

Improving Sepsis Treatment Strategies by Combining Deep and Kernel-Based Reinforcement Learning

Xuefeng Peng, MSc.¹, Yi Ding, MSc.², David Wihl, A.L.B.¹, Omer Gottesman, PhD¹,
Matthieu Komorowski, MD³, Li-wei H. Lehman, PhD⁴, Andrew Ross, MSE¹,
Aldo Faisal, PhD³, Finale Doshi-Velez, PhD¹

¹Harvard University, Paulson School of Engineering and Applied Sciences, Cambridge, MA

²Harvard University, T.H. Chan School of Public Health, Cambridge, MA

³Imperial College London, London, UK

⁴MIT, Institute for Medical Engineering & Science, Cambridge, MA

Abstract

Sepsis is the leading cause of mortality in the ICU. It is challenging to manage because individual patients respond differently to treatment. Thus, tailoring treatment to the individual patient is essential for the best outcomes. In this paper, we take steps toward this goal by applying a mixture-of-experts framework to personalize sepsis treatment. The mixture model selectively alternates between neighbor-based (kernel) and deep reinforcement learning (DRL) experts depending on patient's current history. On a large retrospective cohort, this mixture-based approach outperforms physician, kernel only, and DRL-only experts.

Introduction

Sepsis is a medical emergency requiring rapid treatment.¹ It is the cause of 6.0% of hospital admissions but 15.0% of hospital mortality.² It is also costly: in 2011 alone, the US spent \$20.3 billion dollars on hospital care for patients with sepsis.³ Managing sepsis remains challenging, in part because there exists large variation in patient response to existing sepsis management strategies.⁴

In this work, we focus on two interventions in the context of sepsis management: intravenous (IV) fluid (adjusted for fluid tonicity) and vasopressors (VP). These two drugs are respectively used to correct the hypovolemia and counteract sepsis-induced vasodilation. While hypovolemia and vasodilation are common among patients with sepsis, there exists little clinical consensus about when and how these should be treated.⁵ However, these choices can have large implications for patient mortality:⁴ vasopressors are known to have harmful effects in certain patients, and recent studies have also demonstrated the association between fluid-overload and negative outcomes in the ICUs.⁶

Thus, it is essential to identify ways to personalize treatment. The availability of large observational critical care data sets⁷ has made it possible to hypothesize improved sepsis management strategies, and prior studies^{8,9} have used this resource to suggest optimal treatment strategies for patients with sepsis. As with those earlier works, we personalize strategies by using reinforcement learning (RL), a technique for optimizing sequences of decisions given patient context. However, we use a mixture-of-experts approach to combine two very different RL techniques with very different strengths—a model-free deep RL approach (DRL) and a model-based kernel RL approach (KRL)—to improve the quality of the recommended treatment policy. Specifically, our work extends prior efforts in three important ways:

1. *Recurrent encoding of the patient's history.* To date, work in this domain has assumed that the patient's current measurements are sufficient to summarize their history. To retain potentially decision-relevant information from the patient's past, we use a recurrent autoencoder to represent the patient's entire history.
2. *Safe-guards on the Deep RL.* DRL-based approaches can be particularly poor at extrapolating, and even in areas of dense data, they can suggest non-sensical actions. We explicitly restrict the DRL to only suggest actions commonly taken by clinicians, moving us toward more clinically-credible policies.
3. *Combining DRL with Kernel RL.* Finally, we use a mixture-of-experts (MoE) to combine the restricted DRL approach with a kernel RL approach selectively based on the context. DRL is more flexible but can be prone to various pathologies; KRL is guaranteed to stay close to the data but as a result can extrapolate poorly. The MoE combines their strengths.

Related Work

Reinforcement Learning has been applied to a number of applications in healthcare, ranging from emergency decision support,¹⁰ treating malaria,¹¹ and managing HIV.¹² Prasad et al.¹³ use RL to identify when to the wean patients from mechanical ventilation in ICUs.

With respect to fluid and vasopressor use in sepsis management, Komorowski et al.⁹ model a discrete Markov decision process from data and then utilize it to learn a treatment strategy. Raghu et al.⁸ extend this work by considering a much more expressive continuous representation of patient state. They use a traditional, non-recurrent autoencoder to first compress measurements from each time step into a continuous state representation, and then they learn a mapping from the state representation to an appropriate treatment via a Dueling Double-Deep Q Network (Dueling DDQN). Our work uses an even more expressive state representation that represents the patient’s entire history, and we also add safe-guards against inappropriate actions and develop richer policies through our mixture of experts.

