A novel method to identify the start and end of the winter surge in demand for pediatric intensive care in real time

Christina Pagel[1] PhD, Padmanabhan Ramnarayan[2] MD, FRCPCH, Samiran Ray[2] MD, Mark J. Peters[2,3] FRCPCH, PhD

1. UCL Clinical Operational Research Unit, University College London, London

2. Children's Acute Transport Service (CATS), Great Ormond Street Hospital NHS Foundation Trust, 26-27, Boswell Street, London, WC1N 3JZ, UK

3. Respiratory, Anaesthesia and Critical Care Group, Institute of Child Health, University College London, London

Corresponding author: Christina Pagel

Email:c.pagel@ucl.ac.uk

Phone number: +44 (0)20 7679 4501

KEYWORDS

Emergency Medical Services; Pediatric intensive care units; Decision Making (Organizational); Planning techniques; Informatics; Algorithms

ABSTRACT

OBJECTIVE

Implementation of winter surge management in intensive care is hampered by the annual variability in the start and duration of the winter surge. We aimed to develop a real-time monitoring system that could identify the start promptly and accurately predict the end of the winter surge in a pediatric intensive care (PIC) setting.

DESIGN

We adapted a method from the stock market called "Bollinger bands" to compare current levels of demand for PIC services to thresholds based on medium term average demand. Algorithms to identify the start and end of the surge were developed using Bollinger bands and pragmatic considerations. The method was applied to a specific PIC service: the North Thames Children's Acute Transport Service (CATS) using eight winters of data (2005-2012) to tune the algorithms and one winter to test the final method (2013/14).

SETTING

A regional specialised pediatric retrieval service based in London, UK.

RESULTS

The optimal Bollinger band thresholds were 1.2 and 1 standard deviations above and below a 41-day moving average of demand respectively. A simple linear model was found to predict the end of the surge and overall surge demand volume as soon as the start had been identified. Applying the method to the validation winter of 2013/14 showed excellent performance, with the surge identified from 18th November 2013 to 4th January 2014.

CONCLUSIONS

We have developed and tested a novel method to identify the start and predict the end of the winter surge in emergency demand for pediatric intensive care.

INTRODUCTION

Each winter in the UK, there is a significant increase in the number of emergency admissions to hospitals, particularly in patients with respiratory illness. In the United Kingdom, this surge in demand for hospital beds regularly outstrips availability and results in frequent bed crises. The British Medical Association, NHS England, and many hospitals, have plans for winter preparedness to cope with these pressures (e.g. 1-4). However, these plans tend to be reactive once bed pressures arise, rather than being proactive attempts to manage capacity. Real-time monitoring of demand and prompt identification of a demand surge offers the potential for more informed management of the system

Winter pressures on pediatric intensive care units (PICUs) and retrieval services can be particularly acute. Referrals of infants with acute respiratory failure approximately double during this period and almost all these infants are admitted to PICU following retrieval from district general hospitals [5]. Plans for surge management have been published recently by the Paediatric Intensive Care Society (PICS) and NHS England [6]. They identify the winter surge as lasting an average of 6 weeks between mid-November and January, and recommend various measures during this time (see also 7), such as to increase capacity, streamline work processes and reschedule elective workload.

Despite these plans, their implementation in practice is hampered by variability in timing of the start of the winter surge from year to year. In the UK there is some real time surveillance during winter for factors "upstream" of general emergency demand, such as general practitioner visits for seasonal influenza, but these factors may not be relevant to pediatric ICU services and are usually published weekly and not daily. Surveillance of other potential data streams such as reported cases of RSV infection is currently not possible in real time. An accurate method to identify the start and end of a surge, and to predict the overall level of demand during it, will be useful for PICU services to implement targeted surge management plans in a clinically relevant and cost effective manner. The start of the surge should ideally be identified as soon as possible based on monitoring demand levels on a daily basis rather than weekly or monthly.

