Personalized Discovery and Treatment Recommendation

Jinsung Yoon
UCLA

Mihaela van der Schaar
UCLA

Abstract

In the era of precision medicine, treatment regimens are selected based on the patients individual characteristics to improve the response. However, given the high dimensionality of clinical data, discovering which data is relevant to consider when selecting treatments can be challenging. We present two novel approaches that discover the information (patient features) from the electronic health record that is most relevant for predicting the best treatment regimen for each specific patient. We describe two approaches: the first assumes that features are independent (DE-IF), the second allows for features to be correlated (DE-CF). We evaluate the performance of these approaches against other methods using a data set of 50,000 cases. Different patient features are relevant to predicting the success of different treatments; DE-IF and DE-CF can discover these different relevant features and use them to make personalized treatment recommendations. We determine rewards based on the results of external knowledge (published literature and clinical practice guidelines) describing the effect of the same treatments on similar patients as the ones considered in our study. In general, our approach achieves a 16.6% improvement over other methods in terms of kappa statistics. We also demonstrate that the performance of our methods is robust against missing information.

1 INTRODUCTION

A key aspect of precision medicine is the ability to personalize the selection of a treatment regimen based on the personal characteristic of an individual patient (Richmond, 2008). The wealth of information being routinely collected as part of the electronic health record (EHR) provides an unprecedented opportunity to discover appropriate treatments for patients given historical information about the treatments administered to similar patients and their outcomes (Shortliffe and Cimino, 2006). However, using this information is difficult precisely because there is too much of it; what is needed is to extract the information that is actually relevant for the particular patient and treatment among the wealth of available information.

In this paper, we present a novel method called Discovery Engine (DE) for optimizing treatment selection by identifying the features in the patient record that differentiate the individuals who receive a treatment and positively respond to it versus those who do not. We describe two variants of this approach: a simplified method that assumes all features are independent (DE-IF) and one that takes into account feature correlations (DE-CF). Our approaches leverage available contextual information about patients and enable learning from the large quantities of observational clinical data to inform treatment recommendations and make better decisions by learning from similar patients. We show that our approaches consistently perform better than existing approaches in matching individual patients to relevant treatment regimens.

One of the biggest challenges faced by this research or any other predictive algorithms is that the rewards/true outcomes of treatment (e.g. five-year progression free survival) are not available in the data we have and indeed, are usually not available in most data (Bennett and Hauser, 2012). Moreover, even if rewards/true outcomes of treatment were available, the counterfactuals rewards/outcomes of alternative treatments that were not used are never available. What is available is a large medical literature that reports the results of a wide range of clinical references including different types of patients, different types of treatments, and the outcomes of these treatments. We use the results of some of these references to construct a transfer reward, which we use as a proxy for rewards.

Preliminary work. Under review by AISTATS 2016. Do not distribute.

This allows us to train the DE algorithms and also to assess their performance in comparison to existing methods. The three primary contributions of this paper are:

- We describe a method for discovering the relevant information from the EHR that distinguishes between patients that should receive one particular chemotherapy and the patients who should receive another. For instance, patients with observed mutations (e.g., EGFR mutation and HER-2 overexpression) are more likely to respond to a specific type of chemotherapy (e.g., erlotinib) (Zhou et al. 2011).
- Using the retrospective cases in the EHR, our approach discovers the optimal treatment based on the available information (e.g. their clinical findings, treatment history, and outcomes).
- In lieu of having actual reward values associated with treatment decisions, we define the transfer rewards, a method for ranking chemotherapies described in external knowledge based on their relevance to individual patients based on reported characteristics.

2 RELATED WORKS

Current medical practice relies on manually curated systematic reviews and clinical guidelines that provide treatment recommendations for large groups of patients rather than personalized treatments that are tailored to individual patients. Clinical decision support systems have been proposed before, but many of them do not consider the specific characteristics of patients and do not provide personalized treatment recommendations; hence, they are not very accurate and have only limited applicability in practice (Tsumoto, 1998, Greenspan and Syeda-Mahmood, 2012, and Xu et al. 2014). Moreover, clinicians often refer to the medical literature available through Medline/PubMed, VisualDX, and UpToDate to help them associate observed finding with possible conditions and recommended actions; but, these resources are not customized to a specific patients case.

Another strand of literature related to this work is that on machine learning techniques, including Support Vector Machines (SVMs), AdaBoost, logistic regression etc. are not able to accurately capture the nuanced relationships between specific patient characteristics and specific treatments, thereby not being able to issue accurate personalized treatment recommendations. (Comparisons against these methods are performed in the experiments section (Peterson, Doughty, and Kann, 2013).) One of the reasons for inaccurate recommendation is that those techniques cannot learn the relevant features of patients that make them respond to specific treatments.

Our method also exhibits similarities to the contextual multi-armed bandit problem (MAB). However, contextual MABs are very inefficient when the number of contexts (in our case patient characteristics) is large (Tekin and van der Schaar, 2015). DE is able to successfully deal with the curse of dimensionality by discovering what information is relevant and making decisions based on relevant contexts rather than the entire EHR, much of which is irrelevant to the decision of whether to administer a specific treatment or not.

