Applying Methods for Signal Detection in Spontaneous Reports to Electronic Patient Records

Jing Zhao*
Department of Computer and Systems Sciences
Stockholm University
Stockholm, Sweden
jingzhao@dsv.su.se

Isak Karlsson
Department of Computer and Systems Sciences
Stockholm University
Stockholm, Sweden
isak-kar@dsv.su.se

Lars Asker
Department of Computer and Systems Sciences
Stockholm University
Stockholm, Sweden
asker@dsv.su.se

Henrik Boström
Department of Computer and Systems Sciences
Stockholm University
Stockholm, Sweden
henrik.bostrom@dsv.su.se

ABSTRACT

Currently, pharmacovigilance relies mainly on disproportionality analysis of spontaneous reports. However, the analysis of spontaneous reports is concerned with several problems, such as reliability, under-reporting and insufficient patient information. Longitudinal healthcare data, such as Electronic Patient Records (EPRs) in which comprehensive information of each patient is covered, is a complementary source of information to detect Adverse Drug Events (ADEs). A wide set of disproportionality methods has been developed for analyzing spontaneous reports to assess the risk of reported events being ADEs. This study aims to investigate the use of such methods for detecting ADEs when analyzing EPRs. The data used in this study was extracted from Stockholm EPR Corpus. Four disproportionality methods (proportional reporting rate, reporting odds ratio, Bayesian confidence propagation neural network, and Gamma-Poisson shrinker) were applied in two different ways to analyze EPRs: creating pseudo spontaneous reports based on all observed drug-event pairs (event-level analysis) or analyzing distinct patients who experienced a drug-event pair (patient-level analysis). The methods were evaluated in a case study on safety surveillance of Celecoxib. The results showed that, among the top 200 signals, more ADEs were detected by the event-level analysis than by the patient-level analysis. Moreover, the event-level analysis also resulted in a higher mean average precision. The main conclusion of this study is that the way in which the disproportionality analysis is applied, the event-level or patient-level analysis, can have a much higher impact on the performance than which disproportionality method is employed.

*Corresponding author.

KDD-DMH'13, August 11, 2013, Chicago, Illinois, USA

Categories and Subject Descriptors

H.2.8 [**Database Applications**]: Data mining; J.3 [**Life and Medical Sciences**]: Medical information systems

Keywords

Pharmacovigilance, disproportionality analysis, drug safety, adverse drug events, electronic patient records

1. INTRODUCTION

Drug safety has always been a crucial issue in the pharmaceutical industry. Although the benefit-risk analysis of newly developed drugs is already conducted in clinical trials before they are released to the market, post-marketing detection and surveillance is still necessary, since clinical trials are normally done with limited samples of subjects and within limited periods of time. Several drugs have been withdrawn from the market due to safety issues. For example, Cerivastatin was withdrawn worldwide in 2001 because of causing fatal rhabdomyolysis [1]. During post-marketing surveillance, potential safety hazards of a certain drug are uncovered through analyzing different types of data source, e.g., spontaneous reporting system (SRS) data, longitudinal administrative or claim data, or electronic patient records (EPRs) [2].

Currently, post-marketing drug surveillance largely relies on SRS data, such as the Yellow Card scheme[1] in the UK, Adverse Event Reporting System (AERS)[2] used by the Food and Drug Administration (FDA) in the USA, and the World Health Organization (WHO) International Database[3] maintained at Uppsala Monitoring Center (UMC) in Uppsala, Sweden. These SRSs have been developed to collect information about adverse drug events (ADEs) voluntarily reported

[1]Information available at https://yellowcard.mhra.gov.uk
[2]Information available at http://www.fda.gov/Drugs
[3]Information available at http://www.who-umc.org

by both clinical professionals and patients. They typically contain drugs being taken, adverse events that occurred, and occasionally some basic patient information. However, SRS data has some obvious limitations, such as under-reporting, problems with reliability and compliance, insufficient information of patients' medical history, and lacking the total number of patients taking the drug of interest. Recently increasingly available large databases of EPRs have been used to complement such SRS data. The EPRs typically contain longitudinal observational data of large samples of patients including demographic information, medical history, drug consumption with exposure time and dosage information, clinical measurements, lifestyle factors, and clinical narratives, which are also rich in information for detecting ADEs [2].