In this study, we describe the novel use of a statistical process control method adopted from the stock market to signal alerts to mark the start and end of the winter surge accurately and illustrate

the method for a large regional PICU retrieval service based in London. Although the timing of seasonal surges may vary with geographical location [7,8], this does not affect the applicability of the method.

METHODS

Statistical control method

Standard definition of Bollinger bands

Bollinger bands were introduced by stock market investor John Bollinger in 1992 [9] as way to trigger buy and sell signals on shares by comparing current prices to the medium term moving average.

For a daily time series, Bollinger bands are traditionally defined as two standard deviations either side of a 20-day moving average (i.e. a "window size" of 20 days and a "band width" of 2 standard deviations). However, the choice of window size and band width is flexible and depends on the purpose for which it is designed. Although very simple, Bollinger bands have been shown to be profitable rules of thumb and continue to be used today in the stock market [10,11]. They have also been used in completely different contexts such as identifying fabric defects during manufacturing [12], although we were able to find only a few documented uses outside of finance, and none in health care.

The advantage of using a technique like Bollinger bands for monitoring demand in pediatric intensive care is that they rely only on knowing recent demand and are simple to implement. Bollinger Bands can be thought of as another type of statistical process control method.

Using Bollinger bands to identify the start and end of a winter surge

Use of Bollinger bands in the stock market is focused on spotting when a time series has deviated significantly from a recent medium term average. The underlying assumption is that such deviations cannot be predicted. In the stock market, all deviations are important regardless of when in the year they occur. However the situation for pediatric intensive care services is somewhat different: there is a surge every winter that places significant pressure on PICU services. Thus, if looking at daily demand over the year, demand *will* breach the upper Bollinger band at some point in the autumn as demand

4

increases during the surge and it *will* breach the lower Bollinger band as demand decreases once the surge is over. Those breaches could then act as surge identifiers, signalling the start and end of the winter surge. Our aim was to choose a window size and band width that respond quickly to changes in demand (to provide timely warning that the surge has started or ended) but not so quickly that a transient spike or dip in demand will result in a 'false positive' identification. Allowing the band width and window size to be determined by tuning the identification algorithms to historical data is a key difference in our approach from both the use of Bollinger Bands in the stock market and other statistical process control (SPC) methods in health care.

Additionally, the likelihood of a false positive identification can be reduced by requiring that daily demand must consistently breach either the upper Bollinger band (to define the start of the surge) or the lower Bollinger band (to define the end of the surge). An example would be requiring daily demand to breech the Bollinger band for several days in a row before the start of the surge is formally identified (e.g. see 13).

Finally, daily demand breaches the upper Bollinger Band whenever current demand is significantly higher than the recent average. However, while this does represent a surge in demand, it only matters to a service if demand is high enough to strain capacity. For instance, in a service that has experienced consistently low demand (e.g. during August), a surge to medium demand is unlikely to present any problems in meeting that demand. To be useful for winter planning purposes, the identification of the start of the winter surge should occur *only* when demand has reached a level that strains available capacity.

The most suitable choices for window size, band width, consistency and absolute demand requirements are likely to depend on the particular pediatric intensive care service. The level of demand that is considered to strain available capacity must be defined by the local team. Choices for window size, band width and consistency can then be optimised by defining start and end dates for previous winter surges in that service.

Once Bollinger bands, consistency and demand thresholds have been identified to identify the start and end of the winter surge, relationships between identified start and end dates and overall demand volume can be explored to enable prediction of the end and size of the winter surge as soon as the start has been identified.

The method is not fitting a statistical model to the data. The approach adopted is deliberately pragmatic, making the most of a single time series of recent demand that is likely to be readily available, to determine when the winter surge has started and ended within a particular service.