3 PROBLEM FORMULATION

In this section, we introduce our method - discovery engine (DE). Figure 1 depicts the proposed system, which issues a treatment recommendation to the physician that is personalized to the patient. While the proposed system is applicable in general, we illustrate its use and performance in the context of breast cancer. Let $\mathbf{x} = (x_1, x_2, ..., x_D)$ denote the patient information where D is the total number of features of each patient such as age, tumor size, estrogen receptor information etc.; $a \in A \equiv \{a_1, a_2, ..., a_K\}$ denotes the action (i.e., treatment) that is performed on the patient. Each feature is denoted as $f \in F \equiv \{f_1, f_2, ..., f_D\}$. The reward y would ideally be derived as the probability of five-year survival of a patient given the chemotherapy regimen. However, obtaining this reward value is difficult in practice; instead, we use a surrogate reward measure which is discussed in a subsequent section. Let $\mathbf{x}(n), a(n), y(n)$ be the patient information, action and reward of n^{th} patient and $H_N = \{\mathbf{x}(n), a(n), y(n)\}_{n=1}^{N}$ be the information available for the N previously seen patients. (This represents the training set.)

The success of a treatment a does not depend on all the features (Goldstein et al. 2008): we assume that the success of a treatment a depends only on a subset of features $R(a) \subseteq F$ which we call the relevant features and let $R = \bigcup_{a \in A} R(a)$ be the set of all relevant features. A key challenge is that the feature that are relevant for predicting the success of a treatment are not known a priori; they need to be discovered/learned. Note that such learning/discovery is very different from existing feature selection algorithms which focus on the patients characteristics and not on how these characteristics differently impact the success of different treatments: our approach is capable of discovering different features that are relevant to different treatments. We say that $R(a)$ is rele-

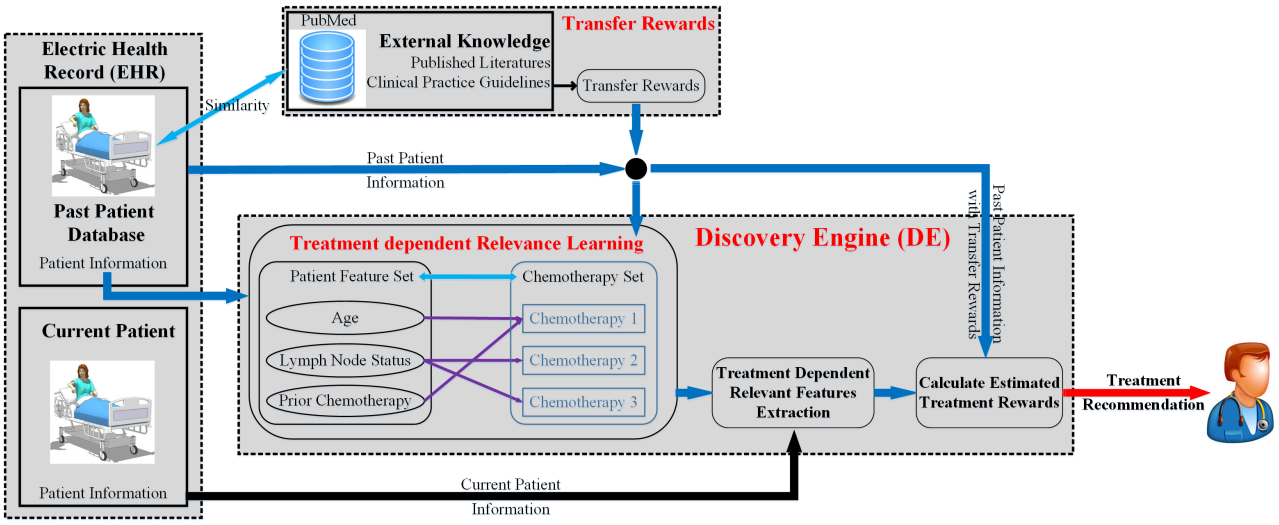

Figure 1: Personalized treatment recommendation system using DE

vant/informative for treatment a if the expected reward only depends on the information contained in $R(a)$.

Figure 1 illustrates the aforementioned formulation using an example, which discovers $R(Chemotherapy1)$={Age, Prior Chemotherapy}, i.e. the patients response to Chemotherapy 1 depends on the patient age and the prior chemotherapy but not on the lymph node status.

Our goal is to discover the relevant features of each treatment a (this may be different for each treatment) and determine the optimal treatment that corresponds to the information contained within the relevant features. The optimal recommended treatment is given by,

$$a^*(\mathbf{x}_R) \equiv arg\max_a E_{(y|a,\mathbf{x}_{R(a)})}(y|a,\mathbf{x}_{R(a)})$$

where $a^*(\mathbf{x}_R)$ is the treatment that yields the highest reward (e.g., best patient outcome) for a patient characterized by the relevant features (age, prior chemotherapy) $\mathbf{x}_R$.

