

Reinforcement Learning in Healthcare: A Survey

Chao Yu, Jiming Liu, *Fellow, IEEE*, and Shamim Nemati

Abstract—As a subfield of machine learning, *reinforcement learning* (RL) aims at empowering one’s capabilities in behavioural decision making by using interaction experience with the world and an evaluative feedback. Unlike traditional supervised learning methods that usually rely on one-shot, exhaustive and supervised reward signals, RL tackles with sequential decision making problems with sampled, evaluative and delayed feedback simultaneously. Such distinctive features make RL technique a suitable candidate for developing powerful solutions in a variety of healthcare domains, where diagnosing decisions or treatment regimes are usually characterized by a prolonged and sequential procedure. This survey discusses the broad applications of RL techniques in healthcare domains, in order to provide the research community with systematic understanding of theoretical foundations, enabling methods and techniques, existing challenges, and new insights of this emerging paradigm. By first briefly examining theoretical foundations and key techniques in RL research from efficient and representational directions, we then provide an overview of RL applications in healthcare domains ranging from dynamic treatment regimes in chronic diseases and critical care, automated medical diagnosis from both unstructured and structured clinical data, as well as many other control or scheduling domains that have infiltrated many aspects of a healthcare system. Finally, we summarize the challenges and open issues in current research, and point out some potential solutions and directions for future research.

Index Terms—Reinforcement Learning, Healthcare, Dynamic Treatment Regimes, Critical Care, Chronic Disease, Automated Diagnosis.

I. INTRODUCTION

Driven by the increasing availability of massive multimodality data, and developed computational models and algorithms, the role of AI techniques in healthcare has grown rapidly in the past decade [1], [2], [3], [4]. This emerging trend has promoted increasing interests in the proposal of advanced data analytical methods and machine learning approaches in a variety of healthcare applications [5], [6], [7], [8], [9]. As a subfield in machine learning, *reinforcement learning* (RL) has achieved tremendous theoretical and technical achievements in generalization, representation and efficiency in recent years, leading to its increasing applicability to real-life problems in playing games, robotics control, financial and business management, autonomous driving, natural language processing, computer vision, biological data analysis, and art creation, just to name a few [10], [11], [12], [13], [14].

In RL problems, an agent chooses an action at each time step based on its current state, and receives an evaluative

Chao Yu is with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China. (Email: yuchao3@mail.sysu.edu.cn). Jiming Liu is with the Computer Science Department, Hong Kong Baptist University, Kowloon Tong, Hong Kong. (Email: jiming@Comp.HKBU.Edu.HK). Shamim Nemati is with the Department of Biomedical Informatics, UC San Diego, La Jolla, CA, USA. (Email: snemati@health.ucsd.edu).

feedback and the new state from the environment. The goal of the agent is to learn an optimal policy (i.e., a mapping from the states to the actions) that maximizes the accumulated reward

it receives over time. Therefore, agents in RL do not receive direct instructions regarding which action they should take, instead they must learn which actions are the best through trial-and-error interactions with the environment. This adaptive closed-loop feature renders RL distinct from traditional supervised learning methods for regression or classification, in which a list of correct labels must be provided, or from unsupervised learning approaches to dimensionality reduction or density estimation, which aim at finding hidden structures in a collection of example data [11]. Moreover, in comparison with other traditional control-based methods, RL does not require a well-represented mathematical model of the environment, but develops a control policy directly from experience to predict states and rewards during a learning procedure. Since the design of RL is letting an agent controller interact with the system, unknown and time-varying dynamics as well as changing performance requirements can be naturally accounted for by the controller [15]. Lastly, RL is uniquely suited to systems with inherent time delays, in which decisions are performed without immediate knowledge of effectiveness, but evaluated by a long-term future reward.

The above features naturally make RL an attractive solution to constructing efficient policies in various healthcare domains, where the decision making process is usually characterized by a prolonged period or sequential procedure [16]. Typically, a medical or clinical treatment regime is composed of a sequence of decision to determine the course of decisions such as treatment type, drug dosage, or re-examination timing at a time point according to the current health status and prior treatment history of an individual patient, with a goal of promoting the patient’s long-term benefits. Unlike the common procedure in traditional randomized controlled trials that derive treatment regimes from the average population response, RL can be tailored for achieving precise treatment for individual patients who may possess high heterogeneity in response to the treatment due to variety in disease severity, personal characteristics and drug sensitivity. Moreover, RL is able to find optimal policies using only previous experiences, without requiring any prior knowledge about the mathematical model of the biological systems. This makes RL more appealing than many existing control-based approaches in healthcare domains since it could be usually difficult or even impossible to build an accurate model for the complex human body system and the responses to administered treatments, due to nonlinear, varying and delayed interaction between treatments and human bodies.

Thus far, a plethora of theoretical or experimental studies have applied RL techniques and models in a variety of healthcare domains, achieving performance exceeding that of

Acronym	Description
AC	<i>Actor-Critic</i>
A3C	<i>Asynchronous Advantage Actor Critic</i>
BRL	<i>Batch Reinforcement Learning</i>
DDPG	<i>Deep Deterministic Policy Gradient</i>
DRL	<i>Deep Reinforcement Learning</i>
DP	<i>Dynamic programming</i>
DQN	<i>Deep Q Network</i>
DDQN	<i>Dueling DQN</i>
DDDQN	<i>Double Dueling DQN</i>
FQI-SVG/ERT	<i>Fitted Q Iteration with Support Vector Regression/Extremely Randomized Trees</i>
GAN	<i>Generative Adversarial Net</i>
HRL	<i>Hierarchical Reinforcement Learning</i>
IRL	<i>Inverse Reinforcement Learning</i>
LSPI	<i>Least-Squares Policy Iteration</i>
MDP	<i>Markov Decision Process</i>
MC	<i>Monte Carlo</i>
NAC	<i>Natural Actor Critic</i>
PAC	<i>Probably Approximately Correct</i>
PI	<i>Policy Iteration</i>
PS	<i>Policy Search</i>
POMDP	<i>Partially Observed Markov Decision Process</i>
PORL	<i>Partially Observed Reinforcement Learning</i>
PPO	<i>Proximal Policy Optimization</i>
PRL	<i>Preference-based Reinforcement Learning</i>
RRL	<i>Relational Reinforcement Learning</i>
TD	<i>Temporal Difference</i>
TRL	<i>Transfer Reinforcement Learning</i>
TRPO	<i>Trust Region Policy Optimization</i>
VI	<i>Value Iteration</i>

alternative techniques in many cases. This survey aims at providing an overview of such successful RL applications, covering adaptive treatment regimes in chronic diseases and critical care, automated clinical diagnosis, as well as many other healthcare domains such as clinical resource allocation/scheduling and optimal process control. We also discuss the challenges, open issues and future directions of research necessary to advance further successful applications of RL in healthcare. By this, we hope this survey can provide the research community with systematic understanding of foundations, enabling methods and techniques, challenges, and new insights of this emerging paradigm. Section II provides a structured summarization of the theoretical foundations and key techniques in RL research from two main directions: *efficient* directions that mainly aim at improving learning efficiency by making best use of past experience or knowledge, and *representational* directions that focus on constructive or relational representation problems in RL. Then, Section III gives a global picture of application domains of RL in healthcare, each of which is discussed in more detail in the following sections. Section IV discusses dynamic treatment regimes in both chronic disease and critical care, and Section V describes automated medical diagnosis using either structured or unstructured medical data. In addition, VI talks about other more broad application domains including health resources allocation and scheduling, optimal process control, drug discovery and development, as well as health management. Section VII describes several challenges and

open issues in current research. Finally, Section VIII discusses potential directions that are necessary in the future research. For convenience, Tables I and II summarize the main acronyms in RL and healthcare domains, respectively.

TABLE II
SUMMARY OF ABBREVIATIONS IN HEALTHCARE

Acronym	Description
AMP	<i>Anemia Management Protocol</i>
BCSC	<i>Breast Cancer Surveillance Consortium</i>
CATIE	<i>Clinical Antipsychotic Trials of Intervention Effectiveness</i>
CBASP	<i>Cognitive Behavioral-Analysis System of Psychotherapy</i>
CIBMTR	<i>Center for International Blood and Marrow Transplant Research</i>
CT	<i>Computed Tomography</i>
DTRs	<i>Dynamic Treatment Regimes</i>
EEG	<i>Electroencephalograph</i>
ESAs	<i>Erythropoiesis-Stimulating Agents</i>
EPO	<i>Endogenous Erythropoietin</i>
FES	<i>Functional Electrical Stimulation</i>
HGB	<i>Hemoglobin</i>
ICU	<i>Intensive Care Unit</i>
MAP	<i>Mean Arterial Pressure</i>
MDD	<i>Major Depressive Disorder</i>
MIMIC	<i>Multiparameter Intelligent Monitoring in Intensive Care</i>
MRI	<i>Magnetic Resonance Images</i>
NSCLC	<i>Non-small Cell Lung Cancer</i>
ODE	<i>Ordinary Difference Equations</i>
PK/PD	<i>PharmacoKinetic/PharmacoDynamic</i>
SC	<i>Symptom Checking</i>
SMARTs	<i>Sequential Multiple Assignment Randomized Trials</i>
STAR*D	<i>Sequenced Treatment Alternatives to Relieve Depression</i>
STI	<i>Structured Treatment Interruption</i>
SUD	<i>Sub-stance Use Disorder</i>
TREC-CDS	<i>Text REtrieval Conference-Clinical Decision Support</i>
UI	<i>Ultrasound Images</i>

II. THEORETICAL FOUNDATIONS AND KEY TECHNIQUES IN RL

This section serves as a brief introduction to the theoretical models, basic solutions and advanced techniques in RL. The goal is to provide a quick overview of what constitutes the main components of RL methods. Some fundamental concepts and major theoretical problems are also clarified. Subsection II-A first discusses the general decision making framework for RL, and its model-based and model-free solutions. Then, Subsection II-B describes some advanced RL techniques from perspectives of facilitating learning efficiency and enriching constructive or relational representation capabilities of RL approaches. Fig. 1 provides a diagram outlining the main components and sub-field research topics in RL.

A. Theoretical Foundations of RL

RL enables an agent to learn effective strategies in sequential decision making problems by trial-and-error interactions with its environment [14]. The *Markov decision process (MDP)*, which has a long history in the research of *theoretic decision making* in stochastic settings, has been used as a general framework to formalize an RL problem. Approaching an MDP can be conducted in various forms

depending on what information of the targeted problem can be specified a priori.

1) *MDP*: Formally, an MDP can be defined by a 5-tuple $M = (S, A, P, R, \gamma)$, where S is a finite *state* space, and $s_t \in S$ denotes the state of an agent at time t ; A is a set of *actions* available to the agent, and $a_t \in A$ denotes the action that the agent performs at time t ; $P(s, a, s') : S \times A \times S \rightarrow [0, 1]$ is a Markovian *transition function* when the agent transits from state s to state s' after taking action a ; $R : S \times A \rightarrow \mathbb{R}$ is a *reward function* that returns the immediate reward $R(s, a)$ to

that $B^\pi V^\pi = V^\pi$ in the limit. The goal of an MDP problem is to compute an *optimal policy* π^* such that $V^{\pi^*}(s) \geq V^\pi(s)$ for every policy π and every state $s \in S$. To involve the action information, Q-value is used to represent the optimal value of each state-action pair by Equation 3.

$$Q^*(s, a) = \mathcal{R}(s, a) + \gamma \sum_{s' \in S} \mathcal{P}(s, a, s') \max_{a' \in A} Q(s', a') \quad (3)$$

2) *Basic Solutions and Challenging Issues*: Many solution techniques are available to compute an optimal policy for a

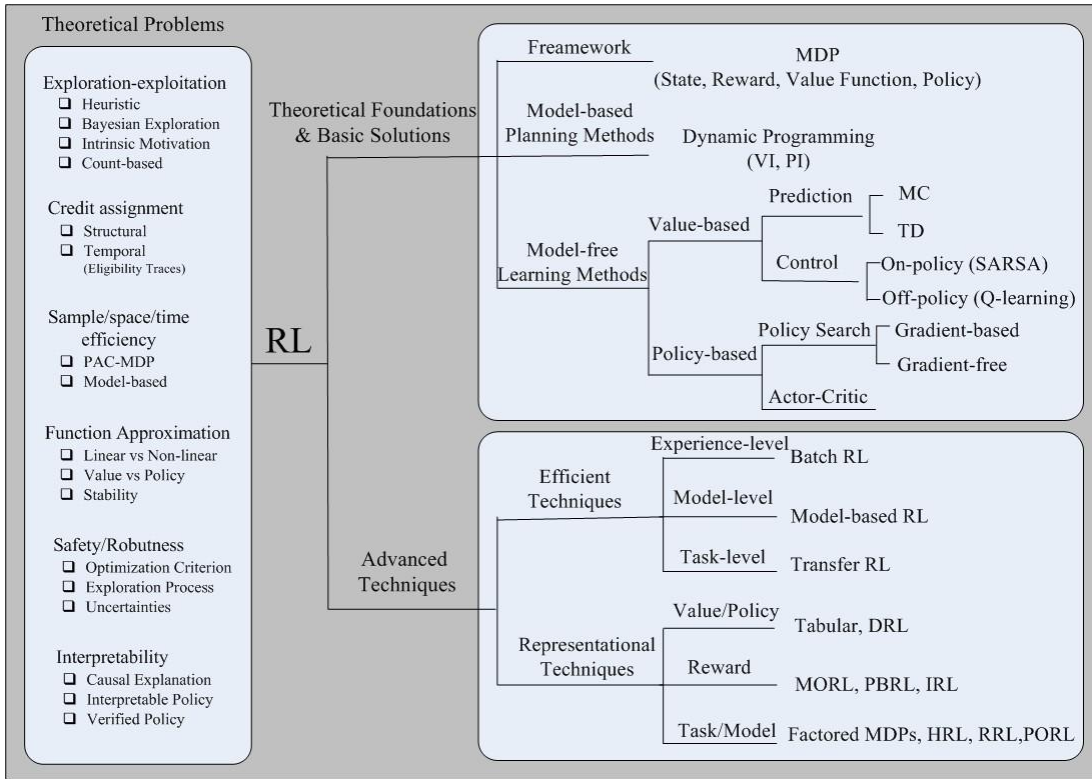


Fig. 1. The summarization of theoretical foundations, basic solutions, challenging issues and advanced techniques in RL.

the agent after taking action a in state s ; and $\gamma \in [0, 1]$ is a *discount factor*.

An agent's *policy* $\pi : S \times A \rightarrow [0, 1]$ is a probability distribution that maps an action $a \in A$ to a state $s \in S$. When given an MDP and a policy π , the expected reward of following this policy when starting in state s , $V^\pi(s)$, can be defined as follows:

$$V^\pi(s) \triangleq E_\pi \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, \pi(s_t)) | s_0 = s \right] \quad (1)$$

The value function can also be defined recursively using the *Bellman operator* B^π :

$$B^\pi V^\pi(s) \triangleq \mathcal{R}(s, \pi(s)) + \gamma \sum_{s' \in S} \mathcal{P}(s, \pi(s), s') V^\pi(s') \quad (2)$$

B^π

Since the *Bellman operator* B^π is a contraction mapping of value function V , there exists a fixed point of value V^π such

given MDP. Broadly, these techniques can be categorized as *model-based* or *model-free* methods, based on whether a complete knowledge of the MDP model can be specified a priori. Model-based methods, also referred to as *planning* methods, require a complete description of the model in terms of the transition and reward functions, while model-free methods, also referred to as *learning* methods, learn an optimal policy simply based on received observations and rewards.

Dynamic programming (DP) [17] is a collection of *modelbased* techniques to compute an optimal policy given a complete description of an MDP model. DP includes two main different approaches: *Value Iteration* (VI) and *Policy Iteration*

$$Q_{t+1}(s, a) = \mathcal{R}(s, a) + \gamma \sum_{s' \in S} \mathcal{P}(s, a, s') \max_{a' \in A} Q_t(s', a') \quad (4)$$

(PI). VI specifies the optimal policy in terms of value function $Q^*(s, a)$ by iterating the *Bellman updating* as follows:

For each iteration, the value function of every state s is updated one step further into the future based on the current estimate. The concept of updating an estimate based on the basis of other estimates is often referred to as *bootstrapping*. The value function is updated until the difference between two iterations, Q_t and Q_{t+1} , is less than a small threshold. The optimal policy is then derived using $\pi^*(s) = \operatorname{argmax}_{a \in A} Q^*$. Unlike VI, PI learns the policy directly. It starts with an initial random policy π , and iteratively updates the policy by first computing the associated value function Q^π (*policy evaluation* or *prediction*) and then improving the policy using $\pi(s) = \operatorname{argmax}_{a \in A} Q(s,a)$ (*policy improvement* or *control*).

Despite being mathematically sound, DP methods require a complete and accurate description of the environment model, which is unrealistic in most applications. When a model of the problem is not available, the problem can then be solved by using *direct RL* methods, in which an agent learns its optimal policy while interacting with the environment. *Monte Carlo* (MC) methods and *Temporal difference* (TD) methods are two main such methods, with the difference of using episode-by-episode update in MC or step-by-step update in TD. Denote

$$\mathcal{R}_t^{(n)} = \mathcal{R}_{t+1} + \gamma \mathcal{R}_{t+2} + \dots + \gamma^{n-1} \mathcal{R}_{t+n} + \gamma^n V_t(s_{t+n})$$

step return at time t , then the general n -step update rule in TD methods is defined by $\Delta V_t(s_t) = \alpha[\mathcal{R}_t^{(n)} - V_t(s_t)]$, in which $\alpha \in (0,1]$ is an appropriate learning rate controlling the contribution of the new experience to the current estimate. MC methods then can be considered as an extreme case of TD methods when the update is conducted after the whole episode of steps. In spite of having higher complexity in analyzing the efficiency and speed of convergence, TD methods usually require less memory for estimates and less computation, thus are easier to implement.

If the value function of a policy π is estimated by using samples that are generated by strictly following this policy, the RL algorithm is called *on-policy*, while *off-policy* algorithms can learn the value of a policy that is different from the one being followed. One of the most important and widely used RL approach is Q-learning [18], which is an *off-policy* TD algorithm. Its one-step updating rule is given by Equation 5,

$$Q_{t+1}(s,a) = Q_t(s,a) + \alpha_t [R(s,a) + \gamma \max_{a^0} Q_t(s^0, a^0) - Q_t(s,a)] \quad (5)$$

where $\alpha \in (0,1]$ is an appropriate learning rate which controls the contribution of the new experience to the current estimate.

Likewise, the SARSA algorithm [19] is an representation for *on-policy* TD approaches given by Equation 6:

$$Q_{t+1}(s,a) = Q_t(s,a) + \alpha_t [R(s,a) + \gamma Q_t(s^0, \pi(s^0)) - Q_t(s,a)] \quad (6)$$

The idea is that each experienced sample brings the current estimate $Q(s,a)$ closer to the optimal value $Q^*(s,a)$. Qlearning starts with an initial estimate for each state-action pair. When an action a is taken in state s , resulting in the next state s^0 , the corresponding Q-value $Q(s,a)$ is updated with a combination of its current value and the TD error ($R(s,a) + \gamma \max_{a^0} Q_t(s^0, a^0) - Q_t(s,a)$ for Q-learning, or $R(s,a) + \gamma Q_t(s^0, \pi(s^0)) - Q_t(s,a)$ for SARSA). The TD error is the difference between the current estimate $Q(s,a)$ and the expected discounted return based on the experienced sample. The Q value of each state-action pair is stored in a table for a discrete state-action space. It has been proved that this tabular Q-learning converges to the optimal $Q^*(s,a)$ w.p.1 when all state-action pairs are visited infinitely often and an appropriate exploration strategy and learning rate are chosen [18].

Besides the above *value-function* based methods that maintain a value function whereby a policy can be derived, *direct policy-search* (PS) algorithms [20] try to estimate the policy directly without representing a value function explicitly, whereas the *actor-critic* (AC) methods [21] keep separate, explicit representations of both value functions and policies.

In AC methods, the *actor* is the policy to select actions, and the *critic* is an estimated value function to criticize the actions chosen by the actor. After each action execution, the critic evaluates the performance of action using the TD error. The advantages of AC methods include that they are more appealing in dealing with large scale or even continuous actions and learning stochastic policies, and more easier in integrating domain specific constraints on policies.

In order to learn optimal policies, an RL agent should make a balance between exploiting the knowledge obtained so far by acting optimally, and exploring the unknown space in order to find new efficient actions. Such an *exploration-exploitation trade-off* dilemma is one of the most fundamental theoretical issues in RL, since an effective *exploration strategy* enables the agent to make an elegant balance between these two processes by choosing explorative actions only when this behavior can potentially bring a higher expected return. A large amount of effort has been devoted to this issue in the traditional RL community, proposing a wealth of exploration strategies including simple heuristics such as ϵ -greedy and Boltzmann exploration, Bayesian learning [22], [23], countbased methods with *Probably Approximately Correct* (PAC) guarantees [24], [25], as well as more expressive methods of intrinsic motivation such as novelty, curiosity and surprise [26]. For example, the ϵ -greedy strategy selects the greedy action, $\operatorname{argmax}_a Q_t(s,a)$, with a high probability, and, occasionally, with a small probability selects an action uniformly at random. This ensures that all actions and their effects are experienced. The ϵ -greedy exploration policy can be given by Equation 7.

$$\pi(a) = \begin{cases} 0 & 1 - \varepsilon \text{ if } a = \operatorname{argmax}_{a_0} \\ Q(s,a), & \\ \varepsilon & \text{otherwise.} \end{cases} \quad (7)$$

where $\varepsilon \in [0,1]$ is an exploration rate.

Other fundamental issues in RL research include but are not limited to the credit assignment problem [14], [27], the sample/space/time complexity [28], [29], function approximation [30], [31], safety [32], [33], robustness [34], [35], and interpretability [36], [37]. A more comprehensive and in-depth review on these issues can be found in [38], and more recently in [12], [14].

B. Key Techniques in RL

This section discusses some key techniques used in contemporary RL, most of which can be understood in the light of the framework and solutions defined in the section ahead, yet these new techniques emphasize more sophisticated use of samples, models of the world and learned knowledge of previous tasks for efficiency purpose, as well as what should be represented and how things should be represented during an RL problem. Note that the classification of these two kinds of techniques are not mutually exclusive, which means that some representation techniques are also used for improving the learning efficiency, and vice versa.

1) *Efficient Techniques*: The purpose of using efficient techniques is to improve the learning performance in terms of, for example, convergence ratio, sample efficient, computation cost or generalization capabilities of an RL method. This improvement can be achieved by using different levels of knowledge: the *Experience-level* techniques focus on utilizing the past experience for more stable and data-efficient learning; the *Model-level* techniques focus on building and planning over a model of the environment in order to improve sample efficiency; while the *Task-level* techniques aim at generalizing the learning experience from past tasks to new relevant ones.

a) *Experience-level*: In traditional pure on-line TD learning methods such as Q-learning and SARSA, an agent immediately conducts a DP-like update of the value functions every step interacting with the environment and then disregards the experienced state transition tuple afterwards. In spite of guaranteed convergence and great success in solving simple toy problems, this kind of *local updates* poses several severe performance problems when applied to more realistic systems with larger and possibly continuous settings. Since each experience tuple is used only for one update and then forgotten immediately, a larger number of samples are required to enable an optimal solution, causing the so called *exploration overhead* problem. Moreover, it has been shown that directly combining function approximation methods with pure on-line TD methods can cause instable or even diverged performance [30], [31]. These inefficiency and instability

problems become even more pronounced in real environments, particularly in healthcare systems, where physical interactions between patients and environments call for more efficient sampling and stable learning methods.

