ORIGINAL REPORT

A modular, prospective, semi-automated drug safety monitoring system for use in a distributed data environment

Joshua J. Gagne*, Shirley V. Wang, Jeremy A. Rassen and Sebastian Schneeweiss

Division of Pharmacoepidemiology and Pharmacoeconomics, Department of Medicine, Brigham and Women's Hospital and Harvard Medical School, Boston, MA, 02120, USA

ABSTRACT

Purpose   The aim of this study was to develop and test a semi-automated process for conducting routine active safety monitoring for new drugs in a network of electronic healthcare databases.
Methods   We built a modular program that semi-automatically performs cohort identification, confounding adjustment, diagnostic checks, aggregation and effect estimation across multiple databases, and application of a sequential alerting algorithm. During beta-testing, we applied the system to five databases to evaluate nine examples emulating prospective monitoring with retrospective data (five pairs for which we expected signals, two negative controls, and two examples for which it was uncertain whether a signal would be expected): cerivastatin versus atorvastatin and rhabdomyolysis; paroxetine versus tricyclic antidepressants and gastrointestinal bleed; lisinopril versus angiotensin receptor blockers and angioedema; ciprofloxacin versus macrolide antibiotics and Achilles tendon rupture; rofecoxib versus non-selective non-steroidal anti-inflammatory drugs (ns-NSAIDs) and myocardial infarction; telithromycin versus azithromycin and hepatotoxicity; rosuvastatin versus atorvastatin and diabetes and rhabdomyolysis; and celecoxib versus ns-NSAIDs and myocardial infarction.
Results   We describe the program, the necessary inputs, and the assumed data environment. In beta-testing, the system generated four alerts, all among positive control examples (i.e., lisinopril and angioedema; rofecoxib and myocardial infarction; ciprofloxacin and tendon rupture; and cerivastatin and rhabdomyolysis). Sequential effect estimates for each example were consistent in direction and magnitude with existing literature.
Conclusions   Beta-testing across nine drug-outcome examples demonstrated the feasibility of the proposed semi-automated prospective monitoring approach. In retrospective assessments, the system identified an increased risk of myocardial infarction with rofecoxib and an increased risk of rhabdomyolysis with cerivastatin years before these drugs were withdrawn from the market. Copyright © 2014 John Wiley & Sons, Ltd.

KEY WORDS—prospective; active; surveillance; monitoring; distributed database; pharmacoepidemiology

Received 9 September 2013; Revised 22 February 2014; Accepted 24 February 2014

INTRODUCTION

Approximately one drug per year is withdrawn from the US market for safety reasons, with a mean time from approval to withdrawal of 6 years.[1] During this time, many millions of individuals can be exposed to medications with unknown safety profiles. Rofecoxib, a cyclooxygenase 2 inhibitor approved in 1999, was withdrawn from the US market in 2004—after 80 million people worldwide had used the drug—because of its association with myocardial infarction.[2] This association was subsequently confirmed by a meta-analysis of 11 observational database studies. However, nine of these studies were conducted after rofecoxib's withdrawal even though the drug and outcome data used in each study were routinely collected in near real-time in the electronic healthcare databases.[3] Prospective assessments analyzing the data as they accrued could have identified this association years before rofecoxib was withdrawn.[4]

Identifying potential drug safety concerns of new drugs as quickly as possible requires three changes to the traditional paradigm of single-database retrospect drug safety assessments.[5] It requires analyses that are simultaneously conducted across multiple databases to maximize sample size and in a sequential and semi-automated fashion to identify potential safety concerns as quickly as possible.[5] Several initiatives

*Correspondence to: J. J. Gagne, Division of Pharmacoepidemiology and Pharmacoeconomics, Department of Medicine, Brigham and Women's Hospital and Harvard Medical School, 1620 Tremont Street, Suite 3030, Boston, MA 02120, USA. E-mail: jgagne1@partners.org

around the world—including the Exploring and Understanding Adverse Drug Reactions project,[6] the US Food and Drug Administration's Mini-Sentinel program[7] and the Observational Medical Outcomes Partnership[8]—are building large networks of electronic healthcare databases in which to assess drug safety. The OMOP and Mini-Sentinel databases, e.g., each comprise more than 100 million covered lives.[9,10] However, little attention has been paid to the development of scalable programs specifically designed to analyze data across these networks to validly identify adverse drug effects as quickly as possible as data accrue prospectively in these networks after new drug approval.

We describe a semi-automated and scalable process for conducting routine active drug safety monitoring to rapidly and validly assess potential associations between pre-specified drug-outcome pairs across a network of electronic healthcare databases.

METHODS

Conceptual epidemiologic framework

The program is built on accepted epidemiological and statistical principles and is designed to perform semi-automated, distributed, sequential, propensity score (PS)-matched, new user, parallel, active comparator cohort analyses (Figure A1).[11–14] The cohort design is commonly used in pharmacoepidemiology to assess the safety of drugs and has distinct advantages for a wide range of drug safety questions.[15] Comprehensive guidance on selection of design and analysis methods for a given routine monitoring question can be found elsewhere.[16–18] Details about the epidemiologic and statistical principles used in this approach can be found in the Appendix and elsewhere.[12–16,19–23]

Distributed data environment

The system is compatible with electronic healthcare databases that have been converted into a widely used common data model (CDM).[24] To maximize data for analysis at any given time after a drug enters the market, the system is designed to perform analyses that are both distributed across multiple databases and sequential as more patients exposed to the drug accrue in the databases. The standardized code incorporating pre-specified inputs (see succeeding text) can be sent to multiple database holders who separately analyze their data behind their own firewall. The code can be run iteratively on each database as new data become available. Data can then be aggregated both over time and across databases at a central hub.