The mixture of experts aspect of our work builds from ideas developed by Parbhoo et al.¹⁴ in the context of HIV management. In their case, they switch between a kernel-based policy and a discrete Bayesian Partially Observable Markov Decision Process (POMDP). We follow the idea of combining experts, but use the DDQN as an expert rather than a discrete POMDP, as Raghu et al.⁸ have already demonstrated that a continuous expressive state representation is valuable for the sepsis management task. Because we use a recurrent encoding to summarize the entire patient history, our state can be thought of as a sufficient statistic, much like the POMDP belief state in Parbhoo et al.¹⁴

Background

The reinforcement learning framework models a sequence of decisions as an agent interacting with an environment over time. At each time step t , the RL agent observes a state s from the state space S and selects an action a from the action space A based on some policy $\pi(s, a)$, which assigns a probability to action in each state. Upon taking the action, the agent receives some reward r and transitions to a new state s' . The agent’s goal is to maximize their expected longterm discounted return $\mathbb{E}[\sum_t \gamma^t r_t]$. The optimal value function is defined as $V^*(s) = \max_{\pi} \mathbb{E}[\sum_t \gamma^t r_t | s_0 = s, \pi]$, and the optimal state-action value function $Q^*(s, a) = \max_{\pi} \mathbb{E}[\sum_t \gamma^t r_t | s_0 = s, a_0 = a, \pi]$. The latter satisfies the Bellman equation $Q^*(s, a) = r(s, a) + \gamma \max_{a'} \mathbb{E}[Q^*(s', a')]$, where γ is the discount factor determines the trade-off between immediate and future rewards. Q-learning methods aim to learn an optimal policy by minimizing the temporal difference (TD) error, defined as $r(s, a) + \gamma Q^*(s', a') - Q^*(s, a)$.

Cohort and Data Processing

Cohort. We used the same patient set as in Raghu et al.⁸ which applied the Sepsis-3 criteria to the Multi-parameter Intelligent Monitoring in Intensive Care (MIMIC-III v1.4) database.⁷ Our cohort consisted of 15,415 adults (age range of 18 to 91), summarized in Table 1.

Table 1: Comparison of cohort statistics for subjects that fulfilled the Sepsis-3 criteria

	% Female	Mean Age	Total Population
Survivors	44.1%	63.9	13,535
Non-survivors	44.3%	67.4	1,880

Cleaning and Preprocessing. As in Raghu et al.⁸, patient histories were partitioned into 4-hour windows each containing fifty attributes, ranging from vitals (e.g. heart rate, mean blood pressure) to clinician assessments of the patient’s conditions (e.g. sequential organ failure assessment (SOFA) score). Patients with missing values were excluded. The observations, which range from demographics, lab values and vital signs, all have different scales. Following Raghu et al.,⁸ we performed log transformations of observations with large values and standardized the remaining values. After the standardization and log transformation, all values were rescaled into $[0 - 1]$. (See [Appendices](#) for details of attributes and preprocessing.) The data set was split into a fixed 75% training and validation set and a 25% test set.

Treatment Discretization. In this work, we focus on administering two drugs: intravenous (IV) fluid and vasopressor (VP). In the cohort, the usage of IV and VP for each patient are recorded at each 4-hour window. Following Raghu et al.⁸, the dosages for each drug are discretized into 5 bins, resulting in a 5×5 action space indexed from 0 to 24. Note

that the first action ($a = 0$) means “no action”—neither IV nor VP are prescribed.

Method Overview

Applying RL to the sepsis management problem involves several pieces. The first is defining our inputs and treatments (sections above). Next we describe how we compress patient histories into a state via a recurrent autoencoder, how we attribute rewards to each state, and also how we determine the quality of some policy given observational data. With these pieces in place, we finally describe how to derive treatment policies that optimize for our rewards, including our mixture-of-experts (MoE) approach.

