Using Contextual Learning to Improve Diagnostic Accuracy: Application in Breast Cancer Screening

Linqi Song, William Hsu, *Member, IEEE*, Jie Xu, and Mihaela van der Schaar, *Fellow, IEEE*

***Abstract*—Clinicians need to routinely make management decisions about patients who are at risk for a disease such as breast cancer. This paper presents a novel clinical decision support system that is capable of helping physicians make diagnostic decisions. We apply this support system to improve the specificity of breast cancer screening. The system utilizes clinical context (e.g., demographics, medical history) to minimize the false positive rates while avoiding false negatives. An online contextual learning algorithm is used to update the diagnostic strategy presented to the physicians over time. We analytically evaluate the diagnostic performance loss of the proposed algorithm, in which the true patient distribution is not known and needs to be learned, as compared with the optimal strategy where all information is assumed known, and prove that the false positive rate of the proposed learning algorithm asymptotically converges to the optimum. Moreover, the relevancy of each contextual information is assessed, enabling the approach to identify specific contexts that provide the most value of information in reducing diagnostic errors. Experiments were conducted using patient data collected from a large academic medical center. Our proposed approach outperforms the current clinical practice by 36% in terms of false positive rate given a 2% false negative rate.**

***Index Terms*—Breast cancer, cancer screening, computer-aided diagnosis system, online learning, contextual learning, multi-armed bandit.**

I. Introduction

CLINICAL decision support systems (CDSSs) help clinicians make detection and diagnostic decisions for complex diseases such as lung cancer [1], breast cancer [2][3], and diabetes [4]. There are a number of advantages to integrate CDSSs as part of the clinical workflow instead of solely relying on human intuition. First, the diagnosis accuracy of clinicians varies widely. A previous study has shown that false positive rates for breast cancer detection range from 2.6% to 15.9% among different radiologists, and younger and more recently trained radiologists have higher false-positive rates than experienced radiologists [5]; the deployment of CDSSs may reduce this variability. Second, although clinicians provide the correct diagnostic result in most cases, room for improvement exists in cases where discerning the difference between a benign or malignant mass is difficult [5]-[7]. CDSSs may provide better diagnostic recommendations in these cases by exploiting past knowledge of prior cases and their outcomes. Third, CDSSs help reduce fluctuations in diagnostic performance due to human factors (e.g., fatigue, negligence), offering consistent recommendations. Nevertheless, while CDSSs have been widely advocated to improve the performance of diagnosis, their adoption has remained limited due to various design and usability reasons. We highlight three major challenges here.

Current CDSSs are developed using existing training data, and the strategies of diagnostic recommendations are fixed when being used in the clinic [2][3]. The disadvantage of a fixed pre-determined strategy is that the training data is based on a limited static number of patient cases, which represents an approximation of the real distribution of patients. However, as more patient cases are gathered, an opportunity exists to refine the approximation of the actual distribution to improve the accuracy of diagnostic strategies. Another issue is that the training data may not be available beforehand in some cases. We propose an adaptive online strategy as an integral part of the design framework that is updated continuously based on new cases.

The accelerated pace of technological development for conducting diagnostic tests (e.g., molecular assays) and collecting data (e.g., electronic health records) has yielded a wealth of information that could potentially be used in clinical decision making [6]. Nevertheless, open questions remain on how to process and understand large quantities of heterogeneous data in a computationally efficient manner. In particular, the ability to harness the available data collected about a large population to improve the diagnosis of an individual patient has not been achieved. In this work, we explore how personalization can be taken into account as part of our CDSS framework.

Given the large amount of information available to the CDSS, distinguishing relevant information is important to give effective and efficient diagnostic recommendations [6]. Therefore, the CDSS design framework should be able to identify the relevant data to reach appropriate diagnostic decisions.

These challenges are not fully addressed in previous CDSSs [2][3][6]. In our paper, an online algorithm is proposed for generating diagnostic recommendations to physicians with the objective of minimizing the false positive rate of diagnosis given a prescribed false negative rate. To demonstrate our approach, we focus on the invasive breast cancer diagnosis, since it is one of the common and leading cancers among women with an estimated 232,670 new cases among United States women in 2014 [8][9]. The proposed CDSS is designed to aid physicians with making management decisions on patients undergoing breast cancer screening. Specifically, in borderline cases where the BI-RADS (*Breast Imaging Report and Data System*) score is 3 or 4, it is hard for the physician to determine a diagnostic strategy (e.g., return for a short interval follow-up

L. Song, J. Xu, and M. van der Schaar are with the Department of Electrical Engineering, UCLA, Los Angeles, CA 90095, USA. Email: songlinqi@ucla.edu, jiexu@ucla.edu, mihaela@ee.ucla.edu.

W. Hsu is with the Department of Radiological Sciences, UCLA, Los Angeles, CA 90024, USA. Email: willhsu@mii.ucla.edu.

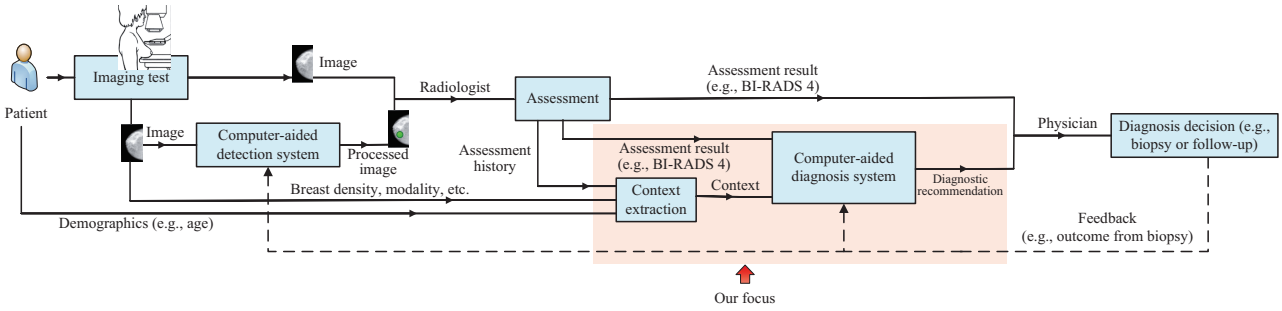

Fig. 1. The breast cancer diagnostic process overview.

or obtain a biopsy). To aid in this determination, we explicitly consider the contextual information of the patient (also known as the situational information) that affects diagnostic errors for breast cancer. The contextual information is captured as the current state of a patient, including demographics (age, disease history, etc.), the breast density (based on the BI-RADS breast density scale), the assessment history, whether the opposite breast has been diagnosed with a mass, and the imaging modality that was used to collect the data. We hypothesize that the incorporation of contextual information will help provide more specific personalized diagnostic recommendations to patients. The key contributions of this paper are summarized below:

- Breast cancer diagnosis is represented as a sequential decision making and online learning problem. Our solution, the Diagnostic Recommendation Algorithm (DRA), makes diagnostic recommendations that vary over time, but converge fast to the optimal strategy. The algorithm exploits the dynamically changing contextual information of the patient to minimize the false positive rate of diagnosis given a false negative rate (e.g., $< 2\%$).
- The term "regret" is used to evaluate the performance of the CDSS, which is defined as the performance gap in terms of the false positive rate between the online learning algorithm and the best oracle diagnostic strategy, which is unknown in practice. The regret associated with the proposed DRA algorithm is analytically quantified, showing that the false positive rate asymptotically converges to the optimal strategy, and the convergence rate is fast (i.e., sublinear).
- The relevancy of contexts is determined in relation to minimizing diagnostic errors. The relevant context analysis identifies what knowledge or information is most influencing the correct diagnosis. This information is provided to the physician who can decide what information to exploit so as to make efficient and effective diagnostic decisions.
- The proposed algorithm's performance is measured through experiments that incorporate clinical, imaging, and pathology data on 4,640 breast screening patients seen at our institution. Results show that an improvement in specificity can be achieved by exploiting the contextual information associated with the patient for the breast cancer diagnosis. Specifically, the proposed algorithm outperforms the current clinical approach by 36% in terms of the false positive rate given a 2% false negative rate.

The rest of the paper is organized as follows. Section II discusses related works. In Section III, we describe the system model and rigorously formulate the computer-aided breast cancer diagnosis system design problem. Section IV presents a systematic methodology for determining the optimal diagnostic recommendation strategy. Section V discusses practical issues related to the system: relevant context selection, prior information, and clinical regret. In Section VI, we present the experimental results and our findings. Section VII concludes the paper.

TABLE I
A COMPARISON WITH PREVIOUSLY PUBLISHED WORKS.

	Employs BI-RADS	Adaptive strategy	Performance guarantee	Trade-off between FPR and FNR
[12]-[29]	No	No	No	No
[30]	Yes	No	No	No
Our work	Yes	Yes	Yes	Yes

II. RELATED WORKS

A. Computer-Aided Detection and Diagnosis System for Breast Cancer

Various signal processing and machine learning techniques have been introduced to perform the computer-aided detection and diagnosis system. Early works have focused on image processing and classification techniques to extract features of the image and predict the outcome (i.e., whether it is benign or malignant) in the image [12]-[20]. A neural network-based algorithm [14] and a linear discriminant approach [2] are proposed to solve the diagnosis problem.