Program architecture

The standardized program comprises five main modules: (i) a cohort identification module;[25] (ii) an adjustment module; (iii) a diagnostics module; (iv) an aggregation module; and (v) an alerting algorithm module. The first three modules are run in a distributed fashion behind each database holder's firewall. Analytic diagnostic information and aggregated and de-identified information are transmitted from each database holder to the central hub, which further aggregates data across sites by estimating a site-stratified summary point estimate (aggregation module) and applies a sequential alerting algorithm (alerting algorithm module). Steps involved in each of these modules are described in more detail in the succeeding text. For practical purposes, we designed the program in a "one-stop-shopping" fashion such that all information is obtained from each database in a single query in each sequential period, limiting the amount of time and resources required by the database holder. The first four modules are written as SAS macros, and the current alerting algorithm module is written in R. As such, the full program requires each database holder to run SAS and the central hub to run both SAS and R. SAS is an industry and regulatory standard for electronic healthcare database analyses.

The five modular components of the program perform 10 key steps to implement the epidemiologic and statistical methods (Figure 1):

1. *Investigator(s) specifies clinical and epidemiological inputs*

The first two steps outlined in Figure 1 require expert determination of appropriate clinical and epidemiological inputs and some manual operation, hence the term "semi-automated" active surveillance. The basic cohort design requires specification of the drug of interest, the comparator of interest, the duration of the washout period to ensure that patients are "new users," codes for pre-defined covariates (including both potential confounders and indicators of subgroups of interest), the risk window following treatment initiation over which outcomes will be ascertained, codes for the outcome(s) of interest, and the calendar period over which new drug users will be identified.

2. *Central hub transmits code to each database holder for execution*

Once the inputs have been specified, the central hub sends a package with the modular code and the drug, outcome, and covariate files to each database holder.

Each database holder executes the same standardized SAS code on their CDM-formatted data.

3. Automatic cohort creation and outcome identification

The code automatically implements the first three modules, which involves steps 3–7. First, the code implements a cohort identification module developed by the Mini-Sentinel pilot project.[25] This program draws in the drug files to identify new users of the drug of interest and new users of a comparator drug. This cohort is then subject to the adjustment module, which comprises steps 4–7.

4. Automatic covariate ascertainment

The adjustment module queries each cohort patient's electronic data history to identify pre-specified and empirically identified covariates. Pre-specified covariates are determined by codes indicated in the pre-specified covariate SAS data files. These will typically include demographic variables, such as age and sex, and algorithms to identify known risk factors for the outcome(s) of interest. The program also automatically identifies empirical covariates using the high dimensional PS (hd-PS) algorithm.[13] The adjustment module can also identify whether and when cohort patients experience the outcome of interest. The program identifies outcomes with a "plug-in" macro that can accommodate any outcome definition that uses elements in the CDM. It also allows for both as-treated and intention-to-treat exposure definitions.

The program automatically computes the combined comorbidity score for each patient. This score has been shown to outperform both the Charlson index and the Elixhauser comorbidity system.[26] The program also identifies measures of health service utilization intensity, including number of unique drug codes dispensed, number of physician visits, and number of hospitalizations in the baseline period.[27]

5. Automatic PS estimation

Once the program has identified variables for adjustment, it includes them as independent variables in logistic regression models that estimate three database-specific PSs for each patient. PS_1 is based only on pre-defined variables, PS_2 based on only hd-PS-identified empirical covariates, and PS_3 includes both sets of covariates. Importantly, the latter two PS models will include different variables across each database. The hd-PS algorithm will empirically identify the most relevant confounders within a dataset and incorporate them into the site-specific PS model.

This will improve adjustment because confounding can manifest differently across sites due to differences in treatment decision processes and data content.[20,28]

6. Automatic PS matching

The program then uses each set of PS values to match patients within each data base Initiators of the drug of interest are matched to initiators of the comparator using multiple matching strategies. The first version of the program conducts 1:1 matching using each of three calipers of 0.025, 0.05, and 0.10 units on the PS scale. The program also conducts variable ratio matching with a ratio of up to 100:1. As the program is iterated as new data accrue in each database, it re-estimates the PS on all initiators that have accrued in each database by the time of each data update, but matches only initiators in each new batch of data.[29] That is, once patients are matched and included in an analysis, they remain as such regardless of how the future data change their PS values.

7. Automatic diagnostic and data file preparation

Once the database-specific analytic steps are complete, the program creates a de-identified patient-level aggregated transfer data file for each database. This transfer file includes randomly generated patient identification numbers, an indicator of exposure status, an indicator of whether patients experience the outcome of interest, the person-time of follow-up between the index date and date of censoring (i.e., first of end of the exposure period, outcome occurrence date, death, disenrollment, or end of analysis period), the three PS values, and a matched set indicator for each matching strategy. This transfer file is considered de-identified and anonymous according to the Health Insurance Portability and Accountability Act (HIPAA) and does not require HIPAA waivers.[30] When subgroup analyses are anticipated, the file will also include indicators for each patient's subgroup status. An identical file that also contains patient identifiers remains behind each database holder's firewall.

The program generates diagnostic information describing the discrimination of the PS models (i.e., c-statistics), plots depicting the overlap in PS distributions between treatment groups before and after matching (Figure A2), and summary tables and figures describing the baseline demographic and clinical characteristics, and the extent to which these variables are balanced individually and overall.