Compressing Patient Histories

Prior efforts^{8,9} assumed that the patient’s current measures were sufficient to summarize their history; however, past and trend patient information is often valuable to deciding the appropriate course of action. To capture more of this temporal information, we encoded patient states recurrently using an LSTM autoencoder representing the cumulative history for each patient. The LSTM had a single layer of 128 hidden units for both encoder and decoder—resulting in a state s that consisted of 128 real-valued vector. The autoencoder was trained to minimize the reconstruction loss (MSE) between the original measurements and the decoded measurements. We trained the model with mini-batches of 128 and the Adam optimizer for 50 epochs until its convergence.

Reward Formulation

Broadly, we are interested in reducing mortality among patients with sepsis. However, mortality is a challenging objective to optimize for because it is only observed after a long sequence of decisions; it can be hard to ascertain which action was responsible for a good or bad outcome. Thus, for the purposes of our training, we introduce an *intermediate* reward that can give a preliminary signal for whether our sequence of treatment decisions is likely to reduce mortality.

Specifically, we first train a regressor that predicts the probability of mortality given a patient’s current observations (implications of this choice in [Discussion](#)). Next, we define the reward as the change in the negative mortality log-odds of mortality between the current observations and the next observations. (Log-odds were used because the probabilities of mortality actually vary over a relatively small range.) Let $f(o)$ be the probability of mortality given current observations o . Then we define the reward $r(o, a, o')$ as

$$r(o, a, o') = -\log \frac{f(o')}{1 - f(o')} f(o') + \log \frac{f(o)}{1 - f(o)} \quad (1)$$

For a sense of scale, among those over 186K patient state transitions from both training and testing sets, the rewards are in the interval $[-3, 3]$.

The mortality predictor $f(o)$ itself was a two-layer neural network with 64 and 32 units for each layer and L1-regularized gradients (see Ross et al.¹⁵ for details). The L1 regularization encourages sparse local linear approximations, and thus makes its behaviors more interpretable. We resampled balanced batches (between survivors and non-survivors) during training with batch sizes of 128 observations for 50 epochs, and our predictor $f(o)$ achieved a test accuracy of 73.1%. (See [Appendices](#) for log-odds distributions for each class.)

Off-Policy Evaluation via the WDR Estimator

The natural question is how to evaluate the quality of a proposed policy π_e given only retrospective data collected according to a clinician policy π_b . The weighted doubly robust (WDR) estimator¹⁶ is widely used for off-policy evaluation in RL. It uses estimated value \hat{V} and action-value \hat{Q} functions as control variates to reduce the variance of the off-policy estimation. In the following, we used the value function of the DRL policy as our control variate; we also explored using the value of the clinician policy, estimated as the mean, instead of the max, of the DRL Q-values over action space.

$$\text{WDR}(D) := \sum_{i=1}^I \sum_{t=0}^T \gamma^t w_t^i r_t^{H_i} - \sum_{i=1}^I \sum_{t=0}^T \gamma^t (w_t^i \hat{Q}^{\pi_e}(S_t^{H_i}, A_t^{H_i}) - w_{t-1}^i \hat{V}^{\pi_e}(S_t^{H_i})) \quad (2)$$

Here, I is the number of patients, t is the time step, and H_i refers to the i_{th} patient's ICU-stay state trajectory. The importance weight of the state of the patient i at time step t is defined as $w_t^i = \frac{\rho_t^i}{\sum_{j=1}^I \rho_t^j}$, where $\rho_t^i = \prod_{t'=0}^t \frac{\pi_e(A_{t'}^{H_i} | S_{t'}^{H_i})}{\pi_b(A_{t'}^{H_i} | S_{t'}^{H_i})}$. The reward r , value function \hat{V} , and action-value function \hat{Q} are as defined above. Finally, we estimate the clinician policy π_b as the empirical distribution over actions of the 300 neighbors in the training set with states closest to s , as research shows that clinicians typically make decisions based on their experience treating similar patients¹⁷.

Deriving Policies

With a means of representing patient history, rewards, and a metric for evaluating policies, we can now start optimizing treatment strategies. Below we describe each expert—the DRL and the KRL—and how we combine them.

Kernel Policy. One simple way to derive a treatment decision rule is to look at the nearest neighbors to the current state s , identify the survivors, and choose actions that correspond to the distribution of treatments performed on these nearby survivors (see cartoon in Figure 1). Specifically, we

1. Get the encoding s from the patient's history h up to time t via the LSTM autoencoder.
2. Search k nearest neighbors in the training set in this encoded representation space using Euclidean distance.
3. The kernel policy π_k is the distribution of actions taken over the surviving nearest neighbors.