In breast screening, two types of CDSSs can be integrated in the diagnostic workflow, as depicted in Fig.1: 1) the computer-aided detection system, which helps the radiologist identify important features in the image that are abnormal [21]-[23]; and 2) the computer-aided diagnosis system, which helps physicians determine the diagnostic strategy for the patient (e.g., whether a patient should receive a biopsy or continue follow-up imaging) [24]-[30]. Our approach utilizes contextual information (e.g., breast density, modality) to make diagnostic recommendations to the physician with the goal of reducing the number of biopsies performed while maintaining a low number of false negatives. Clinically, management decisions that involve further interventions are indicated by a BI-RADS score of 0 (the imaging study cannot be interpreted and must be retaken), 4 (or 4A, 4B, and 4C, which represent different levels of suspicion), or 5 (high suspicion) [10][11][31]. A BI-RADS score of 3 means that the mass is likely benign with a short interval follow-up recommended. On the other hand, a BI-RADS 4 or 5 indicates that a biopsy is recommended. These decisions are made in the context of other clinical

TABLE II
A COMPARISON WITH EXISTING MAB ALGORITHMS

	Explores patients	Considers prior information	Considers relevant context	Optimizes under constraint
Prior works [38]-[40]	Yes	No	No	No
Our approach	No	Yes	Yes	Yes

variables (e.g., if the mass is palpable). The difficulty lies in borderline cases (e.g., BI-RADS 3 or 4A) where indications are unclear whether a biopsy is truly necessary. For example, despite the cost and risk associated with biopsies, the positive predictive value for BI-RADS 4A is only 9% [34]. More efficient and accurate approaches are needed to reduce unnecessary biopsies [35]. In [30], a neural-fuzzy rule has been proposed to implement the computer-aided diagnosis system. However, these rule-based algorithms [2][3][30][31] cannot be easily updated. Such approaches incur some performance loss because the underlying distribution of patient (outcome and context) is not known, and limited training data gives a limited estimation of the actual distribution, resulting in suboptimal diagnostic strategies. A PO-MDP based algorithm has been proposed in [36] to solve a similar screening related question. However, the algorithm [36] does not learn unknown distributions. Our paper focuses on creating a system that aids physicians with making decisions on borderline cases such as BI-RADS 4A, as highlighted in Fig. 1. Compared with existing works [2][3][30][31], our proposed framework employs an online learning approach, which accumulates data and adaptively updates the diagnostic strategy over time. In this way, our approach unravels the optimal strategy over time. Moreover, the proposed learning algorithm learns (provably) fast this optimal strategy. A detailed comparison of our framework with existing frameworks is shown in Table I.

B. Contextual Multi-Armed Bandit

Our diagnostic recommendation algorithm is based on the contextual multi-armed bandit (MAB) framework [38][39][40] and incorporates the following innovations. First, prior information is considered, allowing the system to learn directly from prior information or other learners. Second, in existing works [38][39][40], the estimated error of an action can be updated only after the action is selected, and the algorithm needs to explore patients (by recommending different diagnostic actions to different patients under the same context) in order to get information about every action. However, the exploration of patients is considered immoral. In our algorithm, the diagnostic error of any action can be updated each time, and we do not need to explore patients. Third, our algorithm considers minimizing the false positive rate, given a false negative rate constraint. Existing works do not consider such a constraint. A summary that compares our work with existing MAB works is shown in Table II.

III. SYSTEM MODEL

A. Computer-Aided Breast Cancer Diagnosis System

We consider a computer-aided breast cancer diagnosis system (CABCDS) as shown in Fig. 2. We consider discrete time slots $t = 1, 2, ...$ Each time, one patient arrives with a borderline test result (e.g., a mass is seen, but suspicion level for malignancy is low). The system collects the contextual information $x_t \in \mathcal{X}$ of the patient, where $\mathcal{X}$ is the d_X-dimensional bounded context space. The bounded distribution of context is denoted by $f(x)$. Then, the system makes the diagnostic decision $a_t \in \mathcal{A}$ and recommends it to the physician. The decision space is denoted by $\mathcal{A} = \{0, 1\}$, where 0 represents a short interval follow-up (e.g., patient is re-imaged in 6 months) and 1 represents the recommendation of a biopsy. This can easily be extended to incorporate more decisions.

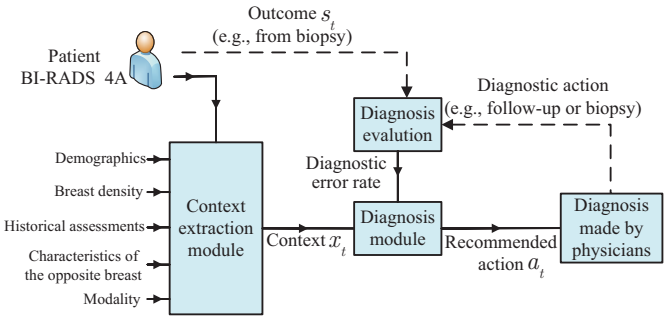

Fig. 2. Computer-aided breast cancer diagnosis system model.

B. Contextual Information Extraction

To better assist physicians, the CABCDS system considers a diverse set of contextual information to make sufficiently accurate recommendations. As shown in Table III, the following types of contexts are considered: patient demographics (e.g., age, race), breast density, assessment history, whether the opposite breast has a high BI-RADS score previously (e.g., achieve BI-RADS 3), and imaging modality. Regardless of whether the context variable is discrete or continuous, each context is modeled as discrete values that have numerical values between 0 and 1[1]. In this way, we normalize the context space as $\mathcal{X} = [0, 1]^{d_X}$. For example, a context can be represented as $x = (0.7, 0.2, 0, 1, 0)$, where $x_1 = 0.7$ represents a 70-year old patient with a scattered fibroglandular breast density ($x_2 = 0.2$), an assessment of only BI-RADS 1 or 2 ($x_3 = 0$) in the preceding screening study, an assessment of BI-RADS 3 for the opposite breast ($x_4 = 1$), and mammography as the imaging modality used ($x_5 = 0$).

C. Diagnostic Recommendation Problem

Before describing the diagnostic recommendation problem, several key concepts are introduced.

Diagnostic strategy: A diagnostic strategy is the approach to selecting an action, either to undergo a biopsy or to follow up, based on the observed contextual information. Formally, the diagnostic strategy is defined as mapping the observation of contextual information to a diagnostic decision, namely, $\pi : \mathcal{X} \to \mathcal{A}$. $\pi(x)$ represents the action selected by the diagnostic strategy when the patient has context x. The strategy set is denoted by Π.

Patient outcome: The outcome of the patient at time t is $s_t(x_t) \in \mathcal{S}$. The outcome space is defined as $\mathcal{S} = \{0, 1\}$,

[1]Note that the modeling of discrete contexts in different categories can also be considered as several learners, each corresponding to one category.

TABLE III
TYPES OF CONTEXTS AND DESCRIPTIONS

Context	Description
Demographics	The characteristics of a patient, age, race, disease history, family medical history, etc.
Breast Density	**Group 1**: The breast is almost entirely fat (fibrous and glandular tissue <25%). **Group 2**: There are scattered fibroglandular densities (fibrous and glandular tissue 25% to 50%). **Group 3**: The breast tissue is heterogeneously dense (fibrous and glandular tissue 50% to 75%). **Group 4**: The breast tissue is extremely dense (fibrous and glandular tissue > 75%).
Historical assessments	The information contained in previous imaging exam assessments (e.g., whether findings in BI-RADS 3 or higher appear in the past, or whether there is a significant change in the past year).
Characteristic of the opposite breast	The information of the opposite breast (e.g., whether findings in BI-RADS 3 or higher appear for the opposite breast).
Modality	The modality used for imaging: mammography (MG), ultrasound (US), magnetic resonance imaging (MRI) or computer radiography (CR).

where 0 represents benign, and 1 represents malignant. If a patient undergoes a biopsy or returns for a short-term follow-up, the patient's outcome is revealed, where if the patient has been followed up for a certain time and the condition is stable, then the outcome is considered benign.

Patient outcome distribution: The underlying distribution of patient outcomes is represented as $\sigma(x) = \Pr\{s(x) = 1\}$, where $\sigma(x)$ represents the probability of being malignant for a patient with context x.

Diagnostic errors: Two types of diagnostic errors are considered: false positive (e.g., if the outcome $s_t(x_t)$ is benign, and the recommended action is to undergo a biopsy) and false negative (e.g., if the outcome $s_t(x_t)$ is malignant, and the recommended action is a short-term follow-up).

False positive/negative indicator: The false positive indicator is defined as $c_1(\pi(x), s(x)) = I\{\pi(x) = 1, s(x) = 0\}$, where $I\{\cdot\}$ is the indicator function. When a false positive occurs (i.e., the outcome is benign and the action is to undergo a biopsy), then $c_1(\pi(x), s(x)) = 1$; otherwise $c_1(\pi(x), s(x)) = 0$. When an action $\pi(x) = 1$ (undergo biopsy) is chosen, the expectation of $c_1(\pi(x), s(x))$ is $\Pr(s(x) = 0) = 1 - \sigma(x)$ (i.e., the probability of being benign). When an action $\pi(x) = 0$ (follow-up) is chosen, the expectation of $c_1(\pi(x), s(x))$ is 0. Similarly, the false negative indicator is defined as $c_0(\pi(x), s(x)) = I\{\pi(x) = 0, s(x) = 1\}$, and its expectation is the probability of being malignant $\sigma(x)$ when action $\pi(x) = 0$ (follow-up) is chosen and is 0 when the action is $\pi(x) = 1$ (biopsy).

