Interpretable Identification of Comorbidities Associated with Recurrent ED and Inpatient Visits

Luoluo Liu[1], Dennis Swearingen[2], Eran Simhon[1], Chaitanya Kulkarni[3], David Noren[1], Ronny Mans[4]

Abstract— In the hospital setting, a small percentage of recurrent frequent patients contribute to a disproportional amount of healthcare resource utilization. Moreover, in many of these cases, patient outcomes can be greatly improved by reducing re-occurring visits, especially when they are associated with substance abuse, mental health, and medical factors that could be improved by social-behavioral interventions, outpatient or preventative care. Additionally, health care costs can be reduced significantly with fewer preventable recurrent visits.

To address this, we developed a novel, interpretable framework that both identifies recurrent patients with high utilization and determines which comorbidities contribute most to their recurrent visits. Specifically, we present a novel algorithm, called the minimum similarity association rules (MSAR), which balances the confidence-support trade-off, to determine the conditions most associated with re-occurring Emergency department and inpatient visits. We validate MSAR on a large Electronic Health Record dataset, demonstrating the effectiveness and consistency in ability to find low-support comorbidities with high likelihood of being associated with recurrent visits, which is challenging for other algorithms such as XGBoost.

Clinical relevance— In the era of value-based care and population health management, the proposal could be used for decision making to help reduce future recurrent admissions, improve patient outcomes and reduce the cost of healthcare.

Index Terms— Recurrent patients, Emergency Department, inpatient re-admissions, confidence-support trade-off, Association Rules

I. INTRODUCTION

Recurrent patients, also known as "frequent flyers", "high utilizers", or "super users" in hospitals, are a small group; however, they impose a disproportionately high utilization of resources [1] [2]. In fact, the top one percent contribute to 22% of health care spending, and the top 5% account for around 50% of overall costs [1]. Further, the top 15% of utilizers contribute to around 85% of total health care costs [2]. Accordingly, identifying potentially preventable visits among this group of patients could reduce hospital and health care costs significantly.

There are two major types of recurrent patients. One type includes those with mental health or substance (drug and/or alcohol) abuse conditions. It has been shown that these patients contribute to high utilization of Emergency Departments (EDs) [3]–[8] and inpatient visits in community hospitals [3], [9]. In 2008, mental health and substance abuse (MHSA) were the principal reasons behind 1.8 million patient hospitalizations, accounting for 4.5% of all hospitaliza-

tions in the U.S and costing $9.7 billion [9]. Additionally, a large percentage of MHSA patients lack insurance coverage, are sometimes unemployed, and have low average education levels [10] [11]. Fortunately, it has been reported that rehabilitation programs [12], accommodations [10], social support in the community [13] and other outpatient programs [10] improve patient outcomes and reduce re-admissions.

A second type of potentially preventable recurrent admission relates to chronic diseases, such as Acquired Immunodeficiency Syndrome (AIDS) and diabetes mellitus. This type of recurrent visits could be reduced through proactive planning and care. About half of AIDS re-admissions could be prevented by providing better access to medical supplies, assuring adherence to follow-up tests and visits, improving treatment compliance, and satisfying patients' psychological and social needs [14]. Similarly, robust diabetes management, discharge planning, and post-hospital instructions could lower ED visits for diabetes patients [15].

One of the biggest challenges in addressing recurrent visits is the lack of a standard criteria for identifying these patients. In fact, over 180 criteria from about 100 sources are surveyed [16]. An existing work combines 180 previously proposed rules using a clustering algorithm [16]. As a result, the final set of rules takes many factors, across clinical and operational elements, into account. This procedure makes the algorithm very complex and lacking in interpretability, creating a barrier for practical usage.

Due to the lack of standard criteria for identifying recurrent patients, our first aim is to identify this group of patients. Our second aim is to develop an interpretable method to provide insight into the pertinent factors contributing frequent visits. In the Electronic Health Record, the underlying reason for a particular patient visit can be inferred from the chief complaint or directly obtained if documented. However, for recurrent patients, the reason for any given visit might not be representative enough to explain the recurrent visits.

In identifying MHSA and chronic conditions associated with high utilization, the third key challenge is that some conditions (for example drug abuse, AIDS) have **low prevalence/support** but potentially **high likelihood/confidence** compared to others (such as hypertension with low confidence but high support). It is difficult for conventional machine learning methods such to detect those conditions, which is validated in Section III-B.5. Aiming at developing a explainable model, we investigated on a interpretable statistical method called Association Rules [17] (AR). The conventional AR method takes the maximum confidence/likelihood from all candidate rules exceeding a

[1]Luoluo Liu is with Philips Research North America (PRNA), Cambridge, USA. Email: luoluo.liu@philips.com

[2]D. Swearingen, MD is with Banner Health, Phoenix, USA. [1] E. Simhon, D. Noren are with PRNA; [3] C. Kulkarni is with Philips Innovation India; [4] R. Mans is with Philips Research Europe, Netherlands.

minimum support/prevalence [17]. Even though AR is advantageous due to its high interpretability, the main drawback is: **a conventional AR method with a fixed minimum support threshold could potentially eliminate those high confidence but low support conditions**.

