

Using the Electronic Health Record User Context in Clinical Decision Support Criteria

Hojjat Salmasian^{1,2} David Rubins^{1,2} David W. Bates^{1,2}

¹Division of General Internal Medicine, Brigham and Women's Hospital, Boston, Massachusetts, United States

²Mass General Brigham, Somerville, Massachusetts, United States

Address for correspondence Hojjat Salmasian, MD, MPH, PhD, Division of General Internal Medicine, Brigham and Women's Hospital, 75 Francis St, Boston, MA 02115-6195, United States (e-mail: hsalmasian@bwh.harvard.edu).

Appl Clin Inform 2022;13:910–915.

Abstract

Background Computerized clinical decision support (CDS) used in electronic health record systems (EHRs) has led to positive outcomes as well as unintended consequences, such as alert fatigue. Characteristics of the EHR session can be used to restrict CDS tools and increase their relevance, but implications of this approach are not rigorously studied.

Objectives To assess the utility of using “login location” of EHR users—that is, the location they chose on the login screen—as a variable in the CDS logic.

Methods We measured concordance between user's login location and the location of the patients they placed orders for and conducted stratified analyses by user groups. We also estimated how often login location data may be stale or inaccurate.

Results One in five CDS alerts incorporated the EHR users' login location into their logic. Analysis of nearly 2 million orders placed by nearly 8,000 users showed that concordance between login location and patient location was high for nurses, nurse practitioners, and physician assistance (all >95%), but lower for fellows (77%) and residents (55%). When providers switched between patients in the EHR, they usually did not update their login location accordingly.

Conclusion CDS alerts commonly incorporate user's login location into their logic. User's login location is often the same as the location of the patient the user is providing care for, but substantial discordance can be observed for certain user groups. While this may provide additional information that could be useful to the CDS logic, a substantial amount of discordance happened in specific user groups or when users appeared not to change their login location across different sessions. Those who design CDS alerts should consider a data-driven approach to evaluate the appropriateness of login location for each use case.

Keywords

- ▶ clinical decision support
- ▶ electronic health record
- ▶ alert fatigue
- ▶ health care quality

Background and Significance

Computerized clinical decision support (CDS) has been widely used in electronic health record (EHR) systems and its use has been associated with significant improvements in health care quality and safety.^{1–3} However, CDS has also been criticized for its unintended consequences, including alert

fatigue.^{4–6} Despite their clinical impact, there has been little standardization in terms of how CDS should be designed, or how CDS rules should be optimized to reduce alert fatigue. One approach to increase the efficacy and accuracy of CDS tools and reduce unnecessary alerting and alert fatigue is to increase their “relevance” to the context in which the patient care and provider workflow take place.⁷ Modern EHR vendor

received
December 6, 2021
accepted after revision
July 28, 2022

© 2022. Thieme. All rights reserved.
Georg Thieme Verlag KG,
Rüdigerstraße 14,
70469 Stuttgart, Germany

DOI <https://doi.org/10.1055/s-0042-1756426>.
ISSN 1869-0327.

systems support CDS rules that could be used to determine the context in which the user is interacting with the system, and this information might be used to restrict CDS that is not relevant in that context. Different variables can be used to determine the relevance of an alert, including patient's location, provider's role, or provider's location. This obviates the need for studies that evaluate the usefulness of such context-based CDS rules.

There is no gold standard as to how the EHR user's actual clinical "context" can be defined. It may be possible to capture the user's interaction with the EHR through direct observation such as time–motion studies,^{8,9} but those studies are time-consuming and require lots of technological and financial resources. Another approach is to use geo-tagging, but even user physical location may not correspond to the care a user is delivering, for example, when a surgeon is consulting on a patient in the medical ward.

A common approach to narrow down the application of a CDS tool and increase its relevance is to restrict it based on characteristics of the user's EHR session. EpicCare (Epic Systems, Verona, Wisconsin, United States) particularly allows users to select a "login location" at the beginning of each EHR session as a surrogate for the context of their activity in the EHR. When users log into the EHR system, they choose the location of care where they are working, e.g., using a drop-down menu in the EHR login screen (→ Fig. 1). The recommendation is that outpatient providers and nurses log in to the appropriate location where they will be caring for the patient. This may be a clinic (such as "cardiology clinic"), an inpatient unit (such as "unit 12 north"), or any other physical location (such as "emergency room"), but it may also be a virtual location (such as "cardiology consult") that represents the type of service being provided. A CDS rule can check the login location and restrict the display of a specific CDS tool—such as an alert or a reminder—to users that have logged into specific locations, or conversely, to exclude users who have logged into specific locations. For

instance, an anesthesiologist may spend part of their time in the operating room and part of it in the pain management clinic, and in both these contexts, they may place orders for an opioid such as fentanyl. Although their role in the system would be the same, their login location may help control when they receive alerts that are only relevant to the use of fentanyl for pain management.

