

## On Intelligence Augmentation and Visual Analytics to Enhance Clinical Decision Support Systems

Tsipi Heart  
Ono Academic College,  
Kiryat Ono, Israel  
[tsipi.h@ono.ac.il](mailto:tsipi.h@ono.ac.il)

Rema Padman  
Carnegie Mellon University,  
Pittsburgh, Pennsylvania  
[rpadman@cmu.edu](mailto:rpadman@cmu.edu)

Ofir Ben-Assuli  
Ono Academic College,  
Kiryat Ono, Israel  
[ofir.benassuli@gmail.com](mailto:ofir.benassuli@gmail.com)

David Gefen  
Drexel University,  
Philadelphia, Pennsylvania  
[gefend@drexel.edu](mailto:gefend@drexel.edu)

Robert Klempfner  
Sheba Medical Center,  
Ramat Gan, Israel  
[Robert.Klempfner@sheba.health.gov.il](mailto:Robert.Klempfner@sheba.health.gov.il)

### Abstract

*Human-in-the-loop intelligence augmentation (IA) methods combined with visual analytics (VA) have the potential to provide additional functional capability and cognitively driven interpretability to Decision Support Systems (DSS) for health risk assessment and patient-clinician shared decision making. This paper presents some key ideas underlying the synthesis of IA with VA (IA/VA) and the challenges in the design, implementation, and use of IA/VA-enabled clinical decision support systems (CDSS) in the practice of medicine through data driven analytical models. An illustrative IA/VA solution provides a visualization of the distribution of health risk, and the impact of various parameters on the assessment, at the population and individual levels. It also allows the clinician to ask “what-if” questions using interactive visualizations that change actionable risk factors of the patient and visually assess their impact. This approach holds promise in enhancing decision support systems design, deployment and use outside the medical sphere as well.*

### 1. Introduction

Decision support systems (DSS), clinical decision support systems (CDSS) included, are designed to provide guidance to decision makers as they analyze new situations and require assistance in assessing situations and in answering “what if” type questions. CDSS do so mostly by recording existing knowledge and processing techniques, sometimes even applying a straightforward rule-based dataset that applies current thinking to new data. Those objectives are not new [1], dating back to Ledley and Lusted [2]. Sheppard and Kouchoukos [3] added automated rules and tests to CDSS architecture, resulting in a consensus of the need to analyze patient data to create case-specific advice [4]. More recently, machine learning has also

been added [5], a necessity considering the size and complexity of healthcare datasets [6]. CDSS currently support many medical fields including patient care at the bedside, medical imaging, pharmacology, pharmacogenomics, and pathology, among others [7].

CDSS have been shown to improve medical best practices [8], increase quality of care and patient safety [9], reduce medical errors [10], and support diagnosis [11]. Indeed, CDSS that assess renal function, pregnancy complications, duplicate order entry, drug allergy, and drug recommendation, have a long history in reducing rates of medical errors [12], with early studies on CDSS claiming decrease of 12.7% decrease in total charges and a 0.9 day decrease in length of stay [13]. Based on the experiences of some leading, large, highly computerized, and internationally recognized medical centers, CDSS are under-utilized if used at all for diagnosing patients [1].

Apart from reducing errors, the objective of CDSS is to improve clinicians’ decision-making processes and provide medical insights at the point of care [14], a hard to achieve objective [7]. In part, that is because designing a medical treatment path often requires real-time, interactive processing of patient data by the clinicians involved. That interactive processing is necessary to attain more accurate clinical predictions of what a course of treatment may entail based on the analysis of many other patients’ data while considering a multitude of varying parameters. Such interactive processes could also decrease the incidence of errors. Errors are a major obstacle to the voluntary adoption of CDSS by physicians [7]. Moreover, and perhaps contributing to that error rate, are patients with multi-morbid conditions.

Interactive CDSS are at a rather extreme point in current classifications. Belard & Buchman [1] suggested that CDSS can be classified into knowledge-driven, rule-based support systems and data-driven, probability analysis support systems.

Rule-based systems include CDSS that alert about medication hazards [14] and drug interactions based on existing guidelines (even if these ignore patient-specific situations) [10]. Closer to probability analysis, CDSS are systems such as those that support dementia diagnosis by applying reasoning strategies, some even based on ambiguous and incomplete information [15, 16], diagnose prostate cancer [11] and predict whether a person is a carrier of toxic gene mutations that cause breast cancer [12]. Adding the degree of user interaction, Liedgren & Elvhage [17] classified DSS on a grid with descriptive or normative on one axis, and static or user interactive on the other (Figure 1). Interactive CDSS with a data-driven probability analysis engine fall into the top of the left quadrant being designed to address very complex medical conditions without relying on a set of rules only.

