# Anomaly Detection in Clinical Data of Patients Undergoing Heart Surgery

Presbitero, A.; Quax, R.; Krzhizhanovskaya, V.; Sloot, P.

**DOI**
10.1016/j.procs.2017.05.002

**Publication date**
2017

**Document Version**
Final published version

**Published in**
Procedia Computer Science

Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 99–108

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

# Anomaly Detection in Clinical Data of Patients Undergoing Heart Surgery

Alva Presbitero[1*], Rick Quax[2], Valeria Krzhizhanovskaya[1,2] and Peter Sloot[1,2,3]

[1]*ITMO University, St. Petersburg.*

[2]*University of Amsterdam, The Netherlands*

[3]*Nanyang Technological University, Singapore*

[*] avpresbitero@gmail.com, r.quax@uva.nl, V.Krzhizhanovskaya@uva.nl, p.m.a.sloot@uva.nl

**Abstract**

We describe two approaches to detecting anomalies in time series of multi-parameter clinical data: (1) *metric and model-based indicators* and (2) *information surprise*. (1) *Metric and model-based indicators* are commonly used as early warning signals to detect transitions between alternate states based on individual time series. Here we explore the applicability of existing indicators to distinguish critical (anomalies) from non-critical conditions in patients undergoing cardiac surgery, based on a small anonymized clinical trial dataset. We find that a combination of time-varying autoregressive model, kurtosis, and skewness indicators correctly distinguished critical from non-critical patients in 5 out of 36 blood parameters at a window size of 0.3 (average of 37 hours) or higher. (2) *Information surprise* quantifies how the progression of one patient's condition differs from that of rest of the population based on the cross-section of time series. With the maximum surprise and slope features we detect all critical patients at the 0.05 significance level. Moreover we show that a naive outlier detection does not work, demonstrating the need for the more sophisticated approaches explored here. Our preliminary results suggest that future developments in early warning systems for patient condition monitoring may predict the onset of critical transition and allow medical intervention preventing patient death. Further method development is needed to avoid overfitting and spurious results, and verification on large clinical datasets.



Peer-review under responsibility of the scientific committee of the International Conference on Computational Science

*Keywords:* Anomaly Detection, Information Surprise, Early Warning Signal, Time series Analysis

## 1 Introduction

Anomaly detection is the process of pinpointing and segregating items in a population exhibiting behaviors that deviate from the norm. These are referred to as “anomalies” or “outliers”. Anomaly detection is extensively used in detecting fraud credit cards [1], cyber security [2, 3], health insurance [4], and patient monitoring using electrocardiography (ECG) signals [5]. For instance, technologies in

1877-0509 
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science
10.1016/j.procs.2017.05.002

anomaly detection, especially for medical applications, are essential for estimating physical conditions or states of patients from health to demise. Often times, when anomalies occur, significant changes in time series patterns become evident. These anomalies (critical conditions) designated by pattern changes could serve as indicators of transitions from healthy to critical state that leads to death in patients.

Numerous complex dynamical systems have been found to exhibit transitions or tipping points where systems abruptly shift from one stable state to another. The specific case of disease can be regarded as a sudden shift in system state from health to disease [6, 7]. For instance, the onset of depression is explored by looking at fluctuations of emotions as indicators of transition from a normal to a state of depression [8]. Other examples include financial systems, which exhibit systemic market crashes [9, 10], climate shifts preceded by the slowing down of fluctuations [11–13], decline of population leading to extinction [14–16], flood early warning systems [17–19] and dams [20].

Early warning signals (EWS) are used as indicators for loss of system resilience prior to transitions based on more subtle statistical properties of the measurements [21]. This can sometimes be detected through changes in correlations, standard deviation, or skewness of the series through time [22]. We utilize indicators used in EWS to segregate critical from non-critical patients with the assumption that critical patients exhibit pattern changes in their time series when anomalies or transitions from health to demise occur. We do not detect how much time these transitions occur in advance. We aim to incorporate this in the future version of our work.

In the present work we explore the applicability of using four classical EWSs on blood parameter concentration time series from 53 patients undergoing complicated cardiac surgery to detect the transitions of patient death; an approach that has not been done in literature before. The most important motivation of using EWS is its potential of real-time use as early warning for increased risk of patient death, with eventually the goal of improved prevention.