This approach assumes that the EHR user's login location is a reliable approximation of the actual context in which they are providing care to patients. While this makes sense intuitively—e.g., providers who work at the neonatal intensive care unit (NICU) would typically choose the NICU in the login screen—it is possible for users to log into one location but interact with charts of patients from a different location without resetting their login location. Examples include moonlighting physicians, pharmacists supporting multiple hospital units, and specialty providers called in to work in different specialties—as was the case during the COVID-19 pandemic. Failure to update the login location could result in CDS malfunctions for such users. In summary, incorporating user's login location into CDS rules may be associated with inaccurate CDS logic, but we were unable to find literature studying the usage and accuracy of login location in CDS rules.

Objectives

Our objective was to investigate the utility of an EHR user's login location as a variable used in CDS criteria. Specifically, we aimed to determine how often the user's login location is in concordance with the patient's location and whether this concordance varied among different user groups.

Methods

This study was conducted by descriptive analysis of data from a single installation of EpicCare EHR at a large health care system with more than 3,000 inpatient beds where EpicCare was implemented in 2015.

We used EHR user's login location as an indirect method to approximate the context of the user's clinical activity and compared this with a patient's physical location as denoted in the EHR. For the purposes of this study, we defined "concordance" as a provider logging into the same location that a patient was being seen at (e.g., Clinic A); if there was discrepancy (e.g., provider logged into Clinic B and the patient was being seen in Clinic A), we defined this as "discordance." For inpatients, defining the meaning of concordance or discordance is more difficult because while some providers (especially nurses and primary teams) may use a physical location at login, many other providers (such as consultants and pharmacists) may use virtual locations which are by definition not the same as the physical location of each patient they are consulting on. In fact, our preliminary analysis showed a much lower concordance between patient location and login location in the inpatient setting. Therefore, we excluded inpatients from this analysis.

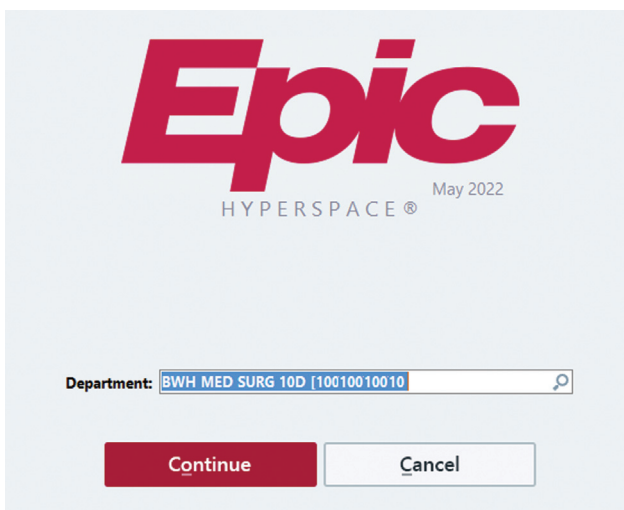


Fig. 1 Login screen in an installation of EpicCare EHR, where the user can select their login location after they have successfully entered their username and password. The screenshot is used as permitted by Epic Systems Corporation.

EpicCare records both the login location of the ordering provider and the patient location for all orders. Therefore, we analyzed orders placed over a 2-month period (January and February of 2020, the most recent months not affected by COVID-19-related changes to workflow). For each order, we compared the user's login location and the patient's location to determine concordance or discordance.

