COmputatioNal Imaging as a training network for Smart biomedical dEvices

COmputatioNal Imaging as training network for Smart biomedical dEvices

Project No. 101072354

Deliverable 3.3

Report on software for and simulations of deep-learning-based adaptive computational imaging

WP3 – SMART-2PM

Authors	Marcos Antonio Obando (CWI), Felix Lucka (CWI)
Lead participant	UJI
Delivery date	23/01/2025
Dissemination level	Public
Type	Document, report

Version 01

Funded by the European Union

Engineering and Physical Sciences Research Council

Funded by the European Union (GA 101072354). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. UK participants are funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee grant number EP/X030733/1).

Revision history

Author(s)	Description	Date
Marcos Antonio Obando (CWI), Felix Lucka (CWI)	Draft deliverable	08/01/2025
Enrique Tajahuerce (UJI), Jörn-Frederik Wotzlaw (accelCH)	Internal review and editing	22/01/2025
Andrea Farina (CNR)	Final review and submission	23/01/2025

Contents

Partner short names

CNR	Consiglio Nazionale delle Ricerche
UJI	Universitat Jaume I
CWI	Centrum Wiskunde & Informatica
accelCH	accelopment Schweiz AG

Abbreviations

aCS	Adaptive Compressed Sensing
aSPI	Adaptive Single Pixel Imaging
AI	Artificial Intelligence
AO	Adaptive Optics
CM	Conditional Mean
CNN	Convolutional Neural Network
CS	Compressed Sensing
D	Deliverable
DAQ	Data Acquisition
DL	Deep Learning
DC	Doctoral Candidate
DMD	Digital Micromirror Device
DNN	Deep Neural Network
EC	European Commission
EU	European Union
GPU	Graphics Processing Unit
HEU	Horizon Europe
KL	Kullback–Leibler
M	Month
MCMC	Markov chain Monte Carlo
ML	Machine Learning
MS	Milestone
OED	Optical Experimental Design
RL	Reinforcement Learning
sOED	sequential Optical Experimental Design
SPI	Single Pixel Imaging
WP	Work Package

Executive summary

In this deliverable, we report on the progress made in WP3 (SMART-2PM) with respect to developing a software framework for nonlinear microscopy of biological tissues with patterned illumination (task 3.5) and the design and analysis of single-pixel adaptive computational imaging strategies (task 3.4). It reflects the work of DC9 Marcos Obando (CWI, Adaptive compressed sensing using deep learning in non-linear microscopy) during months 1-14 in collaboration with UJI and CNR.

Objectives

The software framework (task 3.5) should provide the central collaborative platform where all experimental and methodical developments in WP3 come together. As such, it needs to take all properties and aberrations of the optical system into account and be able to integrate modules for uploading illumination pattern onto the digital micromirror device (DMD), steering the adaptive optics (AO), controlling data acquisition (DAQ), performing image reconstruction and adapting the illumination pattern to the object being imaged (adaptive single pixel imaging, aSPI)

The goal of aSP (task 3.4) is to recover images using fewer measurements, finer details, and shorter reconstruction times than conventional, non-adaptive SPI techniques. The key idea is to adapt the illumination pattern to the object that is imaged during the measurement instead of choosing a predefined set of illumination pattern before the measurement. To do this in a fast and efficient way, we aim to develop aSPI algorithms based on machine learning (ML, the methodical approach dominating Artificial Intelligence advances of the past years), in particular algorithms that use deep neural networks (deep learning, DL) trained with reinforcement learning (RL) techniques.

Methodology and implementation

The requirements and specification of the software framework were discussed through regular online meetings of the SMART-2PM team and prepared the secondment at UJI, during which the first version of the Python-based framework (working title: smart-2pm-napari) was designed and implemented. The characteristics of the current SMART-2PM system were embedded into smart-2pm-napari, its interaction with the DMD and DAQ modules of the current version of the SMART-2PM system were tested, and its ability to reconstruct images from SPI measurements demonstrated.

In parallel, the use of RL methods for aSPI was explored. We started with a simple, artificial aSPI scenario in which the validity and effectiveness of different methodical choices is easy to examine and verify. After these proof-of-concept studies, we considered a more complex, higher dimensional scenario, where the aSPI algorithm incorporates state-of-the-art generative AI components. We examined which methodical components scale well, i.e., are still efficient in large scale imaging scenarios and which ones need to be modified. The algorithms are implemented in a software module that is designed to be integrated into the smart-2pm-napari software described above later in the project.

Outcomes

The swift collaboration between CWI and UJI resulted in a first version of the smart-2pm-napari software. It is a plug-in for napari, a Python-based image visualization and analysis framework with a graphical user interface. The current version of smart-2pm-napari already allows to flexibly connect to the DMD and DAQ devices of the SMART-2PM system and can reconstruct an image from SP measurements. Different sets of illumination pattern can be either generated or loaded. As the whole implementation is done in Python, it is easy to extend the framework to incorporate more components.

Our proof-of-concept studies show that RL-based algorithms with generative AI components are promising approaches for aSPI. In addition, we identified which components need further development to scale to large images. We are preparing a first paper publication on the results.

Next steps

Together with the SMART-2PM partners, the smart-2pm-napari software will be extended to communicate with the adaptive optics module. This way, we can integrate the ongoing developments on defocus correction, aberrations measurement, and higher aberration correction (see Deliverable 3.4). Furthermore, an interface between smart-2pm-napari and the aSPI software will be developed such that a proof-of-concept demonstration of an adaptive measurement can be carried out in the future.

There are several next developments of the aSPI algorithms directed at either improving their performance or making training and run times faster. For instance, their mapping onto graphics processing units (GPUs) will be improved to speed up both their training and run times and the efficiency of different algorithmic components will be examined and improved. The final challenge will be to transition from simulated data to real experimental data.

1 Introduction

In single pixel imaging (SPI), the whole object of interest is illuminated with a series of spatially modulated light pattern and the reflected or transmitted light is integrated in a single photo detector (Gibson et al., 2020). A common way to realize SPI is to illuminate a digital micromirror device (DMD) with a pulsed laser, which allows to use binary light pattern. To achieve faster acquisition times than conventional point-scanning techniques, SPI typically tries to use less patterns than the number of pixels in the desired image resolution, i.e., to stay below the Nyquist-Shannon sampling rate (Foucart & Rauhut, 2013). Since its introduction in the seminal paper by Duarte et al. (2008), SIP has been explored for a wide range of applications (Gibson et al., 2020). It is appealing when sensitivity at non-visible wavelengths or very precise timing resolution is crucial, but impractical or prohibitively costly with conventional, pixelated detectors such as silicon-based charge-coupled devices (CCD) and complementary metal-oxide-semiconductors (CMOS).