8. Central hub evaluates diagnostics and determines whether to aggregate

Before looking at any effect estimates characterizing the association between the drug and outcome, the

Central hub

1. Define parameters

2a. Populate modular code and feeder files

2b. Select alerting algorithm(s)

3. Transmit modular code and feeder files to data holders

Study drug file	Covariate file	Outcome file

Multiple databases

Common data model

Common data model

Common data model

4. Identify cohort and outcomes

5. Ascertain covariates

6. Estimate propensity score

7. Summarize data

8. Transmit SAS log, aggregate data and diagnostic information to coordinating center

9. Evaluate diagnostics

10. Aggregate across data partners

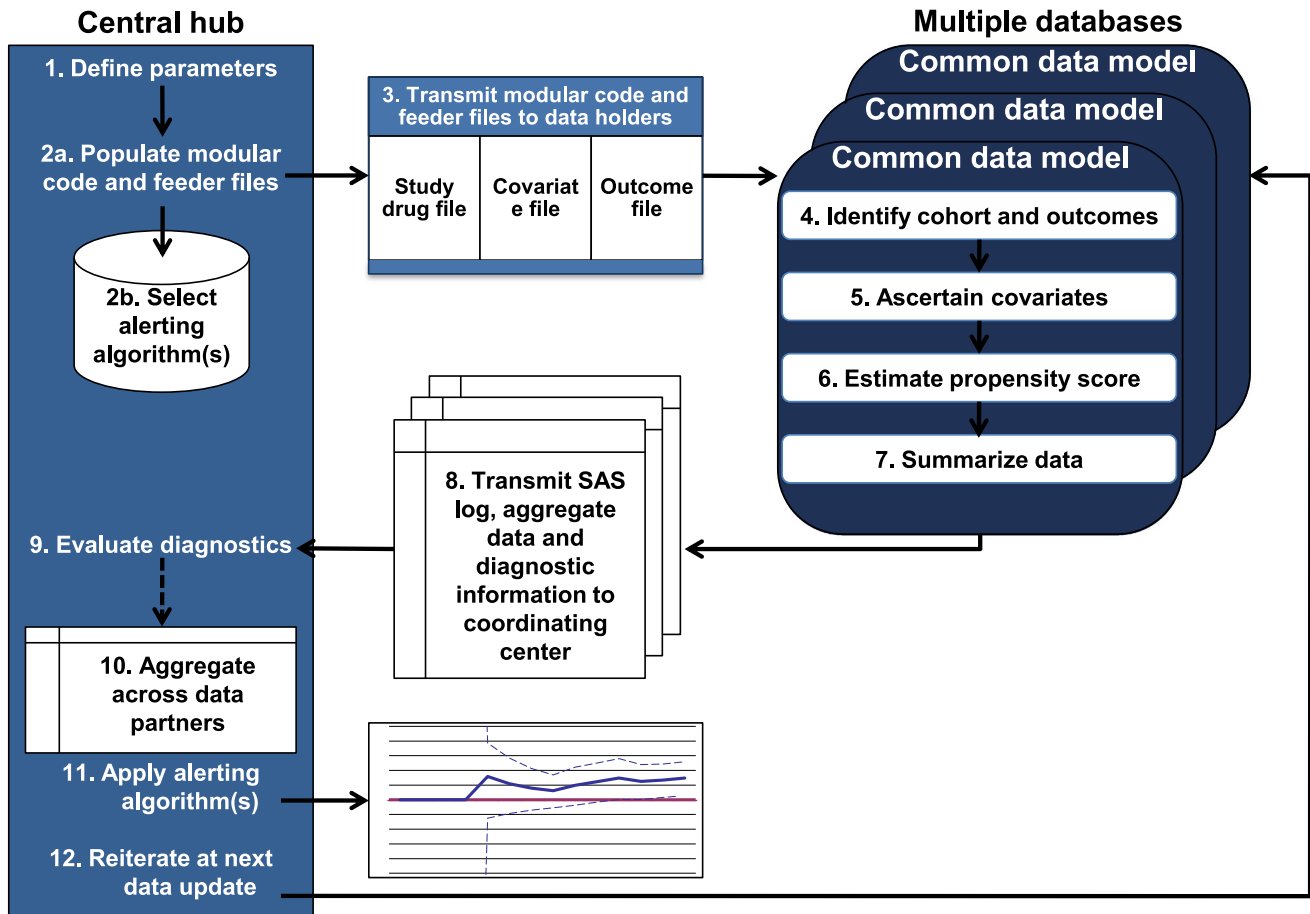11. Apply alerting algorithm(s)

12. Reiterate at next data update

Figure 1. Operational steps to implement modular program to perform semi-automated propensity score-matched new user cohort analyses

central hub can review the diagnostics from each database and compare descriptive data across the databases. This enables identification of issues such as under- or over-ascertainment of covariates, exposures, or outcomes in one or more databases, identification of substantial differences in demographic and clinical characteristics of cohort patients across databases, and comparisons of the extent to which covariate balance is achieved in each database. As with step 1, this step requires human input. The central hub can then decide whether and which data to aggregate into a single summary effect estimate. Although these diagnostic steps are recommended, they are not required to performing monitoring with the system.

9. *Central hub aggregates data and applies code for sequential alerting algorithm*

After deciding to combine data, the central hub uses the aggregation module to automatically aggregate data across databases and over time. Multiple measures of association can be calculated, including the hazard ratio and rate difference. The central hub

can further stratify the outcome models by subgroup indicators.[31] The aggregation module then prepares the inputs for assessing whether a safety alert should be raised.

Currently, the sequential alerting process uses R code developed by Kulldorff and Silva to perform formal statistical tests using the maximized sequential probability ratio test (maxSPRT) developed by Kulldorff *et al*.[4,32] The maxSPRT is a continuous monitoring algorithm that was developed for monitoring vaccine safety in observational data and easily accommodates the potentially irregular database updating schedule.[33] The sequential matched cohort approach also accommodates other alerting algorithms that can enable ongoing monitoring beyond the pre-specified runtime of formal statistical hypothesis tests.[12,34]

10. *Central hub uses results to determine whether to iterate*

The outputs of the aggregation and sequential alerting programs include effect estimates and alerting criteria. The central hub uses this information to determine

whether to continue prospective monitoring, which would involve rerunning steps 3 through 10 when the databases are refreshed with new data. Each time the program is iterated, the database holders rerun the program on their growing databases and create a de-identified aggregated dataset that includes all cohort members to date. The central hub then aggregates and analyzes the PS-balanced appended dataset from each databases when conducting sequential monitoring.