# 2 Data Preparation and Analysis

The raw data consists of concentrations of 36 different blood parameters from 53 patients (total of 878 sample data points) undergoing complicated cardiac surgery (including timestamps), three of whom died after the operation [23]. Patients are composed of male or non-pregnant, non-lactating female of any race with an age over 18. The period to which the data was collected comprised a time interval of 24 hours prior to the time of surgery up until 30 days after surgery. Several blood samples were usually taken within 24 hours after the surgery as this time period is the most critical for patients who have recently undergone surgery. Patients then usually stabilize after this point so the rate at which blood is sampled is reduced.

Raw data also contains a substantial fraction of missing data points of 62.5% because not all 36 parameters were always tested for in all blood samples. However the missing values are reasonably well distributed over the parameters: the 95% confidence interval (CI) of missing values is 43.8%-88.8%. Missing data will be inevitable in clinical trial data so any EWS method must be capable of dealing with it. There are various techniques to deal with missing values but most importantly the technique should not significantly increase the rate of false negatives (labeling healthy patients as critical) because this would render the signal noisy and make it impractical for medical practitioners to act upon it.

Figure 1 shows a sample of the bootstrapped raw data for IL10, one of the blood parameters in the raw data. Red corresponds to the critical patients while blue corresponds to non-critical patients. Increasing concentrations of IL10, a type of anti-inflammatory cytokine that aids the human body against foreign attacks, in patients usually indicate the presence of inflammation. We see an increasing concentration for critical patients. However, we note that we only have three samples for the critical patients, hence bootstrapping is only limited to at most three points per time step.

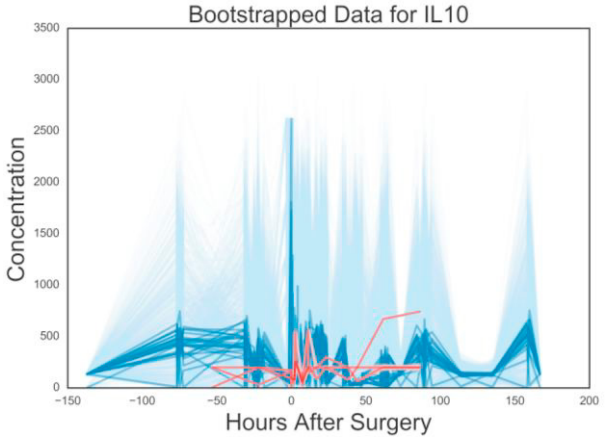


**Figure 1.** Bootstrapped data for blood parameter IL10. Dark red corresponds to the median concentrations of critical patients while dark blue corresponds to non-critical patients in the placebo treatment. Light red and light blue corresponds to the actual concentrations of critical and non-critical patients respectively.

Data collection was done separately in two different hospitals namely Catharina Hospital Eindhoven (The Netherlands), and Zuid Oost-Limburg Hospital (Belgium). The conditions to which the patients were exposed to safety procedures, methods used to obtain the measurements, and time intervals chosen for data collection were standardized across the two hospitals.

## 2.1 Data Preparation

The data contains missing values because sometimes only a subset of the variables was measured for a patient. We use a non-parametric resampling (bootstrap) approach to replace a missing value for a particular variable by a randomly selected value from all other patients who do have a recorded value for that variable.

*Metric and Model-based Indicators.* We interpolate the bootstrapped data to assure regularity in the time series prior to using metric-based indicators (*standard deviation, kurtosis, and skewness*). Non-stationary in the mean of the time series may indicate the onset of transitions. Furthermore, seasonal periodicities present in the time series may lead to the detection of strong correlation [22]. In order to remove trends and filter out high frequencies in the time series prior to applying metric-based indicators in EWS, we utilized a Gaussian filter with standard deviation a sixth the size of the data for smoothing out the time series. Then we subtract the filtered time series to the original record to obtain the residual time series. We did not use filtering prior to the application of the model-based indicator (*time-varying autoregressive model*) since we need these variations in the data for autocorrelation to work.

*Information Surprise.* This procedure is chosen because it prevents overestimating the surprise: if some measurements of a patient are equal to that of some other patients then the surprise will tend to be lower compared to substituting a non-existent value by extrapolation for instance. Indeed, in the hypothetical limit of replacing all measurements of a patient the surprise value will tend towards the entropy (average surprise) of the sample points in the population. Consequently, if a patient is found to have a high surprise value then we can be reasonably confident that it measures a property of the patient and not of our algorithm. We repeat the bootstrap procedure 50 times in order to account for variability induced by computing averages and standard errors for all subsequent calculations.

# 3 Metric and Model-Based Indicators

Leading indicators for transitions can either be metric or model-based indicators. Both methods quantify changes in memory or correlation structure, and variability of time series as the system transitions between alternate regimes [22].