The aim of work package 3 (WP3 called “SMART-2PM”) of the CONcISE project is to develop a SPI-based system for 2-Photon Microscopy (also known as “non-linear microscopy”) that surpasses the fundamental trade-off between penetration depth, sensitivity, and imaging speed of current nonlinear microscopes and allow for fast, high-resolution imaging of biological samples. To achieve this, the SMART-2PM system will combine several innovative components, such as a new, all-fiber fs-laser in the mid-infrared (see Deliverable 3.2), an adaptive optics (AO) module (see Deliverable 3.4) and the ability to adapt the illumination pattern during the measurement. An overview of the system design is given in Deliverable 3.1. In this report, we describe the development of a software framework that is able to integrate and control the different components of the SMART-2PM system in Section 2, and the development of computational algorithms for reconstructing images from the SPI measurements and for adapting the illumination pattern (adaptive single pixel imaging, aSPI) in Section 3. For the latter, we give a more detailed and technical introduction in the following.

The field of Compressed Sensing (CS) developed and analysed algorithms that can reconstruct high-quality images free of aliasing artifacts from under-sampled SPI measurements if it is known a-priori that the image is sparse, i.e., exhibits low complexity in some suitable spatial representation (Eldar & Kutyniok, 2009; Foucart & Rauhut, 2013). The first generation of CS methods used explicit, hand-crafted models of sparsity, e.g., by constraining the sum of the absolute values (i.e., the L1-norm) of the wavelet coefficients of the image. In recent years, learning more expressive models from examples of typical images has become a major research direction and state-of-the-art (SOTA) methods for CS image reconstruction rely on machine learning (ML) techniques, in particular deep learning (DL, (Goodfellow et al., 2016)) using deep neuronal networks (DNNs) with convolutional layers (CNNs). For instance, powerful diffusion-based image generators can be used to solve the inverse problem behind CS by casting it as a Bayesian inference problem (Zhao et al., 2023): The learned image generator defines an implicit prior that, combined with the likelihood distribution given by measured data leads to an implicit posterior distribution that can be sampled from by sophisticated Markov chain Monte Carlo (MCMC) methods.

A crucial decision in the setup of SPI is the choice of the set of binary sensing patterns. Generic error bounds from CS theory together with algorithmic considerations lead to scrambled Hadamard matrices as the most employed choice that works well over a wide range of images (Foucart & Rauhut, 2013; Gibson et al., 2020). If one knows which type of images to expect a-priori, i.e., before the measurement, one can try to further increase performance or lower acquisition time by choosing a more optimal set of patterns. A principled, Bayesian approach to do this is given by optimal experimental design (OED), which determines a fixed set of optimal sensing patterns for a given image prior (Rainforth et al., 2024; Ryan et al., 2016). In sequential OED (sOED) or adaptive Compressed Sensing (aCS), one tries to go even one step further by choosing each new sensing pattern sequentially, i.e., during the experiment and based on the measurements about the image collected so far, i.e., a-posteriori. Depending on the image priors used, different algorithmic approaches have been considered to solve the optimization problems corresponding to OED or sOED: While using Gaussian priors can lead to tractable problems with closed form solutions (Burger et al., 2021), using more sophisticated priors leads to very complex bi-level optimization problems (Rainforth et al., 2024; Ruthotto et al., 2018). Despite considerable recent progress (De los Reyes & Villacís, 2023), these problems cannot be solved fast enough computationally to enable sOED in many imaging applications such as SPI.

Recent work started exploring a different idea to solve sOED problems by using reinforcement learning (RL) (W. Shen & Huan, 2023; Z. Shen et al., 2022; Wang, Florian, et al., 2024; Wang, Lucka, et al., 2024): RL is a ML paradigm that tries to develop AI systems that learn to solve sequential decision-making problems in uncertain, dynamic environments to maximise a predefined reward. In particular, deep RL has achieved remarkable results on a range of complex problems such as robotics, autonomous driving or gaming in the past decade (Li, 2023). Applied to sOED, the idea is that instead of making the design choice by solving a computational optimization problem, RL learns a choice strategy or policy that is parameterized by a neural network. The network receives a representation of the current state of information (i.e., moments or samples of the posterior distribution) as an input and outputs the design decision. During a computationally extensive training phase, a given policy network is iteratively applied to test sOED problems, evaluated and improved based on a cumulative reward assigned to all the decisions taken. Once trained, the policy network can typically be executed fast on new sOED problems, making this approach viable for imaging applications. Compared to classical approaches that focus only on the sOED problem given and try to resolve its full complexity, RL completely relies on its experience in solving similar sOED problems during training.

2 Software framework for SMART-2PM

Figure 1 illustrates the main hardware components of the SMART-2PM system and how the software platform should interact with them. We chose to implement the software as plugins for napari (https://napari.org/), a Python-based image analyser that allows processing of large data sets. This way, many components that we needed were already provided and we could focus on designing modules for the hardware control.

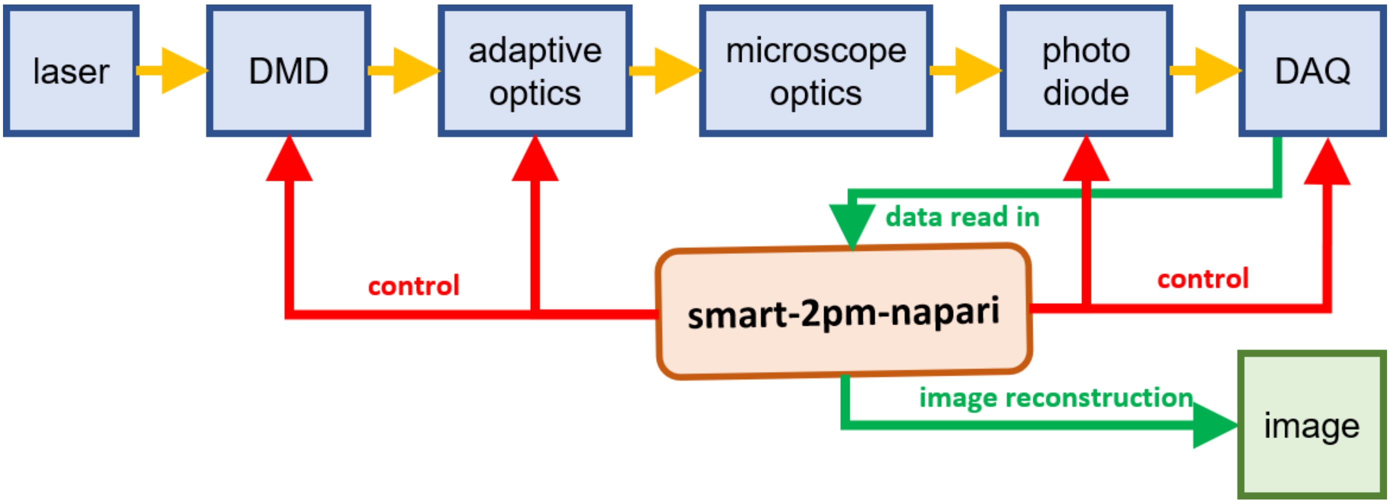

Figure 1: Schematic of the order of the different hardware components (top row) in the SMART-2PM system, the software platform smart-2pm-napari, which components it controls and how it reconstructs images.