Variants of AR methods have been developed for balancing confidence and support to help resolve this issue. Authors in [18] propose a modification on confidence calculation by adding an extra parameter. However, requiring the additional term makes it hard to interpret in deployment. Authors in [19] use prediction of future data to determine optimal weights. The underlining assumption that the distribution of comorbidities associated with frequent visits are consistent over time, is not likely true especially if some preventative efforts has been taken over the years.

Instead, we propose an algorithm named minimum similarity association rules (MSAR) that improves on the conventional AR method by **balancing the support and confidence trade-off** with optimal weights learned from data which are easily adjustable to varying deployment sites. In addition, the MSAR algorithm for identifying top comorbidities associated with recurrent visits is **highly interpretable, efficient, consistent, customizable, and deployment-friendly**. The proposed algorithm is particularly well suited to handle cases with a large variance in support distribution, which is crucial given that comorbidities vary widely in prevalence. Additionally, MSAR successfully identifies challenging case of **low-support comorbidities with high likelihood of recurrent visits**. The comparison between different algorithms is listed in Table I.

Algorithms	Association Rules (AR)	XGBoost [20] +Shapley [21]	MSAR (Ours)
interpretability	High	Medium	**High**
Consistency	High	Medium	**High**
balances conf-supp trade-offs	No	No	**Yes**
ability to select high-conf, low-supp factors	Limited	Limited	**Yes**
ability to distinguish across factors	Limited	Limited	**Yes**

TABLE I: Algorithm comparisons for selecting top factors.

II. METHOD

The flowchart of the proposal is illustrated in Figure 1 with two corresponding modules. The recurrent patient identifier module uses admission timestamps within 1 year to determine the frequency of admission, and then determines whether the visitor is a recurrent patient. The comorbidities explainer module, gives the top three comorbidities associated with their frequent visits.

A. Recurrent Patients Identifier

One of the challenges in reducing preventable recurrent visits is the lack of a standard definition for recurrent patients. Previous works have defined recurrence based on utilization, physiological states, and cost factors [16]. Among these choices, utilization elements are the most easily accessible from the EHR, and the algorithm can be executed at the time of admission. We therefore decided to use top utilization data elements from the literature and then combine them.

Our proposed framework identifies recurrent patients by satisfying at least one of the following criteria: i) readmission within 30-days, which is accepted by the Centers for Medicare and Medicaid Services [22] in measuring operational performance; ii) more than 4 non-elective inpatient visits within a year (44% out of 180 criteria reviewed in [16]); iii) more than 4 emergency department visits within 1 year (41% out of 180 criteria reviewed in [16]).

B. Comorbidity Explainer via Min-similarity Association Rules (MSAR)

We introduce the proposed algorithm: Min-similarity Association Rules (MSAR), for selecting the top comorbidities associated with recurrent visits. This algorithm seeks optimal solutions that balance confidence and support, learned from retrospective data. It also addresses the difficulty in identifying high-confidence but low-support comorbidity combinations. Below we describe key parts related to MSAR.

1) Derivation of Comorbidities: From the Electronic Health Records (EHR) database, we first take International Classification of Diseases (ICD)-9 and ICD-10 diagnosis codes, which, in total, contain approximately 83000 codes [23]. To reduce the dimensionality when encoding such a large number of diagnosis codes, we use the Elixhauser comorbidity index [24] to map ICD codes into comorbidity categories. Elixhauser comorbidity is a good choice because it is an updated version of the widely adopted Charlson comorbidity index [25] [26] with newly added categories for mental disorders, drug and alcohol abuse, obesity, and weight loss, which are the factors that could potentially be improved via social-behavioral interventions and outpatient care. Since most recurrent patients has more than 4 visits within a year, we take ICD codes from previous 3 visits within 1 year and map them to Elixhauser comorbidities using the Hcuppy 0.0.7 [27] Python package.

2) Number of comorbidities for recurrent patients: On average, recurrent patients have $7.5(\pm3.7)$ comorbidities from their past three visits within a year. Displaying all comorbidities associated with recurrent visits is not efficient for visualization nor for identifying the root cause for recurrent admissions. Therefore, this motivates us to develop a comorbidity explainer to pick the top comorbidities associated with frequent visits.

