

# A Deep Bayesian Bandits Approach for Anticancer Drug Screening : Exploration via Functional Prior

MingYu Lu  
Yifang Chen  
Su-In Lee

*Department of Computer Science  
University of Washington  
Seattle, WA 98195-4322, USA*

MINGYULU@CS.WASHINGTON.EDU  
YFCHEN@CS.WASHINGTON.EDU  
SUINLEE@CS.WASHINGTON.EDU

## Abstract

Identifying potential cancer treatments with machine learning holds great promise to precision oncology. However, anticancer drugs screening is a lengthy and cost-intensive process. To address this challenge, we formulate drug screening studies as a *contextual bandit* problem, in which an algorithm selects anticancer therapeutics based on contextual information about cancer cell lines while adapting its treatment strategy to maximize treatment response. We propose using a novel *deep Bayesian bandits* framework that uses functional prior to approximate posterior for drug response prediction based on genomic features and drug structure and leverage inherent information between actions. We empirically evaluate our method on three large-scale *in-vitro* pharmacogenomic datasets and show that our approach outperforms several benchmarks in identifying the optimal treatment for a given cell line.

**Keywords:** Contextual Bandits, Deep Bayesian Bandits, Drug Discovery, Drug Screening

## 1. Introduction

By using genomic signatures and identifying druggable targets, the landscape of precision oncology has progressed rapidly.(Yang et al., 2012; Jordi et al., 2012; Iorio et al., 2016) One approach is to conduct anticancer therapeutics screening but this approach remains challenging as drug screening study is often an expensive, time-consuming process. Moreover, most of the studies rely on exhaustively experimenting all available drugs without considering cancer heterogeneity. To solve *in-vitro* identification and treatment allocation in anticancer therapeutics screening more efficiently, a natural attempt is to cast it as a *contextual bandit* problem, a well-known tool for personalization (Li et al., 2010). Contextual bandits framework enables effective personalized intervention by taking appropriate actions based on sample’s characteristics instead of random or uniform treatment assignment. Recently, contextual bandits methods, especially Thompson Sampling (Thompson, 1933), are being tailored to healthcare applications in a different domains, including mobile health (Tomkins et al., 2020) and risk stratification for multiple myeloma.(Zhou et al., 2019)

However, naively applying conventional bandit algorithms to drug screen studies is not feasible. First of all, the performance relies on accurate assumptions about the reward environment, which can be highly complex and non-linear. Second, finding a proper set

of candidate models to map genomic features to drug sensitivity may also be difficult. Moreover, due to inherent heterogeneity across and within different cancer types and drug metabolism (Dagogo-Jack and Shaw., 2018), modeling uncertainty plays an important role in determining the right decision. Third, the exploration-exploitation dilemma adds another layer of complexity, meaning that more action-context pairs would be required to obtain a better approximation of different policies. Insufficient sample size or model of less expressive power would lead to a wrong decision. (Zhou et al., 2019) Finally, most conventional contextual bandit algorithms treat each action as an unrelated discrete point and fail to leverage the similarity and inherent structure between each action. For instance, in cancer trial, OTX015, a BET inhibitor (BETi) that exhibits anti-tumor activity in non-small cell and small cell lung cancer, shares a similar structure with JQ1 (Riveiro et al., 2016). Under a conventional bandit setting, without knowing the inherent relationship between two similar drugs, unnecessary exploration would be conducted.

In this work, our contributions are as follows. First, we adapt a *deep Bayesian bandits* approach in drug screening studies as it provides a natural solution to uncertainty estimation with high-dimensional data. Second, under this settings, to find an approach to approximate posterior distribution of drug response prediction function, we use ‘functional prior’ (Hafner et al., 2019; Sun et al., 2019; Tran et al., 2020) and, following Sun et al. (2019), update functional posterior with variational inference through the Stein gradient estimator (SSGE)(Liu and Wang, 2019). Third, to avoid unnecessary exploration, for functional posterior update, to select proper measurement data points, we use a bootstrap procedure to sample from both perturbed genomics contexts and drug compound features. We empirically evaluated our method and compared it with other treatment allocation strategies using three large-scale drug sensitivity screens. Our method delivers the best performance against several benchmark algorithms.