The first hardware control module that we implemented controls the DMD, the main component that allows to spatially structure the laser illumination light by only projecting part of it onto the sample. Figure 2 illustrates the DMD module that allows to setup several binary pattern sequences and upload them to the DMD, such that they will be projected sequentially once the measurement starts. It allows to modify basic settings of this projection such as its physical position and orientation on the micromirror array. The projection time per pattern can also be configured from this menu (specified in milliseconds), as well as the option to invert the whole pattern sequence. In a second tab of the DMD module, the user can compile a list of sequences of pattern to be projected by the DMD. One option is to add predefined sequences such as Hadamard or scrambled Hadamard pattern of different orders or random sets of patterns. Another option is to load a user-defined sequence of binary patterns from a folder.

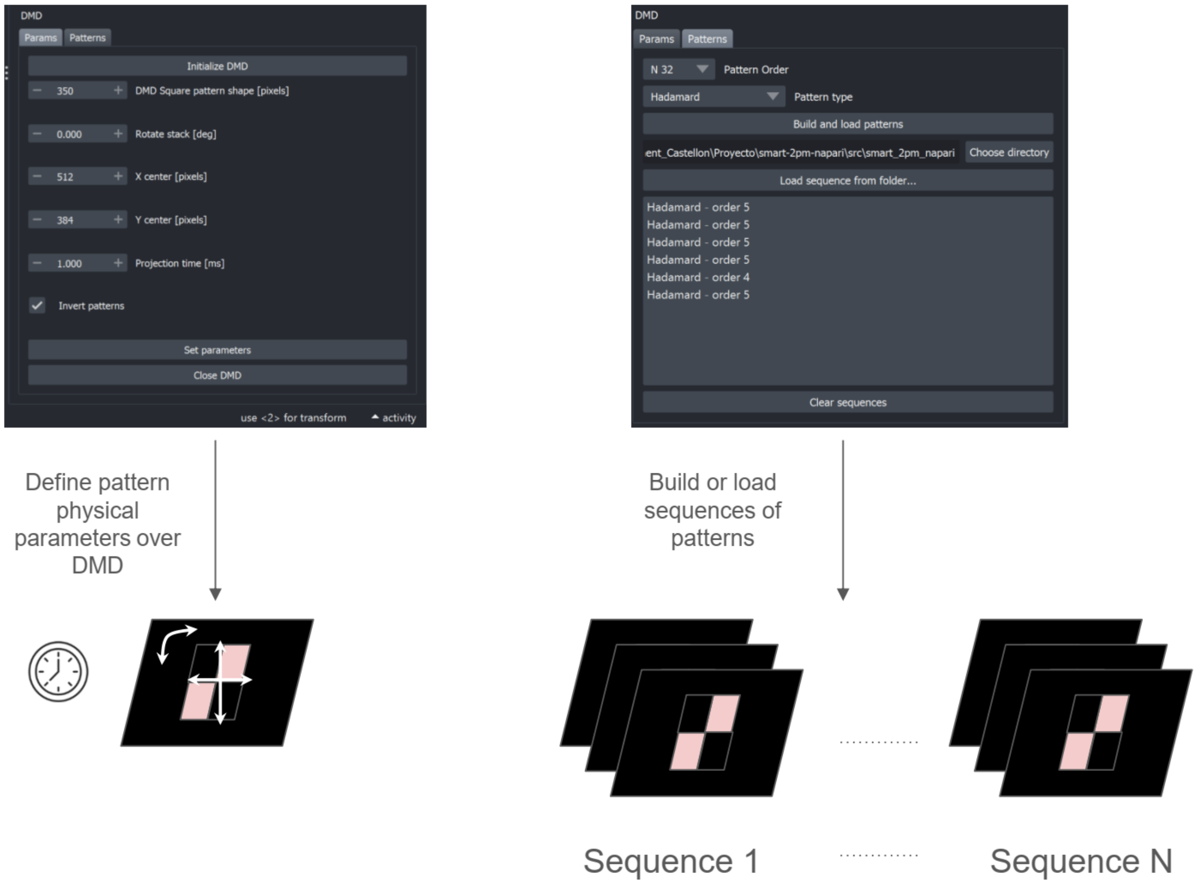

Figure 2. The DMD module allows to define the sequence of binary pattern, sets up their physical parameters (translational shifts and rotation respect to the center, binary inversion and projection time per pattern) and loads them onto the DMD. Over the patterns tab, one can compose lists of sequences of patterns to be projected after each other. The user can add commonly used sequences like Hadamard pattern or load any self-defined set of binary patterns.

The second hardware control module that we implemented is the DAQ module that controls a National Instruments device consisting of the both the photodiode and the DAQ hardware component. We use on the nidaqmx Python interface to set up the communication. The data acquisition allows to measure two channels: an analogic signal coming from the voltage conversion of the photon counting detector in the photodiode and a trigger signal from the DMD indicating the time at which each pattern is projected onto the sample. Figure 3 shows how these two signals are used to generate the SPI measurements. Within the time of the projection trigger, the collected photodiode voltage samples are processed as follows: A fixed number of the first samples are discarded, as they were taken while the micromirror array was flipping from the old to the new pattern. After this, the remaining samples are averaged to obtain a resulting single pixel measurement for the pattern being projected. The photodiode sampling rate and number of samples to acquire can be defined using a slider button and can, in principle, be faster than the pattern projection time that was defined in the DMD module (see above).

The interplay between DMD and DAQ module determines the SPI data collection as illustrated in Figure 4. Once they are configured, the measurement tab allows to acquire a set of single-pixel measurements, illustrated in Figure 5. This module furthermore saving the measurements, displaying a 1D plot of them and has already implemented a simple image reconstruction technique using the back-projection algorithm. Figures 5-7 illustrate its features and usage applied to experimental data acquired from a test object.