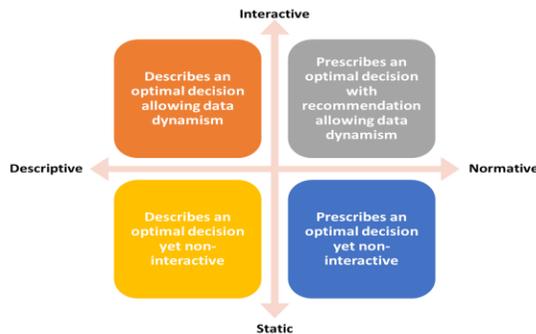


Figure 1. DSS classification based on [17]

### 1.1. Barriers to CDSS implementation

Among the major obstacles to successful deployment of CDSS are perceptions that the CDSS obscures physician autonomy and obligation and that the recommendations are insufficiently relevant and current [1]. The result is a lukewarm attitude towards current CDSS because it is perceived that the analysis and the rules behind the CDSS are too simple, and so physicians, specifically experts, are reluctant to rely on the CDSS because it makes it seem, wrongly, that their profession can be automated. Simple CDSS that record existing rule-bases just do not provide a convincing incentive to potential adopters. However, when the decision is very complex and the analysis is not perfectly clear to the cardiologists, then using the CDSS is too risky. One way of overcoming this conundrum is by allowing the design and subsequent interaction with the CDSS to be user-centered [18]. This also increases its perceived utility [19].

An alternative approach, examined in this study, is to replace the DSS rule-base with a dynamic AI-based, intelligence augmentation and visual analytics that provides additional functional capability and

cognitively-driven interpretability to DSS-enabled risk assessment. This AI-driven, visually augmented approach stands in contrast to typical DSS that freeze knowledge at the point where its rule-base has been created, potentially resulting in the dangerous ossification of knowledge in a rapidly changing context. What is needed is a dynamic support system that is not based on a rule base alone, but rather a system that lets the human expert explore options based on a real-time analysis of existing relevant data without being forced into a rigid pattern of thought that had been prevalent when the DSS or CDSS was created. Such a system could allow decision makers to consider “what if” questions dynamically and allow any kind of analysis that existing data had measured. A dynamic visual representation of data analysis results is called Visual Analytics (VA), and is currently perceived as the preferred mode of informing users about complex results [20]. VA can be static or collaborative, with the latter allowing users to harness their expertise to improve the system’s performance. Thus, the computer and user work in collaboration, with each partner contributing their strengths and unique capabilities: the computer in analyzing complex, big data and creating visual data presentations, and the user in pinpointing the search and analysis to points of interest, employing the holistic view, which is part of human expertise that a computer lacks. This approach is termed Intelligence Augmentation (IA) [21].

We describe a CDSS prototype that is composed of an AI-based back-end analysis engine for predictive analytics and results visualization, and a VA-IA based front-end component, for the user and computer to interact with and collaborate in a human-in-the-loop format, for optimal decision making. The AI-based engine dynamically and without resorting to a predefined rule-base analyzes all the medical tests, diagnostics, prognosis, and treatments, including the ICD codes of thousands of patients with the same overall medical condition as observed over many years. (ICD is the International Statistical Classification of Diseases and Related Health Problems. It was created and is continuously updated by the World Health Organization (WHO) as a classification system. More about ICD can be found at <https://www.who.int/classifications/icd/en/>).

In the next two sections, we briefly describe the IA and VA concepts.

### 1.2. Intelligence augmentation (IA)

With the prevalence of large amount of data and ubiquity of computing power, we recently see more tasks being performed by computers, profoundly

changing the modern way of life at work and beyond [22]. The promise of AI has also given rise to the growing debate on whether a thinking machine is possible and whether the notion of such a machine provides a suitable model of the human mind [23]. Specifically, problems in artificial intelligence are often framed in terms of matching or surpassing human performance [24]. While there is an increased concern that smart machines may replace many humans in decision making, current evidence is that, with greater computational information-processing capacity and an analytical approach, machines can amplify human cognition when addressing complexity, whereas humans offer a more holistic, intuitive approach in dealing with uncertainty in decision making [25]. This new concept is termed Intelligence Augmentation (IA), in contrast to human-computer interaction (HCI) where man and machine are considered separate entities [26].

IA focuses on building systems that allow humans and machines to work together, as opposed to AI, which is focused on building intelligence in machines. Intelligence Augmentation provides an efficient theoretical framework for understanding what humans can accomplish and how artifacts and tools can be designed and evaluated to empower human beings and to alter tasks [27].

The goal of IA, in the context of this study, is to use computational systems to help physicians in the decision making process by means of an ongoing human-machine dialogue [28]. The belief underlying this human-machine collaboration is that it potentially allows the humans to gain the time to concentrate on other parts of the process and thereby expanding their perspectives and augmenting their intellect [29]. IA has been recently proposed as assisting clinical assessment workflow, to sequentially augment physician assessment of patients' symptoms, while integrating their socio-demographic determinants and heterogeneous biological measures to accurately predict treatment outcomes using machine learning methods [30]. Yet, cumbersome human-computer communication and challenging user interfaces pose barriers to wide utilization of these IA platforms in highly stressed and time-constrained clinical settings. Recent research has suggested interactive visualizations leveraging advanced analytics, which are yet relatively easy to use and comprehend, as a plausible human-computer communication platform [31, 32] suitable for environments such as clinical care delivery.