We hypothesized that one possible cause for discordance between patient location and login location could be that a user may not “reset” their login location throughout an EHR session or between consecutive EHR sessions. For instance, a provider who works in two hospitals that both use the same EHR system may always use the same login location from one of the hospitals, which makes the login location data inaccurate for some of the EHR sessions. This can cause problems not only for CDS rules that incorporate login location in their criteria, but also for other EHR features that are location-dependent (such as order lists for tests or medications). To study this, we conducted two analyses. First, we sorted the orders placed by each provider from oldest to newest and looked for situations in which a provider stopped ordering for one patient whose location matched their login location and started ordering for another patient in a different location. Of all these cases, we determined how often the provider “reset” their login location to match the location of the new patient. A lower percentage here would indicate that much of discordance between patient and login location is due to “inertia” in the login location information (i.e., it is not updated by providers through resetting their login context, and the login location information becomes stale). Second, we hypothesized that if a provider were to keep their login location accurate, then they would login to a larger number of distinct locations; therefore, we calculated the correlation between number of distinct login locations and the average concordance between login location and patient location, by provider. Here, a lower amount of correlation would indicate that changes in the login location are not following the changes in the patient location.

We limited our primary analysis to orders placed during in-person or virtual outpatient visits. Understanding the clinical situation for inpatient orders is more difficult as

multiple factors can affect patient location (e.g., patient preference, hospital capacity) and ordering provider login location (e.g., consultants, proceduralist) leading to high rates of discordance that may not be meaningful.

To estimate the pervasiveness of the use of login location in CDS tools, we identified all CDS alerts that factored in the user's login location as part of their logic. To contextualize the usage of login location in CDS alerts, we compared the number of CDS alerts using such criteria with the overall number of CDS alerts that were active in the EHR system.

This analysis was conducted primarily as a quality improvement initiative and the Mass General Brigham Institutional Review Board review considered it exempt from review. In the dataset used for the analysis, all unique identifiers for patients and users had been replaced with pseudo-identifiers. All analyses were done using R version 4.1.0 including the *tidyverse* family of software packages.¹⁰

Results

Analysis of Order-Level Data

A total of 143,981 orders were considered for this study, of which 644 (0.4%) were excluded because provider type or patient location was missing for them. The included 143,337 orders were placed by 1,257 distinct providers of 19 distinct provider types. Attending physicians were the largest contributors to orders and placed 99,853 (70%) of all orders; the next provider types, based on total number of orders, included physician assistants (9.1%), nurse practitioners (8.4%), clinical fellows (fellows, 4.5%), resident physicians (residents, 3.7%), and registered nurses (RNs; 3.6%). All other provider types contributed to less than 1% of all orders each and we combined them into a single group called “Other” (→ **Table 1**). The most common order types included laboratory orders (49%), medications (21%), imaging (7.1%), referrals (6.1%), and microbiology (5.8%); all other order types composed less than 5% of all orders each.

Overall, providers had selected their login location from one of 274 distinct options, and there were 106 distinct patient locations in our data. Expectedly, login locations were more diverse than patient locations, because patient

Table 1 Volume of orders placed by each provider type and the percentage “concordance rate” for outpatient orders

Provider type	Number of orders (percent total)	Direct concordance	Direct or indirect concordance
Attending physician	99,853 (70%)	79%	93%
Resident physician	5,261 (3.7%)	49%	55%
Physician assistant	13,033 (9.1%)	92%	97%
Registered nurse	5,185 (3.6%)	98%	99%
Fellow	6,483 (4.5%)	56%	77%
Nurse practitioner	12,080 (8.4%)	96%	100%
Other	1,442 (1.0%)	76%	93%

Note: Provider types whose order volume was less than 1% of all orders were grouped together into the *Other* group. A 100% direct concordance would mean that provider's login location and patient's location were identical for all orders. A 90% direct or indirect concordance would mean that for 90% of the orders, provider's login location and patient's location were either identical or the patient location would “roll up” to the provider's choice of login location. Please refer to the text for a detailed definition of “concordance.”

locations only include physical patient care locations while providers can log into virtual locations as well as physical patient care locations. Of these location identifiers, 97 were common, i.e., both used for patient location and used by providers as login location. Examples of patient locations that were not found in provider login location included interpreter services and executive locations. Examples of login locations that were not found in patient locations included virtual departments focused on specific inpatient service lines (e.g., medicine, medical oncology) or outpatient clinic types (e.g., breast oncology, leukemia).

Most providers (834; 66%) only used a single login location throughout the study period, irrespective of the patients' location. The average number of distinct login locations used by providers was 1.0 (median = 1, interquartile range [IQR] = 1–2, max = 5). The average number of distinct login locations per user was highest among fellows (1.7) and residents (1.6) and lowest among nurses (1.1). There was a weak correlation between the distinct number of patient locations a provider placed orders for and the distinct number of login locations that provider used (Pearson correlation coefficient = 0.23; ▶ Fig. 2).