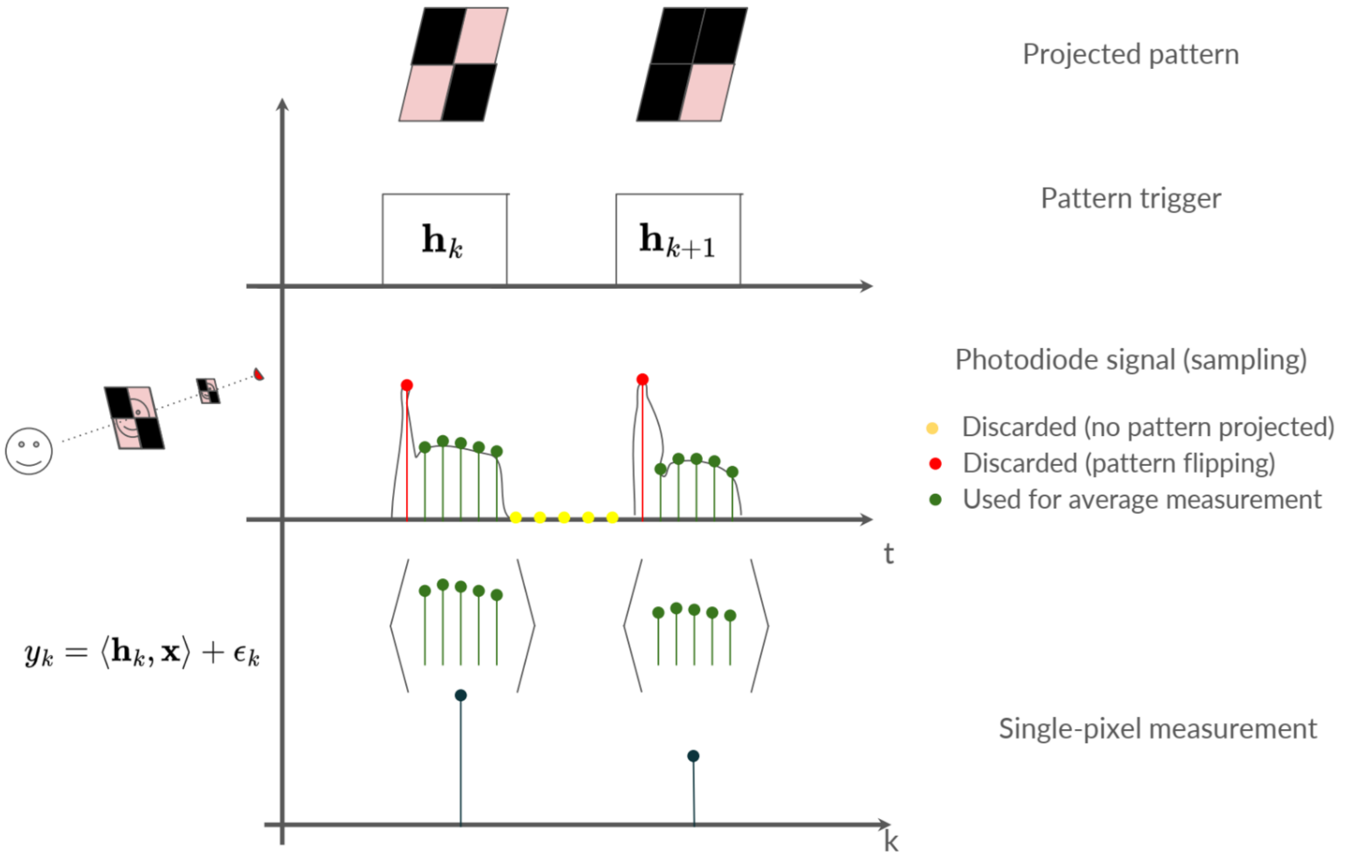

Figure 3. Acquisition signals for the DMD and photodiode measurements. From top to bottom: Each pattern is projected for a fixed amount of time during which the DMD produces a trigger signal that allows to associate the voltage samples of the photodiode over this time to the pattern displayed. The first photodiode signal after the trigger is discarded as it reflects the flipping of the mirror. The signals thereafter are averaged to produce the measurement.

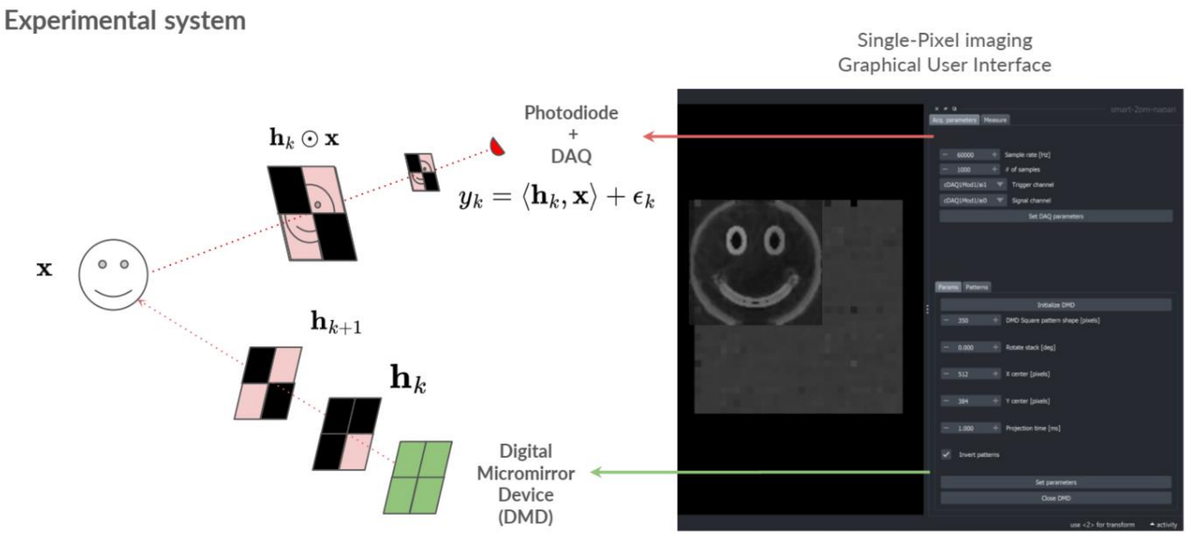

Figure 4: Schematic of the interplay between DMD module and DAQ module steering the SPI measurement.