Example application in the North Thames Children's Acute Transport Service (CATS)

The Children's Acute Transport Service (CATS) is a regional PICU retrieval service based at Great Ormond Street Hospital in London, UK. CATS transports patients from over 50 district general hospitals covering a large geographical region in the South East of England to any one of three tertiary intensive care units in North London (Great Ormond Street Hospital, the Royal Brompton Hospital and St Mary's Hospital). The CATS service does not sit within an ICU and has two dedicated retrieval teams. In this work, we applied the method to demand for CATS retrieval.

Data on all CATS referrals are checked and entered onto a dedicated database daily so that the database is up to date and of high quality in almost real time. In practice, at 9 am on any given day, data exist up until the day before. The data used to develop and test the method were anonymised reports generated from routinely collected data. The study was discussed with the local Independent Review Board Chair who confirmed that ethical approval was not required.

To tune the method, we used data on all calls to CATS from April 2005 to July 2013. We quarantined data from July 2013 to February 2014 to validate the final choices of window size, band width and consistency requirements. We defined "demand for retrievals" as all retrievals performed plus retrievals refused due to lack of a retrieval team or PICU capacity.

Since raw daily demand data is highly variable, we used the rolling 7-day total demand as our daily time series of interest. The clinical team defined a demand level of "28 retrievals in the last week" as the threshold beyond which capacity was strained. To prevent 'false positive' identification, we required demand to breach the upper or lower Bollinger band three days in a row to identify the start or end of the surge. Finally, since we are only concerned with the winter surge, we added the constraint that the start could only be identified on or after the 1st October and that the end could only be triggered on or after the 1st December and more than 31 days after the start of the surge. These fixed date and duration constraints were based on historical data (the earliest start of the winter surge was 8th October in 2010) but we acknowledge that they are somewhat arbitrary.

Identifying the start of winter surge

The start of the winter surge would be identified if it was after September and EITHER the rolling 7-day total demand had breached the upper Bollinger band three days in a row and the most recent 7-day total had reached 28 OR we had seen demand at or above 28 for four consecutive days. The latter was included to allow for the (rare) possibility of a rising tide of demand that was slow enough to stay within the Bollinger bands (based on a medium-term moving average) but would nonetheless result in consistently high demand.

Identifying the end of winter surge

The end of the winter surge would be identified if it was after November, more than 31 days since the start and the rolling 7-day total demand has breached the lower Bollinger band three days in a row.

To define the historical winter surges, two authors (CP, PR) visually picked out the start and end dates for the surge each year using the "28 a week" demand threshold as a guide. The choices of window size and band width were optimised by minimising the sum of squared differences between the automatically identified dates and the manually chosen dates. We allowed the window size and band width to differ for identifying the start and end of the winter surge.

Finally, the chosen Bollinger bands were tested on the quarantined 2013/2014 winter data as a performance check in "out-of-sample" data.

RESULTS

Between 1 April 2005 and 15 July 2013, CATS received 17,527 calls. Of these, 9,731 represented genuine demand for retrieval, with a daily mean of 3.2 and standard deviation 1.7. The overall referral breakdown by outcome is given in Table 1, along with an indication of which referrals corresponded to genuine demand for retrieval.

The seasonal nature of CATS activity can be clearly seen from the monthly time series of demand shown in Figure 2.

Results of the optimisation

The optimal Bollinger bands for both the start and end of the surge use a window size of 41 days. The Bollinger Band widths for the start and end of the winter surge are 1.2 and 1 standard deviations respectively. The corresponding automatically identified start and end dates, along with the manually identified dates and the sum of square errors, are shown in Table 2. The dates identified using these bands matched the manually identified start and end dates very well, particularly for the start of the surge.

