Interpretable Clustering for Prototypical Patient Understanding: A Case Study of Hypertension and Depression Subgroup Behavioral Profiling in National Health and Nutrition Examination Survey Data

Pei-Yun S. Hsueh, PhD[1], Subhro Das, PhD[2]
[1]IBM Watson Research Center, Yorktown Heights, NY, USA

Introduction

Behavioral factors are the key contributors to morbidity and mental health risk [1]. For cardiovascular diseases alone, behavioral factors have contributed to 41 percent of its global disease burden [2]. In recent years, clinical guidelines (e.g., the 2017 Type II Diabetes management guideline [3]) have established the need to better understand patient behavior and identify patient-centered goals for care plan personalization. However, such guidelines for personalization are usually not formalized and will not grow organically with the increasing knowledge of patients. The key question is, therefore, how to further inform personalized care plan by accurately and reliably quantifying the behavioral patterns from target user-generated health data?

The recent emergence of health consumers and person health IT technologies made available an unprecedented amount of patient-generated health data (PGHD) [4]. To augment the current guideline-based care plan with PGHD insights, we introduce an interpretable clustering method for uncovering outcome-differential behavioral patterns in distinctive patient subcohorts. In addition, to further enable case-based reasoning for subcohort behavioral profiling, the proposed method also explores adding the identification of prototypical users for sense-making explanation and persuasion. Curating the learned subcohort profiles has been expected to facilitate a hybrid approach of PGHD-augmented guideline deployment.

Methods

Using the self-reported outcome survey and sensor measurement data in National Health and Nutrition Examination Survey (NHANES) during 2005 and 2006, we deployed an interpretable clustering method to identify distinctive behavioral profiles related to two types of proxy outcome indicator: blood pressure control (BP) and depression. For the first BP study, we follow [5]'s footstep to use a set of 64 features generated from the accelerometer measurements, blood pressure readings and questionnaire items of 5,695 subjects. The set of features include physical activity intensity, duration, variance of activity, systolic and diastolic BP, and so on. Similarly, for the second Depression study, we use only the set of 38 features generated from the accelerometer sensor readings.

We evaluate the proposed Interpretable Clustering (IC) method that segments subjects into different behavioral segments based on their behavioral patterns and proxy outcomes. Our hypothesis is that although each feature contains only weak behavioral signals for the whole population, when considered collectively in certain sub-cohorts, these features can explain rich behavioral patterns that matter to some proxy health outcome indicator, e.g., whether a user scores less than 4 (minimal depression or none) using the Patient Health Questionnaire (PHQ-9).

In particular, we first apply the Locally Supervised Metric Learner (LSML) [6] of patient similarity analytics to estimate the outcome-adjusted behavioral distances between the users. Then, based on the adjusted behavioral distances, hierarchical clustering is employed to generate sub-cohorts and learn the key features (which contain behavioral signals about implicit user preferences and barriers) that drive the differential outcomes. An automatic tuning algorithm is applied to determine the optimal number of segments. In addition, as we expect it to be easier to interpret patient need from behavior profiles by examples, the proposed method also includes a component to identify prototypical examples (i.e., a set of top 10 subjects who are the closest to the centroid of each behavioral segment in the outcome-adjusted distance space).

Results

For both the BP and Depression datasets, five distinctive sub-cohorts are identified, each appearing in more than 1% of the subjects. Statistical and information theoretic measures are applied to infer cross- and within-subcohort behavioral patterns among subjects and identify 3-10 distinctively informative patterns for each subcohort.