The *Experience-level* techniques focus on how to make the best of the past learning experience for more stable and efficient learning, and are the major driving force behind the proposal of modern *Batch RL (BRL)* [39]. In BRL, two basic techniques are used: storing the experience in a buffer and reusing it as if it were new (the idea of *experience replay* for addressing the inefficiency problem), and separating the DP step from the function approximation step by using a supervised learning to fit the function approximator over the sampled experience (the idea of *fitting* for addressing the instability problem). There are several famous BRL approaches in the literature, such as the non-linear approximator cases of *Neural Fitted Q Iteration (NFQI)* [40]), the Tree-based FQI [41], and robust linear approximation techniques for policy learning such as *Least-Squares Policy Iteration (LSPI)* [42]). As will be discovered later, these BRL methods have enjoyed wide and successful applications in clinical decision makings, due to their promise in greatly improving learning speed and approximation accuracy, particularly from limited amounts of clinical data.

b) *Model-level*: Unlike *Experience-level* techniques that emphasize the efficient use of experience tuples, the *Model-level* techniques try to build a model of the environment (in terms of the transition and reward functions) and then derive optimal policies from the environment model when it is approximately correct. This kind of *model-based RL (MRL)* approaches is rather different from the model-free RL methods such as TD methods or MC methods that directly estimate value functions without building a model of the environment [43]. Using some advanced exploration strategies and planning methods such as DP or *Monte Carlo Tree Search (MCTS)* [44], *MRL* methods are usually able to learn an accurate model quickly and then use this model to plan multi-step actions. Therefore, *MRL* methods normally have better sample efficiency than model-free methods [28].

c) *Task-level*: A higher task-level of efficient approaches focuses on the development of methods to transfer knowledge from a set of source tasks to a target task. *Transfer RL (TRL)* uses the transferred knowledge to significantly improve the learning performance in the target task, e.g., by reducing the samples needed for a nearly optimal performance, or increasing the final convergence level [45]. Taylor and Stone [46] provided a thorough review on TRL approaches by five transfer dimensions: how the source task and target task may differ (e.g., in terms of action, state, reward or transition functions), how to select the source task (e.g., all previously seen tasks, or only one task specified by human or modified automatically), how to define task mappings (e.g., specified by human or learned from experience), what knowledge to transferred (from experience

instances to higher level of models or rules), and allowed RL methods (e.g., MRL, PS, or BRL).

2) *Representational Techniques*: Unlike traditional machine learning research that simply focuses on feature engineering for function approximation, representational techniques in RL can be in a broader perspective, paying attention to constructive or relational representation problems relevant not only to function approximation for state/action, policies and value functions, but also to more exogenous aspects regarding agents, tasks or models [12].

a) *Representation for Value Functions or Policies*: Many traditional RL algorithms have been mainly designed for problems with small discrete state and action spaces, which can be explicitly stored in tables. Despite the inherent challenges, applying these RL algorithms to continuous or highly dimensional domains would cause extra difficulties. A major aspect of representational techniques is to represent structures of policies and value functions in a more compact form for an efficient approximation of solutions and thus scaling up to larger domains. Broadly, three categories of approximation methods can be clarified [31]: *model-approximation* methods that approximate the model and compute the desired policy on this approximated model; *value-approximation* methods that approximate a value function whereby a policy can be inferred, and *policy-approximation* methods that search in policy space directly and update this policy to approximate the optimal policy, or keep separate, explicit representations of both value functions and policies.

The value functions or policies can be parameterized using either *linear* or *non-linear* function approximation presentations. Whereas the *linear function approximation* is better understood, simple to implement and usually has better convergence guarantees, it needs explicit knowledge about domain features, and also prohibits the representation of interactions between features. On the contrary, *non-linear function approximation* methods do not need for good informative features and usually obtain better accuracy and performance in practice, but with less convergence guarantees.

A notable success of RL in addressing real world complex problems is the recent integration of deep neural networks into RL [47], [48], fostering a new flourishing research area of *Deep RL* (DRL) [12]. A key factor in this success is that deep learning can automatically abstract and extract high-level features and semantic interpretation directly from the input data, avoiding complex feature engineering or delicate feature hand-crafting and selection for an individual task [49].

b) *Representation for Reward Functions*: In a general RL setting, the reward function is represented in the form of an evaluative scalar signal, which encodes a single objective for the learning agent. In spite of its wide applicability, this kind of quantifying reward functions has its limits inevitably. For example, real life problems usually involve two or more

objectives at the same time, each with its own associated reward signal. This has motivated the emerging research topic of *multi-objective RL* (MORL) [50], in which a policy must try to make a trade-off between distinct objectives in order to achieve a Pareto optimal solution. Moreover, it is often difficult or even impossible to obtain feedback signals that can be expressed in numerical rewards in some real-world domains. Instead, qualitative reward signals such as being better or higher may be readily available and thus can be directly used by the learner. *Preference-based RL* (PRL) [51] is a novel research direction combining RL and preference learning [52] to equip an RL agent with a capability to learn desired policies from qualitative feedback that is expressed by various ranking functions. Last but not the least, all the existing RL methods are grounded on an available feedback function, either in an explicitly numerical or a qualitative form. However, when such feedback information is not readily available or the reward function is difficult to specify manually, it is then necessary to consider an approach to RL whereby the reward function can be learned from a set of presumably optimal trajectories so that the reward is consistent with the observed behaviors. The problem of deriving a reward function from observed behavior is referred to as *Inverse RL* (IRL) [53], [54], which has received an increasingly high interest by researchers in the past few years. Numerous IRL methods have been proposed, including the Maximum Entropy IRL [55], the Apprenticeship Learning [56], nonlinear representations of the reward function using Gaussian processes [57], and Bayesian IRL [58].

c) *Representation for Tasks or Models*: Much recent research on RL has focused on representing the tasks or models in a compact way to facilitate construction of an efficient policy. Factored MDPs [59] are one of such approaches to representing large structured MDPs compactly, by using a *dynamic Bayesian network* (DBN) to represent the transition model among states that involve only some set of state variables, and the decomposition of global task reward to individual variables or small clusters of variables. This representation often allows an exponential reduction in the representation size of structured MDPs, but the complexity of exact solution algorithms for such MDPs also grows exponentially in the representation size. A large number of methods has been proposed to employ factored representation of MDP models for improving learning efficiency for either model-based [60], [61] or model-free RL problems [62]. A more challenging issues is how to learn this compact structure dynamically during on-line learning [63].

Besides the factored representation of states, a more general method is to decompose large complex tasks into smaller sets of sub-tasks, which can be solved separately. *Hierarchical RL* (HRL) [64] formalizes hierarchical methods that use abstract states or actions over a hierarchy of subtasks to decompose the original problem, potentially reducing its computational complexity. Hengst [65] discussed the various concepts and approaches in HRL, including algorithms that can

automatically learn the hierarchical structure from interactions with the domain. Unlike HRL that focuses on hierarchical decomposition of tasks, *Relational RL* (RRL) [66] provides a new representational paradigm to RL in worlds explicitly modeled in terms of objects and their relations. Using expressive data structures that represent the objects and relations in an explicit way, RRL aims at generalizing or facilitating learning over worlds with the same or different objects and relations. The main representation methods and techniques in RRL have been surveyed in detail in [66].

Last but not the least, *Partially Observable MDP* (POMDP) is widely adopted to represent models when the states are not fully observable, or the observations are noisy. Learning in POMDP, denoted as *Partially Observable RL* (PORL), can be rather difficult due to extra uncertainties caused by the mappings from observations to hidden states [67]. Since environmental states in many real life applications, notably in healthcare systems, are only partially observable, PORL then becomes a suitable technique to derive a meaningful policy in such realistic environments.

III. APPLICATIONS OF RL IN HEALTHCARE

On account of its unique features against traditional machine learning, statistic learning and control-based methods, RL-related models and approaches have been widely applied in healthcare domains since decades ago. The early days of focus has been devoted to the application of DP methods in various pharmacotherapeutic decision making problems using *pharmacokinetic/pharmacodynamic* (PK/PD) models [68], [69]. Hu *et al.*, [70] used POMDP to model drug infusion problem for the administration of anesthesia, and proposed efficient heuristics to compute suboptimal though useful treatment strategies. Schaeffer *et al.* [71] discussed the benefits and associated challenges of MDP modeling in the context of medical treatment, and reviewed several instances of medical applications of MDPs, such as spherocytosis treatment and breast cancer screening and treatment.

With the tremendous theoretical and technical achievements in generalization, representation and efficiency in recent years, RL approaches have been successfully applied in a number of healthcare domains to date. Broadly, these application domains can be categorized into three main types: dynamic treatment regimes in chronic disease or critical care, automated medical diagnosis, and other general domains such as health resources allocation and scheduling, optimal process control, drug discovery and development, as well as health management. Figure 2 provides a diagram outlining the application domains, illustrating how this survey is organized along the lines of the three broad domains in the field.

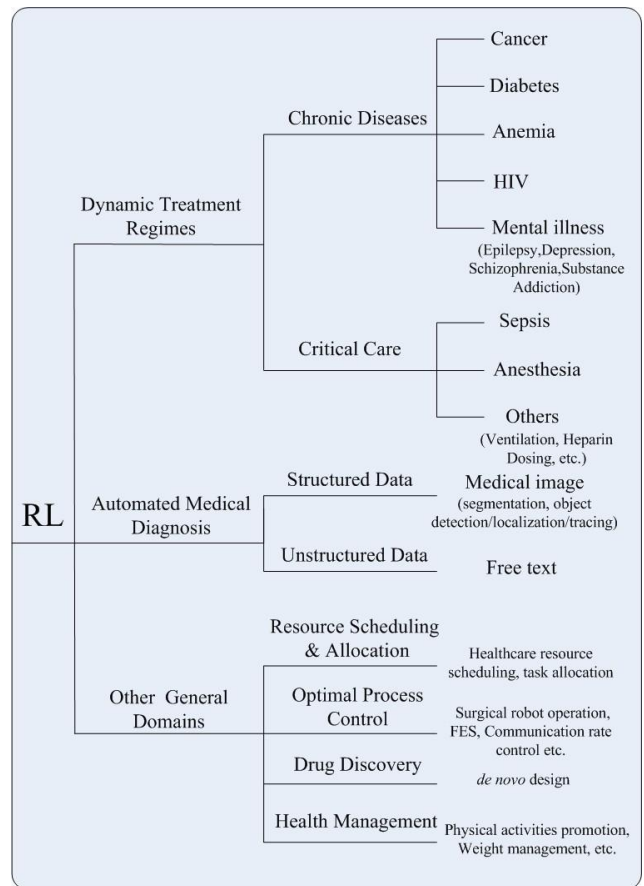


Fig. 2. The outline of application domains of RL in healthcare.

IV. DYNAMIC TREATMENT REGIMES

One goal of healthcare decision-making is to develop effective treatment regimes that can dynamically adapt to the varying clinical states and improve the long-term benefits of patients. *Dynamic treatment regimes* (DTRs) [72], [73], alternatively named as *dynamic treatment policies* [74], *adaptive interventions* [75], or *adaptive treatment strategies* [76], provide a new paradigm to automate the process of developing new effective treatment regimes for individual patients with long-term care [77]. A DTR is composed of a sequence of decision rules to determine the course of actions (e.g., treatment type, drug dosage, or reexamination timing) at a time point according to the current health status and prior treatment history of an individual patient. Unlike traditional randomized controlled trials that are mainly used as an evaluative tool for confirming the efficacy of a newly developed treatment, DTRs are tailored for generating new scientific hypotheses and developing optimal treatments across or within groups of patients [77]. Utilizing valid data generated, for instance, from the *Sequential Multiple Assignment Randomized Trial* (SMART) [78], [79], an optimal DTR that is capable of optimizing the final clinical outcome of particular interest can be derived.

The design of DTRs can be viewed as a sequential decision making problem that fits into the RL framework well. The series of decision rules in DTRs are equivalent to the policies in RL, while the treatment outcomes are expressed by the reward functions. The inputs in DTRs are a set of clinical observations and assessments of patients, and the outputs are the treatments options at each stage, equivalent to the states and actions in RL, respectively. Apparently, applying RL methods to solve DTR problems demonstrates several benefits. RL is capable of achieving time-dependent decisions on the best treatment for each patient at each decision time, thus accounting for heterogeneity across patients. This precise treatment can be achieved even without relying on the identification of any accurate mathematical models or explicit relationship between treatments and outcomes. Furthermore, RL driven solutions enable to improve long-term outcomes by considering delayed effect of treatments, which is the major characteristic of medical treatment. Finally, by careful engineering the reward function using expert or domain knowledge, RL provides an elegant way to multi-objective optimization of treatment between efficacy and the raised side effect.

Due to these benefits, RL naturally becomes an appealing tool for constructing optimal DTRs in healthcare. In fact, solving DTR problems accounts for a large proportion of RL studies in healthcare applications, which can be supported by the dominantly large volume of references in this area. The domains of applying RL in DTRs can be classified into two main categories: chronic diseases and critical care.

A. Chronic Diseases

Chronic diseases are now becoming the most pressing public health issue worldwide, constituting a considerable portion of death every year [80]. Chronic diseases normally feature a long period lasting three months or more, expected to require continuous clinical observation and medical care. The widely prevailing chronic diseases include endocrine diseases (e.g., diabetes and hyperthyroidism), cardiovascular diseases (e.g., heart attacks and hypertension), various mental illnesses (e.g., depression and schizophrenia), cancer, HIV infection, obesity, and other oral health problems [81]. Long-term treatment of these illnesses is often made up of a sequence of medical intervention that must take into account the changing health status of a patient and adverse effects occurring from previous treatment. In general, the relationship of treatment duration, dosage and type against the patient's response is too complex to be explicitly specified. As such, practitioners usually resort to some protocols following the *Chronic Care Model* (CCM) [82] to facilitate decision making in chronic disease conditions. Since such protocols are derived from average responses to treatment in populations of patients, selecting the best sequence of treatments for an individual patient poses significant challenges due to the diversity across or within the population. RL has been utilized to automate the discovery and generation of optimal DTRs in a variety of

chronic diseases including cancer, diabetes, anemia, HIV and several common mental illnesses.

1) *Cancer*: Cancer is one of the main chronic diseases that causes death. About 90.5 million people had cancer in 2015 and approximately 14 million new cases are occurring each year, causing about 8.8 million annual deaths that account for

TABLE III
SUMMARY OF RL APPLICATION EXAMPLES IN THE DEVELOPMENT OF DTRs IN CANCER

Applications	References	Base Methods	Efficient Techniques	Representational Techniques	Data Acquisition	Highlights or Limits
Optimal chemotherapy drug dosage for cancer treatment	Zhao <i>et al.</i> [83]	Q-learning	BRL	N/A	ODE model	Using SVR or ERT to fit Q values; simplistic reward function structure with integer values to assess the tradeoff between efficacy and toxicity.
	Hassani <i>et al.</i> [84]	Q-learning	N/A	N/A	ODE model	Naive discrete formulation of states and actions.
	Ahn & Park [85]	NAC	N/A	N/A	ODE model	Discovering the strategy of performing continuous treatment from the beginning.
	Humphrey [86]	Q-learning	BRL	N/A	ODE model proposed in [83]	Using three machine learning methods to fit Q values, in high dimensional and subgroup scenarios.
	Padmanabhan [87]	Q-learning	N/A	N/A	ODE model	Using different reward functions to model different constraints in cancer treatment.
	Zhao <i>et al.</i> [88]	Q-learning	BRL (FQSVR)	N/A	ODE model driven by real NSCLC data	Considering censoring problem in multiple lines of treatment in advanced NSCLC; using overall survival time as the net reward.
	Furnkranz <i>et al.</i> [52], Cheng <i>et al.</i> [89]	PI	N/A	PRL	ODE model proposed in [83]	Combining preference learning and RL for optimal therapy design in cancer treatment, but only in model-based DP settings.
	Akrour <i>et al.</i> [90], Busa-Fekete <i>et al.</i> [91]	PS	N/A	PRL	ODE model proposed in [83]	Using active ranking mechanism to reduce the number of needed ranking queries to the expert to yield a satisfactory policy without a generated model.
Optimal fractionation scheduling of radiation therapy for cancer treatment	Vincent [92]	Q-learning, SARSA(λ), TD(λ), PS	BRL (FQIERT)	N/A	Linear model, ODE model	Extended ODE model for radiation therapy; using hard constraints in the reward function and simple exploration strategy.
	Tseng <i>et al.</i> [93]	Q-learning	N/A	DRL (DQN)	Data from 114 NSCLC patients	Addressing limited sample size problem using GAN and approximating the transition probability using DNN.
	Jalalimanesh <i>et al.</i> [94]	Q-learning	N/A	N/A	Agent-based model	Using agent-based simulation to model the dynamics of tumor growth.
	Jalalimanesh <i>et al.</i> [95]	Q-learning	N/A	MORL	Agent-based model	Formulated as a multi-objective problem by considering conflicting objective of minimising tumour therapy period and unavoidable side effects.
Hypothetical or generic cancer clinical trial	Goldberg & Kosorok [96], Soliman [97]	Q-learning	N/A	N/A	Linear model	Addressing problems with censored data and a flexible number of stages.
	Yauney & Shah [98]	Q-learning	N/A	DRL (DDQN)	ODE model	Addressing the problem of unstructured outcome rewards using action-driven rewards.

16.7% of total deaths worldwide [99]. The primary treatment options for cancer include surgery, chemotherapy, and radiation therapy. To analyze the dynamics between tumor and immune systems, numerous computational models for spatiotemporal or non-spatial tumor-immune dynamics have been proposed and analyzed by researchers over the past decades [100]. Building on these models, control policies have been put forward to obtain efficient drug administration (see [85], [101] and references therein).

Being a sequential evolutionary process by nature, cancer treatment is a major objective of RL in DTR applications [102], [103]. Table III summarizes the major studies of applying RL in various aspects of cancer treatment, from the perspectives of application scenarios (chemotherapy, radiotherapy or generic cancer treatment simulation), basic RL methods, the efficient and representational techniques applied (if applicable), the learning data (retrospective clinical data, or generated from simulation models or computational models), and the main highlights and limits of the study.

RL methods have been extensively studied in deriving efficient treatment strategies for cancer chemotherapy. Zhao *et al.* [83] first applied model-free TD method, Q-learning, for decision making of agent dosage in chemotherapy. Drawing on

the chemotherapy mathematical model expressed by several *Ordinary Difference Equations* (ODE), virtual clinical trial data from *in vivo* tumor growth patterns was quantitatively generated. Two explicit machine learning approaches, *support vector regression* (SVR) [104] and *extremely randomized trees* (ERT) [41], were applied to fit the approximated Q-functions to the generated trial data. Using this kind of batch learning methods, it was demonstrated that optimal strategies could be extracted directly from clinical trial data in simulation. Ahn and Park [85] studied the applicability of the *Natural AC* (NAC) approach [21] to the drug scheduling of cancer chemotherapy based on an ODE-based tumor growth model proposed by de Pillis and Radunskaya [105]. Targeting at minimizing the tumor cell population and the drug amount while maximizing the populations of normal and immune cells, the NAC approach could discover an effective drug scheduling policy by injecting drug continuously from the beginning until an appropriate time. This policy showed better performance than traditional pulsed chemotherapy protocol that administers the drug in a periodical manner, typically on an order of several hours. The superiority of using continuous dosing treatment over a burst of dosing treatment was also supported by the work [84], where naive discrete Q-learning was applied. More recently,

Padmanabhan *et al.* [87] proposed different formulations of reward function in Q-learning to generate effective drug dosing policies for patient groups with different characteristics. Humphrey [86] investigated several supervised learning approaches (*Classification And Regression Trees* (CART), random forests, and modified version of *Multivariate Adaptive Regression Splines* (MARS)) to estimate Q values in a simulation of an advanced generic cancer trial.

Radiotherapy is another major option of treating cancer, and a number of studies have applied RL approaches for developing automated radiation adaptation protocols [106]. Jalalimanesh *et al.* [94] proposed an agent-based simulation model and Q-learning algorithm to optimize dose calculation in radiotherapy by varying the fraction size during the treatment. Vincent [92] described preliminary efforts in investigating a variety of RL methods to find optimal scheduling algorithms for radiation therapy, including the exhaustive PS [20], FQI [40], SARSA(λ) [19] and K-Nearest NeighborsTD(λ) [107]. The preliminary findings suggest that there may be an advantage in using non-uniform fractionation schedules for some tissue types.

As the goal of radiotherapy is in essence a multi-objective problem to erase the tumour with radiation while not impacting normal cells as much as possible, Jalalimanesh *et al.* [95] proposed a multi-objective distributed Q-learning algorithm to find the Pareto-optimal solutions for calculating radiotherapy dose. Each objective was optimized by an individual learning agent and all the agents compromised their individual solutions in order to derive a Pareto-optimal solution. Under the multiobjective formulation, three different clinical behaviors could be properly modeled (i.e., aggressive, conservative or moderate), by paying different degree of attention to eliminating cancer cells or taking care of normal cells.

A recent study [93] proposed a multi-component DRL framework to automate adaptive radiotherapy decision making for *non-small cell lung cancer* (NSCLC) patients. Aiming at reproducing or mimicking the decisions that have been previously made by clinicians, three neural network components, namely *Generative Adversarial Net* (GAN), transition *Deep Neural Networks* (DNN) and *Deep Q Network* (DQN), were applied: the GAN component was used to generate sufficiently large synthetic patient data from historical small-sized real clinical data; the transition DNN component was employed to learn how states would transit under different actions of dose fractions, based on the data synthesized from the GAN and available real clinical data; once the whole MDP model has been provided, the DQN component was then responsible for mapping the state into possible dose strategies, in order to optimize future radiotherapy outcomes. The whole framework was evaluated in a retrospective dataset of 114 NSCLC patients who received radiotherapy under a successful dose escalation protocol. It was demonstrated that the DRL framework was able to learn effective dose adaptation policies between 1.5 and 3.8 Gy,

which complied with the original dose range used by the clinicians.

The treatment of cancer poses several significant theoretical problems for applying existing RL approaches. Patients may drop out the treatment anytime due to various uncontrolled reasons, causing the final treatment outcome (e.g., survival time in cancer treatment) unobserved. This *data censoring problem* [96] complicates the practical use of RL in discovering individualized optimal regimens. Moreover, in general cancer treatment, the initiation and timing of the next line of therapy depend on the disease progression, and thus the number of treatment stage can be flexible. For instance, NSCLC patients usually receive one to three treatment lines, and the necessity and timing of the second and third lines of treatment vary from person to person. Developing valid methodology for computing optimal DTRs in such a flexible setting is currently a premier challenge. Zhao *et al.* [88] presented an adaptive Q-learning approach to discover optimal DTRs for the first and second lines of treatment in Stage IIIB/IV NSCLC. The trial was conducted by randomizing the different compounds for first and second-line treatments, as well as the timing of initiating the second-line therapy. In order to successfully handle the complex censored survival data, a modification of SVG approach, $\epsilon - SVR - C$, was proposed to estimate the optimal Q values. A simulation study showed that the approach could select optimal compounds for two lines of treatment directly from clinical data, and the best initial time for second-line therapy could be derived while taking into account the heterogeneity across patients. Other studies [96], [97] presented the novel *censored-Q-learning* algorithm that is adjusted for a multi-stage decision problem with a flexible number of stages in which the rewards are survival times that are subject to censoring.

To tackle the problem that a numerical reward function should be specified beforehand in standard RL techniques, several studies investigated the possibility of formulating rewards using qualitative preference or simply based on past actions in the treatment of cancer [89], [52], [98]. Akrou *et al.* [90] proposed a PRL method combined with active ranking in order to decrease the number of ranking queries to the expert needed to yield a satisfactory policy. Experiments on the cancer treatment testbeds showed that a very limited external information in terms of expert's ranking feedbacks might be sufficient to reach state-of-the-art results. Busa-Fekete *et al.* [91] introduced a preference-based variant of a direct PS method in the medical treatment design for cancer clinical trials. A novel approach based on action-driven rewards was first proposed in [98]. It was showed that new dosing regimes in cancer chemotherapy could be learned using action-derived penalties, suggesting the possibility of using RL methods in situations when final outcomes are not available, but priors on beneficial actions can be more easily specified.

2) *Diabetes: Diabetes mellitus*, or simply called *diabetes*, is one of the most serious chronic diseases in the world. According to a recent report released by International Diabetes Federation (IDF), there are 451 million people living with diabetes in 2017, causing approximately 5 million deaths worldwide and USD 850 billion global healthcare expenditure. It is expected that by 2045, the total number of adults with diabetes would increase to near 700 million, accounting for 9.9% of the adult population.

. Since the high prevalence of diabetes presents significant social influence and financial burdens, there has been an increasing urgency to ensure effective treatment to diabetes across the world.

Intensive research concern has been devoted to the development of effective blood glucose control strategies in treatment of insulin-dependent diabetes (i.e., type 1 diabetes). Since its first proposal in the 1970s [109], *artificial pancreas* (AP) have been widely used in the blood glucose control process to compute and administrate a precise insulin dose, by using a *continuous glucose monitoring system* (CGMS) and a closed-loop controller [110]. Tremendous progress has been made towards insulin infusion rate automation in AP using traditional control strategies such as *Proportional-IntegralDerivative* (PID), *Model Predictive Control* (MPC), and *Fuzzy Logic* (FL) [111], [112]. A major concern is the inter- and intra- variability of the diabetic population which raises the demand for a personalized, patient specific approach of the glucose regulation. Moreover, the complexity of the physiological system, the variety of disturbances such as meal, exercise, stress and sickness, along with the difficulty in modelling accurately the glucose-insulin regulation system all raise the need in the development of more advanced adaptive algorithms for the glucose regulation.

