pySTED : A STED Microscopy Simulation Tool for Machine Learning Training

Benoit Turcotte[1,2,*] **Anthony Bilodeau**[1,2,*] **Flavie Lavoie-Cardinal**[1,2], **Audrey Durand**[1,2,4]

[1]CERVO Brain Research Center; [2]Université Laval; [3]Institut Intelligence et Données (IID); [4]Canada CIFAR AI chair;
[*]Authors contributed equally
{betur57, anbil106}@ulaval.ca, flavie.lavoie-cardinal@cervo.ulaval.ca, audrey.durand@ift.ulaval.ca

Abstract

Despite the advantages in resolution granted by super-resolution fluorescence microscopy, the techniques remain challenging to use for non-expert users due to the large number of objectives which need to be optimized to obtain high quality images. Artificial intelligence, in particular reinforcement learning, could prove useful in assisting or controlling image acquisition. However, reinforcement learning approaches are data-hungry in training, rendering their application to super-resolution microscopy infeasible due to the large amount of sample waste training would require. We present pySTED, a STED microscopy simulation tool designed for development and application of reinforcement learning approaches to super-resolution microscopy. We show how pySTED allows to train reinforcement learning agents to resolve nanostructures in simulated synapses.

1 Introduction

Most of the super-resolution (SR) fluorescence microscopy techniques rely on our ability to switch fluorophores between an ON and an OFF state to overcome the diffraction limit of light (Nobel Prize of Chemistry, 2014) (Schermelleh et al. 2019). These approaches are used to characterize fluorescently-tagged sub-cellular structures in fixed and living cells allowing to resolve two objects that are a few nanometers apart (Schermelleh et al. 2019). STimulated Emission Depletion (STED) microscopy is a point-scanning SR technique which provides a ten fold resolution improvement in comparison to diffraction-limited confocal microscopy (Schermelleh et al. 2019). This SR technique remains challenging to use for microscopists, due to the required simultaneous optimization of conflicting objectives such as the spatio-temporal resolution, the signal-to-noise ratio (SNR), and the photobleaching. The optimization of those objectives depends on the control of several microscope parameters (e.g. excitation and STED beam intensities, pixel dwell time, pixel size). Gains in resolution and SNR typically come at the cost of increased photobleaching or temporal resolution (Lemon and McDole 2020).

Machine learning (ML) methods, more specifically reinforcement learning (RL) approaches, present themselves as promising solutions for assisting microscopists by learning to optimize STED imaging processes under these multi-objective trade-offs (Durand et al. 2018). In an RL problem, an agent interacts with an environment (e.g. the microscope), playing actions (e.g. modifying the acquisition parameters) which impact the environment's state (e.g. the sample) and generate a reward signal (e.g. the imaging objectives). Through training, the agent learns by attempting to maximize its earned reward. For example, imaging parameters optimization has been tackled under the most simple RL instance of multi-armed bandits, allowing the acquisition of high quality STED images of diverse biological structures in a completely autonomous fashion (Durand et al. 2018). These results opened the door to more applications falling under the more complex RL setting of sequential decision-making. However, sequential RL methods are known to be data-hungry, which limits their applicability to microscopy experiments. Developing and training such approaches would likely require large number of samples combined with time and cost intensive microscopy experiments. In other fields sharing similar concerns, such as robotics, this issue is usually tackled by developing and training RL algorithms within simulated environments before deploying and fine-tuning these methods in the real environment (Kober, Bagnell, and Peters 2013; Kadian et al. 2020).

In this work we introduce pySTED (Section 2), a simulation environment designed with the aim of supporting the development and training of ML methods for STED microscopy. More specifically, to facilitate the development of RL approaches, the pySTED simulator is integrated within the OpenAI gym framework (Brockman et al. 2016), allowing an RL agent to control acquisitions and learn through interactions with the simulated microscope. We showcase pySTED using an application of RL methods for STED imaging of biological structures (Section 3).

2 The pySTED simulator

STED microscopy relies on the combination of two laser beams that are co-aligned: a Gaussian excitation beam to excite fluorescent molecules (fluorophores) and a donut-shaped beam depleting the fluorophores to their ground state by the mean of stimulated emission. The laser beams are focused at the sample using the objective lens and the emitted light is collected for a specific dwell time on a single-pixel

[*]These authors contributed equally.

