The 35th Canadian Conference on Artificial Intelligence
DOI: 0
Canadian AI 2022
30 May–3 June 2022, Toronto

Contextual bandit optimization of super-resolution microscopy

Anthony Bilodeau[†, ‡], Renaud Bernatchez[†, ‡], Albert Michaud-Gagnon[†, ‡], Flavie Lavoie-Cardinal[†, ‡, ◇, *], Audrey Durand[‡, •, *]

[†] CERVO Brain Research Center, Université Laval [‡] IID, Université Laval
[◇] Psychiatry and Neuroscience, Université Laval [•] Canada CIFAR AI chair, Mila

Abstract

The online optimization of optical microscopy parameters aims at learning the set of imaging parameters while the experiment unfolds. The success of this task relies on the trade-off between multiple confounding objectives and may vary depending on the experimental setting. In this work, we frame the optimization problem as a multi-armed bandit framework with contextual information about the sample to identify optimal sample-dependant imaging parameters. This allows to take into consideration the current state of the sample and choose the imaging parameters accordingly.

Keywords: Online optimization, Multi-objective, Contextual-Bandit, Super-Resolution Microscopy

1. Introduction

Optical microscopy is a very powerful tool to visualize biological structures with sub-cellular resolution (∼250 nm). Unfortunately, smaller structures (nanostructures) such as synapses cannot be properly resolved by conventional optical microscopy techniques due to the diffraction limit [1]. In contrast, super-resolution optical microscopy can overcome the diffraction limit (<50 nm), but comes at the expense of more challenging experimental protocols [1]. Indeed, trade-offs between conflicting objectives (e.g. spatio-temporal resolution, signal-to-noise ratio, photobleaching, acquisition speed, sample damage) must be made by the microscopist during the imaging tasks. For instance, in STimulated Emission Depletion (STED) nanoscopy, a higher resolution image is reached by scanning two focused and co-aligned laser beams over the sample: an excitation Gaussian-shaped laser beam and a high intensity donut-shaped depletion beam (or STED beam) that turns off the emitters through stimulated emission [2]. For example, to detect a given nanostructure, the microscopists need to find the optimal scanning speed, collect a sufficient amount of fluorescence photons, while limiting the amount of light reaching the sample to reduce sample damage [1]. Hence, they must tune the imaging parameters (e.g. excitation/depletion laser power, pixel dwell time, pixel size) by trading off between the imaging objectives [3, 4], often by considering the experimental setting (e.g. type of structure, sample density, sample intensity).

Finding the set of optimal imaging parameters for a given sample is a non-trivial task. The optimal parameter set can vary from sample to sample and can even vary across days for a given type of sample due to biological variability [3]. Manual parameter tuning is thus required for each new experimental setting. When working with biological samples, the number of negative attempts has to be limited to maximize the information that can be extracted. This can be naturally formulated using the exploration-exploitation trade-off. The exploration stems from exploring the imaging parameter space in search of optimal parameters, while the exploitation aims to maximize the number of acquired high quality images from a given sample. Since imaging parameters need to be adapted to the sample, we tackle how we can leverage this information in the decision making process. We build on existing work [3] by formulating the optimization problem of sample-dependant parameters as a contextual bandit [5] problem with structured actions.

[*] flavie.lavoie-cardinal@cervo.ulaval.ca & audrey.durand@ift.ulaval.ca

2. Related Works

Tools have been developed to assist the microscopists in specific imaging task. For instance, DyMIN (Dynamic intensity Minimum) [6] dynamically adapts the imaging scheme of STED microscopy to reduce photobleaching and sample damage by switching on the depletion laser only when necessary, *i.e.* when a structure is detected. To exploit DyMIN to the fullest however, microscopists must manually tune the detection parameters to the experimental setting which can be difficult for non-initiated users and a limiting factor to a broad deployment of the technique.

The online automated optimization of optical microscopy parameters to simultaneously improve multiple objectives was previously formulated under a multi-armed bandit framework using Kernel-TS, i.e. posterior sampling from a Gaussian Process regression [3]. This approach considered one regression model per objective to induce a notion of structure onto the set of possible actions. Using this method the authors were able to optimize up to three imaging parameters in a multi-objectives setting. Unfortunately, this method requires a fine discretization of the input space, reducing its capability to scale to higher dimensional parameter spaces, such as images.

