

Exploiting time in electronic health record correlations

George Hripcsak, David J Albers, Adler Perotte

► Additional materials are published online only. To view these files please visit the journal online (www.jamia.org/content/18/Suppl_1.toc).

Department of Biomedical Informatics, Columbia University, New York, New York, USA

Correspondence to

Dr George Hripcsak, Department of Biomedical Informatics, Columbia University Medical Center, 622 West 168th Street, VC5, New York, NY 10032, USA; hripcsak@columbia.edu

Received 1 July 2011
Accepted 29 October 2011
Published Online First
23 November 2011

ABSTRACT

Objective To demonstrate that a large, heterogeneous clinical database can reveal fine temporal patterns in clinical associations; to illustrate several types of associations; and to ascertain the value of exploiting time.

Materials and methods Lagged linear correlation was calculated between seven clinical laboratory values and 30 clinical concepts extracted from resident signout notes from a 22-year, 3-million-patient database of electronic health records. Time points were interpolated, and patients were normalized to reduce inter-patient effects.

Results The method revealed several types of associations with detailed temporal patterns. Definitional associations included low blood potassium preceding 'hypokalemia.' Low potassium preceding the drug spironolactone with high potassium following spironolactone exemplified intentional and physiologic associations, respectively. Counterintuitive results such as the fact that diseases appeared to follow their effects may be due to the workflow of healthcare, in which clinical findings precede the clinician's diagnosis of a disease even though the disease actually preceded the findings. Fully exploiting time by interpolating time points produced less noisy results.

Discussion Electronic health records are not direct reflections of the patient state, but rather reflections of the healthcare process and the recording process. With proper techniques and understanding, and with proper incorporation of time, interpretable associations can be derived from a large clinical database.

Conclusion A large, heterogeneous clinical database can reveal clinical associations, time is an important feature, and care must be taken to interpret the results.

INTRODUCTION

With the national push for the adoption of electronic health records,¹ health record data are becoming increasingly available and have the potential to facilitate clinical research.² Electronic health records come with their own challenges, however. The data do not simply represent patient physiology but also represent clinical processes and workflow, implying that there may be several types of associations and causes. There may be a physiologic basis for a change, or a change may occur because a human being—based on seeing a clinical picture—decided to take an action. For example, a clinical finding may prompt a physician to order a drug, which should then correct the finding. Teasing out these distinctions can benefit from not just the presence but also the timing of relationships.

Several different methods have been developed to uncover temporal properties and relationships among clinical variables, including temporal

abstraction,^{3–16} causal inference,^{17–24} and numeric time series analysis.^{25–28} Nevertheless, many recent large-scale studies of electronic health record data do not exploit time in any detail.^{29–34} A relatively simple method, which we use in this study, is to measure linear correlation between co-occurrences of pairs of variables, lagging one variable with respect to the other to assess the change in correlation as variables are shifted in time.³⁵ Such a method may require a fairly large database to estimate correlation, but it can be applied to many problems in a generic way, and, because it does not rely on previous domain or temporal knowledge, it may reveal unexpected types of temporal relationships.

A second challenge with using electronic health record data instead of a research-specific study cohort is that the electronic health record patients—whose appearance in the database is unpredictable—tend to be heterogeneous, potentially increasing the variance of measured parameters and increasing opportunities for bias. One approach that may ameliorate this effect is to reduce inter-patient effects so that patients serve as their own controls. In the case of linear correlation, the effects can be reduced by normalizing each patient's data (eg, to standard mean and variance).

In this paper, we study the timing of clinical associations to address three goals: (1) to demonstrate that a large, heterogeneous database can reveal fine temporal patterns in associations; (2) to illustrate several types of associations; and (3) to ascertain the value of exploiting time. We attempt to classify the resulting associations into categories: definitional, physiologic, and intentional. Definitional associations are those in which a concept found in the record, such as 'hyperkalemia,' is defined as a specific physical finding, such as elevated potassium. Physiologic associations are linked by physiologic cause and effect, such as a response to a drug. Intentional associations are linked by human decision and action, such as ordering a drug based on seeing a physical finding. Before the study, we expected that definitional associations should be easily detectable but possibly have poor temporal resolution, and that other associations might require a more focused sample. In this paper, we used linear correlation because it is simple and easily repeated, it is directional, and it accommodates continuous values. Proving causality was out of scope for this study.

METHODS

We used New York-Presbyterian Hospital's clinical data warehouse, which contains 22 years of electronic health record data on 3 million patients. We abstracted seven blood laboratory tests from the warehouse (sodium, potassium, bicarbonate,

creatinine, urea nitrogen, glucose, and hemoglobin), and we abstracted 30 clinical concepts from clinician signout notes.³⁶ Signout notes are written by clinicians (mostly residents) caring for inpatients to facilitate the transfer of care for overnight coverage. They are updated frequently and contain clinically relevant information. While the notes are primarily inpatient documentation, the laboratory data span inpatient and ambulatory care and the signout notes include chronic ambulatory diagnoses; therefore, associations that are relevant only to ambulatory care may still be detected. The concepts included diseases, symptoms, medications, and procedures. The laboratory tests were chosen because they were commonly available. The concepts were chosen such that they were somewhat common (among the 250 most common diseases, symptoms, or procedures in the signout notes, or among the 250 most common medications in the signout notes) and such that one of the two physician authors (GH and AP) expected an association between the concept and one of the seven laboratory tests (eg, hyperkalemia and potassium) or expected no particularly strong association (eg, atelectasis or osteomyelitis).