4 DE ALGORITHM

4.1 Independent Features (DE-IF)

We begin by describing our approach under the simplifying assumption that the features are independent. The main idea of this algorithm is to find the relevant features for a specific treatment that distinguish the expected rewards for patients with specific patient information (i.e., age 50+ patients) and those obtained for the entire population. Features that create a large reward separation are considered relevant to that specific treatment. Then, DE-IF simultaneously learns the relevant features that predict the success of a treatment and the best treatment given a patients characteristic. The key steps of DE-IF are outlined below:

Step 1: Define $\hat{y}_a^S(\mathbf{x}_S)$ and $N_a^S(\mathbf{x}_S)$ as the sample mean reward estimator and the number of patients whose information contains $\mathbf{x}_S$ for treatment a. Intuitively, the sample mean reward estimator can be thought of as the average reward of treatment a for the patients with information $\mathbf{x}_S$ (such as age= 50+, tumor size= $4mm$ patients). Then, let $\hat{y}_a$ be the sample mean reward estimator for treatment a and N_a be the number of patients to whom treatment a was recommended.

Step 2: Using the reward estimates for all (a, f), the algorithm computes a relevance metric $h_f(a)$ for each treatment feature pair (a, f), i.e.,

$$h_f(a) \equiv \sum_{x_f} \frac{N_a^f(x_f)}{N_a} |\hat{y}_a^f(x_f) - \hat{y}_a|$$

that measures the weighted variation of the reward estimates. Recall that f was defined earlier as a feature itself.

Step 3: Using the relevance metric for each pair (a, f), we discover the relevant feature $(\hat{r}(a))$ for each treatment a.

$$\hat{r}(a) = arg\max_f h_f(a)$$

Define $m(a)$ as the number of relevant features and $\hat{R}(a)$ as the set of the relevant features for treatment a. We sequentially discover $m(a)$ number of relevant features $(\hat{r}(a))$ for each treatment a

Initialize: $\hat{R}(a) = \emptyset, F(a) = \{f_1, ..., f_D\}$ for each a
for each treatment-feature pair (a, f) **do**
  Compute $h_f(a) = \sum_{x_f} \frac{N_a^f(x_f)}{N_a} |\hat{y}_a^f(x_f) - \hat{y}_a|$
end for
for each treatment a **do**
  do
  $\hat{r}(a) = arg\max_{f \in F(a)} h_f(a)$
  $\hat{R}(a) = \hat{R}(a) \bigcup \hat{r}(a)$
  $F(a) = F(a) \backslash \hat{r}(a)$
  while $(|\hat{R}(a)| < m(a))$
end for
for any patient information $\mathbf{x}_{\hat{R}}$ **do**
  Compute $\hat{a}(\mathbf{x}_{\hat{R}}) = arg\max_a \hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)})$
end for

Figure 2: Psuedocode of DE-IF

Step 4: Define $\hat{R} = \bigcup_{a \in A} \hat{R}(a)$ as the relevant features for all treatments. The optimal treatment with respect to relevant features $\hat{R}$ is determined as

$$\hat{a}(\mathbf{x}_{\hat{R}}) = arg\max_a \hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)})$$

This optimization selects the treatment with the highest estimated reward given a patient with patient information $\mathbf{x}_{\hat{R}}$. The pseudo-code is given in Figure 2.

4.2 Correlated Features (DE-CF)

We now describe our approach for the more general case in which features are allowed to be correlated. For example, for some cancers, tumor size and tumor grade may be correlated, not independent. Therefore, we provide a generalization of the DE algorithm called DE-CF that properly treats correlated features. While the basic steps are similar to DE-IF, DE-CF uses a different relevance metric:

$$h_S(a) \equiv \sum_{\mathbf{x}_S} \frac{N_a^S(\mathbf{x}_S)}{N_a} |\hat{y}_a^S(\mathbf{x}_S) - \hat{y}_a|$$

Then, the most relevant tuple (consisting of $m(a)$ features) associated with treatment a can be found as

$$\hat{R}(a) = arg\max_{S:|S|=m(a)} h_S(a)$$

The pseudo-code of DE-CF is given in Figure 3.

4.3 Contrasting DE-IF and DE-CF

DE-CF suffers from the curse of dimensionality because it needs to estimate rewards for each tuple of features and hence it needs to see many tuples, not just many features. For instance, it needs to see many cases with age 50+ **and** tumor size 4cm not just many cases with age 50+ **or** tumor size 4cm. Hence, if the data set

Initialize: $\hat{R}(a) = \emptyset$ for each a
for each treatment-feature set (a, S) **do**
  Compute $h_S(a) \equiv \sum_{\mathbf{x}_S} \frac{N_a^S(\mathbf{x}_S)}{N_a} |\hat{y}_a^S(\mathbf{x}_S) - \hat{y}_a|$
end for
for each treatment a **do**
  Compute $\hat{R}(a) = arg\max_{S:|S|=m(a)} h_S(a)$
end for
for any patient information $\mathbf{x}_{\hat{R}}$ **do**
  Compute $\hat{a}(\mathbf{x}_{\hat{R}}) = arg\max_a \hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)})$
end for

Figure 3: Psuedocode of DE-CF

is small, DE-CF may not yield high confidence. For this reason, we suggest that for small data sets, it is likely better to make the perhaps incorrect assumption of independent features and use DE-IF rather than DE-CF.