Disproportionality analysis, widely used in post-marketing drug surveillance, was initially designed to analyze data from SRS for detecting drug safety signals that correspond to clinical events associated with the use of one or more drugs. These events can be any undesired medical occurrence in a patient that has taken the drug. A signal indicating an ADE is detected if the frequency of such an event, e.g., a skin problem, deviates substantially from what is expected, i.e., when not taking the drug. The most common disproportionality methods include Proportional Reporting Ratio (PRR) [3], Reporting Odds Ratio (ROR) [4], Bayesian Confidence Propagation Neural Network (BCPNN) [5], and Gamma Poisson Shrinker (GPS) [6]. They count the drug-event pairs from the reports and calculate the "observed-expected" ratio of each drug-event pair.

Many studies applied disproportionality methods to investigate the safety of specific drugs with SRS data [3, 7, 8, 9]. However, the best methodology to detect ADE signals with data from EPRs is still nascent. According to Madigan and Ryan [10], the only way to find out is through extensive empirical experimentation. So far, only a few studies have investigated how to analyze longitudinal data using a disproportionality method [11, 12] or several such methods [13]. Zorych el al. [13] made the first attempt to extend disproportionality methods to the analysis of administrative claim data, where they proposed three counting approaches to represent the events: *distinct patient*, *pseudo SRS*, and *modified SRS*. In the first approach, all events are represented by distinct patients who experienced them; in the second approach, all valid drug-event pairs, i.e., events associated to one or more drugs, are counted; and in the third approach, besides valid drug-event pairs, events that are not associated to any drug are also counted. The last two approaches lead to very similar results in their study. They focused on the relative difference among those disproportionality methods instead of the difference on their overall performance between different counting approaches. However, it is crucial to investigate the most suitable way to apply disproportionality methods to EPRs. Therefore, this study aims at evaluating two ways of adapting disproportionality methods to EPRs: *pseudo SRS* and *distinct patient*. These strategies are here referred to as event-level analysis and patient-level analysis, respectively, to more clearly indicate that the former is based on analyzing all events that occur in conjunction with a drug, while the latter is based on analyzing at most one event pair per patient. The data was extracted from the Stockholm EPR corpus[4] [14] and the safety of Celecoxib was used as a case study.

This paper is organized as follows: Section 2 discusses in more detail the differences between SRS and EPRs. Section 3 introduces the methods evaluated and compared in this study, including details of the four disproportionality methods, event-level and patient-level analysis, implementation of the case study, and performance metrics. The results are presented in Section 4. Finally, concluding remarks are given in Section 5.

2. SPONTANEOUS REPORTS AND ELECTRONIC PATIENT RECORDS

During the last decade, the interest in pharmacovigilance has been increasing, along with which a wide variety of data sources have been collected and used, including SRS data, captured data of drug dispensing, longitudinal administrative or claims databases, EPRs databases, clinical trials or lab measurements, medical internet forums, and biomedical literatures [15]. Among these data sources, SRS data has been most commonly used for drug safety surveillance, while EPRs data is recently emerging as a potentially useful resource.

The SRS data contains valuable information voluntarily reported by both clinical professionals and drug users. Each spontaneous report typically comprises of one or more drugs, one or more ADEs, and possibly some basic demographic information [13]. Many national and international organizations, such as FDA and WHO, are using such SRSs to detect drug safety signals. However, analyzing SRS data is associated with some problems. First of all, the proportion of under-reported events is typically not known and varies for different types of events and drugs. Second, SRS data is limited in providing drug users' other information including disease history and previous drug prescriptions. Furthermore, as the information reported in SRS is not longitudinal, but is based on a single event only, the analysis of such data only gives a small part of the whole picture [2]. Last but not least, it is not possible to calculate the incidence rate, since the denominator, i.e., the total number of times a drug has been taken, cannot be estimated from the SRS data. Given these limitations of the spontaneous reports, it is necessary to look for additional sources of information to improve ADE signal detection.