We used case-insensitive pattern matching to extract the clinical concepts from the notes. Signout notes were collected and words in the notes were matched to stemmed phrases. No attempt was made to detect negation or other modifiers. Each patient had a list of laboratory values and collection times, and a list of signout note times and an indication whether each of the 30 concepts appeared or not.

We used lagged linear correlation³⁵ to characterize the associations between laboratory values and signout concepts. The laboratory values were continuous. For each patient, the laboratory values were normalized to a mean of 0 and variance of 1 in order to reduce inter-patient effects. Patients with fewer than three laboratory values were eliminated from the analysis for that test. For every time point where there was a concept, a laboratory value was interpolated as the weighted mean of the two surrounding measurements (ie, the internal time point was interpolated linearly), or as the closest measurement if there was no measurement on one side (figure 1). Concepts were represented as 0 (concept absent) and 1 (concept present). We emphasize that while failure to mention a concept does not necessarily mean that the condition is absent in the patient, we treated it as absent in this initial study to avoid incorporating domain knowledge that might create associations as an artifact. For every time point where there was a laboratory value, the concept value was interpolated from the surrounding signout values. Therefore, the concepts took on continuous values between signout notes. The result was a set of original and interpolated values for laboratory values and for signout concepts such that a real or interpolated laboratory value and a real or interpolated signout concept were paired for every time point.

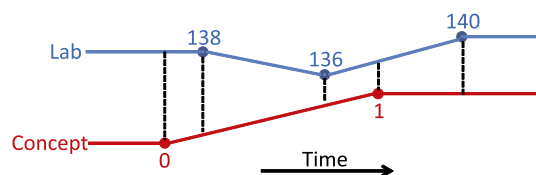


Figure 1 Temporal interpolation. For a given patient, each measured point (solid circles on both curves) is mapped to an interpolated point on the opposite curve. Concepts are mapped to 0 and 1 (absent and present), and laboratory values are continuous. Interpolation is linear between points and constant at the ends of the curves.

We calculated Pearson's linear correlation over all pairs for all patients and recorded the result as the 0 day lag for that laboratory–concept pair. We repeated this process 121 times, first lagging all laboratory values by -60 to 60 days. The result was a curve from -60 to 60 days with positive and negative correlations. We generated curves for every laboratory–concept pair. We repeated the entire experiment using mutual information instead of linear correlation to look for non-linear correlations, at the expense of losing the direction of the association. For this data set, the results were similar other than the loss of direction.

The full list of 30 stemmed phrases and the concepts to which they map are shown in the online data supplement. Common trade names were used to find medications, and common abbreviations were used. Associations shown in figures 2–4 were first selected based on the size of the correlation, relative absence of noise, and characteristic shape of the curve, and additional curves were added based on expected associations. A subset was then selected to illustrate data mining concepts.

The experiment was approved by the Institutional Review Board.

RESULTS

Figure 2 illustrates the results, and the online data supplement shows the full cohort of seven laboratory tests and 30 concepts. *Definitional associations* are shown in figure 2A, where higher than average potassium precedes mention of hyperkalemia and lower than average potassium precedes mention of hypokalemia. This relationship is largely definitional (eg, hyperkalemia means high potassium), and it occurs because the clinician mentions the concept after seeing the potassium test results. After the mention of either concept, the signal falls closer to zero, presumably because the conditions were subsequently corrected for many patients. Similar profiles can be seen for hypernatremia and hyponatremia with sodium (figure 2B), and hyperglycemia and hypoglycemia with glucose (figure 2C).

Intentional and physiologic associations are also shown in figure 2A, where mention of the drug spironolactone, which generally signifies that the drug was given to the patient, is associated with low potassium beforehand and high potassium afterward. This is appropriate because clinicians are likely to order spironolactone, a potassium-sparing diuretic, for patients with low potassium, and patients are likely to experience a rise in potassium after the drug is initiated. This example illustrates human intention, in the case of the ordering of a drug based on observing physiologic parameters, as well as human physiology, in the case of potassium rising after the drug is initiated. Error bands are not shown on the graphs, but to illustrate the variance, the width of the error band (calculated using bootstrap resampling on the patient variable) for the potassium–spironolactone association in figure 2A is 0.0044, or almost 10 times smaller than the size of the effect.

Hypomagnesemia shows a similar profile. Blood magnesium is measured infrequently, but because it can be a cause of hypokalemia,³⁷ it is sometimes measured when hypokalemia is found or when attempts to address hypokalemia have failed. Therefore, one would expect to find low blood potassium followed by measurement of magnesium and mention of hypomagnesemia, followed by magnesium replacement and successful potassium replacement with a rise in potassium. Figure 2A confirms the pattern.

Figure 2B shows physiologic effects for blood sodium. Spironolactone, furosemide, and hydrochlorothiazide have the bulk of their signals to the right of 0 days with a negative correlation, implying that they all precede low blood sodium, and, in fact, all