4.4 DE with Missing Information

Electronic health records are often missing hence, DE must be able to operate efficiently even when information is missing. When information is missing, the feature information vector $\mathbf{x}$ can be divided into two components: the available features ($\mathbf{x}^{av}$) and the missing features ($\mathbf{x}^m$). Thus, $\mathbf{x} = \{\mathbf{x}^{av}, \mathbf{x}^m\}$. When features are missing, the relevance metric of DE-IF is solely computed based on the available information:

$$h_f(a) \equiv \sum_{x_f^{av}} \frac{N_a^f(x_f^{av})}{N_a} |\hat{y}_a^f(x_f^{av}) - \hat{y}_a|$$

If the feature f is frequently missing, $h_f(a)$ decreases, and as a result the feature f is rarely selected as a relevant feature. The relevance metric of DE-CF is similarly modified. It should be also note that the estimated reward ($\hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)}^{av})$) given a patients available relevant information, $\mathbf{x}_{\hat{R}(a)}^{av}$, for each treatment a is computed as:

$$\hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)}^{av}) = \mathbf{E}(\hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)}) | \mathbf{x}_{\hat{R}(a)}^{av})$$

$$= \sum_{\mathbf{x}_{\hat{R}(a)}^m} P(\mathbf{x}_{\hat{R}(a)}^m | \mathbf{x}_{\hat{R}(a)}^{av}) \cdot \hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)}^{av}, \mathbf{x}_{\hat{R}(a)}^m)$$

$$= \sum_{\mathbf{x}_{\hat{R}(a)}^m} P(\mathbf{x}_{\hat{R}(a)}^m | \mathbf{x}_{\hat{R}(a)}^{av}) \cdot \hat{y}_a^{\hat{R}(a)}(\mathbf{x}_{\hat{R}(a)})$$

We can estimate the conditional probability, $P(x_{\hat{R}(a)}^m | x_{\hat{R}(a)}^{av})$, based on the probability distribution of the features in training set.

For example, let us assume that there are two relevant features (menopausal status and tumor stage) for the treatment a. Suppose now that a patients tumor stage

is G1 and her menopausal status is missing. Then to compute the estimated reward for selecting a treatment for that patient, we use the conditional probability of being premenopausal given that the patient has a G1 tumor stage. Let us assume that this is 0.7 (as computed based on the patients in the training set), then DE estimates the reward of treatment a for this patient as:

$$\hat{y}_a^{\hat{R}(a)}(G1) = 0.7 \cdot \hat{y}_a^{\hat{R}(a)}(Pre, G1) + 0.3 \cdot \hat{y}_a^{\hat{R}(a)}(Post, G1)$$

5 TRANSFER REWARDS: ESTIMATING REWARDS BASED ON MEDICAL LITERATURE

As we have said in the introduction, the most natural rewards would be patient outcomes (survival rates) but these outcomes are very difficult to obtain in practice. Instead, we use a proxy for outcomes based on external knowledge which consists of published literature and clinical practice guidelines. Let us refer to all external knowledge as the term *reference* for the rest of the paper from here on. The idea is to match patients to relevant *reference* given patient characteristics. For each patient (represented by each features) and each *reference*, we define the amount of information that *reference* provides about that patient; we then aggregate this information across all *references*. From this information, we compute the posterior probability that the treatment used in a given *reference* is the best treatment for that patient. This defines the transfer rewards, which we take as the reward for that treatment when applied to that patient

To compute the transfer rewards, we first find relevant patient features for each reference; we do this using a modified feature selection algorithm based on the mutual information (Peng, Long, and Ding, 2005). The definition of the mutual information between the i^{th} *reference* (E_i) and the k^{th} feature (f_k) is:

$$I(f_k; E_i) = \sum_{x \in \chi_k} P(x|E_i) \log \frac{P(x|E_i)}{P(x)}$$

where $P(x|E_i)$ is the probability of feature x in *reference* i, $P(x)$ is the probability of feature x across the entire set of *references*, and χ_k is context space of f_k.