Predicting the end of the winter surge and total volume of demand

Exploring graphically the relationship between the identified start and end dates and volume of demand, the duration and overall volume of demand were linearly related to the start of the surge. We defined a new variable s_{year} as the number of days after the 1st October (the first possible identification date) that the winter surge is identified each year. Plotting the duration of the surge, defined as the number days between the identified start and end, versus s_{year}, we found an excellent linear relationship (Figure 2). Using a simple linear regression we get a fitted R^2 value of 0.91 and the fitted equation:

$$duration_{year} = 108.7 - 1.26s_{year} \ .$$

We also found an excellent linear relationship exploring the relationship of total volume of demand with s_{year}. This is a result of the fact that during the surge each year, despite the peaks and troughs, the average daily volume of demand was very close to 4 for all historical years. Thus the overall volume during the surge was simply (approximately) 4 multiplied by the number of days it lasted. Using this, we were able to calculate a reasonable estimate of the total volume of demand over the surge. The fitted linear relationship for volume had an R^2 of 0.91 and equation:

$volume_{year} = 452.7 - 5.12s_{year}$. However, we note that while overall demand during the signalled surges average around "4 a day", the daily demand experienced by CATS remained highly variable, fluctuating anywhere between 0 and 10 on any given day.

Testing the automated identification method on the winter of 2013/4

Running the optimised Bollinger bands on the quarantined dataset, the start of the winter surge was identified as the 18th November 2013 (Figure 3). When using a fitted equation to predict a future

value, there is some uncertainty in the predicted estimate which can be represented by a prediction interval. Applying the fitted equations to predict the duration and volume (with 60% prediction intervals given in the square brackets), we expected the end of surge to be 5th January 2014 [29 December, 11th January] (shown as the grey dashed line in Figure 3) and an expected overall volume of demand 207 [181, 241] over the winter surge. Thus the end of the surge and the overall volume of the surge have a 60% chance of falling within their respective prediction intervals. The automatically identified end of the surge was the 4th January 2014 and the overall volume of demand was 235 (see Figure 4), both well within their respective prediction intervals. As tested on the winter of 2013/4, our method for identifying the start and end of the winter surge is fit for purpose.

DISCUSSION

Using methods adapted from the stock market, we have developed a system for monitoring daily demand that can be used to identify the start and end of the winter surge for emergency pediatric intensive care services in real time. Importantly, the system also provides robust predictions of the duration of the winter surge and the total volume of demand during that time at the beginning of the surge. We suspect that the accuracy of these predictions is due at least partly to the fact that the winter surge almost always ends during the ten days of January, regardless of when the winter surge started (i.e. starting earlier does not mean it will end sooner!). The method performed very well when tested on data from a single pediatric retrieval service.

The method identifies the start and end of the winter surge by comparing current demand to a medium-term moving average, where the optimal window size for the moving average was 41 days. The relatively large window size (almost 6 weeks) is probably required to prevent false positive identification from transient fluctuations in demand. In almost all cases, the automatically identified date for both the start and the end of the surge was after the manually identified date. This is not surprising since a delay of at least three days was built into the method by requiring a consistent breach. This is not a weakness of the method since when looking at historical data with the benefit of hindsight (as was done to identify the target dates), it is relatively easy to identify the start and end of the surge. However, if using this method to automatically identify the start and end in real time, it is

important to be relatively sure that that we are identifying a real phenomenon and not simply a transient spike (or dip).

Although demand for pediatric ICU is not exclusively from external emergency referrals, it is such external referrals that drive the increase in demand every winter. Retrieval services cover an entire geographical area and, from an emergency demand point of view, sit "upstream" of individual PICUs. Applying this method within a retrieval service thus provides a practical way to alert clinical teams within local PICUs and local commissioners when the winter surge has started. The ability to know when the surge has started provides significant advantages over the current crude fixed dates used by several national bodies (e.g. 6). A key advantage is that the method also predicts the end date of the surge once the start has been signalled; this can be simply translated into a more approximate rule of thumb "the surge for PIC services will end during the first two weeks of January regardless of how early or late it starts".