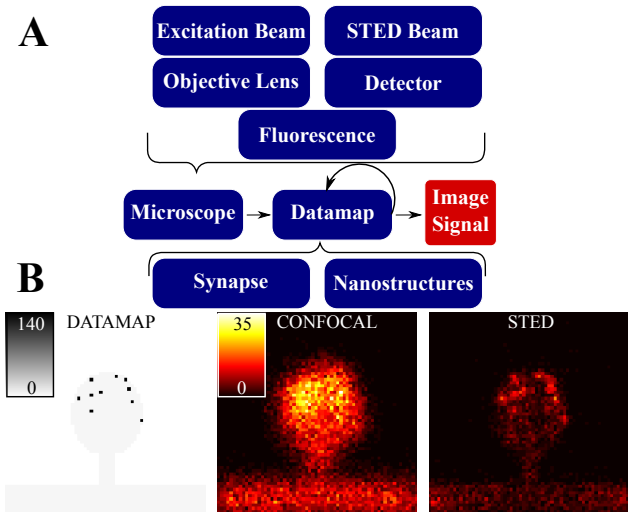

Figure 1: pySTED workflow and example acquisitions. **A** Simulator inputs (blue) and output (red). The datamap is updated in real time during acquisition by removing molecules which get photobleached. **B** (Left) Datamap displaying the nanostructures (NST). Colorbar shows the number of molecules. Confocal (Center, 25 μW Exc. power, 0 W STED power, and 10 μs pixel dwell time) and STED (Right, 0.15 mW Exc. power, 87.5 mW STED power, and 10 μs pixel dwell time) acquisitions from the datamap use the same color range (number of photons). Datamap and acquired images are 64×64 pixels (pixel size: 20 nm).

detector. A complete image is formed by scanning the sample point by point (pixel by pixel).

STED principles

The smallest distance (d) between two objects that can be resolved with a STED microscope is given by (Schermelleh et al. 2019)

$$d = \frac{\lambda}{2NA\sqrt{1 + \frac{I_{\text{STED}}}{I_{\text{sat}}}}}, \tag{1}$$

where λ is the wavelength, NA $\simeq 1$ is the numerical aperture of the objective, I_{STED} is the intensity of the STED beam, and I_{sat} is the beam intensity at which the decay rates of stimulated and spontaneous emission are equal. Fluorophore de-excitation to the ground state can take place either through the emission of a fluorescence photon or through stimulated emission (induced by the STED donut-shaped beam). Thus, for increasing I_{STED}, the effective size of the area where fluorescence photons are emitted is reduced (Schermelleh et al. 2019). When $I_{\text{STED}} = 0$ (i.e. a confocal acquisition), Eq. 1 reduces back to the diffraction limit of light : $d = \frac{\lambda}{2\text{NA}} \simeq 250$ nm for visible light (Schermelleh et al. 2019). In this paper, we refer to nanostructures (NST) for biological structures that are closer than the distance d when $I_{\text{STED}} = 0$ and cannot be resolved by confocal microscopy. Increasing the excitation intensity or the pixel dwell time generally increases the amount of emitted