The second step is to compute the transfer rewards as a posterior probability. Given the n^{th} patient, characterized by the feature vector $x_n = \{x_1(n), ..., x_D(n)\}$, we compute the posterior probability that the treatment in *reference* i would be successful when applied to a patient with these features; we abuse notation by writing this as, $P(E_i|X_1 = x_1(n), ..., X_D = x_D(n))$. We compute this posterior probability via Bayes rule; it is computationally convenient to take logarithms:

$$\begin{aligned}
&\log\left(P(E_i|X_1 = x_1(n), ..., X_D = x_D(n)\right) \\
&= \log \tfrac{(P(X_1=x_1(n),...,X_D=x_D(n)|E_i)}{P(X_1=x_1(n),...,X_D=x_D(n))} \cdot P(E_i)) \\
&\cong \log P(E_i) + \textstyle\sum_{l \in Rel} \log \frac{P(X_l=x_l(n)|E_i)}{P(X_l=x_l(n))}
\end{aligned}$$

where with similar abuse of notation we write $P(E_i)$ as the probability of selecting the i^{th} *reference* as the best treatment for the entire population.

We computed all the transfer rewards for each *reference*/treatment. These transfer rewards provide a complete ranking of each treatment for each patient (characterized by the relevance features). The treatment with the highest transfer reward is the best treatment for the given patient.

6 EXPERIMENTS

6.1 Dataset

We perform experiments on a de-identified data set of 50,000 patient cases which underwent screening and diagnostic testing at a large academic medical center. The patient data is characterized by 15 features summarized in Table 1.

Table 1: Summary of patient information type

Features	**Range**
Age	30s - 60+
Menopausal	Pos/Neg
Race	White/Black/Other
Estrogen Receptor	Pos/Neg
Progesteron Receptor	Pos/Neg
HER2NEU	Pos/Neg/Neu
Tumor Stage	T1 - T4
Tumor Grade	G1 - G3
PLNC	0 - 10+
Lymph Node Status	Pos/Neg
WHO Score	0 - 5
Surgery Type	BCT/MRM/No
Prior Radiotherapy	Experience / No
Prior Chemotherapy	Experience/No
Histology	Ductal/Mix/Lobular

From an initial set of 2,353 references (performing a narrow search of therapies using PubMed Clinical Queries), six chemotherapies listed in Table 2, the standard chemotherapy regimen for breast cancer (Fleeger et al. 2015, Hamel and Johnson, 2015, and Levine, 2001), described in 32 references were selected for further analysis. The sample size of reported references ranged from 50 to 3,934 with a total population of 27,118 individuals.

We utilize a training set of 4,000 patients. The remaining patients provide the test set. Standard 10-fold stratified cross-validation was applied, and therefore, no training data were used during testing of the model, but 10 different models were used to derive the final test results.

Table 2: Code for each chemotherapy regimen

Code	Specific Chemotherapy Regimen
AC	Doxorubicin + Cyclophosphamide
ACT	Doxorubicin + Cyclophosphamide + Taxanes
AT	Doxorubicin + Taxanes
CAF	Cyclophosphamide + Doxorubicin + 5-Fluorouracil
CEF	Cyclophosphamide + Epirubicin + 5-Fluorouracil
CMF	Cyclophosphamide + Methotrexate + 5-Fluorouracil

6.2 Benchmarks

We compare the performance of DE-IF with six existing popular machine learning algorithms described below:

- Correlation Feature Selection (CFS): a well-known feature selection algorithm (Hall, 1998);
- All Contextual Learning (ACL): a well-known contextual learning algorithm which uses all features. This is a modified offline version of the contextual bandit algorithm of Slivkins. (2014);
- Logistic Regression;
- Linear Regression;
- Support Vector Machines (SVM); we use a radial basis function (RBF) kernel SVM;
- Adaptive Boosting (AdaBoost);

6.3 Measuring Success

Given a patient, our algorithm or any of these other algorithms recommend a course of treatment corresponding to a particular *reference*. If a patient has exactly the same feature information as the patient in that *reference*, we regard the algorithm in question as making the correct recommendation for that patient; i.e. it has recommended the best course of treatment. (Notice that the best course of treatment may not promise a good outcome: some cancers are not treatable.) We take the fraction/percentage of correct recommendations to be the success rate for the algorithm in question.

Given the success rate for the algorithm, we apply two performance metrics: simple percent agreement (p_0) and Cohens kappa coefficient (κ). Simple percent agreement (p_o) is the success rate (the fraction of times the personalized treatment prediction coincides with the recommendation provided in the medical literature for the patients with the same characteristics). Cohens kappa coefficient (κ) is a metric which measures inter-rater agreement. It is usually considered a more robust measure than a simple percent agreement, because κ measures the improvement over chance agreements. If p_e is the probability of agreement by chance, then, $\kappa = \frac{p_o - p_e}{1 - p_e}$.

As the bar graphs in Figure 4 show, the first chemotherapy recommendation of DE-IF is successful (as defined above) 73.4% of the time and one of the first two recommendations is successful 88.4% of the time. This is 7.7% better than the next best approach (i.e., SVM) in terms of selecting the optimal chemotherapy on its first choice, and 5.6% better in terms of matching the optimal chemotherapy within the first 2 choices. This is already a significant improvement. In terms of kappa statistics, the improvement is even greater: our algorithm works 16.6% better than SVM. This is because SVM indiscriminately recommends the popular chemotherapies and is not robust when classifying the less popular chemotherapies. Given robustness considerations, which are essential in medical treatment recommendations, kappa statistics is more often used as a performance metric in medical informatics.