### **1.3. Visual analytics (VA)**

Visual analytics (VA) combines computational analysis and interactive visualization-based user interfaces to support analytical reasoning and human cognition, incorporating disciplines including cognitive science, data mining and machine learning [33]. This enables interactive, analytical interpretation of very large and complex datasets through the integration of human intuition and machine learning [20, 34, 35]. VA is especially applicable where the data are too large, and the decisions too complex, for purely visual methods, and therefore require data processing and mining capabilities in exploratory setting. This necessitates an interdisciplinary approach [29]. VA is particularly useful when dealing with multivariable items, aka "dimensions" that are sourced from multiple systems, are sparsely populated, and might even be of questionable quality. Comprehending multi-dimensional information is difficult, known as 'the curse of dimensionality' [36]. Health data is an example of multivariable, multi-sourced data. Policy and decision makers in healthcare have long valued the capacity of healthcare information to improve clinical decision making by assuming that such data, if available to the right people in an appropriate format and at the right time, could significantly drive healthcare effectiveness and efficiency [37]. It is therefore not surprising that many recent studies have examined the application of VA in the healthcare domain [38-40].

## **2. Methods**

### **2.1. CDSS requirements analysis**

Requirements for the CDSS design were gathered from a thorough literature survey and interactions with clinical experts that resulted in the development of an early prototype (not displayed due to space limitations). This prototype was then demonstrated to experts in diabetes mellitus and heart disease [31], whose comments were used for the generation of the second version of the prototype which is the focus of this study. This version was reviewed by three cardiologists, who saw a demonstration of the prototype using patient data from their own setting and commented on the displayed features.

### **2.2. Data set for the AI-driven engine**

The research dataset, focused on the cardiovascular condition, was obtained from the Sheba Medical Center, which is the largest hospital in Israel,

located in the center of the country [41]. It is an extensively computerized hospital, where an integrated EHR system is used in the Emergency Department (ED) and all inpatient departments, fully replacing paper-based medical records. In recent years, Sheba has established a research-focused data warehouse that collects data from the various information systems installed in the hospital. Patient socio-demographic data is obtained from the Israel Population Registry system and integrated into the EHR. Each patient admitted to the hospital, via the ED or otherwise, who is diagnosed with a cardiovascular condition as a primary diagnosis, has their data uploaded to the data warehouse. All patient information is then retrieved from the other systems using the individual's unique national ID available in the Israeli Population Registry. This unique identifier allows accurate location and integration of patient data. We removed variables with more than 20% missing values based on clinical judgment and used the Predictive Mean Matching imputation method [42] with the mice package in R [43] to impute missing values in key attributes that could not be eliminated from consideration. Diagnoses were coded based on the ICD-9 standard. The resulting dataset included 10,763 patients who visited the Sheba Medical Center between 2010 and 2017 and were diagnosed with Congestive Heart Failure (CHF) as the primary diagnosis. Each patient record included 211 attributes.

### 2.3. The feature selection process

Two senior cardiologists selected 30 features they considered most effective to predict 30-day readmission or 30-day mortality based on extant scientific knowledge and local practices and outcomes. Next, we evaluated multiple feature selection criteria and identified the Information Gain (IG) [44] method to select a set of 30 features, to investigate if the method may delineate features not commonly selected by cardiologists. Information Gain (IG) favors splits with small counts but many unique values.

The human and machine generated feature sets were then merged to a larger feature set consisting of 45 variables. This was evaluated by two different experts (senior cardiologists), to mark those features they regarded as most important for mortality and readmission (two separate sets), respectively. We tested the lists for inter-rater agreement using Cohen's Kappa [45], resulting in moderate to poor agreement among the raters (0.688 for mortality and 0.474 for readmission, before adjusting for chance agreement). Next, we extracted the features disagreed upon and returned to the same physicians two weeks after the

first session, asking them to re-rate only those items, without indicating prior choices. This second round yielded high agreement (0.893 and 0.938 for readmission and mortality, respectively, and 0.777 and 0.858 after adjusting for chance agreement), resulting in a final set of 29 variables for mortality and 28 variables for readmission. We termed the sets "Human-Machine Collaborative" sets.