One such source of information is EPRs, which contain longitudinal data. Although EPRs databases were historically implemented for administrative purposes, they are becoming valuable resources for different kinds of clinically related research, especially in pharmacovigilance. Compare to SRS, EPRs can potentially provide valuable additional information not only because it contains a wider variety of information, but also because each EPR typically covers all the medical information of one patient since the first day he or she enters the system. Any clinical event of importance is recorded, such as newly developed diseases and prescribed drugs or additional measurements from clinical testing. Hence, it can shed light on the whole picture of a patient's medical history. In addition, the problems of under-reporting and compliance are diminished, especially

[4]This research has been approved by the Regional Ethical Review Board in Stockholm (Etikprövningsnämnden i Stockholm), permission number 2012/834-31/5.

for the hospitalized patients, as the information reporting is under a fairly good control of clinical professionals. Even the patients who are not exposed to any drug or who experience no adverse events are also included in the EPR database. Therefore, it is possible to calculate incidence rates. As such data is longitudinal, the information of time frame is available, which was considered in recently proposed methods for drug safety surveillance [16]. In disproportionality analysis, although SRS contains a large amount of records, the EPRs involve much more data as each patient may contribute to plenty of records over time [13].

3. METHODS

3.1 Disproportionality methods

Disproportionality analysis detects drug safety signals by measuring the disproportionality based on a contingency table with two variables corresponding to the drug of interest and the clinical event of interest, see Table 1. Each cell represents the number of drug-event pairs: a is the number of observations for which both event y and drug x are present; b is the number of observations where other events and drug x are present; c is the number of observations where event y and other drugs are present; and d is the number of observations where other drug-event pairs are present involving neither drug x nor event y. There are four commonly applied disproportionality methods, PRR, ROR, BCPNN, and GPS. The first two are basic measurements calculating the observed and expected ratio directly from the contingency table, while the other two are Bayesian based techniques estimating the prior and posterior probability distribution of the relative reporting ratio (RRR). $RRR = \frac{P(drug\ \mathrm{x}, event\ \mathrm{y})}{P(drug\ \mathrm{x}) \cdot P(event\ \mathrm{y})}$, where $P(drug\ \mathrm{x}, event\ \mathrm{y})$ is the joint probability of drug x and event y and $P(drug\ \mathrm{x})$ and $P(event\ \mathrm{y})$ are the unconditional probabilities of drug x and event y, respectively.

Table 1: A two-by-two contingency table for drug x and event y

	Event y	**Other events**
Drug x	a	b
Other drugs	c	d

- $PRR = \frac{a}{a+b} / \frac{c}{c+d}$
- $ROR = \frac{a}{b} / \frac{c}{d}$
- BCPNN [5] is based on the estimation of the information component (IC) value, which measures the mutual information between two variables, such as drug x and event y. In this case, $IC = \log_2(RRR) = \log_2 \frac{a \cdot (a+b+c+d)}{(a+b) \cdot (a+c)}$. The method derives a posterior expectation and the variance of the IC value.
- GPS [6] defines RRR as $\lambda = \frac{\mu}{E}$, where μ is the mean of Poisson distribution of a and E is the expected event count, estimated as $E = \frac{(a+b) \cdot (a+c)}{a+b+c+d}$. The method then estimates the empirical Bayesian geometric mean (EBGM) defined as $EBGM = \mathrm{e}^{\mathbb{E}(\log(\lambda))}$.

PRR and ROR have the advantage of simplicity, but when it comes to rare events, they are rather unstable with wide confidence intervals and often result in many false positive signals [15, 17]. BCPNN and GPS adopt a Bayesian approach to address this variety issue.