In 115,389 (81%) of orders, the ordering provider's login location was concordant with the patient's location. There was a high level of variability in the percentage of orders placed by each provider type in which the provider's login location exactly matched the patient's location; we described this as "direct concordance." When also considering the specialty of the department the provider logged in to and the patient was seen in (which we call "indirect concordance"), the overall concordance increased to 132,292 (92%) of orders (▶ Table 1).

For 612 providers (49%), the login location always matched the patient location, and for 252 (20%) it never matched the patient location. For the remaining 393 providers, the percentage of orders for which the provider's login location was concordant with the patient's location was highly variable (mean = 65%, median = 71%, IQR = 46%). Most of the providers whose login location always matched the patient location were attending physicians (363; 59%).




Fig. 2 Scatter plot showing the relationship between the distinct number of patient locations a provider placed orders for and the distinct number of login locations that provider used. Each dot represents one provider.

Of 36,894 incidents in which a provider was placing orders for a patient whose location matched their login location and then they switched to a different patient at a different location, in 31,147 (84%), they did not reset their login location to match the location of the new patient. The correlation between distinct number of locations and the average concordance between login and patient locations was weak (Pearson $R = -0.03$, p -value = 0.006). These two findings, collectively, indicate that providers are not resetting their login location based on the transitions of their care to patients in different locations.

Analysis of CDS Alerts

A total of 277 CDS alerts used login location in their logic and were in active use at the time of our study. In comparison, there existed 652 active CDS alerts which did not use any rules that checked the user's login location. In summary, of a total of 929 active CDS alerts, 277 (29%) used at least one rule that checked the user's login location. Looking at a subset of these alerts, they were primarily clinical alerts (e.g., discouraging unnecessary *Clostridium difficile* testing, or adding plans of care) targeted at both nurse and providers. The login location was used in the alerts either: to exclude users (e.g., nursing alerts targeted at the primary nurse for a patient that were intended to exclude other nurses on the team who may have been treating wounds, placing IVs, or doing administrative/quality improvement work in the chart); to target a particular role type (e.g., restricting to physicians logged into an anesthesia location to show anesthesia-specific alerts); or to restrict the CDS to the "primary team" by comparing the login location of the user with the admitting service of the patient.

Discussion

Our findings suggest that in most cases, a provider's EHR login location was concordant with the patient's location, but there was variability among provider types. RNs and attending physicians have higher rates of concordance with patient location, while residents and fellows have lower rates. Without time-motion studies and/or interviews, the reasons for this cannot be known for certain, though we think this is likely explained by the fact that residents and fellows have a larger number of clinical roles (i.e., seeing patients inpatient, rotating through different consult services, covering for other residents) and may not update their login location correctly with each transition. As CDS alerts rely on this being updated, it could lead to CDS malfunctions, such as a CDS rule not being triggered when it should be.

Our analysis of CDS tools suggested that a user's login location is used in CDS alerts. In our institution, all CDS tools that used login location in their criteria also used patient location or provider role (or both) in their criteria. Given that hundreds of CDS tools are using login location in their criteria, this means CDS tool designers believe that patient location and provider role alone are not sufficient to restrict the CDS tool to the most relevant context.

Our study has several limitations. We used data from a single health care institution and our data are at best representative for one EHR system. Other EHR systems, or other configurations of the same EHR system, could yield different results. Our analyses are limited by using patient location as a surrogate for true clinical context; only prospective observational studies can capture the true clinical context of the EHR user accurately. To capture which CDS tools incorporate user context in their logic, we only looked at those elements of logic that use standard EHR capabilities; if CDS tools use backend code (e.g., “Extensions” in EpicCare) and access the user’s login location in nonstandard ways, they would not be captured in our analysis, and we would have underestimated the prevalence of incorporating user context in CDS logic.

Finally, while our data suggest that a large portion of CDS alerts use login location in their logic and this information may be inaccurate in many cases, it does not provide a point of reference, i.e., we cannot make claims as to whether this type of CDS rules is more or less accurate than other types, and we cannot offer specific alternatives to context-based rules either. It might be that for some types of user concordance is reliably high, so context could be useful in these instances. Future research can focus on comparing the CDS tools that use login location with those that do not use it in their criteria, in terms of what types of alerts they are or whether they have higher relevance or higher acceptance rates by the user.

Overall, our findings provide an initial insight into the accuracy of this method of determining user context and a reference for future studies on other types of CDS rules. Specifically, our findings suggest that CDS tools that use user context should account for both direct and indirect concordance of context with clinical workflow.