3) Confidence and support: In the context of our problem, **confidence** is the **likelihood of being a recurrent patient given a specific combination of comorbidities**; and **support** is the **prevalence of a given comorbidity combination**. We first visualize the confidence and support of a single comorbidity for all types of comorbidities in Figure 2. Drug abuse has the highest confidence overall. Additionally, the variance of support across different comorbidities is large, with values ranging from 0.0045 to 0.72. There are comorbidities with high confidence but low support, such
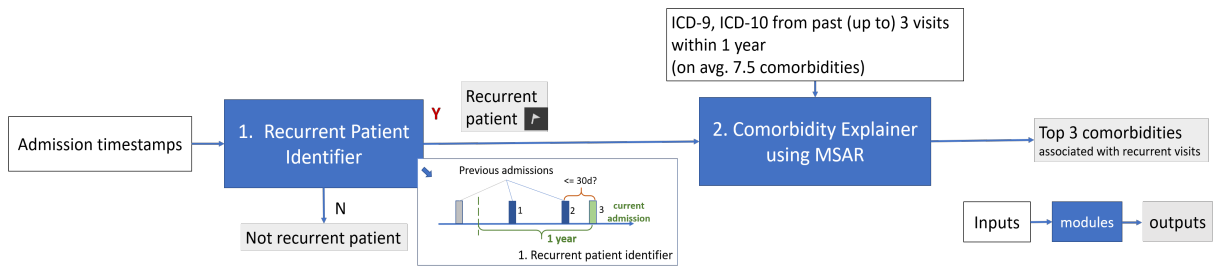
Fig. 1: The first module identifies a recurrent patient and the Comorbidities Explainer outputs the top 3 comorbidities associated with frequent visits for that patient. Rules are trained on retrospective data with both recurrent and non-recurrent patients. The algorithm can be executed for recurrent patients on an individual patient level.
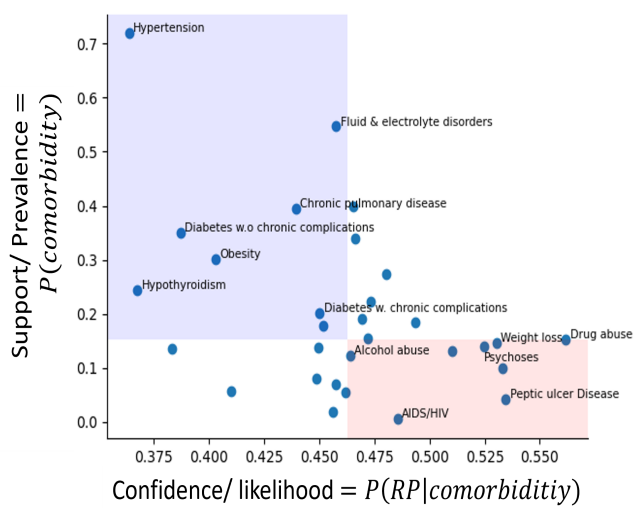

Fig. 2: Confidence and support of Elixhauser comorbidities. The blue and red boundaries are medians of confidence and support, respectively. Comorbidities in the blue shaded area have low-confidence and high-support. Comorbidities in the red shaded area have high-confidence and low-support.
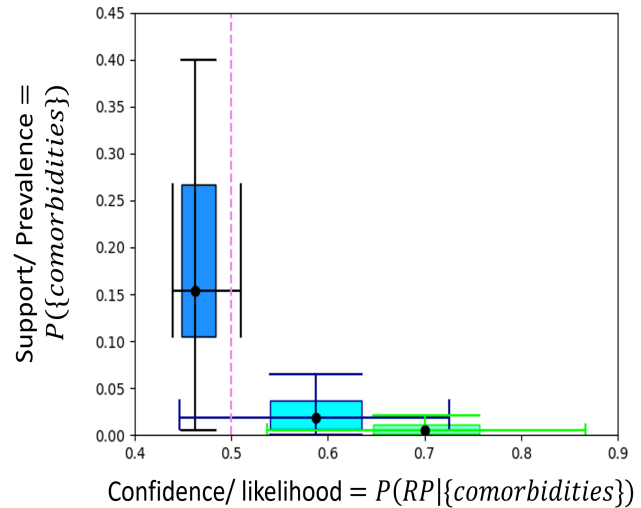

Fig. 3: Increasing the number of comorbidities n will increase confidence while decreasing support. Box plots show differences in ranges of confidence and support with varying numbers of comorbidities (Left $n = 1$, middle $n = 2$, right $n = 3$). Violet vertical line is a confidence marker at 0.5.

as drug abuse, psychoses, weight loss and AIDS, which posts challenges for algorithms in deciding on key factors associated with recurrent visits.

4) Confidence increases and support decreases with increasing number of comorbidities: From Figure 2, the range of confidence for single comorbidities is between 0.36 to 0.56; this range is not very wide for discriminating recurrent patients versus non-recurrent patients. To increase the discriminability, we use a combination of comorbidities rather than a single one from past visits. We plot the confidence and support ranges in Figure 3 to find a reasonable parameter for the number of comorbidities n. As n increases, the confidence range increases, and the range of supports reduces. We take the comorbidities size to be $n = 3$, as the majority of their confidence are over 0.5. We therefore use association rules on combinations of 3 past comorbidities from the previous ≤ 3 visits within last one year.

