Intelligent Health Systems – From Technology to Data and Knowledge
E. Andrikopoulou et al. (Eds.)

doi:10.3233/SHTI250283

Subgroup Discovery to Identify Determinants of Influence on CDSS Medication Alert Handling: A Feasibility Study

Kimmy RAVEN[a,1], Stephanie MEDLOCK[a], Linda W. PEUTE[a], and Iacopo VAGLIANO[a]
[a] *Department of Medical Informatics Amsterdam UMC location University of Amsterdam, Meibergdreef 9, Amsterdam, The Netherlands*
ORCiD ID: Kimmy Raven https://orcid.org/0000-0002-5504-636X,
Stephanie Medlock https://orcid.org/0000-0002-2679-8095,
Linda W Peute https://orcid.org/0000-0002-9859-0912,
Iacopo Vagliano https://orcid.org/0000-0002-3066-9464

Abstract. Clinical decision support systems (CDSSs) are designed to enhance patient safety by providing alerts to prescribers about potential medication issues. However, a significant proportion of these alerts are ignored, which can compromise patient safety. This study explores the feasibility of using subgroup discovery, a machine learning method, to identify determinants influencing physicians' medication-related CDSS alert handling. By analyzing CDSS log data from the electronic health record, this research shows the feasibility of the use of subgroup discovery on this data, and its potential to uncover behavioral patterns and factors that affect how alerts are managed. This can ultimately contribute to the design of more effective CDSS alerts and improving patient safety.

Keywords. Subgroup discovery, machine learning, clinical decision support system, medication alerts, human factors

1. Introduction

Medication-related alerts in Clinical Decision Support Systems (CDSSs) are crucial for preventing adverse drug events (ADEs) [1]. CDSSs increase the quality and safety of provided care and decrease the number of ADEs with the presentation of alerts during the prescribing process [1,2]. Despite their importance, studies have shown that 46.2% to 96.2% of these alerts are overridden by prescribers [1]. These high override rates may be attributed to alert fatigue, and may lead to inadequate alert handling [3]. This raises concerns about alert fatigue and the overall effectiveness of CDSSs. While studies have extensively examined the clinical relevance and the content of alerts [1,2], we believe that understanding the behavioral determinants that influence alert handling is essential for improving and understanding effective CDSS design, and ensuring patient safety.

[1] Corresponding Author: Kimmy Raven; E-mail: k.e.raven@amsterdamumc.nl.

Research suggests that machine learning (ML) can predict user interaction with alerts, and reduce alert fatigue [3]. To date, research on ML in medication alerts has mainly focused on supervised ML to predict the acceptance, and/or filter alerts to increase the signal-to-noise ratio [4-7]. Many factors have been shown to influence alert handling, but we do not have a complete list of factors, nor do we understand how they interact with each other [1-3,8].

Subgroup discovery (SD) is an ML technique which has not yet been used for the identification of determinants of influence of medication-related alert handling. SD aims to identify interesting relationships between different objects in a dataset, while taking into account a specific outcome [9]. The relations and patterns that emerge from this analysis are called subgroups and presented as “if-then” rules. An example of an “if-then” rule could be: if provider experience is low and provider gender is male, then the alert will be accepted. SD contains elements of both supervised and unsupervised ML and can be seen as a method that lies between the extraction of association rules and obtaining classification rules [9]. Classification rules classify new objects or predict outcomes, while association rules identify relationships between unlabeled objects. SD, however, aims to find these relationships while taking into account a specific outcome variable.

The objective of this study is to explore the feasibility of applying SD on Electronic Health Record (EHR), CDSS, and CDSS log data. CDSS log data includes the alerts that were triggered and user responses to alerts, including clinical actions such as ordering medications. Using SD, we aim to identify behavioral patterns and determinants that influence physicians’ medication-related alert handling. This research can provide insights into user behavior and inform the development of more effective CDSS alerts.

2. Methods

2.1. Data and Setting

This explorative research was conducted within the Amsterdam University Medical Centre, location University of Amsterdam. The data used was extracted from the testing environment of the Epic EHR and contains EHR, CDSS and CDSS log data from both centers of the Amsterdam UMC (University of Amsterdam and Free University), from the period January 1, 2017 until May 17, 2023. The data set contains user log data from fictional providers created for testing purposes, including interactions with alerts, and reflects the activity of developers and testers.

2.2. Definitions and Data Preprocessing

The main outcome studied is the response of the prescriber to an alert, represented as a binary outcome with the values accepted or declined. The data was cleaned and preprocessed. Within cleaning, we inferred the meaning of unknown values by identifying and interpreting these values with expert input (SM). E.g., missing pregnancy data for males was marked not applicable, and missing data links were labeled “not available”. Data preprocessing included the removal of duplicate warnings, the coupling of different tables, and aggregating data into new variables. Afterward, variables were excluded that still had 40% or more missing values, or had all identical values.

2.3. Subgroup Discovery

SD algorithms can be divided into two groups: exhaustive search and heuristic search. Exhaustive search generates all possible rules and then checks if these rules meet certain conditions [10].