## 2. Background and Related Work

### 2.1 Contextual Bandits

Contextual bandits (Li et al., 2010) is a generalization of the multi-armed bandit in which the policy for choosing future actions is dependent on *context* information at each step  $t$ . For example, in clinical trial, the context could be patients’ symptoms or laboratory observations. It is characterized by a sequence of a context-reward pairs  $\{x_t, r_t\}_{t=1,2,\dots,T} \in \mathcal{X} \times \mathbb{R}^{|\mathcal{A}|}$ , where  $\mathcal{X}, \mathcal{A}$  are some arbitrary context space and action space (e.g., a set of treatment choices), respectively. Given a candidate policy set  $\Pi$ , the contextual bandits problem aims to learn a policy  $\pi : \mathcal{X} \rightarrow \mathcal{A}$ , or a mapping from context to actions that maximizes cumulative rewards, or equivalently to minimize the total regret compared to the best policy in hindsight.

### 2.2 Thompson Sampling & Deep Bayesian Bandits

Thompson Sampling is a Bayesian approach (Thompson, 1933; Agrawal and Goyal, 2011) that allows sampling from the posterior distribution,  $P(\theta)$ , over plausible problem instances such as rewards or model parameters. At each round  $t$ , it observes the context,  $x_t$ , and an arm  $a$  is chosen according to the probability that it maximizes the expected reward. The

posterior distribution is then updated after the result of an action is observed. Through the lens of Thompson Sampling, deep Bayesian bandits refer to tackling contextual bandits by parameterizing the action-value function and a prior distribution over  $P_0(\theta)$  placed over the parameters. Although the weight posterior for neural networks is analytically intractable, it can be approximated by methods such as Dropout (Kingma et al., 2015; Srivastava et al., 2014) or Monte Carlo methods (Betancourt, 2017). Riquelme et al. (2018) conduct a comprehensive analysis of different approximation methods in deep Bayesian bandit settings.

### 3. Methods

#### 3.1 Problem Setup

We formulate treatment allocation in anti-cancer therapeutics as a Bayesian contextual bandit problem. That is, we have a context set  $\mathcal{X} \subset \mathbb{R}^{d_1}$ , defined as the individual genomics information with  $d_1$ -dimension; a discrete action set  $\mathcal{A} \subset \mathbb{R}^{d_2}$ , defined as a set of given drugs whose compounds can be expressed as a  $d_2$ -dimension feature vector; and a reward set  $R \subset \mathbb{R}^{|\mathcal{A}|}$ , where each reward vector is a possible collection of corresponding drug responses (pIC-50/Activity Area) for all drugs. Finally, we are given a policy set  $\Pi : \mathcal{X} \rightarrow \mathcal{A}$ , which is deep Bayesian neural net here. The protocol is as follows:

- At time  $t$ , we are given a context,  $x_t \in \mathcal{X}$ , and a *non-observable* reward vector  $r_t \in R$ , that are jointly drawn from some unknown *i.i.d* distribution  $\mathcal{D}$ . In our case, this scenario can be regarded as randomly sampled cancer cell line at time  $t$  whose actual responses to various drugs are unknown.
- Then we play  $a_t \in \mathcal{A}$  based on previous history and observe *only* the corresponding  $r_t(a_t)$ , for example, the drug response in our case, not in the whole  $r_t$  vector.