3.2 Data

The data source used in this study is the Stockholm EPR Corpus [14]. It includes medical records from more than 700,000 patients living in Stockholm, Sweden during 2009 and 2010. Approximately 9,863 different diagnoses, encoded using ICD-10-SE[5], and 1,312 different drugs, encoded by the anatomical therapeutic chemical (ATC) classification system, are covered in this EPR database.

Celecoxib was taken as an example for safety signal detection in this case study. It is a non-steroidal anti-inflammatory drug mainly treating oesteoarthritis, rheumatoid arthritis, acute pain, etc. In the Stockholm EPR Corpus, 7,380 patients were prescribed with Celecoxib.

3.3 Event-level vs. patient-level analysis

Two ways of applying the disproportionality analysis methods were investigated: event-level (EL) and patient-level (PL) analysis. In the EL analysis, a pseudo SRS was created using the data from the EPR database, mimicking the way that spontaneous reports are reported. For this representation the assumption is that events are more likely to be reported together with the closest drug taken before. Therefore, events that occurred after drug x was prescribed but before the next drug (drug x or any other drug) was prescribed were considered as *occurred during the exposure to drug x*. In this way, large number of drug-event pairs were generated from the database. For example, if three events occurred after drug x, three pairs were counted for drug x with each of the event; if event y occurred after three different drugs, one pair was counted for it with the closest drug. These drug-event pairs were then counted to fill the contingency table based on the elements in the pair, i.e., pairs that included both drug x and event y were added to a, pairs that included drug x and other events were added to b, pairs that included other drugs and event y were added to c, and pairs that neither included drug x and drug y were added to d. The upper limit of exposure time is 12 weeks.

- a is the number of event of interests that occurred during the exposure to Celecoxib;
- b is the number of other events that occurred during the exposure to Celecoxib;
- c is the number of event of interests that occurred during the exposure to other drugs;
- d is the number of other events that occurred during the exposure to other drugs.

In the PL analysis, instead of drug-event pairs, the content of the contingency table switched to numbers of unique patients. Patients who experienced event y after they took drug x within 12 weeks were considered as *experienced event y during their exposure to drug x*. Each patient was counted only once in this analysis, meaning that the patient who experienced event y during his or her exposure to drug x was added to a only, even if this patient also took other drugs

[5]International Classification of Disease, Version 10, Swedish Modification

or experienced other events during the exposure to drug x. For patients who did not meet the condition to be added to a, their qualifications of being added to b were assessed at first, then to c, otherwise they were added to d.

- a is the number of patients who experienced event of interests during their exposure to Celecoxib;
- b is the number of patients who experienced other events during their exposure to Celecoxib;
- c is the number of patients who experienced event of interests during their exposure to other drugs;
- d is the number of patients who were neither prescribed with Celecoxib nor experienced the event of interests, which equals to the total number of patients in the database subtracted by the sum of a, b and c.

The main difference between these two ways of analyzing the information is the way in which the number of events is counted. In EL analysis, events that occur several times for a single patient can be counted several times, while in PL analysis, each patient is counted only once no matter how many events he or she experienced. This will affect the scores calculated by the disproportionality analysis methods, as d is much larger in EL analysis than it is in PL analysis. Moreover, in PL analysis, all patients were included, while in EL analysis, information from patients that have not taken any drug or did not experience any event were excluded from the analysis.

3.4 Implementation

The four aforementioned disproportionality methods were applied for each of the above two strategies, hence there were eight different ways of detecting signals. They were implemented using the R[6] package "PhViD" [18]. To guarantee a non-zero a, the events that were not reported at least once for a drug were excluded. For PRR and ROR, detected signals were ranked by the PRR and ROR score itself. For BCPNN, detected signals were ranked by the lower bound of the 95% credibility interval of IC, estimated through Monte-Carlo simulation with 10,000 iterations, following Noren et al. [19]. For GPS, detected signals were ranked by the $EBGM$.

3.5 Evaluation

The evaluation was based on the ADE information of Celecoxib extracted from SIDER[7] [20], a drug side effect information system containing information on marketed medicines and their recorded ADEs extracted from public documents and package inserts. If a detected signal was found in the ADE list from SIDER, it was considered as an ADE. In each analysis, the top 50, 100, and 200 signals were investigated separately. Overall precision (OP) and mean average precision (MAP) were used to compare the different measurements.