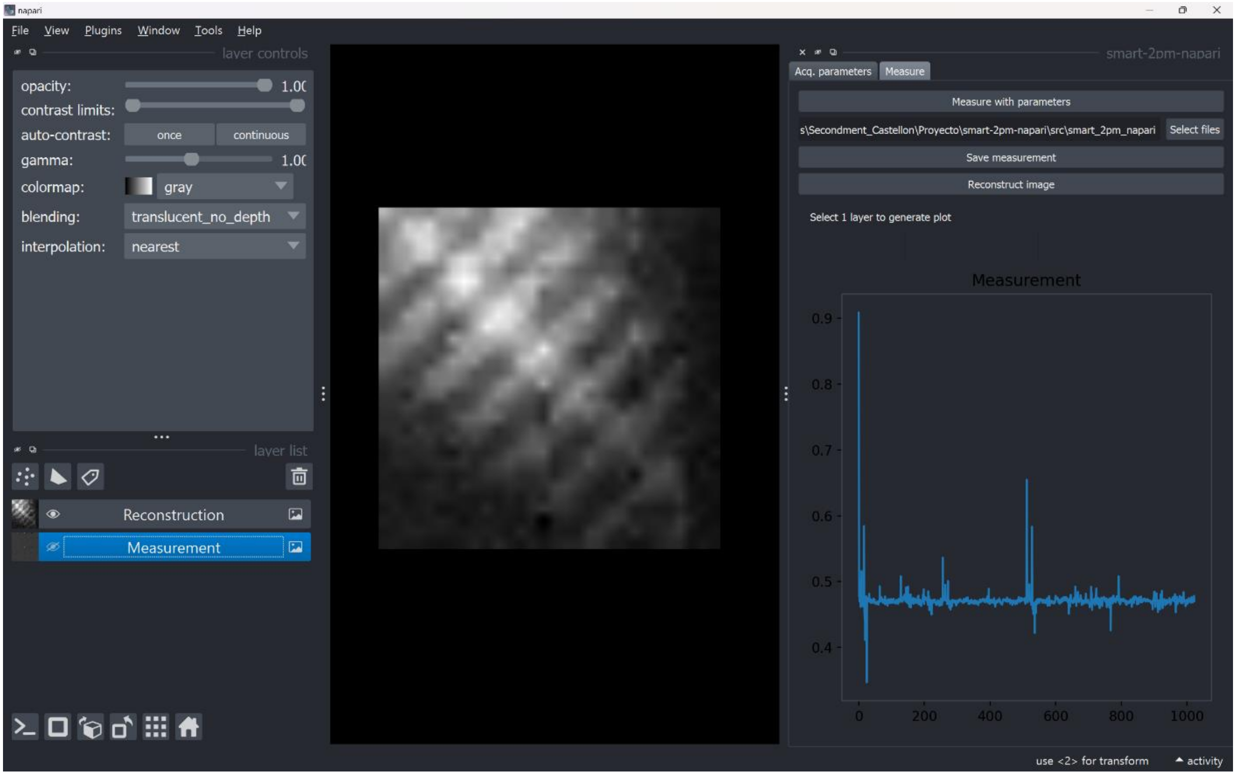

Figure 5. The measurement tab in smart-2pm-napari initiates the measurement and the acquired single pixel data can be saved, displayed (right 1D plot) and reconstructed into an image (middle plane).

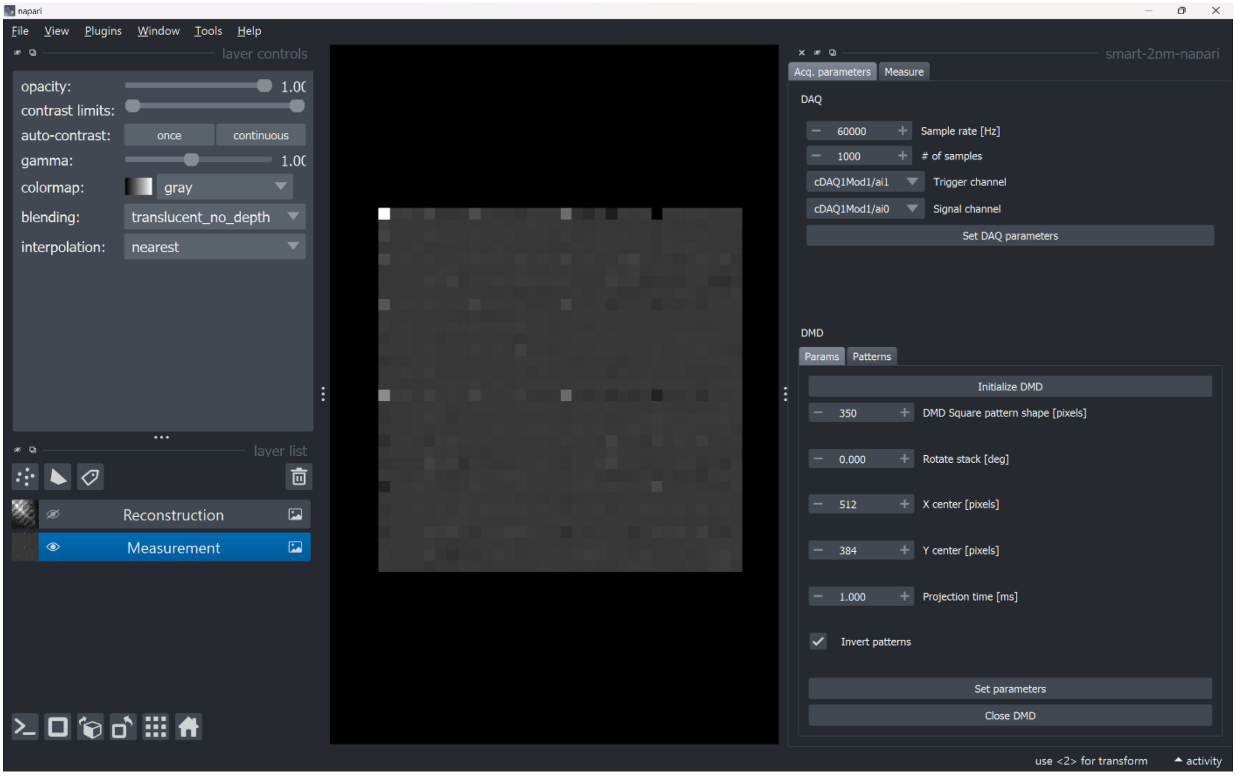

Figure 6. If Hadamard pattern are chosen (here of order 32), the acquired data can also be visualized in 2D (middle plane).