It is especially useful to compare our algorithm with other algorithms that rely on feature selection. Again, note that other algorithms use feature selection, but they do not select relevant features for specific treatments and it is by doing this that our algorithm achieves improvement. CFS achieves only a 48% of simple percent agreement because it cannot use the efficacy of the treatment to discriminate the relevant features and hence the technique is entirely unsupervised. ACL succumbs to the curse of dimensionality because there are 15 features, resulting in over 500 million combinations to test. Logistic regression, linear regression, and SVM perform worse than DE-IF because they do not consider the relevant patient information for selecting treatments.

DE-CF performs 9.4% worse than DE-IF in terms of simple percent agreement. This is not surprising: 4,000 cases are simply not enough for DE-CF to estimate reward of feature tuples. The 95% confidence intervals of DE-IF and DE-CF are 4.4% and 5.1% re-

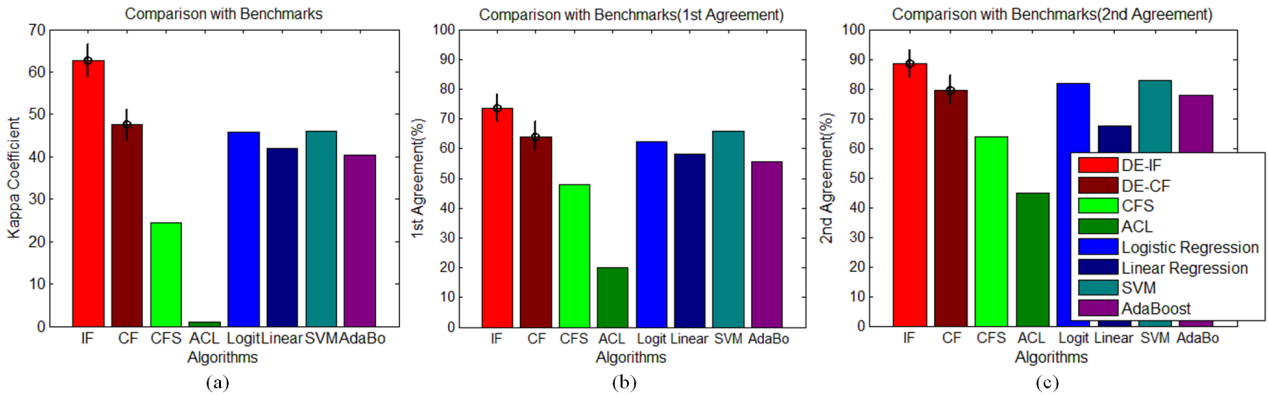

Figure 4: Performance analysis with benchmark algorithms (a) Kappa statistics, (b), (c) Simple percent agreement (1st and 2nd)

spectively. Note also that all the other algorithms with the exception of DE-IF and DE-CF cannot issue such confidence estimates.

Because a large academic medical center may be able to obtain a sufficient cohort to train treatment recommendation algorithms such as the one presented, it is often informative to know the amount of cases needed to ensure that the near-optimal treatment is recommended with a high probability. Our simulation results (Table 3) show that: DE-IF requires 4,000 cases and DE-CF requires 40,000 cases to reach near-optimal treatment recommendation. After 35,000 cases, DE-CF works better than DE-IF because DE-CF has more practical assumption for patients features and 35,000 cases are enough to discover relevance in this application.

Table 3: Simple percent agreement analysis with various number of cases

	Past Patients Number ($\times 10^3$)					
%	**2**	**4**	**10**	**30**	**40**	**45**
DE-IF	**67.4**	**73.4**	**73.0**	**73.3**	73.3	73.4
DE-CF	57.9	64.0	69.4	72.7	**74.5**	**74.9**
AdaBo	54.3	55.4	56.3	57.9	56.9	57.1

6.4 Relevant Features for Each Chemotherapy

Table 4 shows the top 4 ranked relevant features discovered by DE-IF - tumor stage, positive axillary lymph node number (PLNC), estrogen receptor etc.- for recommending AC, ACT, AT, CAF, CEF and CMF chemotherapy. As it can be seen from Table 4, DE-IF is able to discover the different relevant features that are relevant for different chemotherapy.