When sample-dependant parameters are considered, such as in the case of the detection parameters of DyMIN, the decision making process can be facilitated by incorporating prior information about the experimental conditions, *e.g.* the current state of the sample. Extending the multi-armed bandit problem with such contextual information results in the contextual bandit problem [5]. This framework has been extensively studied in the literature and applied to recommendation systems and medical treatment planning [7]. It is often assumed that there exists a finite, discrete number of possible actions/contexts from which the agent can choose at each time step. In super-resolution microscopy optimization, the space of possible actions is continuous for each imaging parameters which requires to consider contextual bandits with large action spaces [8]. We extend previous work [8] to simultaneously optimize multiple objectives under contextual bandits with large action spaces using Thompson Sampling (TS) based methods as previously applied in online optimization of optical microscopy parameters [3].

The optimization of STED imaging parameters was recently tackled using sequential reinforcement learning (RL) algorithms [9]. A proximal policy optimization algorithm was trained to automatically perform the parameter selection in given sequences. This translates in a RL agent that does not aim to maximize its performance during the training phase but rather at deployment which is akin to supervised learning.

3. Problem setting

In this work, we aim to optimize the parameters of DyMIN, which are highly dependant on the state of the current sample [6], as shown in Figure 1[1]. DyMIN uses intensity thresholds on the detected fluorescence photons to activate the STED beam only when necessary and limit the light dose. At every pixel, the fluorescence signal is measured; if the intensity is below a chosen threshold, both lasers are turned off until the next pixel. If the intensity is above the threshold, the STED laser power is increased to turn off the emitters under the depletion beam and increase the resolution of the image. Multiple thresholds, typically two [6], are applied until the final step at which the photons are accumulated for the entire pixel dwell time. For a given sample, increasing the detection thresholds reduces photobleaching but leads to possible acquisition artifacts (Figure 1c, dashed line) thus creating non-trivial trade-offs between the objectives. We implemented the DyMIN acquisition scheme in pySTED, a STED microscopy simulation tool designed to realistically emulate a STED acquisition [9].

[1]We thank Abberior Instruments for their help with the DyMIN implementation.

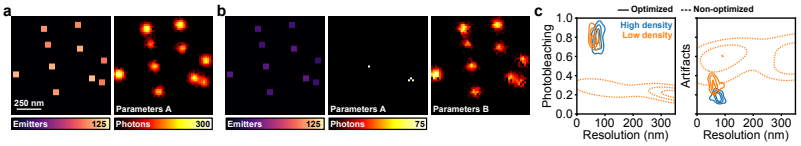

Figure 1. DyMIN parameters are sample dependant. a) Datamap for 900 emitters per point source (*left*) and its corresponding pySTED stimulated image using DyMIN parameters adapted to the number of emitters (*right*). b) Datamap for 225 emitters per point source (*left*), its corresponding pySTED simulated image using i) the same parameters as in **a** (*center*), and ii) DyMIN parameters optimized for this new sample (*right*). c) Kernel Density Estimation (KDE) over the imaging objectives of the simulated STED images in high (a, blue) and low (b, orange) emitter density with optimized (solid) and non-optimized (dashed) parameters. N=25 images per condition.

While the simulations are realistic, the parameters of the simulator could not be tuned to perfectly match a real sample thus requiring online optimization. pySTED uses as input a datamap which consists of a matrix indicating the positions of emitters in the field of view. A low/high fluorescence signal can be simulated by varying the number of emitters at each positions.

We frame the problem of selecting the DyMIN parameters based on the state of the current sample as a contextual bandit problem that goes as follows. At every time step t, a datamap of 10 point sources with a random number of emitters ($[90, 900]$ emitters per point source) positioned randomly within the field of view is generated, and a confocal image of the datamap is simulated using the pySTED simulator. We define the context s_t as i) the average foreground intensity in the confocal image (scalar value) or ii) the confocal image. The context is revealed to the agent which proposes combinations of imaging parameters $\mathcal{X}_t$ associated with their respective predicted objective values. Given that each parameter combination is associated with multiple objectives, a human expert selects the preferred set of parameters $x_t \in \mathcal{X}_t$ according to their expected objectives. A STED image is simulated from the datamap using the selected parameters. The objectives are calculated from this image and are used as reward r_t to update the agent. By sequentially learning from these trials, the agent should select parameters less prone to failure (bad imaging objectives) as time goes, thus leading to a sub-linear number of cumulative failures.