By repeating this process for a total of  $T$  times (or  $T$  incoming samples), we can obtain the total cumulative reward  $\sum_{t=1}^T r_t(a_t)$ . We measure the difference between our algorithm performance and the benchmark performance by the notion of *cumulative regret*  $R(T)$ , defined as

$$R(T) = \max_{\pi^*} \sum_{t=1}^T r_t(\pi^*(x_t)) - \sum_{t=1}^T r_t(a_t). \quad (1)$$

#### 3.2 Drug Response Prediction with Genomics and Drug Contexts

In precision oncology, treatment efficacy depends not only on patients' genomics features but also on therapeutics information which have shown to improve the predictive performance. (Kuenzi et al., 2020; Preuer et al., 2017) Here, our solution is to treat both genomics features *and* therapeutics context as a combined input into our neural network and show that the trained neural net can automatically capture inherent relationships between the two, genomics information (RNA expression) and drug features (MorganFingePrint); therefore, explore more efficiently. Specifically, on each round  $t$  for each action  $a$ , we denote such concatenated feature vector as  $x_t = (x_g, x_d) \in R^{d_1+d_2}$ , where  $d_1$  and  $d_2$  are the dimension of genomics and drug features, respectively. When the agent selects an action  $a_d$  based on  $x_g$ , corresponding drug information  $x_d$  will be provided for reward estimation.

### 3.3 Functional Prior and Posterior Approximation for multi-modal Data

Even though drug features can provide additional information in reward prediction, multimodal data of different data types complicates determining the prior distribution. The weight distribution would be highly variable and difficult to express with some commonly used prior distributions, such as the Gaussian distribution. (Hafner et al., 2019; Tran et al., 2020) Therefore, we are not aiming to learn a posterior distribution of weights in weight space but a distribution of *functional* posterior (Hafner et al., 2019; Sun et al., 2019; Tran et al., 2020) which is a function space that maps pharmacogenomics data to drug response. Sun et al. (2019) introduce function space variance, analogous to weight space variational inference.

$$f(x) = g_\phi(x, \xi) \quad (2)$$

Here  $f(x)$  is a sampled function for some function  $g_\phi$ , a neural network with stochastic weights and stochastic inputs, and  $\xi$ , a random noise vector, corresponding to randomness in *functional* space. For any given dataset  $D$ , the functional variance inference maximizes functional ELBO (fELBO) defined as:

$$\mathcal{L}(q) = \mathbb{E}_q[\log(p(D|f))] - \sup_{n \in N, x \in \mathcal{X}^n} KL[q(f^X)||p(f^X)]. \quad (3)$$

As there is no analytical form for the functional KL divergence, following Sun et al. (2019), we also utilize the spectral Stein gradient estimator (SSGE) to approximate log density derivatives. SSGE is proposed by Liu and Wang (2019) is an estimator that only requires measurement samples from the targeted implicit distribution i.e., distributions without tractable densities, functional space for drug response prediction in our case.

### 3.4 Sampling Measurement Sets for Cancer Treatment Exploration

To estimate KL divergence properly with SSGE, we need unbiased samples from the implicit distribution of genomic contexts and drug features. Typically, one can sample from the training set where noise is injected into the data or from region of testing data. (Hafner et al., 2019). Here we use a combination of two approaches (Tran et al., 2020): to encourage exploration, we use sample not only from history  $D_t$  but also from random drug feature  $x_{a'}$  in the action space that we are interested in exploring. To address this problem, we use a simple bootstrapped method (Osband and Roy, 2015) to sample measurement sets. Accordingly, we sample a new history set  $H$  from both training history  $D_t$  and  $\tilde{D}_t$  with perturbed action-context-concatenated features  $\tilde{x}_t$ .

$$\tilde{D}_t = (\tilde{x}_1, \dots, \tilde{x}_t); \tilde{x}_t = (x_g, x_{\tilde{a}}); H_t \sim D_t \cup \tilde{D}_t. \quad (4)$$

Therefore, with  $H$  as our measurement set, fELBO can be re-written as

$$\mathcal{L}_H(q) = \mathbb{E}_q[\log(p(D|f))] - \sup_{n \in N, x \in \mathcal{X}^n} KL [q(f^H)||p(f^H)]. \quad (5)$$

With SSGE and proper measurement samples, we can approximate the entropy of both  $q$  and  $p$ , which are both from implicit distribution; then, we can sample from the implicit posterior of drug response prediction for different compounds in bandits settings.