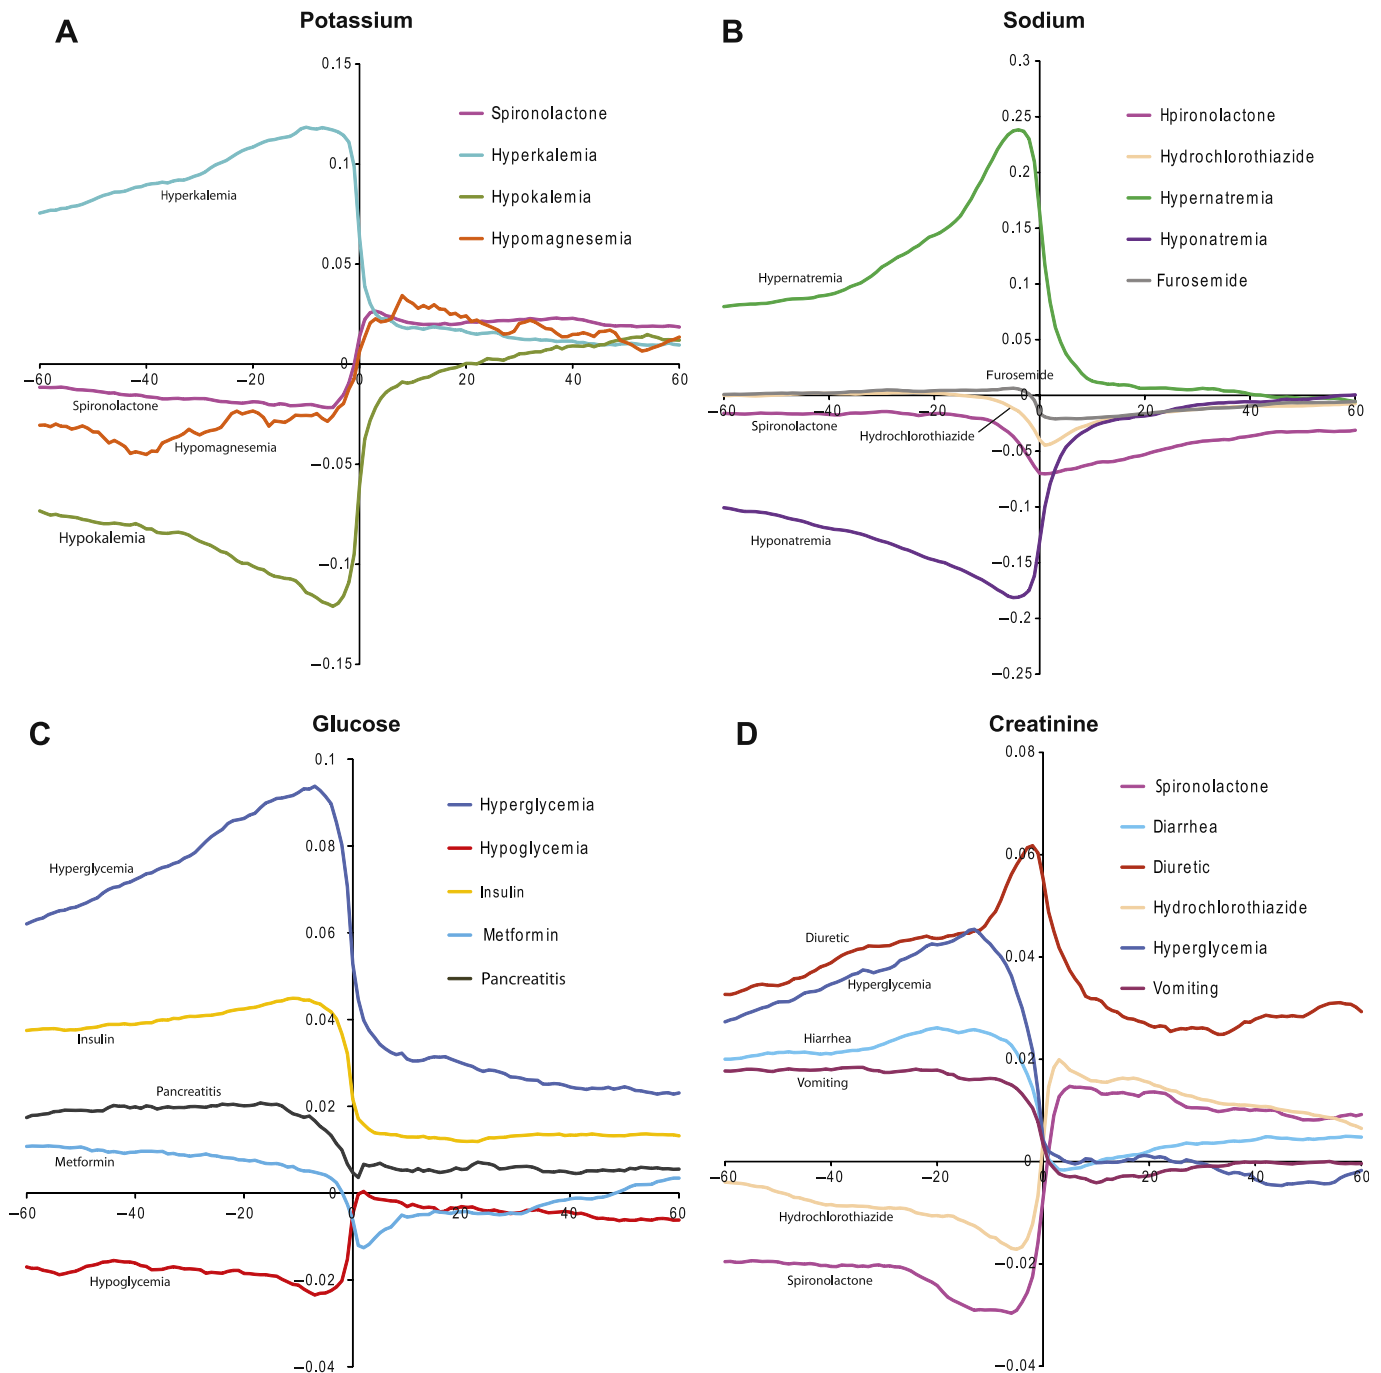


Figure 2 Correlation of laboratory values and signout note concepts. Four blood laboratory values (graphs A–D) are correlated with clinical concepts extracted from physician signout notes (respective graph legends). Linear correlation (y axis) is plotted against time lag in days (x axis). Signal to the left of 0 days implies that changes in the laboratory value preceded changes in the concept, and signal to the right implies that changes in the laboratory value followed changes in the concept. Correlation greater than zero implies that a higher laboratory value was associated with presence of the concept. See text for interpretations.