False positive/negative rate: The false positive rate is defined as the probability of selecting an action $\pi(x) = 1$ (biopsy) when the patient outcome $s(x) = 0$ (benign), denoted by $\mu_1(\pi(x), s(x)) = \Pr(\pi(x) = 1 | s(x) = 0)$. Similarly, the false negative rate is defined as $\mu_0(\pi(x), s(x)) = \Pr(\pi(x) = 0 | s(x) = 1)$.

The diagnostic recommendation problem is formulated as minimizing the false positive rate given a maximal false negative rate η (e.g., $< 2\%$). The physician or the policy of the clinic has control over the trade-off between false positive and false negative rates. Therefore, the diagnostic recommendation problem can be formally written as:

$$\begin{aligned} &\min_{\pi \in \Pi} E_x \mu_1(\pi(x), s(x)) \\ &s.t.\ E_x \mu_0(\pi(x), s(x)) \leq \eta \end{aligned} . \tag{1}$$

IV. DIAGNOSTIC RECOMMENDATION ALGORITHM

In this section, the algorithm used for the diagnostic recommendation is elaborated. The main idea is to adaptively cluster patients based on associated contexts and to learn the best action for each patient cluster.

A. Optimal Strategy

In order to solve the diagnostic recommendation problem, we first analyze the structure of the optimal solution where all information (i.e., the distribution of context $f(x)$ and the distribution of patient outcome $\sigma(x)$) is known. We observe that the underlying distribution $\sigma(x)$ varies for different contexts x, and hence the solution is to recommend undergoing a biopsy for patients with a sufficiently high probability of being malignant, and to recommend a short interval follow-up for patients with a sufficiently low probability of being malignant in the diagnostic recommendation problem. Hence, we first prove that the optimal strategy is a threshold-based strategy that is described by the following proposition.

Proposition 1: The optimal strategy $\pi^*(x)$ for the diagnostic recommendation problem in eq. (1) is a threshold strategy: there exists a threshold σ_η, such that the optimal strategy satisfies $\pi^*(x) = 1$, if $\sigma(x) \geq \sigma_\eta$, and $\pi^*(x) = 0$, otherwise.

Proof: See Appendix A. ■

Note that the context distribution $f(x)$ and outcome distribution $\sigma(x)$ are not known in practice. As such, the algorithm needs to learn the distribution of contexts and outcomes.

B. Algorithm Description

As introduced in Section III, the probability of being malignant is different for patients with different contexts. However, mining the entire contextual information space which is large is challenging. In our proposed algorithm, we cluster patients with similar contexts to accumulate enough information for a specific patient and make recommendations based on the accumulated information. More importantly, the algorithm adaptively shrinks the cluster size when information of similar patients becomes available, allowing more specific recommendations to be made.

The diagnostic recommendation algorithm is formally presented in Table IV, and an illustration is given in Fig. 3. The algorithm consists of two steps: the context space refinement and the optimal action selection. In the refinement process, the context space is clustered into 2^{ld_X} clusters, if the context space is uniformly partitioned on each dimension by 2^l. This clustering method is called a level l partition $\mathcal{P}_l$, and $\mathcal{P} = \cup_{l=0}^{\infty} \mathcal{P}_l$ denotes the set of all possible clusters. In each time point, the algorithm keeps a set of mutually exclusive clusters (referred to as *active clusters*) that consist of the entire population of patients. The set of active clusters at

TABLE IV
DIAGNOSTIC RECOMMENDATION ALGORITHM (DRA)

Initialization: $\mathcal{C}^1 = \mathcal{P}_0$, $M_{\mathcal{X}} = 0$, $\bar{\sigma}_{\mathcal{X}} = 1$.

for $t = 1, 2, \ldots$ **do**
  Context x_t arrives. Determine $C \in \mathcal{C}^t$, such that $x_t \in C$.
  Determine the threshold σ_t (step 4-5): initialize $\sigma_t = 1$.
  Determine strategy for each $C' \in \mathcal{C}^t$: ▷ Pre-action selection
    set pre-action $\pi^t(C') = \begin{cases} 0, \ if \ \bar{\sigma}_{C'} < \sigma_t \\ 1, \ otherwise \end{cases}$.
  Check the average false negative rate: ▷ Constraint checking
    Calculate the average false negative rate:
    $$\bar{\mu}_0 = \frac{\sum_{C' \in \mathcal{C}^t} M_{C'} I\{\pi^t(C')=0\}}{\max\{1, t-1\}}.$$
    if $\bar{\mu}_0 \leq \eta$, select the action $\pi^t(C)$, and go to step 6;
    otherwise, $\sigma_t = \sigma_t - t^{-1}$ and go to step 4.
  The outcome of the patient s_t is revealed.
  Update $\bar{\sigma}_C = [M_C \bar{\sigma}_C + I\{s_t = 1\}]/(M_C + 1)$.
  Update $M_C = M_C + 1$.
  if $M_C \geq 2^{pl}$ **then** ▷ Context space refinement
    Set $M_{C'} = \sum_{\tau \leq t} I\{x_\tau \in C'\}$ for all $C' \in \mathcal{P}_{l+1}(C)$.
    Set $\bar{\sigma}_{C'} = \frac{\sum_{\tau \leq t} I\{x_\tau \in C', s_\tau = 1\}}{\sum_{\tau \leq t} I\{x_\tau \in C'\}}$ for all $C' \in \mathcal{P}_{l+1}(C)$.
    Set $\mathcal{C}^{t+1} = (\mathcal{C}^t \backslash \{C\}) \cup \mathcal{P}_{l+1}(C)$.
  else
    Set $\mathcal{C}^{t+1} = \mathcal{C}^t$.
  end if
end for

time t is denoted by $\mathcal{C}^t$. For example, consider a patient cluster as patients aged 40 to 49 with breast density of Group 1 or 2. As new patients are added that fit this cluster, an accurate estimation of the patient outcome distribution can be characterized for this cluster. The cluster can be partitioned into finer clusters in order to make more specific diagnostic recommendations: patients aged 40 to 44 with breast density Group 1, patients aged 40 to 44 with breast density Group 2, patients aged 45 to 49 with breast density Group 1, and patients aged 45 to 49 with breast density Group 2. When the number of arrivals in a cluster $C \in \mathcal{C}^t$ exceeds some level-varying threshold 2^{pl}, where $p > 0$ is an empirical design parameter[2], the cluster C is partitioned into 2^d level $l+1$ clusters, as illustrated in Fig. 3.

Each time a patient with context x_t arrives, the algorithm assigns the patient to an active cluster $C \in \mathcal{C}^t$ with matching contextual information. The algorithm keeps an estimation $\bar{\sigma}_C$ of the probability of a patient in that cluster having a malignancy and uses an adaptive threshold-based selection rule to select the action, where the action is 1 if $\bar{\sigma}_C$ exceeds the threshold, and is 0 otherwise. The threshold is selected adaptively as large as possible each time to make sure that the average false negative rate is within a pre-defined level.

C. Evaluation of Algorithm Performance

In this subsection, the performance of the proposed DRA algorithm is analyzed in terms of the learning regret, which is the expected false positive rate of our learning algorithm compared with the optimal strategy $\pi^*(x)$ that assumes all information (i.e., the distribution of context $f(x)$ and the distribution of patient outcome $\sigma(x)$) is known. In practice, the information is not known and needs to be learned. Formally,

[2]The empirical parameter p depends on how fast the clusters are to be partitioned. A smaller p results in a faster partition process of the context space, and a larger p results in a slower partition process of the context space.

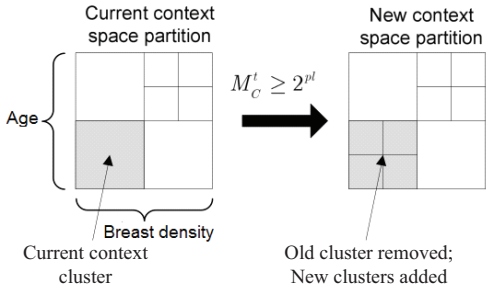

Fig. 3. An illustrative example of the DRA algorithm.

the regret is defined as

$$R_\pi(T) = \sum_{t=1}^{T} [E_{x_t} \mu_1(\pi^t(x_t), s(x_t)) - E_{x_t} \mu_1(\pi^*(x_t), s(x_t))]. \quad (2)$$

Similarity between two patients is used as a measure to effectively cluster the patients into groups. For example, the probability of being malignant for a patient aged 46 is similar to a patient aged 47, given that other types of contexts are the same. The Lipschitz condition is used to characterize the patient outcome distribution and the context distribution. Importantly, the Lipschitz condition is only used to characterize the system performance, and the algorithm can be performed even if this condition is not satisfied perfectly in practice.

Lipschitz Condition:

(1) **Outcome distribution**: there exists a Lipschitz constant $L_1 > 0$, such that for all $x, x' \in \mathcal{X}$, we have $|\sigma(x) - \sigma(x')| \leq L_1 ||x - x'||$.

(2) **Context distribution**: there exists a Lipschitz constant $L_2 > 0$, such that for all $x, x' \in \mathcal{X}$, we have $|f(x) - f(x')| \leq L_2 ||x - x'||$.

The regret is proven in the following theorem.