---

**Algorithm 1** Functional Posterior Update

---

- 1: **Universal parameters** :  $KL$  weight  $\lambda$
  - 2: **Input**:  $D_t$ , functional variational posterior  $g(\cdot|\phi)$
  - 3:  $\tilde{D}_t \sim D_t$ ; ▷ Sample  $\tilde{D}_t$  from  $D$  with perturbed genomics and action contexts.
  - 4:  $f = g([\tilde{D}_t, D_t], \mathcal{E}|\phi)$  ▷ Sample a new function  $f(\cdot)$  with noise vector  $\mathcal{E}$ .
  - 5:  $\Delta_1 = \frac{1}{|D_t|} \nabla \phi \log_p(y|f_i)$  ▷ Compute reconstruction loss.
  - 6:  $\Delta_2 = SSGE(\tilde{D}_t, D_t)$  ▷ Estimate KL divergence from implicit posterior.
  - 7:  $\phi \leftarrow \text{optimize}(\phi, \Delta_1 - \lambda \Delta_2)$
  - 8: **Return** functional variational posterior  $g(\cdot|\phi)$
- 

## 4. Experimental Design and Data

### 4.1 Simulated Drug Screen Design

In this work, we simulate anticancer treatment trials and evaluate our algorithms in three large-scale cancer cell line screenings, including Genomics in Drug Sensitivity in Cancer (GDSC-1 and GDSC-2) (Yang et al., 2012; Iorio et al., 2016) and the Cancer Cell Line Encyclopedia (CCLE) (Jordi et al., 2012). For cancer cell lines, we used RNA expressions and performed principal component analysis (PCA) to reduce the RNA expression dimension to 500. For drug context, we converted the drug SMILE string to Morgan fingerprint with 256 bits. (Capecchi et al., 2020) For drug-dose response, in GDSC, we adapted negative log of half-maximal inhibitory concentration,  $pIC50$ . In CCLE, since some IC50 measurements are missing, we used activity area (AA) instead. All cancer cell lines and drug responses were scaled into the range of  $[0,1]$ . We run the bandits/simulated trials for 5000 rounds.

### 4.2 Evaluation Metrics and Benchmark Algorithms

As described in Equation 1, we reported cumulative regrets, the reward difference between competing algorithm,  $\pi$ , and an oracle,  $\pi^*$ , that would have access to the (unknown) treatment effect functions. We test our algorithm against competing algorithms (**Appendix A.4**) of different posterior approximations that demonstrated promising results. (Riquelme et al., 2018)

## 5. Results

### 5.1 Effect of Different Measurement Sets on the fBNN posterior

Here we show the effectiveness of using drug compound features as input contexts and adding action perturbation. From Figure 1, we show that our strategy gives much lower cumulative regret compared to ones without them. To demonstrate that context-action pair perturbation as proxy of functional prior improves uncertainty estimation for drug response prediction, we compared the performance of models updated with different samples for SSGE, including (i) perturbing only genomics context,  $x'_g = x_g + \epsilon; \epsilon \sim \mathcal{N}(0, 1)$  (ii) perturbing only action (drug) context  $x'_d \in R^{d_2}$  and (iii) random sampling from multivariate Gaussian distribution,  $x'_g \in R^{d_1}; x_g \sim \mathcal{N}(0, 1)$  and random selected drugs from the action set,  $x'_d \sim U(R^{d_2})$ . Our sampling method, compared to the other proxy samples as in Figure 1, improves the quality of approximation of the desired prior and achieves lower regret.