5) Conventional association rules: Our rule candidates are comorbidity combinations with 3 different comorbidities,

taken from ICD codes from the most recent 3 visits. Thus, for a comorbidity combination of 3 comorbidities $\{A, B, C\}$, the confidence/likelihood of recurrent patients (RP) is:

$$\text{Confidence:} \quad c(\{A, B, C\}) = P(RP|A, B, C); \quad (1)$$

and its support/prevalence is:

$$\text{Support:} \quad s(\{A, B, C\}) = P(A, B, C). \quad (2)$$

The conventional association rules method applies a user-specified threshold (τ) on support and then takes the highest confidence rule remaining. To simplify the notation, we use v to denote a combination of comorbidities.

$$v^{\star}_{AR} = \arg\max \text{Confidence}(v), \text{s.t. } \forall v, \text{Support}(v) \geq \tau. \quad (3)$$

6) Drawbacks of conventional AR with a fixed threshold: Figure 2 and Figure 3 show that there is a trade-off between confidence and support. First, in clinical applications, some diseases are less prevalent than others. However these may have higher confidences associated with recurrent visits,

rule	Confidence P(RP|rule)	Support P(rule)
A,B,C	High	Low
A,C,D	Low	High

MSAR aims at deciding between similar high-confidence low-support, and low-confidence high-support rules.

Combidity examples
A – Alcohol abuse
B – Blood loss
C – Congestive Heart Failure
D – Drug abuse

(A,B,C) connects to (A,C,D)
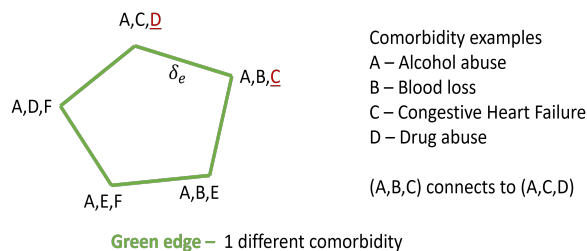
Green edge – 1 different comorbidity

Fig. 4: **Top**: The goal of MSAR is to choose between similar rules when one is higher confidence, lower support and the other is lower confidence but higher support.
Bottom: Illustration of min-similarity rule graph construction. The graph is built with each node being a triplet of comorbidities. An edge exists for similar rules with node pairs that differ by exactly one comorbidity. The optimization to balance confidence and support is only on pairs of similar rules with 1 comorbidity difference.

shown in Figure 2. Secondly, while increasing the number of elements (comorbidities) can boost confidence, it largely decreases the support.

The main drawbacks for conventional AR at a fixed threshold are: i) A large threshold could remove rare diseases, including the ones with high confidence. ii) A small threshold will make the algorithm very sensitive to outliers and potentially risks revealing the identity of a patient with rare conditions. These drawbacks motivate us to develop an algorithm to balance confidence and support automatically.

7) Motivation and definition of similar rules: For similar comorbidity combinations, selecting between high-confidence or high-support rules motivates us to balance confidence-support trade-off. For example, in Table II, we have two sets of comorbidities: the first is weight loss (WL), peripheral vascular diseases (PVD), hypertension, and the second set is WL, PVD, and fluid-electrolyte disorders. Those two sets have similar confidences, but they have very different supports. For a recurrent patient with all four of these comorbidities during their past three visits, the ability to differentiate which three comorbidities to prioritize as factors of recurrent visits is important, and serves as a key motivation of this solution. We therefore define pairs of similar rules and optimize the weights of confidence and support among them.

Similar rules are defined as the following: for rule set v_1 and rule set v_2, both of size 3, they are similar rules if and only if exactly one element is different between two sets. For example, for $v_1 = \{A, B, C\}, v_2 = \{A, B, E\}, v_3 = \{A, E, F\}$, v_1 and v_2 are similar rules, whereas v_1 and v_3 are not similar rules. In graph representation, we model each

rule as a vertex in a graph, and the edge is only present if there is only one different element between the two rule sets. Figure 4 illustrates this graph structure.

8) MSAR balances confidence and support for similar rules: Then, we formulate the combined confidence and support rule: for weights $w_c, w_s \geq 0$, where the sum of weights is one: $w_c + w_s = 1$, the rule score function on a vertex v_i is as follows:

$$r(v_i) = w_c \tilde{c}(v_i) + w_s \tilde{s}(v_i), \tag{4}$$

where $\tilde{c}(v_i), \tilde{s}(v_i)$ are **z−normalized confidence and support**, respectively. Now, for an edge e between vertices v_i and v_j, we define the difference by the following:

$$
\begin{aligned}
\delta(e) &= r(v_i) - r(v_j) \\
&= w_c \tilde{c}(v_i) + w_s \tilde{s}(v_i) - w_c \tilde{c}(v_j) - w_s \tilde{s}(v_j) \\
&= w_c(\tilde{c}(v_i) - w_c \tilde{s}(v_j)) + w_s(\tilde{s}(v_i) - \tilde{s}(v_j)).
\end{aligned}
\tag{5}
$$