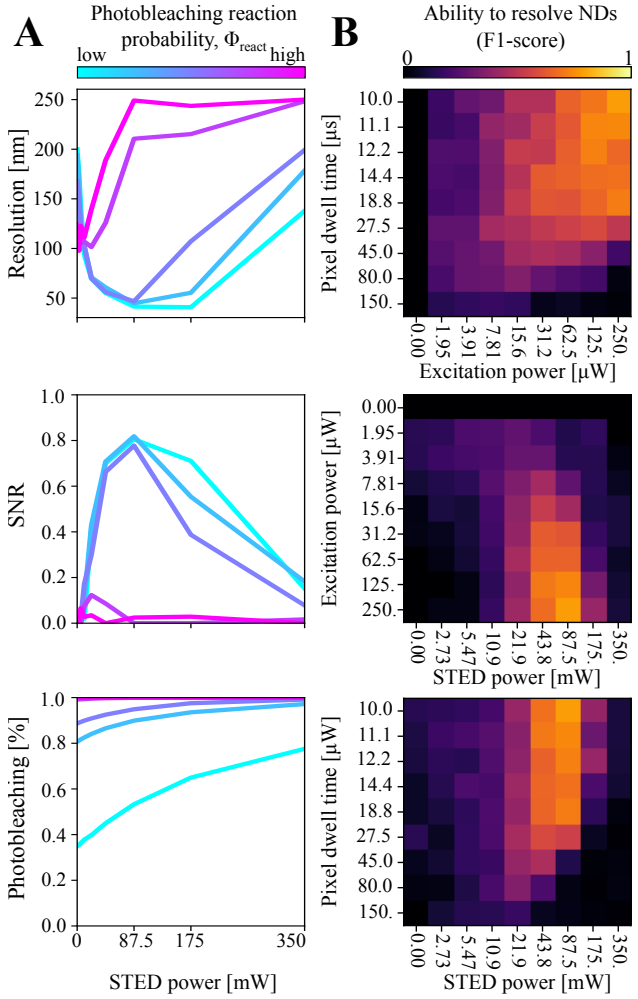

Figure 2: Effect of fluorophore parameters on imaging objectives and imaging parameters. **A** Imaging objectives as a function of STED power. **B** Maximum F_1-score projection for three parameters (Exc. and STED power, and pixel dwell time). F_1-score shows the ability to resolve distinct NST. Optimal imaging parameters are around 0.25 mW Exc. power, 87.5 mW STED power, and 10 μs pixel dwell time.

or detected fluorescence photons respectively, while increasing the STED intensity increases the resolution. In practice however, this improvement in resolution or detected fluorescence signal is limited by photobleaching of the fluorophores. Staudt (2009) derived the photobleaching rate from the excitation beam:

$$k_{\text{bleach}} = \text{T1}(I) I \sigma_T \Phi_{\text{react}}, \tag{2}$$

where $\text{T1}(I)$ is the population of fluorophores in the excited state, I is the excitation intensity, σ_T is the cross-section for triplet-triplet absorption, and Φ_{react} is the probability of a photobleaching reaction once the molecule is in the reactive state.

STED simulation

The pySTED simulator aims to realistically simulate STED imaging acquisitions. The implementation follows from the analytical descriptions of STED microscopy mechanisms (Leutenegger, Eggeling, and Hell 2010; Willig et al. 2006; Xie et al. 2013; Staudt 2009). [1] The simulator consists in a microscope that acquires images on a datamap describing the structure to image (Figure 1A). The microscope is comprised of five objects : the excitation and STED beams, an objective lens, a detector, and the photophysical parameters of the fluorophores. Each object is characterized by adjustable parameters which will affect the image signal, such as the fluorophore's quantum yield, the detector's efficiency, and the background signal. The datamap is represented by a 2D array in which each element of the structure of interest indicates the number of fluorescent molecules at that position (Figure1B, left).

pySTED allows to simulate corresponding confocal and STED acquisitions from a given datamap (e.g. NST inside a neuronal dendritic spine, Figure 1B). The photobleaching induced by STED acquisitions is simulated by multiplying a STED-specific constant to Eq. 2, proportional to the number of state transitions which the beam forces (Staudt 2009). Importantly, diverse photophysical constants associated with fluorescence transition (e.g. photobleaching reactivity, quantum yield, absorption cross-section) can be adapted to the simulated experimental conditions. The resolution, SNR and photobleaching objectives are strongly modulated by increasing STED power and rate of photobleaching reactions (k_{bleach}) (Figure 2A). Figure 2B displays the ability to resolve the NST using various combinations of microscope parameters, showing the impact of imaging parameters on a downstream task.

The pySTED simulator is integrated within an OpenAI gym framework (Brockman et al. 2016) to facilitate the formulation of STED imaging tasks as RL tasks.[2,3]

3 RL for resolving biological nanostructures

This section provides an example of RL applied to STED control for the specific task of resolving NST in simulated dendritic spines. A RL agent interacts with an environment by sequentially making decisions based on its observations. At every time step t, the agent receives the current state s_t of the environment and chooses an action a_t according to its internal policy π (a probability distribution mapping states to actions). As a result of taking action a_t, the agent transitions in a new state s_{t+1} and receives a reward r_{t+1}. The goal of the agent is to maximise the expected return $\mathbb{E}[R]$ in episodes of length T, where $R = \sum_{t=0}^{T-1} \gamma^t r_{t+1}$ is the sum of discounted reward with discount factor $\gamma \in [0, 1]$ (Sutton and Barto 2018).