## 5.2 Molecular Therapeutics Recommendation in Cancer Treatment

Now we empirically evaluate our functional Bayesian neural bandit with benchmark algorithms in deep bayesian bandits in selecting drug compounds that would maximize drug response. Selecting optimal treatment requires not only learning drug response function but also uncertainty estimation for exploration. As shown in Table 1, models that can only optimize both, for example, NeuralGreedy and BBB performs poorly. (Riquelme et al., 2018) Bootstrap and ParameterNoise seem to provide better posterior approximations for drug response prediction but suffer from under exploration, sometimes committing to the suboptimal drug in the experiments (Riquelme et al., 2018; Sun et al., 2019). Our method, on the other hand, by using genomics-action pair as proxy for functional prior, provides a better exploration strategy.

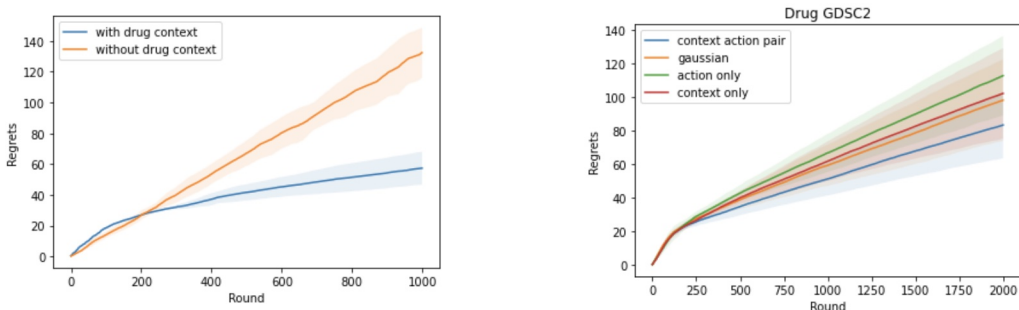


Figure 1: Cumulative regrets between (a) with and without drug context in the GDSC dataset, and (b) Functional Bayesian Neural Network with different measurement set samples. The confidential interval is generated through 20 trials.

Table 1: Cumulative regret distribution at the final step in 3 drug screen studies. We report the mean and standard error of the mean over 20 trials.

Algorithm	GDSC-1	GDSC-2	CLE
Uniform	2011 ± 10.2	995 ± 6.7	2509 ± 12.3
NeuralGreedy	531.4 ± 194.3	262.2 ± 152	1361.4 ± 361.7
BayesByBackprop(BBB)	1500.2 ± 396.4	497.3 ± 187.3	1682.3 ± 390.0
DropOut	508.3 ± 33.4	378.2 ± 84.0	1176.2 ± 342.1
BootstrappedNN	429.7 ± 162.2	179.1 ± 70.5	705.5 ± 409.8
ParameterNoise	467.5 ± 287.5	170.6 ± 65.4	785.51 ± 585.5
<b>FunctionalPosterior</b>	<b>202.4 ± 70.0</b>	<b>98.46 ± 30.3</b>	<b>252.8 ± 35.5</b>

## 6. Conclusion

In this work, we introduce a new deep Bayesian bandit approach for *in-vitro* anticancer therapeutics selection. To enable efficient exploration in a contextual bandit setting with drug compound features and gene expression, we use functional variational approach to approximate this pharmacogenomics posterior. Empirically, we demonstrate that our approach explores efficiently and performs the best. The result is consistent across three publicly available pharmacogenomics datasets in cancer treatment recommendations. We hope that our work will serve as a stepstone towards improving the efficacy and quality of treatment allocation strategy in both pre-clinical therapeutic discovery and clinical study.