### 3. The new CDSS approach

A key characteristic of the approach utilized by the new CDSS is that it allows the clinician to apply a data-driven, exploratory approach. This analysis encompasses data discovery by assessing patterns in the medical records of the other patients to inform the decision maker. This is a bottom-up approach, in contrast to the objectives of some CDSS that attempt to enforce policy and standardization of decision making through a policy driven, top-down process. The new CDSS is about empowering the clinicians with additional visual tools that can address new topics and questions, supported in the backend by AI, allowing the clinician to review medical records of other patients that inform the visual presentation. Moreover, as the visual presentation displays numbers based on predictive analytics, the recommendation is mostly quantitative and leaves the actionable interpretation to the clinician. For example, showing the probability that changing the systolic blood pressure of the patient by so many units will change the likelihood of expected mortality by so many percentage points. This type of analysis can, providing the dataset it is based upon is large enough, also allow answering many complex what-if questions concurrently, allowing interaction analyses. For example, a physician may ask about the result of both reducing systolic blood pressure by 10% but only among patients with a certain condition (including one or more existing comorbidities). This ad hoc type of investigation is unique to current clinical decision making due to the availability of highly granular data on many patients collected as part of routine medical practice and applies without relying on existing rules and research. The approach also allows studying patient-reported adverse reactions that have not yet been identified in the literature, and hence are absent from any rule-base a CDSS could possibly apply. This ad-hoc and evidence-based analysis also allows an investigation of the what-if questions in the context of a specific demographic or social determinant, such as women only or living in a particular neighborhood type. Such an assignment of numeric values to the risk associated with the co-occurrence of medical conditions, tests, and treatments, could also open the

opportunity of adding hypothesis testing and complex multilayer path analysis models. This ability to test hypotheses is another characteristic that sets this CDSS approach in a new direction. Moreover, this approach could provide insights also about multi-morbid conditions that have not yet been incorporated into existing CDSS models, but about which there is data available from other patients, past and present, with an equivalent profile. Being exploratory also means that the proposed CDSS approach could reveal data-derived previously unknown patterns, and that these patterns may vary across populations.

Basically, the new approach stands in contrast to existing CDSS approaches that are beneficial in training based on known medical case-studies and guidelines, and in enforcement of rules as well as testing for mistakes such as drug interactions. However, based on the extant literature, and on our preliminary interviews with cardiologists to whom the new CDSS has been demonstrated, the approach of common CDSS is generally too rigid and holds limited benefits to them once the rules are known and training is over. In contrast, the new CDSS approach, because it is about empowering the decision maker, rather than checking or prescribing their actions, was accepted enthusiastically. That is important also because one of the major holdbacks in the adoption of healthcare IT is the reluctance of clinicians to adopt the new IT and its organizational solution, partly because it infringes on rather than augments their capabilities [46]. There needs to be a viable value proposition in new IT if it is to be adopted. The intelligence augmentation enabled CDSS, the easy-to-use visual presentation, and its AI-based analysis of the population provides such a value proposition. That value proposition is increased by the ability to examine specific patient profiles and where the patient fits into each such profile, such as by varying demographics, social determinants and clinical history parameters. Applying such a data-driven exploratory approach to AI augmented decision making, we extended an existing prototype, developed by one of the co-authors and collaborators in a prior study [31], to include additional requirements for a visual CDSS prototype and presented it to cardiologists, implemented as an RShiny application (<https://shiny.rstudio.com/>).

## 4. Results

### 4.1. The CDSS functions

Figures 2 to 7 show the pertinent aspects of this CDSS as it was demonstrated to cardiologists on a subset of the Sheba CHF data. Creating and managing

appropriate expectations is a crucial first set in this process. To date, many cardiologists among those we interviewed were disappointed by other approaches to CDSS that had solidified and verified current best-practices, tested drug interactions, workflows, and billing procedures. Those aspects are beyond doubt very important in managing standard processes and enforcing quality controls, many of which are mandated. The current CDSS approach is not intended to replace those processes, but, rather, to build on top of those processes another layer to augment their clinical decision-making capability. The initial screen, Figure 2, is intended to set those expectations in place.

Based on those expectations, the cardiologist is then provided with tools to run descriptive analysis on the data. This is not the core contribution of the new CDSS, but it serves an important purpose of increasing the transparency and explainability of the CDSS, as well as exploration of the population under investigation. Figure 3 is an example of one such variable, age. This aspect is not unique or new. It is merely showing that the population at hand is indeed quite elderly, as expected, although the number of relatively younger patients is not negligible. Thus, a user can decide to investigate only a sub-population of older or younger patients. The data show that CHF increases gradually with age, reaching the peak number of people diagnosed with it in this sample to be between 84 and 87 years of age. The fact that this is aligned with common medical knowledge may strengthen the cardiologists' trust in the system. For this purpose, this CDSS allows having all prediction performance scores available to the decision makers (e.g. Precision, Recall, Confusion matrix and Accuracy) on any selected population, so that a user who is knowledgeable on predictive analytics can assess the predictive power of the system.

The next typical stage in data analysis is to run multivariate analyses. The simplest of these is a Pearson correlation, showing the linear correlation coefficients among pairs of terms in the data, a prerequisite for good feature selection. The matrix shows the correlations among the variables or factors through a heatmap where the darker the blue cell, the higher the positive correlation with the target variable; the darker the red cell, the higher the negative correlation. Heatmaps are a standard tool to visualize such correlations and allow a quick eyeballing of the data to identify correlations of most interest. As evident in Figure 4, most factors are not highly correlated, indicating that each factor indeed contributes independently to the risk prediction. The

correlation coefficients are listed in the table next to the correlation graph for easy access and verification.