Task In the task of resolving NST, a state corresponds to the confocal image of a new spine and actions correspond to choosing imaging parameters. An episode of this task therefore corresponds to acquiring a sequence of T images. At each time step $t < T$, the agent receives the confocal image of a new spine, selects parameters, an image is acquired with these parameters, and the imaging objectives (Resolution, SNR, and Photobleaching) are evaluated. The agent also has a memory of the previous confocal/STED pair and of the history of the episode, i.e. previously selected parameters and the resulting imaging objectives. At the end of an episode, NST locations are predicted by looking at the sequence of acquired images. On each image, local intensity maxima are identified and fitted with a 2D Gaussian. If the standard deviation along both of the axis of the Gaussian is below 250 nm (diffraction limit), a NST is predicted at that location. Predicted NSTs are then associated to the ground truth (GT) positions (from the datamap) by solving a linear sum assignment problem with the Hungarian algorithm. An associated pair (< 2 pixels apart) is a true positive association, a non-associated GT is a false negative, and a non-associated prediction a false positive. This performance is summarized by the F_1-score.

Methods

We considered Proximal Policy Optimization (PPO) to resolve NST in simulated dendritic spines using the Atari convolutional network as backbone. PPO is a state-of-the-art RL algorithm, achieving high performance on problems such as MoJuCo and Atari (Schulman et al. 2017).

We trained two PPO agents using different reward functions. The first agent uses the F_1-score on the detected NST as reward. However, computing the F_1-score requires knowing the ground-truth position of the NST, which is not available in real microscopy experiments. We therefore propose a second agent which uses an improvement score as reward. For this agent, a bank of 50 images was acquired in the simulation platform by an Expert microscopist. At time step t, the image acquired by the agent is ranked against the bank of images using the F_1-score and a reward $r_{t+1} \in [0, 1]$ is generated based on the ranking (1: highest rank, i.e. better than all Expert images; 0: worst rank). While the F_1-score is still required for ranking, the agent does not have access to these values during training, which is more similar to a real-life task. For both reward functions, the maximum episodic cumulative reward is of 10.

The agent trains in episodes of 10 time steps. Both agents train for 200 000 steps (20 000 episodes), meaning they each acquire 200 000 image (repeated 5 times with different random seeds). This would be infeasible to do on a real microscope, highlighting the need for a robust STED simulation platform to train RL agents. After training, the agents are evaluated on 100 episodes of 10 time steps. For the experiments we limit our agent's parameter selection to image-wide values for excitation power ($[0, 250]$ μW), STED power ($[0, 350]$ mW), and pixel dwell time ($[10, 150]$ μs).

Results

Figure 3A displays the performance of the agents trained using the F_1-score (top, blue) and ranking (bottom, orange) as

[1] https://github.com/FLClab/pySTED

[2] https://github.com/FLClab/gym-sted

[3] https://github.com/FLClab/gym-sted-pfrl

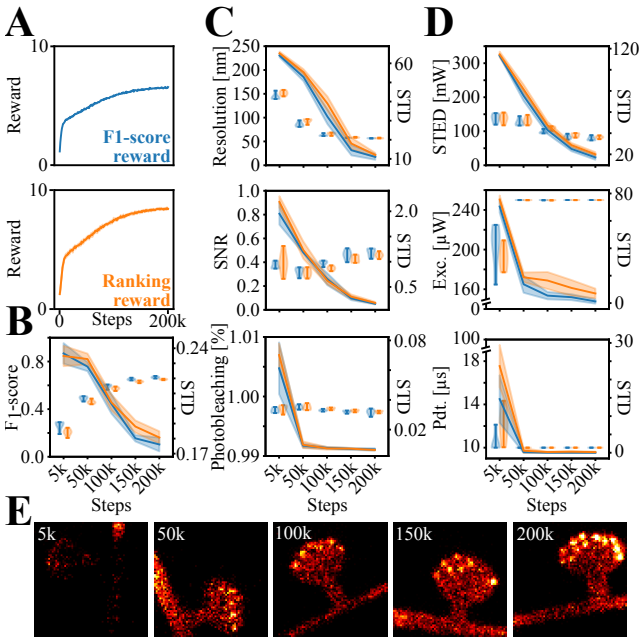

Figure 3: RL agent performances. **A** Training curve for agent using F_1-score as reward (top, blue), and ranking as reward (bottom, orange). Evolution of the F_1-score (**B**), the imaging objectives (**C**) and the imaging parameters (**D**) at various steps during training (violinplot). Mean and standard deviation (line and shade) of the evolution of standard deviation (STD) across training is presented (right axis). **E** Typical images acquired at various steps during training (1.28μm $\times$ 1.28μm, same intensity scale). 5 repetitions.

