Int. J. Production Economics 177 (2016) 139–148

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Forecasting supply chain sporadic demand with nearest neighbor approaches

Konstantinos I. Nikolopoulos [a,*], M. Zied Babai [b], Konstantinos Bozos [c]

[a] forLAB, Bangor Business School, Prifysgol Bangor University, Bangor, Gwynedd, United Kingdom
[b] Kedge Business School, Bordeaux, France
[c] Leeds University Business School, University of Leeds, Leeds, United Kingdom

ARTICLE INFO

Article history:
Received 6 March 2015
Received in revised form
24 February 2016
Accepted 18 April 2016
Available online 22 April 2016

Keywords:
Supply chain
Logistics
Demand forecasting
Exponential smoothing
Nearest neighbors

ABSTRACT

One of the biggest challenges in Supply Chain Management (SCM) is to forecast sporadic demand. Our forecasting methods' arsenal includes Croston's method, SBA and TSB as well as some more recent non-parametric advances, but none of these can identify and extrapolate patterns existing in data; this is essential as these patterns do appear quite often, driven by infrequent but nevertheless repetitive managerial practices. One could claim such patterns could be picked up by Artificial Intelligence approaches, however these do need large training datasets, unfortunately non-existent in industrial time series. Nearest neighbors (NN) can however operate in these latter contexts, and pick up patterns even in short series. In this research we propose applying NN for supply chain data and we investigate the conditions under which these perform adequately through an extensive simulation. Furthermore, via an empirical investigation in automotive data we provide evidence that practitioners could benefit from employing supervised NN approaches. The contribution of this research is not in the development of a new theory, but in the proposition of a new conceptual framework that brings existing theory (i.e. NN) from Computer Science and Statistics and applies it successfully in an SCM setting.

1. Introduction

One of the biggest challenges in Operations Management (OM) is to forecast as accurately as possible the sporadic/intermittent demand faced in most supply chain and logistics operations. Assuming we are not in a made-to-order supply context, then forecasters are facing true uncertain stochastic demand. The reason that makes the forecasting task so challenging has to be attributed to the dual uncertainty faced by decision makers due to the sporadic nature of both the actual demand volume as well as the timing of demand arrivals; the latter do not occur during each and every period of time and as such there are many periods in time with zero demand.

Especially for logistics operations where the spatial dimension comes into the equation, the appearance of sporadic demand is more often, as prima facie regular demand at the aggregate manufacturing or distribution-center level becomes intermittent when demand is realized through the alternative geographical channels. Furthermore the alternative direct retailing channels through online and digital media create more 'channel products' for which demand needs to be forecasted and stock to be replenished.

There have been a few forecasting methods tailored for such data that have been developed over the last forty years (Petropoulos et al., 2014), mostly exponential smoothing approaches like Croston (1972), Syntetos and Boylan approximation–SBA (2001) and more recently a method developed from Teunter, Syntetos and Babai–TSB (2011). A handful of more advanced but at the same more complex non-parametric methods has also been proposed over the years but these do not necessarily outperform the former (Syntetos et al., 2015).

In general the assumption is that the aforementioned intermittent series present no trend or autocorrelation, no seasonality or cycles – in essence none whatsoever structural component that could be identified via a formal statistical procedure. However, this for practitioners just does not make sense as they do know that very often they apply the same tactics when they order, replenish or produce products and services; it is just that these are happening in non-periodic lags and thus becomes very difficult to spot them in the past of a time series. Attesting to that, Altay et al. (2008, 2012) have provided evidence of existence of trend and correlation in intermittent data and proved that these features significantly affect forecasting and stock control performance.

* Corresponding author.
E-mail addresses: k.nikolopoulos@bangor.ac.uk (K.I. Nikolopoulos), mohamed-zied.babai@kedgebs.com (M.Z. Babai), k.bozos@lubs.leeds.ac.uk (K. Bozos).

http://dx.doi.org/10.1016/j.ijpe.2016.04.013
0925-5273/

Unfortunately, none of the proposed to date methods can pick up one-off (or more frequent) patterns existing in the past of a series, patterns that may appear quite often due to irregular but repetitive managerial practices. Artificial Intelligence and other computationally intensive approaches (Haykin, 1998) that could pick up such patterns have been successfully applied in other disciplines like finance, economics, marketing and computer science. These methods however do rely on very large training datasets in order to predict satisfactorily, thus are not fit-for-purpose in our context where lengthy series are perceived as quite a luxury. Nevertheless there is one non-parametric regression smoothing method that does not require large training data sets: nearest neighbors (NN), which are capable of producing forecasts even when only just two or three observations are available.