three drugs are known causes of hyponatremia.^{38,39} In figure 2C, metformin, an oral drug used to treat diabetes mellitus, is associated with high glucose beforehand and low glucose afterward. Insulin shows more markedly high glucose beforehand, but it does not drop below zero afterward. Metformin is used chronically, whereas insulin is used both chronically and in acute episodes. During acute episodes, while insulin will lower the glucose, it will usually remain higher than average for the patient.

The timing of treatment causes differs from the timing of disease causes. Pancreatitis can precipitate hyperglycemia, implying that

the high blood glucose should follow the concept pancreatitis in figure 2C. The workflow of healthcare flips the relationship, however. High blood glucose may be uncovered during a diagnostic workup, then pancreatitis is diagnosed (and mentioned), and then the glucose is corrected. High blood glucose therefore appears to precede its cause, pancreatitis. Figure 2C confirms this latter pattern. This illustrates the differing temporal patterns uncovered for disease-related causes like pancreatitis versus treatment-related causes like medications (eg, spironolactone and sodium in figure 2B). A disease appears to follow its effect because it is the process of diagnosis, not the disease itself, that

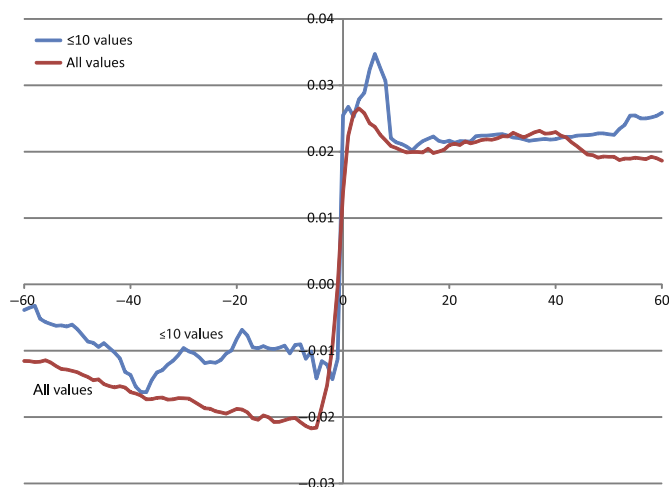


Figure 3 Patients with few values. The linear correlation between blood potassium level and mention of spironolactone is shown, comparing all patients with relevant data versus only patients with 10 or fewer laboratory values. Linear correlation (y axis) is plotted against time lag in days (x axis).

is measured. Thus, correctable effects appear to precede the disease but follow the treatment.

The *shape of the curve* may help to distinguish types of associations. In figure 2D, vomiting, diarrhea, and severe hyperglycemia cause dehydration and elevated creatinine, and as noted above, high creatinine appears to precede mention of the causative conditions. Note also the similar shapes of the curves of the three conditions, as well as the shape of the curve of pancreatitis in figure 2C. This shape may be indicative of a disease-related cause, and it may be distinct from the shape of a definitional attribute like hyper- and hypokalemia in figure 2A, hyper- and hyponatremia in figure 1B, and hyper- and hypoglycemia in figure 2C, which all share similarly shaped curves. Thus, shape may be indicative of type of association, with disease causes looking different from definitions.

The difference in shape may simply be a difference in *chronicity of the association*. Definitional associations appear to be more acute (sharper peak closer to 0 days) than the causative associations included in our study. When an electrolyte is measured—especially in inpatients—review, documentation, and possible treatment are likely to be rapid. Diagnosis is a more complex process, so the time between the appearance of laboratory values and documenting a condition is likely to be longer.

The *specificity of the concept* also affects the result. In figure 2D, mention of specific drugs (spironolactone and hydrochlorothiazide) is preceded by low creatinine and followed by high creatinine, likely reflecting the treatment of patients in fluid overload, with lower creatinine associated with fluid overload beforehand and relatively higher creatinine after being treated. Mention of specific drugs likely reflects actual use of the drug. Mention of the general concept ‘diuretic’ may not signify any specific treatment at the present time, so the more balanced association found in figure 2D may simply reflect reduced renal function at times when patients require diuretics.

We demonstrate the *value of aggregation* in figure 3, which shows the correlation between blood potassium and mention of spironolactone. Whereas figure 2 is based on 20 years of data, the ‘≤10 values’ curve in figure 3 is limited to patients with 10 or fewer potassium values. This limits the data set to 444 patients instead of the 5424 patients with relevant data (8%), and only