Three different SD algorithms were compared: Apriori-SD [11], Patient Rule Induction Method (PRIM) [12], and Subgroup Set Discovery (SSD++) [13]. The algorithms were selected to form a diverse mix of algorithms based on association rules (APRIORI-SD), decision trees (PRIM) and inductive inference (SSD++) as well as heuristic (SSD++, PRIM) and exhaustive (APRIORI-SD) algorithms. APRIORI-SD can also be considered heuristic depending on whether the characterization of its search is based on the rule generation or on the rule post-processing. The performance of the three algorithms was assessed using coverage, support, confidence, Weighted Relative Accuracy (WRAcc), significance and redundancy, which are common measures to assess quality of SD algorithms and identified subgroups [9,10].

2.4. Analysis

We explored the feasibility of the application of SD on CDSS log data from the EHR testing environment, to see if different subgroup discovery algorithms were able to identify subgroups for behavioral patterns. Furthermore, we evaluated the performance of these algorithms and the relevance of discovered patterns. Data preprocessing was performed using R version 4.2.1. and the following packages: dplyr, stringr, tidyr, lubridate, feather. Python was used for implementing the SD algorithms. The algorithms were derived from a framework used for the identification of subgroups based on Covid-19 patient mortality available at https://bitbucket.org/aumc-kik/subgroup-discovery/ [14].

3. Results

The final dataset contained nineteen variables for 7,216 unique alerts, of which 469 (6.50%) were accepted by the user. Table 1 provides an overview of the variables available in the dataset. Table 2 provides the outcomes of the performance-indicators of the SD-algorithms. Table 3 provides the number of identified subgroups and the mean number of determinants per subgroup.

Table 1. Variables used for subgroup discovery, sorted by user-, patient-, alert-, and setting-related.

User	Patient	Alert	Setting
Gender	Gender	Type of alert	Month alert was shown
Age category	Age category	Alert priority level	Part of day alert was shown
Year of first login	Pregnant	Place alert was shown	Alert shown during weekend
Specialty	Department	Context of prescription	
Type of caregiver		Multiple alerts shown	
Total number of prescriptions			
Total number of patients			

The number of identified subgroups varied between the three algorithms (PRIM: 3; Apriori-SD: 5; SSD++: 30), and the average number of determinants per subgroup ranged from 1.67 to 2.17. Redundancy of subgroups varied between 0.1 to 0.6, where SSD++ produced the lowest number of redundant subgroups. Apriori-SD performed best regarding highest mean coverage (0.2731) and highest mean support (0.0372). SSD++

scored the highest total coverage (0.9997) and support (0.0647). PRIM performed best in terms of confidence (0.7575), average and maximum WRAcc (resp. 0.01962 and 0.057447), and significance (556.53).

Table 2. Evaluation of performance indicators for SD-algorithms. Coverage and support are average, others are maximum or total values. Best outcomes are marked bold.

Algorithm	Coverage	Support	Confidence	WRAcc	Significance	Redundancy
Apriori-SD	**.2731**	**.0372**	.2958	.0196	245.41	0.60
PRIM	.0337	.0218	**.7575**	**.0575**	**556.53**	0.33
SSD++	.0770	.0048	.4251	.0193	106.12	**0.10**

Table 3. The number or subgroups and average number of items per subgroup for the studied SD-algorithms.

Algorithm	Number of subgroups	Average number of items per subgroup
Apriori-SD	5	1.80
PRIM	3	1.67
SSD++	30	2.17

4. Discussion

The findings from this feasibility study highlight the potential of SD to uncover patterns in EHR, CDSS and CDSS log data. It suggests that the results of SD can be used to derive valuable insights into the determinants of CDSS alert handling. By identifying specific subgroups with distinct behaviors, healthcare organizations can tailor CDSS alerts to better meet the needs of different prescriber groups, thereby reducing alert fatigue and improving patient safety.

Current research identifies the importance of human factors and user characteristics on physicians' alert handling [2,8]. However, the relationships between factors are not well studied or are studied with traditional statistical methods [8]. SD can further extend the knowledge on factors influencing alert handling by identifying specific behavioral handling patterns and specifying user-interaction groups. To the best of our knowledge, this is the first study to use SD with CDSS data. Other SD studies used EHR data [14,15], supporting our feasibility statement for EHR data, but lacked CDSS data.

Further understanding of the determinants influencing alert handling can inform the design and presentation of more effective CDSS alerts. For example, alerts can be customized based on a combination of the prescriber's age, experience, and/or the time of prescription. Additionally, reducing the number of non-critical alerts can help mitigate alert fatigue and ensure that important alerts receive the necessary attention.

4.1. Limitations and Future Works

This study has two main limitations. The analysis was conducted on test data from the Epic EHR system; the method needs to be used on clinical data to determine if the findings indeed provide valuable insights. Additionally, the study focused on a limited set of factors, and further research is needed to explore other potential determinants of alert handling.