We cross-validated the k values ranging from 200 to 500 using WDR, and proceeded with $k = 300$.

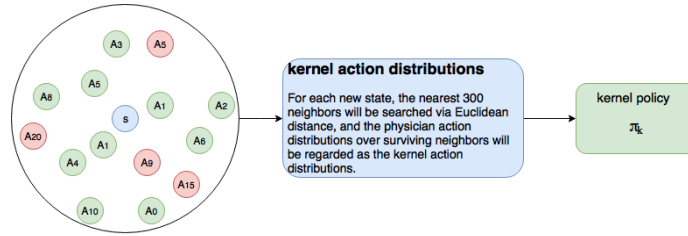


Figure 1: The circle in the left shows an example of the neighborhoods of a new state s , red and green marks the mortality and surviving states respectively, and each of these states is associated with a physician action A_i .

DQN Policy. Double DQN (DDQN) with dueling structure¹⁸, which is a variant of DQN, has been applied to derive a policy that outperforms the physician policy⁸. The structure of dueling DDQN is particularly suitable for sepsis treatment strategy learning, as it differentiates the value function V into the value of the patient's underlying physiological condition, called the *Value* stream, and the value of the treatment given, called the *Advantage* stream.

We train the dueling DDQN for 200,000 steps with *batch size* = 30 to minimize the TD-error. At each given state, the agent is trained to take an action with the highest Q-value, in order to achieve the ultimate goal of improving the overall survival rate. To stabilize the training process and improve the performance, we applied regularization term λ to the Q-network loss to penalize output Q-values which exceeded the maximum observed rewards $r_{max} = 3$

and used prioritized experience replay¹⁹ to balance the training sets with high-value and high-error states.

$$\mathcal{L}(\theta) = E[(Q_{double-target} - Q(s, a; \theta))^2] + \lambda \max(|Q(s, a; \theta) - r_{max}|, 0) \quad (3)$$

where

$$Q_{double-target} = r + \gamma Q(s, \arg \max_{a'} Q(s, a'; \theta')) \quad (4)$$

and θ, θ' are parameters of the DQN networks.

Finally, given a set of action-values $Q(s, a)$, we must still define a policy π_d . Typically, these actions are chosen by the $\max Q$ -value, but this ignores the fact that two actions may have very similar values—and given the limitations of our learning, it may not be possible which is actually the best. We define the DRL policy π_d as the softmax of the action-value or *Advantage* stream, giving higher probability to actions with higher values but not forcing us to take (what might be a brittle) best value.

Mixture-of-Experts (MoE)

The two approaches above—KRL and DRL—have different strengths. For patient states which are atypical, i.e. farther Euclidean distance away from any neighbors, the kernel policy may end up relying on neighbors that are not really that similar to the patient. In contrast, DQN, in trying to fit a value function to the whole state space at once, may still underfit in regions with plentiful data. Our mixture-of-experts (MoE) uses properties of the patient’s current state, and the relationship between the patient’s current state and states observed during training, to switch between the kernel and DQN policies (see Figure 2 for a cartoon).

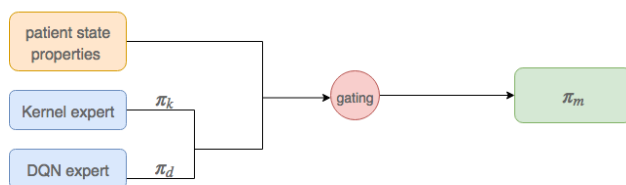


Figure 2: The architecture of MoE, it produces a mixed policy via combining kernel.

Action Restriction. The DDQN, as a complex function approximator, comes with relatively few guarantees. Sometimes, it can place high value on actions that were rarely or never performed by clinicians. To safe guard against these rare (and likely dangerous) actions, we restricted the DDQN to actions only taken more than 1% of the time by the physicians among its 300 nearest neighbors. Specifically, let $\pi_d(s, a)$ be the DDQN policy, and $\pi_b(s, a)$ be the physician policy. If $\pi_b(s, a) < 0.01$, then we set $\pi_d(s, a) \leftarrow 0$ and then normalize $\pi_d(s, a)$ to be a valid probability distribution. (We note that the kernel policy, which is derived directly from clinician actions, cannot deviate in this way; once we restrict the DDQN actions, the MoE will also never take rare actions.)