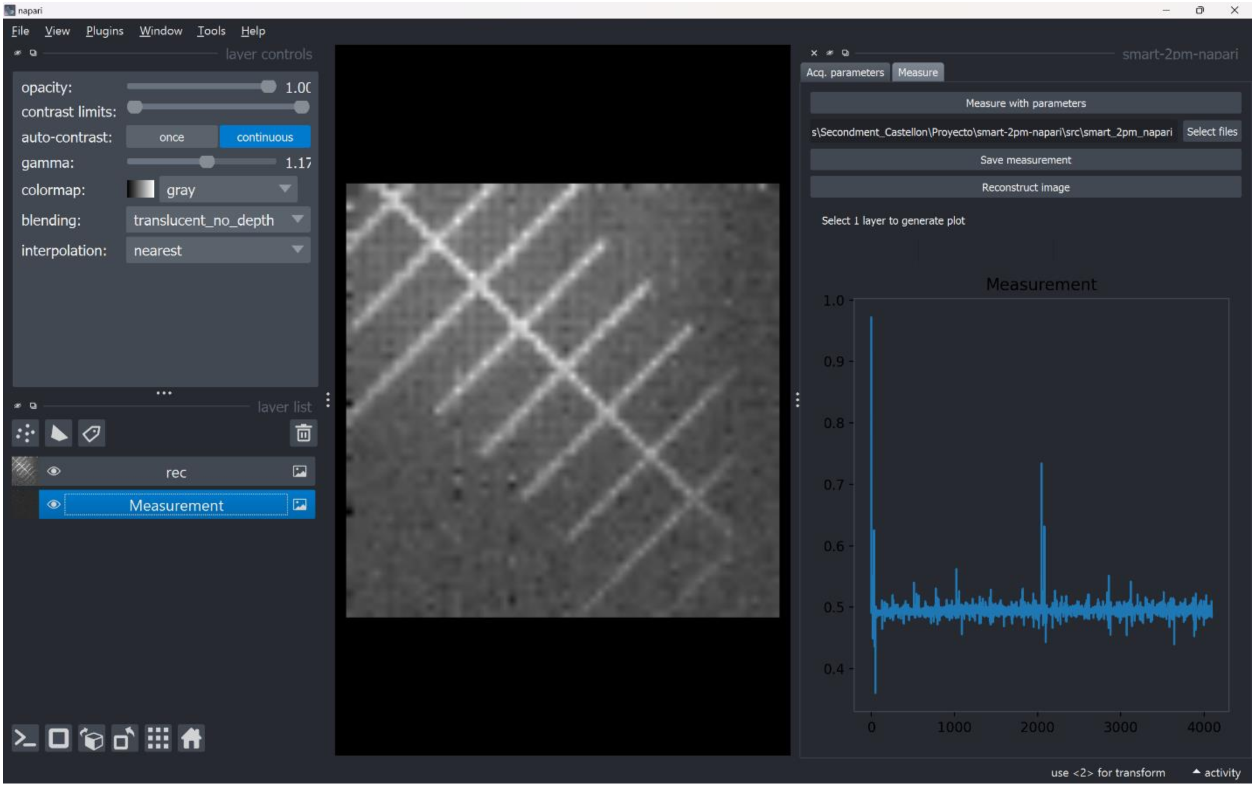

Figure 7. Example of the image of a test object reconstructed from SPI with Hadamard pattern of order 64 instead of 32 as in Figure 5.

3 Bayesian adaptive Single Pixel Imaging via Reinforcement Learning

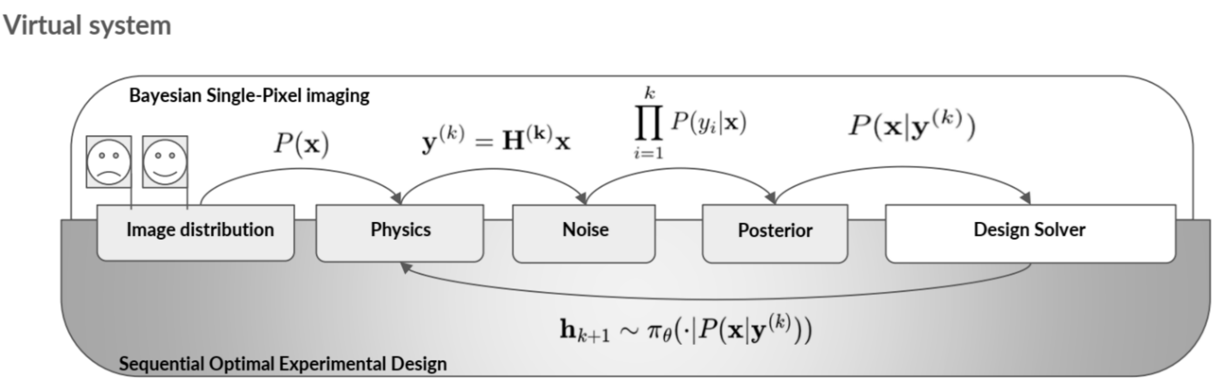

Figure 8. Overview of the key components of the Bayesian framework for adaptive single pixel imaging.

As described in the introduction, the problem of aSPI amounts to solving a sOED problem. We decided to formulate the sOED problem using the framework of Bayesian inversion, i.e., Bayesian inference applied to inverse problems such as SPI (Kaipio & Somersalo, 2005; Stuart, 2010). It allows for a mathematically rigorous modelling and quantification of all relevant uncertainties. Figure 8 illustrates its key components:

- The a-priori or prior distribution P(**x**) is a probability distribution that models our believes about how typical images that we may measure look like before any measurement takes place. One choice is to use explicit prior models like multidimensional Gaussians or sparsity-inducing L1-type priors from classical CS theory. While using such priors has computational advantages, they are typically not very expressive, i.e., they fail to concentrate probability mass on realistic images. We decided to also consider very expressive but implicit priors defined through state-of-the-art diffusion-based image generators (Zhao et al., 2023).
- The forward operator H describes the imaging physics by mapping image **x** to measurements **y**. We can model all relevant properties of the microscope (e.g., include the adaptive optics later in the project) and the single pixel sensing mechanism. Together with a stochastic model of the measurement noise, it defines the likelihood P(**y**|**x**), a distribution describing how likely an image **x** would lead to measurements **y**. As such, it encapsules a digital twin of the SMART-2PM system, and we will improve its accuracy throughout the project.
- The a-posteriori or posterior distribution P(**x**|**y**) combines the information given by prior and likelihood and represents everything we know about the image **x** after the measurement. It is the central object in Bayesian inversion but often not tractable, which means that computing quantities of interest such as the posterior mean (conditional mean estimate, CM) or covariance is computationally challenging. We adapted and evaluated sophisticated Markov chain Monte Carlo (MCMC) algorithms described in (Chung et al., 2022) to sample the posterior resulting from using implicit priors based on diffusion-based image generators.

Given the posterior P($\mathbf{x}|\mathbf{y}^{(k)}$) after measuring with k sensing pattern, the sOED problem consists of deciding which subsequent experiment design, i.e., which sensing pattern to measure next, would be optimal (cf. “Design Solver” in Figure 8). It comes with various challenges:

- To formulate the sOED problem properly, one needs to decide what “optimal” means: A “D-optimal” design tries to maximize the information gained from the (k+1)-th measurement, which is reflected in how much it would change the posterior distribution P($\mathbf{x}|\mathbf{y}^{(k)}$). Mathematically, this is expressed as the the Kullback–Leibler (KL) divergence between P($\mathbf{x}|\mathbf{y}^{(k)}$) and P($\mathbf{x}|\mathbf{y}^{(k+1)}$). In practice, this quantity is hard to compute and needs to be approximated. An alternative is given by “A-optimal” design, which minimizes the expected error between the CM estimate and **x**. While A optimality is easier to implement, it does not capture all aspects of the information gained by the additional measurement.
- As the whole posterior is a continuous function of **x** and typically not tractable, we need to use design solvers that only work with a discrete representation of it, e.g., an estimate of the posterior mean and variance.
- Solving the sOED problem leads to a numerical optimization problem with a complex structure, for which no efficient algorithms exist that are guaranteed to find the best solution. Even if ad-hoc heuristics are used that produce approximate solutions, the design decision often needs to be computed very fast.
- In many applications, the scalability of sOED algorithms is a major challenge. In aSPI for instance, the number of possible design choices in each step scales as 2^n with the number of DMD-pixels n. This renders naive approaches that go through all possible designs useless for realistic image resolutions.

For many applications in inverse problems and imaging, the above difficulties prevented the adoption of sOED so far.