Empirical examples

We have used empirical data from nine exposure-outcome pairs to beta-test the program at various points throughout development.[11,12,14] Here, we summarize the results to describe the overall performance of the program. For each example, we emulated prospective monitoring of each drug of interest. For six examples (cerivastatin, telithromycin, rofecoxib, celecoxib, and the two rosuvastatin examples), we began monitoring at the time the drug of interest entered the market. Table 1 summarizes these examples and provides our *a prior* expectation about whether the example was one in which we expected to observe a signal. Five pairs were examples for which we expected signals, two pairs were negative controls, and two pairs were examples for which it was uncertain whether a signal would be expected.

We used data from five sources, which are described in the Appendix. For simplicity, we used a conservative maxSPRT critical value based on an upper limit of surveillance defined by the occurrence of 2000 events. We used a 96-terabyte, 96-processor IBM-Netezza parallel-computing database supercomputer with a Unix pre-processing unit to test and run the modular program using SAS 9.2 and R.

RESULTS

In emulated prospective monitoring, we observed an increased rate of myocardial infarction among rofecoxib

Table 1. Overview of empirical examples

Drug of interest	Comparator	Outcome(s)	Follow-up period	Data source(s)*	Signal expected?
Cerivastatin	Atorvastatin	Rhabdomyolysis	Duration of index treatment	(1) New Jersey Medicare data linked to pharmacy assistance program; (2) Pennsylvania Medicare data linked to pharmacy assistance program	Yes
Paroxetine	Tricyclic antidepressants with low affinity for serotonin receptors	Gastrointestinal bleed	90 days following treatment initiation	HealthCore Integrated Research Database (HIRD)	Yes
Lisinopril	Angiotensin receptor blockers	Angioedema	Duration of index treatment plus 30 days	HIRD	Yes
Ciprofloxacin	Macrolide antibiotics	Achilles tendon rupture	183 days following treatment initiation	HIRD	Yes
Rofecoxib	Non-selective non-steroidal anti-inflammatory drugs	Myocardial infarction	180 days following treatment initiation	(1) New Jersey Medicare data linked to pharmacy assistance program; (2) Pennsylvania Medicare data linked to pharmacy assistance program; (3) Medicaid Analytic eXtract (MAX) (covering Medicaid beneficiaries in 48 states).	Yes
Telithromycin	Azithromycin	Hepatotoxicity	60 days following treatment initiation	(1) HIRD; (2) New Jersey Medicare data linked to pharmacy assistance program; (3) Pennsylvania Medicare data linked to pharmacy assistance program	Uncertain
Rosuvastatin	Atorvastatin	Diabetes	Duration of index treatment	(1) HIRD; (2) New Jersey Medicare data linked to pharmacy assistance program; (3) Pennsylvania Medicare data linked to pharmacy assistance program	Uncertain
Rosuvastatin	Atorvastatin	Rhabdomyolysis	Duration of index treatment plus 60 days	(1) HIRD; (2) New Jersey Medicare data linked to pharmacy assistance program; (3) Pennsylvania Medicare data linked to pharmacy assistance program	No
Celecoxib	Non-selective non-steroidal anti-inflammatory drugs	Myocardial infarction	180 days following treatment initiation	(1) New Jersey Medicare data linked to pharmacy assistance program; (2) Pennsylvania Medicare data linked to pharmacy assistance program	No

*All data sources comprise administrative claims data including demographic data, medical claims from healthcare providers and facilities, and outpatient pharmacy dispensing records.

initiators as compared to PS-matched initiators of ns-NSAIDs (Figure 2). The overall rate difference was 2.24 myocardial infarction events (nominal 95% confidence interval, 1.10–3.38) per 1000 person-years comparing rofecoxib initiators versus ns-NSAID initiators. The corresponding hazard ratio was 1.19 (nominal 95% confidence interval, 1.09–1.30). This is consistent with previous formal pharmacoepidemiologic studies.[35,36] Overall hazard ratios for each example are listed in Table 2. For each example, we observed results that were consistent in direction and magnitude with our *a prior* expectation based on previous pharmacoepidemiologic studies.[3,12,14,19,37–40]

The maxSPRT generated timely alerts for lisinopril, rofecoxib, ciprfloxaxin, and cerivastatin, all true positives (Figure 3). No other alerts were generated during the monitoring time frame, but the maxSPRT had not yet reached the end of its pre-specified runtime, meaning that the algorithm had not yet achieved its full statistical power.

Table A1 presents an example table describing baseline characteristics of matched rofecoxib and ns-NSAID initiators from the first monitoring period from the MAX database. As part of the diagnostic output, Figure A2 presents example histograms and smoothed densities of the PS values among rofecoxib initiators and ns-NSAID initiators, both before and after matching by the PS. Panel A of Figure A2 illustrates substantial overlap in PS distributions between patients in the two treatment groups, which is a prerequisite for valid effect estimation. As expected, the PS distributions are highly overlapping after matching (Panel B).