Theorem 1: If 1) the Lipschitz condition holds for the patient outcome distribution and context distribution; 2) patients arrive with independently identically distributed bounded context distribution (i.e., $0 < f_{\min} \leq f(x) \leq f_{\max}$); and 3) $\Pr\{x : \sigma(x) = \sigma_\eta\} = 0$; then the regret of the DRA algorithm up to time T can be bounded by

$$R(T) = O(T^{g(d_X)}), \quad (3)$$

where $g(d_X) = \frac{d_X + 1/2 + \sqrt{9 + 8d_X}/2}{d_X + 3/2 + \sqrt{9 + 8d_X}/2}$.

Proof: See Appendix B. ■

Note that $0 < g(d_X) < 1$, and therefore, the average short-term performance $R(T)/T$ goes to 0, as T goes to infinity, implying that the proposed DRA algorithm will converge to the optimal diagnostic performance. Accordingly, the following corollary is introduced.

Corollary 1: The performance of the proposed DRA algorithm converges to the optimal performance, in terms of the false positive rate.

In addition, from Theorem 1, the convergence speed is fast, at least in a sublinear[3] rate.

D. Algorithm Performance with Known Threshold

In previous sections, the algorithm is shown to adaptively learn the optimal threshold σ_η (probability of being malignant) over time, given a false negative rate constraint. However, in

[3]A sublinear rate indicates that the expected performance loss is $O(1/t^\gamma)$ for time period t, where $0 < \gamma < 1$.

some clinical contexts, a fixed false negative rate may already exist (e.g., based on physician preference or clinical practice guidelines) [10] [11]. In this situation, the DRA algorithm degrades to a fixed threshold-based algorithm that does not need to learn the threshold value over time. Hence, step 5 of the DRA algorithm can be omitted, and the algorithm only needs to learn the distribution of patient outcomes $\sigma(x)$. We consider a weighted false positive and false negative error in this setting:

$$c(\pi(x), s(x)) = \sigma_\eta c_1(\pi(x), s(x)) + (1 - \sigma_\eta)c_0(\pi(x), s(x)). \tag{4}$$

The strategy for minimizing the weighted error $c(\pi(x), s(x))$ is obtained by

$$\min_{\pi \in \Pi} Ec(\pi(x), s(x)). \tag{5}$$

The optimal strategy is denoted by $\pi^\dagger(x)$, which assumes all information is known, leading to the following proposition:

Proposition 2: The optimal strategy $\pi^\dagger(x)$ is equivalent to the optimal strategy $\pi^*(x)$ for the same σ_η.

Proof: Appendix C. ■

Intuitively, Proposition 2 shows that the optimal weighted error minimization strategy is equivalent to the optimal threshold-based strategy. Hence, we consider the regret, difference in the weighted error, comparing our fixed threshold-based DRA algorithm to the optimal strategy $\pi^*(x)$. Formally, the regret is defined as

$$R_\pi(T) = \sum_{t=1}^{T} [Ec(\pi^t(x_t), s(x_t)) - Ec(\pi^*(x_t), s(x_t))]. \tag{6}$$

We have the following theorem to bound this regret:

Theorem 2: The regret of the fixed threshold-based DRA algorithm is bounded by

$$R(T) = O(T^{g(d_X)}). \tag{7}$$

Proof: See Appendix D. ■

The regret in terms of weighted error of the fixed threshold-based DRA algorithm has the same sublinear order as the regret in terms of false positive rate of the DRA algorithm. This finding implies that the fixed threshold-based DRA algorithm will converge to the optimal diagnostic strategy, summarized by the following corollary.

Corollary 2: The performance of the fixed threshold-based DRA algorithm converges to the optimal performance, in terms of the expected weighted error in eq. (4).

In addition, from Theorem 2, the convergence speed is fast, at least in a sublinear rate.

V. Practical Considerations

In this section, we discuss some practical issues related to the system implementation, and give appropriate approaches to address these issues.

A. Relevant Context Analysis

While large amounts of data are routinely captured about a patient as part of clinical care, some information is inherently more relevant to assessing the probability of breast cancer than others. In an online learning setting, identifying which contextual information is more relevant to making a clinical recommendation based on retrospective data is important. For the d_X-dimensional context space $\mathcal{X}$, the chance of making diagnostic errors may be correlated with missing or noisy data. For example, the chance of making a diagnostic error may be low when results of a molecular assay is available, but the chance of making a diagnostic error may be high when only information about a patient's previous BI-RADS assessment is known.

For a d_X-dimensional context space, 2^{dx} DRA learning instances can be executed at the same time. At time t, the average false positive rates of all the learning instances are evaluated. We denote by instance 1 the learning instance using all d_X-dimensional contextual information. If for another learning instance i, the difference of false positive rate compared with instance 1 is below some level given by physicians, then the contextual information used by learning instance i is relevant. Hence, the system can select the more relevant context to make diagnostic recommendations, as shown in Fig. 4. As can be seen from Theorem 1, when less contextual information is used, the convergence rate improves. This property is demonstrated in practice using actual data in Section VI-C.

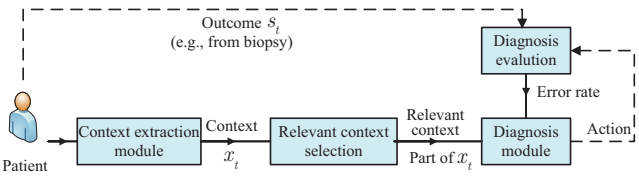

Fig. 4. The CABCDS system with relevant context.

B. Learning with Prior Information

Although the relevant context analysis module can help to identify significant predictors from the entire information space, the selection process can be challenging when little information is known about the underlying patient distribution, a problem known as "cold start". One approach to solve this is to introduce prior contextual information, such as probabilistic statements that have been previously reported in other research studies, showing the relationship between contexts and the probability of cancer [37]. This prior information can be represented as an input parameter into the relevant context analysis module to derive an initial distribution.

Another type of prior information, the statistical data, can be former patient cases with known outcomes or based on previous researches. The prior statistical information includes the distribution of patients with cancer, the false negative rate and false positive rate obtained from other studies. The effect of using the prior statistical information is equivalent to a number of N training patient cases before running the system. In this case, the regret up to time T can be bounded by $R(T) = O((T-N)^{g(d_X)})$. This shows that the performance of the system is greatly improved at the beginning of its operation since it can successfully capitalize on the prior knowledge. This system, which capitalizes on the availability of prior information (knowledge).

C. Clinical Regret Analysis

In previous sections, we design the algorithm for the CABCDS system and analyze the performance of the algorithm in terms of learning regret. However, in practice, physicians may not exactly take the actions recommended by the system. In this section, we assess the clinical regret by considering the actions of physicians.

Note that the regret we have analyzed in previous sections is used to evaluate the performance of the designed system. In order to evaluate the real diagnostic errors in clinic, we use *clinical regret* to distinguish from the regret used in previous sections.

We first consider that physicians have a certain but fixed chance of not following the recommended strategy, due to a difference in opinion between the experience of the physician and the recommended diagnostic strategy, or caused by human diagnostic error. In this scenario, we denote the chance of not following the recommended strategy by ε. That is, when the diagnostic strategy recommended by the CABCDS system is a_t, the physician has a probability ε of selecting a strategy $\hat{a}_t \neq a_t$.

Theorem 3: Given the fixed chance ε of not following the recommended strategy, the clinical regret of up to time T can be bounded by

$$R(T) = O(T^{g(d_X)}) + O(\varepsilon T). \tag{8}$$

Proof: See Appendix E. ■

In this case, there is a linear regret term in T. In fact, since the learning regret of the DRA algorithm is sublinear in T, and converges to the optimal strategy, a constant chance of deviation will result in a linear clinical regret in T. Next, we consider that physicians have a decreasing chance of not following the recommended strategy. This is caused by the interactions between the physicians and the system. In this scenario, we denote the chance of not following the recommended strategy by $\varepsilon_t = \frac{1}{t^\beta}$ $(0 < \beta < 1)$. That is, when the diagnostic strategy recommended by the CABCDS system is a_t, the physician has a probability ε_t of selecting a strategy $\hat{a}_t \neq a_t$.

Theorem 4: Given the decreasing chance $\varepsilon_t = \frac{1}{t^\beta}$ of not following the recommended strategy, the clinical regret of up to time T can be bounded by

$$R(T) = O(T^{g(d_X)}) + O(T^{1-\beta}). \tag{9}$$

Proof: See Appendix E. ■

In this case, the clinical regret has another sublinear term $O(T^{1-\beta})$. In fact, both the learning regret of the DRA algorithm and the clinical regret are sublinear in T, and hence converge to the optimal strategy.

We show the difference in convergence between fixed and decreasing chance of not following the recommended strategy.

Corollary 3: The clinical performance in terms of false positive rate does not converge to the optimal performance when physicians have a fixed chance of not following recommendations provided by the DRA algorithm, and converges to the optimal performance when physicians have a decreasing chance of not following recommendations provided by the DRA algorithm.

VI. Experimental Results

In this section, the performance of the designed system is shown using our proposed algorithm. First, the breast cancer dataset used to evaluate the system performance is described. Then, our proposed online learning algorithm is evaluated and compared with other exiting algorithms. Finally, the impact of relevant contexts on the system performance in terms of diagnostic error rate and convergence rate is discussed.