Choice of Gating Function. We examined several medical sources^{20–22} to determine which features might be most useful for selecting between experts. Our final set of features were: age, Elixhauser, SOFA, FiO_2 , BUN, GCS, Albumin, trajectory length, and max distance from neighbors. For our MoE gating function we combined these features x linearly via weights w , along with a bias term b , and passed them through a logit to get the probability of choosing each policy:

$$p_k = \text{sigmoid}(w \cdot x + b) \quad \text{and} \quad p_d = 1 - p_k \quad (5)$$

where p_k and p_d denote the assigned probability for choosing the kernel and DQN policy respectively.

Optimizing the Gating Function. Of course, the core question is how to choose the gating parameters so as to maximize long-term rewards. Given a set of weights w and bias b , the MoE policy π_m is defined as $\pi_m(s, a) = p_k \pi_k(s, a) + p_d \pi_d(s, a)$. We can estimate the expected discounted return of the policy π_m WDR from above; we again perform gradient descent on the gating parameters (with a minibatch of 256 samples). Due to the nonconvexity of the WDR-based objective, we take the best of 1000 random restarts.

Results

The estimates of the discounted expected return for each policy are presented in Table 2. We provide two columns for the mixture of experts policy because it is challenging to derive accurate value estimates \hat{V} and \hat{Q} for the WDR

estimator in this case. Thus, we consider two sensible options: using the value estimates \hat{V} and \hat{Q} from the clinician policy, and using the estimates from the DDQN policy.

Regardless of the choice of evaluation covariate, both kernel and DQN policy improve over the physician policy, and the MoE policy projects a further improvement. Using a recurrent state representation that compresses the entire history results in a further improvement for all RL policies (we used a sparse autoencoder²³ for the non-recurrent encoding). The [Appendices](#) include bootstrap intervals for these values²⁴; while the variances are large, we find the MoE still averages values higher than the other policies.

Table 2: Estimate of the discounted expected return for policies over test set, $\gamma = 0.99$. V_d indicates approximating the MoE V by DQN V function, V_b indicates approximating the MoE V by behavioral policy, namely, physician V function

	Physician	Kernel	DQN	MoE_{V_d, Q_d}	MoE_{V_b, Q_b}
non-recurrent encoded	3.76	3.73	4.06	3.93	4.31
recurrent encoded	3.76	4.46	4.23	5.03	5.72

Analysis of discovered policies. Figure 3a shows the action distributions for the KRL, DRL, clinician, and MoE policies over the test set. The no treatment $a = 0$ action dominates policies. Actions which are favored by physicians, such as IV but no vasopressor are also (as expected) favored by the kernel policy. That said, the kernel policy tends to be more conservative than the clinicians as it suggests nonaction at approximately twice the clinicians’ frequency—perhaps reflecting a bias toward the fact that those patients who were not treated were somehow healthier and thus survived. Perhaps in a similar vein, the kernel expert prescribes more fluid alone than the clinicians: while both suggest high fluids and no vasopressor almost in the same frequency, the kernel expert very rarely suggests actions with vasopressor.

The DRL policy, like clinician and the kernel policies, favors giving actions a range of fluid values. But, over the test set, DQN expert prescribes more extreme values than clinicians. Clinicians frequently prescribe high fluids and low vasopressor; however, the DQN policy also tends to give more high vasopressor dosage actions in addition to fluids.