Metric-based indicators quantify changes in statistical properties of time series without attempting to fit the data to a specific model. We utilize *standard deviation*, *skewness*, and *kurtosis* as indicators of patient state transitions. *Standard deviation.* When a system is close to a tipping point, the rate at which the system returns back to equilibrium slows down. A phenomenon that causes the system state to drift to and from boundaries of alternative states, also called flickering, is observed especially upon exposure to greater disturbances. The combination of slowing down and flickering, leads to an increase in variations [24]. *Skewness.* Perturbations may push the state of the system to values close to the boundary that separates alternative states. Slowing down of the return rate of the system towards equilibrium results in asymmetry of the distribution of the time series [25]. *Skewness* either increases or decreases depending on the direction of transition. *Kurtosis.* Strong perturbations resulting in flickering pushes the system to reach extreme values that are close to transition, which increases the occurrence of rare values in the time series [26]. This in turn leads to the increase of *kurtosis*, or "bulging," of the time series prior to the tipping point.

Model-based indicators capture changes in time series quantitatively by attempting at fitting the data to a model. Autocorrelation presents a simple method to quantitatively describe how a system slows down. For instance, increasing autocorrelation implies that consecutive points in the time series have become increasingly similar. *Time-varying Autoregressive models (AR)* with $p$ time lags provide ways to estimate the local dynamics in a time series [27]. This is done by determining the inverse of the characteristic root ($\lambda$), through estimating autoregressive. Values of $\lambda$ approaching 0 imply that the system returns quickly to the mean while values approaching 1 imply instability. We use a time lag equal to one to indicate that the current value is based on the value immediately preceding it. The equation for time-varying AR(1) model is given in equation (**1**).

$$y(t) = a(t)y(t-1) + \varepsilon(t) \quad (1)$$

where $a(t)$ corresponds to the autoregressive coefficient, and $\varepsilon(t)$ corresponds to the environmental variability [22].

**Detection Test for Metric and Model-Based Indicators**

Increasing or decreasing trends detected by leading indicators are evaluated via Mann-Kendall trend test. The Mann-Kendall trend test is a non-parametric test used to analyze data series for consistent increasing or decreasing patterns. Null hypothesis assumes that a monotonic trend in the series does not exist, while the alternate hypothesis assumes that a trend exists. We test the strength of the trends to a significance level of 5% (one-sided hypothesis test). We only consider positive trends as these correspond to increasing kurtosis, skewness, and standard deviations, which indicate transition in system state.

**Evaluation of Metric and Model-Based Indicators**

We evaluate the performance of leading indicators in detecting critical and non-critical patients, we calculate precision and recall with definitions summarized in Table 1.

**Table 1.** Definition of terms used for assigning critical and non-critical patients.

| Symbol | Interpretation | Definition |
|---|---|---|
| $T_P$ | True Positive | assigning critical patients as ***critical*** |
| $F_P$ | False Positive | assigning non-critical patients as ***critical*** |
| $F_N$ | False Negative | assigning critical patients as ***non-critical*** |
| $T_N$ | True Negative | assigning non-critical patients as ***non-critical*** |

Precision ($P$) provides a measure of the relevance of the results. It is formally defined as in equation (**2**).

$$P = \frac{T_P}{T_P + F_P} \tag{2}$$

While Recall ($R$) measures the fraction of relevant instances retrieved.

$$R = \frac{T_P}{T_P + F_N} \tag{3}$$

High precision implies the unlikely occurrence of non-critical patients being labeled critical (false positives), while high recall indicates low occurrence of critical patients being labeled as non-critical (false negatives). A system having high precision but low recall reports less number of critical patients, but most of these patients are labeled correctly. A system having low precision but high recall, on the other hand, reports a greater number of critical patients, but most of these patients are labeled incorrectly.

We also look at Accuracy(A), which measures the proportion of correctly predicted observations in the population.

$$A = \frac{T_P + T_N}{T_P + T_N + F_P + F_N} \tag{4}$$

# 4 Information Surprise Method

The previous measures were developed for individual time series. This makes sense for example for climate data such as global CO2 level over time of which only one time series exists. However our dataset consists of a cross-section of time series, namely patients who underwent the same surgery divided into a placebo, treated, and deceased group. This provides us the opportunity to compare one patient's time series to that of the other patients and compute an intrinsic measure of how much a patient's state deviates over time from 'what is expected'.