4. **Investigated approaches**

Similarly to previous work [3] involving a human expert for trading-off between the objectives predicted by a learning agent, we restrict ourselves to algorithms based on Thompson Sampling (TS). This is related to the fact that values sampled from the posterior of a model are more informative for a microscopist about the true underlying objective value than upper confidence bounds. We therefore investigate different variations of TS for contextual bandits with structured actions. First, we consider Kernel-TS [3] to map the joint parameters and context (scalar value) spaces to the expected objective values. Similarly, we consider LinTS [10] using 3^{rd} order polynomial features, a common contextual bandit algorithm which can be thought of a parametric equivalent of Kernel-TS. We also investigate approaches that rely on a neural network (NN) to extract their own relevant features from the parameters and context (scalar value or image). Our last candidates are thus LinTS Diag [11], which samples from the posterior using the gradients of the last linear layer of the network, and NeuralTS [11], which uses the gradient through the entire network.

At each time step, a function is sampled from the posterior of each regression models (one per objective) to generate the possible imaging parameter combinations. Using Kernel-TS, since it evaluates the objective functions over the discretized parameter space, all parameters

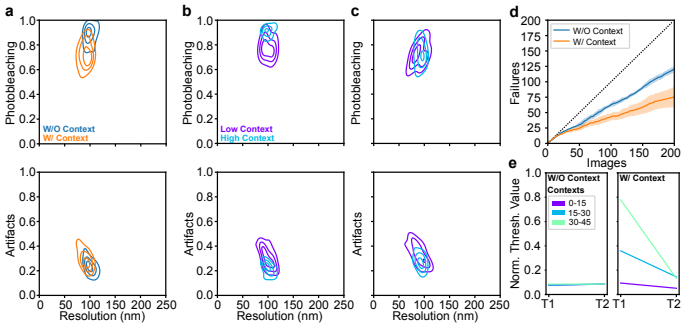

Figure 2. Contextual information improves optimization of imaging parameters. a) Typical KDE of the last 50 acquired images for Kernel-TS trained without (blue) and with (orange) contextual information. b, c) Typical KDE of the last 50 acquired images segregated in low and high context (mean foreground intensity, low: <15 and high: >15) for Kernel-TS trained without (b) and with (c) contextual information. d) Cumulative number of failures during the optimization for both Kernel-TS experiments (mean ± std, N = 3 repetitions). e) Average selected parameters (Threshold 1/2: T1/2) for each range of context for an optimization without (left) and with contextual information (right) (N = 3 repetitions).

associated with non-dominated objective values are considered for presentation to human experts. Using other methods, a Pareto front of parameter combinations is extracted using NSGA-II [12]. The experts (microscopists) must trade off between the presented possibilities and select one combination according to their preferences. In order to alleviate this task, we automated this process by training a NN on past expert objective preferences to perform the trade-off [3].

5. **Experiments**

We define two DyMIN thresholds as the parameters to be optimized[2]. We measure three imaging objectives which are indicative for the task of optimizing DyMIN microscopy: Spatial resolution, detection artifact using the Squirrel score [13], and Photobleaching. The resolution is calculated from the simulated image using a decorrelation analysis [14]. The Squirrel score aims at detecting missing structures in the super-resolved image that appear in the ground truth low-resolution image [13]. This can be the case in DyMIN microscopy when the thresholds are set too high since the final acquisition step is not reached. We approximate the photobleaching of the sample by comparing the foreground signal between low-resolution images acquired before and after the DyMIN acquisition [3]. All objectives should be minimized. When the context (mean foreground intensity) is greater than 15 (high-context), an expert considers an image as a success if the spatial resolution is below 120 nm, the artifact score is below 0.3, and photobleaching is below 0.8. Otherwise (mean foreground intensity below 15 defined as low-context), the image is a success if the spatial resolution is below 120 nm, the artifact score is below 0.5, and photobleaching is below 1.0.