RL approaches have attracted increasingly high attention in personalized, patient specific glucose regulation in AP systems [113]. Yasini *et al.* [114] made an initial study on using RL to control an AP to maintain normoglycemic around 80 mg/dl. Specifically, model-free TD Q-learning algorithm was applied to compute the insulin delivery rate, without relying on an explicit model of the glucose-insulin dynamics. Daskalaki *et al.* [115] presented an AC controller for the estimation of insulin infusion rate *in silico* trial based on the University of Virginia/Padova type 1 diabetes simulator [116]. In an evaluation of 12 day meal scenario for 10 adults, results showed that the approach could prevent hypoglycaemia well, but hyperglycaemia could not be properly solved due to the static behaviors of the *Actor* component. The authors then proposed using daily updates of the average *basal rate* (BR) and the *insulin-to-carbohydrate* (IC) ratio in order to optimize glucose regulation [117], and using estimation of *information transfer* (IT) from insulin to glucose for automatic and

personalized tuning of the AC approach [118]. This idea was motivated by the fact that small adaptation of insulin in the *Actor* component may be sufficient in case of large amount of IT from insulin to glucose, whereas more dramatic updates may be required for low IT. The results from the *Control Variability Grid Analysis* (CVGA) showed that the approach could achieve higher performance in all three groups of patients, with 100% percentages in the A+B zones for adults, and 93% for both adolescents and children, compared to approaches with random initialization and zero initial values. The AC approach was significantly extended to directly link to patient-specific characteristics, and evaluated more extensively under a complex meal protocol, meal uncertainty and insulin sensitivity variation [119], [120].

A number of studies used certain mathematical models to simulate the glucose-insulin dynamic system in patients. Based on the Palumbo mathematical model [121], the onpolicy SARSA was used for insulin delivery rate [122]. Ngo *et al.* applied model-based VI method [123] and AC method [124] to reduce the fluctuation of the blood glucose in both fasting and post-meal scenarios, drawing on the Bergman's minimal insulin-glucose kinetics model [125] and the Hovorka model [126] to simulate a patient. De Paula *et al.* [127], [128] proposed policy learning algorithms that integrates RL with Gaussian processes to take into account glycemic variability under uncertainty, using the Ito's stochastic model of the glucose-insulin dynamics [129].

There are also several data-driven studies carried out to analyze RL in diabetes treatment based on real data from diabetes patients. Utilizing the data extracted from the medical records of over 10,000 patients in the University of Tokyo Hospital, Asoh *et al.* [130] estimated the MDP model underlying the progression of patient state and evaluated the value of treatment using the VI method. The opinions of a doctor were used to define the reward for each treatment. The preassumption of this predefined reward function then motivated the application of IRL approach to reveal the reward function that doctors were using during their treatments [131]. Using observational data on the effect of food intake and physical activity in an outpatient setting using mobile technology, Lockett *et al.* [132] proposed the *V-learning* method that directly estimates a policy which maximizes the value over a class of policies and requires minimal assumptions on the data-generating process. The method has been applied to estimate treatment regimes to reduce the number of hypo and hyperglycemic episodes in patients with type 1 diabetes.

3) *Anemia*: Anemia is a common comorbidity in chronic renal failure that occurs in more than 90% of patients with *endstage renal disease* (ESRD) who are undertaking hemodialysis. Caused by a failure of adequately producing *endogenous erythropoietin* (EPO) and thus red blood cells, anemia can have significant impact on organ functions, giving rise to a number of severe consequences such as heart disease or even increased mortality. Currently, anemia can be successfully treated by administering *erythropoiesis-*

stimulating agents (ESAs), in order to maintain the *hemoglobin* (HGB) level within a narrow range of 11-12 g/dL. To achieve this, professional clinicians must carry out a labor intensive process of dosing ESAs to assess monthly HGB and iron levels before making adjustments accordingly. However, since the existing *Anemia Management Protocol* (AMP) does not account for the high inter- and intraindividual variability in the patient's response, the HGB level of some patients usually oscillates around the target range, causing several risks and side-effects.

As early as in 2005, Gaweda *et al.* [133] first proposed using RL to perform individualized treatment in the management of renal anemia. The target under control is the HGB, whereas the control input is the amount of EPO administered by the physician. As the iron storage in the patient, determined by *Transferrin Saturation* (TSAT), also has an impact on the process of red blood cell creation, it is considered as a state component together with HGB. To model distinct dose-response relationship within a patient population, a fuzzy model was estimated first by using real records of 186 hemodialysis patients from the Division of Nephrology, University of Louisville. On-policy TD method, SARSA, was then performed on the sample trajectories generated by the model. Results show that the proposed approach generates adequate dosing strategies for representative individuals from different response groups. The authors then proposed a combination of MPC approach with SARSA for decision support in anemia management [134], with the MPC component used for simulation of patient response and SARSA for optimization of the dosing strategy. However, the automated RL approaches in these studies could only achieve a policy with a comparable outcome against the existing AMP. Other studies applied various kinds of Q-learning, such as Q-learning with function approximation, or directly based on state-aggregation [135], [136], [137], in providing effective treatment regimes in anemia.

Several studies resorted to BRL methods to derive optimal ESA dosing strategies for anemia treatment. By performing a retrospective study of a cohort of 209 hemodialysis patients, Malof and Gaweda [138] adopted the batch FQI method to achieve dosing strategies that were superior to a standard AMP. The FQI method was also applied by Escandell *et al.* [139] for discovering efficient dosing strategies based on the historical treatment data of 195 patients in nephrology centers allocated around Italy and Portugal. An evaluation of the FQI method on a computational model that describes the effect of ESAs on the hemoglobin level showed that FQI could achieve an increment of 27.6% in the proportion of patients that are within the targeted range of hemoglobin during the period of treatment. In addition, the quantity of drug needed is reduced by 5.13%, which indicates a more efficient use of ESAs [140].

4) *HIV*: Discovering effective treatment strategies for HIV-infected individuals remains one of the most significant challenges in medical research. To date, the effective way to treat HIV makes use of a combination of anti-HIV drugs (i.e.,

antiretrovirals) in the form of *Highly Active Antiretroviral Therapy* (HAART) to inhibit the development of drug-resistant HIV strains [141]. Patients suffering from HIV are typically prescribed a series of treatments over time in order to maximize the long-term positive outcomes of reducing patients' treatment burden and improving adherence to medication. However, due to the differences between individuals in their immune responses to treatment, discovering the optimal drug combinations and scheduling strategy is still a difficult task in both medical research and clinical trials.

Ernst *et al.* [142] first introduced RL techniques in computing *Structured Treatment Interruption* (STI) strategies for HIV infected patients. Using a mathematical model [141] to artificially generate the clinical data, the BRL method FIQERT was applied to learn an optimal drug prescription strategy in an off-line manner. The derived STI strategy is featured with a cycling between the two main anti-HIV drugs: *Reverse Transcriptase Inhibitors* (RTI) and *Protease Inhibitors* (PI), before bringing the patient to the healthy drug-free steadystate. Using the same mathematical model, Parbhoo [143] further implemented three kinds of BRL methods, FQI-ERT, neural FQI and LSPI, to the problem of HIV treatment, indicating that each learning technique had its own advantages and disadvantages. Moreover, a testing based on a ten-year period of real clinical data from 250 HIV-infected patients in Charlotte Maxeke Johannesburg Academic Hospital, South Africa verified that the RL methods were capable of suggesting treatments that were reasonably compliant with those suggested by clinicians.

A mixture-of-experts approach was proposed in [144] to combine the strengths of both kernel-based regression methods (i.e., history-alignment model) and RL (i.e., model-based Bayesian PORL) for HIV therapy selection. Since kernelbased regression methods are more suitable for modeling more related patients in history, while model-based RL methods are more suitable for reasoning about the future outcomes, automatically selecting an appropriate model for a particular patient between these two methods thus tends to provide simpler yet more robust patterns of response to the treatment. Making use of a subset of the EuResist database consisting of HIV genotype and treatment response data for 32,960 patients, together with the 312 most common drug combinations in the cohort, the treatment therapy derived by the mixture-of-experts approach outperformed those derived by each method alone.

Since the treatment of HIV highly depends the patient's immune system that varies from person to person, it is thus necessary to derive efficient learning strategies that can address and identify the variations across subpopulations. Marivate *et al.* [145] formalized a routine to accommodate multiple sources of uncertainty in BRL methods to better evaluate the effectiveness of treatments across a subpopulations of patients. Other approaches applied various kinds of TRL techniques so as to take advantage of the prior

information from previously learned transition models [146], [147] or learned policy [148]. More recently, Yu *et al.* [149] proposed a causal policy gradient algorithm and evaluated it in the treatment of HIV in order to facilitate the final learning performance and increase explanations of learned strategies.

The treatment of HIV provides a well-known testbed for evaluation of exploration mechanisms in RL research. Simulations show that the basin of attraction of the healthy steady-state is rather small compared to that of the nonhealthy steady state [141]. Thus, general exploration methods are unable to yield meaningful performance improvement as they can only obtain samples in the vicinity of the “nonhealthy” steady state. To solve this issue, several studies have proposed more advanced exploration strategies in order to increase the learning performance in HIV treatment. Pazis *et al.* [150] introduced an algorithm for PAC optimal exploration in continuous state spaces. Kawaguchi considered the time bound in a PAC exploration process [151]. Results in both studies showed that the exploration algorithm could achieve far better strategies than other existing exploration strategies in HIV treatment.

5) *Mental Disease*: Mental diseases are characterized by a long-term period of clinical treatments that usually require adaptation in the duration, dose, or type of treatment over time [152]. Given that the brain is a complex system and thus extremely challenging to model, applying traditional control-based methods that rely on accurate brain models in mental disease treatment is proved infeasible. Well suited to the problem at hand, RL has been widely applied to DTRs in a wide range of mental illness including epilepsy, depression, schizophrenia and various kinds of substance addiction.

a) *Epilepsy*: Epilepsy is one of the most common severe neurological disorders, affecting around 1% of the world population. When happening, epilepsy is manifested in the form of intermittent and intense seizures that are recognized as abnormal synchronized firing of neural populations. Implantable electrical deep-brain stimulation devices are now an important treatment option for drug-resistant epileptic patients. Researchers from nonlinear dynamic systems analysis and control have proposed promising prediction and detection algorithms to suppress the frequency, duration and amplitude of seizures [153]. However, due to lack of full understanding of seizure and its associated neural dynamics, designing optimal seizure suppression algorithms via minimal electrical stimulation has been for a long time a challenging task in treatment of epilepsy.

RL enables direct closed-loop optimizations of deep-brain stimulation strategies by adapting control policies to patients’ unique neural dynamics, without necessarily relying on having accurate prediction or detection of seizures. The goal is to explicitly maximize the effectiveness of stimulation, while simultaneously minimizing the overall amount of stimulation applied thus reducing cell damage and preserving cognitive and neurological functions [154]. Guez *et al.* [155], [156], [157] applied the BRL method, FQI-ERT, to optimize a deepbrain

stimulation strategy for the treatment of epilepsy. Encoding the observed *Electroencephalograph* (EEG) signal as a 114-dimensional continuous feature vector, and four different simulation frequencies as the actions, the RL approach was applied to learn an optimal stimulation policy using data from an *in vitro* animal model of epilepsy (i.e., field potential recordings of seizure-like activity in slices of rat brains). Results showed that RL strategies substantially outperformed the current best stimulation strategies in the literature, reducing the incidence of seizures by 25% and total amount of electrical stimulation to the brain by a factor of about 10. Subsequent validation work [158] showed generally similar results that RL-based policy could prevent epilepsy with a significant reduced amount of stimulation, compared to fixed-frequency stimulation strategies. Bush and Pineau [159] applied manifold embeddings to reconstruct the observable state space in MRL, and applied the proposed approach to tackle the high complexity of nonlinearity and partially observability in real-life systems. The learned neurostimulation policy was evaluated to suppress epileptic seizures on animal brain slices and results showed that seizures could be effectively suppressed after a short transient period.

While the above *in vitro* biological models of epilepsy are useful for research, they are nonetheless time-consuming and associated with high cost. In contrast, computational models can provide large amounts of reproducible and cheap data that may permit precise manipulations and deeper investigations. Vincent [92] proposed an *in silico* computational model of epileptiform behavior in brain slices, which was verified by using biological data from rat brain slices *in vitro*. Nagaraj *et al.* [160] proposed the first computational model that captures the transition from inter-ictal to ictal activity, and applied naive Q-learning method to optimize stimulation frequency for controlling seizures with minimum stimulations. It was shown that even such simple RL methods could converge on the optimal solution in simulation with slow and fast interseizure intervals.

b) *Depression*: *Major depressive disorder* (MDD), also known simply as depression, is a mental disorder characterized by at least two weeks of low mood that is present across most situations. Using data from the *Sequenced Treatment Alternatives to Relieve Depression* (STAR*D) trial [161], which is a sequenced four-stage randomized clinical trial of patients with MDD, Pineau *et al.* [162] first applied Kernel-based BRL [163] for constructing useful DTRs for patients with MDD. Other work tries to address the problem of nonsmooth of decision rules as well as nonregularity of the parameter estimations in traditional RL methods by proposing various extensions over default Q-learning procedure in order to increase the robustness of learning [164]. Lauer *et al.* [165] proposed a new version of Q-learning, *interactive Qlearning* (IQ-learning), by interchanging the order of certain steps in traditional Q-learning, and showed that IQ-learning improved on Q-learning in terms of integrated mean squared error in a study of MDD. The IQ-learning framework was then extended

to optimize functionals of the outcome distribution other than the expected value [166], [167]. Schulte *et al.* [168] provided systematic empirical studies of Q-learning and *Advantage-learning* (A-learning) [169] methods and illustrated their performance using data from an MDD study. Other approaches include the *penalized Q-learning* [170], the *Augmented Multistage Outcome-Weighted Learning* (AMOL) [171], the budgeted learning algorithm [172], and the *Censored Q-learning* algorithm [97].

c) *Schizophrenia*: RL methods have been also used to derive optimal DTRs in treatment of schizophrenia, using data from the *Clinical Antipsychotic Trials of Intervention Effectiveness* (CATIE) study [173], which was an 18-month study divided into two main phases of treatment. An in-depth case study of using BRL, FQI, to optimize treatment choices for patients with schizophrenia using data from CATIE was given by [174]. Key technical challenges of applying RL in typically continuous, highly variable, and high-dimensional clinical trials with missing data were outlined. To address these issues, the authors proposed the use of multiple imputation to overcome the missing data problem, and then presented two methods, *bootstrap voting* and *adaptive confidence intervals*, for quantifying the evidence in the data for the choices made by the learned optimal policy. Ertefaie *et al.* [175] accommodated residual analyses into Q-learning in order to increase the accuracy of model fit and demonstrated its superiority over standard Q-learning using data from CATIE.

Some studies have focused on optimizing multiple treatment objectives in dealing with schizophrenia. Lizotte *et al.* [176] extended the FQI algorithm by considering multiple rewards of symptom reduction, side-effects and quality of life simultaneously in sequential treatments for schizophrenia. However, it was assumed that end-users had a true reward function that was linear in the objectives and all future actions could be chosen optimally with respect to the same true reward function over time. To solve these issues, the authors then proposed the non-deterministic multi-objective FIQ algorithm, which computed policies for all preference functions simultaneously from continuous-state, finite-horizon data [177]. When patients do not know or cannot communicate their preferences, and there is heterogeneity across patient preferences for these outcomes, formation of a single composite outcome that correctly balances the competing outcomes for all patients is not possible. Laber *et al.* [178] then proposed a method for constructing DTRs for schizophrenia that accommodates competing outcomes and preference heterogeneity across both patients and time by recommending sets of treatments at each decision point. Butler *et al.* [179] derived a preference sensitive optimal DTR for schizophrenia patient by directly eliciting patients' preferences overtime.

d) *Substance Addiction*: *Substance addiction*, or *substance use disorder* (SUD), often involves a chronic course of repeated cycles of cessation followed by relapse [180], [75]. There has been great interest in the development of DTRs by

investigators to deliver in-time interventions or preventions to end-users using RL methods, guiding them to lead healthier lives. For example, Murphy *et al.* [181] applied AC algorithm to reduce heavy drinking and smoking for university students. Chakraborty *et al.* [77], [182], [183] used Q-learning with linear models to identify DTRs for smoking cessation treatment regimes. Tao *et al.* [184] proposed a tree-based RL method to directly estimate optimal DTRs, and identify dynamic SUD treatment regimes for adolescents.

B. Critical Care

Unlike the treatment of chronic diseases, which usually requires a long period of constant monitoring and medication, critical care is dedicated to more seriously ill or injured patients that are in need of special medical treatments and nursing care. Usually, such patients are provided with separate geographical area, or formally named the *intensive care unit* (ICU), for intensive monitoring and close attention, so as to improve the treatment outcomes [185]. ICUs will play a major role in the new era of healthcare systems. It is estimated that the ratio of ICU beds to hospital beds would increase from 3-5% in the past to 20-30% in the future [186].

Significant attempts have been devoted to the development of clearer guidelines and standardizing approaches to various aspects of interventions in ICUs, such as sedation, nutrition, administration of blood products, fluid and vasoactive drug therapy, haemodynamic endpoints, glucose control, and mechanical ventilation [185]. Unfortunately, only a few of these interventions could be supported by high quality evidence from randomised controlled trials or meta-analyses [187], especially when it comes to development of potentially new therapies for complex ICU syndromes, such as sepsis [188] and acute respiratory distress syndrome [189].

Thanks to the development in ubiquitous monitoring and censoring techniques, it is now possible to generate rich ICU data in a variety of formats such as free-text clinical notes, images, physiological waveforms, and vital sign time series, suggesting a great deal of opportunities for the applications of machine learning and particularly RL techniques in critical care [190], [191]. However, the inherent 3C (*Compartmentalization, Corruption, and Complexity*) features indicate that critical care data are usually noisy, biased and incomplete [5]. Properly processing and interpreting this data in a way that can be used by existing machine learning methods is the premier challenge of data analysis in critical care. To date, RL has been widely applied in the treatment of sepsis (Section IV-B1), regulation of sedation (Section IV-B2), and some other decision making problems in ICUs such as mechanical ventilation and heparin dosing (Section IV-B3). Table IV summarizes these applications according to the applied RL techniques and the sources of data acquired during learning.

1) *Sepsis*: Sepsis, which is defined as severe infection causing life-threatening acute organ failure, is a leading cause

of mortality and associated healthcare costs in critical care [226]. While numbers of international organizations have devoted significant efforts to provide general guidance for treating sepsis over the past 20 years, physicians at practice still lack universally agreed-upon decision support for sepsis [188]. With the available data obtained from freely accessible critical care databases such as the *Multiparameter Intelligent Monitoring in Intensive Care* (MIMIC) [227], recent years have seen an increasing number of studies that applied RL techniques to the problem of deducing optimal treatment policies for patients with sepsis [228].

The administration of *intravenous* (IV) and maximum *vasopressor* (VP) is a key research and clinical challenge in sepsis. A number of studies have been carried out to tackle this issue in the past years. Komorowski *et al.* [192], [193] directly applied the on-policy SARSA algorithm and modelbased PI method in a discretized state and action-space. Raghu *et al.* [194], [195] examined fully continuous state and action space, where policies are learned directly from the physiological state data. To this end, the authors proposed the fully-connected *Dueling Double DQN* to learn an approximation for the optimal action-value function, which combines three state-of-the-art efficiency and stability boosting techniques in DRL, i.e., *Double DQN* [229], *Dueling DQN* [230] and *Prioritized Experience Replay* (PER) [231]. Experimental results demonstrated that using continuous state-space modeling could identify interpretable policies with improved patient outcomes, potentially reducing patient mortality in the hospital by 1.8 - 3.6%. The authors also directly estimated the transition model in continuous state-space, and applied two PS methods, the direct policy gradient and *Proximal Policy Optimization* PPO [232], to derive a treatment strategy [196]. Utomo *et al.* [197] proposed a graphical model that was able to show transitions of patient health conditions and treatments for better explainability, and applied MC to generate a realtime treatment recommendation. Li *et al.* [201] provided an online POMDP solution to take into account uncertainty and history information in sepsis clinical applications. Futoma *et al.* [199] used multi-output Gaussian processes and DRL to directly learn from sparsely sampled and frequently missing multivariate time series ICU data. Peng *et al.* [198] applied the mixture-of-experts framework [144] in sepsis treatment by automatically switching between kernel learning and DRL depending on patient's current history. Results showed that this kind of mixed learning could achieve better performance than the strategies by physicians, Kernel learning and DQN learning alone. Most recently, Yu *et al.* [200] addressed IRL problems in sepsis treatment.

Targeting at glycemic regulation problems for severely ill septic patients, Weng *et al.* [202] applied PI to learn the optimal targeted blood glucose levels from real data trajectories. Petersen *et al.* [203] investigated the cytokine mediation problem in sepsis treatment, using the DRL method, *Deep Deterministic Policy Gradient* (DDPG) [233], to tackle the hidimensional continuous states and actions, and potential-

based reward shaping [234] to facilitate the learning efficiency. The proposed approach was evaluated using an agent-based model, the *Innate Immune Response Agent-Based Model* (IIRABM), that simulates the immune response to infection. The learned treatment strategy was showed to achieve 0.8% mortality over 500 randomly selected patient parameterizations with mortalities average of 49%, suggesting that adaptive, person-

TABLE IV
SUMMARY OF RL APPLICATION EXAMPLES IN THE DEVELOPMENT OF DTRS IN CRITICAL CARE

Domain	Application	Reference	Base method	Efficient Techniques	Representational Techniques	Data Acquisition	Highlights and Limits
Sepsis	Administration of IV fluid and maximum VP	Komorowski <i>et al.</i> [192], [193]	SARSA,PI	N/A	N/A	MIMIC-III	Naive application of SARSA and PI in a discrete state and action-space.
		Raghu <i>et al.</i> [194], [195]	Q-learning	N/A	DRL (DDQN)	MIMIC-III	Application of DRL in a fully continuous state but discrete action space.
		Raghu <i>et al.</i> [196]	PS	MRL	N/A	MIMIC-III	Model-based learning with continuous state-space; integrating clinician's policies into RL policies.
		Utomo <i>et al.</i> [197]	MC	N/A	N/A	MIMIC-III	Estimating transitions of patient health conditions and treatments to increase its explainability.
		Peng <i>et al.</i> [198]	Q-learning	N/A	DRL (DDQN)	MIMIC-III	Adaptive switching between kernel learning and DRL.
		Futoma <i>et al.</i> [199]	Q-learning	N/A	DRL	Clinical data at university hospital	Tackling sparsely sampled and frequently missing multivariate time series data.
		Yu <i>et al.</i> [200]	Q-learning	BRL(FQI)	DRL, IRL	MIMIC-III	Inferring the best reward functions using deep IRL.
	Li <i>et al.</i> [201]	AC	N/A	PORL	MIMIC-III	Taking into account uncertainty and history information of sepsis patients.	
	Targeted blood glucose regulation	Weng <i>et al.</i> [202]	PI	N/A	N/A	MIMIC-III	Learning the optimal targeted blood glucose levels for sepsis patients
Cytokine mediation	Petersen <i>et al.</i> [203]	AC	N/A	DRL (DDPG)	Agent-based model	Using reward shaping to facilitate the learning efficiency; significantly reducing mortality from 49% to 0.8%.	
Anesthesia	Regulation and automation of sedation and analgesia to maintain physiological stability and lowering pains of patients	Moore <i>et al.</i> [204], [205]	$Q(\lambda)$	N/A	N/A	PK/PD model	Achieving superior stability compared to a well-tuned PID controller.
		Moore <i>et al.</i> [206], [207]	Q-learning	N/A	N/A	PK/PD model	Using the change of BIS as the state representation.
		Moore <i>et al.</i> [208], [209]	Q-learning	N/A	N/A	<i>In vivo</i> study	First clinical trial for anesthesia administration using RL on human volunteers.
		Sadati <i>et al.</i> [210]	Unclear	N/A	N/A	PK/PD model	Expert knowledge can be used to realize reasonable initial dosage and keep drug inputs in safe values.
		Borera <i>et al.</i> [211]	Q-learning	N/A	N/A	PK/PD model	Using an adaptive filter to eliminate the delays when estimating patient state.
		Lowery & Faisal [212]	AC	N/A	N/A	PK/PD model	Considering the continuous state and action spaces.
		Padmanabhan <i>et al.</i> [213]	Q-learning	N/A	N/A	PK/PD model	Regulating sedation and hemodynamic parameters simultaneously.
		Humbert <i>et al.</i> [214]	N/A	N/A	POMDP, IRL	Clinical data	Training an RL agent to mimic decisions by expert anesthesiologists.
Others	Heparin Dosing	Nemati <i>et al.</i> [215]	Q-learning	BRL	PORL	MIMIC II	End-to-end learning with hidden states of patients.
		Lin <i>et al.</i> [216]	AC	N/A	DRL(DDPG)	MIMIC, Emory Healthcare data	Addressing dosing problems in continuous state-action spaces.
	General medication recommendation	Wang <i>et al.</i> [217]	AC	N/A	DRL (DDPG)	MIMIC-III	Combining supervised and reinforcement learning for medication dosing covering a large number of diseases.
	Mechanical ventilation and sedative dosing	Prasad <i>et al.</i> [218]	Q-learning	BRL(FQI)	N/A	MIMIC-III	Optimal decision making for the weaning time of mechanical ventilation and personalized sedation dosage.
		Yu <i>et al.</i> [219]	Q-learning	BRL(FQI)	IRL	MIMIC-III	Applying IRL in inferring the reward functions.
		Yu <i>et al.</i> [220]	AC	N/A	N/A	MIMIC-III	Combining supervised learning and AC for more efficient decision making.
		Jagannatha <i>et al.</i> [221]	Q-learning, PS	BRL(FQI)	N/A	MIMIC-III	Analyzing limitations of off-policy policy evaluation methods in ICU settings.
	Ordering of lab tests	Cheng <i>et al.</i> [222]	Q-learning	BRL(FQI)	MORL	MIMIC III	Designing a multi-objective reward function that reflects clinical considerations when ordering labs.
		Chang <i>et al.</i> [223]	Q-learning	N/A	DRL (Dueling DQN)	MIMIC III	The first RL application on multi-measurement scheduling problem in the clinical setting.
	Prevention and treatments for GVHD	Krakov <i>et al.</i> [224]	Q-learning	N/A	N/A	CIBMTR data	First proposal of DTRS for acute GVHD prophylaxis and treatment.
		Liu <i>et al.</i> [225]	Q-learning	N/A	DRL (DQN)	CIBMTR data	Incorporation of a supervised learning step into RL.

alized multi-cytokine mediation therapy could be promising for treating sepsis.