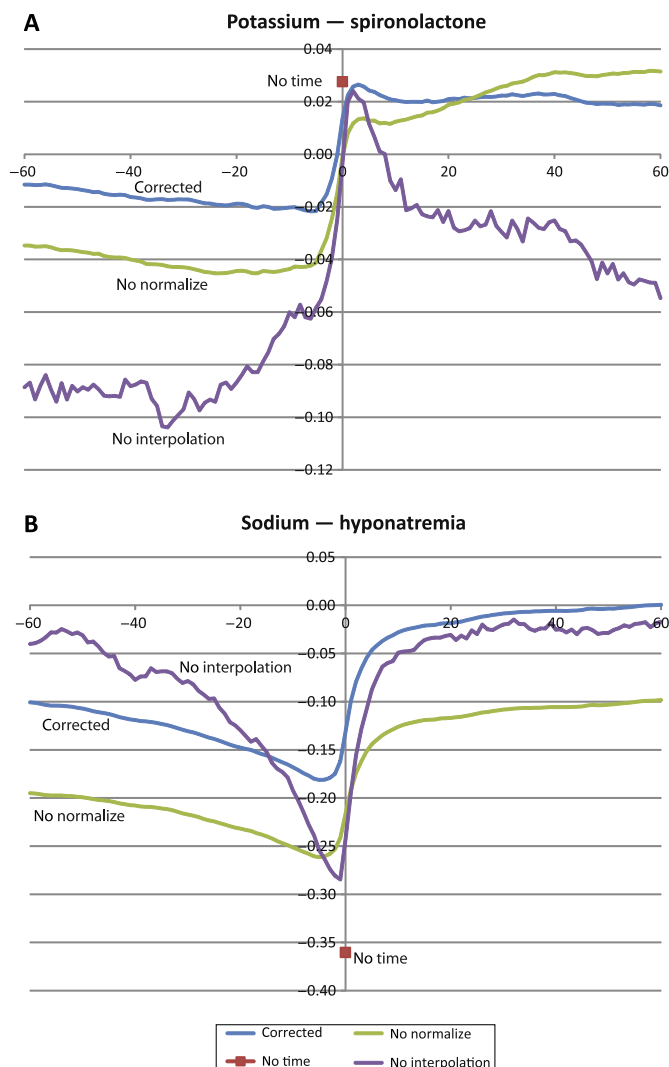


Figure 4 Comparison of different temporal algorithms. Linear correlation of blood potassium with mention of spironolactone (A) and blood sodium with mention of hyponatremia (B) is shown, comparing four temporal algorithms (see text). Linear correlation (y axis) is plotted against time lag in days (x axis).

2534 values instead of 570 000 (0.4%), with an average of fewer than six values per patient. The result is grossly similar to the full set, with a similar magnitude of effect, but with more noise and a slight positive shift. Despite the likelihood that this cohort is medically distinct from patients with more potassium values, the result is similar. This result could never have been obtained based on a single patient with only 10 data values, yet the aggregation produced a result that mirrors the result for a much more substantial sample.

Figure 4 shows the *value of using all the time* and correcting for *inter-patient effects*, comparing our approach to alternative, simpler algorithms. The ‘corrected’ curves are the algorithms used in the corresponding curves in figure 2. ‘No time’ uses a single value for each patient, where laboratory tests are the mean for the patient and concepts are the proportion of signout notes that mention the concept. This discards temporal information within patients, eliminating the lag. Therefore, only a simple relationship can be represented—that a pair of variables is somehow directly or inversely associated—and the subtleties of their temporal relationship are lost. For example, in figure 4A, the ‘no time’ version of blood potassium and spironolactone

shows positive correlation, completely losing the fact that spironolactone is correlated with low potassium before it is given.

Failure to normalize in figure 4B (*'no normalize'*) incorporates an inter-patient effect, which shifts the curve negatively, implying that patients with lower average sodium tend to have hyponatremia, and obscuring the physiologic and intentional effects within patients. Failure to interpolate (*'no interpolation'*) loses temporal information because only points that happen to occur on the same day are included in the correlation calculation and all other points are discarded; figure 4A,B both show significantly more noise and figure 4A shows a differently shaped curve.

Despite the correlations in figure 2 being highly statistically significant, they are very small. For example, spironolactone and potassium peaks at only 0.025. To attempt to explain the small size, we carried out a scrambling experiment on one of the larger correlations, hyponatremia and sodium, with the results shown in figure 5. When 70% of the patients have their values scrambled in time, the signal becomes approximately 70% smaller, and when 100% of the patients have their values scrambled, the signal disappears. This demonstrates that it is possible to create a small but significant correlation by mixing patients that have a real signal (the 30% unscrambled ones) with patients that have no signal (the 70% scrambled). Therefore, if the database is large enough, there may be enough cases that display a relationship even if the majority of cases do not display the relationship; the result would be a small but statistically significant signal.

DISCUSSION

These generated graphs thus illustrate several basic electronic health record data mining concepts. Few single patients have sufficient data to carry out such an analysis. Therefore, aggregation is critical to address the incompleteness of records, although it must be done with caution to account for inter-patient effects. Despite the complex format of the data, simple pattern matching and temporal interpolation are sufficient to reveal associations. Using all available temporal information can

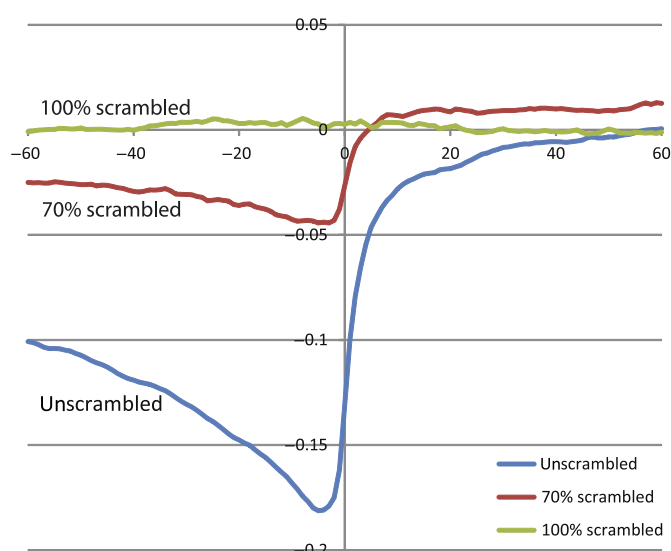


Figure 5 Inducing low correlation. For the hyponatremia–sodium correlation, a lower correlation is induced by scrambling a proportion of the patients. When 70% of the patients have their values scrambled in time, the height of the correlation drops approximately 70%. When 100% of the patients have their values scrambled, the correlation disappears.

reveal relationships that would be missed if information is rolled into a single value per patient or if a lag was not used. The relationships reveal information not just about human physiology, but also about human intention, so the results must be interpreted appropriately. Nevertheless, clues such as curve shape may facilitate interpretation.