In order to allow for evaluation, both global and personalized predictive models are constructed. While the former is trained on the whole population, the latter is trained for each sub-cohort. In addition, we also compare the outcome-differentiating Interpretable Clustering (IC) method with K-means clustering and Agglomerative clustering method that do not directly account for outcome differentiation during the clustering process. The evaluation includes clustering quality, i.e. Sihouette score that indicate internal consistency by measuring how similar an object is to its own cluster (cohesion) compared to other clusters (separation), and outcome-differentiating F-statistics and z-scores. Results show that the IC-based personalized prediction model yields significantly better outcome differentiation results in terms of the sub-cohorts it produced. Although the personalized IC models did not produce more accurate risk scores than the general model, it captures more intuitively understandable patterns. Behavioral profile evaluation metrics show that further zooming into the prototypical examples helps identify distinctively more different profiles across the sub-cohorts, as demonstrated in the differences in the top behavioral patterns and those between the individual and global risk factors.

Discussion

The consumer and pervasive health informatics community is increasingly handicapped by the problem of not being able to interpret patient need. Previously, a family of patient phenotyping approaches (including the use of deep learning) has been developed to extract distinctive phenotypes from medical records [7,8] and "digital phenotypes" from digital data (including device/sensor data and self-reported outcomes) [9,10]. One missing key is a behavioral learning mechanism that can sift through user-generated health data to identify outcome-differential patient behavioral patterns.

To address this important problem, this presentation has conducted a 3-fold investigation: First, it has reviewed existing guidelines and data-driven methods of patient subgroup identification and digital phenotyping and identified gaps for realizing the personalization goal of guidelines. Secondly, this presentation also proposed an interpretable clustering approach that adds behavioral response pattern understanding into patient subgroup analysis and informs guideline deployment. Lastly, this presentation provided evaluation by scanning through NHANES data to identify behavioral profiles and prototypical examples related to blood pressure control and depression.

The result demonstrates a large potential for learning methods to derive outcome-differentiating patient behavioral insights, which in turn lend support to better decision tools for care team and facilitate positive changes in patient behavior. In addition, the proposed method also improves on the interpretability of the learning results by providing more distinctively different behavioral patterns and prototypical examples. Finally, by curating the detected sub-cohort and differential behavioral patterns, a natural next step would be to investigate whether we can better deploy guideline-based personalized care plans with better individualized interpretations.

References

1. Ezzati M, Riboli E. Behavioral and Dietary Risk Factors for Noncommunicable Diseases. *New England Journal of Medicine* 2013;*369*(10):954–964.
2. Lloyd-Jones DM, Hong Y, Labarthe D, et al. Defining and setting national goals for cardiovascular health promotion and disease reduction: The American Heart Association's strategic impact goal through 2020 and beyond. Circulation. 2010;121:586-613.
3. Reusch JEB, Manson JE. Management of Type 2 Diabetes in 2017: Getting to Goal. *JAMA.* 2017.
4. Marschollek M, Woodcock J, et al. A Semi-Quantitative Method to Denote Generic Physical Activity Phenotypes from Long-Term Accelerometer Data – The ATLAS Index. *PLoS ONE*, 2013;*8*(5):e63522.
5. Hsueh PS, Chang H, Ramakrishnan S. Next-Generation Wellness: A Technology Model for Personalizing Healthcare. Healthcare information Management, 4th Edition. Springer Publishing, 2016.
6. Ng K, Sun J, Hu J, Wang F. Personalized Predictive Modeling and Risk Factor Identification using Patient Similarity. *AMIA Summits on Translational Science Proceedings*. 2015;2015:132-136.
7. Shivade C, Raghavan P, Fosler-Lussier E, Embi PJ, Elhadad N, Johnson SB, Lai AM. A review of approaches to identifying patient phenotype cohorts using electronic health records. *JAMIA*, 2014;*21*(2):221–230.
8. Delude CM. Deep phenotyping: The details of disease. *Nature* 2015;*527*(7576):S14–S15.
9. Jain SH, Powers BW, Hawkins JB, Brownstein JS. The digital phenotype. Nat Biotechnol 2015;33(5):462-463.
10. Onnela J, Rauch SL. Harnessing smartphone-based digital phenotyping to enhance behavioral and mental health. *Neuropsychopharmacology 2016;41*(7):1691-1696.