2) *Anesthesia*: Another major drug dosing problem in ICUs is the regulation and automation of sedation and analgesia, which is essential in maintaining physiological stability and lowering pains of patients. Whereas surgical patients typically require deep sedation over a short duration of time, sedation for ICU patients, especially when using mechanical ventilation, can be more challenging [218]. Critically ill patients who are supported by mechanical ventilation require adequate sedation for several days to guarantee safe treatment in the ICU [235]. A misdosing of sedation or under sedation is not acceptable since over

sedation can cause hypotension, prolonged recovery time, delayed weaning from mechanical ventilation, and other related negative outcomes, whereas under sedation can cause symptoms such as anxiety, agitation and hyperoxia [213].

The regulation of sedation in ICUs using RL methods has attracted attention of researcher for decades. As early as in 1994, Hu *et al.* [70] studied the problem of anesthesia control by applying some of the founding principles of RL (the MDP formulation and its planning solutions). More recently, RLbased control methods, using surrogate measures of anesthetic effect, e.g., the *bispectral* (BIS) index, as the controlled variable, has enhanced individualized anesthetic management, resulting in the overall improvement of patient

outcomes when compared with traditional controlled administration. Moore *et al.* [204], [205] applied TD $Q(\lambda)$ in administration of intravenous propofol in ICU settings, using the well-studied MarshSchnider pharmacokinetic model to estimate the distribution of drug within the patient, and a pharmacodynamic model for estimating drug effect. The RL method adopted the error of BIS and estimation of the four compartmental propofol concentrations as the input state, different propofol dose as control actions, and the BIS error as the reward. The method demonstrated superior stability and responsiveness when compared to a well-tuned PID controller. The authors then modeled the drug disposition system as three states corresponding to the change of BIS, and applied basic Q-learning method to solving this problem [206], [207]. They also presented the first clinical *in vivo* trial for closed-loop control of anesthesia administration using RL on 15 human volunteers [208], [209]. It was demonstrated that patient specific control of anesthesia administration with improved control accuracy as compared to other studies in the literature could be achieved both in simulation and the clinical study.

Targeting at both muscle relaxation (paralysis) and *Mean Arterial Pressure* (MAP), Sadati *et al.* [210] proposed an RL-based fuzzy controllers architecture in automation of the clinical anesthesia. A multivariable anesthetic mathematical model was presented to achieve an anesthetic state using two anesthetic drugs of *Atracurium* and *Isoflurane*. The highlight was that the physician's clinical experience could be incorporated into the design and implementation of the architecture, to realize reasonable initial dosage and keep drug inputs in safe values. Padmanabhan *et al.* [213] used a closed-loop anesthesia controller to regulate the BIS and MAP within a desired range. Specifically, a weighted combination of the error of the BIS and MAP signals is considered in the proposed RL algorithm. This reduces the computational complexity of the RL algorithm and consequently the controller processing time. Borera *et al.* [211] proposed an *Adaptive Neural Network Filter* (ANNF) to improve RL control of propofol hypnosis.

Lowery and Faisal [212] used a continuous AC method to first learn a generic effective control strategy based on average patient data and then fine-tune itself to individual patients in a personalization stage. The results showed that the reinforcement learner could reduce the dose of administered anesthetic agent by 9.4% as compared to a fixed controller, and keep the BIS error within a narrow, clinically acceptable range 93.9% of the time. More recently, an IRL method has been proposed that used expert trajectories provided by anesthesiologists to train an RL agent for controlling the concentration of drugs during a global anesthesia [214].

3) *Other Applications in Critical Care:* While the previous sections are devoted to two topic-specific applications of RL methods in critical care domains, there are many other more general medical problems that perhaps have received less attention by researchers. One such problem is regarding the medication dosing, particularly, heparin dosing, in ICUs. A recent study by Ghassemi *et al.* [236] highlighted that the

mis-dosing of medications in the ICU is both problematic and preventable, e.g., up to two-thirds of patients at the study institution received a non-optimal initial dose of heparin, due to the highly personal and complex factors that affect the dose-response relationship. To address this issue, Nemati *et al.* [215] inferred hidden states of patients via discriminative hidden Markov model and applied neural FQI to learn optimal heparin dosages. Lin *et al.* [216] applied DDPG in continuous state-action spaces to learn a better policy for heparin dosing from observational data in MIMIC and the Emory University clinical data. Wang *et al.* [217] combined supervised signals and reinforcement signals to learn recommendations for medication dosing involving a large number of diseases and medications in ICUs.

Another typical application of RL in ICUs is to develop a decision support tool for automating the process of airway and mechanical ventilation. The need for mechanical ventilation is required when patients in ICUs suffer from *acute respiratory failure* (ARF) caused by various conditions such as cardiogenic pulmonary edema, sepsis or weakness after abdominal surgery [237]. The management of mechanical ventilation is particularly challenging in ICUs. On one hand, higher costs occur if unnecessary ventilation is still taking effect, while premature extubation can give rise to increased risk of morbidity and mortality. Optimal decision making regarding when to wean patients off of a ventilator thus becomes nontrivial since there is currently no consistent clinical opinion on the best protocol for weaning of ventilation [238]. Prasad *et al.* [218] applied off-policy RL algorithms, FQI-ERT and with feed forward neural networks, to determine the best weaning time of invasive mechanical ventilation, and the associated personalized sedation dosage. The policies learned showed promise in recommending weaning protocols with improved outcomes, in terms of minimizing rates of reintubation and regulating physiological stability. Targeting at the same problem as [218], Jagannatha *et al.* [221] analyzed the properties and limitations of standard off-policy evaluation methods in RL and discussed possible extensions to them in order to improve their utility in clinical domains. More recently, Yu *et al.* applied Bayesian inverse RL [219] and Supervised-actor-critic [220] to learn a suitable ventilator weaning policy from real trajectories in retrospective ICU data. RL has been also used in the development of optimal policy for the ordering of lab tests in ICUs [222], [223], and prevention and treatments for *graft versus host disease* (GVHD) [224], [225] using data set from the *Center for International Bone Marrow Transplant Research (CIBMTR)* registry database.

V. AUTOMATED MEDICAL DIAGNOSIS

Medical diagnosis is a mapping process from a patient's information such as treatment history, current signs and symptoms to an accurate clarification of a disease. Being a complex task, medical diagnosis often requires ample medical investigation on the clinical situations, causing significant

cognitive burden for clinicians to assimilate valuable information from complex and diverse clinical reports. It has been reported that diagnostic error accounts for as high as 10% of deaths and 17% of adverse events in hospitals [239]. The error-prone process in diagnosis and the necessity to assisting the clinicians for a better and more efficient decision making urgently call for a significant revolution of the diagnostic process, leading to the advent of automated diagnostic era that is fueled by advanced big data analysis and machine learning techniques [240], [241], [242].

Normally formulated as a supervised classification problem, existing machine learning methods on clinical diagnosis heavily rely on a large number of annotated samples in order to infer and predict the possible diagnoses [243], [244], [245]. Moreover, these methods have limits in terms of capturing the underlying dynamics and uncertainties in the diagnosing process and considering only a limited number of prediction labels [246]. To overcome these issues, researchers are increasingly interested in formulating the diagnostic inferencing problem as a sequential decision making process and using RL to leverage a small amount of labeled data with appropriate evidence generated from relevant external resources [246]. The existing research can be classified into two main categories, according to the type of clinical data input into the learning process: the structured medical data such as physiological signals, images, vital signs and lab tests, and the unstructured data of free narrative text such as laboratory reports, clinical notes and summaries.

A. Structured Medical Data

The most successful application of RL in diagnosis using structured data pertains to various processing and analysis tasks in medical image examination, such as feature extracting, image segmentation, and object detection/localization/tracing [247], [248]. Sahba *et al.* [249], [250], [251], [252] applied basic Q-learning to the segmentation of the prostate in transrectal *ultrasound images* (UI). Liu and Jiang [253] used a DRL method, *Trust Region Policy Optimization* (TRPO), for joint surgical gesture segmentation and classification. Ghesu *et al.* [254] applied basic DQN to automatic landmark detection problems, and achieved more efficient, accurate and robust performance than state-of-the-art machine learning and deep learning approaches on 2D *Magnetic Resonance Images* (MRI), UI and 3D *Computed Tomography* (CT) images. This approach was later extended to exploit multi-scale image representations for large 3D CT scans [255], and consider incomplete data [256] or nonlinear multi-dimensional parametric space in MRI scans of the brain region [257].

Alansary *et al.* evaluated different kinds of DRL methods (*DQN*, *Double DQN* (DDQN), *Duel DQN*, and *Duel DDQN*) [12] for anatomical landmark localization in 3D fetal UI [258], and automatic standard view plane detection [259]. Al and Yun [260] applied AC based direct PS method for aortic valve

landmarks localization and left atrial appendage seed localization in 3D CT images. Several researchers also applied DQN methods in 3D medical image registration problems [261], [262], [263], active breast lesion detection from dynamic contrast-enhanced MRI [264], and robust vessel centerline tracing problems in multi-modality 3D medical volumes [265].

Netto *et al.* [266] presented an overview of work applying RL in medical image applications, providing a detailed illustration of particular use of RL for lung nodules classification. The problem of classification is modeled as a sequential decision making problem, in which each state is defined as the combination of five 3D geometric measurements, the actions are random transitions between states, and the final goal is to discover the shortest path from the pattern presented to a known target of a malignant or a benign pattern. Preliminary results demonstrated that the Q-learning method can effectively classify lung nodules from benign and malignant directly based on lung lesions CT images.

Fakih and Das [267] developed a novel RL-based approach, which is capable of suggesting proper diagnostic tests that optimize a multi-objective performance criterion accounting for issues of costs, morbidity, mortality and time expense. To this end, some diagnostic decision rules are first extracted from current medical databases, and then the set of possible testing choices can be identified by comparing the state of patient with the attributes in the decision rules. The testing choices and the combined overall performance criterion then serve as inputs to the core RL module and the VI algorithm is applied for obtaining optimized diagnostic strategies. The approach was evaluated on a sample diagnosing problem of *solitary pulmonary nodule* (SPN) and results verified its success in improving testing strategies in diagnosis, compared with several other fixed testing strategies.

B. Unstructured Medical Data

Unlike the formally structured data that are directly machine understandable, large proportions of clinical information are stored in a format of unstructured free text that contains a relatively more complete picture of associated clinical events [3]. Given their expressive and explanatory power, there is great potential for clinical notes and narratives to play a vital role in assisting diagnosis inference in an underlying clinical scenario. Moreover, limitations such as knowledge incompleteness, sparsity and fixed schema in structured knowledge have motivated researchers to use various kinds of unstructured external resources such as online websites for related medical diagnosing tasks [246].

Motivated by the *Text REtrieval Conference-Clinical Decision Support* (TREC-CDS) track dataset [268], diagnosis inferencing from unstructured clinical text has gained much attention among AI researchers recently. Utilizing particular natural language processing techniques to extract useful information from clinical text, RL has been used to optimize the diagnosis

inference procedure in several studies. Ling *et al.* [246], [269] proposed a novel clinical diagnosis inferencing approach that applied DQN to incrementally learn about the most appropriate clinical concepts that best describe the correct diagnosis by using evidences gathered from relevant external resources (from Wikipedia and MayoClinic). Experiments on the TRECCDS datasets demonstrated the effectiveness of the proposed approach over several non RL-based systems.

Exploiting real datasets from the *Breast Cancer Surveillance Consortium* (BCSC) [270], Chu *et al.* [271] presented an adaptive online learning framework for supporting clinical breast cancer diagnosis. The framework integrates both supervised learning models for breast cancer risk assessment and RL models for decision-making of clinical measurements. The framework can quickly update relevant model parameters based on current diagnosis information during the training process. Additionally, it can build flexible fitted models by integrating different model structures and plugging in the corresponding parameters during the prediction process. The authors demonstrated that the RL models could achieve accurate breast cancer risk assessment from sequential data and incremental features.

In order to facilitate self-diagnosis while maintaining reasonable accuracy, the concept of *symptom checking* (SC) has been proposed recently. SC first inquires a patient with a series of questions about their symptoms, and then attempts to diagnose some potential diseases [245]. Tang *et al.* [272] formulated inquiry and diagnosis policies as an MDP, and adopted DQN to learn to inquire and diagnose based on limited patient data. Kao *et al.* [273] applied context-aware HRL scheme to improve accuracy of SC over traditional systems making a limited number of inquiries. Empirical studies on a simulated dataset showed that the proposed model drastically improved disease prediction accuracy by a significant margin. The SC system was successfully employed in the DeepQ

Tricorder which won the second prize in the Qualcomm Tricorder XPRIZE competition in year 2017 [274], [275].

A dialogue system was proposed in [276] for automatic diagnosis, in which the medical dataset was built from a pediatric department in a Chinese online healthcare community. The dataset consists of self-reports from patients and conversational data between patients and doctors. A DQN approach was then used to train the dialogue policy. Experiment results showed that the RL-based dialogue system was able to collect symptoms from patients via conversation and improve the accuracy for automatic diagnosis. In order to increase the efficiency of the dialogue systems, Tang *et al.* [277] applied DQN framework to train an efficient dialogue agent to sketch disease-specific lexical probability distribution, and thus to converse in a way that maximizes the diagnosis accuracy and minimizes the number of conversation turns. The dialogue system was evaluated on the mild cognitive impairment diagnosis from a real clinical trial, and results showed that the RL-driven framework could significantly

outperform state-of-the-art supervised learning approaches using only a few turns of conversation.

VI. OTHER HEALTHCARE DOMAINS

Besides the above applications of RL in DTR design and automated medical diagnosis, there are many other case applications in broader healthcare domains that focus on problems specifically in health resource scheduling and allocation, optimal process control, drug discovery and development, as well as health management.

(1) *Health Resource Scheduling and Allocation*. The healthcare system is a typical service-oriented system where customers (e.g., patients) are provided with service using limited resources, e.g. the time slots, nursing resources or diagnostic devices [278]. *Business process management* (BPM) plays a key role in such systems as the objective of the service provider is to maximize profit overtime, considering various customer classes and service types with dynamics or uncertainties such as cancellations or no-shows of patients [279], [280]. Since the optimal resource allocation problem in BPM can be seen as a sequential decision making problem, RL is then naturally suitable for offering reasonable solutions. Huang *et al.* [279] formulated the allocation optimization problems in BPM as an MDP and used basic Q-learning algorithm to derive an optimal solution. The RL-based approach was then applied to address the problem of optimizing resource allocation in radiology CT-scan examination process. A heuristic simulationbased approximate DP approach was proposed in [278], which considered both stochastic service times and uncertain future arrival of clients. The experimental investigation using data from the radiological department of a hospital indicated an increases of 6.9% in the average profit of the hospital and 9% in the number of examinations. Gomes [281] applied a DRL method, *Asynchronous Advantage Actor Critic* (A3C) [282], to schedule appointments in a set of increasingly challenging environments in primary care systems.

(2) *Optimal Process Control*. RL has also been widely applied in deriving an optimal control policy in a variety of healthcare situations, ranging from surgical robot operation [283], [284], [285], [286], [287], *functional electrical stimulation* (FES) [288], [289], and adaptive rate control for medical video streaming [290], [291]. Li and Burdick [283] applied RL to learn a control policy for a surgical robot such that the robot can conduct some basic clinical operations automatically. A function approximation based IRL method was used to derive an optimal policy from experts' demonstrations in high dimensional sensory state space. The method was applied to the evaluation of surgical robot operators in three clinical tasks of knot tying, needling passing and suturing. Thananjeyan *et al.* [284] and Nguyen *et al.* [285] applied DRL algorithm, TRPO, in learning tensioning policies effectively for surgical gauze cutting. Chen *et al.* [286] combined programming by demonstration and RL for motion control of flexible manipulators in minimally invasive surgical

performance, while Baek *et al.* [287] proposed the use of RL to perform resection automation of cholecystectomy by planning a path that avoids collisions in a laparoscopic surgical robot system.

FES employs neuroprosthesis controllers to apply electrical current to the nerves and muscles of individuals with spinal cord injuries for rehabilitative movement [292]. RL has been used to calculate stimulation patterns to efficiently adapt the control strategy to a wide range of time varying situations in patients' preferences and reaching dynamics. AC-based control strategies [293], [294], [289] were proposed to evaluate target-oriented task performed using a planar musculoskeletal human arm in FES. To solve the reward learning problem in large state spaces, an IRL approach was proposed in [288] to evaluate the effect of rehabilitative stimulations on patients with spinal cord injuries based on the observed patient motions.

RL-based methods have also been widely applied in adaptive control in mobile health medical video communication systems. For example, Istepanian *et al.* [290] proposed a new rate control algorithm based on Q-learning that satisfies medical quality of service requirements in bandwidth demanding situations of ultrasound video streaming. Alinejad [291] applied Q-learning for cross-layer optimization in real-time medical video streaming.

(3) *Drug Discovery and Development.* Drug discovery and development is a time-consuming and costly process that usually lasts for 10-17 years, but with as low as around 10% overall probability of success [295]. To search an effective molecule that meets the multiple criteria such as bioactivity and synthetic accessibility in a prohibitively huge synthetically feasible molecule space is extremely difficult. By using computational methods to virtually design and test molecules, *de novo* design offers ways to facilitate cycle of drug development [296]. It is until recent years that RL methods have been applied in various aspects of *de novo* design for drug discovery and development. Olivecrona [297] used RL to fine tune the recurrent neural network in order to generate molecules with certain desirable properties through augmented episodic likelihood. Serrano *et al.* [298] applied DQN to solve the protein-ligand docking prediction problem, while Neil *et al.* [299] investigated the PPO method in molecular generation.

More recently, Popova *et al.* [300] applied DRL methods to generate novel targeted chemical libraries with desired properties.

(4) *Health Management.* As a typical application domain, RL has also been used in adaptive interventions to support health management such as promoting physical activities for diabetic patients [301], [302], or weight management for obesity patients [303], [304]. In these applications, throughout continuous monitoring and communication of mobile health, personalized intervention policies can be derived to input the monitored measures and output when, how and which plan to deliver. A notable work was by Yom *et al.* [301], who applied

RL to optimize messages sent to the users, in order to improve their compliance with the activity plan. A study of 27 sedentary diabetes type 2 patients showed that participants who received messages generated by the RL algorithm increased the amount of activity and pace of walking, while the patients using static policy did not. Patients assigned to the RL algorithm group experienced a superior reduction in blood glucose levels compared to the static control policies, and longer participation caused greater reductions in blood glucose levels.

VII. CHALLENGES AND OPEN ISSUES

The content above has summarized the early endeavors and continuous progress of applying RL in healthcare over the past decades. Focus has been given to the vast variety of application domains in healthcare. While notable success has been obtained, the majority of these studies simply applied existing naive RL approaches in solving healthcare problems in a relatively simplified setting, thus exhibiting some common shortcomings and practical limitations. This section discusses several challenges and open issues that have not been properly addressed by the current research, from perspectives of how to deal with the basic components in RL (i.e., formulation of states, actions and rewards, learning with world models, and evaluation of policies), and fundamental theoretical issues in traditional RL research (i.e., the exploration-exploitation tradeoff and credit assignment problem).

A. State/Action Engineering

The first step in applying RL to a healthcare problem is determining how to collect and pre-process proper medical data, and summarize such data into some manageable state representations in a way that sufficient information can be retained for the task at hand. Selecting the appropriate level of descriptive information contained in the states is extremely important. On one hand, it would be better to contain as detailed information as possible in the states, since this complete information can provide a greater distinction among patients. On the other hand, however, increasing the state space makes the model become more difficult to solve. It is thus essential that a good state representation include any compulsory factors or variables that causally affect both treatment decisions and the outcomes. Previous studies have showed that, to learn an effective policy through observational medical data, the states should be defined in a way that to the most approximates the behavior policy that has generated such data [305], [306].

However, data in medical domains often exhibit notable biases or noises that are presumably varying among different clinicians, devices, or even medical institutes, reflecting comparable inter-patient variability [92]. For some complex diseases, clinicians still face inconsistent guides in selecting

exact data as the state in a given case [191]. In addition, the notorious issue of missing or incomplete data can further exaggerate the problem of data collection and state representation in medical settings, where the data can be collected from patients who may fail to complete the whole trial, or the number of treatment stages or timing of initializing the next line of therapy is flexible. This missing or censoring data will tend to increase the variance of estimates of the value function and thus the policy in an RL setting. While the missing data problem can be generally solved using various imputation methods that sample several possible values from the estimated distribution to fill in missing values, the censoring data problem is far more challenging, calling for more sophisticated techniques for state representation and value estimation in such flexible settings [96], [97].

Most existing work defines the states over the processed medical data with raw physiological, pathological, and demographics information, either using simple discretization methods to enable storage of value function in tabular form, or using some kinds of function approximation models (e.g., linear models or deep neural models). While this kind of state representation is simple and easy to implement, the rich temporal dependence or causal information, which is the key feature of medical data, can be largely neglected [307]. To solve this issue, various probabilistic graphical models [308] can be used to allow temporal modeling of time series medical data, such as *dynamic Bayesian networks* (DBNs), in which nodes correspond to the random variables of interest, edges indicate the relationship between these random variables, and additional edges model the time dependency. These kinds of graphical models have the desirable property that allows for interpretation of interactions between state variables or between states and actions, which is not the case for other methods such as SVMs and neural networks.

Coupled with the state representation in RL is the formulation of actions. The majority of existing work has mainly focused on discretization of the action space into limited bins of actions. Although this formulation is quite reasonable in some medical settings, such as choices in between turning on ventilation or weaning off it, there are many other situations where actions are by themselves continuous/multidimensional variables. While the simplification of discretizing medicine dosage is necessary in the early proof-of-concept stage, realizing fully continuous dosing in the original action space is imperative in order to meet the commitments of precision medicine [309]. There has been a significant achievement in the continuous control using AC methods and PS methods in the past years, particularly from the area of robotic control [20] and DRL [12]. While this achievement can provide direct solutions to this problem, selecting the action over large/infinite space is still non-trivial, especially when dealing with any sample complexity guarantees (PAC). An effective method for efficient action selection in continuous and high dimensional action spaces, while at the same time maintaining low exploration complexity

of PAC guarantees would extend the applicability of current methods to more sample-critical medical problems.