Future research will focus on applying SD to clinical EHR, CDSS and CDSS log data to identify determinants of influence on prescribers' medication-related alert handling. Additionally, we will conduct further exploration of known potential determinants of alert handling, such as prescriber workload and patient characteristics.

5. Conclusions

The study concludes that subgroup discovery is a feasible method for analyzing EHR, CDSS and CDSS log data to identify determinants of alert handling. Future research on clinical data should provide further insights into the relationship between different determinants on medication-related alert handling. By gaining a deeper understanding of user behavior through log file data, healthcare organizations can design and present more effective medication-related alerts, ultimately enhancing patient safety.

References

[1] Poly TN, Islam MM, Yang HC, Li YJ. Appropriateness of Overridden Alerts in Computerized Physician Order Entry: Systematic Review. JMIR Med Inform. 2020;8(7):e15653. doi: 10.2196/15653.
[2] Olakotan OO, Mohd Yusof M. The appropriateness of clinical decision support systems alerts in supporting clinical workflows: A systematic review. Health Informatics J. 2021;27(2): 14604582211007536; doi: 10.1177/14604582211007536..
[3] Kane-Gill S, O'Connor M, Rothschild J, Selby N, McLean B, Bonafide C, et al. Technologic Distractions (Part 1): Summary of Approaches to Manage Alert Quantity With Intent to Reduce Alert Fatigue and Suggestions for Alert Fatigue Metrics. Critical Care Medicine. 2017;45 (9), 1481-88; doi: 10.1097/CCM.0000000000002580
[4] Poly TN, Islam MM, Muhtar MS, Yang HC, Nguyen PAA, Li YJ. Machine Learning Approach to Reduce Alert Fatigue Using a Disease Medication-Related Clinical Decision Support System: Model Development and Validation. JMIR Med Inform. 2020;8(11):e19489; doi: 10.2196/19489
[5] Liu S, Kawamoto K, Del Fiol G, Weir C, Malone DC, Reese TJ, et al. The potential for leveraging machine learning to filter medication alerts. J Am Med Inform Assoc. 2022;29(5):891-9; doi: 10.1093/jamia/ocab292
[6] Chen J, Chokshi S, Hegde R, Gonzalez J, Iturrate E, Aphinyanaphongs Y, et al. Development, Implementation, and Evaluation of a Personalized Machine Learning Algorithm for Clinical Decision Support: Case Study With Shingles Vaccination. J Med Internet Res. 2020;22(4):e16848; doi: 10.2196/16848.
[7] Graafsma J, Murphy RM, van de Garde EM, Karapinar-Çarkit F, Derijks HJ, Hoge RH, et al. The use of artificial intelligence to optimize medication alerts generated by clinical decision support systems: A scoping review. J Am Med Inform Assoc. 2024;31(6):1411–22; doi: 10.1093/jamia/ocae076
[8] Bittmann JA, Haefeli WE, Seidling HM. Modulators Influencing Medication Alert Acceptance: An Explorative Review. Appl Clin Inform. 2022;13(2):468-85; doi: 10.1055/s-0042-1748146
[9] Herrera F, Carmona CJ, González P, del Jesus MJ. An overview on subgroup discovery: foundations and applications. Knowl Inf Syst. 2011; 29:495–525; doi: 10.1007/s10115-010-0356-2
[10] Helal S. Subgroup Discovery Algorithms: A Survey and Empirical Evaluation. J. Comput. Sci. Technol. 2016; 31:561–76; doi: 10.1007/s11390-016-1647-1
[11] Kavšek B, Lavrač N, Jovanoski V. APRIORI-SD: Adapting Association Rule Learning to Subgroup Discovery. In: Berthold MR, Lenz HJ, Bradley E, Kruse R, Borgelt C. (eds) IDA 2003: Advances in Intelligent Data Analysis V. Berlin, Heidelberg: Springer-Verlag; 2003. p.230-41. doi: 10.1007/978-3-540-45231-7_22
[12] Friedman JH, Fisher NI. Bump hunting in high-dimensional data. Statistics and Computing. 1999; 9:123–43. doi: 10.1023/A:1008894516817
[13] Proença HM., Grünwald P, Bäck T, Leeuwen Mv. Discovering Outstanding Subgroup Lists for Numeric Targets Using MDL. In: Hutter F, Kersting K, Lijffijt J, Valera I. (eds) Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2020. Lecture Notes in Computer Science(), vol 12457. Springer, Cham. doi: 10.1007/978-3-030-67658-2_2
[14] Vagliano I, Kingma MY, Dongelmans DA, de Lange DW, de Keizer NF, Schut MC; the Dutch COVID-19 ICU Research Consortium. Automated identification of patient subgroups: A case-study on mortality of COVID-19 patients admitted to the ICU. Comput Biol Med. 2023:107146. doi: 10.1016/j.compbiomed.2023.107146
[15] Yang J, Mwangi AW, Kantor R, Dahabreh IJ, Nyambura M, Delong A, et al. Tree-based subgroup discovery using electronic health record data: heterogeneity of treatment effects for DTG-containing therapies. Biostatistics. 2024;25(2):323-35; doi: 10.1093/biostatistics/kxad014.