Table 2. Summary of results of nine examples

Example	Total number of observed events	Hazard ratio (95% confidence interval) at end of monitoring
Cerivastatin	6	∞*
Paroxetine	38	1.72 (0.89, 3.32)
Lisinopril	344	1.92 (1.54, 2.41)
Ciprofloxacin	22	1.45 (0.62, 3.39)
Rofecoxib	1937	1.19 (1.09, 1.30)
Telithromycin	41	1.26 (0.68, 2.33)
Rosuvastatin (diabetes)	1914	0.95 (0.87, 1.04)
Rosuvastatin (rhabdomyolysis)	8	0.39 (0.08, 1.94)
Celecoxib	226	1.16 (0.83, 1.64)

*All events occurred among cerivastatin initiators.

Table A2 presents the run time required for cohort identification module and the adjustment module, and each computationally intensive component of the adjustment module, in the large Optum database. In this representative example, the total time to analyze data for more than 150 000 patients drawn from a pool of about 50 million covered lives was less than 3.5 h. The bulk of the computational time resided with extracting the analytic cohort from the main database.

DISCUSSION

We have developed a program to perform distributed, sequential, PS-matched, new user cohort studies in a reproducible and scalable manner. The SAS-based modular program can be run simultaneous on multiple databases converted into the freely available Mini-Sentinel CDM[24] and iteratively as experience with

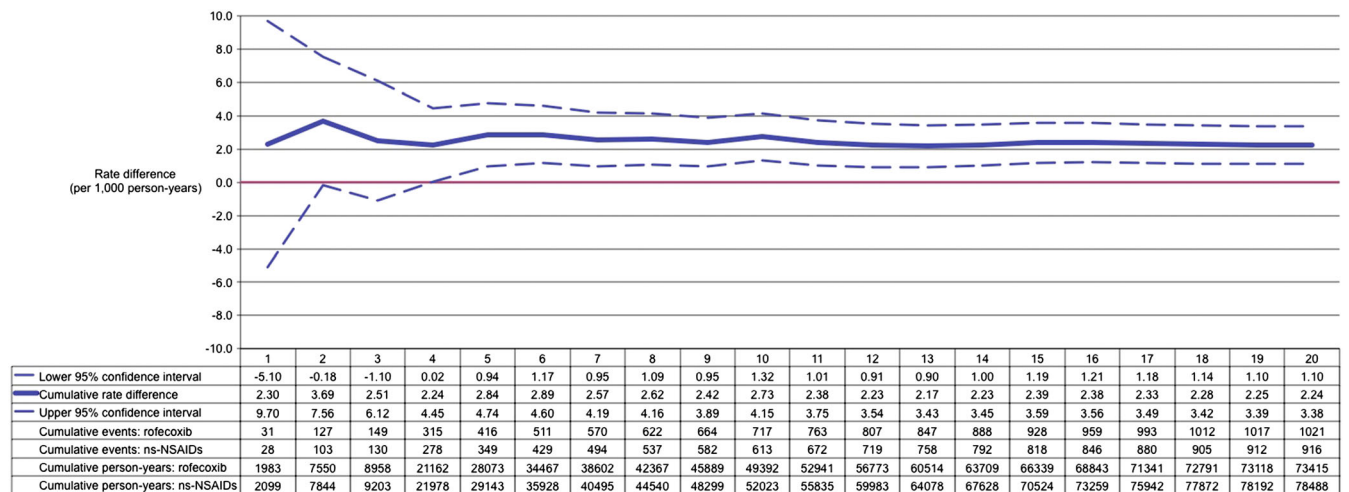|	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Lower 95% confidence interval	-5.10	-0.18	-1.10	0.02	0.94	1.17	0.95	1.09	0.95	1.32	1.01	0.91	0.90	1.00	1.19	1.21	1.18	1.14	1.10	1.10
Cumulative rate difference	2.30	3.69	2.51	2.24	2.84	2.89	2.57	2.62	2.42	2.73	2.38	2.23	2.17	2.23	2.39	2.38	2.33	2.28	2.25	2.24
Upper 95% confidence interval	9.70	7.56	6.12	4.45	4.74	4.60	4.19	4.16	3.89	4.15	3.75	3.54	3.43	3.45	3.59	3.56	3.49	3.42	3.39	3.38
Cumulative events: rofecoxib	31	127	149	315	416	511	570	622	664	717	763	807	847	888	928	959	993	1012	1017	1021
Cumulative events: ns-NSAIDs	28	103	130	278	349	429	494	537	582	613	672	719	758	792	818	846	880	905	912	916
Cumulative person-years: rofecoxib	1983	7550	8958	21162	28073	34467	38602	42367	45889	49392	52941	56773	60514	63709	66339	68843	71341	72791	73118	73415
Cumulative person-years: ns-NSAIDs	2099	7844	9203	21978	29143	35928	40495	44540	48299	52023	55835	59983	64078	67628	70524	73259	75942	77872	78192	78488