A surge in demand for PIC services occurs regularly every winter in the UK and stretches a system that is already running at high bed occupancy levels [14]. The use of specialist retrieval teams has been shown to improve survival [15]. Lack of specialist teams to undertake retrieval may thus worsen outcomes for critically ill children either because children are then transported by non-specialised teams or because children may have to spend longer in local hospitals that are not equipped with a full range of specialist PIC services. In parallel, unavailability of regional PIC beds may result in children being managed at the referring hospital for prolonged periods and/or requiring long distance retrieval to out of region PICUs, further exacerbating the mismatch between supply of PIC resources and demand.

The availability of a system to alert clinical teams to the start of a surge in demand, along with predictions of the end of the surge and the total volume of demand, will have significant implications on how emergency preparedness plans are implemented in PIC in the future. Since emergency demand is inherently unpredictable and cannot be controlled, the main response from the retrieval service could be to increase the number of available teams from the identified date (rather than a fixed date) until the predicted end of the surge, although we acknowledge difficulties with designing rotas that are flexible enough to cope with short term changes. Knowledge of the predicted demand will also help plan for the number of such additional teams required. Further work will include

exploring the pattern of PICU occupancy in the North Thames region before, during and after identified winter surge periods. Other possible responses by retrieval services could include reducing the number of non-essential meetings, restricting annual leave or using non-clinical staff days. Sharing information regarding the surge with the regional PICUs will help them prepare for the excess demand by hiring temporary staff to open more beds or by rescheduling elective surgery cases to after the end of the surge period. While these measures have already been tried in the past, they have either been too late (long after the surge has started) or too early (weeks before any surge), resulting in an ineffective response both clinically and from a cost perspective. The tool was used successfully in real time at CATS for the winter of 2014/15 and its outputs shared with other regional retrieval services and local service commissioners. Next year, we expect this tool to include informing the scheduling of the weekly regional teleconferences during winter and to inform local PICUs when CATS has start experiencing the winter surge. We also note that since the implemented tool runs throughout the year, plotting the current and recent demand with the upper and lower Bollinger Bands, it provides a daily visual check for sudden increases in demand outside of winter and would highlight unexpected demand due to, for instance, a new epidemic.

The method's robustness and simplicity should also allow for a relatively straightforward application to any PIC service, and indeed, other health care environments where external demand for services is both highly variable and unpredictable. Applying it to another PIC service does require local knowledge of what is a suitable threshold of high demand and requires data to be available on refusals that were in scope of the service. At this time, these requirements make this method more suitable for local use than for UK national use (using the PICAnet dataset), although it is conceivable that it could be used nationally in the future using the new PICAnet referrals dataset. The winter surge in demand for PIC services comes almost entirely from children with respiratory problems [6] who need intensive care support but not necessarily other forms of care such as surgical interventions. For very different contexts, the methodology for signalling a surge would be the same but the range of possible responses/implications would differ.

This is also a novel method - to our knowledge, this is the first time that Bollinger bands have been adapted for use in health care services. The methodology has been implemented at CATS for the winter of 2014/15, so that an Excel spreadsheet runs automatically every morning to produce up to date plots of recent demand and the optimal Bollinger bands. Although the method has performed

well on previous years, including the validation year of 2013/4, it is still untested, and will inevitably be different for the team to experience it in real time as opposed to reviewing the surges in hindsight. That said, the algorithm for identifying the start and end of the winter surge can be adapted after each year to take into account evolving experience – inevitably the eight years used to develop this initial method will not provide the full range of possible winter surge activity. We thus expect this method to become ever more accurate over time as more data is collected.

CONCLUSIONS

We have developed and tested a novel method to identify the start and end of the winter surge in emergency demand for pediatric intensive care depending on absolute levels of demand and how these compare to the 41-day moving averages and standard deviation of demand. Prospective studies are required to validate the method when used in real time, and to study the effects of their implementation in clinical practice.