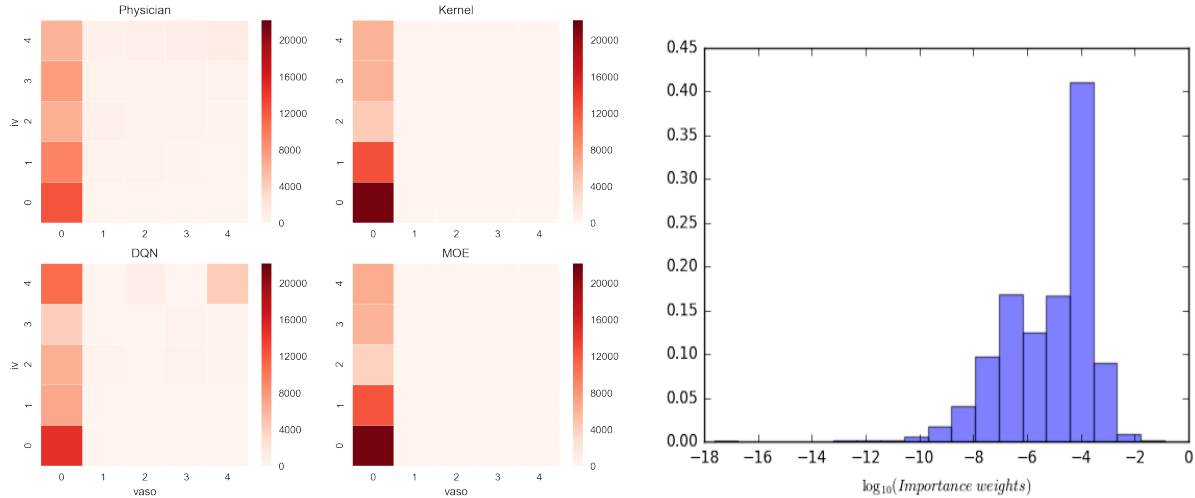
Over the test set, the MoE policy is closer to the kernel policy. But influenced by the DQN policy, MoE prescribes more high dosage actions. Table 3 shows how actions suggested by the experts overlap; the rate is high for the kernel and MoE policies again reflecting the fact that most of the time, patients can find similar neighbors. In 4.4% of circumstances, the gating results in a MoE policy follows neither that of kernel nor that of DQN policy.

Table 3: Percent similarity of different policies over test set patient states

	kernel	DQN	MoE
physician	0.305	0.151	0.296
kernel	-	0.182	0.871
DQN	-	-	0.258

Evaluation Quality Assessment. The WDR estimator of policy quality relies on having a large enough collection of patient histories in the evaluation set having non-zero weight w_t^i . For the MoE policy, 90% of the importance sampling weights are non-zero and 86% final weights in the sequences are non-zero. These high numbers of non-zero importance weights indicates that nearly all of our data was used in the evaluation of the policy. We plot the full distribution of weights in Figure 3b. A significant number of weights lie in the range of $[10^{-4}, 10^{-3}]$ and only very few observations have weights significantly larger than that range (the samples with significantly smaller weights are unlikely to have a significant influence on the estimate). However, the few observations with weights on the order of 10^{-1} could potentially have large influence (see [Appendices](#) for variances computed via bootstrap).

Running Time. Besides the quality of policy, the ability to make recommendations quickly is also important in the ICU (note that it is less important for the initial training time to be fast). We measured the computational time for all components in our framework on a dual-core Intel Core i5 processor. The 2-layer NN for the reward function



(a) Action Distributions for physician and experts over test set (b) WDR importance weights distribution
Figure 3: The left heatmap shows action distribution of each expert. Action assignment starts at bottom left corner (action 0) in the grid and increases from left to right, with action 24 (max level of fluid and vasopressor) corresponding to the top right corner. The MoE is much akin to the conservative kernel expert, suggesting most of the patients can find similar neighbors. The right chart plots the full distribution of importance weights. A significant number of weights lie in the range of $[10^{-4}, 10^{-3}]$ and only very few observations have weights significantly larger than that range.

$f(o)$ took 56.8s to train by 128-sized mini-batch for 50 epochs on dataset with 39856×45 dimension. The recurrent autoencoder for the state space took 491s for 50 epochs with 128 patients per mini-batch.

The kernel policy required 909s to identify policies for the 150720 samples in the training set. Training DQN by sampling 30 transitions for 200,000 times required 5.15×10^4 s. Finally, MoE gating function took 62s to train by a 256-sized mini-batch, for 1 epoch using $1e-4$ learning rate, on dataset with 150720×9 dimension; however, since MoE cannot guarantee global maxima, we conducted 1000 random restarts over the initial parameters, and trained each for 50 epochs.

Most importantly, with regard to prediction at test time, for a patient with 10 timesteps in the ICU, it took only 0.162s to encode all the observations, compute the kernel and DRL policies, and compute the gating function for the final MoE policy.

Discussion

Overall, the DQN policy recommended a treatment strategy with more aggressive use of both vasopressors and fluids. In comparison to the physician policy, DQN recommended 70% more actions involving medium-to-high fluid volume and vasopressor dosage (actions 18,19,23, and 24). Most notably, frequency for the DQN action corresponding to maximum levels of both fluid and vasopressor (action 24) increased by 3.8 fold from the physician policy. These results suggest that despite the recent advances in deep reinforcement learning, further investigations are required, and careful clinical judgment should be exercised to guard against potentially high-risk actions introduced from pathologies in non-linear function approximation.