Importantly, the relevancy of the features discovered by DE is confirmed by clinical studies. First, note that the six considered chemotherapies are commonly recommended to node positive breast cancer patients, i.e. patients where cancer has been found in the lymph nodes (Hamel, and Johnson, 2015). PLNC is the most relevant feature to differentiate among node positive or node negative breast cancer. Hence, PLNC is correctly identified by DE to be relevant. Second, the menopause status is considered as an important factor because medications affect cancer differently for premenopausal and postmenopausal women (Mark, 2001). More specifically, the CEF chemotherapy is only recommended to premenopausal women. Although the menopausal status is not included in this relevant feature set, women over the age of 50 are usually considered as postmenopausal (Mark, 2001). Therefore, age was correctly identified by DE to be a discriminative feature for select-

Table 4: Discovered relevant features for each chemotherapy

Chemo	Relevant Feature1	Relevant Feature2	Relevant Feature3	Relevant Feature4
AC	PLNC	Tumor Stage	Estrogen Receptor	Age
ACT	Tumor Stage	Prior Chemo	PLNC	Estrogen Receptor
AT	Prior Chemo	PLNC	Surgery	Age
CAF	Surgery	Tumor Stage	Age	Tumor Grade
CEF	PLNC	Estrogen Receptor	Tumor Stage	Age
CMF	Estrogen Receptor	PLNC	Radio therapy	Tumor Stage

ing among these chemotherapies. Third, tumor stage is another important feature to discriminate among chemotherapy treatments as confirmed in (Fleeger et al. 2015). Chemicals A(Doxorubicin), T(Taxotere), E(Epirubicin) are highly recommended to advanced breast cancer and our top five chemotherapies include more than one of these chemicals. Therefore, DE has correctly discovered that the features that are relevant for these chemotherapies contain tumor stage information. Finally, chemical T(Taxotere) is usually recommended to patients having breast cancer which do not react to the current chemotherapy. Thus, the prior chemotherapy information is indeed correctly discovered by DE to be relevant for AT and ACT therapies.

6.5 Performance When Patient Information is Missing

As mentioned before, information is often missing from a patients EHR. Moreover, studies have shown that the missing information is often not random (Botsis et al. 2010). For example, the age of the patient is easy to record and is often verified so it is typically neither missing nor incorrect. However, her2neu may not be recorded depending on the diagnostic tests ordered and capabilities of a clinic. Figure 3 shows the performance degradation of DE and benchmark algorithms as a function of the average *degree of incompleteness* (Botsis et al. 2010). (We did not use the percentage of missing features as a metric as the features are not randomly missing. Which features are missing features and the percentages of missing features are described in the supplementary materials.) Figure 5 shows that the performance of DE-IF degrades from 73.4% to 63.0% (when the average *degree of incompleteness* is 50%. However, DE-IF continues to outperform the other methods. DE discovers relevant features with low missing probability, and is able to estimate the missing feature information based on the available feature information, so the impact of missing information is minimized. In fact, DE does better than most other algorithms even when DE misses significant amounts of information from the EHR while the other algorithms make their decision without missing information. Hence DE is robust to missing information.

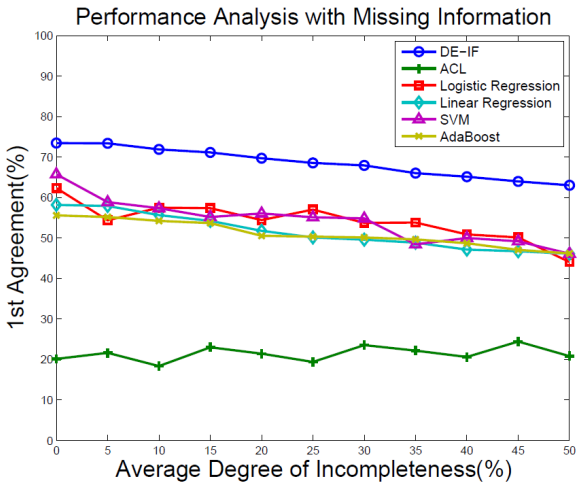

Figure 5: Performance analysis with missing information

7 CONCLUSION AND FUTURE WORK

We describe a novel approach that discovers the relevant information from the EHR that is important in determining which chemotherapy to assign to a patient and use this information to provide personalized treatment recommendations to the physician. Our results demonstrate that DE is capable of outperforming existing machine learning and prediction techniques. We also show that our method is robust against missing information.

Future work will consider that various chemotherapy treatments contain similar chemicals, and will extend DE to incorporate the correlations of the treatments to transfer the knowledge of one to another treatment. In addition, variables such as the tumor size and PLNC number may change over time. These changes may influence the duration of a therapy and the selection of future therapies. Currently, we only consider a single action, identifying what information is important in deciding on a single treatment. Another DE extension will consider the global sequence of treatment decisions that optimize long-term outcomes (e.g., overall survival).

In conclusion, we believe that our proposed contextual learning approach demonstrates promise towards providing personalized treatment recommendations and leveraging the external knowledge to provide reward estimation. As new therapies are evaluated and approved for use, clinicians will have an increasingly difficult time determining which treatment is most effective for an individual patient. DE provides a step towards providing computational methods for personalized treatment recommendation.

References

M. E. Trudeau (2001). Optimizing adjuvant breast cancer chemo-therapy: rationale for the MA. 21 study. *Oncology* 15(5 Suppl 7), 7-13.