reward. The performance of both agents converge within approximately 150k steps. We next compared the capability of the trained agents at resolving the NST. Agents trained using F_1-score as reward achieve a mean F_1-score of 0.65 ± 0.06, while the agents using using ranking as reward achieves a mean F_1-score of 0.63 ± 0.06. The evolution of the F_1-score (Figure 3B) shows that he agent trained with the F_1-score as reward obtained slightly higher performance, as expected. The standard deviation (STD) of the F_1-score at each evaluated steps decreases for both agents implying the convergence of the metric.

The quantification of the imaging objectives during training reveals an increase in the resolution and SNR, while the photobleaching remains constant (Figure 3C). Notably, the imaging objectives converge as depicted by the reduction in STD obtained by the agents (Figure 3C, lines). The quantification of the selected parameters shows that both agent converged to similar imaging parameters even if trained using different reward functions (Figure 3D). Interestingly, the imaging parameters at convergence are similar to the optimal parameters that were extracted from the grid search in Figure 2B (STED power : 82 mW, Exc. power : 0.25 mW and Pdt. : 10 μs). This confirms the capability of both agents to resolve the NST STED imaging task (Figure 3E).

4 Conclusion

The pySTED simulator allows RL (and more broadly ML) methods to be trained and validated in simulation prior to the deployment on a real microscope. This will allow the development of new AI based tools to assist microscopists on complex image acquisition tasks. Here, we have shown how an RL agent could learn to optimally image simulated NST in dendritic spines. The next step will be to simulate the NST dynamics in time-resolved live-cell experiments.

Acknowledgments

F.L.C. is a Canada Research Chair Tier II, Audrey Durand is a CIFAR AI Chair, and A.B. is supported by a PhD scholarship from the Fonds de Recherche du Québec - Nature et Technologie (FRQNT) and an excellence scholarship from the FRQNT strategic cluster UNIQUE. The project was funded by a FRQNT Team Grant awarded to F.L.C. and A.D.

References

Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *arXiv:1606.01540 [cs]*.

Durand, A.; Wiesner, T.; Gardner, M.-A.; Robitaille, L.-E.; Bilodeau, A.; Gagné, C.; De Koninck, P.; and Lavoie-Cardinal, F. 2018. A machine learning approach for online automated optimization of super-resolution optical microscopy. *Nature Communications*, 9(1): 5247.

Kadian, A.; Truong, J.; Gokaslan, A.; Clegg, A.; Wijmans, E.; Lee, S.; Savva, M.; Chernova, S.; and Batra, D. 2020. Sim2Real Predictivity: Does Evaluation in Simulation Predict Real-World Performance? *IEEE Robotics and Automation Letters*, 5(4): 6670–6677.

Kober, J.; Bagnell, J. A.; and Peters, J. 2013. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11): 1238–1274.

Lemon, W. C.; and McDole, K. 2020. Live-cell imaging in the era of too many microscopes. *Current Opinion in Cell Biology*, 66: 34–42.

Leutenegger, M.; Eggeling, C.; and Hell, S. W. 2010. Analytical description of STED microscopy performance. *Optics Express*, 18(25): 26417–26429.

Schermelleh, L.; Ferrand, A.; Huser, T.; Eggeling, C.; Sauer, M.; Biehlmaier, O.; and Drummen, G. P. C. 2019. Super-resolution microscopy demystified. *Nature Cell Biology*, 21(1): 72–84.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *arXiv:1707.06347 [cs]*.

Staudt, T. M. 2009. Strategies to reduce photobleaching, dark state transitions and phototoxicity in subdiffraction optical microscopy.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: an introduction*. Adaptive computation and machine learning series. Cambridge, Massachusetts: The MIT Press, second edition edition.

Willig, K. I.; Keller, J.; Bossi, M.; and Hell, S. W. 2006. STED microscopy resolves nanoparticle assemblies. *New Journal of Physics*, 8(6): 106–106.

Xie, H.; Liu, Y.; Jin, D.; Santangelo, P. J.; and Xi, P. 2013. Analytical description of high-aperture STED resolution with 0–2pi vortex phase modulation. *Journal of the Optical Society of America.*, 30(8): 1640–1645.