We use $\delta_c(e), \delta_s(e)$ to denote the difference of confidence and support between two vertices, respectively. Then, equation (5) reduces to

$$\delta(e) = w_c \delta_c(e) + w_s \delta_s(e). \tag{6}$$

We require similarity to be inversely correlated to the individual differences; specifically, we define the similarity as the overall maximum difference minus each individual difference: $\text{sim}_e = w_c(\delta_{max} - \delta_c(e)) + w_s(\delta_{max} - \delta_s(e))$, where δ_{max} is obtained by taking the max over all δ_c and δ_s in all edges from the similarity graph: $\delta_{max} = \max(\max \delta_c(e), \max \delta_s(e))$. Therefore, the minimum similarity association rules can be derived by the following:

$$
\begin{aligned}
&\min_{w_c, w_s} \sum_{e \in G} \text{sim}_e^2 \\
&= \sum_{e \in G} (w_c(\delta_{max} - \delta_c(e)) + w_s(\delta_{max} - \delta_s(e)))^2
\end{aligned}
\tag{7}
$$
$$\textbf{s.t. } w_s + w_c = 1, w_s, w_c \geq 0$$

Equation (7) is in the form of convex quadratic programming (QP) with only 2 parameters. We use the off-the-shelf MATLAB *quadprog* solver to solve (7). Not only is the solution efficient to obtain, but also this algorithm is easy to deploy. From retrospective data, we can calculate $w_c^\star, w_s^\star$ as the optimal solution from (7), and then the minimum similarity score can be calculated as the following:

$$\text{MSAR score: } r_{\text{MSAR}}^\star(v_i) = w_c^\star \tilde{c}(v_i) + w_s^\star \tilde{s}(v_i). \tag{8}$$

The scores can be added directly to the learned rules along with confidence and support for each rule, and customers will be able to select top ranked rules based on MSAR scores. Table II gives example of scores and how rules are compared during algorithm execution.

III. Numeric results

We use the Banner Health EHR dataset to validate the proposed algorithm. It contains 8 years of retrospective data from 30 community teaching hospitals in the United States.

Rule candidates: comorbidities combinations	Confidence	Support	AR score (Confidence)	MSAR score (8)
1. Weight loss, Peripheral vascular disease, Hypertension	0.681	0.00871	**0.681**	-0.192
2. Weight loss, Peripheral vascular disease, Fluid-electrolyte disorders	0.675	0.0244	0.675	**0.0543**

(a) Case I: the lower confidence rule has a support that is larger enough compared to a similar rule, and it is chosen by MSAR.

Rule candidates: comorbidities combinations	Confidence	Support	AR score (Confidence)	MSAR score (8)
1. AIDS, Coagulopathy, Psychoses	0.819	0.000156	**0.819**	**1.0014**
2. AIDS, Coagulopathy, Renal failure	0.748	0.000198	0.748	0.300

(b) Case II: the lower confidence rule has only slightly greater support compared to a similar rule. AR and MSAR select the higher confidence rule.

TABLE II: MSAR balances confidence-support trade-offs with examples of comparison to the conventional association rules (AR) always selects the highest confidence rule. Bold texts indicate selected rules by AR and MSAR.

A. Statistics for all three recurrent patient criteria

At a population level, the combination of ED and inpatients produces a 30-day readmission rate of 12.3%, which is within a similar range to reported values on the HCUP dataset [28]. Specifically, the rates of recurrent inpatient and ED visits are 7% and 21%, respectively.

The percentage of total recurrent patients is about 25% of all visits. Noticeably, ED recurrent patients contribute to 94% of total recurrent patients (ED and inpatients combined).

B. MSAR results

To obtain past comorbidities associated with recurrent/frequent visits, we first take ICD-9 and ICD-10 codes from a patient's past visits as training data. To avoid duplicate counts, we only take the most recent (up to) 3 visits for both recurrent and non-recurrent patients. One patient is therefore only counted once in the training data, and this process results in more than $400,000$ unique patients in that set.

The constant value in optimization equation (7) δ_{max} is obtained by taking the maximum of $z-$normalized confidences and supports, and the value is 4.34 from our retrospective data. Optimal weight parameters of confidence and support are learned from retrospective data, and the numeric optimal solution is $w_c^\star = 0.778, w_s^\star = 0.221$. The weight of confidence is about 3.5 times than the weight of support. We substitute these values into (8) to obtain MSAR rule scores for all candidate rules.

1) MSAR balances confidence-support trade-offs compared to conventional AR: Table II demonstrates MSAR and AR selections for similar rules while one has high-confidence low-support and the other one with low-confidence high-support. In the top examples, the low confidence rule has larger support than the high confidence rule, and therefore MSAR picks the slighter lower confidence rule with much higher support. The bottom example illustrates the case when the low confidence rule has lower support than the high confidence rule, and MSAR agrees with the conventional highest confidence association rules. This behavior validated that **MSAR balances confidence-support trade-offs.**

2) MSAR successfully selects challenging high-confidence, low-support comorbidities relates to recurrent visits: In total 3860 rules from triplets of Elixhauser comorbidities are learned. To understand the contributions

Optimal weights	Mean (std) from CV
$w_c^\star$	0.785 (0.0120)
$w_s^\star$	0.215 (0.0120)

TABLE III: Distribution of optimal weights from $10-$fold CV with 80% randomly sampled training data in each fold. The standard deviations of optimal weights are fairly small, validating the consistency of learned weights by MSAR.