For each sample point $x_i^t$ of patient $i$ at time $t$ we can compute its "surprise" [28].

$$\sigma(x_i^t) = log_2 \frac{1}{p(x^t)} \tag{5}$$

Here, $p(x^t)$ is a probability distribution of a sample point at time $t$ for a randomly selected patient out of a patient population $X$. Surprise is a fundamental concept in information theory, e.g., the expected surprise $\langle \sigma(x_i^t) \rangle_i$ of the population $X$ itself equals its 'entropy' which can intuitively be understood as a measure of variation capable of handling also non-linear and categorical data (unlike standard deviation).

We take as population $X$ all patients who survived the procedure. Then for each patient $i$ we estimate the probability distribution $p_i(x^t)$ from the data by means of a standard kernel density estimation algorithm whose bandwidth parameter $b = 0.83$ is chosen such that maximizes the $k$-fold cross-validation ($k = 3$). The subscript $i$ indicates that patient $i$ is left out of the population in this estimation, simulating the case where patient $i$ would be a new patient who is to be compared to the database of previous patients.

A low surprise value indicates that a patient's blood measurements are 'as expected', or within the variation of patients who undergone the same surgery. A high value indicates that a patient's measurements deviate from the expectation. Since all patients in the population survived the procedure a high surprise could indicate the onset of critical, unstable conditions which lead to death. On the other hand, high surprise values could also have other causes such as rare co-morbidities, severe complications, or even an extraordinarily fast recovery of health. Nevertheless, it may still pay off to further investigate high surprise patients on the whole to prevent patient deaths, at the expense of occasionally investigating healthy patients (false positive).

**Detection Test for Information Surprise**

We detect "critical" patients based on their surprise curve over time by means of a non-parametric statistical hypothesis test with $p$-value $p = 0.05$ (one-sided). All measurements of survived patients form the nulls hypothesis distributions. As features of the surprise curve we explore both the maximum surprise value as well as the slope of the surprise curve.

# 5 Results and Discussion

## 5.1 Anomaly Detection by Metric and Model-Based Indicators

We look at the sensitivity of the methods used with respect to varying sliding window widths by looking at the precision-recall scatter plots, and accuracy. We set the window frame widths in increments of 0.05 within the interval [0.1,0.5] and test the strength of these trends by comparing $p$-values with significance level $p = 0.05$. Note that we are dealing with a total number of 936 hypotheses at a time. The one-sided hypothesis test works well when a single hypothesis is tested. If we are dealing, however, with multiple hypotheses simultaneously, this in effect leads to the increase in the probability of detecting false positives. In order to account for the numerous tests that we are doing, we corrected the $p$-values through multiple hypothesis testing. We used two known methods namely Familywise error rate (FWER) and False discovery rate (FDR). For FWER, we used Bonferroni procedure, one of the classical solutions for multiple hypothesis testing corrections. For FDR, we used the Benjamini–Hochberg procedure. We find that both result in the same p-value corrections.

Each point on the left-hand side of Figure 2 corresponds to the recall-precision of one method and window size pair for *all* parameters based on corrected p-values via multiple hypothesis testing. We see that as the window size increases, Time-Varying AR, Skewness and Kurtosis exhibit slow increase in precision, but relatively fast increase in recall. This implies that values slowly become asymmetrical and extreme values appear at a slower rate in critical patients. Accuracy (see right-hand side of Figure 2) decreases with increasing window frame width. This means that proportion of predicted critical and non-critical patients is indirectly proportional to the window frame width. It also seems that accuracy becomes relatively stable starting from window size 0.3 to 0.5. Although the number of correctly labeled critical patients increases with increasing timeframe width (indicated by the increase in recall over increasing window width), this effect is wiped out by the decrease in correctly labeled non-critical patients (indicated by the decrease in accuracy over increasing window width) especially since critical patients are much less than non-critical patients. This suggests that values in non-critical patients slowly become asymmetrical and extreme values appear at a slower rate much like in critical patients, which allow detection of trends in bigger window frames. Standard deviation, on the other hand, increases slowly implying that the distribution of values tend to go farther away from the mean at an almost steady rate for both critical and non-critical patients, which is also apparent in the steadiness in accuracy. Although we highlight kurtosis at window size 0.5 (largest area under the PR curve), our results seem to be sensitive towards sliding window widths and the type of leading indicators used.

Our results generally show low precision but medium to high recall. Leading indicators are detecting more critical patients than there actually are in the population. We conjecture that this might be the consequence of our initial assumption stating that critical patients exhibit transitions in *all* blood parameters, while non-critical patients do not. We investigate this further by looking at the performance of specific parameters in terms of recall and precision values. Figure 3 shows that a combination of leading indicators and window frame sizes was able to distinguish critical from non-critical patients in 5 out of 36 blood parameters. We get the same number of indicating parameters if we only consider those where accuracies (window sizes greater than or equal to 0.3) are stable. Frequency distributions corresponding to parameters exhibiting 100% precision and recall for window

sizes (white vertical bars) and methods (colored horizontal bars) are summarized on the right of Figure 3. Our results suggest that these 5 key parameters could be used by medical practitioners to determine which among the patients are critical.