The proposed kernel policy displayed a different kind of bias. It recommended far fewer actions involving vasopressors in comparison to both the physician policy and DQN, perhaps because amongst a patient’s neighbors, the survivors were relatively healthier and thus treated less aggressively. By focusing on survivors the kernel policy also focuses on patients who did not just receive a good treatment or were potentially healthier now, but also patients that received good treatments and remained healthy in the future. If healthier patients are easier to treat in general, then we might expect a bias toward less aggressive treatment from the kernel policy as well.

More broadly, while it appears that our MoE policy significantly outperforms the clinician policy (as well as each

individual expert), and we have ensured that the actions it suggests are at least sensible (that is, often taken by clinicians), there still exist a number of limitations. When encoding the patient clinical course in a recurrent representation, in spite of our high prediction accuracy, we cannot be certain with only these 50 measures that there are no hidden confounding factors; aside from pre-ICU fluid balance, we have no information from prior to their ICU stay. The choice of representation also influences the quality of our off-policy evaluation, as the WDR estimator assumes that the system is Markov in the state. More generally, the WDR estimator requires either the behavior policy estimate π_b to be accurate or the control variate estimates \hat{V} and \hat{Q} to be accurate to be unbiased; something that we could not guarantee. That said, we do demonstrate that our results were at least insensitive to different choices of \hat{V} and \hat{Q} .

Our work also focused on a very specific reward structure. To apply off-policy evaluation (WDR) in a statistically credible way, we considered the accumulation of low mortality risk as the objective, rather than mortality itself (as using the latter, the policies could not be evaluated reliably). To maximize the interpretability of the reward to clinicians, this risk was calculated only from the current observations and not the patient's entire history. Creating reward functions that can both be checked by human experts and accurately convey clinical goals is a direction for future work.

Finally, with respect to sepsis management, there also exist many other interventions, such as antibiotics use and mechanical ventilation, that also affect patient outcomes. Future work remains to investigate policies that incorporated a broader scope of patient history as well as a larger variety of interventions.

Conclusion

We presented a MoE framework to learn improved fluid and vasopressor administration strategies for sepsis patients in ICUs using observational data. We demonstrated that the proposed mixture model approach can automatically adapt to patient states at each time step, and dynamically switch between a conservative kernel policy and a more aggressive deep-RL policy to achieve, under our reward measure, better expected outcomes than clinician, kernel policy, and deep-RL policy. While much further investigation is required to truly validate the efficacy of derived policies, the proposed MoE framework represents a novel approach to take advantage of the strengths of different treatment policies.

Acknowledgments

We would like to thank the other students in Harvard CS282R - Reinforcement Learning for Healthcare, Fall 2017 for their insights, encouragement and feedback. Omer Gottesman was supported by the Harvard Data Science Initiative. L. Lehman was supported by NIH grant 2RO1GM104987.

References

1. Christopher W Seymour, Foster Gesten, Hallie C Prescott, Marcus E Friedrich, Theodore J Iwashyna, Gary S Phillips, Stanley Lemeshow, Tiffany Osborn, Kathleen M Terry, and Mitchell M Levy. Time to treatment and mortality during mandated emergency care for sepsis. *New England Journal of Medicine*, 376(23):2235–2244, 2017.
2. Chanu Rhee, Raymund Dantes, Lauren Epstein, David J Murphy, Christopher W Seymour, Theodore J Iwashyna, Sameer S Kadri, Derek C Angus, Robert L Danner, Anthony E Fiore, et al. Incidence and trends of sepsis in us hospitals using clinical vs claims data, 2009-2014. *JAMA*, 318(13):1241–1249, 2017.
3. Anne Pfuntner, Lauren M Wier, and Claudia Steiner. Costs for hospital stays in the United States, 2011: Statistical brief# 168. 2006.
4. Jason Waechter, Anand Kumar, Stephen E Lapinsky, John Marshall, Peter Dodek, Yaseen Arabi, Joseph E Parrillo, R Phillip Dellinger, Allan Garland, Cooperative Antimicrobial Therapy of Septic Shock Database Research Group, et al. Interaction between fluids and vasoactive agents on mortality in septic shock: a multicenter, observational study. *Critical care medicine*, 42(10):2158–2168, 2014.
5. PE Marik. The demise of early goal-directed therapy for severe sepsis and septic shock. *Acta Anaesthesiologica Scandinavica*, 59(5):561–567, 2015.