Conclusion

We found that a fifth of active CDS alerts incorporated the EHR user’s login location into their logic and, using an analysis of orders, we found that the EHR user’s login location was often compatible with the actual location of the patient they were providing care for, but this was not always the case and the discordance between EHR user’s context and the clinical context of their work varied by user role. This calls for a rigorous consideration and data-driven analyses before using EHR user context as a criterion in CDS tools.

Clinical Relevance Statement

Although it is reasonable to try to increase the relevance of CDS tools by focusing their activation for users only in specific clinical contexts, data from this study suggest that the clinical context of a user may not be properly captured by CDS tools using the conventional way, i.e., by looking at the user’s login location. To increase the usability of CDS tools and reduce alert fatigue, CDS tool designers can thoughtfully incorporate user’s login location in CDS logic after performing a data-driven assessment for each use case.

Multiple Choice Questions

- How frequently is user context used in computerized clinical decision support (CDS) criteria?
 - Never
 - Sometimes
 - Usually
 - Always

Correct Answer: The correct answer is option b. The study estimates that one in five CDS tools used user context in their logic. Practically, user context is only needed when the CDS logic is too broad and must be narrowed to a specific setting to increase its relevance. For these reasons, choices a, c, and d are not correct.

- What is one of the key limitations of incorporating the user’s EHR context into the logic of computerized clinical decision support (CDS) tools?
 - EHR context does not always correctly represent clinical context
 - EHR context data are not available in modern EHR systems
 - EHR systems do not have a way to capture the user’s context
 - Capturing EHR context takes a long time and slows down the EHR system

Correct Answer: The correct answer is option a. The study shows that while there is a high concordance between EHR context and patient location, this concordance is far from ideal, and is particularly lower for certain user groups. Modern EHR systems provide easy ways to capture user context based on the user’s last login location, so choices b and c are incorrect. Once EHR context is calculated based on the user’s login location, this information is readily available for each user session, so choice d is incorrect.

Protection of Human and Animal Subjects

The study protocol was evaluated by the Mass General Brigham Institutional Review Board (IRB) and was deemed exempt from IRB review.

Funding

None.

Conflict of Interest

D.W.B. reports grants and personal fees from EarlySense, personal fees from DCI Negev, equity from ValeraHealth, equity from Clew, equity from MDCLone, personal fees and equity from AESOP, personal fees and equity from FeelBetter, and grants from IBM Watson Health, outside the submitted work. The other authors report no conflict of interest.

References

- Hunt DL, Haynes RB, Hanna SE, Smith K. Effects of computer-based clinical decision support systems on physician performance and

- patient outcomes: a systematic review. *JAMA* 1998;280(15):1339–1346
- 2 Bright TJ, Wong A, Dhurjati R, et al. Effect of clinical decision-support systems: a systematic review. *Ann Intern Med* 2012;157(01):29–43
 - 3 Kwan JL, Lo L, Ferguson J, et al. Computerised clinical decision support systems and absolute improvements in care: meta-analysis of controlled clinical trials. *BMJ* 2020;370:m3216
 - 4 Campbell EM, Sittig DF, Ash JS, Guappone KP, Dykstra RH. Types of unintended consequences related to computerized provider order entry. *J Am Med Inform Assoc* 2006;13(05):547–556
 - 5 Embi PJ, Leonard AC. Evaluating alert fatigue over time to EHR-based clinical trial alerts: findings from a randomized controlled study. *J Am Med Inform Assoc* 2012;19(e1):e145–e148
 - 6 Backman R, Bayliss S, Moore D, Litchfield I. Clinical reminder alert fatigue in healthcare: a systematic literature review protocol using qualitative evidence. *Syst Rev* 2017;6(01):255
 - 7 Shortliffe EH, Sepúlveda MJ. Clinical decision support in the era of artificial intelligence. *JAMA* 2018;320(21):2199–2200
 - 8 Mamykina L, Vawdrey DK, Hripcsak G. How do residents spend their shift time? A time and motion study with a particular focus on the use of computers. *Acad Med* 2016;91(06):827–832
 - 9 Young RA, Burge SK, Kumar KA, Wilson JM, Ortiz DF. A time-motion study of primary care physicians' work in the electronic health record era. *Fam Med* 2018;50(02):91–99
 - 10 R Core Team. R: A language and environment for statistical computing. 2017. Accessed August 6, 2022 at: <https://www.r-project.org>