B. Reward Formulation

Among all the basic components, the reward may be at the core of an RL process. Since it encodes the goal information of a learning task, a proper formulation of reward functions plays the most crucial role in the success of RL. However, the majority of current RL applications in healthcare domains are still grounded on simple numerical reward functions that must be explicitly defined beforehand to indicate the goal of treatments by clinicians. It is true that in some medical settings, the outcomes of treatments can be naturally generated and explicitly represented in a numerical form, for example, the time elapsed, the vitals monitored, or the mortality reduced. In general, however, specifying such a reward function precisely is not only difficult but sometimes even misleading. For instance, in treatment of cancers [83], the reward function was usually decomposed into several independent or contradictory components based on some prior domain knowledge, each of which was mapped into some integer numbers, e.g., -60 as a high penalty for patient death and +15 as a bonus for a cured patient. Several threshold and weighting parameters were needed to provide a way for trading-off efficacy and toxicity, which heavily rely on clinicians' personal experience that varies from one to another. This kind of somewhat arbitrary quantifications might have significant influence on the final learned therapeutic strategies and it is unclear how changing these numbers can affect the resulting strategies.

To conquer the above limitations, one alternative is to provide the learning agent with more qualitative evaluations for actions, turning the learning into a PRL problem [310]. Unlike the standard RL approaches that are restricted to numerical and quantitative feedback, the agent's preferences instead can be represented by more general types of preference models such as ranking functions that sort states, actions, trajectories or even policies from most to least promising [51]. Using such kind of ranking functions has a number of advantages as they are more natural and easier to acquire in many applications in clinical practice, particularly, when it is easier to require comparisons between several, possibly suboptimal actions or trajectories than to explicitly specify their performance. Moreover, considering that the medical decisions always involve two or more related or contradictory aspects during treatments such as benefits versus associated cost, efficacy versus toxicity, and efficiency versus risk, it is natural to shape the learning problem into a multi-objective optimization problem. MORL techniques [50] can be applied to derive a policy that makes a trade-off between distinct objectives in order to achieve a Pareto optimal solution. Currently, there are only very limited studies in the literature that applied PRL [52], [89], [178] and MORL [177], [311], [176], [222] in medical settings, for optimal

therapy design in treatment of cancer, schizophrenia or lab tests ordering in ICUs. However, all these studies still focus on very limited application scenarios where only static preferences or fixed objectives were considered. In a medical context, the reward function is usually not a fixed term but subject to changing with regard to a variety of factors such as the time, the varying clinical situations and the evolving physiopsychic conditions of the patients. Applying PRL and MORL related principles to broader domains and considering the dynamic and evolving process of patients' preferences and treatment objectives is still a challenging issue that needs to be further explored.

A more challenging issue is regarding the inference of reward functions directly from observed behaviors or clinical data. While it is straightforward to formulate a reward function, either quantitatively or qualitatively, and then compute the optimal policy using this function, it is sometimes preferable to directly estimate the reward function of experts from a set of presumably optimal treatment trajectories in retrospective medical data. *Imitation learning*, particularly, IRL [53], [54], is one of the most feasible approaches to infer reward functions given observations of optimal behaviour. However, applying IRL in clinical settings is not straightforward, due to the inherent complexity of clinical data and its associated uncertainties during learning. The variance during the policy learning and reward learning can amplify the bias in each learning process, potentially leading to divergent solutions that can be of little use in practical clinical applications [312], [313].

Last but not the least, while it is possible to define a short-term reward function at each decision step using prior human knowledge, it would be more reasonable to provide a long-term reward only at the end of a learning episode. This is especially the case in healthcare domains where the real evaluation outcomes (e.g., decrease of patients, duration of treatment) can only be observed at the end of treatment. Learning with sparse rewards is a challenging issue that has attracted much attention in recent RL research. A number of effective approaches have been proposed, such as the *hindsight experience replay* [313], the *unsupervised auxiliary learning* [314], the *imagination-augmented learning* [315], and the *reward shaping* [234]. While there have been several studies that address the sparse reward problem in healthcare domains, most of these studies only focus on DTRs with a rather short horizon (typically three or four steps). Moreover, previous work has showed that entirely ignoring short-term rewards (e.g. maintaining hourly physiologic blood pressure for sepsis patients) could prevent from learning crucial relationships between certain states and actions [307]. How to tackle sparse reward learning with a long horizon in highly dynamic clinical environments is still a challenging issue in both theoretical and practical investigations of RL in healthcare.

C. Policy Evaluation

The process of estimating the value of a policy (i.e., target policy) with data collected by another policy (i.e., behavior policy) is called *off-policy evaluation* problem [14]. This problem is critical in healthcare domains because it is usually infeasible to estimate policy value by running the policy directly on the target populations (i.e., patients) due to high cost of experiments, uncontrolled risks of treatments, or simply unethical/illegal humanistic concerns. Thus, it is needed to estimate how the learned policies might perform on retrospective data before testing them in real clinical environments. While there is a large volume of work in RL community that focuses on *importance sampling* (IS) techniques and how to trade off between bias and variance in IS-based offpolicy evaluation estimators (e.g., [316]), simply adopting these estimators in healthcare settings might be unreliable due to issues of sparse rewards or large policy discrepancy between RL learners and physicians. Using sepsis management as a running example, Gottesman *et al.*, [305] discussed in detail why evaluation of policies using retrospective health data is a fundamentally challenging issue. They argued that any inappropriate handling of state representation, variance of IS-based statistical estimators, and confounders in more adhoc measures would result in unreliable or even misleading estimates of the quality of a treatment policy. The estimation quality of the off-policy evaluation is critically dependent on how precisely the behaviour policy is estimated from the data, and whether the probabilities of actions under the approximated behaviour policy model represent the true probabilities [306]. While the main reasons have been largely unveiled, there is still little work on effective policy evaluation methods in healthcare domains. One recent work is by Li *et al.* [201], who provided an off-policy POMDP learning method to take into account uncertainty and history information in clinical applications. Trained on real ICU data, the proposed policy was capable of dictating near-optimal dosages in terms of vasopressor and intravenous fluid in a continuous action space for sepsis patients.

D. Model Learning

In the efficient techniques described in Section II-B, modelbased methods enable improved sample efficiency over modelfree methods by learning a model of the transition and reward functions of the domain on-line and then planing a policy using this model [43]. It is surprising that there are quite limited model-based RL methods applied in healthcare in the current literature [196], [193], [197]. While a number of modelbased RL algorithms have been proposed and investigated in the RL community (e.g., R-max [25], E^3 [317]), most of these algorithms assume that the agent operates in small domains with a discrete state space, which is

contradictory to the healthcare domains usually involving multi-dimensional continuously valued states and actions. Learning and planning over such large scale continuous models would cause additional challenges for existing model-based methods [43]. A more difficult problem is to develop efficient exploration strategies in continuous action/state space [29]. By deriving a finite representation of the system that both allows efficient planning and intelligent exploration, it is potential to solve the challenging model learning tasks in healthcare systems more efficiently than contemporary RL algorithms.

E. Exploration Strategies

Exploration plays a core role in RL, and a large amount of effort has been devoted to this issue in the RL community. A wealth of exploration strategies have been proposed in the past decades. Surprisingly, the majority of existing RL applications in healthcare domains simply adopt simple heuristic-based exploration strategies (i.e., ϵ -greedy strategy). While this kind of handling exploration dilemmas has made notable success, it becomes infeasible in dealing with more complicated dynamics and larger state/action spaces in medical settings, causing either a large sample complexity or an asymptotic performance far from the optimum. Particularly, in cases of an environment where only a rather small percentage of the state space is reachable, naive exploration from the entire space would be quite inefficient. This problem is getting more challenging in continuous state/action space, for instance, in the setting of HIV treatment [142], where the basin of attraction of the healthy state is rather small compared to that of the unhealthy state. It has been shown that traditional exploration methods are unable to obtain obvious performance improvement and generate any meaningful treatment strategy even after a long period of search in the whole space [150], [151]. Therefore, there is a justifiable need for strategies that can identify dynamics during learning or utilize a performance measure to explore smartly in high dimensional spaces. In recent years, several more advanced exploration strategies have been proposed, such as PAC guaranteed exploration methods targeting at continuous spaces [150], [151], concurrent exploration mechanisms [318], [319], [320] and exploration in deep RL [321], [322], [323]. It is thus imperative to incorporate such exploration strategies in more challenging medical settings, not only to decrease the sample complexity significantly, but more importantly to seek out new treatment strategies that have not been discovered before.

Another aspect of applying exploration strategies in healthcare domains is the consideration of true cost of exploration. Within the vanilla RL framework, whenever an agent explores an inappropriate action, the consequent penalty acts as a negative reinforcement in order to discourage the wrong action. Although this procedure is appropriate for most situations, it may be problematic in some environments where the consequences of wrong actions are not limited to bad performance, but can result in unrecoverable effects. This

is obviously true when dealing with patients in healthcare domains: although we can reset a robot when it has fallen down, we cannot bring back to life when a patient has been given a fatal medical treatment. Consequently, methods for safe exploration are of great real world interest in medical settings, in order to preclude unwanted, unsafe actions [32], [33], [324].

F. Credit Assignment

Another important aspect of RL is the credit assignment problem that decides when an action or which actions is responsible for the learning outcome after a sequence of decisions. This problem is critical as the evaluation of whether an action being “good” or “bad” usually cannot be decided upon right away, but until the final goal has been achieved by the agent. As each action at each step contributes more or less to the final performance of success or failure, it is thus necessary to give distinct credit to the actions along the whole path, giving rise to the difficult problem of *temporal credit assignment* problem. A related problem is the *structural credit assignment* problem, in which the problem is to distribute feedback over the multiple candidates (e.g., multiple concurrently learning agents, action choices, or structure representations of the agent’s policy).

The *temporal credit assignment* problem is more prominent in healthcare domains as the effect of treatments can be much varied or delayed. Traditional RL research tackles the credit assignment problem using simple heuristics such as eligibility traces that weigh the past actions according to how far the time has elapsed (i.e., the backward view), or discount factors that weigh the future events according to how far away they will happen (i.e., the forward view) [14]. These kinds of fixed and simplified heuristics are incapable of modelling more complex interaction modes in a medical situation. As a running example of explaining changes in blood glucose of a person with type 1 diabetes mellitus [325], it is difficult to give credit to the two actions of doing exercise in the morning or taking insulin after lunch, both of which can potentially cause hypoglycemia in the afternoon. Since there are many factors to affect blood glucose and the effect can take place after many hours, e.g., moderate exercise can lead to heightened insulin sensitivity for up to 22 hours, simply assigning an eligibility trace that decays with time elapsed is thus unreasonable, misleading or even incorrect. How to model the time-varying causal relationships in healthcare and incorporate them into the learning process is therefore a challenging issue that requires more investigations. The abundant literature in causal explanation [326] and inference [327] can be introduced to provide a more powerful causal reasoning tool to the learning algorithm. By producing hypothesized sequences of causal mechanisms that seek to explain or predict a set of real or counterfactual events which have been observed or manipulated [328], not only can the learning performance be potentially improved, but also more

explainable learned strategies can be derived, which is ultimately important in healthcare domains.

VIII. FUTURE PERSPECTIVES

We have discussed a number of major challenges and open issues raised in the current applications of RL techniques in healthcare domains. Properly addressing these issues are of great importance in facilitating the adoption of any medical procedure or clinical strategy using RL. Looking into the future, there is an urgent need in bringing recent development in both theories and techniques of RL together with the emerging clinical requirements in practice so as to generate novel solutions that are more interpretable, robust, safe, practical and efficient. In this section, we briefly discuss some of the future perspectives that we envision the most critical towards realizing such ambitions. We mainly focus on three theoretical directions: the interpretability of learned strategies, the integration of human or prior domain knowledge, and the capability of learning from small data. Healthcare under ambient intelligence and real-life applications are advocated as two main practical directions for RL applications in the coming age of intelligent healthcare.

A. Interpretable Strategy Learning

Perhaps one of the most profound issues with modern machine learning methods, including RL, is the lack of clear interpretability [329]. Usually functioning as a black box expressed by, for instance, deep neural networks, models using RL methods receive a set of data as input and directly output a policy which is difficult to interpret. Although impressive success has been made in solving challenging problems such as learning to play Go and Atari games, the lack of interpretability renders the policies unable to reveal the real correlation between features in the data and specific actions, and to impose and verify certain desirable policy properties, such as worst-case guarantees or safety constraints, for further policy debugging and improvement [330]. These limits therefore greatly hinder the successful adoption of RL policies for safety-critical applications such as in medical domains as clinicians are unlikely to try new treatments without rigorous validation for safety, correctness and robustness [37], [331].

Recently, there has been growing interest in attempting to address the problem of interpretability in RL algorithms. There are a variety of ways to realize interpretability of learned policy, by either using small, closed-form formulas to compute index-based policies [332], using program synthesis to learn higher-level symbolic interpretable representations of learned policies [333], utilizing genetic programming for interpretable policies represented by compact algebraic equations [36], or using program verification techniques to verify certain properties of the programs which are represented as decision trees [37]. Also, there has been growing attention not only on developing interpretable representations, but also on generating explicit explanations for sequential decision

making problems [334]. While several works specifically focused on interpretability of deep models in healthcare settings [335], [336], how to develop interpretable RL solutions in order to increase the robustness, safety and correctness of learned strategies in healthcare domains is still an unsolved issue that calls for further investigations.

B. Integration of Prior Knowledge

There is a wealth of prior knowledge in healthcare domains that can be used for learning performance improvement. The integration of such prior knowledge can be conducted in different manners, either through configuration or presentation of learning parameters, components or models [337], [135], knowledge transfer from different individual patients, subtypes/sub-populations or clinical domains [147], or enabling human-in-the-loop interactive learning [338].

Gaweda *et al.* [337], [135] presented an approach to management of anemia that incorporates a critical prior knowledge about the doseresponse characteristic into the learning approach, that is, for all patients, it is known that the dose-response curve of HGB vs. EPO is monotonically nonincreasing. Thus, if a patient's response is evaluated as insufficient for a particular dose at a particular state, then the physician knows that the optimal dose for that state is definitely higher than the administered one. Consequently, there is no need to explore the benefit of lower doses at further stages of treatment. To capture this feature, the authors introduced an additional mechanism to the original Q-learning algorithm so that the information about monotonically increasing character of the HGB vs. EPO curve can be incorporated in the update procedure. This modification has been shown to make the EPO dosing faster and more efficiently.

While transfer learning has been extensively studied in the agent learning community [46], there is quite limited work on applying TRL techniques in healthcare settings. The learning performance in the target task can be potentially facilitated by using latent variable models, pre-trained model parameters from past tasks, or directly learning a mapping between past and target tasks, thus extending personalized care to groups of patients with similar diagnoses. Marivate *et al.* [145] highlighted the potential benefit of taking into account individual variability and data limitations when performing batch policy evaluation for new individuals in HIV treatment. A recent approach on TRL using latent variable models was proposed by Killian *et al.* [146], [147], who used a Gaussian Process latent variable model for HIV treatment by both inferring the transition dynamics within a task instance and also in the transfer between task instances.

Another way of integrating prior knowledge into an RL process can be making use of the human cognitive abilities or domain expertise to guide, shape, evaluate or validate the agent's learning process, making the traditional RL into a human-in-the-loop interactive RL problem [339]. Human knowledge-driven RL methods can be of great interest to

problems in healthcare domains, where traditional learning algorithms would possibly fail due to issues such as insufficient training samples, complex and incomplete data or unexplainable learning process [338]. Consequently, the integration of the humans (i.e., doctors) into the learning process, and the interaction of an expert's knowledge with the automatic learning data would greatly enhance the knowledge discovery process [340]. While there is some previous work from other domains, particularly in training of robots [341], [342], human-in-the-loop interactive RL is not yet well established in the healthcare domain. It remains open for future research to transfer the insights from existing studies into the healthcare domain to ensure successful applications of existing RL methods.

C. Learning from Small Data

There is no doubt that the most recent progresses of RL, particularly DRL, are highly dependent on the premise of large number of training samples. While this is quite reasonable conceptually, that is, we cannot learn new things that we have not tried sufficiently enough, there still exist many domains lacking sufficient available training samples, specifically, in some healthcare domains [343]. For example, in diagnose settings, medical images are much more difficult to be annotated with certain lesions in high-quality without specific expertise compared to general images with simple categories. In addition, there are usually few historical data or cases for new diseases and rare illness, making it impossible to obtain sufficient training samples with accurate labels. In such circumstances, directly applying existing RL methods on limited data may result in overly optimistic, or in other extreme, pessimistic about treatments that are rarely performed in practice.

Broadly, there are two different ways of dealing with a small sample learning problem [344]. The direct solution can be using data augmentation strategies such as deformations [345] or GANs [346] to increase samples and then employ conventional learning methods. The other type of solutions can be applying various model modification or domain adaptation methods such as knowledge distillation [347] or meta-learning [348] to enable efficient learning that overcomes the problem of data scarcity. While still in its early stage, significant progress has been made in small sample learning research in recent years [344]. How to build on these achievements and tackle the small data RL problems in healthcare domains thus calls for new methods of future investigations. One initial work is by Tseng *et al.* [93] who developed automated radiation adaptation protocols for NSCLC patients by using GAN to generate synthetic patient data and DQN to learn dose decisions with the synthesized data and the available real clinical data. Results showed that the learned dose strategies by DQN were capable of achieving similar results to those chosen by clinicians, yielding feasible

and quite promising solutions for automatic treatment designs with limited data.

D. Healthcare under Ambient Intelligence

The recent development in sensor networks and wearable devices has facilitated the advent of new era of healthcare systems that are characterized by low-cost mobile sensing and pervasive monitoring within the home and outdoor environments [349]. The *Ambient Intelligence* (Aml) technology, which enables innovative human-machine interactions through unobtrusive and anticipatory communications, has the potential to enhance the healthcare domain dramatically by learning from user interaction, reasoning about users' goals and intentions, and planning activities and future interactions [350]. By using various kinds of sensing devices, such as smart phones, GPS and body sensors monitoring motions and activities, it is now possible to remotely and continuously collect patients' health information such that proper treatment or intervention decisions can be made anytime and anywhere. As an instance of online decision making in a possibly infinite horizon setting involving many stages of interventions, RL plays a key role in achieving the future vision of Aml in healthcare systems through continuous interaction with the environment and adaption to the user needs in a transparent and optimal manner. In fact, the high level of monitoring and sensing provides ample opportunity for RL methods that can fuse estimates of a given physiologic parameter from multiple sources to provide a single measurement, and derive optimal strategies using these data. Currently, there are several studies that have applied RL to achieve Aml in healthcare domains. For example, RL has been used to adapt the intervention strategies of smart phones in order to recommend regular physical activity to people who suffer from diabetes type 2 [301], [302], or who have experienced a cardiac event and been in cardiac rehab [351], [352], [353]. It has been also been used for mobile health intervention for college students who drink heavily and smoke cigarettes [181].

Despite the successes, healthcare under Aml poses some unique challenges that preclude direct application of existing RL methodologies for DTRs. For example, it typically involves a large number of time points or infinite time horizon for each individual; the momentary signal may be weak and may not directly measure the outcome of interest; and estimation of optimal treatment strategies must be done online as data accumulate. How to tackle these issues is of great importance in the successful applications of RL methods in the advent of healthcare systems under Aml.

E. Future in-vivo Studies

To date, the vast volume of research reporting the development of RL techniques in healthcare is built upon certain computational models that leverages mathematical

representation of how a patient responds to given treatment policies, or upon retrospective clinical data to directly derive appropriate treatment strategies. While this kind of *in silico* study is essential as a tool for early stage exploration or direct derivation of adaptive treatment strategies by providing approximate or highly simplified models, future *in vivo* studies of closed-loop RL approaches are urgently required to reliably assess the performance and personalization of the proposed approaches in real-life implementations. However, a number of major issues still remain related to, in particular, data collection and preprocessing in real clinical settings, and high inter-individual differences of the physiological responses, thus calling for careful consideration of safety, efficiency and robustness of RL methods in real-life healthcare applications.

First and foremost, safety is of paramount importance in medical settings, thus it is imperative to ensure that the actions during learning be safe enough when dealing with *in vivo* subjects. In some healthcare domains, the consequences of wrong actions are not merely limited to bad performance, but may include long-term effects that cannot be compensated by more profitable exploitation later on. As one wrong action can result in unrecoverable effects, learning in healthcare domains poses a safety exploration dilemma [324], [32]. It is worth noting that there are substantial ongoing efforts in the computer science community to address precisely these problems, namely in developing risk-directed exploration algorithms that can efficiently learn with formal guarantees regarding the safety (or worst-case performance) of the system [33]. With this consideration, the agent's choice of actions is aided by an appropriate risk metric acting as an exploration bonus toward safer regions of the search space. How to draw on these achievements and develop safe exploration strategies is thus urgently required to implement RL methods in real-life healthcare applications.

Another issue is regarding the sample efficiency of RL methods in *in vivo* studies [83]. While it is possible for the RL algorithms to collect large numbers of samples in simulations, it is unrealistic for sample-critical domains where collecting samples would cause significant cost. This is obviously true when dealing with real patients who would possibly not survive the long-term repeated trial-and-error treatment. Luckily, the wide range of efficient techniques reviewed in Section II-B can provide promising solutions to this problem. Specifically, the sample-level batch learning methods can be applied for more efficient use of past samples, while modelbased methods enable better use of samples by building the model of the environment. Another appealing solution is using task-level transfer methods that can reuse the past treatment or patient information to facilitate learning in new cases, or directly transfer the learned policies in simulations to real environments. To enable efficient transfer, RL algorithms can be provided with initial knowledge that can direct the learning in its initial stage toward more profitable and safer regions of the state space, or with demonstrations and teacher advice from an external expert that can interrupt exploration

and provide expert knowledge when the agent is confronted with unexpected situations.

The last issue in the real-life implementation of RL approaches is regarding the robustness of derived solutions. Despite inherently being suitable for optimizing outcomes in stochastic processes with uncertainty, existing RL methods are still facing difficulties in handling incomplete or noisy state variables in partially observable real healthcare environments, and in providing measures of confidence (e.g. standard errors, confidence sets, hypothesis tests). Uncertainty can also be caused by the MDP parameters themselves, which leads to significant increases in the difficulty of the problem, in terms of both computational complexity and data requirements. While there has been some recent work on robust MDP solutions which accounts for this issue [34], [35], a more general and sound theoretical and empirical evaluation is still lacking. Moreover, most current studies are built upon predefined functions to map states and actions into some integer numbers. It is unclear how changing these numbers would affect the resulting optimal solutions. Understanding the robustness of RL methods in uncertain healthcare settings is the subject of ongoing critical investigations by the statistics, computer science and healthcare communities.

IX. CONCLUSIONS

RL presents a mathematically solid and technically sound solution to optimal decision making in various healthcare tasks challenged with noisy, multi-dimensional and incomplete data, nonlinear and complex dynamics, and particularly, sequential decision procedures with delayed evaluation feedback. This paper aims to provide a state-of-the-art comprehensive survey of RL applications to a variety of decision making problems in the area of healthcare. We have provided a structured summarization of the theoretical foundations and key techniques in the RL research from traditional machine learning perspective, and surveyed the broad-ranging applications of RL methods in solving problems affecting manifold areas of healthcare, from DTRs in chronic diseases and critical care, automated clinical diagnosis, to other healthcare domains such as clinical resource allocation and scheduling. The challenges and open issues in the current research have been discussed in detail from the perspectives of basic components constituting an RL process (i.e., states, actions, rewards, policies and models), and fundamental issues in RL research (i.e., the exploration-exploitation dilemma and credit assignment). It should be emphasized that, although each of these challenging issues has been investigated extensively in the RL community for a long time, achieving remarkably successful solutions, it might be problematic to directly apply these solutions in the healthcare settings due to the inherent complexity in processes of medical data processing and policy learning. In fact, the unique features embodied in the clinical or medical decision making process urgently call for development of more advanced RL

methods that are really suitable for real-life healthcare problems. Apart from the enumerated challenges, we have also pointed out several perspectives that remain comparatively less addressed by the current literature. Interpretable learning, transfer learning as well as small-data learning are the three theoretical directions that require more effort in order to make substantial progress. Moreover, how to tailor the existing RL methods to deal with the pervasive data in the new era of Aml healthcare systems and take into consideration safety, robustness and efficiency caused by real-life applications are two main paradigms that need to be carefully handled in practice.