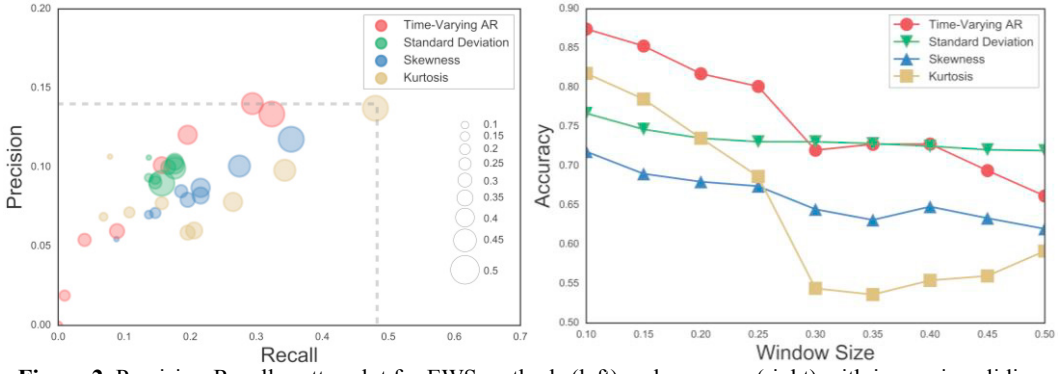


**Figure 2.** Precision-Recall scatter plot for EWS methods (left) and accuracy (right) with increasing sliding window widths for all parameters. Different colors correspond to the EWS method used. Point size increases with increasing sliding window width ranging from [0.1,0.5]. On the left we highlight Kurtosis at window size 0.5, which shows the largest area under the curve (precision= 0.14 and recall=0.48).

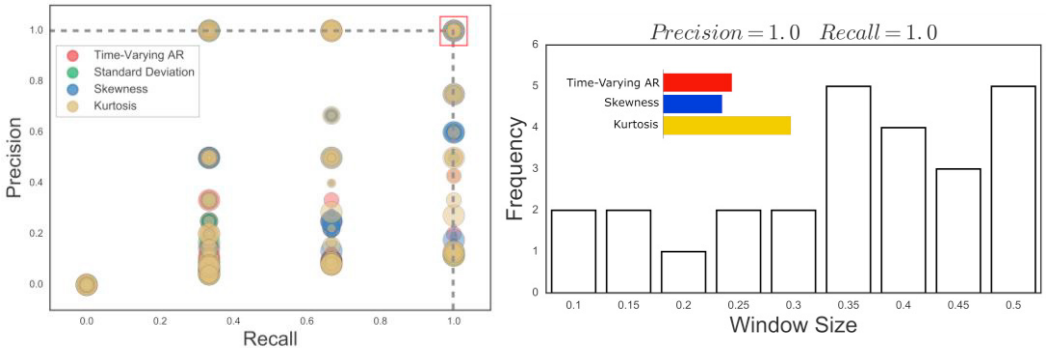


**Figure 3.** Precision-Recall scatter plot for each blood parameter (left) and frequency distributions of window frame widths (white vertical bars) and methods (colored horizontal bars) at 100% precision and recall (right).

One could argue that the results presented could just be spurious especially given the nature of the raw data. Hence, any random assumption of criticality in patients, regardless of whether the patient is indeed critical or not, would still give out a number of indicators (blood parameters with 100% precision and recall). We generated 1150 non-critical patients by bootstrapping from the raw population of non-critical patients and tested their blood parameters for transitions using the leading indicators presented. We find that 93.5% of the population of non-critical patients exhibited non-trends (absence of transitions), hence non-anomalies and correctly labeled as non-critical patients. This implies that our experiment is relatively statistically sound, although ideally, we would want a 95% detection of non-criticality in non-critical patients.

## 5.2 Information Surprise as an Indicator for Patient Demise

The surprise curve for each patient is shown in Figure 4**.** Despite variation due to data sparseness it is easy to spot at least two deceased patients (16, 38) whose surprise contains high peaks at about 1 and 3.5 days after surgery. The third deceased patient (13) has no significant peak but instead has a significant upward trend. Additionally it is evident that some other peaks in surprise were due to patients who were reported to have developed multiple complications of which at least one severe (dashed lines), such as defibrillation or internal bleeding. Although eventually these patients survived

the procedure their high peaks may still be considered meaningful and worthwhile acting upon, however we will restrict ourselves to detecting deceased patients only.

