PEER REVIEW HISTORY

ARTICLE DETAILS

TITLE (PROVISIONAL)	Optimising computerised decision support to transform medication safety and reduce prescriber burden: Study protocol for a mixed-methods evaluation of drug-drug interaction alerts
AUTHORS	Baysari, Melissa; Zheng, Wu Yi; Li, Ling; Westbrook, Johanna; Day, Richard; Hilmer, Sarah; Van Dort, Bethany; Hargreaves, Andrew; Kennedy, Peter; Monaghan, Corey; Doherty, Paula; Draheim, Michael; Nair, Lucy; Samson, Ruby

VERSION 1 – REVIEW

REVIEWER	Eric Shelov The Children's Hospital of Philadelphia
REVIEW RETURNED	15-Oct-2018

GENERAL COMMENTS	Thank you for this submission. This study represents a significant undertaking and I look forward to the results. Very well written and well researched background. There have not been many prospective studies of medication alerts on this scale and we lack a clear roadmap for how to perform this kind of intervention and analysis. This approach is quite reasonable, but I have two points of concern: 1) The concept of a "DDI alert" is presented in a binary nature, they are either present or they are not. There is no mention made of how they will be presented, in particular how strong the level of control will be regarding whether they will be modal vs non-modal, requiring override reasons, etc. While the study team may not be yet aware or be able to control these factors in the implementation, than can have a very significant impact in the effectiveness of the alert. Both of the papers cited by Strom, et al. allude to this issue. In the case of unintended consequences, the primary issue was likely not that the DDI alert was not appropriate, but rather it was the manner in which it was implemented, with an extraordinarily high level of control that likely led to incorrect actions taken and the study being terminated. The authors plan to have human factors analysis performed and I would encourage this analysis to look at the qualitative nature of these alerts and how it may impact their performance. The authors could also consider a third arm of the study, in which some element of the intrusiveness or interruptive nature of the alert are varied and compared. 2) While the idea of algorithms to improve alert specificity and PPV is very appealing, the suggestion that data from this study will be able to develop complete algorithms to improve the relevance of DDI alerts strikes me as a bit overly optimistic. As I'm sure the authors know, the context of a "clinically relevant" alert can be quite complex with temporal, patient and prescriber factors

	involved. Decision Tree and Bayesian modeling would no doubt be important tools in this development, but much more detail needs to be shared regarding this approach to this modeling if the authors are going to so clearly assert that the algorithms will truly improve alert performance.

REVIEWER	Gordon Schiff Brigham and Womens Hospital Havard Medical School
REVIEW RETURNED	19-Nov-2018

GENERAL COMMENTS	This manuscript is a protocol for a study that appears to be funded and beginning in Australia related to decision support and drug-drug interactions. This is the first time I have reviewed a "protocol paper" and it is not clear what sort of revisions would be appropriate to suggest, since I assume the funding and protocol are already in place. Thus I will just raise a few more general questions that perhaps the manuscript discussion could dilate upon (or perhaps be better addressed as the study methods are refined in practice). The protocol and study appear to be a good one. Certainly there is a need to advance the state-of-the-art related to this problem. We recently published a paper showing the marked deterioration in drug-drug interaction alerts with the implementation of a commercial EMR system. The continuing poor specificity of these alerts risks poisoning the effectiveness of the entire ability of clinical alerts to be noted and acted upon. Thus, to the extent that this study can help clarify and clean up and target key drug-drug interactions and better understand alerting and how it impacts the work flow, it would represent a great advance and contribution. The devil of courses in the details. And with drug-drug interactions the details have to do with what particular drug-drug interactions are focused on. The authors cite as a reference Stockly's drug Interaction checker/list, which is apparently a known standard reference list of interactions, however one that I am not personally familiar with. This, I assume is a reasonable choice but would be interested in more details about the "levels" of interactions and how these are categorized in this interaction checker. Does it have a limited list of so-called “level 1” alerts that are particularly evidence-based and high risk? As authors doubtless are aware, there is highly variable and conflicting information from 1 drug interaction reference source to another. How will they be addressing this issue? And critically how are these DDI alerts operationalized in the decision support software that is to be implemented? In other words, is this the Stockley interaction checker to be used by these 6 hospitals, or will each be “doing their own thing” as a “usual care” type of implementation to be studied here? And how are the alerts presented – as interrupted, hard stops, requiring a reason for over-rides, what level of alerting is “turned on,” etc. What about "evolving" implementation - changes over time during the study period in a hospitals' use of alerts (e.g. do they change the level of alerting midway through the study period?) From the paper it was not clear whether the "response" of the clinician was to be inferred from the chart review or if this is