Overall precision was measured by the ratio between the number of ADEs found and the number of signals, which indicates the proportion of true positives among the total detected signals. Given the fact that the more disproportionally an event occurs, the stronger the signal should be, it is desirable to also consider the order of the ADEs presented, hence the MAP score was calculated to take into account the rank of the ADEs.

The MAP score [21], commonly used in information retrieval, for a set of ADEs is the mean of the average precision scores for each ADE. The higher the MAP score is, the better a method performs. For example, a MAP of 0.1 means that every 10th signal is expected to be an ADE, and even if more than half of the signals are ADEs, the MAP score might be below 0.5, if more ADEs are ranked at the second half of the list. Calculating the MAP score starts with ranking all the signals from top to bottom. Let i denote the position (rank) of each signal, $i = 1, 2, ..., N$, and k_i denote the accumulated number of ADEs at position i, $k_i = 1, 2, ..., K$. The average precision (AP) is $\sum_{i=1}^{N} \frac{k_i}{i} \cdot x(i)$, where $x(i)$ is the indication of an ADE ($x(i) = 1$, if the signal at position i is an ADE; otherwise, $x(i) = 0$). Then, $MAP = \frac{AP}{K}$.

4. RESULTS

The results from applying ROR, PRR, BCPNN and GPS in the EL and PL analysis for the top 50, 100, and 200 signals are listed in Table 2, where the OP and MAP scores are presented. In general, the EL analysis gives better results than the PL analysis. For top 100 and 200 signals, the EL analysis dominated the PL analysis with both higher OP and higher MAP scores, which indicates that in the EL analysis, not only were more ADEs found but they were also ranked higher in general. For the top 50 signals, ROR and PRR detected more ADEs in the PL analysis, and BCPNN and GPS detected more ADEs in the EL analysis; however, those ADEs were detected with stronger signals on average in the EL analysis.

Table 2: Performance of four disproportionality methods in two analyses for the top 50, 100, and 200 signals

	Top 50		Top 100		Top 200	
	OP	**MAP**	**OP**	**MAP**	**OP**	**MAP**
PRR (EL)	0.12	0.25	0.14	0.20	0.12	0.17
PRR (PL)	0.14	0.16	0.08	0.15	0.08	0.12
ROR (EL)	0.14	0.26	0.15	0.21	0.12	0.19
ROR (PL)	0.16	0.18	0.10	0.17	0.09	0.13
BCPNN (EL)	0.20	0.26	0.15	0.25	0.12	0.20
BCPNN (PL)	0.14	0.13	0.11	0.14	0.09	0.12
GPS (EL)	0.20	0.27	0.15	0.24	0.12	0.20
GPS (PL)	0.04	0.17	0.10	0.10	0.08	0.10

Table 2 also shows that the performances among the four methods vary in both analyses. In the EL analysis, BCPNN and GPS performed better than ROR and PRR on both OP and MAP. In the PL analysis, ROR performed best in general. The low OP but high MAP from GPS in the PL analysis for the top 50 signals suggests that the few ADEs found had rather high ranks.

After looking into the output signal lists, it was noticed that ADEs such as headache and nausea, which are known as common ADEs of Celecoxib, were not detected as signals at all by any method. This might be due to the fact that

[6]Open source statistical software, available at http://www.r-project.org

[7]Available at http://sideeffects.embl.de/

these more common events are not coded as diagnoses in the EPRs each time they occur.

5. CONCLUDING REMARKS

This study investigated the analysis of EPRs, extracted from the Stockholm EPR Corpus, using disproportionality methods traditionally employed on spontaneous reports. Two different approaches for how to apply the methods were investigated: creating pseudo spontaneous reports based on all observed drug-event pairs (event-level analysis) or analyzing distinct patients that experienced a drug-event pair (patient-level analysis). Four disproportionality methods, PRR, ROR, BCPNN, and GPS, were evaluated for both levels of analysis. It was shown that the way in which the disproportionality analysis is applied, the event-level or patient-level analysis, can have a much higher impact on the performance than which disproportionality method is employed.