## References

- Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem, 2011. URL <https://arxiv.org/abs/1111.1797>.
- Michael Betancourt. A conceptual introduction to hamiltonian monte carlo, 2017. URL <https://arxiv.org/abs/1701.02434>.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks, 2015.
- A. Capecchi, D. Probst, and JL. Reymond. One molecular fingerprint to rule them all: drugs, biomolecules, and the metabolome. *Journal of Cheminformatics*, 2020. doi: <https://doi.org/10.1186/s13321-020-00445-4>.
- I Dago-Jack and AT Shaw. Tumour heterogeneity and resistance to cancer therapies. *Nat Rev Clin Oncol.*, 2018. doi: 10.1038/nrclinonc.2017.166.
- Danijar Hafner, Dustin Tran, Timothy Lillicrap, Alex Irpan, and James Davidson. Noise contrastive priors for functional uncertainty, 2019.
- Francesco Iorio, Theo Knijnenburg, Daniel Vis, Graham R Bignell, Michael Menden, Michael Schubert, Nanne Aben, Emanuel Gonçalves, Syd Barthorpe, Howard Lightfoot, Thomas Cokelaer, Patricia Greninger, Ewald van Dyk, Han Chang, Heshani de Silva, Holger Heyn, Xianming Deng, Regina K Egan, Qingsong Liu, and Mathew Garnett. A landscape of pharmacogenomic interactions in cancer. *Cell*, 166, 07 2016. doi: 10.1016/j.cell.2016.06.017.
- Barretina Jordi, Caponigro G, and Stransky N. The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. *Nature*, 483, 3 2012.
- Diederik P. Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick, 2015. URL <https://arxiv.org/abs/1506.02557>.
- BM Kuenzi, J Park, and et al. Fong SH. Predicting drug response and synergy using a deep learning model of human cancer cells. *Cancer Cell*, 2020. doi: doi:10.1016/j.ccell.2020.09.014.
- Branislav Kveton, Manzil Zaheer, Csaba Szepesvari, Lihong Li, Mohammad Ghavamzadeh, and Craig Boutilier. Randomized exploration in generalized linear bandits, 2019. URL <https://arxiv.org/abs/1906.08947>.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm, 2019.
- Ian Osband and Benjamin Van Roy. Bootstrapped thompson sampling and deep exploration, 2015.

- Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y. Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration, 2018.
- Kristina Preuer, Richard P I Lewis, Sepp Hochreiter, Andreas Bender, Krishna C Bulusu, and Günter Klambauer. Deepsynergy: predicting anti-cancer drug synergy with deep learning. *Bioinformatics*, 34(9):1538–1546, 12 2017. ISSN 1367-4803. doi: 10.1093/bioinformatics/btx806. URL <https://doi.org/10.1093/bioinformatics/btx806>.
- Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling, 2018.
- ME Riveiro, Astorgues-Xerri L, Vazquez R, Frapolli R, Kwee I, Rinaldi A, Odore E, Rezai K, Bekradda M, Inghirami G, D’Incalci M, Noel K, Cvitkovic E, Raymond E, and Bertoni F. Otx015 (mk-8628), a novel bet inhibitor, exhibits antitumor activity in non-small cell and small cell lung cancer models harboring different oncogenic mutations. *Oncotarget*, 2016. doi: 10.18632/oncotarget.13181.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Shengyang Sun, Guodong Zhang, Jiaxin Shi, and Roger Grosse. Functional variational bayesian neural networks, 2019.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25, 1933.
- Sabina Tomkins, Peng Liao, Predrag Klasnja, and Susan Murphy. Intelligentpooling: Practical thompson sampling for mhealth, 2020. URL <https://arxiv.org/abs/2008.01571>.
- Ba-Hien Tran, Simone Rossi, Dimitrios Milios, and Maurizio Filippone. All you need is a good functional prior for, 2020.
- Wanjuan Yang, Jorge Soares, Patricia Greninger, Elena J. Edelman, Howard Lightfoot, Simon Forbes, Nidhi Bindal, Dave Beare, James A. Smith, I. Richard Thompson, Sridhar Ramaswamy, P. Andrew Futreal, Daniel A. Haber, Michael R. Stratton, Cyril Benes, Ultan McDermott, and Mathew J. Garnett. Genomics of Drug Sensitivity in Cancer (GDSC): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic Acids Research*, 41(D1):D955–D961, 11 2012. ISSN 0305-1048. doi: 10.1093/nar/gks1111. URL <https://doi.org/10.1093/nar/gks1111>.
- Weitong Zhang, Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural thompson sampling, 2020. URL <https://arxiv.org/abs/2010.00827>.
- Zhijin Zhou, Yingfei Wang, Hamed Mamani, and David Coffey. How do tumor cytogenetics inform cancer treatments? dynamic risk stratification and precision medicine using multi-armed bandits. *SSRN Electronic Journal*, 01 2019. doi: 10.2139/ssrn.3405082.