	electronically captured and available for the researchers to analyze. The paper mentions that the data would be deidentified. Does this include the prescriber or just the patients? No mention of how study will you deal with clustering by prescriber? And will you be able to link a given prescriber in the database with the interviews with that prescriber about their comments and use of decision support? Big worry in many of the studies is the inability to detect adverse outcomes from patients that received interacting drugs. I will you appropriately question that there is really too much of a beta error to confidently say a particular interaction is not harmful. Given the admonition to "do no harm" we would not want this study to give false reassurance about any particular drug-drug interaction.

REVIEWER	Luca Pasina Istituto di Ricerche Farmacologiche Mario Negri IRCCS
REVIEW RETURNED	05-Feb-2019

GENERAL COMMENTS	The aims of the study are very important for clinical practice, because the clinical outcome of a potential DDI is often not known and studies dealing with this problem are rare. Below my comments for the authors: - a flow chart could be useful to better understand the study design - information about the selection of wards and patients included in the study (geriatric, internal medicine, orthopedic wards) be usefull: how will be selected the 280 patients at each site? - a brief description or an example of Stockley's Interactions Checker should be reported: whch information are given in addition to severity and potential clinicl effects?

VERSION 1 – AUTHOR RESPONSE

Reviewer 1 comments	Responses
The concept of a "DDI alert" is presented in a binary nature, they are either present or they are not. There is no mention made of how they will be presented, in particular how strong the level of control will be regarding whether they will be modal vs non-modal, requiring override reasons, etc. While the study team may not be yet aware or be able to control these factors in the implementation, than can have a very significant impact in the effectiveness of the alert. Both of the papers cited by Strom, et al. allude to this issue. In the case of unintended consequences, the primary issue was likely not that the DDI alert was not appropriate, but rather it was the manner in which it was	We thank the reviewer for this comment and agree that design and implementation of alerts are critical. We have added a sentence under the section 'Research project setting' to clarify that all alerts implemented would be interruptive, all would require an override reason to be entered, but none would be hard-stop alerts preventing the prescriber from continuing. Implementation decisions are made by the hospitals without intervention from the research

implemented, with an extraordinarily high level of control that likely led to incorrect actions taken and the study being terminated. The authors plan to have human factors analysis performed and I would encourage this analysis to look at the qualitative nature of these alerts and how it may impact their performance. The authors could also consider a third arm of the study, in which some element of the intrusiveness or interruptive nature of the alert are varied and compared.	team, precluding a third arm from being included in the study.
While the idea of algorithms to improve alert specificity and PPV is very appealing, the suggestion that data from this study will be able to develop complete algorithms to improve the relevance of DDI alerts strikes me as a bit overly optimistic. As I'm sure the authors know, the context of a "clinically relevant" alert can be quite complex with temporal, patient and prescriber factors involved. Decision Tree and Bayesian modeling would no doubt be important tools in this development, but much more detail needs to be shared regarding this approach to this modeling if the authors are going to so clearly assert that the algorithms will truly improve alert performance.	Our goal is to develop algorithms that can begin to predict clinically relevant DDIs. We agree that accounting for the influence of clinical context and the patient and prescriber factors on the specificity of predictions is ambitious and acknowledge that further work to refine algorithms will likely be needed. We have amended the text on Page 14 to refer to their 'potential' to improve specificity.

Reviewer 2 comments	Responses
The authors cite as a reference Stockly's drug Interaction checker/list, which is apparently a known standard reference list of interactions, however one that I am not personally familiar with. This, I assume is a reasonable choice but would be interested in more details about the "levels" of interactions and how these are categorized in this interaction checker. Does it have a limited list of so-called "level 1" alerts that are particularly evidence-based and high risk? As authors doubtless are aware, there is highly variable and conflicting information from 1 drug interaction reference source to another. How will they be addressing this issue?	Stockley's Interactions Checker is a component of MedicinesComplete which is published by the Royal Pharmaceutical Society of Great Britain. Four levels of severity are used by the interactions checker: (1) Severe (2) Moderate (3) Mild (4) Nothing expected We have now included additional information about Stockley's in an Appendix.