A. Data Description

A de-identified dataset of 4,640 individuals who underwent screening at a large academic medical center is used. Patient outcome is derived from biopsy result, which is typically obtained for individuals with a BI-RADS score of 4 or 5. Our focus is on analyzing cases that are BI-RADS 4A; this category represents patients, whose test results are less suspicious for cancer, raising the concern about unnecessary biopsies. A five dimensional context space is used, which includes patient age, breast density, assessment history (whether or not the immediately preceding exam shows a finding of BI-RADS 3 or above), assessment results for the opposite breast (whether or not the immediately preceding exam shows a finding of BI-RADS 3 or above), and the imaging modality used.

Characteristics of different BI-RADS categories are shown in Table V. The probability of being malignant increases from 9.91% to 78.61% as the BI-RADS category varies from 4A, 4B, to 4C. Prior to the introduction of BI-RADS 4A, 4B, 4C, all suspicious nodules were categorized as BI-RADS 4. The probability of being malignant of BI-RADS 4 is 26.12%, which is between those of BI-RADS 4A and 4B and near the total average probability of being malignant.

TABLE V
Description of Different Categories

BI-RADS	No. instances	Prob. of malignant
4	2282	26.12%
4A	1171	9.91%
4B	827	37.24%
4C	360	78.61%
Total	4640	28.08%

B. Performance Evaluation of the DRA algorithm

To perform the online adaptive learning, the data instances described previously are sequentially fed to the algorithm. Results are compared with the clinical approach, the neural-fuzzy approach, and the linear discriminant analysis approach, which are defined as follows:

- **Clinical approach** [31]: Current clinical practice may be thought of as a threshold-based approach, which recommends a biopsy for all patients that fall in BI-RADS 4, 4A, 4B, 4C and above.
- **Neural-fuzzy approach** [14][30]: The neural-fuzzy approach models the diagnosis system as a three-layered neural network. The first layer represents input variables with various patient features; the hidden layer represents the fuzzy rules for diagnostic decision based on the input variables; and the third layer represents the output diagnostic recommendations.

TABLE VI
FPR AND FNR COMPARISON

Algorithms	$\eta = 5\%$			$\eta = 2\%$		
	FPR	FNR	Accuracy	FPR	FNR	Accuracy
DRA (proposed)	61.0%	4.4%	0.45	64.2%	2.0%	0.42
Neural fuzzy	71.0%	4.9%	0.36	72.8%	1.9%	0.34
Clinical	100%	0	0.10	100%	0	0.10
LDA	92.1%	7.8%	0.16	92.1%	7.8%	0.16

- **Linear discriminant analysis (LDA)** [2][32]: The LDA approach trains a classifier using features extracted from imaging tests and assessment report, and the trained classifier can be used to make diagnostic recommendations.

System performance using different algorithms is shown in Fig. 5, Fig. 6, and Table VI. The false negative rate is empirically given as 5% and 2%, respectively. We assume that the same prior information or training data is available for each algorithm. Results show the relationship between average false positive rate and the percentage of patient arrivals (patient cases). The false positive rate of the proposed DRA algorithm decreases over time. In order to achieve a lower false negative rate, the false positive rate and the accuracy need to be sacrificed. Table VI shows that the false positive rate of the clinical approach is 1 for BI-RADS 4A patients, since it simply recommends all BI-RADS 4A patients to undergo a biopsy. The LDA algorithm cannot satisfy the false negative rate constraint, since it tries to linearly cluster patients based on the contextual information. However, the structure of the contextual information may not be linear, and as a result, a big performance loss is incurred. The neural fuzzy approach results in a performance loss and does not converge to optimum, likely because the trained rule may not be optimal and cannot be adaptively updated in time. Our DRA approach can be updated over time, maintaining a balance between false negative and false positive rates. The DRA algorithm outperforms the clinical approach in terms of the false positive rate by 39% and 36% for $\eta = 5\%$ and $\eta = 2\%$, respectively, the LDA approach in terms of the false positive rate by 34% and 30% for $\eta = 5\%$ and $\eta = 2\%$, respectively, and the neural-fuzzy approach in terms of the false positive rate by 14% and 12% for $\eta = 5\%$ and $\eta = 2\%$, respectively.

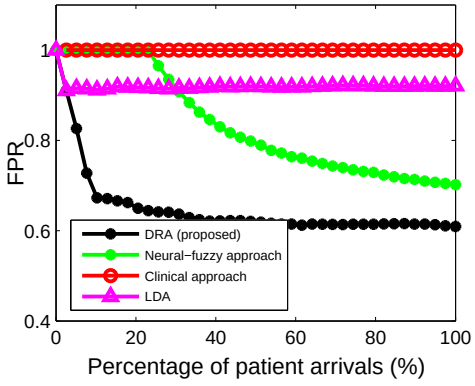

Fig. 5. Comparison of FPR for different algorithms, given tolerable FNR=5%.

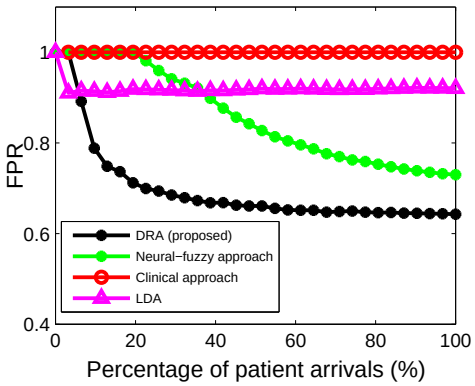

Fig. 6. Comparison of FPR for different algorithms, given tolerable FNR=2%.

C. Relevant Context Analysis

Different combinations of context components result in varying action selections. The false positive rate and false negative rate using different combinations of context components to analyze our proposed DRA algorithm for BI-RADS 4A patients are presented in Table VII. Thirty two different cases are considered. From the table, the relevance of different contextual information to predicting patient outcome is quantitatively described. First, no significant change in false positive rate is seen when the age information is considered, by comparing Case 1 with Case 17 from Table VII. Although women with different ages have significantly different chances of having breast cancer [37], our results imply that the information about patient age plays a less important role in determining the diagnostic strategy than the imaging test information, such as breast density, assessment history, characteristic of opposite breast, and modality. Second, results show the importance of considering breast density and modality in order to achieve a diagnostic recommendation. Case 1 and Case 10 show that without the information about breast density and modality, the false positive rate increases by over 14% for both scenarios of 2% and 5% tolerable false negative rate. In addition, taking into account the information about breast density without knowing the modality can result in a significant increase in false positive rate, as shown by Case 1 and Case 2. In fact, no research has shown that breast density significantly implies the risk of cancer [10], but the breast density may cause lesions to be obscured in mammography [10][33]. Hence, using different modalities, such as mammography, ultrasound, and magnetic resonance imaging, can help reduce the diagnostic error when the patient has dense breasts. Third, by observing Case 1, Case 3, Case 5, and Case 7, assessment history of both breasts plays an important role in determining the diagnostic strategy (resulting in a 16% decrease in false positive rate) when the tolerable false negative rate is low (i.e., 2%), and affects little (less than 7% variation in false positive rate) when the tolerable false negative rate is high (i.e., 5%). Hence, the information about the assessment history of both breasts needs to be considered when a low false negative rate is suggested.

As discussed in Section V, the different selection of context also affects the convergence rate of our online learning algorithm. In Fig. 7 and 8, we show results of convergence rate for the above discussed Case 7, where the assessment history

of both breasts is not considered, Case 17, where the age information is not considered, and Case 1 with all contextual information. We can see from Fig. 7 that for a high tolerable false negative rate (5%), the convergence rate of Case 17 with a low context dimension is higher than that of the Case 7 with a medium context dimension, as well as than that of Case 1 with a high context dimension. However, for the low tolerable false negative rate (2%), the convergence rate of Case 7 is very low. In fact, as we previously showed, Case 7 does not consider the relevant contextual information regarding the assessment history of both breasts. This results in poor performance in terms of both the learning speed as well as the false positive rate in the scenario of a low tolerable false negative rate (2%).

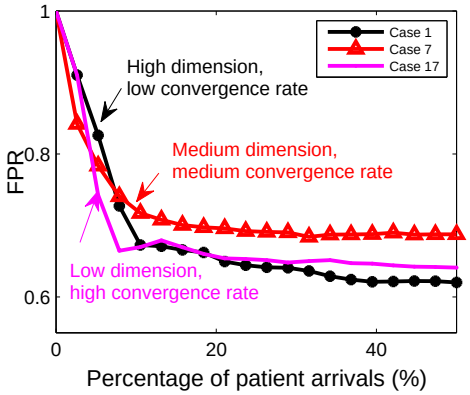

Fig. 7. Comparison of convergence rate for different context selection, tolerable FNR=5%.

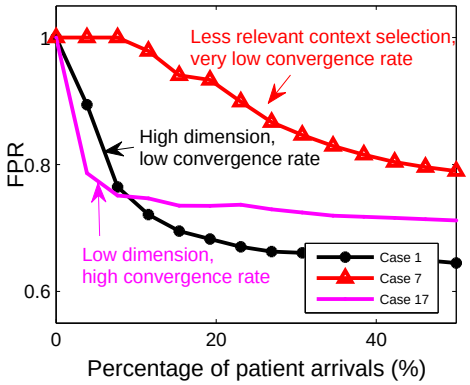

Fig. 8. Comparison of convergence rate for different context selection, tolerable FNR=2%.