Figure 2. Emulated prospective monitoring results for myocardial infarction among initiators of rofecoxib versus initiators of non-selective non-steroidal anti-inflammatory drugs
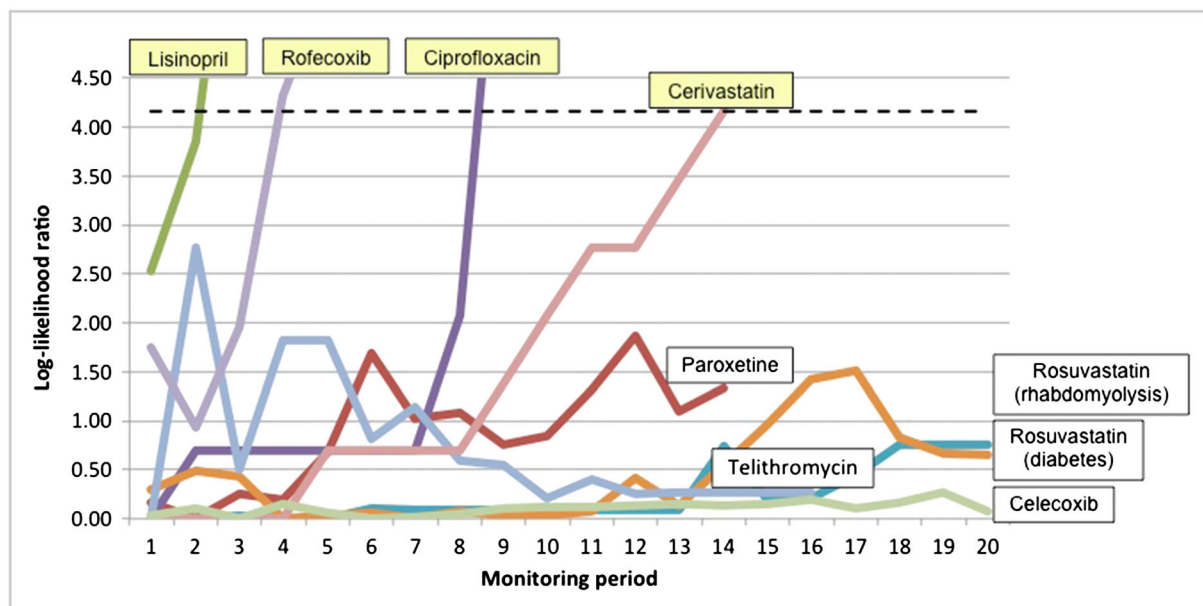
Figure 3.   Results of the maximized sequential probability ratio test for nine drug-outcome monitoring examples

new medical products accumulates in the databases. To improve the robustness of the output and to support decision-making, the program is based on validated design and analysis methods that are commonly used in pharmacoepidemiology to address the limitations of secondary electronic healthcare data. It uses PSs to achieve covariate balance between treatment groups. PSs possess many attractive properties in the active safety monitoring setting, including that they easily address non-exchangeability on measured covariates, which is a prerequisite for causal inference (i.e., matching implicitly excludes patients in areas of non-overlap), they enable evaluation of multiple outcomes per exposure (as with the rosuvastatin example), and they simplify data aggregation and enable application of a wide range of sequential alerting algorithms, such as those for matched data (e.g., maxSPRT) or those for continuous monitoring (e.g., statistical process control rules).

In beta-testing, across nine examples and using various data sources, the modular program consistently produced results in line with expectation with respect to both the direction and magnitude of associations. Our beta-test analyses were not sufficiently powered to generate an alert for paroxetine, but the log-likelihood ratio was trending upwards, suggesting that it might generate an alert with continued monitoring or with larger distributed data networks, such as with the Mini-Sentinel distributed database.[7] Once the system generates an alert for a particular example, formal hypothesis testing with the maxSPRT ends. However,

continued monitoring of the point estimate can provide additional information as part of signal follow-up activities. With ciprofloxacin, the system generated an alert in the ninth monitoring period, but continued monitoring produced a point estimate that was closer to the null and with a nominal 95% confidence interval that included one. Additional analyses would be needed to determine the cause of such a pattern, such as possible changes in drug prescribing over time. Overall, when sufficiently powered, the program generated alerts for known positive associations but did not generate any alerts for any negative control drug-outcome pairs. Subsequent evaluations of the program against a large number of positive and negative controls will provide more comprehensive insight into its overall operating characteristics.

The program is designed to ensure that identifiable patient-level information remains behind each database holder's firewall. Only de-identified, aggregate data are shared with the central hub. However, the data are shared in a way that enables diagnostics and preserves substantial flexibility for the central hub at the aggregation step. This permits the central hub to conduct subgroup analyses based on pre-specified subgroup indicators to exclude data from selected databases in which sufficient covariate balance is not achieved or when data issues arise, and to conduct *post hoc* sensitivity analyses by incorporating the PSs into the analysis in different ways. Additionally, other methods, such as disease risk scores, can be built into the existing cohort

framework and can be used either along with or in place of PSs.[41] Disease risk scores derived in a recent and similar historical population have shown promise for confounding control in the setting of newly approved drugs.

The safety monitoring process that we have described is easily scalable along several dimensions, including (i) the number of drug-outcome pairs monitored, including multiple outcomes per drug; (ii) the number of subgroups evaluated; (iii) the number of sequential analyses conducted; and (iv) the number of databases used. The rate-limiting step in deploying the system will likely be the decision-making process required to initiate each monitoring activity (i.e., Step 1). Importantly, the program requires clinical and epidemiological expertise for determining, first, whether the program should be applied to a specific scenario (i.e., whether it is the right tool for the job) and, second, the most appropriate inputs. While this process can be structured and expedited,[16–18] it requires considerable clinical and epidemiologic input to ensure that each monitoring activity will answer the most relevant clinical and regulatory question. Otherwise, the system can be run automatically with very little investigator input. However, if desired, the system does allow for investigator input at various steps throughout the process, as we have described.

The program has some important limitations. First, the methods that it implements may not be optimal or even appropriate for all monitoring scenarios. Clinical and epidemiologic expertise is required to ensure that the program is applied to scenarios that it is well suited to address. Second, even when the program is appropriately applied, it has inherent limitations. These include, but are not limited to, the need for sufficient numbers of initiators of the new drug of interest to fit propensity scores in the early marketing period and the lack of analytic ability to address time-varying confounding.

In conclusion, we have developed and tested a scalable, semi-automated process with modular programs to conduct rapid prospective drug safety monitoring across a distributed data network and iteratively as new data accrue in the network. The program integrates widely used pharmacoepidemiologic design and analysis tools in a modular fashion. We are currently testing this modular program with live Data Partners in the Mini-Sentinel Distributed Data network. The program, including code and technical specifications, has been made publicly available through the Mini-Sentinel program website—http://mini-sentinel.org/methods/methods_development/details.aspx?ID=1045.