The application of RL in healthcare is at the intersection of computer science and medicine. Such cross-disciplinary research requires a concerted effort from machine learning researchers and clinicians who are directly involved in patient care and medical decision makings. While notable success has been obtained, RL has still received far less attention by researchers, either from computer science or from medicine, compared to other research paradigms in healthcare domains, such as traditional machine learning, deep learning, statistical learning and control-driven methods. Driven by both substantial progress in theories and techniques in the RL research, as well as practical demands from healthcare practitioners and managers, this situation is now changing rapidly and recent years have witnessed a surge of interest in the paradigm of applying RL in healthcare, which can be supported by the dramatic increase in the number of publications on this topic in the past few years. Serving as the first comprehensive survey of RL applications in healthcare, this paper aims at providing the research community with systematic understanding of foundations, broad palette of methods and techniques available, existing challenges, and new insights of this emerging paradigm. By this, we hope that more researchers from various disciplines can utilize their expertise in their own area and work collaboratively to generate more applicable solutions to optimal decision makings in healthcare.

ACKNOWLEDGMENT

This work is supported by the Hongkong Scholar Program under Grant XJ2017028.

REFERENCES

- [1] V. L. Patel, E. H. Shortliffe, M. Stefanelli, P. Szolovits, M. R. Berthold, R. Bellazzi, and A. Abu-Hanna, "The coming of age of artificial intelligence in medicine," *Artificial Intelligence in Medicine*, vol. 46, no. 1, pp. 5–17, 2009.
- [2] S. E. Dilsizian and E. L. Siegel, "Artificial intelligence in medicine and cardiac imaging: harnessing big data and advanced computing to provide personalized medical diagnosis and treatment," *Current Cardiology Reports*, vol. 16, no. 1, p. 441, 2014.
- [3] F. Jiang, Y. Jiang, H. Zhi, Y. Dong, H. Li, S. Ma, Y. Wang, Q. Dong, H. Shen, and Y. Wang, "Artificial intelligence in healthcare: past, present and future," *Stroke and Vascular Neurology*, vol. 2, no. 4, pp. 230–243, 2017.
- [4] J. He, S. L. Baxter, J. Xu, J. Xu, X. Zhou, and K. Zhang, "The practical implementation of artificial intelligence technologies in medicine," *Nature Medicine*, vol. 25, no. 1, p. 30, 2019.
- [5] A. E. Johnson, M. M. Ghassemi, S. Nemati, K. E. Niehaus, D. A. Clifton, and G. D. Clifford, "Machine learning and decision support in critical care," *Proceedings of the IEEE*, vol. 104, no. 2, pp. 444–466, 2016.
- [6] D. Ravì, C. Wong, F. Deligianni, M. Berthelot, J. Andreu-Perez, B. Lo, and G.-Z. Yang, "Deep learning for health informatics," *IEEE Journal of Biomedical and Health Informatics*, vol. 21, no. 1, pp. 4–21, 2017.
- [7] T. Ching, D. S. Himmelstein, B. K. Beaulieu-Jones, A. A. Kalinin, B. T. Do, G. P. Way, E. Ferrero, P.-M. Agapow, M. Zietz, M. M. Hoffman *et al.*, "Opportunities and obstacles for deep learning in biology and medicine," *bioRxiv*, p. 142760, 2018.
- [8] J. Luo, M. Wu, D. Gopukumar, and Y. Zhao, "Big data application in biomedical research and health care: a literature review," *Biomedical Informatics Insights*, vol. 8, pp. BII–S31559, 2016.
- [9] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. Corrado, S. Thrun, and J. Dean, "A guide to deep learning in healthcare," *Nature Medicine*, vol. 25, no. 1, p. 24, 2019.
- [10] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [11] M. L. Littman, "Reinforcement learning improves behaviour from evaluative feedback," *Nature*, vol. 521, no. 7553, p. 445, 2015.
- [12] Y. Li, "Deep reinforcement learning," *arXiv preprint arXiv:1810.06339*, 2018.
- [13] M. Mahmud, M. S. Kaiser, A. Hussain, and S. Vassanelli, "Applications of deep learning and reinforcement learning to biological data," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 6, pp. 2063–2079, 2018.
- [14] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [15] L. Busoniu, T. de Bruin, D. Tolic, J. Kober, and I. Palunko, "Reinforcement learning for control: Performance, stability, and deep approximators," *Annual Reviews in Control*, 2018.
- [16] O. Gottesman, F. Johansson, M. Komorowski, A. Faisal, D. Sontag, F. Doshi-Velez, and L. A. Celi, "Guidelines for reinforcement learning in healthcare," *Nature medicine*, vol. 25, no. 1, p. 16, 2019.
- [17] R. Bellman, *Dynamic programming*. Courier Corporation, 2013.
- [18] C. J. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [19] G. A. Rummery and M. Niranjan, *On-line Q-learning using connectionist systems*. University of Cambridge, Department of Engineering Cambridge, England, 1994, vol. 37.
- [20] J. Kober and J. R. Peters, "Policy search for motor primitives in robotics," in *Advances in Neural Information Processing Systems*, 2009, pp. 849–856.
- [21] J. Peters and S. Schaal, "Natural actor-critic," *Neurocomputing*, vol. 71, no. 7-9, pp. 1180–1190, 2008.
- [22] N. Vlassis, M. Ghavamzadeh, S. Mannor, and P. Poupart, "Bayesian reinforcement learning," in *Reinforcement Learning*. Springer, 2012, pp. 359–386.
- [23] M. Ghavamzadeh, S. Mannor, J. Pineau, A. Tamar *et al.*, "Bayesian reinforcement learning: A survey," *Foundations and Trends in Machine Learning*, vol. 8, no. 5-6, pp. 359–483, 2015.
- [24] A. L. Strehl, L. Li, and M. L. Littman, "Reinforcement learning in finite mdps: Pac analysis," *Journal of Machine Learning Research*, vol. 10, no. Nov, pp. 2413–2444, 2009.
- [25] R. I. Brafman and M. Tennenholtz, "R-max-a general polynomial time algorithm for near-optimal reinforcement learning," *Journal of Machine Learning Research*, vol. 3, no. Oct, pp. 213–231, 2002.
- [26] A. G. Barto, "Intrinsic motivation and reinforcement learning," in *Intrinsically motivated learning in natural and artificial systems*. Springer, 2013, pp. 17–47.
- [27] L. Busoniu, R. Babuska, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," *IEEE Transactions on Systems, Man, And Cybernetics-Part C: Applications and Reviews*, 38 (2), 2008, 2008.
- [28] S. M. Kakade *et al.*, "On the sample complexity of reinforcement learning," Ph.D. dissertation, University of London London, England, 2003.

- [29] L. Li, "Sample complexity bounds of exploration," in *Reinforcement Learning*. Springer, 2012, pp. 175–204.
- [30] L. Busoniu, R. Babuska, B. De Schutter, and D. Ernst, *Reinforcement learning and dynamic programming using function approximators*. CRC press, 2010.
- [31] H. Van Hasselt, "Reinforcement learning in continuous state and action spaces," in *Reinforcement learning*. Springer, 2012, pp. 207–251.
- [32] T. M. Moldovan and P. Abbeel, "Safe exploration in markov decision processes," in *Proceedings of the 29th International Conference on Machine Learning*. Omnipress, 2012, pp. 1451–1458.
- [33] J. Garcia and F. Fernandez, "A comprehensive survey on safe reinforcement learning," *Journal of Machine Learning Research*, vol. 16, no. 1, pp. 1437–1480, 2015.
- [34] W. Wiesemann, D. Kuhn, and B. Rustem, "Robust markov decision processes," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 153–183, 2013.
- [35] H. Xu and S. Mannor, "Distributionally robust markov decision processes," in *Advances in Neural Information Processing Systems*, 2010, pp. 2505–2513.
- [36] D. Hein, S. Udfluft, and T. A. Runkler, "Interpretable policies for reinforcement learning by genetic programming," *Engineering Applications of Artificial Intelligence*, vol. 76, pp. 158–169, 2018.
- [37] O. Bastani, Y. Pu, and A. Solar-Lezama, "Verifiable reinforcement learning via policy extraction," in *Advances in Neural Information Processing Systems*, 2018, pp. 2499–2509.
- [38] M. Wiering and M. Van Otterlo, "Reinforcement learning," *Adaptation, learning, and optimization*, vol. 12, 2012.
- [39] S. Lange, T. Gabel, and M. Riedmiller, "Batch reinforcement learning," in *Reinforcement learning*. Springer, 2012, pp. 45–73.
- [40] M. Riedmiller, "Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method," in *European Conference on Machine Learning*. Springer, 2005, pp. 317–328.
- [41] D. Ernst, P. Geurts, and L. Wehenkel, "Tree-based batch mode reinforcement learning," *Journal of Machine Learning Research*, vol. 6, no. Apr, pp. 503–556, 2005.
- [42] M. G. Lagoudakis and R. Parr, "Least-squares policy iteration," *Journal of Machine Learning Research*, vol. 4, no. Dec, pp. 1107–1149, 2003.
- [43] T. Hester and P. Stone, "Learning and using models," in *Reinforcement learning*. Springer, 2012, pp. 111–141.
- [44] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A survey of monte carlo tree search methods," *IEEE Transactions on Computational Intelligence and AI in games*, vol. 4, no. 1, pp. 1–43, 2012.
- [45] A. Lazaric, "Transfer in reinforcement learning: a framework and a survey," in *Reinforcement Learning*. Springer, 2012, pp. 143–173.
- [46] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research*, vol. 10, no. Jul, pp. 1633–1685, 2009.
- [47] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot et al., "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, p. 484, 2016.
- [48] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi, "A survey of deep neural network architectures and their applications," *Neurocomputing*, vol. 234, pp. 11–26, 2017.
- [49] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [50] C. Liu, X. Xu, and D. Hu, "Multiobjective reinforcement learning: A comprehensive overview," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 3, pp. 385–398, 2015.
- [51] C. Wirth, R. Akrou, G. Neumann, and J. Furnkranz, "A survey" of preference-based reinforcement learning methods," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 4945–4990, 2017.
- [52] J. Furnkranz, E. Hüllermeier, W. Cheng, and S.-H. Park, "Preference-based reinforcement learning: a formal framework and a policy iteration algorithm," *Machine Learning*, vol. 89, no. 1-2, pp. 123–156, 2012.
- [53] A. Y. Ng, S. J. Russell et al., "Algorithms for inverse reinforcement learning," in *ICML*, 2000, pp. 663–670.
- [54] S. Zhifei and E. Meng Joo, "A survey of inverse reinforcement learning techniques," *International Journal of Intelligent Computing and Cybernetics*, vol. 5, no. 3, pp. 293–311, 2012.
- [55] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *AAAI*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [56] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proceedings of the twenty-first international conference on Machine learning*. ACM, 2004, p. 1.
- [57] S. Levine, Z. Popovic, and V. Koltun, "Nonlinear inverse reinforcement learning with gaussian processes," in *Advances in Neural Information Processing Systems*, 2011, pp. 19–27.
- [58] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," *Urbana*, vol. 51, no. 61801, pp. 1–4, 2007.
- [59] C. Guestrin, D. Koller, R. Parr, and S. Venkataraman, "Efficient solution algorithms for factored mdps," *Journal of Artificial Intelligence Research*, vol. 19, pp. 399–468, 2003.
- [60] M. Kearns and D. Koller, "Efficient reinforcement learning in factored mdps," in *IJCAI*, vol. 16, 1999, pp. 740–747.
- [61] C. Guestrin, R. Patrascu, and D. Schuurmans, "Algorithm-directed exploration for model-based reinforcement learning in factored mdps," in *ICML*, 2002, pp. 235–242.
- [62] I. Osband and B. Van Roy, "Near-optimal reinforcement learning in factored mdps," in *Advances in Neural Information Processing Systems*, 2014, pp. 604–612.
- [63] A. L. Strehl, C. Diuk, and M. L. Littman, "Efficient structure learning in factored-state mdps," in *AAAI*, vol. 7, 2007, pp. 645–650.
- [64] A. G. Barto and S. Mahadevan, "Recent advances in hierarchical reinforcement learning," *Discrete Event Dynamic Systems*, vol. 13, no. 1-2, pp. 41–77, 2003.
- [65] B. Hengst, "Hierarchical approaches," in *Reinforcement learning*. Springer, 2012, pp. 293–323.
- [66] M. van Otterlo, "Solving relational and first-order logical markov decision processes: A survey," in *Reinforcement Learning*. Springer, 2012, pp. 253–292.
- [67] T. Jaakkola, S. P. Singh, and M. I. Jordan, "Reinforcement learning algorithm for partially observable markov decision problems," in *Advances in Neural Information Processing Systems*, 1995, pp. 345–352.
- [68] R. W. Jelliffe, J. Buell, R. Kalaba, R. Sridhar, and R. Rockwell, "A computer program for digitalis dosage regimens," *Mathematical Biosciences*, vol. 9, pp. 179–193, 1970.
- [69] R. E. Bellman, *Mathematical methods in medicine*. World Scientific Publishing Co., Inc., 1983.
- [70] C. Hu, W. S. Lovejoy, and S. L. Shafer, "Comparison of some control strategies for three-compartment pk/pd models," *Journal of Pharmacokinetics and Biopharmaceutics*, vol. 22, no. 6, pp. 525–550, 1994.
- [71] A. J. Schaefer, M. D. Bailey, S. M. Shechter, and M. S. Roberts, "Modeling medical treatment using markov decision processes," in *Operations Research and Health Care*. Springer, 2005, pp. 593–612.
- [72] B. Chakraborty and S. A. Murphy, "Dynamic treatment regimes," *Annual Review of Statistics and Its Application*, vol. 1, pp. 447–464, 2014.
- [73] E. B. Laber, D. J. Lizotte, M. Qian, W. E. Pelham, and S. A. Murphy, "Dynamic treatment regimes: Technical challenges and applications," *Electronic Journal of Statistics*, vol. 8, no. 1, p. 1225, 2014.
- [74] J. K. Lunceford, M. Davidian, and A. A. Tsiatis, "Estimation of survival distributions of treatment policies in two-stage randomization designs in clinical trials," *Biometrics*, vol. 58, no. 1, pp. 48–57, 2002.
- [75] D. Almirall and A. Chronis-Tuscano, "Adaptive interventions in child and adolescent mental health," *Journal of Clinical Child & Adolescent Psychology*, vol. 45, no. 4, pp. 383–395, 2016.
- [76] P. W. Lavori and R. Dawson, "Adaptive treatment strategies in chronic disease," *Annu. Rev. Med.*, vol. 59, pp. 443–453, 2008.
- [77] B. Chakraborty and E. E. M. Moodie, *Statistical Reinforcement Learning*. Springer New York, 2013.

- [78] S. A. Murphy, "An experimental design for the development of adaptive treatment strategies," *Statistics in Medicine*, vol. 24, no. 10, pp. 1455–1481, 2005.
- [79] S. A. Murphy, K. G. Lynch, D. Oslin, J. R. McKay, and T. TenHave, "Developing adaptive treatment strategies in substance abuse research," *Drug & Alcohol Dependence*, vol. 88, pp. S24–S30, 2007.
- [80] W. H. Organization, *Preventing chronic diseases: a vital investment*. World Health Organization, 2005.
- [81] B. Chakraborty and E. Moodie, *Statistical methods for dynamic treatment regimes*. Springer, 2013.
- [82] E. H. Wagner, B. T. Austin, C. Davis, M. Hindmarsh, J. Schaefer, and A. Bonomi, "Improving chronic illness care: translating evidence into action," *Health Affairs*, vol. 20, no. 6, pp. 64–78, 2001.
- [83] Y. Zhao, M. R. Kosorok, and D. Zeng, "Reinforcement learning design for cancer clinical trials," *Statistics in Medicine*, vol. 28, no. 26, pp. 3294–3315, 2009.
- [84] A. Hassani *et al.*, "Reinforcement learning based control of tumor growth with chemotherapy," in *2010 International Conference on System Science and Engineering (ICSSSE)*. IEEE, 2010, pp. 185–189.
- [85] I. Ahn and J. Park, "Drug scheduling of cancer chemotherapy based on natural actor-critic approach," *BioSystems*, vol. 106, no. 2-3, pp. 121–129, 2011.
- [86] K. Humphrey, "Using reinforcement learning to personalize dosing strategies in a simulated cancer trial with high dimensional data," 2017.
- [87] R. Padmanabhan, N. Meskin, and W. M. Haddad, "Reinforcement learning-based control of drug dosing for cancer chemotherapy treatment," *Mathematical biosciences*, vol. 293, pp. 11–20, 2017.
- [88] Y. Zhao, D. Zeng, M. A. Socinski, and M. R. Kosorok, "Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer," *Biometrics*, vol. 67, no. 4, pp. 1422–1433, 2011.
- [89] W. Cheng, J. Furnkranz, E. H. Ullermeier, and S.-H. Park, "Preference-based policy iteration: Leveraging preference learning for reinforcement learning," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2011, pp. 312–327.
- [90] R. Akrou, M. Schoenauer, and M. Sebag, "April: Active preference learning-based reinforcement learning," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2012, pp. 116–131.
- [91] R. Busa-Fekete, B. Szorényi, P. Weng, W. Cheng, and E. H. Ullermeier, "Preference-based reinforcement learning: evolutionary direct policy search using a preference-based racing algorithm," *Machine Learning*, vol. 97, no. 3, pp. 327–351, 2014.
- [92] R. Vincent, "Reinforcement learning in models of adaptive medical treatment strategies," Ph.D. dissertation, McGill University Libraries, 2014.
- [93] H. H. Tseng, Y. Luo, S. Cui, J. T. Chien, R. K. Ten Haken, and I. E. Naqa, "Deep reinforcement learning for automated radiation adaptation in lung cancer," *Medical Physics*, vol. 44, no. 12, pp. 6690–6705, 2017.
- [94] A. Jalalimanesh, H. S. Haghighi, A. Ahmadi, and M. Soltani, "Simulation-based optimization of radiotherapy: Agent-based modeling and reinforcement learning," *Mathematics and Computers in Simulation*, vol. 133, pp. 235–248, 2017.
- [95] A. Jalalimanesh, H. S. Haghighi, A. Ahmadi, H. Hejazian, and M. Soltani, "Multi-objective optimization of radiotherapy: distributed q-learning and agent-based simulation," *Journal of Experimental & Theoretical Artificial Intelligence*, pp. 1–16, 2017.
- [96] Y. Goldberg and M. R. Kosorok, "Q-learning with censored data," *Annals of Statistics*, vol. 40, no. 1, p. 529, 2012.
- [97] Y. M. Soliman, "Personalized medical treatments using novel reinforcement learning algorithms," *arXiv preprint arXiv:1406.3922*, 2014.
- [98] G. Yaune and P. Shah, "Reinforcement learning with action-derived rewards for chemotherapy and clinical trial dosing regimen selection," in *Machine Learning for Healthcare Conference*, 2018, pp. 161–226.
- [99] B. Stewart, C. P. Wild *et al.*, "World cancer report 2014," *Health*, 2017.
- [100] R. Eftimie, J. L. Bramson, and D. J. Earn, "Interactions between the immune system and cancer: a brief review of non-spatial mathematical models," *Bulletin of Mathematical Biology*, vol. 73, no. 1, pp. 2–32, 2011.
- [101] J. Shi, O. Alagoz, F. S. Erenay, and Q. Su, "A survey of optimization models on cancer chemotherapy treatment planning," *Annals of Operations Research*, vol. 221, no. 1, pp. 331–356, 2014.
- [102] N. Beerenwinkel, R. F. Schwarz, M. Gerstung, and F. Markowetz, "Cancer evolution: mathematical models and computational inference," *Systematic Biology*, vol. 64, no. 1, pp. e1–e25, 2014.
- [103] M. Tenenbaum, A. Fern, L. Getoor, M. Littman, V. Manasinghka, S. Natarajan, D. Page, J. Shrager, Y. Singer, and P. Tadepalli, "Personalizing cancer therapy via machine learning," in *Workshops of NIPS*, 2010.
- [104] V. Vapnik, S. E. Golowich, and A. J. Smola, "Support vector method for function approximation, regression estimation and signal processing," in *Advances in Neural Information Processing Systems*, 1997, pp. 281–287.
- [105] L. G. De Pillis and A. Radunskaya, "The dynamics of an optimally controlled tumor model: A case study," *Mathematical and Computer Modelling*, vol. 37, no. 11, pp. 1221–1244, 2003.
- [106] M. Feng, G. Valdes, N. Dixit, and T. D. Solberg, "Machine learning in radiation oncology: Opportunities, requirements, and needs," *Frontiers in Oncology*, vol. 8, 2018.
- [107] J. de Lope, D. Maravall *et al.*, "Robust high performance reinforcement learning through weighted k-nearest neighbors," *Neurocomputing*, vol. 74, no. 8, pp. 1251–1259, 2011.
- [108] N. Cho, J. Shaw, S. Karuranga, Y. Huang, J. da Rocha Fernandes, A. Ohlrogge, and B. Malanda, "Idf diabetes atlas: Global estimates of diabetes prevalence for 2017 and projections for 2045," *Diabetes Research and Clinical Practice*, vol. 138, pp. 271–281, 2018.
- [109] A. M. Albisser, B. Leibel, T. Ewart, Z. Davidovac, C. Botz, W. Zingg, H. Schipper, and R. Gander, "Clinical control of diabetes by the artificial pancreas," *Diabetes*, vol. 23, no. 5, pp. 397–404, 1974.
- [110] C. Cobelli, E. Renard, and B. Kovatchev, "Artificial pancreas: past, present, future," *Diabetes*, vol. 60, no. 11, pp. 2672–2682, 2011.
- [111] B. W. Bequette, "A critical assessment of algorithms and challenges in the development of a closed-loop artificial pancreas," *Diabetes Technology & Therapeutics*, vol. 7, no. 1, pp. 28–47, 2005.
- [112] T. Peyser, E. Dassau, M. Breton, and J. S. Skyler, "The artificial pancreas: current status and future prospects in the management of diabetes," *Annals of the New York Academy of Sciences*, vol. 1311, no. 1, pp. 102–123, 2014.
- [113] M. K. Bothe, L. Dickens, K. Reichel, A. Tellmann, B. Ellger, M. Westphal, and A. A. Faisal, "The use of reinforcement learning algorithms to meet the challenges of an artificial pancreas," *Expert Review of Medical Devices*, vol. 10, no. 5, pp. 661–673, 2013.
- [114] S. Yasini, M. B. Naghibi Sistani, and A. Karimpour, "Agent-based simulation for blood glucose," *International Journal of Applied Science, Engineering and Technology*, vol. 5, pp. 89–95, 2009.
- [115] E. Daskalaki, L. Scarnato, P. Diem, and S. G. Mougiakakou, "Preliminary results of a novel approach for glucose regulation using an actor-critic learning based controller," 2010.
- [116] B. P. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli, "In silico preclinical trials: a proof of concept in closed-loop control of type 1 diabetes," 2009.
- [117] E. Daskalaki, P. Diem, and S. G. Mougiakakou, "An actor-critic based controller for glucose regulation in type 1 diabetes," *Computer Methods and Programs in Biomedicine*, vol. 109, no. 2, pp. 116–125, 2013.
- [118] —, "Personalized tuning of a reinforcement learning control algorithm for glucose regulation," in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2013, pp. 3487–3490.
- [119] —, "Model-free machine learning in biomedicine: Feasibility study in type 1 diabetes," *PLoS One*, vol. 11, no. 7, p. e0158722, 2016.
- [120] Q. Sun, M. Jankovic, J. Budzinski, B. Moore, P. Diem, C. Stettler, and S. G. Mougiakakou, "A dual mode adaptive basal-bolus advisor based on reinforcement learning," *IEEE journal of biomedical and health informatics*, 2018.
- [121] P. Palumbo, S. Panunzi, and A. De Gaetano, "Qualitative behavior of a family of delay-differential models of the glucose-insulin system,"