Algorithms	Mean (std) Rank-biased Overlap (RBO) [29]
MSAR	**0.973 (0.0138)**
XGBoost [20]	0.957 (0.0249)

TABLE IV: Pairwise RBOs for ranked comorbidities across 10 folds. MSAR comorbidities are ranked by MSAR scores, and XGBoost comorbidities are ordered by Shapley values [21] on training data. MSAR has more consistent outputs.

of different comorbidities from learned rules, we take the top 1000 (25.9%) rules and count the frequency of each comorbidity, which is depicted in Figure 5. Top categories include drug abuse, psychoses, neurological disorders, depression. Among them, drug abuse, reported as one of the top reasons for high utilization [9], even though with lower support, is identified by MSAR as a majority comorbidity associated with recurrent visits.

3) The consistency of learned weights and outputs of MSAR are validated through cross-validation (CV): To study the consistency of MSAR, we conduct a 10-fold CV with 80% of the training data in each fold. We first calculate optimal weights of confidence and support for each fold. Table III shows the mean and standard deviation of optimal weights, which indicates low variance in optimal weights from various folds. Then, we rank the comorbidities by their MSAR scores in the same way as the previous section III-B.2. Finally, we calculate pair-wise Rank-based Overlap [29] from the outputs of 10 folds, using ranked comorbidities lists

Algorithms	% of indiscriminate across comorbidities
MSAR	**0%**
XGBoost [20]	$33.3\% - 40\%$

TABLE V: The percentages of comorbidities with zero weights. These can't be differentiated from each other while selecting top comorbidities associated with recurrent visits.

Comorbidity	Frequency in top MSAR rules	Frequency in top-3 XGBoost features	non-zero % of frequency in top MSAR rules	non-zero % of frequency in top-3 XGBoost features
Drug abuse	0.35 (0.003)	10^{-6} (10^{-6})	**100%**	40%
Weight loss	0.13 (0.004)	5×10^{-7} (10^{-6})	**100%**	10%
AIDS	0.049 (0.009)	0.0005 (0.0002)	**100%**	**100%**

TABLE VI: Averages and standard deviations of frequency for example comorbidities, learned from MSAR rules, and non-zero percentage of frequency across 10 CV. When the frequency of a certain comorbidity is zero, it means that that comorbidity is not in any of top features associated with recurrent patients learned from algorithms. The non-zero percentage of frequency measures the percentage of frequency of certain comorbidity across 10 folds. MSAR is much more consistent in detecting low-support, high-confidence comorbidities than XGBoost.

ordered by MSAR scores. Table IV shows that the pair-wised overlap across 10 folds is very high with an average of 0.973 overlap, validating the consistency of MSAR outputs.

4) Compare with a popular XGBoost method, MSAR rules distinguishes across comorbidities better: We compare MSAR with a decision tree-based, popular and commonly used method for tabular datasets, called XGBoost [20]. For the comparison, we use the same fold split as for MSAR. We use 30 estimators, maximum depth at 6 and 0.05 as the learning rates, achieving on average 0.733 area under the Receiver Operating Characteristic curve.

Although Table IV shows that the consistency of XGBoost across validations is lower than MSAR, the RBO overlap of outputs is fairly high. However, the main drawbacks of XGBoost for this task is the lack of interpretability and discrimination across different comorbidities. We investigated the frequency of comorbidities selected by the XGBoost model from its top 3 rules and compared them with MSAR, whose numbers are reported in Table V. A large percentage of comorbidities, ranging from $33\% - 40\%$ (corresponding to $10 - 12$ comorbidities) from $10-$fold CV, are not selected as top 3 rules due to their low prevalences in training sets with around $330,000$ patients. On the other hand, MSAR scores for all comorbidities are non-zero, making ranking possible for outputting top choices. **In summary, MSAR can distinguish all** 30 **comorbidities; however** $10 - 12$ **comorbidities trained by XGBoost can't be discriminated.**

5) Compare with XGBoost, MSAR rules select low-support high confidence comorbidities better: Table VI gives frequency from learned rules by MSAR and XGBoost of a few challenging comorbidities with low-support, high confidence. We can conclude that **MSAR is much more consistent at identifying challenging high-confidence low-support comorbidities compared to XGBoost.** Additionally, XGBoost missed low missed some high confidence however low support comorbidities such as drug abuse, weight loss. Especially, **drug abuse** is reported in the literature as an important factor for recurrent patients, **even though having a high confidence, due to a lower support than some chronic conditions, XGBoost fails to select it as top** 3 **contributor in** 6 **out of** 10 **folds, whereas MSAR successfully selects drug abuse out of all folds.**