## Appendix A.

### A.1 Contextual bandits in Healthcare

An instantiation of the contextual bandits framework in treatment recommendations.

---

**Algorithm 2** Contextual Bandits in Healthcare

---

- 1: Available treatment options  $a \in A$ , and the treatment period length  $T$ .
  - 2: **for**  $t = 1, 2, \dots, T$  **do**
  - 3:     Agent receives the context information  $x_t$  for the current patient.
  - 4:     A treatment  $a_t$  is determined based on the history and the context information.
  - 5:     Record the post-treatment health outcome  $r_t$  for the patient.
  - 6: **end for**
- 

### A.2 Thompson Sampling & Deep Bayesian Bandits

Deep Bayesian Bandits refers to posterior approximation with a deep neural network. For example, in Algorithm 3, can be neural neural networks parameters drawn from a weight distribution,  $P(\theta)$ . Riquelme et al. (2018) conduct a comprehensive analysis of different approximation methods in deep Bayesian bandit settings and propose *NeuralLinear*, a Bayesian linear regression built on the representation of the last layer of a neural network. Blundell et al. (2015) introduce *Bayes by Backprop*, which learns the probability distribution of weights. Kveton et al. (2019) propose *DeepFPL*, a neural network that uses ensemble sampling as an approximation to generating new noise in each round. In addition to posterior approximation of model weights, Zhang et al. (2020) propose Neural Thompson Sampling, which focuses on the posterior distribution of reward. Sun et al. (2019) introduce the functional prior, sampling from the functional posterior by concatenating a random noise vector into input.

---

**Algorithm 3** Thompson Sampling

---

- 1:  $f(\cdot | \hat{\theta}_0)$ ; prior distribution over models,  $P_0(\theta) : \hat{\theta}_0 \in \Theta \rightarrow [0, 1]$
  - 2: **for**  $t = 1, 2, \dots, T$  **do**
  - 3:     Observe context  $x_t \in R^d$ .
  - 4:     Sample model parameter  $\hat{\theta}_t \sim P_t(\theta)$ .
  - 5:      $a_t = \arg \max_a f(x, a | \hat{\theta}_t)$ .
  - 6:     Select action  $a_t$  and observe reward  $r_t(a_t)$ .
  - 7:     Update the history  $D_t \leftarrow D_{t-1} \cup (x_t, a_t, r_t(a_t))$ .
  - 8:     Update the posterior distribution  $P_{t+1}(\theta) = \mathbf{PosteriorUpdate}(D_t, P_t)$ .
  - 9: **end for**
- 

### A.3 Dataset

GDSC is the resource for therapeutic biomarker discovery in cancer cells. The GDSC1 dataset was generated between 2009 and 2015 using a matched set of cancer cell lines;

contains 987 cell lines and 367 compounds. GDSC2 was generated in 2015 following improvements to screen design and assay (Iorio et al., 2016); contains 198 Compounds and 809 cell lines. The Cancer Cell Line Encyclopedia (CCLE) project, started in 2008, contains pharmacological profiles for 24 anticancer drugs across 479 cell lines. For all datasets, we first filtered out cancer cell lines with missingness in drug-dose response greater than 70%. We then filtered out drug compounds with missing sensitivity in at least one of the curated cell lines.