Normalizing within patient appears to produce less biased results, as shown in figure 4, but the price is that information is being discarded because there may be legitimate inter-patient effects. It is important to remember, however, that we have lumped together every type of patient in the past 2 decades, inpatient and outpatient, and given such a heterogeneous population, it will be difficult to distinguish legitimate associations from artifacts induced by differential patterns of confounding (eg, the mixing of two populations with different levels of a confounder may induce spurious association patterns).

While the discovered correlations appear to be small, we believe that they reflect real associations for several reasons. They are statistically significantly different from zero (no association). The shape of the curves appears to have a characteristic shape shared among other associations of a similar type. The uncovered associations appear to make clinical and physiological sense. And figure 5 illustrates one sufficient explanation for the small size: a mixture of patients that display the association with patients that do not display the association. In reality, the reason for the small signal may be more complex, including the possibility of patients with opposing effects.

In fact, the small size of the correlations illustrates one of the benefits of incorporating time: in an atemporal experiment (no time), a correlation of 0.05 may be considered inconsequential even if it is statistically significant. The shape of the temporal curve reveals additional information that may allow the researcher to interpret the correlation and decide if it is significant (eg, if the shape is like that of spironolactone, then suspect a real association).

Our results are both encouraging and cautionary. The main lesson is that electronic health records are not direct reflections of the patient state, but rather reflections of the healthcare process and the recording process. Counterintuitive results such as the apparent reordering of cause and effect must be understood and interpreted. The expected appearance of voluminous clinical data should be a boon to research, but only after analysis to understand the record and its biases.

There is a long tradition of mining clinical data in informatics. We touch on four relevant areas here: temporal abstraction,^{3–16} causal inference,^{17–24 40 41} numeric time series analysis,^{25–28} and recent electronic health record-based association studies.^{29–34 42} Temporal abstraction has a similar goal to ours—making sense of complex temporal relationships among clinical variables and concepts—although its methods tend to use symbolic patterns and knowledge engineering rather than purely numeric manipulation of raw data values. Causal inference has a greater focus on distinguishing causes from other associations than on uncovering detailed temporal relationships (with exceptions highlighted below). Numeric time series analysis—which is closer in methodology to our work—has been applied extensively in physiology studies but less so in electronic health record research. Recent electronic health record-based association studies employ primarily numeric methods to uncover associations, but generally do not capture complex temporal information.

In the first area, Bellazzi and coauthors³ used a multi-step approach to identify temporal abstractions in home monitoring

diabetes data to study patient dynamics. Sacchi and coauthors⁴ described a method for inferring relationships based on encoding the time series as a sequence of trends and discovering the relationships among them in terms of Allen's intervals,⁵ quantifying the evidence by confidence and support. They discovered complex temporal patterns, not merely instantaneous events. They ran it on a hemodialysis data set with three quantifiable variables. Concaro and coauthors⁶ demonstrated the extraction of temporal associations on a combined clinical and administrative dataset to address diabetes care. Bellazzi and coauthors⁷ used the same dataset to demonstrate a method that accommodates both point and interval data. They uncovered 24 rules with sufficient support and confidence, and found them to be clinically meaningful. Batal and coauthors⁸ employed temporal abstraction followed by machine learning of a classification model to address anticoagulation therapy.

Jin and coauthors⁹ described a technique to find unexpected temporal associations for the purpose of discovering adverse drug reactions. A key element of this work was operating efficiently despite the infrequency of events. Winarko and Roddick¹⁰ mined temporal associations from interval-based data on simulated datasets.

Shahar and colleagues have studied a critical issue in temporal mining of electronic health records: information is recorded at many levels of abstraction, such as from primary blood tests to diagnoses.^{11–12} For example, Moskovitch and Shahar¹³ turned point data into intervals and presented an algorithm for temporal hierarchical clustering in the area of diabetes. Klimov and coauthors¹⁴ developed a system to visualize and explore time-oriented health data, incorporating aggregation functions (eg, to aggregate a population of values over time) as well as knowledge-based temporal abstraction tools.

Similarly, Guyet and coauthors¹⁵ addressed multilevel abstraction and visualization, creating a collaborative man-machine process for segmentation, classification, and learning. Aigner and coauthors¹⁶ also address visualization and abstraction, adding principal component analysis and clustering to facilitate visualization by summarizing the data.

In general, our work has similar goals to this first area of research, but it uses purely numeric methods that exploit the raw data without imposing a predefined framework for trends and temporal relationships. It may some day complement temporal abstraction by uncovering new types of temporal relationships that would be useful in the first abstraction phase.