ACKNOWLEDGEMENTS

REFERENCES

1.   British Medical Association - Winter pressures. Available at: http://bma.org.uk/winter-pressures. Accessed 2014 Aug 29

2.   University College London Hospital (2013) Winter preparedness 2013-14. Available at: https://www.uclh.nhs.uk/aboutus/wwd/Pages/Winterpreparedness.aspx Accessed 2014 Aug 29

3.   NHS Wales (2013) Winter Plan 2013, Available at: http://www.wales.nhs.uk/sitesplus/documents/862/Item10bHDHBWinterPlan2013.pdf Accessed August 29 2014.

4.   Hanratty B, Robinson M: Coping with winter bed crises. *BMJ* 1999: 319(7224):1511–1512.

5.   O'Donnell DR, Parslow RC, Draper ES: Deprivation, ethnicity and prematurity in infant respiratory failure in PICU in the UK, *Acta Pediatrica* 2010: 99(8):1186-91.

6.   Pediatric Intensive Care Surge Report. Pediatric Intensive Care Society; 2014.

7.   Ong MEH, Ho KK, Tan TP, Koh SK, Almuthar Z, Overton J, Lim SH: Using demand analysis and system status management for predicting ED attendances and rostering. *American Journal of Emergency Medicine* 2009: 27(1):16-22

8.   Wiler JL, Griffey RT, Olsen T: Review of Modeling Approaches for Emergency Department Patient Flow and Crowding Research. *Academic Emergency Medicine* 2011: 18(12): 1371-9.

9.   Bollinger J.: Using Bollinger bands. *Tech Anal Stocks Commod V.* 1992: 10(2):47–51.

10.   Lento C, Gradojevic N: The Profitability Of Technical Trading Rules: A Combined Signal Approach. *J Appl Bus Res* 2011: 23(1)

11.   Leung JM-J, Chong TT-L. An empirical comparison of moving average envelopes and Bollinger bands. *Appl Econ Lett.* 2003: 10(6):339–41.

12.   Ngan HYT, Pang GKH: Novel method for patterned fabric inspection using Bollinger bands. *Opt Eng.* 2006: 45(8):087202–087202–15.

13.   Wheeler DJ. Making Sense of Data. Knoxville, Tenn: SPC Press; 2003.

14.   Pearson GA, Reynolds F, Stickley J: Calculating the need for intensive care beds. *Arch Dis Child* 2012: 97(11):943–6

15.   Ramnarayan P, Thiru K, Parslow RC, Harrison DA, Draper ES, Rowan KM: Effect of specialist retrieval teams on outcomes in children admitted to pediatric intensive care units in England and Wales: a retrospective cohort study. *Lancet* 2010: 376(9742):698-704

TABLE CAPTIONS

Table 1 - Number of calls by type of call received by Children's Acute Transport Service (CATS) and outcome of call between 1 April 2005 and 15 July 2013. The final column indicates which calls were considered to represent "demand for retrieval".

Table 2 –Date for start and finish of the winter surge period as identified by the Bollinger bands method. The squared difference is defined as the square of the difference in days between the automatic and manually identified dates. So the first row, there are 3 days difference between the automatically signalled start of 7 November 2005 and the manually identified start of 4th November and this is then squared to give a value of 9.

FIGURE CAPTIONS

Figure 1 – Monthly demand for retrievals by the Children's Acute Transport Service 2005-2013. The annual winter peak is clearly seen.

Figure 2 - Relationship between the number of days between 1 October and the identified start of the winter surge (s_{year}) and the duration of the automatically identified surge for the eight winters from 2005-2012. The best linear fit is shown by the blue line.

Figure 3 - Applying the automatic identification method (described in to the winter of 2013/4 to identify the start of the winter surge. This data was not used to choose window size, band width or consistency requirements. The solid red lines show the upper and lower Bollinger bands. The red dashed line shows the automatically identified start of the surge and the grey dashed line the predicted end of the surge.