L. J. Goldstein, A. O'Neill, J. A. Sparano, E. A. Perez, L. N. Shulman, S. Martino, N. E. Davidson (2008). Concurrent doxoru-bicin plus docetaxel is not more effective than concurrent doxoru-bicin plus cyclophosphamide in operable breast cancer with 0 to 3 positive axillary nodes: North American Breast Cancer Intergroup Trial E 2197. *Journal of Clinical Oncology* 26(25), 4092-4099.

A. Badanidiyuru, J. Langford, A. Slivkins (2014). Resourceful contextual bandits. *arXiv preprint arXiv*:1402.6779.

H. Peng, F. Long, C. Ding (2005). Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(8), 1226-1238.

K. S. Albain (2004). Adjuvant Chemotherapy for Lymph NodeNegative, Estrogen ReceptorNegative Breast Cancer: A Tale of Three Trials. *Journal of the National Cancer Institute*, 96(24), 1801-1804.

M. Levine (2001). Clinical practice guidelines for the care and treatment of breast cancer: adjuvant systemic therapy for node-positive breast cancer (summary of the 2001 update). *Canadian Medical Association Journal*, 164(5), 644-646.

C. Tekin,M. van der Schaar (2015). RELEAF: An algorithm for learning and exploiting relevance. *IEEE journal of Selected Topics in Signal Processing*, 9(4), 716-727

T. D. Richmond (2008). The current status and future potential of personalized diagnostics: streamlining a customized process. *Bio-technology annual review*, 14, 411-422.

A. X. Garg, N. K. Adhikari, H. McDonald, M. P. Rosas-Arellano, P. J. Devereaux, J. Beyene, R. B Haynes (2005). Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *Jama*, 293(10), 1223-1238.

E. H. Shortliffe, J. J Cimino (2013). Biomedical informat-ics: computer applications in health care and biomedicine. *Springer Science and Business Media*.

C. C. Bennett, K. Hauser (2013). Artificial intelligence framework for simulating clinical decision-making: A Markov decision process approach. *Artificial intelligence in medicine*, 57(1), 9-19.

H. G. H. Mller, T. Syeda-Mahmood (2013). Medical Content-Based Retrieval for Clinical Decision Support.

S. Tsumoto (1998). Automated extraction of medical expert system rules from clinical databases based on rough set theory. *Information sciences*, 112(1), 67-84.

A. Hyvrinen, E. Oja, (2000). Independent component analysis: algorithms and applications. *Neural networks*, 13(4), 411-430.

T. A. Peterson, E. Doughty, M. G. Kann (2013). Towards precision medicine: advances in computational approaches for the analysis of human variants. *Journal of molecular biology*, 425(21), 4047-4063.

C. Zhou, Y. L. Wu, G. Chen, J. Feng, X. Q. Liu, C. Wang, C. You (2011). Erlotinib versus chemotherapy as first-line treatment for patients with advanced EGFR mutation-positive non-small-cell lung cancer (OPTIMAL, CTONG-0802): a multicentre, open-label, randomised, phase 3 study. *The lancet oncology*, 12(8), 735-742.

B. Kaplan (2001). Evaluating informatics applicationsclinical decision support systems literature review. *International journal of medical informatics*, 64(1), 15-37.

J. Xu, D. Sow, D. Turaga, M. van der Schaar (2014). Online Transfer Learning for Differential Diagnosis Determination. *AAAI Workshop on the World Wide Web and Public Health Intelligence*

A. Slivkins (2014). Contextual bandits with similarity information. *The Journal of Machine Learning Research*, 15(1), 2533-2568.

T. Lu, D. Pl, M. Pl (2010). Contextual multi-armed bandits. *In International Conference on Artificial Intelligence and Statistics*, 485-492.

M. A. Hall (1999). Correlation-based feature selection for machine learning . Doctoral dissertation, The University of Waikato

F. M. Fleegler, J. Griggs, B. Reiner, B. Reville, S. F. Schnall, M. C. Weiss, L. Weissmann (2015). Chemotheray Medicine, Re-trieved from http://www.breastcancer.org/treatment/chemotherapy /medicines

L, Mark (2001). Clinical practice guidelines for the care and treatment of breast cancer: adjuvant systemic therapy for node-positive breast cancer (summary of the 2001 update). *Canadian Medical Association Journal*, 644-646.

P. J. Hamel, P. Johnson (2015). Chemotheray Regimen for Breast Cancer, Retrieved from http://www.healthcentral.com/breast-cancer/chemo-

regimen.html

Y. Yan, R. Rosales, G. Fung, F. Farooq, B. Rao, J. G. Dy (2012). Active learning from multiple knowledge sources. *In International Conference on Artificial Intelligence and Statistics*, 1350-1357.

T. Botsis, G. Hartvigsen, F. Chen, C. Weng (2010). Second-ary use of EHR: data quality issues and informatics opportuni-ties. *AMIA summits on translational science proceedings*, 1.