We decided to follow a new line of research (W. Shen & Huan, 2023; Z. Shen et al., 2022; Wang, Florian, et al., 2024; Wang, Lucka, et al., 2024) that tries to approach it from a different perspective: Instead of trying to devise algorithms that resolve the full complexity of the sOED problem in every step, we try to learn a strategy or policy how to solve sOED problems of a particular class by gradually improving the performance on a set of example problems from the same class. The policy is parameterized by a deep neural network and this type of training is called reinforcement learning (RL). To develop this approach for aSPI, we needed to work on several components:

- We reformulated the sOED problem for SPI as a partially observable Markov decision process (POMDP).
- In principle, the policy maps the current posterior distribution P($\mathbf{x}|\mathbf{y}^{(k)}$) - which is an infinite dimensional object - to a probability distribution over the 2^n possible sensing pattern, which can be huge number. We needed to explore efficient ways to parameterize the policy that work with discrete representations of the posterior (e.g., its mean and variance) and map to low-dimensional probability distributions over the next sensing pattern.
- RL relies on defining a reward function that supplies feedback to the policy training. To connect to classical Bayesian sOED, we transcribed A-optimal and D-optimal design to a reward function. In the case of D-optimal, this involves examining how to best approximate the KL divergence based on posterior samples. In

addition, we examined how to modify the reward function to obtain policies that use minimal light exposure of the sample, to minimize photobleaching.

- We adopted an actor-critic approach for policy training. In addition, we realized that using curriculum learning, i.e., curating and changing the set of example problems presented to the RL algorithm is necessary to obtain good results in more challenging SPI scenarios, e.g., when the measurement noise it high.

To examine and validate our methodical developments, we designed two artificial aSPI scenarios:

“One-active-pixel" scenario: The explicit and simple prior is that the image **x** that consists of n pixels is all 0 except for a single pixel at an unknown location that has the value 1. The prior P(**x**) is thus discrete and assigns probability 1/n to all n vectors that have a single active pixel with value 1 and 0 to every other **x**. Solving the imaging problem corresponds to finding the location of the active pixel. Each measurement is given as the scalar product between **x** and a binary sensing pattern **h** whose pixels also consists of 0’s and 1’s (cf. Figure 4). In the noiseless case, the outcome of a measurement can therefore only be 1 if the active pixel in **x** is part of the active pixels in the sensing pattern or 0 otherwise. The posterior P(**x|y**) is an explicit, discrete probability distribution that assigns probability 1/m to all m locations that are still compatible with the measurements performed so far. Intuitively, it means that each SPI measurement asks the question "Is the active pixel part of this set of pixels?” and we are trying to find the location with as few questions as possible. The posterior distribution keeps track of the locations we have not discarded yet. Thereby, the aSPI problem is equivalent to a simple search problem. The naive strategy would be to query each pixel individually, which finds the active pixel after (n+1)/2 questions on average. The optimal strategy splits the set of possible locations in half in each turn (bisection) and uses that set to define the next sensing pattern. It takes log2(n) questions no matter where the active pixel is located and it thus much more efficient for large n. We designed this simple, intuitive scenario so we can validate whether our RL framework can recover this optimal strategy. Figure 9 shows an example of training two policies that rely on different ways to parameterize the probability distribution of the next sensing pattern. Another reason for choosing this scenario is that the whole posterior P(**x|y**) can be represented by a vector **p** of length n that in each entry stores the probability that the active pixel is at that location. This allows to examine how ideal policies that work with the full posterior information compare to those that work with reductions or samples of it, which is the only viable option in real-world aSPI problems. Furthermore, it allows to examine how different reward functions, and their approximations influence the policy training: Quantities that are hard to compute to a high accuracy in real-world aSPI problems such as the KL divergence are easy to compute in this scenario. Overall, the development of this scenario allows us to examine many important aspects of using RL for solving the aSPI problem in way that has not been possible for RL applied to other sOED problems such as in (Wang, Lucka, et al., 2024).

“MNIST” scenario: Next we consider a small scale (32 x 32 pixels) SPI scenario where the prior is given implicitly via a diffusion-based AI image generator that relies on a DNN trained on the MNIST data set of hand-written digits (Deng, 2012). Now, the resulting posterior can only be sampled from using MCMC (Chung et al., 2022; Zhao et al., 2023) and various quantities in the RL framework have to be reduced and approximated. To the best of our knowledge, this challenging combination of generative models for inverse problems with RL-based sOED has not been demonstrated so far. The scenario allows us to identify which components will scale to higher dimensional scenarios and which will have to be replaced by more efficient but more constrained ones. For instance, we needed to develop policies that map directly to probabilities over the Hadamard matrices as sensing pattern, which are commonly used in SPI. Figure 10 shows an example of how the reconstructed aSPI images improve while the policy that chooses the sensing pattern improves during training.

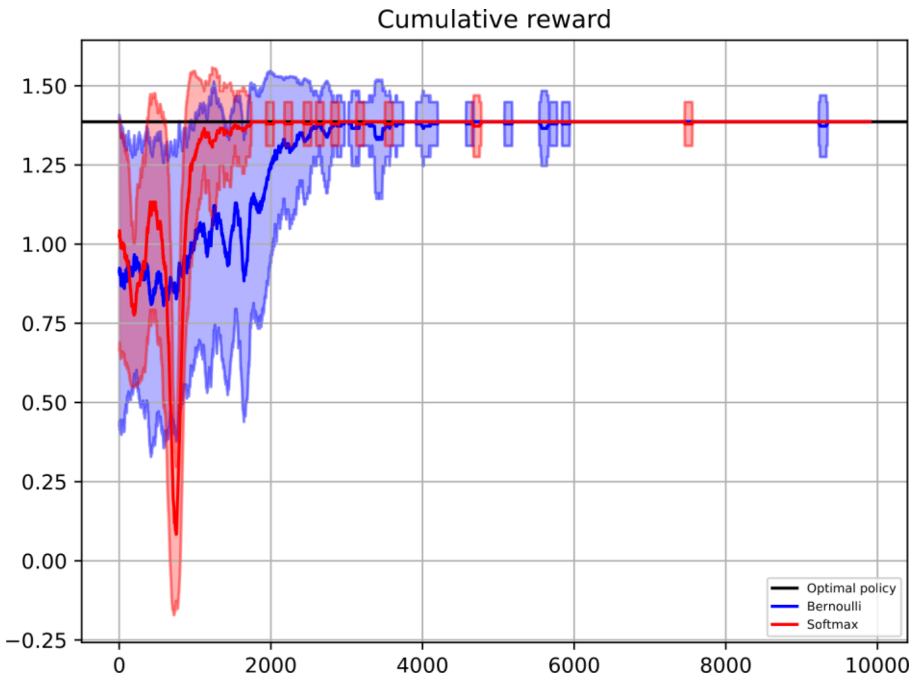

Figure 9. Example of training two sOED policies in the “one-active-pixel" scenario compared to the optimal bisection policy. On the horizontal axis, the training iterations (episodes) are plotted, on the vertical the cumulative reward. The policies differ in the way they parameterize the probability distribution of the next sensing pattern: While the “Softmax” policy maps to an individual probability for each of the 2^n possible pattern, the “Bernoulli” policy maps to the probability that each of the n sensing pixels is active. Both can match the performance of the optimal policy, but the “Bernoulli” policy will also scale to large problem sizes.