The disproportionality methods performed better in the EL analysis than in the PL analysis in terms of both overall precision and mean average precision. It means that the EL analysis detected more ADEs which were also stronger ADE signals compared to the rest detected signals. In this study, event-drug pairs in the EL analysis were counted in the way which mimics how they were reported in SRS, and patients who experienced a drug-event pair in the PL analysis were counted distinctly without duplication. The EL and PL analysis are just two approaches among many other possible ones. For example, in the EL analysis, events can be paired with all drugs that were taken one week before instead of being paired only with the closest drug; and in the PL analysis, patients who experienced both event of interest and other events can be counted twice, i.e., in both a and b (see Table 1). Besides the counting strategy, temporality might have an impact on the results as well. In this study, 12 weeks was chosen as the length of valid period for including an event, meaning that the events occurred after 12 weeks since the drug was taken were not counted. This is an empirical choice for all kinds of ADEs. An alternative is to assign different length of valid period to different ADE, given the fact that some ADEs, such as vomiting, occur rather fast after taking the causing drug, while some ADEs, such as myocardial infarction, might appear after a while until the toxicity accumulates to a certain level to result in an ADE. Therefore, a fixed length of 12 weeks might lead to over-estimated number of ADEs which occur fast and under-estimated number of ADEs which occur slowly.

The results presented in this study provide insight into how the disproportionality methods perform differently with two representations of event counting and demonstrate the potential of using EPRs for pharmacovigilance using disproportionality analysis. They can also serve as a benchmark for comparison when other methods are applied to the same data in the future.

One obvious limitation of this study is that the case study only consists of one specific drug and two ways of representing events, which does not allow for drawing conclusions on what relative performance can be expected in general when comparing disproportionality methods and which way of using these methods on EPRs is the most suitable.

There are several possible directions for future research. First, experiments with several data sources will increase the credibility of arguments about which way of adapting disproportionality methods is more suitable. Second, when the disproportionality methods count the drug-event pairs or distinct patients, it may also be worth investigating that whether all events or patients should be weighted equally or not. For example, in the EL analysis, drugs that are taken closer to the event of interest should potentially have a higher influence than drugs that were taken much earlier; in the PL analysis, patients who are exposed to the drug of interest for a longer time or with a higher consumption in dosage were counted only once like those who only took the drug of interest once. More elaborate approaches which can take into account more information, such as weights described above or have different counting strategies could be proposed and evaluated. Third, it would also be interesting to compare disproportionality methods on EPRs with other means of detecting signals, such as machine learning techniques. Finally, the information required by the current disproportionality analysis is only the tip of the iceberg. The EPRs contains large amounts of information which are under-utilized, especially for the purpose of drug safety surveillance. For example, the demographical information and clinical measurements can be of importance when conducting stratification or adjustment to deal with confounding. Moreover, in reality, when a patient has one or more severe diseases, events such as headache or nausea might be omitted to be encoded as diagnoses by doctors. Instead, doctors might have mentioned such events in the patient's clinical notes. Mining the unstructured narrative of EPRs may bring a large contribution to pharmacovigilance research.

ACKNOWLEDGMENT

This work was partly supported by the project High Performance Data Mining for Drug Effect Detection at Stockholm University, funded by Swedish Foundation for Strategic Research under grant IIS11-0053.

REFERENCES

[1] C.D. Furberg and B. Pitt. Withdrawal of cerivastatin from the world market. *Curr Control Trials Cardiovasc Med*, 2:205–207, 2001.

[2] M. Suling and I. Pigeot. Signal detection and monitoring on longitudinal healthcare data. *Pharmaceutics*, 4:607–640, 2012.

[3] S.J. Evans, P.C. Waller, and S. Davis. Use of proportional reporting ratios (prrs) for signal generation from spontaneous adverse drug reaction reports. *Pharmacoepidemiology and Drug Safety*, 10:483–486, 2001.