Figure 4 - Applying the new automatic identification method to the winter of 2013/4 to identify the end of the winter surge. The red dashed lines shows the identified start and end of the surge and the grey dashed line shows the end of the surge predicted at the start of the surge (19 November 2013).

Table 1 - Number of calls by type of call received by Children's Acute Transport Service (CATS) and outcome of call between 1 April 2005 and 15 July 2013. The final column indicates which calls were considered to represent "demand for retrieval".

Outcome of call	Frequency (% of calls)	Counts as demand for retrieval?
CATS team deployed	9337 (53%)	Yes
Transfer request refused due to no CATS team or PICU bed	394 (2%)	Yes
Transfer request refused but not due to capacity constraint	2434 (14%)	No
Call cancelled by referrer	1397 (8%)	No
Child died before team deployed	119 (1%)	No
Courtesy call	488 (3%)	No
Advice given	3343 (19%)	No
Unknown	15 (0%)	No
Total number of calls	**17527 (100%)**	
Total demand for retrieval	**9731 (56%)**	

Table 2 –Date for start and finish of the winter surge period as identified by the Bollinger bands method. The squared difference is defined as the square of the difference in days between the automatic and manually identified dates. So the first row, there are 3 days difference between the automatically signalled start of 7 November 2005 and the manually identified start of 4th November and this is then squared to give a value of 9.

Year	Manually Identified start	Automatically signalled start	Squared difference for start	Manually identified end	Automatically signalled end	Squared difference for end
2005/6	04-Nov-05	07-Nov-05	9	10-Jan-06	14-Jan-06	16
2006/7	18-Nov-06	23-Nov-06	25	05-Jan-07	08-Jan-07	9
2007/8	25-Nov-07	26-Nov-07	1	03-Jan-08	06-Jan-08	9
2008/9	04-Nov-08	07-Nov-08	9	24-Dec-08	26-Dec-08	4
2009/10	19-Nov-09	20-Nov-09	1	23-Dec-09	29-Dec-09	36
2010/11	08-Oct-10	08-Oct-10	0	11-Jan-11	12-Jan-11	1
2011/12	23-Oct-11	24-Oct-11	1	10-Jan-12	15-Jan-12	25
2012/13	20-Oct-12	21-Oct-12	1	26-Jan-13	19-Jan-13	49
Total			47			75