CONFLICT OF INTEREST

Dr. Gagne is principal investigator of an unrelated investigator-initiated grant to the Brigham and Women's Hospital from Novartis. Dr. Rassen holds shares of Aetion, Inc. Dr. Schneeweiss is principal investigator of the Harvard-Brigham Drug Safety and Risk Management Research Center funded by FDA. Dr. Schneeweiss is a consultant to WHISCON LLC and to Aetion, Inc., of which he also owns shares. He is principal investigator of unrelated investigator-initiated grants to the Brigham and Women's Hospital from Novartis and Boehringer-Ingelheim.

KEY POINTS

- A scalable, semi-automated approach for conducting routine active drug safety monitoring to rapidly and validly assess potential associations between pre-specified drug-outcome pairs in electronic healthcare data is needed.
- The authors built a program that semi-automatically performs cohort identification, confounding adjustment, diagnostic checks, aggregation and effect estimation across multiple databases, and application of a sequential alerting algorithm.
- In beta-testing, the system generated four alerts, all among positive control examples (i.e., lisinopril and angioedema; rofecoxib and myocardial infarction; ciprofloxacin and tendon rupture; and cerivastatin and rhabdomyolysis).
- In retrospective assessments, the system identified an increased risk of myocardial infarction with rofecoxib and an increased risk of rhabdomyolysis with cerivastatin years before these drugs were withdrawn from the market.

ETHICS STATEMENT

Empirical analyses were approved by the Brigham and Women's Hospital Institutional Review Board.