- Discrete and Continuous Dynamical Systems Series B*, vol. 7, no. 2, p. 399, 2007.
- [122] A. Noori, M. A. Sadrnia *et al.*, "Glucose level control using temporal difference methods," in *2017 Iranian Conference on Electrical Engineering (ICEE)*. IEEE, 2017, pp. 895–900.
- [123] P. D. Ngo, S. Wei, A. Holubova, J. Muzik, and F. Godtliebsen, "Reinforcement-learning optimal control for type-1 diabetes," in *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*. IEEE, 2018, pp. 333–336.
- [124] —, "Control of blood glucose for type-1 diabetes by using reinforcement learning with feedforward algorithm," *Computational and Mathematical Methods in Medicine*, vol. 2018, 2018.
- [125] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity." *American Journal of PhysiologyEndocrinology And Metabolism*, vol. 236, no. 6, p. E667, 1979.
- [126] R. Hovorka, V. Canonico, L. J. Chassin, U. Haueter, M. MassiBenedetti, M. O. Federici, T. R. Pieber, H. C. Schaller, L. Schaupp, T. Vering *et al.*, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological Measurement*, vol. 25, no. 4, p. 905, 2004.
- [127] M. De Paula, L. O. Avila, and E. C. Mart´ınez, "Controlling blood glucose variability under uncertainty using reinforcement learning and gaussian processes," *Applied Soft Computing*, vol. 35, pp. 310–332, 2015.
- [128] M. De Paula, G. G. Acosta, and E. C. Mart´ınez, "On-line policy learning and adaptation for real-time personalization of an artificial pancreas," *Expert Systems with Applications*, vol. 42, no. 4, pp. 2234–2255, 2015.
- [129] S. U. Acikgoz and U. M. Diwekar, "Blood glucose regulation with stochastic optimal control for insulin-dependent diabetic patients," *Chemical Engineering Science*, vol. 65, no. 3, pp. 1227–1236, 2010.
- [130] H. Asoh, M. Shiro, S. Akaho, T. Kamishima, K. Hashida, E. Aramaki, and T. Kohro, "Modeling medical records of diabetes using markov decision processes," in *Proceedings of ICML2013 Workshop on Role of Machine Learning in Transforming Healthcare*, 2013.
- [131] H. Asoh, M. S. S. Akaho, T. Kamishima, K. Hasida, E. Aramaki, and T. Kohro, "An application of inverse reinforcement learning to medical records of diabetes treatment," in *ECMLPKDD2013 Workshop on Reinforcement Learning with Generalized Feedback*, 2013.
- [132] D. J. Lockett, E. B. Laber, A. R. Kahkoska, D. M. Maahs, E. MayerDavis, and M. R. Kosorok, "Estimating dynamic treatment regimes in mobile health using v-learning," *Journal of the American Statistical Association*, no. just-accepted, pp. 1–39, 2018.
- [133] A. E. Gaweda, M. K. Muezzinoglu, G. R. Aronoff, A. A. Jacobs, J. M. Zurada, and M. E. Brier, "Reinforcement learning approach to individualization of chronic pharmacotherapy," in *IJCNN'05*, vol. 5. IEEE, 2005, pp. 3290–3295.
- [134] A. E. Gaweda, M. K. Muezzinoglu, A. A. Jacobs, G. R. Aronoff, and M. E. Brier, "Model predictive control with reinforcement learning for drug delivery in renal anemia management," in *IEEE EMBS'06*. IEEE, 2006, pp. 5177–5180.
- [135] A. E. Gaweda, M. K. Muezzinoglu, G. R. Aronoff, A. A. Jacobs, J. M. Zurada, and M. E. Brier, "Individualization of pharmacological anemia management using reinforcement learning," *Neural Networks*, vol. 18, no. 5-6, pp. 826–834, 2005.
- [136] J. D. Mart´ın-Guerrero, F. Gomez, E. Soria-Olivas, J. Schmidhuber, M. Clemente-Mart´ı, and N. V. Jimenez-Torres, "A reinforcement learning approach for individualizing erythropoietin dosages in hemodialysis patients," *Expert Systems with Applications*, vol. 36, no. 6, pp. 9737–9742, 2009.
- [137] J. D. Mart´ın-Guerrero, E. Soria-Olivas, M. Mart´ınez-Sober, M. Clemente-Mart´ı, T. De Diego-Santos, and N. V. Jimenez-Torres, "Validation of a reinforcement learning policy for dosage optimization of erythropoietin," in *Australasian Joint Conference on Artificial Intelligence*. Springer, 2007, pp. 732–738.
- [138] J. M. Malof and A. E. Gaweda, "Optimizing drug therapy with reinforcement learning: The case of anemia management," in *Neural Networks (IJCNN), The 2011 International Joint Conference on*. IEEE, 2011, pp. 2088–2092.
- [139] P. Escandell-Montero, J. M. Mart´ınez-Mart´ınez, J. D. Mart´ın-Guerrero, E. Soria-Olivas, J. Vila-Frances, and R. Magdalena-Benedito, "Adaptive treatment of anemia on hemodialysis patients: A reinforcement learning approach," in *CIDM2011*. IEEE, 2011, pp. 44–49.
- [140] P. Escandell-Montero, M. Chermisi, J. M. Martinez-Martinez, J. Gomez-Sanchis, C. Barbieri, E. Soria-Olivas, F. Mari, J. VilaFrances, A. Stopper, E. Gatti *et al.*, "Optimization of anemia treatment in hemodialysis patients via reinforcement learning," *Artificial Intelligence in Medicine*, vol. 62, no. 1, pp. 47–60, 2014.
- [141] B. M. Adams, H. T. Banks, H.-D. Kwon, and H. T. Tran, "Dynamic multidrug therapies for hiv: Optimal and sti control approaches," *Mathematical Biosciences and Engineering*, vol. 1, no. 2, pp. 223–241, 2004.
- [142] D. Ernst, G.-B. Stan, J. Goncalves, and L. Wehenkel, "Clinical data based optimal sti strategies for hiv: a reinforcement learning approach," in *45th IEEE Conference on Decision and Control*. IEEE, 2006, pp. 667–672.
- [143] S. Parbhoo, "A reinforcement learning design for hiv clinical trials," Ph.D. dissertation, 2014.
- [144] S. Parbhoo, J. Bogojeska, M. Zazzi, V. Roth, and F. Doshi-Velez, "Combining kernel and model based learning for hiv therapy selection," *AMIA Summits on Translational Science Proceedings*, vol. 2017, p. 239, 2017.
- [145] V. N. Marivate, J. Chemali, E. Brunskill, and M. L. Littman, "Quantifying uncertainty in batch personalized sequential decision making." in *AAAI Workshop: Modern Artificial Intelligence for Health Analytics*, 2014.
- [146] T. Killian, G. Konidaris, and F. Doshi-Velez, "Transfer learning across patient variations with hidden parameter markov decision processes," *arXiv preprint arXiv:1612.00475*, 2016.
- [147] T. W. Killian, S. Daulton, G. Konidaris, and F. Doshi-Velez, "Robust and efficient transfer learning with hidden parameter markov decision processes," in *Advances in Neural Information Processing Systems*, 2017, pp. 6250–6261.
- [148] J. Yao, T. Killian, G. Konidaris, and F. Doshi-Velez, "Direct policy transfer via hidden parameter markov decision processes," 2018.
- [149] C. Yu, Y. Dong, J. Liu, and G. Ren, "Incorporating causal factors into reinforcement learning for dynamic treatment regimes in hiv," *BMC medical informatics and decision making*, vol. 19, no. 2, p. 60, 2019.
- [150] J. Papis and R. Parr, "Pac optimal exploration in continuous space markov decision processes." in *AAAI*, 2013.
- [151] K. Kawaguchi, "Bounded optimal exploration in mdp." in *AAAI*, 2016, pp. 1758–1764.
- [152] S. A. Murphy, D. W. Oslin, A. J. Rush, and J. Zhu, "Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders," *Neuropsychopharmacology*, vol. 32, no. 2, p. 257, 2007.
- [153] T. N. Alotaiby, S. A. Alshebeili, T. Alshawi, I. Ahmad, and F. E. A. El-Samie, "Eeg seizure detection and prediction algorithms: a survey," *EURASIP Journal on Advances in Signal Processing*, vol. 2014, no. 1, p. 183, 2014.
- [154] G. Panuccio, M. Semprini, L. Natale, S. Buccelli, I. Colombi, and M. Chiappalone, "Progress in neuroengineering for brain repair: New challenges and open issues," *Brain and Neuroscience Advances*, vol. 2, p. 2398212818776475, 2018.
- [155] A. Guez, R. D. Vincent, M. Avoli, and J. Pineau, "Adaptive treatment of epilepsy via batch-mode reinforcement learning." in *AAAI*, 2008, pp. 1671–1678.
- [156] J. Pineau, A. Guez, R. Vincent, G. Panuccio, and M. Avoli, "Treating epilepsy via adaptive neurostimulation: a reinforcement learning approach," *International Journal of Neural Systems*, vol. 19, no. 04, pp. 227–240, 2009.
- [157] A. Guez, "Adaptive control of epileptic seizures using reinforcement learning," Ph.D. dissertation, McGill University Library, 2010.
- [158] G. Panuccio, A. Guez, R. Vincent, M. Avoli, and J. Pineau, "Adaptive control of epileptiform excitability in an in vitro model of limbic seizures," *Experimental Neurology*, vol. 241, pp. 179–183, 2013.
- [159] K. Bush and J. Pineau, "Manifold embeddings for model-based reinforcement learning under partial observability," in *Advances in Neural Information Processing Systems*, 2009, pp. 189–197.

- [160] V. Nagaraj, A. Lamperski, and T. I. Netoff, "Seizure control in a computational model using a reinforcement learning stimulation paradigm," *International Journal of Neural Systems*, vol. 27, no. 07, p. 1750012, 2017.
- [161] A. J. Rush, M. Fava, S. R. Wisniewski, P. W. Lavori, M. H. Trivedi, H. A. Sackeim, M. E. Thase, A. A. Nierenberg, F. M. Quitkin, T. M. Kashner *et al.*, "Sequenced treatment alternatives to relieve depression (star* d): rationale and design," *Controlled clinical trials*, vol. 25, no. 1, pp. 119–142, 2004.
- [162] J. Pineau, M. G. Bellemare, A. J. Rush, A. Ghizaru, and S. A. Murphy, "Constructing evidence-based treatment strategies using methods from computer science," *Drug & Alcohol Dependence*, vol. 88, pp. S52–S60, 2007.
- [163] D. Ormoneit and S. Sen, "Kernel-based reinforcement learning," *Machine learning*, vol. 49, no. 2-3, pp. 161–178, 2002.
- [164] B. Chakraborty, E. B. Laber, and Y. Zhao, "Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme," *Biometrics*, vol. 69, no. 3, pp. 714–723, 2013.
- [165] E. B. Laber, K. A. Linn, and L. A. Stefanski, "Interactive model building for q-learning," *Biometrika*, vol. 101, no. 4, pp. 831–847, 2014.
- [166] K. A. Linn, E. B. Laber, and L. A. Stefanski, "Interactive q-learning for probabilities and quantiles," *arXiv preprint arXiv:1407.3414*, 2014.
- [167] —, "Interactive q-learning for quantiles," *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 638–649, 2017.
- [168] P. J. Schulte, A. A. Tsiatis, E. B. Laber, and M. Davidian, "Q-and alearning methods for estimating optimal dynamic treatment regimes," *Statistical science: a review journal of the Institute of Mathematical Statistics*, vol. 29, no. 4, p. 640, 2014.
- [169] S. A. Murphy, "Optimal dynamic treatment regimes," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 65, no. 2, pp. 331–355, 2003.
- [170] R. Song, W. Wang, D. Zeng, and M. R. Kosorok, "Penalized q-learning for dynamic treatment regimes," *Statistica Sinica*, vol. 25, no. 3, p. 901, 2015.
- [171] Y. Liu, Y. Wang, M. R. Kosorok, Y. Zhao, and D. Zeng, "Robust hybrid learning for estimating personalized dynamic treatment regimes," *arXiv preprint arXiv:1611.02314*, 2016.
- [172] K. Deng, R. Greiner, and S. Murphy, "Budgeted learning for developing personalized treatment," in *ICMLA2014*. IEEE, 2014, pp. 7–14.
- [173] R. S. Keefe, R. M. Bilder, S. M. Davis, P. D. Harvey, B. W. Palmer, J. M. Gold, H. Y. Meltzer, M. F. Green, G. Capuano, T. S. Stroup *et al.*, "Neurocognitive effects of antipsychotic medications in patients with chronic schizophrenia in the catie trial," *Archives of General Psychiatry*, vol. 64, no. 6, pp. 633–647, 2007.
- [174] S. M. Shortreed, E. Laber, D. J. Lizotte, T. S. Stroup, J. Pineau, and S. A. Murphy, "Informing sequential clinical decision-making through reinforcement learning: an empirical study," *Machine Learning*, vol. 84, no. 1-2, pp. 109–136, 2011.
- [175] A. Ertefaie, S. Shortreed, and B. Chakraborty, "Q-learning residual analysis: application to the effectiveness of sequences of antipsychotic medications for patients with schizophrenia," *Statistics in Medicine*, vol. 35, no. 13, pp. 2221–2234, 2016.
- [176] D. J. Lizotte, M. Bowling, and S. A. Murphy, "Linear fitted-q iteration with multiple reward functions," *Journal of Machine Learning Research*, vol. 13, no. Nov, pp. 3253–3295, 2012.
- [177] D. J. Lizotte and E. B. Laber, "Multi-objective markov decision processes for data-driven decision support," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 7378–7405, 2016.
- [178] E. B. Laber, D. J. Lizotte, and B. Ferguson, "Set-valued dynamic treatment regimes for competing outcomes," *Biometrics*, vol. 70, no. 1, pp. 53–61, 2014.
- [179] E. L. Butler, E. B. Laber, S. M. Davis, and M. R. Kosorok, "Incorporating patient preferences into estimation of optimal individualized treatment rules," *Biometrics*, 2017.
- [180] M. Dennis and C. K. Scott, "Managing addiction as a chronic condition," *Addiction Science & Clinical Practice*, vol. 4, no. 1, p. 45, 2007.
- [181] S. A. Murphy, Y. Deng, E. B. Laber, H. R. Maei, R. S. Sutton, and K. Witkiewitz, "A batch, off-policy, actor-critic algorithm for optimizing the average reward," *arXiv preprint arXiv:1607.05047*, 2016.
- [182] B. Chakraborty, S. Murphy, and V. Strecher, "Inference for non-regular parameters in optimal dynamic treatment regimes," *Statistical Methods in Medical Research*, vol. 19, no. 3, pp. 317–343, 2010.
- [183] B. Chakraborty, V. Strecher, and S. Murphy, "Bias correction and confidence intervals for fitted q-iteration," in *Workshop on Model Uncertainty and Risk in Reinforcement Learning, NIPS, Whistler, Canada*. Citeseer, 2008.
- [184] Y. Tao, L. Wang, D. Almirall *et al.*, "Tree-based reinforcement learning for estimating optimal dynamic treatment regimes," *The Annals of Applied Statistics*, vol. 12, no. 3, pp. 1914–1938, 2018.
- [185] J.-L. Vincent, "Critical care—where have we been and where are we going?" *Critical Care*, vol. 17, no. 1, p. S2, 2013.
- [186] K. Krell, "Critical care workforce," *Critical Care Medicine*, vol. 36, no. 4, pp. 1350–1353, 2008.
- [187] M. Ghassemi, L. A. Celi, and D. J. Stone, "State of the art review: the data revolution in critical care," *Critical Care*, vol. 19, no. 1, p. 118, 2015.
- [188] A. Rhodes, L. E. Evans, W. Alhazzani, M. M. Levy, M. Antonelli, R. Ferrer, A. Kumar, J. E. Sevransky, C. L. Sprung, M. E. Nunnally *et al.*, "Surviving sepsis campaign: international guidelines for management of sepsis and septic shock: 2016," *Intensive Care Medicine*, vol. 43, no. 3, pp. 304–377, 2017.
- [189] A. D. T. Force, V. Ranieri, G. Rubenfeld *et al.*, "Acute respiratory distress syndrome," *Jama*, vol. 307, no. 23, pp. 2526–2533, 2012.
- [190] T. Kamio, T. Van, and K. Masamune, "Use of machine-learning approaches to predict clinical deterioration in critically ill patients: A systematic review," *International Journal of Medical Research and Health Sciences*, vol. 6, no. 6, pp. 1–7, 2017.
- [191] A. Vellido, V. Ribas, C. Morales, A. R. Sanmartín, and J. C. R. Rodríguez, "Machine learning in critical care: state-of-the-art and a sepsis case study," *Biomedical engineering online*, vol. 17, no. 1, p. 135, 2018.
- [192] M. Komorowski, A. Gordon, L. Celi, and A. Faisal, "A markov decision process to suggest optimal treatment of severe infections in intensive care," in *Neural Information Processing Systems Workshop on Machine Learning for Health*, 2016.
- [193] M. Komorowski, L. A. Celi, O. Badawi, A. C. Gordon, and A. A. Faisal, "The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care," *Nature Medicine*, vol. 24, no. 11, p. 1716, 2018.
- [194] A. Raghu, M. Komorowski, I. Ahmed, L. Celi, P. Szolovits, and M. Ghassemi, "Deep reinforcement learning for sepsis treatment," *arXiv preprint arXiv:1711.09602*, 2017.
- [195] A. Raghu, M. Komorowski, L. A. Celi, P. Szolovits, and M. Ghassemi, "Continuous state-space models for optimal sepsis treatment: a deep reinforcement learning approach," in *Machine Learning for Healthcare Conference*, 2017, pp. 147–163.
- [196] A. Raghu, M. Komorowski, and S. Singh, "Model-based reinforcement learning for sepsis treatment," *arXiv preprint arXiv:1811.09602*, 2018.
- [197] C. P. Utomo, X. Li, and W. Chen, "Treatment recommendation in critical care: A scalable and interpretable approach in partially observable health states," 2018.
- [198] X. Peng, Y. Ding, D. Wihl, O. Gottesman, M. Komorowski, L.-w. H. Lehman, A. Ross, A. Faisal, and F. Doshi-Velez, "Improving sepsis treatment strategies by combining deep and kernel-based reinforcement learning," *arXiv preprint arXiv:1901.04670*, 2019.
- [199] J. Futoma, A. Lin, M. Sendak, A. Bedoya, M. Clement, C. O'Brien, and K. Heller, "Learning to treat sepsis with multi-output gaussian process deep recurrent q-networks," 2018.
- [200] C. Yu, G. Ren, and J. Liu, "Deep inverse reinforcement learning for sepsis treatment," in *2019 IEEE ICHI*, 2019, pp. 1–3.
- [201] L. Li, M. Komorowski, and A. A. Faisal, "The actor search tree critic (astc) for off-policy pomdp learning in medical decision making," *arXiv preprint arXiv:1805.11548*, 2018.
- [202] W.-H. Weng, M. Gao, Z. He, S. Yan, and P. Szolovits, "Representation and reinforcement learning for personalized glycemic control in septic patients," *arXiv preprint arXiv:1712.00654*, 2017.
- [203] B. K. Petersen, J. Yang, W. S. Grathwohl, C. Cockrell, C. Santiago, G. An, and D. M. Faissol, "Precision medicine as a control problem: Using simulation and deep reinforcement learning to discover adaptive,

- personalized multi-cytokine therapy for sepsis,” *arXiv preprint arXiv:1802.10440*, 2018.
- [204] B. L. Moore, E. D. Sinzinger, T. M. Quasny, and L. D. Pyeatt, “Intelligent control of closed-loop sedation in simulated icu patients.” in *FLAIRS Conference*, 2004, pp. 109–114.
- [205] E. D. Sinzinger and B. Moore, “Sedation of simulated icu patients using reinforcement learning based control,” *International Journal on Artificial Intelligence Tools*, vol. 14, no. 01n02, pp. 137–156, 2005.
- [206] B. L. Moore, A. G. Doufas, and L. D. Pyeatt, “Reinforcement learning: a novel method for optimal control of propofol-induced hypnosis,” *Anesthesia & Analgesia*, vol. 112, no. 2, pp. 360–367, 2011.
- [207] B. L. Moore, T. M. Quasny, and A. G. Doufas, “Reinforcement learning versus proportional–integral–derivative control of hypnosis in a simulated intraoperative patient,” *Anesthesia & Analgesia*, vol. 112, no. 2, pp. 350–359, 2011.
- [208] B. L. Moore, L. D. Pyeatt, V. Kulkarni, P. Panousis, K. Padrez, and A. G. Doufas, “Reinforcement learning for closed-loop propofol anesthesia: a study in human volunteers,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 655–696, 2014.
- [209] B. L. Moore, P. Panousis, V. Kulkarni, L. D. Pyeatt, and A. G. Doufas, “Reinforcement learning for closed-loop propofol anesthesia: A human volunteer study.” in *IAAI*, 2010.
- [210] N. Sadati, A. Aflaki, and M. Jahed, “Multivariable anesthesia control using reinforcement learning,” in *IEEE SMC’06*, vol. 6. IEEE, 2006, pp. 4563–4568.
- [211] E. C. Borera, B. L. Moore, A. G. Doufas, and L. D. Pyeatt, “An adaptive neural network filter for improved patient state estimation in closedloop anesthesia control,” in *IEEE ICTAI’11*. IEEE, 2011, pp. 41–46.
- [212] C. Lowery and A. A. Faisal, “Towards efficient, personalized anesthesia using continuous reinforcement learning for propofol infusion control,” in *IEEE/EMBS NER’13*. IEEE, 2013, pp. 1414–1417.
- [213] R. Padmanabhan, N. Meskin, and W. M. Haddad, “Closed-loop control of anesthesia and mean arterial pressure using reinforcement learning,” *Biomedical Signal Processing and Control*, vol. 22, pp. 54–64, 2015.
- [214] P. Humbert, J. Audiffren, C. Dubost, and L. Oudre, “Learning from an expert.”
- [215] S. Nemati, M. M. Ghassemi, and G. D. Clifford, “Optimal medication dosing from suboptimal clinical examples: A deep reinforcement learning approach,” in *IEEE 38th Annual International Conference of the Engineering in Medicine and Biology Society*. IEEE, 2016, pp. 2978–2981.
- [216] R. Lin, M. D. Stanley, M. M. Ghassemi, and S. Nemati, “A deep deterministic policy gradient approach to medication dosing and surveillance in the icu,” in *IEEE EMBC’18*. IEEE, 2018, pp. 4927–4931.
- [217] L. Wang, W. Zhang, X. He, and H. Zha, “Supervised reinforcement learning with recurrent neural network for dynamic treatment recommendation,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2018, pp. 2447–2456.
- [218] N. Prasad, L.-F. Cheng, C. Chivers, M. Draugelis, and B. E. Engelhardt, “A reinforcement learning approach to weaning of mechanical ventilation in intensive care units,” *arXiv preprint arXiv:1704.06300*, 2017.
- [219] C. Yu, J. Liu, and H. Zhao, “Inverse reinforcement learning for intelligent mechanical ventilation and sedative dosing in intensive care units,” *BMC medical informatics and decision making*, vol. 19, no. 2, p. 57, 2019.
- [220] C. Yu, G. Ren, and Y. Dong, “Supervised-actor-critic reinforcement learning for intelligent mechanical ventilation and sedative dosing in intensive care units,” *BMC medical informatics and decision making*, 2020.
- [221] A. Jagannatha, P. Thomas, and H. Yu, “Towards high confidence offpolicy reinforcement learning for clinical applications.”
- [222] L.-F. Cheng, N. Prasad, and B. E. Engelhardt, “An optimal policy for patient laboratory tests in intensive care units,” *arXiv preprint arXiv:1808.04679*, 2018.
- [223] C.-H. Chang, M. Mai, and A. Goldenberg, “Dynamic measurement scheduling for adverse event forecasting using deep rl,” *arXiv preprint arXiv:1812.00268*, 2018.
- [224] E. F. Krakow, M. Hemmer, T. Wang, B. Logan, M. Arora, S. Spellman, D. Couriel, A. Alousi, J. Pidala, M. Last *et al.*, “Tools for the precision medicine era: How to develop highly personalized treatment recommendations from cohort and registry data using q-learning,” *American journal of epidemiology*, vol. 186, no. 2, pp. 160–172, 2017.
- [225] Y. Liu, B. Logan, N. Liu, Z. Xu, J. Tang, and Y. Wang, “Deep reinforcement learning for dynamic treatment regimes on medical registry data,” in *IEEE ICHI’17*. IEEE, 2017, pp. 380–385.
- [226] J. E. Gotts and M. A. Matthay, “Sepsis: pathophysiology and clinical management,” *Bmj*, vol. 353, p. i1585, 2016.
- [227] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, “Mimic-iii, a freely accessible critical care database,” *Scientific Data*, vol. 3, p. 160035, 2016.
- [228] S. Saria, “Individualized sepsis treatment using reinforcement learning,” *Nature medicine*, vol. 24, no. 11, p. 1641, 2018.
- [229] H. Van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning.” in *AAAI*, vol. 2. Phoenix, AZ, 2016, p. 5.
- [230] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, “Dueling network architectures for deep reinforcement learning,” in *International Conference on Machine Learning*, 2016, pp. 1995–2003.
- [231] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, “Prioritized experience replay,” *arXiv preprint arXiv:1511.05952*, 2015.
- [232] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [233] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015.
- [234] A. Y. Ng, D. Harada, and S. Russell, “Policy invariance under reward transformations: Theory and application to reward shaping,” in *ICML*, vol. 99, 1999, pp. 278–287.
- [235] W. M. Haddad, J. M. Bailey, B. Gholami, and A. R. Tannenbaum, “Clinical decision support and closed-loop control for intensive care unit sedation,” *Asian Journal of Control*, vol. 20, no. 5, pp. 1343–1350, 2012.
- [236] M. M. Ghassemi, S. E. Richter, I. M. Eche, T. W. Chen, J. Danziger, and L. A. Celi, “A data-driven approach to optimized medication dosing: a focus on heparin,” *Intensive Care Medicine*, vol. 40, no. 9, pp. 1332–1339, 2014.
- [237] S. Jaber, G. Bellani, L. Blanch, A. Demoule, A. Esteban, L. Gattinoni, C. Guerin, N. Hill, J. G. Laffey, S. M. Maggiore *et al.*, “The intensive care medicine research agenda for airways, invasive and noninvasive mechanical ventilation,” *Intensive Care Medicine*, vol. 43, no. 9, pp. 1352–1365, 2017.
- [238] A. De Jong, G. Citerio, and S. Jaber, “Focus on ventilation and airway management in the icu,” *Intensive Care Medicine*, vol. 43, no. 12, pp. 1912–1915, 2017.
- [239] E. National Academies of Sciences, Medicine *et al.*, *Improving diagnosis in health care*. National Academies Press, 2016.
- [240] S. K. Rai and K. Sowmya, “A review on use of machine learning techniques in diagnostic health-care,” *Artificial Intelligent Systems and Machine Learning*, vol. 10, no. 4, pp. 102–107, 2018.
- [241] M. Fatima and M. Pasha, “Survey of machine learning algorithms for disease diagnostic,” *Journal of Intelligent Learning Systems and Applications*, vol. 9, no. 01, p. 1, 2017.
- [242] K. T. Chui, W. Alhalabi, S. S. H. Pang, P. O. d. Pablos, R. W. Liu, and M. Zhao, “Disease diagnosis in smart healthcare: Innovation, technologies and applications,” *Sustainability*, vol. 9, no. 12, p. 2309, 2017.
- [243] Z. C. Lipton, D. C. Kale, C. Elkan, and R. Wetzell, “Learning to diagnose with lstm recurrent neural networks,” *arXiv preprint arXiv:1511.03677*, 2015.
- [244] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, “Retain: An interpretable predictive model for healthcare using reverse time attention mechanism,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3504–3512.
- [245] T. R. Goodwin and S. M. Harabagiu, “Medical question answering for clinical decision support,” in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. ACM, 2016, pp. 297–306.