Figure 1

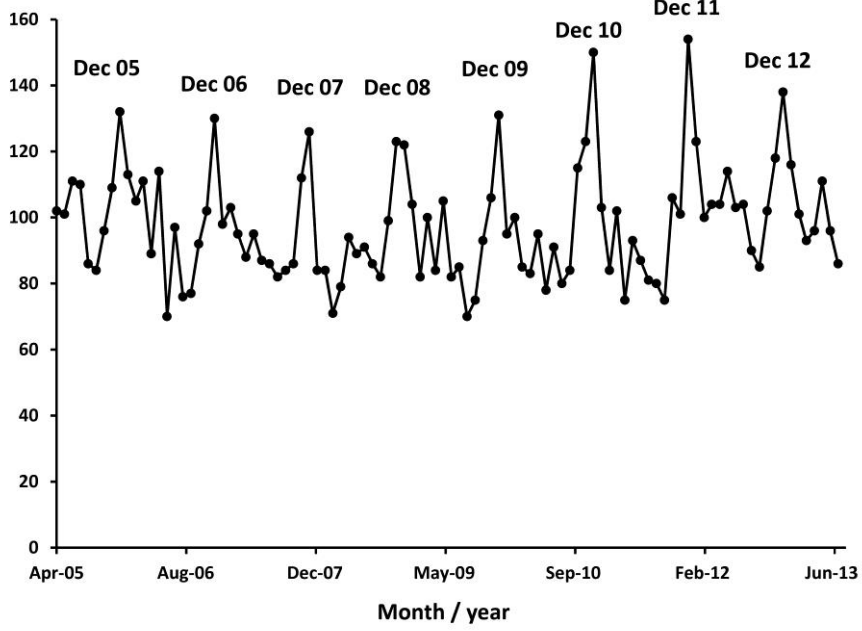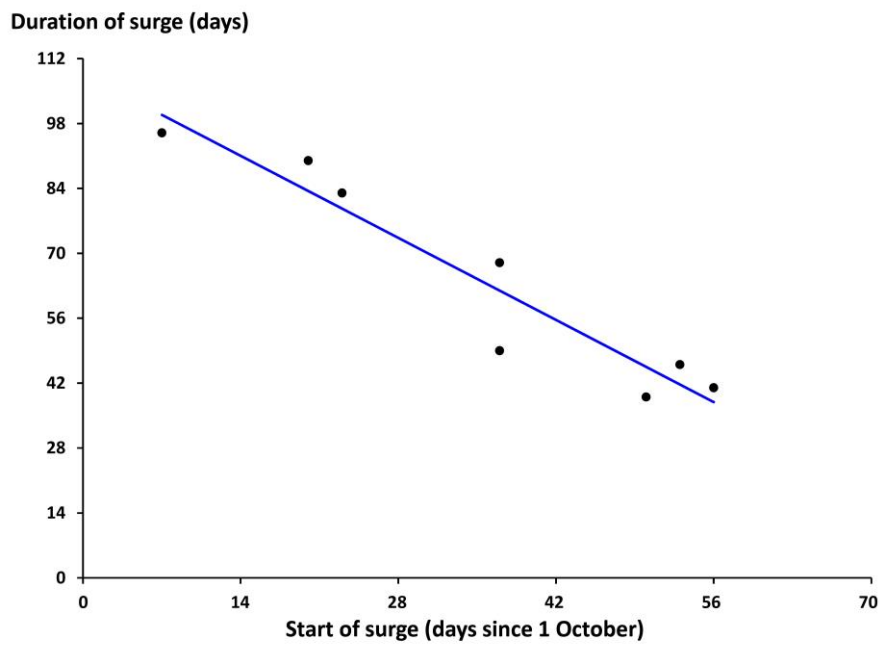Total monthly demand for retrievals

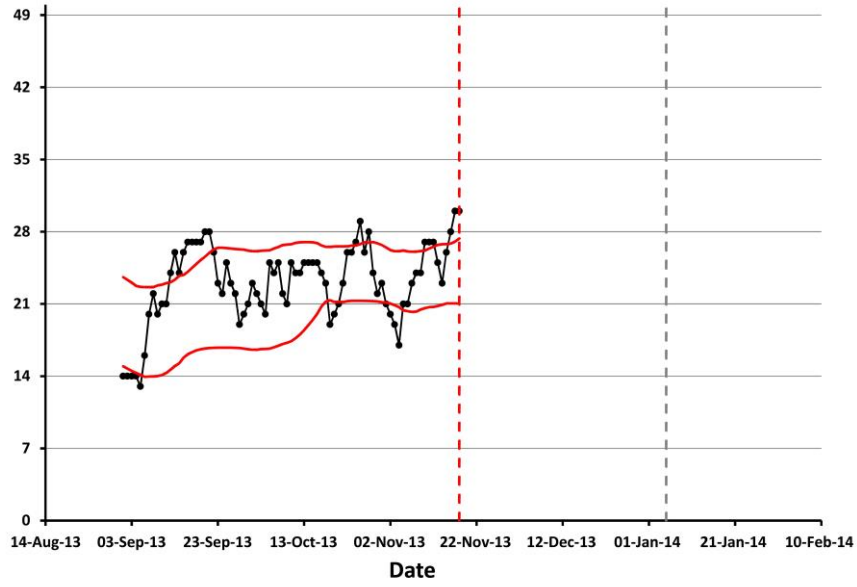**Figure 2**

Figure 3

Running 7-day total demand
for retrieval

Figure 4

Running 7-day total demand for retrieval