In this research we propose applying NN for supply chain data and we investigate the conditions under which these perform adequately through an extensive simulation on 480 time series where we control for intermittence levels, frequency, plurality and type of patterns, and level of noise. Furthermore via an empirical investigation in automotive data over 3000 SKUs we provide evidence that practitioners could benefit from employing supervised NN approaches – supervised in the sense that you only allow the patterns to be identified if there is some evidence that these patterns do exist – this in this paper is achieved through an adjusted cumulative autocorrelation function.

The contribution of this research is not in the development of a new theory, but in the proposition of a new conceptual framework that brings existing theory (i.e. NN) from Computer Science and Statistics and applies it successfully in an SCM setting. The proposed framework comes with a series of limitations as it should be applied only when evidence of existence of patterns is apparent, and even so not left unsupervised rather than applied only in the part of the dataset that exhibits higher autocorrelations. Ceteris paribus, we still believe that we live in a world where humans drive data, through their action and behaviors eventually do introduce infrequent patterns in time series data; and as such we do believe there is merit for practitioners and academics to consider our proposition and further research it in the future.

The remaining of the paper is structured as follows: Section 2 revisits the background literature while Section 3 presents the conceptual framework of the NN method and gives an illustrative example of how the method could be used in a SCM context; Section 4 compares the newly proposed method with other classical parametric approaches. Section 5 gives the results of an extensive empirical evaluation across a real large dataset from RAF. Section 6 discusses the implications for SCM theory and practice, while the Section 7 concludes and provides the roadmaps for future research.

2. Background literature

The first part of the literature review revisits the most important methods for forecasting intermittent demand. The second part focuses on successful applications of nearest neighbors in various contexts.

2.1. Forecasting supply chain sporadic demand

For intermittent data, simple techniques such as Naïve, Moving Averages and Simple Exponential Smoothing (SES) have been quite popular among practitioners over the years due to their simplicity and accuracy but most importantly the ability to handle non-demand observations without the need of time series transformations (Petropoulos et al., 2013). The first method tailored to intermittent data came from John Croston in his seminal article Croston (1972) proposing a decomposition of the data into two subseries, one for the positive demands (excluding the zeros) and one for the arrival intervals. Syntetos and Boylan (2001) proposed and evaluated successfully (Syntetos and Boylan, 2005) SBA: a bias-correction approximation to the Croston method. More recently, Teunter et al. (2011) suggested the TSB decomposition method that builds on the separate extrapolation of the non-zero demands and the probability to have a demand; with this method being very useful in cases of obsolescence.

A handful of non-parametric methods has been proposed over the years, most notably from Willemain et al. (2004) with a bootstrapping approach that captures potential auto-correlations of the underlying demand patterns and simultaneously accounts for variability not observed in the original demand sample through the patented 'Jittering' process. AI approaches have been also proposed; most often Artificial Neural Networks (Gutierrez et al., 2008) but these do need very large training datasets. In any case, there is no sufficient empirical evidence that these methods are more accurate from the simpler ones (Syntetos et al., 2015).

More recently, Nikolopoulos et al. (2011) proposed the ADIDA non-overlapping temporal aggregation forecasting framework that was successfully evaluated both in terms of forecasting accuracy as well as of stock control performance (Babai et al., 2012). The proposed framework is now perceived as a "forecasting-method improving" mechanism that through frequency transformations helps methods achieve better accuracy performance. The first theoretical developments for the framework appeared recently in the literature (Spithourakis et al., 2014; Rostami-Tabar et al., 2013, 2014). Kourentzes et al. (2014) extended this idea by means of estimating time series structural components across multiple frequencies and optimally extrapolate and combine them; with empirical results being quite promising for long-term forecasting. Petropoulos and Kourentzes (2014) also proposed forecasting method combinations on the aforementioned context with improved forecasting performance.

2.2. Nearest neighbors

Nearest neighbor approaches (NN/NNs, Härdle, 1992)[1] are quite popular in the forecasting literature (Green, 2002), largely because of their intuitively appealing simplistic nature and theoretical attributes (Yakowitz, 1987). They are generally found to present distinct advantages versus their alternatives for non-linear fluctuations, since, while their parameters are linear in nature, NNs can capture the complex non-linear patterns among neighbors (Yankov et al., 2006), thus predicting composite non-linear behaviors in a fairly accurate manner. Hence, NN methods have been applied to time-series and cross-sectional data in a wide variety of domains and have often been found to outperform alternative, most of the time far more complex approaches. However to the best of our knowledge these have never been applied in an SCM sporadic demand context. We hereafter review some of the most indicative works in the most common domains.

2.2.1. Economics and finance

In his attempt to beat the 'tenacious' random walk model in forecasting exchange rates, Mizrach (1992) combines k-NN estimators (using the outcomes from 'k' closest past cases = the neighbors) with a locally weighted regression procedure. He finds an improvement in the forecast accuracy in just one out of the three rates examined, a result which is however not robust to

[1] This is by many academics considered the most complete reference book describing the theoretical foundations as well as many practical applications of nearest neighbors for smoothing, classification and forecasting.