Figure 2. Home page with a simple menu bar and short system description

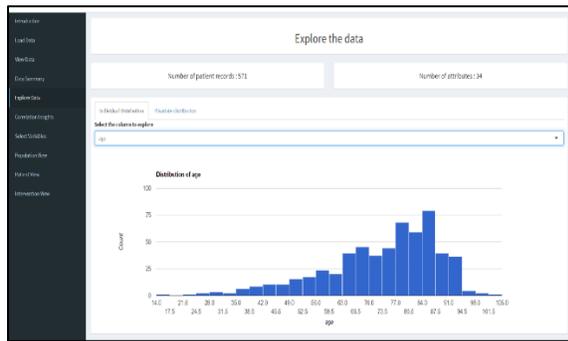


Figure 3. Data exploration

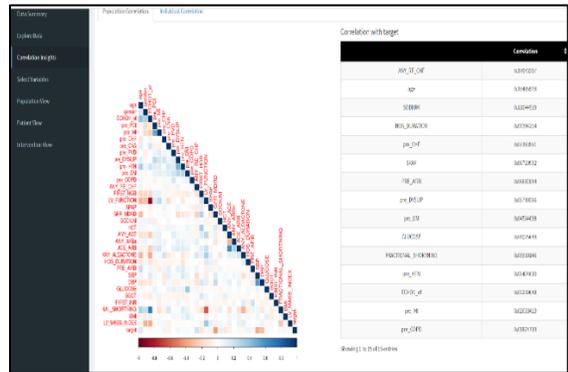


Figure 4. Correlation matrix and heatmap

Once the data is reviewed and feature selection validated, statistical and machine learning methods are used to build the risk assessment models [31]. Figure 5 displays the results of applying these methods to a subset of the Sheba data since the interpretation and explainability value of visualization is reduced when large amounts of data is compressed into a small 2-d display. The figure includes several components to be noted, such as the scatter plot in the center of the graph, the circumscribing circle with the features, the methods associated with the x and y axes, the

boundary line separating the risk dimension, and others.

The circumscribing circle shows the factors/features/variables included in the current analysis, each one positioned by the algorithm in the quadrant where it is most significant. These risk factors, or features, are shown in blue, whereas those that the user selected to monitor are shown in pink. The larger the circle, the more weight that feature has in determining the positive or negative risk level. The position of that feature on the x and the y axes is derived based on the two analytical methods the CDSS applies. The pattern shown in the circumscribing circle in Figure 5 is the combination of two dimensionality-reduction methods: principal component analysis (PCA) and Fisher’s Linear Discriminant Analysis (LDA) [31].

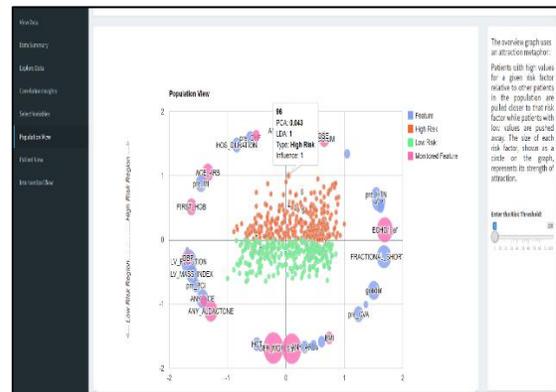


Figure 5. Population view, filter slider on the right hand side, patient details displayed while hovering over an individual patient

A PCA calculates an orthogonal transformation of the original data matrix into principal components so that features in the original data that are highly correlated with each other will be highly associated (and hence correlated) with the same principal component. Each principal component is the sum of all the features with a weight associated with each feature for each principal component. These weights are estimated to maximize the level of explained variance. Each principal component is accordingly a latent variable, or high-level abstraction, that represents the features that combine to create it. Features may have high weights on more than one principal component, suggesting that the same feature may contribute highly, with a larger weight, to more than one latent variable. As a result, a PCA combines the original data points by how close the features

expressed in those data points are to each other on a set of data derived principal components, so that features that have a high weight for the same principal component are treated as if they represent the same latent construct. The PCA algorithm constructs as many principal components as there are features but retains only as many as instructed, allowing data reduction. Data reduction means that only a few abstract latent variables are retained out of the many features that combine to create them. As a result, rather than having many features, it is their shared abstractions, principal components, that are retained. In this analysis, we took only the first principal component. This principal component has the most variance associated with it. The position of the feature on this first principal component is shown graphically as its position on the x axis.