In the second area, while a number of statistical methods can be used to uncover associations among clinical variables in electronic health records, determining which ones are actually causal is more complex, and is an area of informatics research, with investigators using methods such as regular and dynamic Bayesian networks,^{17–20 40 41} Granger causality,^{21 22} and logic formalisms.^{23 24} This research has not focused on distinguishing physiologic causes from workflow-related intentional causes, and only recent work^{23 24} has produced finely detailed temporal information.

In the third area, physiologists have used time series analysis extensively to study health and disease,²⁵ and the relationships are complex. For example, in a study of signal variability over time, Hanss and coauthors showed that low heart rate variability is associated with greater risk of cardiac ischemia during general anesthesia,²⁶ whereas Anderson and coauthors showed that high breathing variability is associated with elevated blood pressure and inversely associated with heart rate variability.²⁷ Glucose time series have been studied for some time with the hope of predicting glucose levels and improving diabetic

control.²⁸ Our work is similar because we use non-linear numeric methods to analyze the data, but our data source is the electronic health record including clinical notes, so we must account not only for physiology, but also for clinical workflow.

In the fourth area, researchers are increasingly carrying out traditional association studies on electronic health record data.^{29–34 42} For the most part, the full temporal record is not exploited, often using simultaneous values³⁰ or using aggregation methods such as taking the median.³⁴ Sometimes gross temporal information is incorporated by filtering to ensure an allowable ordering of events,³² but without fully exploiting detailed relationships.

In summary, our goals are similar to those of the temporal abstraction field, but we employ numeric methods that more directly operate on the raw data. In the long run, the temporal abstraction approach may turn out to be more efficient, but our approach carries the fewest assumptions about the types of temporal relationships that should be expected and is thus a useful exploratory technique. We are not attempting in our work to prove causality, and we focus instead on the detailed timing of associations. Our methods (both lagged linear correlation and mutual information) are drawn primarily from the time series field, but we apply them not to a purely natural physiologic system but to the complex healthcare process that is represented in the electronic health record (especially in its narrative notes), and the discovered relationships must be interpreted in that context. Our work parallels association studies in electronic health records, but we seek to uncover more complex temporal relationships.

Although lagged linear correlation is used frequently to find the temporal precedence between variables (eg, in syndromic surveillance⁴³), we believe that this is the first use of it for characterizing the timing of associations in a large-scale study of electronic health records. Furthermore, we believe that this is the first addition of within-patient normalization to reduce inter-patient effects in the lagged correlations and the first application of linear interpolation in the context of electronic health record lagged correlations. The mutual information version of the algorithm builds on our previous work.⁴⁴ Our hope is that our work will help inform the temporal abstraction field and expand the scope of electronic health record association studies.

Our study was carried out in a single medical center with an emphasis on inpatient care. Future work should address a broader range of care settings, improved concept abstraction, and more advanced time series techniques.

CONCLUSION

We demonstrated that a relatively simple set of methods can extract detailed temporal properties of clinical associations from a heterogeneous database of electronic health record data. We classified associations into three types—definitional, physiologic, and intentional—and showed that care must be taken in interpreting the associations because the health record represents the clinical workflow and not just patient physiology. We found that fully exploiting time in the record revealed the most detailed and reliable information.

Funding This work was funded by grants from the National Library of Medicine, 'Discovering and applying knowledge in clinical databases' (R01 LM006910), and 'Training in Biomedical Informatics at Columbia University' (T15 LM007079).

Competing interests None.

Ethics approval Ethics approval was provided by Columbia University Institutional Review Board.

Provenance and peer review Not commissioned; externally peer reviewed.