ACKNOWLEDGEMENTS

REFERENCES

1. Qureshi ZP, Seoane-Vazquez E, Rodriguez-Monguio R, Stevenson KB, Szeinbach SL. Market withdrawal of new molecular entities approved in the United States from 1980 to 2009. *Pharmacoepidemiol Drug Saf* 2011; **20**: 772–772.
2. Topol EJ. Failing the public health--rofecoxib, Merck, and the FDA. *N Engl J Med* 2004; **351**: 1707–1709.
3. McGettigan P, Henry D. Cardiovascular risk and inhibition of cyclooxygenase: a systematic review of the observational studies of selective and nonselective inhibitors of cyclooxygenase 2. *JAMA* 2006; **296**: 1633–1644.
4. Brown JS, Kulldorff M, Chan KA, et al. Early detection of adverse drug events within population-based health networks: application of sequential testing methods. *Pharmacoepidemiol Drug Saf* 2007; **16**: 1275–1284.
5. Gagne JJ. You can observe a lot (about medical products) by watching (those who use them). *Epidemiology* 2013; **24**: 700–702.
6. Trifiro G, Pariente A, Coloma PM, et al. Data mining on electronic health record databases for signal detection in pharmacovigilance: which events to monitor? *Pharmacoepidemiol Drug Saf* 2009; **18**: 1176–1184.
7. Platt R, Carnahan RM, Brown JS, et al. The U.S. Food and Drug Administration's Mini-Sentinel program: status and direction. *Pharmacoepidemiol Drug Saf* 2012; **21**(Suppl 1): 1–8.
8. Stang PE, Ryan PB, Racoosin JA, et al. Advancing the science for active surveillance: rationale and design for the Observational Medical Outcomes Partnership. *Ann Intern Med* 2010; **153**: 600–606.
9. Toh S, Baker MA, Brown JS, Kornegay C, Platt R, Mini-Sentinel I. Rapid assessment of cardiovascular risk among users of smoking cessation drugs within the US Food and Drug Administration's Mini-Sentinel program. *JAMA Intern Med* 2013; **173**: 817–819.
10. Observational Medical Outcomes Partnership: OMOP Web RL. 2013. [Accessed on February 22, 2014] at http://omop.fnih.org/.
11. Gagne JJ, Glynn RJ, Rassen JA, et al. Active safety monitoring of newly marketed medications in a distributed data network: application of a semi-automated monitoring system. *Clin pharmacol Therapeut* 2012; **92**: 80–86.
12. Gagne JJ, Rassen JA, Walker AM, Glynn RJ, Schneeweiss S. Active safety monitoring of new medical products using electronic healthcare data: selecting alerting rules. *Epidemiology* 2012; **23**: 238–246.
13. Schneeweiss S, Gagne JJ, Glynn RJ, Ruhl M, Rassen JA. Assessing the comparative effectiveness of newly marketed medications: methodological challenges and implications for drug development. *Clin Pharmacol Ther* 2011; **90**: 777–790.
14. Wahl PM, Gagne JJ, Wasser TE, et al. Early steps in the development of a claims-based targeted healthcare safety monitoring system and application to three empirical examples. *Drug Saf* 2012; **35**: 407–416.
15. Schneeweiss S. A basic study design for expedited safety signal evaluation based on electronic healthcare data. *Pharmacoepidemiol Drug Saf* 2010; **19**: 858–868.
16. Gagne JJ, Fireman B, Ryan PB, et al. Design considerations in an active medical product safety monitoring system. *Pharmacoepidemiol Drug Saf* 2012; **21**(Suppl 1): 32–40.
17. Taxonomy for monitoring methods within a medical product safety surveillance system: report of the Mini-Sentinel Taxonomy Project Work Group. 2010. Available at http://www.mini-sentinel.org/work_products/Statistical_Methods/Mini-Sentinel_FinalTaxonomyReport.pdf [9 September 2013].
18. Taxonomy for monitoring methods within a medical product safety surveillance system: year two report of the Mini-Sentinel Taxonomy Project Workgroup. 2012. Available at http://www.mini-sentinel.org/work_products/Statistical_Methods/Mini-Sentinel_Methods_Taxonomy-Year-2-Report.pdf [9 September 2013].
19. Gagne JJ, Glynn RJ, Rassen JA, et al. Active safety monitoring of newly marketed medications in a distributed data network: application of a semi-automated monitoring system. *Clin Pharmacol Ther* 2012; **92**: 80–86.
20. Rassen JA, Avorn J, Schneeweiss S. Multivariate-adjusted pharmaco-epidemiologic analyses of confidential information pooled from multiple health care utilization databases. *Pharmacoepidemiol Drug Saf* 2010; **19**: 848–857.
21. Rassen JA, Glynn RJ, Brookhart MA, Schneeweiss S. Covariate selection in high-dimensional propensity score analyses of treatment effects in small samples. *Am J Epidemiol* 2011; **173**: 1404–1413.
22. Rassen JA, Solomon DH, Curtis JR, Herrinton L, Schneeweiss S. Privacy-maintaining propensity score-based pooling of multiple databases applied to a study of biologics. *Med Care* 2010; **48**: S83–S89.
23. Schneeweiss S, Rassen JA, Glynn RJ, Avorn J, Mogun H, Brookhart MA. High-dimensional propensity score adjustment in studies of treatment effects using health care claims data. *Epidemiology* 2009; **20**: 512–522.
24. Curtis LH, Weiner MG, Boudreau DM, et al. Design considerations, architecture, and use of the Mini-Sentinel distributed data system. *Pharmacoepidemiol Drug Saf* 2012; **21**(Suppl 1): 23–31.
25. Modular Program 3: frequency of select events during exposure to a drug/procedure group of interest. 2013. Available at http://www.mini-sentinel.org/work_products/Data_Activities/Mini-Sentinel-Modular_Program_3-Documentation.pdf'
26. Gagne JJ, Glynn RJ, Avorn J, Levin R, Schneeweiss S. A combined comorbidity score predicted mortality in elderly patients better than existing scores. *J Clin Epidemiol* 2011; **64**: 749–759.
27. Schneeweiss S, Seeger JD, Maclure M, Wang PS, Avorn J, Glynn RJ. Performance of comorbidity scores to control for confounding in epidemiologic studies using claims data. *Am J Epidemiol* 2001; **154**: 854–864.
28. Rassen JA, Schneeweiss S. Using high-dimensional propensity scores to automate confounding control in a distributed medical product safety surveillance system. *Pharmacoepidemiol Drug Saf* 2012; **21**(Suppl 1): 41–49.
29. Gagne JJ, Wang S, Rassen JA, Glynn RJ, Schneeweiss S. Propensity scores in sequential monitoring of new drugs: evaluation of dynamic matching. *Pharmacoepidemiol Drug Saf* 2012; **21**: 328.
30. Evaluating strategies for data sharing and analyses in distributed data settings. 2012. Available at http://www.mini-sentinel.org/work_products/Statistical_Methods/Mini-Sentinel_Methods_Evaluating-Strategies-for-Data-Sharing-and-Analyses.pdf [9 September 2013].
31. Rassen JA, Glynn RJ, Rothman KJ, Setoguchi S, Schneeweiss S. Applying propensity scores estimated in a full cohort to adjust for confounding in subgroup analyses. *Pharmacoepidemiol Drug Saf* 2012; **21**: 697–709.
32. Kulldorff M, Davis RL, Kolczak M, Lewis E, Lieu T, Platt R. A maximized sequential probability ratio test for drug and vaccin safety surveillance. *Seq Anal* 2011; **3**: 58–78.
33. Gagne JJ, Bykov K, Willke RJ, Kahler KH, Subedi P, Schneeweiss S. Treatment dynamics of newly marketed drugs and implications for comparative effectiveness research. *Value Health* 2013; **16**: 1054–1062.
34. Gagne JJ, Walker AM, Glynn RJ, Rassen JA, Schneeweiss S. An event-based approach for comparing the performance of methods for prospective medical product monitoring. *Pharmacoepidemiol Drug Saf* 2012; **21**: 631–639.
35. Solomon DH, Avorn J, Sturmer T, Glynn RJ, Mogun H, Schneeweiss S. Cardiovascular outcomes in new users of coxibs and nonsteroidal antiinflammatory drugs: high-risk subgroups and time course of risk. *Arthritis Rheum* 2006; **54**: 1378–1389.
36. Solomon DH, Schneeweiss S, Glynn RJ, et al. Relationship between selective cyclooxygenase-2 inhibitors and acute myocardial infarction in older adults. *Circulation* 2004; **109**: 2068–2073.
37. Toh S, Reichman ME, Houstoun M, et al. Comparative risk for angioedema associated with the use of drugs that target the renin-angiotensin-aldosterone system. *Arch Intern Med* 2012; **172**: 1582–1589.
38. Seeger JD, West WA, Fife D, Noel GJ, Johnson LN, Walker AM. Achilles tendon rupture and its association with fluoroquinolone antibiotics and other potential risk factors in a managed care population. *Pharmacoepidemiol Drug Saf* 2006; **15**: 784–792.
39. Graham DJ, Staffa JA, Shatin D, et al. Incidence of hospitalized rhabdomyolysis in patients treated with lipid-lowering drugs. *JAMA* 2004; **292**: 2585–2590.
40. Graham DJ, Campen D, Hui R, et al. Risk of acute myocardial infarction and sudden cardiac death in patients treated with cyclo-oxygenase 2 selective and non-selective non-steroidal anti-inflammatory drugs: nested case-control study. *Lancet* 2005; **365**: 475–481.
41. Glynn RJ, Gagne JJ, Schneeweiss S. Role of disease risk scores in comparative effectiveness research with emerging therapies. *Pharmacoepidemiol Drug Saf* 2012; **21**(Suppl 2): 138–147.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publishers web-site.