An LDA runs a different process to group the features. LDA is a machine learning algorithm that identifies the best linear combination of independent variables (predictors, in this case these are the features) that predicts the classification value of one dependent variable (in this case, the health risk level). While a PCA optimizes the estimation of as many principal components as instructed without regard to the classification of the values in each component, an LDA is optimized to classify one dependent variable into one or many classes. The LDA values of each feature are shown on the y axis. These are the risk levels. The combination of a PCA and LDA allows the contrasting of two methods: a PCA that seeks to group features by how correlated they are to each other, and an LDA that groups them based on how well they predict an explicit criterion (in this case, risk of death or readmission). This method is based on existing functionality [31].

The red and green dots in the population view (Figure 5) show the scatter plot of risk distribution of actual patients within that data induced pattern of PCA-based x axis and LDA-based y axis values. Red circles are patients with higher risk, green with lower risk. The further the patient is from the x axis, the higher or the lower that risk is. Hovering the cursor over a specific patient circle will open a popup box with some details about that patient. That popup box shows the patient ID in the dataset, PCA and LDA values, a derived textual classification of risk level based on the LDA, and a measure of statistical influence. This population view presents the overall profile of what determines a feature of interest, in this case that is risk level, as a function of other features of

interest, and the position of each data point, in this case that is each patient, in that classification scheme. The slide on the right-hand side of the screen allows the user to filter out the population into a sub-group, based on risk level. Thus, the cardiologist can, for example, filter out low risk patients, and display only the 20% highest-risk patients, for deeper examination. This feature is part of the interactive user interface.

Clicking on a patient circle in Figure 5 opens a new screen with details about that specific patient. The selected patient's features are shown as a black dot in the left pane, and its details in the right pane. The user can then select specific features to monitor for this patient, which are displayed in the right-hand side of the screen in either graphical or table display. The graphical visualization shows the values of the features for this specific patient, and how they relate to the normal range. Thus, in Figure 6, for patient number 96 who is at high risk for mortality within 30 days, the sodium level in the blood is very high, whereas the ejection fraction (ECHO EF) is very low, as are the systolic and diastolic blood pressure (SBP and DBP, respectively).

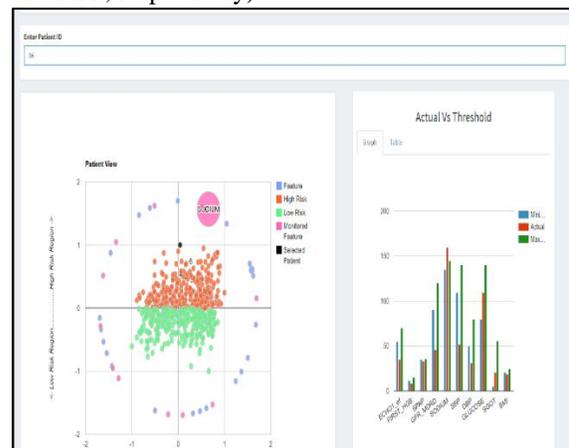


Figure 6. Patient view

The two previous screens allowed an analytics-guided understanding of the population and a specific patient's risk profile as selected by the user. The next stage in the analysis is to show if and how an evidence-based intervention can reduce the patient's risk level. Figure 7 demonstrates this approach graphically. The user can change the values of actionable features to those achieved after medical and/or behavioral intervention. For example, appropriate diet and medication can reduce the sodium level and improve SBP and DBP. When the desired values are reached, the patient's likelihood to die will decrease as shown



drive CDSS adoption, by placing the expert in the decision-making loop alongside the computer, facilitating an intelligence augmenting environment [50, 51]. Thus, this CDSS conceptualization addresses recent calls to harness AI, VA and human-computer collaboration for IA, leading to more effective and efficient healthcare processes and point of care decision making [33, 47]. Such systems are advocated to eventually leverage the huge amount of patient-level healthcare data for precision care and more informed clinical decision making [52]. Moreover, it is believed that the collaborative nature of such systems is likely to alleviate barriers to CDSS adoption due to concerns about loss of autonomy and devaluation of expertise [53]. These hypotheses need to be evaluated by more experts, preferably via randomized experiments in varied clinical settings such as inpatient, ambulatory and emergency care, and with different groups of clinical specialties.