	We selected Stockley's as this is recognised as the gold standard in Australia. This tool will be used by study pharmacists to identify potential DDIs in each patient's medication list. We recognise that drug interaction references are highly variable in the information presented, but elected to use one tool to ensure consistency in this initial classification. Note that all potential DDIs detected by Stockley's are then also reviewed by experienced pharmacists to determine clinical relevance. We have included some additional text on our reason for selecting Stockley's under Part 1.
And critically how are these DDI alerts operationalized in the decision support software that is to be implemented? In other words, is this the Stockley interaction checker to be used by these 6 hospitals, or will each be "doing their own thing" as a "usual care" type of implementation to be studied here? And how are the alerts presented – as interrupted, hard stops, requiring a reason for over-rides, what level of alerting is "turned on," etc. What about "evolving" implementation - changes over time during the study period in a hospitals' use of alerts (e.g. do they change the level of alerting midway through the study period?)	Stockley's Interaction Checker is not used by the hospital EMM systems. All intervention hospital EMM systems will utilise the Cerner Multum™ DDI knowledge-base for DDI detection, although some local customisation is expected. A list of all DDI alerts which have been incorporated into EMM systems in study hospitals will be provided to researchers following implementation. This is now stated under 'Research Project setting'. A list of any changes made to alerts during the data collection period will also be provided to researchers. These changes are likely to be minor, as sites have agreed to minimise any modifications during the trial period (6 months). We have added a sentence under the section 'Research project setting' to clarify that all alerts implemented will be interruptive, all will require an override reason to be entered, but none will be hard-stop alerts preventing the prescriber from continuing.
From the paper it was not clear whether the "response" of the clinician was to be inferred from the chart review or if this is electronically captured and available for the researchers to analyze.	Thank you for this comment. At the time of writing, we believed this would be possible, but on commencing data collection, we discovered that it would not be possible to infer prescriber responses from chart review. We have modified this statement to make this clear (see Page 11).

	As explained under section 3a, data will be extracted from systems to determine 'the number DDI alerts encountered and overridden'
The paper mentions that the data would be deidentified. Does this include the prescriber or just the patients? No mention of how study will you deal with clustering by prescriber? And will you be able to link a given prescriber in the database with the interviews with that prescriber about their comments and use of decision support?	Patient data will be de-identified. Prescriber data will be available. A mixed effect model will be applied to consider the correlation between medications ordered by the same prescribers. We have now included this in the analysis section (page 13-14). This study will not link prescriber interview data with database data, however interviews will be continued until data saturation is achieved, ensuring the 'majority' view is captured.
Big worry in many of the studies is the inability to detect adverse outcomes from patients that received interacting drugs. I will you appropriately question that there is really too much of a beta error to confidently say a particular interaction is not harmful. Given the admonition to "do no harm" we would not want this study to give false reassurance about any particular drug-drug interaction.	We acknowledge the challenges associated with identification of DDI related harm in section 'Part 1: Chart review', but as stated, we will be "using multiple sources of information from records, and both pharmacists and expert clinical pharmacologists to assess clinical outcomes". We have added 'and their link to DDIs' (Page 11) to this sentence to highlight that a causality assessment is being undertaken.

Reviewer 3 comments	Responses
A flow chart could be useful to better understand the study design	Thank you for this suggestion. We have now included a flow diagram as Figure 1.
Information about the selection of wards and patients included in the study (geriatric, internal medicine, orthopedic wards) be usefull: how will be selected the 280 patients at each site?	Patients will be randomly selected from all patients admitted to study hospitals during a 1-week period 6 months before and 6 months after EMM. All patients will be included except for those who visited the ED but were not admitted to wards, and those in wards where a different EMM system was in use (i.e. the ICU and oncology). This information has now been included under section 'Part 1: Retrospective chart review'
A brief description or an example of Stockley's Interactions Checker should be reported: whch information are given in addition to severity and potential clinicl effects?	Information on Stockley's now appears in Appendix 1.