VII. Discussion and Future Works

This paper presents a novel design framework for a computer-aided breast cancer diagnosis system. The system incorporates contextual information and makes diagnostic recommendations to physicians, aiming to minimize the false positive rate of diagnosis, given a pre-defined false negative rate. The proposed algorithm is an online algorithm that allows the system to update the diagnosis strategy over time. We analytically show that the performance of our proposed algorithm converges to the optimal performance and quantify the rate of convergence. In addition, our framework can learn the relevancy of contexts to a specific diagnosis, thereby helping physicians become aware of what contextual information they should pay special attention to when making diagnostic decisions. Experiments conducted using data from the Department of Radiological Sciences show that our proposed method outperforms the current clinical approach by 36% in terms of the false positive rate, if a 2% tolerable false negative rate is used. The relevant context simulations show that age information for a patient is less relevant than breast density and imaging modality, which are two important context components in determining the diagnostic strategy. The assessment history of both breasts is more relevant if a higher false negative rate (5%) is acceptable.

One future work is to continue evaluating the design framework and explore its implementation in the clinic. Understanding the utility and impact of the proposed approach in the current practice of breast screening requires further study. Nevertheless, the initial experimental results demonstrate that our online contextual learning algorithm is adaptable, and potentially could be used in other disease domains. Each of these domains has its own unique set of contextual information and desired patient outcomes. For example, in lung cancer screening patients, results of the low-dose computed tomography study (e.g., characterization of the nodule), pulmonary function tests, smoking and medical history, and environmental exposures are potential contexts. The contextual online learning algorithm can be adapted to handle these types of variables, helping physicians leverage available clinical big data to inform clinical decisions in each of these respective disease domains.

Appendix A
Proof of Proposition 1

The recommended strategy should be consistent, i.e., $\pi(x') \leq \pi(x'')$ if $\sigma(x') \leq \sigma(x'')$. In this case, the optimal solution will be among the threshold-based strategies: $\pi_\sigma(x) = 1$, if $\sigma(x) \geq \sigma$, and $\pi_\sigma(x) = 0$, otherwise. We also notice that for threshold-based strategies, the monotone property is satisfied: $E_x\mu_0(\pi_\sigma(x), s(x)) \leq E_x\mu_0(\pi_{\sigma'}(x), s(x))$ and $E_x\mu_1(\pi_\sigma(x), s(x)) \geq E_x\mu_1(\pi_{\sigma'}(x), s(x))$, when $\sigma \geq \sigma'$. This basically indicates that when a higher threshold is chosen, actions for some contexts will change from undergoing a biopsy to follow-up. Therefore, the false negative rate is reduced and the false positive rate is increased. Hence, a threshold σ_η exists, such that for any threshold-based strategy $\pi_\sigma(x)$ with $\sigma > \sigma_\eta$, the following property holds: $E_x\mu_1(\pi_\sigma(x), s(x)) > \eta$, and $E_x\mu_1(\pi_{\sigma_\eta}(x), s(x)) \leq \eta$. Obviously, the optimal solution is $\pi_{\sigma_\eta}(x)$, and we write the optimal solution as $\pi^*(x)$ for short.

Appendix B
Proof of Theorem 1

We first describe the main idea of the proof, and then prove it. The main idea is to cluster the contexts into small clusters over time. Within each context cluster, if the outcome of patient $\sigma(x)$ within the context cluster has a gap from σ_η, and the estimation of $\sigma(x)$ is accurate enough, then the actions in these context clusters are the same as the optimal strategy. The

TABLE VII
IMPACT OF RELEVANT CONTEXTS

Case	Selection of context components					$\eta = 5\%$			$\eta = 2\%$		
	Age	Breast density	Assessment history	Opposite breast	Modality	FPR	FNR	Accuracy	FPR	FNR	Accuracy
1	X	X	X	X	X	61.0%	4.4%	0.45	64.2%	2.0%	0.42
2	X	X	X	X		72.1%	4.7%	0.35	81.0%	2.0%	0.27
3	X	X	X		X	65.6%	4.7%	0.40	77.7%	1.8%	0.30
4	X	X	X			72.7%	4.8%	0.34	80.5%	1.8%	0.28
5	X	X		X	X	64.0%	4.8%	0.42	73.2%	2.0%	0.34
6	X	X		X		74.4%	4.5%	0.33	88.9%	1.8%	0.20
7	X	X			X	67.9%	4.4%	0.38	79.9%	1.9%	0.28
8	X	X				76.1%	4.7%	0.31	82.0%	2.0%	0.26
9	X		X	X	X	62.8%	4.7%	0.43	67.4%	1.9%	0.40
10	X		X	X		74.8%	4.3%	0.32	79.6%	2.0%	0.29
11	X		X		X	69.6%	4.8%	0.37	73.7%	1.9%	0.33
12	X		X			74.0%	4.5%	0.33	83.1%	1.8%	0.26
13	X			X	X	66.6%	5.0%	0.39	75.3%	1.9%	0.32
14	X			X		76.2%	5.0%	0.31	90.9%	1.8%	0.18
15	X				X	72.2%	4.6%	0.35	80.5%	2.0%	0.28
16	X					76.7%	4.4%	0.31	82.4%	1.9%	0.26
17		X	X	X	X	63.1%	4.7%	0.43	67.2%	1.7%	0.40
18		X	X	X		73.4%	4.4%	0.33	84.4%	2.0%	0.24
19		X	X		X	68.0%	4.8%	0.38	84.9%	2.0%	0.24
20		X	X			75.0%	4.4%	0.32	85.6%	2.0%	0.23
21		X		X	X	69.3%	5.0%	0.37	76.1%	1.8%	0.32
22		X		X		76.8%	4.7%	0.30	91.0%	1.8%	0.18
23		X			X	73.2%	4.8%	0.34	86.1%	1.9%	0.23
24		X				78.0%	4.7%	0.29	83.8%	2.0%	0.25
25			X	X	X	65.0%	5.0%	0.41	69.8%	2.0%	0.37
26			X	X		77.2%	4.4%	0.30	81.4%	2.0%	0.27
27			X		X	71.2%	4.8%	0.36	86.2%	2.0%	0.23
28			X			75.7%	4.7%	0.31	87.2%	2.0%	0.22
29				X	X	70.4%	4.7%	0.36	81.9%	2.0%	0.26
30				X		78.2%	4.4%	0.29	83.8%	1.9%	0.25
31					X	73.9%	4.4%	0.33	87.4%	1.9%	0.22
32						-	-	-	-	-	-

probability that the estimation of $\sigma(x)$ has a large deviation from the true value tends to 0. For the context clusters with $\sigma(x)$ close to σ_η, the actions selected by the algorithm may not be the same as the optimal strategy, however, the probability of context arrivals in these clusters will tend to 0, as noticed by the condition $\Pr\{x : \sigma(x) = \sigma_\eta\} = 0$. The strategy of the learning algorithm tends to the optimal strategy except some context clusters whose arrival probability tends to 0.

Formally, we consider the per-period regret r_t at some sufficiently large time t. Then we can see that the per-period regret can be decomposed into two terms $r_t = r_{t,1} + r_{t,2}$, where $r_{t,1}$ is the regret caused by clusters that have a $\sigma(x)$ near the threshold σ_η, and $r_{t,2}$ is the regret caused by clusters that have a $\sigma(x)$ far from the threshold σ_η, but have a wrong estimation of $f(x)$. We denote by $\sigma_{\min}(C) = \min_{x \in C} \sigma(x)$ the minimum probability of being malignant for cluster C, and denote by $\sigma_{\max}(C) = \max_{x \in C} \sigma(x)$ the maximum probability of being malignant for cluster C. We define three types of clusters:

Type I cluster: the cluster at time t that has a $\sigma_{\max}(C)$ smaller than or equal to the threshold minus a small value, i.e., $\{C : C \in \mathcal{C}^t, \sigma_{\max}(C) \leq \sigma_\eta - bt^{-\alpha}\}$, where $b > 0$, $0 < \alpha < 1$ are parameters.

Type II cluster: the cluster at time t that has a $\sigma_{\min}(C)$ greater than or equal the threshold plus a small value, i.e., $\{C : C \in \mathcal{C}^t, \sigma_{\min}(C) \geq \sigma_\eta + bt^{-\alpha}\}$.

Type III cluster: the remaining clusters that have a $\sigma(x)$ near σ_η, i.e., $\{C : C \in \mathcal{C}^t, \sigma_{\min}(C) - bt^{-\alpha} < \sigma_\eta < \sigma_{\max}(C) + bt^{-\alpha}\}$.

Due to Berstein's inequality, we have that the estimation for the context arrival at a cluster C has the following property:

$$\Pr\{|\frac{M_C}{t} - f(C)| > b_1 t^{1-\alpha}\} \leq b_{11} e^{-t^\alpha},$$

where b_1 and b_{11} are positive constants. And the realized $\bar{\sigma}_C$ in a cluster C has the following property:

$$\Pr\{\bar{\sigma}_C > \sigma_{\max}(C) + b_2 t^{1-\alpha} \text{ or } \bar{\sigma}_C < \sigma_{\min}(C) - b_2 t^{1-\alpha}\} \leq b_{22} e^{-t^\alpha},$$

where b_2 and b_{22} are positive constants. We define the normal state as the event that the estimations of $f(x)$ and $\sigma(x)$ are accurate enough. The set of normal states are denoted by $\mathcal{N}_C = \{|\frac{M_C}{t} - f(C)| \leq b_1 t^{1-\alpha}, \sigma_{\min}(C) - b_2 t^{1-\alpha} \leq \bar{\sigma}_C \leq \sigma_{\max}(C) + b_2 t^{1-\alpha}\}$. And we denote the set of abnormal state by $\bar{\mathcal{N}}_C$, which is the complementary set of $\mathcal{N}_C$. The

probability of an abnormal state happens for one of the active cluster is bounded by $\sum_{C\in\mathcal{C}^t}\Pr\{\bar{\mathcal{N}}_C\}$.