## 6. References

- [1] Belard, A., et al., *Precision diagnosis: a view of the clinical decision support systems (CDSS) landscape through the lens of critical care*. Journal of clinical monitoring and computing, 2017. **31**(2): pp. 261-271.
- [2] Ledley, R.S. and L.B. Lusted, *Reasoning Foundations of Medical Diagnosis*. Science, 1959. **130**(3366): pp. 9-21.
- [3] Sheppard, L.C. and N. Kouchoukos, *Automation of measurements and interventions in the systematic care of postoperative cardiac surgical patients*. Medical instrumentation, 1977. **11**(5): pp. 296-301.
- [4] Wyatt, J. and D. Spiegelhalter. *Field trials of medical decision-aids: potential problems and solutions*. in *Proceedings of the annual symposium on computer application in medical care*. 1991. American Medical Informatics Association.
- [5] Naylor, C.D., *On the prospects for a (deep) learning health care system*. Jama, 2018. **320**(11): pp. 1099-1100.
- [6] Shortliffe, E.H. and M.J. Sepúlveda, *Clinical Decision Support in the Era of Artificial Intelligence*. JAMA, 2018. **320**(21): pp. 2199-2200.
- [7] Castaneda, C., et al., *Clinical decision support systems for improving diagnostic accuracy and achieving precision medicine*. Journal of Clinical Bioinformatics, 2015. **5**(1): pp. Article 4.
- [8] McCoy, A.B., et al., *Clinical decision support for colon and rectal surgery: An overview*. Clinics in colon and rectal surgery, 2013. **26**(01): pp. 023-030.
- [9] Ingraham, A.M., et al., *Association of surgical care improvement project infection-related process measure compliance with risk-adjusted outcomes: implications for quality measurement*. Journal of the American College of Surgeons, 2010. **211**(6): pp. 705-714.
- [10] Vermeulen, K.M., et al., *Cost-effectiveness of an electronic medication ordering system (CPOE/CDSS) in hospitalized patients*. International Journal of Medical Informatics, 2014. **83**(8): pp. 572-580.
- [11] Eminaga, O., et al., *CMDX©-based single source information system for simplified quality management and clinical research in prostate cancer*. BMC Medical Informatics and Decision Making, 2012. **12**(1): pp. Article 141.
- [12] Berry, D.A., et al., *BRCAPRO validation, sensitivity of genetic testing of BRCA1/BRCA2, and prevalence of other breast cancer susceptibility genes*. Journal of Clinical Oncology, 2002. **20**(11): pp. 2701-2712.
- [13] Bates, D.W., et al., *Reducing the Frequency of Errors in Medicine Using Information Technology*. Journal of the American Medical Informatics Association, 2001. **8**(4): pp. 299-308.
- [14] Shojania, K.G., et al., *Effect of point-of-care computer reminders on physician behaviour: a systematic review*. Canadian Medical Association Journal, 2010. **182**(5): pp. E216-E225.
- [15] Kaplan, B., *Evaluating informatics applications—clinical decision support systems literature review*. International Journal of Medical Informatics, 2001. **64**(1): pp. 15-37.
- [16] Lindgren, H. and C. Yan, *Detecting Learning and Reasoning Patterns in a CDSS for Dementia Investigation*, in *Digital Healthcare Empowering Europeans*, Ronald Cornet, et al., Editors. 2015, IOS Press. p. 739-742.
- [17] Liedgren, P., et al., *The Use of Decision Support Systems in Social Work: A Scoping Study Literature Review*. Journal of Evidence-Informed Social Work, 2016. **13**(1): pp. 1-20.
- [18] Kilsdonk, E., L.W. Peute, and M.W.M. Jaspers, *Factors influencing implementation success of guideline-based clinical decision support systems: A systematic review and gaps analysis*. International Journal of Medical Informatics, 2017. **98**: pp. 56-64.
- [19] Brunner, J., et al., *User-centered design to improve clinical decision support in primary care*. International Journal of Medical Informatics, 2017. **104**: pp. 56-64.
- [20] Keim, D., et al., *Visual analytics: Definition, process, and challenges*, in *Information visualization, Lecture Notes in Computer Science, vol 4950* , Kerren A., et al., Editors. 2008, Springer: Berlin, Heidelberg. p. 154-175.
- [21] Hassani, H., et al., *Artificial Intelligence (AI) or Intelligence Augmentation (IA): What Is the Future?* AI, 2020. **1**(2): pp. 143-155.
- [22] Markoff, J., *Machines of Loving Grace: The Quest for Common Ground Between Humans and Robots*. 2015, New York: Ecco Books.
- [23] Skagestad, P., *Thinking with machines: Intelligence augmentation, evolutionary epistemology, and semiotic*. Journal of Social and Evolutionary Systems, 1993. **16**(2): pp. 157-180.
- [24] Carter, S. and M. Nielsen *Using artificial intelligence to augment human intelligence*. Distill, 2017. **2**, <https://distill.pub/2017/aia/?ref=webdesignernews.co>