Dataset	number of cell lines	number of drugs	Drug Response
GDSC1	472	133	IC-50
GDSC2	602	116	IC-50
CCLE	411	21	Activity Area

Table 2: Summary Statistics of drug response dataset used in the work

## A.4 Baselines & Benchmark Algorithms

**Uniform** a baseline for which an agent takes each action at random with equal probability.

**NeuralGreedy** is a algorithm trained with a neural network and acts greedily (take action that predicted with the highest drug response.) where a random action is selected with probability  $\epsilon$ .

**BayesByBackprop (BBB)** is proposed by Blundell et al. (2015) and is one of the commonly used variational approaches that approximate posterior by minimizing KL divergence. It is a Bayesian neural network with sampled weights from the variational posterior,  $w \sim q(w|\theta)$ . Variational parameters  $\theta$  are updated with KL divergence between a mixture of two Gaussian densities prior and model weights distribution. BBB is a commonly used baseline in deep Bayesian bandits.

**Dropout** is a training technique where the output of each neuron is independently zeroed out with probability  $p$  at each forward pass (Srivastava et al., 2014; Kingma et al., 2015). Following the best action with respect to the random dropout prediction can be interpreted as an implicit form of Thompson sampling.

**BootstrappedNN** is an empirical approach to approximate the posterior sampling distribution (Osband and Roy, 2015). The idea is to train  $q$  models with different dataset  $D_i$  where  $D_i$  is typically created by sampling with replacement from an original dataset,  $D$ . In our case, we train  $q$  models with the same structure as *NeuralGreedy* and  $D$  is the history context that the agent has seen. In action selection, we sample a model uniformly at random (with probability  $1/q$ ) and take action predicted reward to be the best by the sampled model.

**DirectNoiseInjection** is a recently proposed method by Plappert et al. (2018). Parameter-noise injection is an exploration technique proved successful in RL. Through noise injection in parameter space:  $\tilde{\theta} = \theta + \mathcal{N}(0, \sigma^2)$ , we can perform state-dependent exploration. When selecting action, model weights are perturbed with isotropic Gaussian noise so that we get

$a_t = \pi_{\hat{\theta}}(s_t)$ . The magnitude of noise is adaptively adjusted based on distance measure between perturbed and non-perturbed policy in action space.(Plappert et al., 2018)

**Network Structure and Hyper-Parameter Tuning** For our oracle policy, we trained a 2-layer fully connected feedforward network with full access to genomics features, drug structure and corresponding drug response in the dataset. For each experiment, we ran 20 trials with 5,000 steps in each trials. All algorithms were updated every 30 steps. The size of each batch is 32. For all neural network structures, we adopted the same number of layers and hidden dimensions. as our oracle policy. We performed a grid search over parameters within a plausible range and reported the best results for each algorithm. We did a grid search over  $\{1e^{-1}, 1e^{-2}, 1e^{-3}, 1e^{-4}, 1e^{-5}\}$  for learning rate. For NeuralGreedy, we performed grid search for  $\epsilon$  over  $\{0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}$ . For DirectNoiseInjection, we searched noise range over  $\{1e^{-1}, 1e^{-2}, 1e^{-3}, 1e^{-4}, 1e^{-5}\}$ . For Dropout, we searched over probability  $\{0.2, 0.4, 0.6, 0.8\}$ . For BayesByBackprop, we searched prior noise over  $\{1, 1e^{-1}, 1e^{-2}\}$  and prior mean over  $\{1, 1e^{-1}, 1e^{-2}\}$ . For BootstrappedNN, we searched the number of models over  $\{2, 3, 5, 10\}$ .

## A.5 Extra Results




Figure 2: Comparison of (mean) cumulative regrets at each step in maximizing cancer treatment effect for 5,000 rounds over 20 trials in GDSC-1, GDSC-2, & CCLE. (left to right.)