Hence, we can see that the regret for $r_{t,1}$ is caused by Type III clusters, and can be bounded by

$$\begin{aligned} &r_{t,1} \leq \Pr\{x : x\in C, \sigma_{\min}(C) - bt^{-\alpha} < \sigma_\eta < \sigma_{\max}(C) \\ &+bt^{-\alpha}\} \leq \tfrac{K2^{(d-1)l}}{2^{dl}} \leq K2^{-l} \end{aligned},$$

where K is a constant, and the inequality is due to the covering property that the $d-1$ dimensional surface $\sigma(x)=\sigma_\eta$ and the Lipschitz condition that $\sigma_{\max}(C)-\sigma_{\min}(C)\leq L_1 2^{-l}$.

The regret $r_{t,2}$ can be bounded by the probability that an abnormal state $\sum_{C\in\mathcal{C}^t}\Pr\{\bar{\mathcal{N}}_C\}$ occurs. Hence, for time t, the regret can be bounded by

$$\begin{aligned} &r_t = r_{t,1} + r_{t,2} \\ &\leq \textstyle\sum_{C\in\mathcal{C}^t} K2^{-l} + b_{11}e^{-t^\alpha} + b_{22}e^{-t^\alpha} \leq O(t^{-1+g(d_X)}), \end{aligned}$$

where $g(d_X) = \frac{d_X+1/2+\sqrt{9+8d_X}/2}{d_X+3/2+\sqrt{9+8d_X}/2}$, and the worst case context arrival and partition is considered as in Appendix D. Hence, we obtain the regret up to time T:

$$R(T) = \sum_{t=1}^{T} r_t \leq O(T^{g(d_X)}).$$

Appendix C
Proof of Proposition 2

In order to show the equivalence of the two optimal strategies, we consider the weighted error of choosing different actions for context x. If the action $\pi(x)=1$ is chosen, then

$$\begin{aligned} &Ec(\pi(x)=1, s(x)) = (1-\sigma_\eta)Ec_1(\pi(x)=1, s(x)) \\ &= (1-\sigma_\eta)\sigma(x). \end{aligned} \tag{10}$$

If the action $\pi(x)=0$ is chosen, then

$$\begin{aligned} &Ec(\pi(x)=0, s(x)) = \sigma_\eta Ec_0(\pi(x)=0, s(x)) \\ &= \sigma_\eta(1-\sigma(x)). \end{aligned} \tag{11}$$

Hence, the optimal strategy $\pi^\dagger(x)$ satisfies:

$$\pi^\dagger(x) = \begin{cases} 1, \ if\ Ec(\pi(x)=1, s(x)) \leq Ec(\pi(x)=0, s(x)) \\ 0, \ otherwise \end{cases}. \tag{12}$$

By plugging (10) and (11) into (12), we have the optimal strategy $\pi^\dagger(x)$:

$$\pi^\dagger(x) = \begin{cases} 1, \ if\ \sigma(x) \geq \sigma_\eta \\ 0, \ otherwise \end{cases}. \tag{13}$$

Therefore, the proposition follows.

Appendix D
Proof of Theorem 2

To prove Theorem 2, we first introduce some important notions to characterize the properties of the cost. Let us define a_C^* as the best action corresponding to the context at the center of the subspace C. Let us also define $\mu_{x,a}$ as the expected weighted error, $\bar{\mu}_{C,a} = \max_{x\in C}\mu_{x,a}$, and $\mu_{C,a} = \min_{x\in C}\mu_{x,a}$. For a level l subspace C, the *suboptimal action set* is defined as $\mathcal{L}_C(B) = \{a : \bar{\mu}_{C,a_C^*} - \mu_{C,a} > BLd_X^{\alpha/2}2^{-l\alpha}\}$. We then can decompose the regret into three terms: the regret caused by virtual exploration $R_e(T)$, the regret caused by suboptimal arm selection $R_s(T)$, and the regret caused by near optimal arm selection $R_n(T)$. We first introduce three lemmas to show useful properties of the DRA algorithm.

Lemma 1. The active subspace level at time t can be at most $(\log_2 t)/p + 1$.

Proof: According to the context space partition process, we have $\sum_{j=1}^{l} 2^{pj} \leq t$, where $l+1$ is the maximum level at time t. Hence, the result follows.

Lemma 2. The regret caused by virtual exploration in one subspace up to time t is bounded by $2t^z \log t$.

Proof: Since the virtual exploration number can be bounded by $t^z \log t$ for each action, the result follows.

Lemma 3. If $B = \frac{2}{Ld_X^{\alpha/2}2^{-\alpha}} + 2$, and $2\alpha p < z < 1$, then the regret caused by suboptimal action selection in one subspace up to time t is bounded by $\frac{2\pi^2}{3}$.

Proof: Let $\mathcal{W}_C^t$ denote the event that the current phase is an exploitation phase in the context subspace C, and let $\mathcal{V}_C^t(a)$ be the event that the suboptimal action a is selected in at time t. Then, we have

$$\begin{aligned} &R_{C,s}(T) \leq \sum_{t=1}^{T}\sum_{a\in\mathcal{L}_C(B)} \Pr\{\mathcal{W}_C^t, \mathcal{V}_C^t(a)\} \\ &\leq \sum_{t=1}^{T}\sum_{a\in\mathcal{L}_C(B)} \Pr\{\bar{r}_{C,a} \geq \bar{\mu}_{C,a} + H_t, \mathcal{W}_C^t\} \\ &+\Pr\{\bar{r}_{C,a_C^*} \leq \bar{\mu}_{C,a_C^*} - H_t, \mathcal{W}_C^t\} + \Pr\{\bar{r}_{C,a} \geq \bar{r}_{C,a_C^*}, \\ &\bar{r}_{C,a} < \bar{\mu}_{C,a} + H_t, \bar{r}_{C,a_C^*} > \bar{\mu}_{C,a_C^*} + H_t, \mathcal{W}_C^t\} \end{aligned}, \tag{14}$$

where $H_t = t^{-z/2}$, $z \geq 2\alpha/p$. The third term on the right hand side of (14) is 0. Hence, we can bound the regret by

$$\begin{aligned} &R_{C,s}(T) \leq \sum_{t=1}^{T}\sum_{a\in\mathcal{L}_C(B)} \Pr\{\bar{r}_{C,a} \geq E[\bar{r}_{C,a}] + Ld_X^{\alpha/2}2^{-l\alpha}\} \\ &+\Pr\{\bar{r}_{C,a_C^*} \leq E[\bar{r}_{C,a_C^*}] - Ld_X^{\alpha/2}2^{-l\alpha}\} \leq \sum_{t=1}^{T} 4t^{-2} \leq \tfrac{4\pi^2}{3} \end{aligned}. \tag{15}$$

We can see that the highest level of subspaces is at most $1+\log_{2^{p+d_X}} T$. Then the maximum number of subspaces is bounded by $2^{2d_X}T^{\frac{d_X}{d_X+p}}$.

Therefore, according to Lemmas 2, we can bound the exploration regret by

$$R_e(T) \leq 2^{2d_X+1}T^{\frac{d_X}{d_X+p}}T^z \log T. \tag{16}$$

Accord to Lemma 3, we can bound the suboptimal regret by

$$R_s(T) \leq \frac{2^{2d_X+1}\pi^2 T^{\frac{d_X}{d_X+p}}}{3}. \tag{17}$$

We can also bound the near optimal regret by

$$\begin{aligned} R_n(T) &\leq \sum_{l=0}^{1+\log_{2^{p+d_X}} T} BLd_X^{\alpha/2}2^{-l\alpha} \\ &\leq BLd_X^{\alpha/2}2^{2(d_X+p-\alpha)}T^{\frac{d_X+p-\alpha}{d_X+p}}. \end{aligned} \tag{18}$$

Therefore, the result follows by setting $z = 2\alpha/p$, and $p = \frac{3\alpha+\sqrt{9\alpha^2+8\alpha d_X}}{2}$.

APPENDIX E
PROOF OF THEOREM 3 AND 4

For Theorems 3 and 4, the clinical regret can be calculated by another deviating regret term. For Theorem 3, this term can be bounded by $\sum_{t=1}^{T} \varepsilon = \varepsilon T$. For Theorem 4, this term can be bounded by $\sum_{t=1}^{T} \varepsilon_t \leq \frac{T^{1-\beta}}{1-\beta}$. Hence, Theorems 3 and 4 follow.

REFERENCES

[1] M. N. Gurcan, B. Sahiner, N. Petrick, H. P. Chan, E. A. Kazerooni, P. N. Cascade, et al., "Lung nodule detection on thoracic computed tomography images: preliminary evaluation of a computer-aided diagnosis system," *Medical Physics*, vol. 29, no. 11, pp. 2552-2558, 2002.