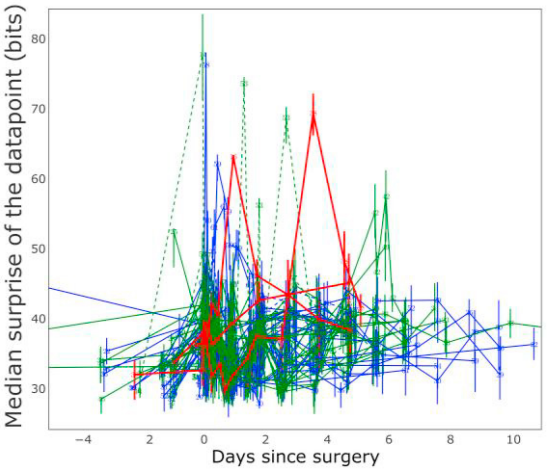


**Figure 4.** The median surprise curves per patient (blue=placebo, green=treated, red=deceased; dashed=severe complications) for the combined dataset. Vertical error bars indicate the standard error of the mean (SEM) at each point.

In Figure 5 we show the histograms of the maximum surprise value as well as the slope of a linear regression of the surprise curves for the placebo, treated, and deceased patients. Although the histograms for deceased patients were constructed from only three data points, this figure suggests that deceased patients may be discriminated reasonably well based on these two features. For the maximum surprise feature the placebo and treated patients appear to follow a right-tailed distribution with a strong peak near 40; the deceased patients however all appear in the tail of this distribution. For the slope feature it is apparent that the placebo and treated patients are distributed around zero whereas the deceased patients all have significant positive slopes.

Indeed we find that these two surprise features would detect all three deceased patients at the 0.05 significance level. The maximum surprise 95th percentile of the survived patients equals 61.71 and the maximum surprises of the deceased patients 13, 16, and 38 equal 47.27, 65.64, and 67.49, respectively, meaning that patients 16 and 38 would be flagged as critical. For the slope feature the 95th percentile equals 1.15 and the deceased patients' values equal 1.69, 0.52, and 2.80, respectively, meaning that patients 13 and 38 would be flagged as critical. Either measure thus detects 2 out of 3 deceased patients and can therefore be said to be relatively successful. On top of this, by combining the two tests naively by an 'either, or' condition all three deceased patients would be detected, leading to a true positive score of 100% in this case. The false positive percentage of either individual test is necessarily 5% due to the significance level; for the combined test it will necessarily lie between 5% and 10% depending on the overlap of the two sets of detected patients. These results suggest that our surprise method may be a viable way to detect critical patients with relatively high precision.

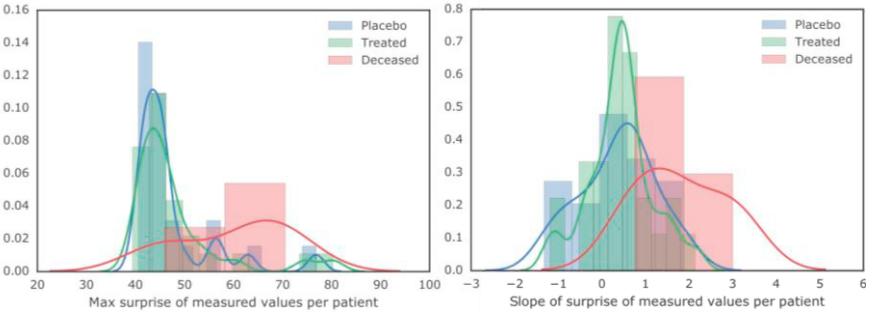


**Figure 5.** Histograms of maximum surprise value (left) and the slope of a linear regression of the surprise curves (right) for the placebo (blue), treated (green), and deceased (red) patients. The histograms shown fit a kernel density estimate.

## 5.3 Comparison to Naive Euclidean Distance Approach

All 878 sample points of all patients form a point cloud in 29-dimensional space. One could hypothesize that critical patients' simply have strongly deviating measurement values for some of the variables, leading to these patients becoming outliers of this point cloud. In this case calculating the surprise measure would be superfluous as critical patients could already be detected by measuring their Euclidean distance to the center of the point cloud formed by the survived patients.