- [246] Y. Ling, S. A. Hasan, V. Datla, A. Qadir, K. Lee, J. Liu, and O. Farri, "Diagnostic inferencing via improving clinical concept extraction with deep reinforcement learning: A preliminary study," in *Machine Learning for Healthcare Conference*, 2017, pp. 271–285.
- [247] A. Bernstein and E. Burnaev, "Reinforcement learning in computer vision," in *CMV'17*, vol. 10696. International Society for Optics and Photonics, 2018, p. 106961S.
- [248] G. W. Taylor, "A reinforcement learning framework for parameter control in computer vision applications," in *Computer and Robot Vision, 2004. Proceedings. First Canadian Conference on*. IEEE, 2004, pp. 496–503.
- [249] F. Sahba, H. R. Tizhoosh, and M. M. Salama, "A reinforcement learning framework for medical image segmentation," in *IJCNN*, vol. 6, 2006, pp. 511–517.
- [250] —, "Application of opposition-based reinforcement learning in image segmentation," in *2007 IEEE Symposium on Computational Intelligence in Image and Signal Processing*. IEEE, 2007, pp. 246–251.
- [251] —, "Application of reinforcement learning for segmentation of transrectal ultrasound images," *BMC Medical Imaging*, vol. 8, no. 1, p. 8, 2008.
- [252] F. Sahba, "Object segmentation in image sequences using reinforcement learning," in *CSCI'16*. IEEE, 2016, pp. 1416–1417.
- [253] D. Liu and T. Jiang, "Deep reinforcement learning for surgical gesture segmentation and classification," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 247–255.
- [254] F. C. Ghesu, B. Georgescu, T. Mansi, D. Neumann, J. Hornegger, and D. Comaniciu, "An artificial agent for anatomical landmark detection in medical images," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2016, pp. 229–237.
- [255] F. C. Ghesu, B. Georgescu, Y. Zheng, S. Grbic, A. Maier, J. Hornegger, and D. Comaniciu, "Multi-scale deep reinforcement learning for realtime 3d-landmark detection in ct scans," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [256] F. C. Ghesu, B. Georgescu, S. Grbic, A. Maier, J. Hornegger, and D. Comaniciu, "Towards intelligent robust detection of anatomical structures in incomplete volumetric data," *Medical Image Analysis*, vol. 48, pp. 203–213, 2018.
- [257] M. Etcheverry, B. Georgescu, B. Odry, T. J. Re, S. Kaushik, B. Geiger, N. Mariappan, S. Grbic, and D. Comaniciu, "Nonlinear adaptively learned optimization for object localization in 3d medical images," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 254–262.
- [258] A. Alansary, O. Oktay, Y. Li, L. Le Folgoc, B. Hou, G. Vaillant, B. Glocker, B. Kainz, and D. Rueckert, "Evaluating reinforcement learning agents for anatomical landmark detection," 2018.
- [259] A. Alansary, L. L. Folgoc, G. Vaillant, O. Oktay, Y. Li, W. Bai, J. Passerat-Palmbach, R. Guerrero, K. Kamnitsas, B. Hou *et al.*, "Automatic view planning with multi-scale deep reinforcement learning agents," *arXiv preprint arXiv:1806.03228*, 2018.
- [260] W. A. Al and I. D. Yun, "Partial policy-based reinforcement learning for anatomical landmark localization in 3d medical images," *arXiv preprint arXiv:1807.02908*, 2018.
- [261] R. Liao, S. Miao, P. de Tournemire, S. Grbic, A. Kamen, T. Mansi, and D. Comaniciu, "An artificial agent for robust image registration." in *AAAI*, 2017, pp. 4168–4175.
- [262] K. Ma, J. Wang, V. Singh, B. Tamersoy, Y.-J. Chang, A. Wimmer, and T. Chen, "Multimodal image registration with deep context reinforcement learning," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 240–248.
- [263] J. Krebs, T. Mansi, H. Delingette, L. Zhang, F. C. Ghesu, S. Miao, A. K. Maier, N. Ayache, R. Liao, and A. Kamen, "Robust non-rigid registration through agent-based action learning," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 344–352.
- [264] G. Maicas, G. Carneiro, A. P. Bradley, J. C. Nascimento, and I. Reid, "Deep reinforcement learning for active breast lesion detection from dce-mri," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 665–673.
- [265] P. Zhang, F. Wang, and Y. Zheng, "Deep reinforcement learning for vessel centerline tracing in multi-modality 3d volumes," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 755–763.
- [266] S. M. B. Netto, V. R. C. Leite, A. C. Silva, A. C. de Paiva, and A. de Almeida Neto, "Application on reinforcement learning for diagnosis based on medical image," in *Reinforcement Learning*. InTech, 2008.
- [267] S. J. Fakhri and T. K. Das, "Lead: a methodology for learning efficient approaches to medical diagnosis," *IEEE Transactions on Information Technology in Biomedicine*, vol. 10, no. 2, pp. 220–228, 2006.
- [268] K. Roberts, M. S. Simpson, E. M. Voorhees, and W. R. Hersh, "Overview of the trec 2016 clinical decision support track." in *TREC*, 2016.
- [269] Y. Ling, S. A. Hasan, V. Datla, A. Qadir, K. Lee, J. Liu, and O. Farri, "Learning to diagnose: Assimilating clinical narratives using deep reinforcement learning," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, vol. 1, 2017, pp. 895–905.
- [270] R. Ballard-Barbash, S. H. Taplin, B. C. Yankaskas, V. L. Ernster, R. D. Rosenberg, P. A. Carney, W. E. Barlow, B. M. Geller, K. Kerlikowske, B. K. Edwards *et al.*, "Breast cancer surveillance consortium: a national mammography screening and outcomes database." *American Journal of Roentgenology*, vol. 169, no. 4, pp. 1001–1008, 1997.
- [271] T. Chu, J. Wang, and J. Chen, "An adaptive online learning framework for practical breast cancer diagnosis," in *Medical Imaging 2016: Computer-Aided Diagnosis*, vol. 9785. International Society for Optics and Photonics, 2016, p. 978524.
- [272] K.-F. Tang, H.-C. Kao, C.-N. Chou, and E. Y. Chang, "Inquire and diagnose: Neural symptom checking ensemble using deep reinforcement learning," in *Proceedings of NIPS Workshop on Deep Reinforcement Learning*, 2016.
- [273] H.-C. Kao, K.-F. Tang, and E. Y. Chang, "Context-aware symptom checking for disease diagnosis using hierarchical reinforcement learning," 2018.
- [274] E. Y. Chang, M.-H. Wu, K.-F. T. Tang, H.-C. Kao, and C.-N. Chou, "Artificial intelligence in xprize deepq tricorder," in *Proceedings of the 2nd International Workshop on Multimedia for Personal Health and Health Care*. ACM, 2017, pp. 11–18.
- [275] E. Y. Chang, "Deepq: Advancing healthcare through artificial intelligence and virtual reality," in *Proceedings of the 2017 ACM on Multimedia Conference*. ACM, 2017, pp. 1068–1068.
- [276] Z. Wei, Q. Liu, B. Peng, H. Tou, T. Chen, X. Huang, K.-F. Wong, and X. Dai, "Task-oriented dialogue system for automatic diagnosis," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, vol. 2, 2018, pp. 201–207.
- [277] F. Tang, K. Lin, I. Uchendu, H. H. Dodge, and J. Zhou, "Improving mild cognitive impairment prediction via reinforcement learning and dialogue simulation," *arXiv preprint arXiv:1802.06428*, 2018.
- [278] H.-J. Schuetz and R. Kolisch, "Approximate dynamic programming for capacity allocation in the service industry," *European Journal of Operational Research*, vol. 218, no. 1, pp. 239–250, 2012.
- [279] Z. Huang, W. M. van der Aalst, X. Lu, and H. Duan, "Reinforcement learning based resource allocation in business process management," *Data & Knowledge Engineering*, vol. 70, no. 1, pp. 127–145, 2011.
- [280] B. Zeng, A. Turkcan, J. Lin, and M. Lawley, "Clinic scheduling models with overbooking for patients with heterogeneous no-show probabilities," *Annals of Operations Research*, vol. 178, no. 1, pp. 121–144, 2010.
- [281] T. S. M. T. Gomes, "Reinforcement learning for primary care e appointment scheduling," 2017.
- [282] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *ICML*, 2016, pp. 1928–1937.
- [283] K. Li and J. W. Burdick, "A function approximation method for modelbased high-dimensional inverse reinforcement learning," *arXiv preprint arXiv:1708.07738*, 2017.
- [284] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral surgical pattern cutting in 2d orthotropic gauze with deep

- reinforcement learning policies for tensioning," in *IEEE ICRA'17*. IEEE, 2017, pp. 2371–2378.
- [285] T. T. Nguyen, N. D. Nguyen, F. Bello, and S. Nahavandi, "A new tensioning method using deep reinforcement learning for surgical pattern cutting," *arXiv preprint arXiv:1901.03327*, 2019.
- [286] J. Chen, H. Y. Lau, W. Xu, and H. Ren, "Towards transferring skills to flexible surgical robots with programming by demonstration and reinforcement learning," in *ICACI'16*. IEEE, 2016, pp. 378–384.
- [287] D. Baek, M. Hwang, H. Kim, and D.-S. Kwon, "Path planning for automation of surgery robot based on probabilistic roadmap and reinforcement learning," in *2018 15th International Conference on Ubiquitous Robots (UR)*. IEEE, 2018, pp. 342–347.
- [288] K. Li, M. Rath, and J. W. Burdick, "Inverse reinforcement learning via function approximation for clinical motion analysis," in *IEEE ICRA'18*. IEEE, 2018, pp. 610–617.
- [289] K. M. Jagodnik, P. S. Thomas, A. J. van den Bogert, M. S. Branicky, and R. F. Kirsch, "Training an actor-critic reinforcement learning controller for arm movement using human-generated rewards," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 10, pp. 1892–1905, 2017.
- [290] R. S. Istepanian, N. Y. Philip, and M. G. Martini, "Medical qos provision based on reinforcement learning in ultrasound streaming over 3.5 g wireless systems," *IEEE Journal on Selected areas in Communications*, vol. 27, no. 4, 2009.
- [291] A. Alinejad, N. Y. Philip, and R. S. Istepanian, "Cross-layer ultrasound video streaming over mobile wimax and hsupa networks," *IEEE transactions on Information Technology in Biomedicine*, vol. 16, no. 1, pp. 31–39, 2012.
- [292] K. Ragnarsson, "Functional electrical stimulation after spinal cord injury: current use, therapeutic effects and future directions," *Spinal cord*, vol. 46, no. 4, p. 255, 2008.
- [293] P. S. Thomas, M. Branicky, A. Van Den Bogert, and K. Jagodnik, "Creating a reinforcement learning controller for functional electrical stimulation of a human arm," in *The Yale Workshop on Adaptive and Learning Systems*, vol. 49326. NIH Public Access, 2008, p. 1.
- [294] P. S. Thomas, A. J. van den Bogert, K. M. Jagodnik, and M. S. Branicky, "Application of the actor-critic architecture to functional electrical stimulation control of a human arm," in *IAAI*, 2009.
- [295] I. Kola and J. Landis, "Can the pharmaceutical industry reduce attrition rates?" *Nature reviews Drug discovery*, vol. 3, no. 8, p. 711, 2004.
- [296] G. Schneider, *De novo molecular design*. John Wiley & Sons, 2013.
- [297] M. Olivecrona, T. Blaschke, O. Engkvist, and H. Chen, "Molecular de-novo design through deep reinforcement learning," *Journal of Cheminformatics*, vol. 9, no. 1, p. 48, 2017.
- [298] A. Serrano, B. Imbernon, H. P´erez-S´anchez, J. M. Cecilia, A. Bueno-Crespo, and J. L. Abellan, "Accelerating drugs discovery with deep reinforcement learning: An early approach," in *Proceedings of the 47th International Conference on Parallel Processing Companion*. ACM, 2018, p. 6.
- [299] D. Neil, M. Segler, L. Guasch, M. Ahmed, D. Plumbley, M. Sellwood, and N. Brown, "Exploring deep recurrent models with reinforcement learning for molecule design," 2018.
- [300] M. Popova, O. Isayev, and A. Tropsha, "Deep reinforcement learning for de novo drug design," *Science Advances*, vol. 4, no. 7, p. eaap7885, 2018.
- [301] E. Yom-Tov, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and I. Hochberg, "Encouraging physical activity in patients with diabetes: Intervention using a reinforcement learning system," *Journal of Medical Internet Research*, vol. 19, no. 10, 2017.
- [302] I. Hochberg, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and E. Yom-Tov, "A reinforcement learning system to encourage physical activity in diabetes patients," *arXiv preprint arXiv:1605.04070*, 2016.
- [303] A. Baniya, S. Herrmann, Q. Qiao, and H. Lu, "Adaptive interventions treatment modelling and regimen optimization using sequential multiple assignment randomized trials (smart) and q-learning," in *Proceedings of IIE Annual Conference*, 2017, pp. 1187–1192.
- [304] E. M. Forman, S. G. Kerrigan, M. L. Butryn, A. S. Juarascio, S. M. Manasse, S. Ontan˜on, D. H. Dallal, R. J. Crochiere, and D. Moskow, "Can the artificial intelligence technique of reinforcement learning use continuously-monitored digital data to optimize treatment for weight loss?" *Journal of Behavioral Medicine*, pp. 1–15, 2018.
- [305] O. Gottesman, F. Johansson, J. Meier, J. Dent, D. Lee, S. Srinivasan, L. Zhang, Y. Ding, D. Wihl, X. Peng *et al.*, "Evaluating reinforcement learning algorithms in observational health settings," *arXiv preprint arXiv:1805.12298*, 2018.
- [306] A. Raghu, O. Gottesman, Y. Liu, M. Komorowski, A. Faisal, F. DoshiVelez, and E. Brunskill, "Behaviour policy estimation in off-policy policy evaluation: Calibration matters," *arXiv preprint arXiv:1807.01066*, 2018.
- [307] R. Jeter, C. Josef, S. Shashikumar, and S. Nemati, "Does the" artificial intelligence clinician" learn optimal treatment strategies for sepsis in intensive care?" *arXiv preprint arXiv:1902.03271*, 2019.
- [308] D. Koller, N. Friedman, and F. Bach, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [309] J. L. Jameson and D. L. Longo, "Precision medicinepersonalized, problematic, and promising," *Obstetrical & Gynecological Survey*, vol. 70, no. 10, pp. 612–614, 2015.
- [310] J. Furnkranz and E. H. Ullermeier, "Preference learning," in *Encyclopedia of Machine Learning*. Springer, 2011, pp. 789–795.
- [311] D. J. Lizotte, M. H. Bowling, and S. A. Murphy, "Efficient reinforcement learning with multiple reward functions for randomized controlled trial analysis," in *ICML'10*. Citeseer, 2010, pp. 695–702.
- [312] M. Herman, T. Gindele, J. Wagner, F. Schmitt, and W. Burgard, "Inverse reinforcement learning with simultaneous estimation of rewards and dynamics," in *Artificial Intelligence and Statistics*, 2016, pp. 102–110.
- [313] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *NIPS'17*, 2017, pp. 5048–5058.
- [314] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu, "Reinforcement learning with unsupervised auxiliary tasks," *arXiv preprint arXiv:1611.05397*, 2016.
- [315] S. Racaniere, T. Weber, D. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. Heess, Y. Li *et al.*, "Imaginationaugmented agents for deep reinforcement learning," in *Advances in neural information processing systems*, 2017, pp. 5690–5701.
- [316] N. Jiang and L. Li, "Doubly robust off-policy value evaluation for reinforcement learning," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*. JMLR. org, 2016, pp. 652–661.
- [317] M. Kearns and S. Singh, "Near-optimal reinforcement learning in polynomial time," *Machine learning*, vol. 49, no. 2-3, pp. 209–232, 2002.
- [318] J. Papis and R. Parr, "Efficient pac-optimal exploration in concurrent, continuous state mdps with delayed updates." in *AAAI*, 2016, pp. 1977–1985.
- [319] M. Dimakopoulou and B. Van Roy, "Coordinated exploration in concurrent reinforcement learning," in *International Conference on Machine Learning*, 2018, pp. 1270–1278.
- [320] Z. Guo and E. Brunskill, "Concurrent pac rl." in *AAAI*, 2015, pp. 2624–2630.
- [321] J. Fu, J. Co-Reyes, and S. Levine, "Ex2: Exploration with exemplar models for deep reinforcement learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 2574–2584.
- [322] H. Tang, R. Houthoofd, D. Foote, A. Stooke, O. X. Chen, Y. Duan, J. Schulman, F. DeTurck, and P. Abbeel, "# exploration: A study of count-based exploration for deep reinforcement learning," in *NIPS'17*, 2017, pp. 2750–2759.
- [323] B. C. Stadie, S. Levine, and P. Abbeel, "Incentivizing exploration in reinforcement learning with deep predictive models," *arXiv preprint arXiv:1507.00814*, 2015.
- [324] T. Mannucci, E.-J. van Kampen, C. de Visser, and Q. Chu, "Safe exploration algorithms for reinforcement learning controllers," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 4, pp. 1069–1081, 2018.
- [325] C. A. Merck and S. Kleinberg, "Causal explanation under indeterminism: A sampling approach." in *AAAI*, 2016, pp. 1037–1043.
- [326] J. Woodward, *Making things happen: A theory of causal explanation*. Oxford university press, 2005.

- [327] S. L. Morgan and C. Winship, *Counterfactuals and causal inference*. Cambridge University Press, 2015.
- [328] D. Dash, M. Voortman, and M. De Jongh, "Sequences of mechanisms for causal reasoning in artificial intelligence." in *IJCAI*, 2013, pp. 839–845.
- [329] Z. C. Lipton, "The mythos of model interpretability," *Communications of the ACM*, vol. 61, no. 10, pp. 36–43, 2018.
- [330] S. Bhupatiraju, K. K. Agrawal, and R. Singh, "Towards mixed optimization for reinforcement learning with program synthesis," *arXiv preprint arXiv:1807.00403*, 2018.
- [331] Z. C. Lipton, "The doctor just won't accept that!" *arXiv preprint arXiv:1711.08037*, 2017.
- [332] F. Maes, R. Fonteneau, L. Wehenkel, and D. Ernst, "Policy search in a space of simple closed-form formulas: towards interpretability of reinforcement learning," in *International Conference on Discovery Science*. Springer, 2012, pp. 37–51.
- [333] A. Verma, V. Murali, R. Singh, P. Kohli, and S. Chaudhuri, "Programmatically interpretable reinforcement learning," in *International Conference on Machine Learning*, 2018, pp. 5052–5061.
- [334] F. Elizalde, E. Sucar, J. Noguez, and A. Reyes, "Generating explanations based on markov decision processes," in *Mexican International Conference on Artificial Intelligence*. Springer, 2009, pp. 51–62.
- [335] Z. Che, S. Purushotham, R. Khemani, and Y. Liu, "Distilling knowledge from deep networks with applications to healthcare domain," *arXiv preprint arXiv:1512.03542*, 2015.
- [336] M. Wu, M. C. Hughes, S. Parbhoo, M. Zazzi, V. Roth, and F. DoshiVelez, "Beyond sparsity: Tree regularization of deep models for interpretability," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [337] A. E. Gaweda, M. K. Muezzinoglu, G. R. Aronoff, A. A. Jacobs, J. M. Zurada, and M. E. Brier, "Incorporating prior knowledge into q-learning for drug delivery individualization," in *Fourth International Conference on Machine Learning and Applications*. IEEE, 2005, pp. 6–pp.
- [338] A. Holzinger, "Interactive machine learning for health informatics: when do we need the human-in-the-loop?" *Brain Informatics*, vol. 3, no. 2, pp. 119–131, 2016.
- [339] D. Abel, J. Salvatier, A. Stuhlmuller, and O. Evans, "Agentagnostic human-in-the-loop reinforcement learning," *arXiv preprint arXiv:1701.04079*, 2017.
- [340] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," *Nature medicine*, vol. 25, no. 1, p. 44, 2019.
- [353] H. Lei, A. Tewari, and S. Murphy, "An actor-critic contextual bandit algorithm for personalized interventions using mobile devices," *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [341] C. Yu, D. Wang, T. Yang, W. Zhu, Y. Li, H. Ge, and J. Ren, "Adaptively shaping reinforcement learning agents via human reward," in *Pacific Rim International Conference on Artificial Intelligence*. Springer, 2018, pp. 85–97.
- [342] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Advances in neural information processing systems*, 2013, pp. 2625–2633.
- [343] J. Shu, Z. Xu, and D. Meng, "Small sample learning in big data era," *arXiv preprint arXiv:1808.04572*, 2018.
- [344] S. W. Carden and J. Livsey, "Small-sample reinforcement learning: Improving policies using synthetic data 1," *Intelligent Decision Technologies*, vol. 11, no. 2, pp. 167–175, 2017.
- [345] J. Salamon and J. P. Bello, "Deep convolutional neural networks and data augmentation for environmental sound classification," *IEEE Signal Processing Letters*, vol. 24, no. 3, pp. 279–283, 2017.
- [346] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [347] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *stat*, vol. 1050, p. 9, 2015.
- [348] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and Brain Sciences*, vol. 40, 2017.
- [349] Y.-L. Zheng, X.-R. Ding, C. C. Y. Poon, B. P. L. Lo, H. Zhang, X.-L. Zhou, G.-Z. Yang, N. Zhao, and Y.-T. Zhang, "Unobtrusive sensing and wearable devices for health informatics," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 5, pp. 1538–1554, 2014.
- [350] G. Acampora, D. J. Cook, P. Rashidi, and A. V. Vasilakos, "A survey on ambient intelligence in healthcare," *Proceedings of the IEEE*, vol. 101, no. 12, pp. 2470–2494, 2013.
- [351] F. Zhu, J. Guo, R. Li, and J. Huang, "Robust actor-critic contextual bandit for mobile health (mhealth) interventions," in *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*. ACM, 2018, pp. 492–501.
- [352] F. Zhu, J. Guo, Z. Xu, P. Liao, L. Yang, and J. Huang, "Groupdriven reinforcement learning for personalized mhealth intervention," in *International Conference on Medical Image Computing and Computer Assisted Intervention*. Springer, 2018, pp. 590–598.