We first evaluated the impact of leveraging the contextual information for Kernel-TS. As shown by the kernel density estimation (KDE) of the last 50 acquired images in Figure 2a, a blind Kernel-TS model converges to a region of the objective space that is higher for the photobleaching (≈0.9) compared to Kernel-TS (≈0.7), while the value stayed similar for the artifact score and resolution. The value of the imaging objectives segregated by their context was also compared (low: <15 and high: >15). As shown in Figure 2b,c, a blind Kernel-TS (Figure 2b) obtains higher photobleaching in high context images resulting from

[2] https://github.com/FLClab/pysted-contextual-bandit

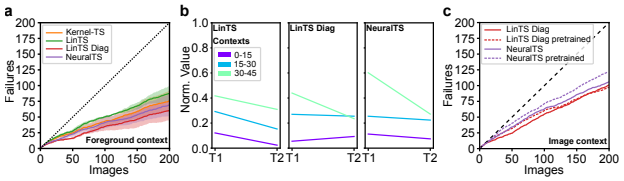

Figure 3. Contextual bandit baseline comparison. a) Cumulative number of failures during the optimization (mean ± std, N = 3 repetitions) using mean foreground intensity as contextual information. b) Average selected parameters (Threshold 1/2: T1/2) for each range of context during an optimization (N = 3 repetitions) for LinTS (left), LinTS Diag (center) and NeuralTS (right). c) Cumulative number of failures during the optimization using the confocal image as contextual information. (N = 3 repetitions, std not shown for vizualisation purposes).

an increased probability of reaching the final acquisition step of DyMIN. This is not the case for Kernel-TS having access to contextual information as depicted in Figure 2c which shows overlapping KDE. We calculated the number of cumulative failures (Figure 2d) which shows a sub-linear trend for both models as expected by a learning algorithm. Using the contextual information reduces the number of cumulative failures for Kernel-TS trained with context (Figure 2d). Figure 2e shows the average DyMIN threshold values that were selected during the second half of the optimization (>100 acquired images) in different range of contexts. As expected, a blind Kernel-TS yields DyMIN threshold values which are constant to accommodate over the range of contexts. By leveraging the contextual information Kernel-TS is able to infer context-specific imaging parameters that limits the number of failures (Figure 2d,e). In low context, the algorithm reduces the thresholds to accommodate the low number of photons, while in high context, the thresholds are increased to reduce the light exposure to the sample (Figure 2e).

We next aimed to compare Kernel-TS with baselines on the imaging task using the mean foreground intensity as contextual information. As shown by the cumulative number of failures in Figure 3a, LinTS with a 3rd order polynomial performed worse than Kernel-TS. Kernel-TS is only slightly outperformed by LinTS Diag and NeuralTS which extract their own features. Both NN models achieved a reduced number of failures (∼50 images, Figure 3a) and a similar selection of parameters (Figure 3b).

One advantage of using NN models is their capacity to automatically extract features to perform a specific task, even from a high dimensional parameter space. Hence, we crafted a last experiment in which the model had access to the confocal image as contextual information. It is important to note that this particular task is more complex since the model has to simultaneously learn i) meaningful features to extract from the image, and ii) the imaging parameters while compromising between exploration and exploitation. Nevertheless, the NN models were able to limit the number of failures but to a lesser extent than when trained on a simpler context (Figure 3c). Since extracting meaningful information from the images is not trivial, we wondered whether pretraining the vision layers using an auto-encoder-like model would improve the performance of the model. Interestingly, the performance of the model was similar or worse indicating that pretraining the vision layer, at least for this specific task, does not help the model infer the objective function faster (Figure 3c).

6. **Discussion**

We have shown how contextual information can help in the online optimization of super-resolution microscopy. The optimization of the thresholds of DyMIN were tackled as a use case since they should vary as a function of available knowledge, *i.e.* the intensity of the sample. Using the contextual information, the models could modify their parameters

dependant on the experimental conditions and outperform their counterpart blindly trained. Our results are in line with the literature in the sense that the methods which extract their own representation from the input features perform better than the methods relying on engineered features. We believe that using contextual information can be leveraged more generally and broadly for the optimization of imaging tasks. Indeed, contextual information provides more information about the sample and helps sharing features across optimization runs thereby facilitating the decision making process and possibly improving the convergence of real world imaging tasks.