We show in Figure 6 that distribution of distances of deceased patients are not distinguishable from the other patients. In fact, one-sided statistical hypothesis testing ($p = 0.05$) leads to detecting only 2.66% (SEM 0.32) of the sample points of deceased patients, versus 4.21% (SEM 0.10) and 5.63% (SEM 0.080) of the placebo and treated patients, respectively. This result suggests that the point cloud is far from having a uniform spherical shape, which is implicitly assumed by comparing Euclidean distances. This means that criticality of patients cannot be detected by a naive outlier detection of the measurement vectors themselves, and indeed a more general approach such as surprise is needed.

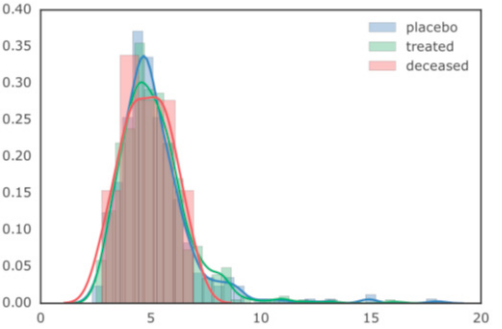


**Figure 6.** Euclidean distance of 29-dimensional sample points to the center of non-critical patients.

# 6 Conclusion

We utilized leading indicators commonly used in early warning signals (EWS) and information surprise as indicators for criticalities in patients where we showed that a combination of the leading indicators (Time-Varying Autoregressive Model, kurtosis, and skewness) was able to perfectly distinguish critical from non-critical patients in 5 out of 36 blood parameters in stable accuracy at window sizes greater than or equal to 0.3 (37 hours). Our results imply that the 5 parameters detected could serve as indicators of patient demise through the presence of transitions. However, we realize the potential for overfitting and spurious results. The future version of our work will mitigate these effects by utilizing a calibrated model like the innate immune system model [29]. We introduced two surprise features namely maximum surprise and slope, where we detected all three critical patients at the 0.05 significance level. The false positive (labeling non-critical patients as critical) percentages lie between 5% and 10%, which suggests that our surprise method could be applicable to detect critical patients with high precision. More importantly, we also showed that the distribution of the critical patients is not distinguishable from the non-critical patients. Therefore, critical patients cannot be detected by a naive outlier detection, and indeed a more general approach such as surprise is needed.

**Acknowledgements**. We thank Ruud Brands of Alloksys Life Sciences BV for providing the data. This research is financially supported by the Russian Science Foundation, Agreement #14-11-00823 (15.07.2014).

# References

1. Bolton, R.J., Hand, D.J., H, D.J.: Unsupervised Profiling Methods for Fraud Detection. Proc. Credit Scoring Credit Control VII. 5–7 (2001).
2. Ten, C.W., Hong, J., Liu, C.C.: Anomaly detection for cybersecurity of the substations. IEEE Trans. Smart Grid. 2, 865–873 (2011).
3. Chandola, V., Eilertson, E., Ert, L.: Data Mining for Cyber Security. Data Warehous. Data Min. Tech. Comput. Secur. 2 (2006).
4. Kirlidog, M., Asuk, C.: A Fraud Detection Approach with Data Mining in Health Insurance. Procedia - Soc.