6. Diana Kelm, Jared T. Perrin, Rodrigo Cartin-Ceba, Ognjen Gajic, Louis Schenck, and Cassie Kenned. Fluid overload in patients with severe sepsis and septic shock treated with early-goal directed therapy is associated with increased acute need for fluid-related medical interventions and hospital death. *Shock*, 43(1):68–73, 2015.
7. Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3, 2016.
8. Aniruddh Raghu, Matthieu Komorowski, Leo Anthony Celi, Peter Szolovits, and Marzyeh Ghassemi. Continuous state-space models for optimal sepsis treatment - a deep reinforcement learning approach. *CoRR*, abs/1705.08422, 2017.
9. M. Komorowski, L. A. Celi, O. Badawi, A. Faisal, and A. Gordon. The intensive care ai clinician learns optimal treatment strategies for sepsis. In *Nature Medicine*, 2018.
10. Devinder Thapa, In-Sung Jung, and Gi-Nam Wang. Agent based decision support system using reinforcement learning under emergency circumstances. In *International Conference on Natural Computation*, pages 888–892. Springer, 2005.
11. Pranav Rajpurkar, Vinaya Polamreddi, and Anusha Balakrishnan. Malaria likelihood prediction by effectively surveying households using deep reinforcement learning. *arXiv preprint arXiv:1711.09223*, 2017.
12. Sonali Parbhoo. *A reinforcement learning design for HIV clinical trials*. PhD thesis, 2014.
13. Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, and Barbara E Engelhardt. A reinforcement learning approach to weaning of mechanical ventilation in intensive care units. *arXiv preprint arXiv:1704.06300*, 2017.
14. Sonali Parbhoo, Jasmina Bogojeska, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Combining kernel and model based learning for hiv therapy selection. *AMIA Summits on Translational Science Proceedings*, 2017:239, 2017.
15. Andrew Ross. Transparent and interpretable machine learning in safety critical environments. NIPS, 2017.
16. Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.
17. Geoffrey Norman. Research in clinical reasoning: past history and current trends. *Medical education*, 39(4):418–427, 2005.
18. Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Van Hasselt, Marc Lanctot, and Nando De Freitas. Dueling network architectures for deep reinforcement learning. *arXiv preprint arXiv:1511.06581*, 2015.
19. Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.
20. Alan E Jones, Stephen Trzeciak, and Jeffrey A Kline. The sequential organ failure assessment score for predicting outcome in patients with severe sepsis and evidence of hypoperfusion at the time of emergency department presentation. *Critical care medicine*, 37(5):1649, 2009.
21. Kevin Beier, Sabitha Eppanapally, Heidi S Bazick, Domingo Chang, Karthik Mahadevappa, Fiona K Gibbons, and Kenneth B Christopher. Elevation of bun is predictive of long-term mortality in critically ill patients independent of ‘normal’ creatinine. *Critical care medicine*, 39(2):305, 2011.
22. F Tamion. Albumin in sepsis. In *Annales francaises d’anesthesie et de reanimation*, volume 29, pages 629–634, 2010.

23. Andrew Ng and Sparse Autoencoder. Cs294a lecture notes. *Dosegljivo: https://web.stanford.edu/class/cs294a/sparseAutoencoder_2011new.pdf.[Dostopano 20. 7. 2016]*, 2011.
24. Omer Gottesman, Fredrik D. Johansson, Joshua Meier, Jack Dent, Donghun Lee, Srivatsan Srinivasan, Linying Zhang, Yi Ding, David Wihl, Xuefeng Peng, Jiayu Yao, Isaac Lage, Christopher Mosch, Li-wei H. Lehman, Matthieu Komorowski, Aldo Faisal, Leo Anthony Celi, David Sontag, and Finale Doshi-Velez. Evaluating reinforcement learning algorithms in observational health settings. *CoRR*, abs/1805.12298, 2018.

Appendices

Due to space constraints, appendices can be found at

https://dtak.github.io/sepsis-DRL-mixture-of-experts/sepsis_moe_appendix.pdf