VERSION 2 – REVIEW

REVIEWER	Gordon Schiff Brigham and Women's Hospital Harvard Medical School U.S.
REVIEW RETURNED	24-Mar-2019

GENERAL COMMENTS	This is an impressive and important study, and it appears to me that the researchers have thought through many key and critical aspects of trying to answer their research questions about DDI CDS as well as in conducting such a study. However, it appears that the researchers are trying to do two things at once here. I wonder if this method best serves the second goal ideally and most powerfully or not. While the first goal is to understand work flow and behaviors and outcomes related to implementation of a DDI decision support program, and is well suited to this task, the second seems to be to establish which drug drug interactions are clinically meaningful and impactful. For this I worry that the sample size is maybe too small to truly rule out a significant problem. Especially if we are thinking a % of a % (% of patients will have DDIs, and only a small % will have potential problems). Figuring out this latter issue has been a daunting one in this area and I fear my still haunt and vex this impressive effort. Why are the cases with a question of clinical important drug-drug interactions reviewed only by pharmacologist rather than physicians? I would think clinicians would be important to include here. Especially since drug drug interaction outcomes are often so ambiguous, and (as suggested above) problematic to determine cause and effect. Will the human factors researcher(s) who accompanies and observes the clinicians on their rounds be known to the clinicians in terms of their role, what aspect of their interactions with the computer that are in question, know (vs. be blinded) to the DDI question of the study? Does your journal typically publish research protocols prior to the study data being presented? While it would be useful for others in the field to know that this study was going on and the methodology is well described I question if this warrants publication at this stage. As a researcher in this area I would certainly welcome the opportunity to read it, but would have to defer to the editors on how they view its appropriateness for publication at this stage

REVIEWER	Luca Pasina Istituto di Ricerche Farmacologiche Mario Negri IRCCS, Italy
REVIEW RETURNED	20-Mar-2019

GENERAL COMMENTS	The manuscript is improved and is suitable for publication

VERSION 2 – AUTHOR RESPONSE

Optimising computerised decision support to transform medication safety and reduce prescriber burden: Study protocol for a mixed-methods evaluation of drug-drug interaction alerts

Please see below our response to reviewer comments.

Comment	**Response**
The researchers are trying to do two things at once here. I wonder if this method best serves the second goal ideally and most powerfully or not. While the first goal is to understand work flow and behaviors and outcomes related to implementation of a DDI decision support program, and is well suited to this task, the second seems to be to establish which drug drug interactions are clinically meaningful and impactful. For this I worry that the sample size is maybe too small to truly rule out a significant problem. Especially if we are thinking a % of a % (% of patients will have DDIs, and only a small % will have potential problems). Figuring out this latter issue has been a daunting one in this area and I fear my still haunt and vex this impressive effort.	Thank you for this comment. The study was powered to determine the impact of DDI alerts on DDI rates. We agree that we are likely to find a small proportion of DDIs that result in harm to patients and our study was not powered to establish the impact of DDI alerts on patient harm. However, we expect to find many potential DDIs and that a large proportion of these will be clinically relevant, in only *some situations*. We will use the data collected to develop algorithms to improve the positive predictive value of alerts.
Why are the cases with a question of clinical important drug-drug interactions reviewed only by pharmacologist rather than physicians? I would think clinicians would be important to include here. Especially since drug drug interaction outcomes are often so ambiguous, and (as suggested above) problematic to determine cause and effect.	We apologise for this confusion. In Australia, clinical pharmacologists are practicing physicians with specialist training in clinical pharmacology We have amended the text to read ‘Clinical pharmacologist physicians’.
Will the human factors researcher(s) who accompanies and observes the clinicians on their rounds be known to the clinicians in terms of their role, what aspect of their interactions with the computer that are in question, know (vs. be blinded) to the DDI question of the study?	Participants will be aware that the observer is noting user interactions with the system and alerts. We are highly experienced in this methodology and have found that knowledge of this does not appear to result in participants reading all alerts or adhering to alert recommendations. For example see: Baysari MT, Westbrook JI, Richardson KL, Day RO. The influence of computerized decision support on prescribing during ward-rounds: are the decision-makers targeted? JAMIA. 2011;18:754-9. Jaensch SL, Baysari MT, Day RO, Westbrook JI. Junior doctors' prescribing work after-hours

	and the impact of computerized decision support. Int J Med Inform. 2013 Jul 24;82:980-6. Santucci W, Day RO, Baysari MT. Evaluation of hospital-wide computerised decision support in an intensive care unit: an observational study. Anaesth Intensive Care. 2016 Jul;44(4):507-12.