Behav. Sci. 62, 989–994 (2012).
5. Keogh, E., Lin, J., Fu, A.W., Van Herle, H.: Finding unusual medical time-series subsequences: Algorithms and applications. IEEE Trans. Inf. Technol. Biomed. 10, 429–439 (2006).
6. Trefois, C., Antony, P.M.A., Goncalves, J., Skupin, A., Balling, R.: Critical transitions in chronic disease: Transferring concepts from ecology to systems medicine, (2015).
7. Liu, R., Yu, X., Liu, X., Xu, D., Aihara, K., Chen, L.: Identifying critical transitions of complex diseases based on a single sample. Bioinformatics. 30, 1579–86 (2014).
8. van de Leemput, I. a, Wichers, M., Cramer, A.O.J., Borsboom, D., Tuerlinckx, F., Kuppens, P., van Nes, E.H., Viechtbauer, W., Giltay, E.J., Aggen, S.H., Derom, C., Jacobs, N., Kendler, K.S., van der Maas, H.L.J., Neale, M.C., Peeters, F., Thiery, E., Zachar, P., Scheffer, M.: Critical slowing down as early warning for the onset and termination of depression. Proc. Natl. Acad. Sci. U. S. A. 111, 87–92 (2014).
9. May, R.M., Levin, S. a, Sugihara, G.: Complex systems: ecology for bankers. Nature. 451, 893–895 (2008).
10. Quax, R., Kandhai, D., Sloot, P.M. a: Information dissipation as an early-warning signal for the Lehman Brothers collapse in financial time series. Sci. Rep. 3, 1898 (2013).
11. Dakos, V., Scheffer, M., van Nes, E.H., Brovkin, V., Petoukhov, V., Held, H.: Slowing down as an early warning signal for abrupt climate change. Proc. Natl. Acad. Sci. U. S. A. 105, 14308–12 (2008).
12. Lenton, T.M., Livina, V.N., Dakos, V., van Nes, E.H., Scheffer, M.: Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. Philos. Trans. R. Soc. A Math. Phys. Eng. Sci. 370, 1185–1204 (2012).
13. Lenton, T.M., Held, H., Kriegler, E., Hall, J.W., Lucht, W., Rahmstorf, S., Schellnhuber, H.J.: Tipping elements in the Earth's climate system. Proc. Natl. Acad. Sci. 105, 1786–1793 (2008).
14. Krkošek, M., Drake, J.M.: On signals of phase transitions in salmon population dynamics. Proc. R. Soc. B Biol. Sci. 281, 20133221 (2014).
15. Clements, C.F., Ozgul, A.: Including trait-based early warning signals helps predict population collapse. Nat. Commun. 7, 10984 (2016).
16. Drake, J.M., Griffen, B.D.: Early warning signals of extinction in deteriorating environments. Nature. 467, 456–459 (2010).
17. Pyayt, A.L., Shevchenko, D.V., Kozionov, A.P., Mokhov, I.I., Lang, B., Krzhizhanovskaya, V.V., Sloot, P.M.A.: Combining Data-Driven Methods with Finite Element Analysis for Flood Early Warning Systems. Procedia Comput. Sci. 51, 2347–2356 (2015).
18. Krzhizhanovskaya, V. V., Shirshov, G.S., Melnikova, N.B., Belleman, R.G., Rusadi, F.I., Broekhuijsen, B.J., Gouldby, B.P., Lhomme, J., Balis, B., Bubak, M., Pyayt, A.L., Mokhov, I.I., Ozhigin, A. V., Lang, B., Meijer, R.J.: Flood early warning system: Design, implementation and computational modules. In: Procedia Computer Science. pp. 106–115 (2011).
19. Melnikova, N.B., Jordan, D., Krzhizhanovskaya, V. V.: Experience of using FEM for real-time flood early warning systems: Monitoring and modeling Boston levee instability. J. Comput. Sci. 10, 13–25 (2015).
20. Fisher, W.D., Camp, T.K., Krzhizhanovskaya, V. V.: Crack detection in earth dam and levee passive seismic data using support vector machines. Procedia Comput. Sci. 80, 577–586 (2016).
21. DeAngelis, D.L.: Energy flow, nutrient cycling, and ecosystem resilience. Ecology. 61, 764–771 (1980).
22. Dakos, V., Carpenter, S.R., Brock, W.A., Ellison, A.M., Guttal, V., Ives, A.R., Kéfi, S., Livina, V., Seekell, D.A., van Nes, E.H., Scheffer, M.: Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. PLoS One. 7, (2012).
23. Kats, S., Brands, R., Hamad, M.A.S., Seinen, W., Scharnhorst, V., Wulkan, R.W., Sch??nberger, J.P., van Oeveren, W.: Prophylactic treatment with alkaline phosphatase in cardiac surgery induces endogenous alkaline phosphatase release. Int. J. Artif. Organs. 35, 144–151 (2012).
24. Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., van Nes, E.H., Rietkerk, M., Sugihara, G.: Early-warning signals for critical transitions. Nature. 461, 53–59 (2009).
25. Guttal, V., Jayaprakash, C.: Changing skewness: An early warning signal of regime shifts in ecosystems. Ecol. Lett. 11, 450–460 (2008).
26. Biggs, R., Carpenter, S.R., Brock, W.A.: Turning back from the brink: detecting an impending regime shift in time to avert it. Proc. Natl. Acad. Sci. U. S. A. 106, 826–31 (2009).
27. Ives, A.R., Dakos, V.: Detecting dynamical changes in nonlinear time series using locally linear state-space models. Ecosphere. 3, art58 (2012).
28. Cover, T.M., Thomas, J.A.: Elements of Information Theory 2nd Edition. (2006).
29. Presbitero, A., Krzhizhanovskaya, V., Mancini, E., Brands, R., Sloot, P.: Immune System Model Calibration by Genetic Algorithm. Procedia Comput. Sci. 101, 161–171 (2016).