IV. CONCLUSION AND DISCUSSION

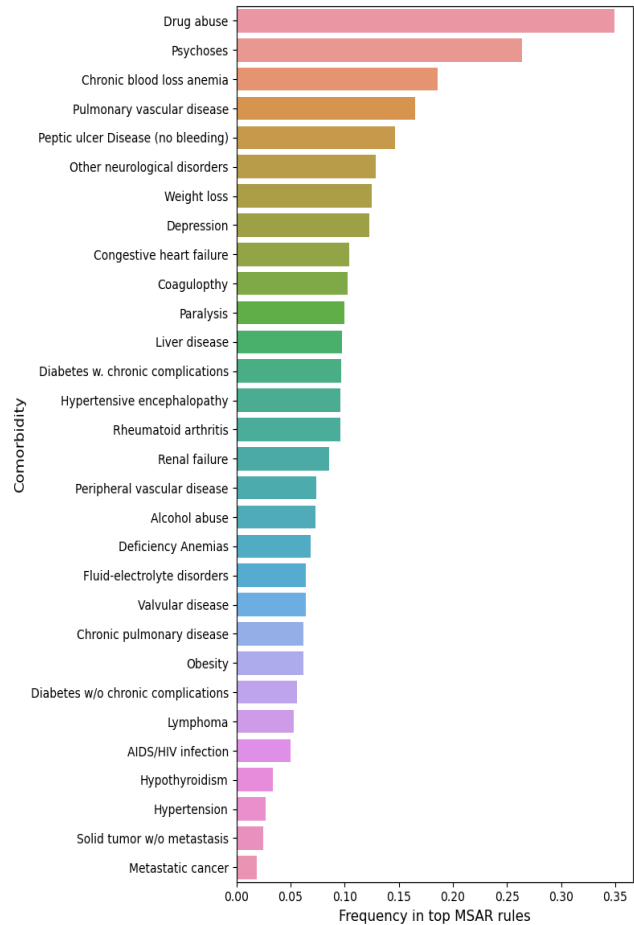In this work, we propose a framework for the healthcare system's need to identify recurrent patients. It identifies

Fig. 5: Frequency of comorbidities from the top 25% of learned MSAR rules. Drug abuse, even with low prevalence compared to other comorbidities, has been successfully identified as being in 35% of those top rules for recurrent visits.

recurrent patients, is executable at admission, and finds the top comorbidities associated with their recurrent visits. This feat is accomplished via a novel proposed algorithm MSAR that balances the trade-off between confidence and support with optimal weights learnable from retrospective data. MSAR successfully learn high-confidence but low-support comorbidities, which is challenging for other algorithms such as XGBoost. The proposed framework can be used to assist future decisions to avoid future recurrent visits, such as recommending social-behavioral interventions and other

relevant outpatient care.

The proposed MSAR algorithm has the benefits of being customizable and easy to deploy in potential future products. Those who may find interest in such a proposal include leaders in the operations and population health space, along with front-line staff such as ED nurses and physicians, inpatient charge nurses, and patient flow coordinators. This tool could help the care team recommend potential interventions (e,g, rehabilitation, education programs, follow-up visits with primary care, medication monitoring, etc.) that help the patient reduce the risk of recurring ED and inpatient visits.

COMPLIANCE WITH ETHICAL STANDARDS

Banner Health data use was a part of an ongoing retrospective deterioration detection study approved by the Institutional Review Board of Banner Health and by the Philips Internal Committee for Biomedical Experiments. Requirement for individual patient consent was waived because the project did not impact clinical care, was no greater than minimal risk, and all protected health information was deidentified to limited dataset.

REFERENCES

[1] "One percent of americans account for 22 percent of health care spending." [Online]. Available: https://www.medscape.com/viewarticle/541137.

[2] S. H.-X. Ng, N. Rahman, I. Y. H. Ang, S. Sridharan, S. Ramachandran, D. D. Wang, C. S. Tan, S.-A. Toh, and X. Q. Tan, "Characterization of high healthcare utilizer groups using administrative data from an electronic medical record database," *BMC health services research*, vol. 19, no. 1, pp. 1–14, 2019.

[3] M. W. Smith, C. Stocks, and P. B. Santora, "Hospital readmission rates and emergency department visits for mental health and substance abuse conditions," *Community mental health journal*, vol. 51, no. 2, pp. 190–197, 2015.

[4] M. B. Doupe, W. Palatnick, S. Day, D. Chateau, R.-A. Soodeen, C. Burchill, and S. Derksen, "Frequent users of emergency departments: developing standard definitions and defining prominent risk factors," *Annals of emergency medicine*, vol. 60, no. 1, pp. 24–32, 2012.

[5] S. Crane, L. Collins, J. Hall, D. Rochester, and S. Patch, "Reducing utilization by uninsured frequent users of the emergency department: combining case management and drop-in group medical appointments," *The Journal of the American Board of Family Medicine*, vol. 25, no. 2, pp. 184–191, 2012.

[6] P. L. Owens, R. Mutter, and C. Stocks, "Mental health and substance abuse-related emergency department visits among adults, 2007," *UMBC School of Public Policy Collection*, 2010.

[7] J. M. Pines, B. R. Asplin, A. H. Kaji, R. A. Lowe, D. J. Magid, M. Raven, E. J. Weber, and D. M. Yealy, "Frequent users of emergency department services: gaps in knowledge and a proposed research agenda," *Academic Emergency Medicine*, vol. 18, no. 6, pp. e64–e69, 2011.