We have limited our experiments to 2 imaging parameters and the contextual information (scalar value or complete image) but many real-world experiments may require more parameters. While Kernel-TS is limited by the discretization of the parameter space ($<$4 imaging parameters), the scalability of parametric and NN approaches will need to be tested for larger imaging parameter space. A benefit of NN implementations (LinTS Diag and NeuralTS) is that they approximate the posterior using the gradients calculated from the model. These methods would allow to learn all objectives in the same architecture. Hence, the model could learn a more efficient mapping between the imaging parameters/context and the objectives.

In this work, experiments were conducted on simulated STED imaging data. The methods described should later be translated to real STED acquisitions.

References

[1] L. Schermelleh, A. Ferrand, T. Huser, C. Eggeling, M. Sauer, O. Biehlmaier, and G. P. C. Drummen. "Super-Resolution Microscopy Demystified". In: *Nat. Cell Biol.* 21.1 (2019), pp. 72–84.

[2] S. W. Hell and J. Wichmann. "Breaking the Diffraction Resolution Limit by Stimulated Emission: Stimulated-Emission-Depletion Fluorescence Microscopy". In: *Opt. Lett.* 19.11 (1994), pp. 780–782.

[3] A. Durand, T. Wiesner, M.-A. Gardner, L.-É. Robitaille, A. Bilodeau, C. Gagné, P. D. Koninck, and F. Lavoie-Cardinal. "A Machine Learning Approach for Online Automated Optimization of Super-Resolution Optical Microscopy". In: *Nat. Commun.* 9.1 (2018), pp. 1–16.

[4] W. C. Lemon and K. McDole. "Live-Cell Imaging in the Era of Too Many Microscopes". In: *Curr. Opin. Cell Biol.* 66 (2020), pp. 34–42.

[5] P. Auer. "Using Confidence Bounds for Exploitation-Exploration Trade-Offs". In: *JMLR* 3.Nov (2002), pp. 397–442.

[6] J. Heine, M. Reuss, B. Harke, E. D'Este, S. J. Sahl, and S. W. Hell. "Adaptive-Illumination STED Nanoscopy". In: *P. Natl. Acad. Sci.* 114.37 (2017), pp. 9797–9802.

[7] A. Bietti, A. Agarwal, and J. Langford. "A Contextual Bandit Bake-Off". In: *arXiv:1802.04064* (2020).

[8] A. Krause and C. Ong. "Contextual Gaussian Process Bandit Optimization". In: *Advances in Neural Information Processing Systems*. Vol. 24. Curran Associates, Inc., 2011.

[9] B. Turcotte, A. Bilodeau, F. Lavoie-Cardinal, and A. Durand. "pySTED : A STED Microscopy Simulation Tool for Machine Learning Training". In: *AAAI2022* (2022).

[10] S. Agrawal and N. Goyal. "Thompson Sampling for Contextual Bandits with Linear Payoffs". In: *ICML'13*. Vol. 28. PMLR, 2013, pp. 127–135.

[11] W. Zhang, D. Zhou, L. Li, and Q. Gu. "Neural Thompson Sampling". In: *arXiv:2010.00827* (2020).

[12] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. "A fast and elitist multiobjective genetic algorithm: NSGA-II". In: *IEEE T. Evolut. Comput.* 6.2 (2002), pp. 182–197.

[13] S. Culley, D. Albrecht, C. Jacobs, P. M. Pereira, C. Leterrier, J. Mercer, and R. Henriques. "Quantitative Mapping and Minimization of Super-Resolution Optical Imaging Artifacts". In: *Nat. Meth.* 15.4 (2018), pp. 263–266.

[14] A. Descloux, K. S. Grußmayer, and A. Radenovic. "Parameter-Free Image Resolution Estimation Based on Decorrelation Analysis". In: *Nat. Meth.* 16.9 (2019), pp. 918–924.