REFERENCES

1. Blumenthal D, Tavenner M. The "meaningful use" regulation for electronic health records. *N Engl J Med* 2010;**363**:501–4.
2. Friedman CP, Wong AK, Blumenthal D. Achieving a nationwide learning health system. *Sci Transl Med* 2010;**2**:57cm29.
3. Bellazzi R, Larizza C, Magni P, et al. Intelligent analysis of clinical time series: an application in the diabetes mellitus domain. *Artif Intell Med* 2000;**20**:37–57.
4. Sacchi L, Larizza C, Combi C, et al. Data mining with temporal abstractions: learning rules from time series. *Data Min Knowl Discov* 2007;**15**:217–47.
5. Allen JF. Maintaining knowledge about temporal intervals. *Commun ACM* 1983;**26**:832–43.
6. Concaro S, Sacchi L, Cerra C, et al. Temporal data mining for the assessment of the costs related to diabetes mellitus pharmacological treatment. *AMIA Annu Symp Proc* 2009;**2009**:119–23.
7. Bellazzi R, Sacchi L, Concaro S. Methods and tools for mining multivariate temporal data in clinical and biomedical applications. *Conf Proc IEEE Eng Med Biol Soc* 2009;**2009**:5629–32.
8. Batal I, Sacchi L, Bellazzi R, et al. A temporal abstraction framework for classifying clinical temporal data. *AMIA Annu Symp Proc* 2009;**2009**:29–33.
9. Jin HW, Chen J, He H, et al. Mining unexpected temporal associations: applications in detecting adverse drug reactions. *IEEE Trans Inf Technol Biomed* 2008;**12**:488–500.
10. Winarko E, Roddick JF. ARMADA—an algorithm for discovering richer relative temporal association rules from interval-based data. *Data Knowl Eng* 2007;**63**:76–90.
11. Shahar Y. A framework for knowledge-based temporal abstraction. *Artif Intell* 1997;**90**:79–133.
12. Shahar Y, Musen MA. Knowledge-based temporal abstraction in clinical domains. *Artif Intell Med* 1996;**8**:267–98.
13. Moskovich R, Shahar Y. Medical temporal-knowledge discovery via temporal abstraction. *AMIA Annu Symp Proc* 2009;**2009**:452–6.
14. Klimov D, Shahar Y, Taieb-Maimon M. Intelligent visualization and exploration of time-oriented data of multiple patients. *Artif Intell Med* 2010;**49**:11–31.
15. Guyet T, Garbay C, Dojat M. Knowledge construction from time series data using a collaborative exploration system. *J Biomed Inform* 2007;**40**:672–87.
16. Aigner W, Miksch S, Muller W, et al. Visual methods for analyzing time-oriented data. *IEEE Trans Vis Comput Graph* 2008;**14**:47–60.
17. Murphy K. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis. Berkeley: University of California, 2002.
18. Charitos T, van der Gaag L, Visscher S, et al. A dynamic Bayesian network for diagnosing ventilator-associated pneumonia in ICU patients. *Expert Syst Appl* 2009;**36**:1249–58.
19. van Gerven M, Taal B, Lucas P. Dynamic Bayesian networks as prognostic models for clinical patient management. *J Biomed Inform* 2008;**41**:515–29.
20. Nachimuthu SK, Wong A, Haug PJ. Modeling glucose homeostasis and insulin dosing in an intensive care unit using dynamic bayesian networks. *AMIA Annu Symp Proc* 2010;**2010**:532–6.
21. Granger CW. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 1969;**37**:424–38.
22. Kaminski M, Ding M, Truccolo WA, et al. Evaluating causal relations in neural systems: granger causality, directed transfer function and statistical assessment of significance. *Biol Cybern* 2001;**85**:145–57.
23. Kleinberg S, Mishra B. The temporal logic of causal structures. *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI-09)*; Corvallis: Oregon, 2009:303–12.
24. Kleinberg S, Mishra B. Multiple testing of causal hypotheses. In: Illari PM, Russo F, Williamson J, eds. *Causality in the Sciences*. Oxford: Oxford University Press, 2011.
25. Kantz H, Kurths J, Mayer-Kress G. *Nonlinear Analysis of Physiological Data*. Berlin: Springer, 1998.
26. Hanss R, Block D, Bauer M, et al. Use of heart rate variability analysis to determine the risk of cardiac ischaemia in high-risk patients undergoing general anaesthesia. *Anaesthesia* 2008;**63**:1167–73.
27. Anderson DE, McNeely JD, Chesney MA, et al. Breathing variability at rest is positively associated with 24-h blood pressure level. *Am J Hypertens* 2008;**21**:1324–9.
28. Sparacino G, Facchinetti A, Maran A, et al. Continuous glucose monitoring time series and hypo/hyperglycemia prevention: requirements, methods, open problems. *Curr Diabetes Rev* 2008;**4**:181–92.
29. Himes BE, Dai Y, Kohane IS, et al. Prediction of chronic obstructive pulmonary disease (COPD) in asthma patients using electronic medical records. *J Am Med Inform Assoc* 2009;**16**:371–9.
30. Cohen MJ, Grossman AD, Morabito D, et al. Identification of complex metabolic states in critically injured patients using bioinformatic cluster analysis. *Crit Care* 2010;**14**:R10.
31. Chen DP, Weber SC, Constantinou PS, et al. Clinical arrays of laboratory measures, or "clinarrays", built from an electronic health record enable disease subtyping by severity. *AMIA Annu Symp Proc* 2007:115–19.
32. Wang X, Hripcsak G, Markatou M, et al. Active computerized pharmacovigilance using natural language processing, statistics, and electronic health records: a feasibility study. *J Am Med Inform Assoc* 2009;**16**:328–37.
33. Brownstein JS, Murphy SN, Goldfine AB, et al. Rapid identification of myocardial infarction risk associated with diabetes medications using electronic medical records. *Diabetes Care* 2010;**33**:526–31.
34. Chen DP, Dudley JT, Butte AJ. Latent physiological factors of complex human diseases revealed by independent component analysis of clinarrays. *BMC Bioinformatics* 2010;**11** Suppl 9:S4.
35. Chatfield C. *The Analysis of Time Series*, 4th edn. New York: Chapman and Hall, 1991.
36. Stein DM, Wrenn JO, Johnson SB, et al. Signout: a collaborative document with implications for the future of clinical information systems. *AMIA Annu Symp Proc* 2007:696–700.
37. Huang CL, Kuo E. Mechanism of hypokalemia in magnesium deficiency. *J Am Soc Nephrol* 2007;**18**:2649–52.
38. Abramow M, Cogan E. Clinical aspects and pathophysiology of diuretic-induced hyponatremia. *Adv Nephrol Necker Hosp* 1984;**13**:1–28.
39. Stigl W, McAlister FA, Ezekowitz J, et al. Usefulness of spironolactone in a specialized heart failure clinic. *Am J Cardiol* 2004;**94**:443–7.
40. Pearl J. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press, 2000.
41. Mani S, Cooper G. Causal discovery using a Bayesian local causal discovery algorithm. *Stud Health Technol Inform* 2004;**107**:731–5.
42. Hripcsak G, Austin JH, Alderson PO, et al. Use of natural language processing to translate clinical information from a database of 889,921 chest radiographic reports. *Radiology* 2002;**224**:157–63.
43. Hripcsak G, Soulakis ND, Li L, et al. Syndromic surveillance using ambulatory electronic health records. *J Am Med Inform Assoc* 2009;**16**:354–61.
44. Albers DJ, Hripcsak G. A statistical dynamics approach to the study of human health data: resolving population scale diurnal variation in laboratory data. *Phys Lett A* 2010;**374**:1159–64.