[2] J. Tang, R. M. Rangayyan, J. Xu, I. E. Naqa, and Y. Yang, "Computer-aided detection and diagnosis of breast cancer with mammography: recent advances," *IEEE Transactions on Information Technology in Biomedicine,* vol. 13, no. 2, pp. 236-251, 2009.

[3] K. Ganesan, U. Acharya, C. K. Chua, L. C. Min, K. Abraham, and K. Ng, "Computer-aided breast cancer detection using mammograms: a review," *IEEE Reviews in Biomedical Engineering*, vol. 6, pp. 77-98, 2013.

[4] M. F. Ganji and M. S. Abadeh, "A fuzzy classification system based on Ant Colony Optimization for diabetes disease diagnosis," *Expert Systems with Applications*, vol. 38, no. 12, pp. 14650-14659, 2011.

[5] J. G. Elmore, D. L. Miglioretti, L. M. Reisch, M. B. Barton, W. Kreuter, C. L. Christiansen, et al., "Screening mammograms by community radiologists: variability in false-positive rates," *Journal of the National Cancer Institute*, vol. 94, no. 18, pp. 1373-1380, 2002.

[6] M. A. Musen, B. Middleton, and R. A. Greenes, "Clinical decision-support systems," *Biomedical informatics*, Springer, London, pp. 643-674, 2014.

[7] K. Kerlikowske, P. A. Carney, B. Geller, M. T. Mandelson, S. H. Taplin, K. Malvin, et al., "Performance of screening mammography among women with and without a first-degree relative with breast cancer," *Annals of Internal Medicine*, vol. 133, pp. 855–863, 2000.

[8] " Cancer Facts & Figures 2014," American Cancer Society, 2014.

[9] R. Seigel, D. Naishadham, and A. Jemal, "Cancer Statistics," American Cancer Society, 2013.

[10] "ACR BI-RADS breast imaging and reporting data system: breast imaging Atlas 5th Edition," American College of Radiology, 2013.

[11] L. Liberman and J. H. Menell, "Breast imaging reporting and data system (BI-RADS)," *Radiologic Clinics of North America*, vol. 40, no. 3, pp. 409-430, 2002.

[12] M. L. Giger, Z. Huo, M. A. Kupinski, and C. J. Vyborny, "Computer-aided diagnosis in mammography," *Handbook of medical imaging*, vol. 2, pp. 915-1004, 2000.

[13] D. Gur, J. H. Sumkin, H. E. Rockette, M. Ganott, C. Hakim, L. Hardesty, et al., "Changes in breast cancer detection and mammography recall rates after the introduction of a computer-aided detection system," *Journal of the National Cancer Institute*, vol. 96, no. 3, pp. 185-190, 2004.

[14] D. Tsai, H. Fujita, K. Horita, T. Endo, C. Kido, and T. Ishigaki, "Classification of breast tumors in mammograms using a neural network: utilization of selected features," In *Proc. IEEE International Joint Conference on Neural Networks*, pp. 967-970, 1993.

[15] F. Schnorrenberg, C. S. Pattichis, K. C. Kyriacou, and C. N. Schizas, "Computer-aided detection of breast cancer nuclei," *IEEE Transactions on Information Technology in Biomedicine*, vol. 1, no. 2, pp. 128-140, 1997.

[16] W. E. Polakowski, D. A. Cournoyer, S. K. Rogers, M. P. DeSimio, D. W. Ruck, J. W. Hoffmeister, et al., "Computer-aided breast cancer detection and diagnosis of masses using difference of Gaussians and derivative-based feature saliency," *IEEE Transactions on Medical Imaging*, vol. 16, no. 6, pp. 811-819, 1997.

[17] N. R. Mudigonda, R. M. Rangayyan, and J. L. Desautels, "Gradient and texture analysis for the classification of mammographic masses," *IEEE Transactions on Medical Imaging*, vol. 19, no. 10, pp. 1032-1043, 2000.

[18] Y.-H. Chang, L. A. Hardesty, C. M. Hakim, T. S. Chang, B. Zheng, W. F. Good, et al., "Knowledge-based computer-aided detection of masses on digitized mammograms: A preliminary assessment," *Medical physics*, vol. 28, no. 4, pp. 455-461, 2001.

[19] C. M. Kocur, S. K. Rogers, L. R. Myers, T. Burns, M. Kabrisky, J. W. Hoffmeister, et al., "Using neural networks to select wavelet features for breast cancer diagnosis," *IEEE Engineering in Medicine and Biology Magazine*, vol. 15, no. 3, pp. 95-102, 1996.

[20] A. Urmaliya and J. Singhai, "Sequential minimal optimization for support vector machine with feature selection in breast cancer diagnosis," In *IEEE Second International Conference on Image Information Processing (ICIIP),* pp. 481-486, 2013.

[21] M. Sameti, R. K. Ward, J. Morgan-Parkes, and B. Palcic, "Image feature extraction in the last screening mammograms prior to detection of breast cancer," *IEEE Journal Selected Topics in Signal Processing*, vol. 3, no. 1, pp. 46-52, 2009.

[22] M. Donelli, I. J. Craddock, D. Gibbins, and M. Sarafianou, "A three-dimensional time domain microwave imaging method for breast cancer detection based on an evolutionary algorithm," *Progress In Electromagnetics Research M*, vol. 18, pp. 179-195, 2011.

[23] P. S. Pawar and D. R. Patil, "Breast Cancer Detection Using Neural Network Models," In *IEEE International Conference on Communication Systems and Network Technologies (CSNT),* pp. 568-572, 2013.

[24] S. Timp, C. Varela, and N. Karssemeijer, "Computer-aided diagnosis with temporal analysis to improve radiologists' interpretation of mammographic mass lesions," *IEEE Transactions on Information Technology in Biomedicine*, vol. 14, no. 3, pp. 803-808, 2010.

[25] S. Ghosh, S. Mondal, and B. Ghosh, "A comparative study of breast cancer detection based on SVM and MLP BPN classifier," In *IEEE First International Conference on Automation, Control, Energy and Systems (ACES),* pp. 1-4, 2014.

[26] S. Singh and K. Bovis, "An evaluation of contrast enhancement techniques for mammographic breast masses," *IEEE Transactions on Information Technology in Biomedicine*, vol. 9, no. 1, pp. 109-119, 2005.

[27] K. Panetta, Y. Zhou, S. Agaian, and H. Jia, "Nonlinear unsharp masking for mammogram enhancement," *IEEE Transactions on Information Technology in Biomedicine,* vol. 15, no. 6, pp. 918-928, 2011.

[28] A. N. Karahaliou, I. S. Boniatis, S. G. Skiadopoulos, F. N. Sakellaropoulos, N. S. Arikidis, E. A. Likaki, et al., "Breast cancer diagnosis: analyzing texture of tissue surrounding microcalcifications, *IEEE Transactions on Information Technology in Biomedicine*, vol. 12, no. 6, pp. 731-738, 2008.

[29] G. Bozza, M. Brignone, and M. Pastorino, "Application of the no-sampling linear sampling method to breast cancer detection," *IEEE Transactions on Biomedical Engineering*, vol. 57, no. 10, pp. 2525-2534, 2010.

[30] A. Keles, A. Keles, and U. Yavuz, "Expert system based on neuro-fuzzy rules for diagnosis breast cancer," *Expert Systems with Applications*, vol. 38, no. 5, pp. 5719-5726, 2011.

[31] W. Hsu, S. Han, C. Arnold, A. Bui, D. R. Enzmann, "RadQA: A data-driven approach to assessing the accuracy and utility of radiologic interpretations," *SIIM Annual Meeting*, 2014.

[32] K. Armstrong, E. A. Handorf, J. Chen, and M. N. B. Demeter, "Breast cancer risk prediction and mammography biopsy decisions: a model-based study," *American Journal of Preventive Medicine*, vol. 44, no. 1, pp. 15-22, 2013.

[33] S. Venkataraman, P. J. Slanetz, "Breast imaging: mammography and ultrasonography," *UptoDate*, 2014.

[34] C. Wiratkapun, W. Bunyapaiboonsri, B. Wibulpolprasert, and P. Lertsithichai, "Biopsy rate and positive predictive value for breast cancer in BI-RADS category 4 breast lesions," *Medical Journal of the Medical Association of Thailand*, vol. 93, no. 7, pp. 830, 2010.

[35] C. I. Flowers, C. O'Donoghue, D. Moore, A. Goss, D. Kim, J.-H. Kim, et al., "Reducing false-positive biopsies: a pilot study to reduce benign biopsy rates for BI-RADS 4A/B assessments through testing risk stratification and new thresholds for intervention," *Breast Cancer Research and Treatment*, vol. 139, no. 3, pp. 769-777, 2013.

[36] T. Ayer, O. Alagoz, and N. K. Stout, "OR forum-a POMDP approach to personalize mammography screening decisions," *Operations Research*, vol. 60, no. 5, pp. 1019-1034, 2012.

[37] S. W. Fletcher, "Risk prediction models for breast cancer screening," *UptoDate*, 2014.

[38] A. Slivkins, "Contextual bandits with similarity information," arXiv preprint arXiv:0907.3986, 2009.

[39] J. Langford and T. Zhang, "The epoch-greedy algorithm for contextual multi-armed bandits," *Advances in Neural Information Processing Systems*, vol. 20, pp. 1096-1103, 2007.

[40] T. Lu, D. Pal, and M Pal, "Contextual multi-armed bandits," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 485-492, 2010.