology and medical statistics. In Epidemiology and Medical Statistics J. Estimating the causal effects of marketing interventions using propensity score methodology. Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. Predicting the present with Bayesian structural time series. International Journal of Mathematical Modeling and Optimization 5 4â€" Measurement of return on marketing investment: A conceptual framework and the future of marketing metrics. Estimating Autocorrelations in Fixed-Effects Models. Linking marketing actions to financial results. Multiple time series analysis of competitive marketing behavior. Measuring ad effectiveness using geo experiments. Technical report, Google Inc. Periodic measurement of advertising effectiveness using multiple-test-period geo experiments. Bayesian Forecasting and Dynamic Models, 2nd ed. The estimation of causal effects from observational data. In Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti P.

DOWNLOAD PDF INFERRING CAUSAL IMPACT USING BAYESIAN STRUCTURAL TIME-SERIES MODELS

Chapter 6 : Inferring causal impact using Bayesian structural time-series models â€“ Google AI

Structural time-series models allow us to examine the time series at hand and ï¬‚exibly choose appropriate components for trend, seasonality, and either static or dynamic re gression for the controls.