[4] E.P. van Puijenbroek, A. Bate, L.G. Leufkens, M. Lindquist, R. Orre, and A.C. Egberts. A comparison of measures of disproportionality for signal detection in spontaneous reporting systems for adverse drug reactions. *Pharmacoepidemioloy and Drug Safety*, 11:3–10, 2002.

[5] A. Bate, M. Lindquist, I.R. Edwards, and et al. A bayesian neural network method for adverse drug reaction signal generation. *Eur J Clin Pharmac*, 54:315–321, 1998.

[6] W. DuMouchel. Bayesian data mining in large frequency tables, with an application to the fda

spontaneous reporting system. *The American Statistician*, 53:190–196, 1999.

[7] J.S. Almenoff, W. DuMouchel, L.A. Kindman, H. Yang, X, and D. Fram. Disproportionality analysis using empirical bayes data mining: a tool for the evaluation of drug interactions in the post-marketing setting. *Pharmacoepidemiology and Drug Safety*, 12:517–521, 2003.

[8] Nam-Kyong Choi, Yoosoo Chang, Ju-Young Kim, Yu-Kyong Choi, and Byung-Joo Park. Comparison and validaton of data-mining indices for signal detection: using the korean national health insurance claims database. *Pharmacoepidemiology and Drug Safety*, 20:1278–1286, 2011.

[9] M. Hauben, S. Horn, and L. Reich. Potential use of data-mining algorithms for the detection of 'surprise' adverse drug reactions. *Drug Safety*, 30:143–155, 2007.

[10] D. Madigan and P. Ryan. Commentary: What can we really learn from observational studies? the need for empirical assessment of methodology for active drug safety surveillance and comparative effectiveness research. *Epidemiology*, 22:629, 2011.

[11] J.R. Curtis, H. Cheng, E Delzell, and el al. Adaptation of bayesian data mining algorithms to longitudinal claims data. *Medical Care*, 46:969–975, 2008.

[12] G.N. Noren, A. Bate, J. Hopstadius, K. Star, and I.R. Edwards. Temporal pattern discovery for trends and transient effects: its application to patient records. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 963–971, Las Vegas, Nevada, USA, 2008.

[13] I. Zorych, D. Madigan, P. Ryan, and A. Bate. Disproportionality methods for pharmacovigilance in longitudinal observational databases. *Statistical methods in medical research*, 0:1–18, 2011.

[14] H. Dalianis, Hassel M., A. Henriksson, and M Skeppstedt. Stockholm epr corpus: A clinical database used to improve health care,. In *Swedish Language Technology Conference*, 2012.

[15] A Shibata and M. Hauben. Pharmacovigilance, signal detection and signal intelligence overview. In *Information Fusion (FUSION), Proceedings of the 14th International Conference*, Chicago, IL, USA, July 2011.

[16] M.J. Schuemie. Methods for drug safety signal detection in longitudinal observational databases: Lgps and leopard. *Pharmacoepidemiology and Drug Safety*, 20:292–299, 2011.

[17] M. Hauben and A. Bate. Decision support methods for the detection of adverse events in post-marketing data. *Drug discovery today*, 14:343–357, 2009.

[18] I. Ahmed and A. Poncet. Phvid, version 1.0.4; a r package for pharmacovigilance signal detection. *Available online: http://cran.r-project.org/web/packages/PhViD.*

[19] G.N. Noren, A. Bate, R. Orre, and I.R. Edwards. Extending the methods used to screen the who drug safety database towards analysis of complex associations and improve accuracy for rare events. *Statistics in Medicine*, 25:3740–3757, 2006.

[20] M. Kuhn, M. Campllos, I. Letunic, L.J. Jensen, and P. Bork. A side effect recsouce to capture phenotypic effects of the drugs. *Molecular Systems Biology*, 6:343, 2010.

[21] Christopher D Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to information retrieval*. Cambridge University Press Cambridge, 2008.