[8] G. L. Larkin, C. A. Claassen, J. A. Emond, A. J. Pelletier, and C. A. Camargo, "Trends in us emergency department visits for mental health conditions, 1992 to 2001," *Psychiatric services*, vol. 56, no. 6, pp. 671–677, 2005.

[9] E. Stranges, K. Levit, C. Stocks, and P. Santora, "State variation in inpatient hospitalizations for mental health and substance abuse conditions, 2002–2008: Statistical brief# 117," 2011.

[10] T. P. O'toole, A. Conde-Martel, J. H. Young, J. Price, G. Bigelow, and D. E. Ford, "Managing acutely ill substance-abusing patients in an integrated day hospital outpatient program: Medical therapies, complications, and overall treatment outcomes," *Journal of general internal medicine*, vol. 21, no. 6, pp. 570–576, 2006.

[11] S. Kangovi, N. Mitra, D. Grande, M. L. White, S. McCollum, J. Sellman, R. P. Shannon, and J. A. Long, "Patient-centered community health worker intervention to improve posthospital outcomes: a randomized clinical trial," *JAMA internal medicine*, vol. 174, no. 4, pp. 535–543, 2014.

[12] J. R. McKay, A. I. Alterman, A. T. McLellan, E. C. Snider *et al.*, "Treatment goals, continuity of care, and outcome in a day hospital substance abuse rehabilitation program," *American Journal of Psychiatry*, vol. 151, no. 2, pp. 254–259, 1994.

[13] Z. Martinez, E. Koker, A. Truchil, and H. Balasubramanian, "Time and effort in care coordination for patients with complex health and social needs: Lessons from a community-based intervention," *Journal of Interprofessional Education & Practice*, vol. 15, pp. 142–148, 2019.

[14] A. E. Nijhawan, E. Kitchell, S. S. Etherton, P. Duarte, E. A. Halm, and M. K. Jain, "Half of 30-day hospital readmissions among hiv-infected patients are potentially preventable," *AIDS patient care and STDs*, vol. 29, no. 9, pp. 465–473, 2015.

[15] S. Ostling, J. Wyckoff, S. L. Ciarkowski, C.-W. Pai, H. M. Choe, V. Bahl, and R. Gianchandani, "The relationship between diabetes mellitus and 30-day readmission rates," *Clinical diabetes and endocrinology*, vol. 3, no. 1, pp. 1–8, 2017.

[16] C. J. Grafe, R. Z. Horth, N. Clayton, A. Dunn, and N. Forsythe, "How to classify super-utilizers: A methodological review of super-utilizer criteria applied to the utah medicaid population, 2016–2017," *Population health management*, vol. 23, no. 2, pp. 165–173, 2020.

[17] R. Agrawal, T. Imieliński, and A. Swami, "Mining association rules between sets of items in large databases," in *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, 1993, pp. 207–216.

[18] C. Rudin, B. Letham, and D. B. Madigan, "Learning theory analysis for association rules and sequential event prediction," 2013.

[19] T. Scheffer, "Finding association rules that trade support optimally against confidence," in *European conference on principles of data mining and knowledge discovery*. Springer, 2001, pp. 424–435.

[20] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.

[21] L. S. Shapley, *A Value for N-Person Games*. Santa Monica, CA: RAND Corporation, 1952.

[22] "Services, centers for medicare and medicaid, hospital readmission reduction." [Online]. Available: https://www.cms.gov/Medicare/Quality-Initiatives-Patient-Assessment-Instruments/Value-Based-Programs/HRRP/Hospital-Readmission-Reduction-Program

[23] "ICD-10," https://en.wikipedia.org/wiki/ICD-10, 2010-09-30.

[24] A. Elixhauser, C. Steiner, D. R. Harris, and R. M. Coffey, "Comorbidity measures for use with administrative data," *Medical care*, pp. 8–27, 1998.

[25] M. E. Charlson, P. Pompei, K. L. Ales, and C. R. MacKenzie, "A new method of classifying prognostic comorbidity in longitudinal studies: development and validation," *Journal of chronic diseases*, vol. 40, no. 5, pp. 373–383, 1987.

[26] R. A. Deyo, D. C. Cherkin, and M. A. Ciol, "Adapting a clinical comorbidity index for use with icd-9-cm administrative databases," *Journal of clinical epidemiology*, vol. 45, no. 6, pp. 613–619, 1992.

[27] "Hcuppy 0.0.7 package." [Online]. Available: https://pypi.org/project/hcuppy/.

[28] "Healthcare cost and utilization project (hcup) dataset." [Online]. Available: https://www.hcup-us.ahrq.gov/reports/statbriefs/sb248-Hospital-Readmissions-2010-2016.jsp

[29] W. Webber, A. Moffat, and J. Zobel, "A similarity measure for indefinite rankings," *ACM Transactions on Information Systems (TOIS)*, vol. 28, no. 4, pp. 1–38, 2010.