- [m.](#)
- [25] Jarrahi, M.H., *Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making*. Business Horizons, 2018. **61**(4): pp. 577-586.
- [26] Mann, S., *Wearable computing: toward humanistic intelligence*. IEEE Intelligent Systems, 2001. **16**(3): pp. 10-15.
- [27] Fischer, G., *Distributed intelligence: extending the power of the unaided, individual human mind*, in *Proceedings of the working conference on Advanced visual interfaces*. 2006, Association for Computing Machinery: Venezia, Italy. pp. 7–14.
- [28] Paraense, A.L.O., R.R. Gudwin, and R. de Almeida Goncalves. *Brainmerge: a Semiotic-Oriented Software Development Process for Intelligence Augmentation Systems*. in *2007 International Conference on Integration of Knowledge Intensive Multi-Agent Systems*. 2007. Waltham, MA, USA.
- [29] Endert, A., et al., *The human is the loop: new directions for visual analytics*. Journal of Intelligent Information Systems, 2014. **43**(3): pp. 411-435.
- [30] Athreya, A., et al., *Augmentation of Physician Assessments with Multi-Omics Enhances Predictability of Drug Response: A Case Study of Major Depressive Disorder*. IEEE Computational Intelligence Magazine, 2018. **13**(3): pp. 20-31.
- [31] Harle, C.A., D.B. Neill, and R. Padman, *Information visualization for chronic disease risk assessment*. IEEE Intelligent Systems, 2012. **27**(6): pp. 81-85.
- [32] Zhang, Y. and R. Padman, *An Interactive Platform to Visualize Data-Driven Clinical Pathways for the Management of Multiple Chronic Conditions*. Studies in health technology and informatics, 2017. **245**: pp. 672-676.
- [33] Wu, D.T.Y., et al., *Evaluating visual analytics for health informatics applications: a systematic review from the American Medical Informatics Association Visual Analytics Working Group Task Force on Evaluation*. Journal of the American Medical Informatics Association, 2019. **26**(4): pp. 314-323.
- [34] Thomas, J.J. and K.A. Cook *Illuminating the path: The research and development agenda for visual analytics*. 2005. <https://www.hsdl.org/?view&did=485291>.
- [35] Sacha, D., et al., *Knowledge Generation Model for Visual Analytics*. IEEE Transactions on Visualization and Computer Graphics, 2014. **20**(12): pp. 1604-1613.
- [36] Najjar, A., et al., *A two-step approach for mining patient treatment pathways in administrative healthcare databases*. Artificial Intelligence in Medicine, 2018. **87**: pp. 34-48.
- [37] Wachter, R.M., *Making IT Work: Harnessing the Power of Health Information Technology to Improve Care in England. Report of the National Advisory Group on Health Information Technology in England*. 2016, the National Advisory Group on Health Information Technology in England: England. pp. 1-69.
- [38] Gotz, D. and D. Borland, *Data-Driven Healthcare: Challenges and Opportunities for Interactive Visualization*. IEEE Computer Graphics and Applications, 2016. **36**(3): pp. 90-96.
- [39] Mane, K.K., et al., *VisualDecisionLinc: A visual analytics approach for comparative effectiveness-based clinical decision support in psychiatry*. Journal of Biomedical Informatics, 2012. **45**(1): pp. 101-106.
- [40] Ratwani, R.M. and A. Fong, *'Connecting the dots': leveraging visual analytics to make sense of patient safety event reports*. Journal of the American Medical Informatics Association, 2015. **22**(2): pp. 312-317.
- [41] Miller, N., *The 10 Best Hospitals in the World*, in *Newsweek*. 2019, Newsweek Magazine: USA.
- [42] Landerman, L.R., K.C. Land, and C.F. Pieper, *An empirical evaluation of the predictive mean matching method for imputing missing values*. Sociological Methods & Research, 1997. **26**(1): pp. 3-33.
- [43] van Buuren, S., et al., *Package 'mice'*. Computer software, 2015.
- [44] Azhagusundari, B. and A.S. Thanamani, *Feature selection based on information gain*. International Journal of Innovative Technology and Exploring Engineering (IJITEE), 2013. **2**(2): pp. 18-21.
- [45] McHugh, M.L., *Interrater reliability: the kappa statistic*. Biochemia medica, 2012. **22**(3): pp. 276-282.
- [46] Varonen, H., et al., *What may help or hinder the implementation of computerized decision support systems (CDSSs): a focus group study with physicians*. Family Practice, 2008. **25**(3): pp. 162-167.
- [47] Buchlak, Q.D., et al., *Machine learning applications to clinical decision support in neurosurgery: an artificial intelligence augmented systematic review*. Neurosurg Rev, 2019. **August**: pp. 1-19.
- [48] Venkatesh, V., et al., *User Acceptance of Information Technology: Toward a Unified View*. MIS Quarterly, 2003. **27**(3): pp. 425-478.
- [49] Sjöborg, B., et al., *Design and implementation of a point-of-care computerized system for drug therapy in Stockholm metropolitan health region—Bridging the gap between knowledge and practice*. International Journal of Medical Informatics, 2007. **76**(7): pp. 497-506.
- [50] Grønsund, T. and M. Aanestad, *Augmenting the algorithm: Emerging human-in-the-loop work configurations*. The Journal of Strategic Information Systems, 2020. **29**(2): pp. 1-19.
- [51] Pedersen, M., et al., *Artificial intelligence for clinical decision support in neurology*. Brain Communications, 2020. **2**(2): pp. 1-11.
- [52] Moxey, A., et al., *Computerized clinical decision support for prescribing: provision does not guarantee uptake*. Journal of the American Medical Informatics Association, 2010. **17**(1): pp. 25-33.
- [53] Sittig, D.F., et al., *Grand challenges in clinical decision support*. Journal of Biomedical Informatics, 2008. **41**(2): pp. 387-392.