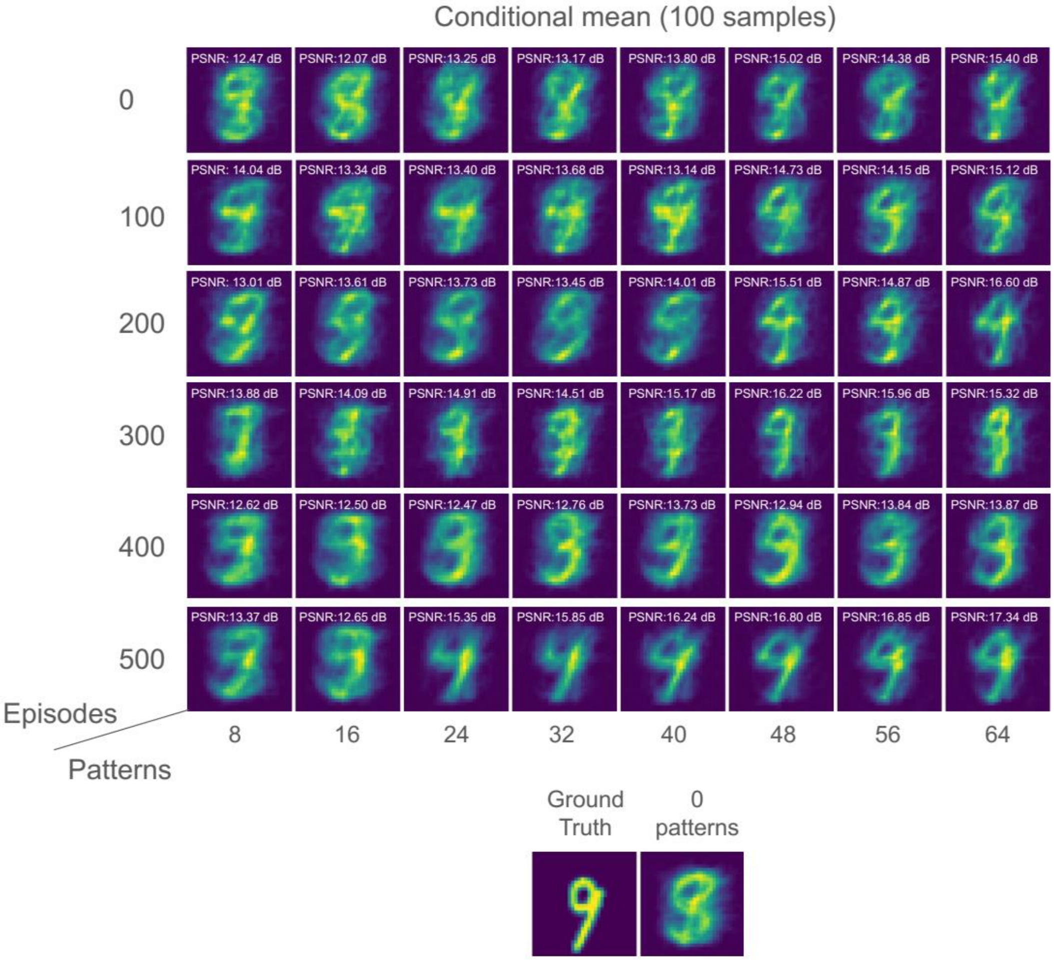

Figure 10. CM estimates for an example from the MNIST scenario. From top to bottom, the number of sensing pattern that the policy proposes is increased. From top to bottom, the time that the policy was trained by the RL framework is increased. By comparing bottom left to bottom right results, one can see that with more training the policy can adapt the sensing pattern to improve the reconstructed image. Given that the image size is 32x32 pixels, the compression ratio for 64 patterns results in 6.25%.

4 Conclusions and Outlook

The smart-2pm software is a good, flexible basis for the integration and control of the different hardware and software components of the SMART-2PM system that will be further developed continuously throughout the CONcISE project. The next concrete extension will be to integrate the adaptive optics module to incorporate ongoing developments on defocus correction, aberrations measurement, and higher aberration correction (see Deliverable 3.4). Furthermore, an interface between smart-2pm-napari and the aSPI software will be developed such that a proof-of-concept demonstration of an adaptive measurement can be carried out in the future.

Our pioneering work on RL-based algorithms with generative AI components shows promising potential for aSPI. In addition, we identified which components need further development to scale to large images. We are preparing a first paper publication about the results at the moment. There are several next developments of the aSPI algorithms directed at either improving their performance or making training and run times faster. For instance, their mapping onto graphics processing units (GPUs) will be improved to speed up both their training and run times and the efficiency of different algorithmic components will be examined and improved. The final challenge will be to transition from simulated data to real experimental data.

5 References

Burger, M., Hauptmann, A., Helin, T., Hyvönen, N., & Puska, J. P. (2021). Sequentially optimized projections in x-ray imaging. Inverse Problems, 37(7). https://doi.org/10.1088/1361-6420/ac01a4

Chung, H., Kim, J., McCann, M. T., Klasky, M. L., & Ye, J. C. (2022). Diffusion Posterior Sampling for General Noisy Inverse Problems. 11th International Conference on Learning Representations, ICLR 2023. https://arxiv.org/abs/2209.14687v4

De los Reyes, J. C., & Villacís, D. (2023). Bilevel Optimization Methods in Imaging. In Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging: Mathematical Imaging and Vision. https://doi.org/10.1007/978-3-030-98661-2_66

Deng, L. (2012). The MNIST database of handwritten digit images for machine learning research. IEEE Signal Processing Magazine, 29(6). https://doi.org/10.1109/MSP.2012.2211477

Duarte, M. F., Davenport, M. A., Takhar, D., Laska, J. N., Sun, T., Kelly, K. F., & Baraniuk, R. G. (2008). Single-pixel imaging via compressive sampling. IEEE Signal Processing Magazine, 25(2). https://doi.org/10.1109/msp.2007.914730

Eldar, Y. C., & Kutyniok, G. (2009). Compressed sensing: Theory and applications. In Compressed Sensing: Theory and Applications. https://doi.org/10.1017/CBO9780511794308

Foucart, S., & Rauhut, H. (2013). A mathematical introduction to compressive sensing. In Applied and Numerical Harmonic Analysis (Issue 9780817649470).

Gibson, G. M., Johnson, S. D., & Padgett, M. J. (2020). Single-pixel imaging 12 years on: a review. Optics Express, 28(19). https://doi.org/10.1364/oe.403195

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. In Nature (Issue 7553). MIT Press.

Kaipio, J. P., & Somersalo, E. (2005). Statistical and computational inverse problems. In Applied Mathematical Sciences (Switzerland) (Vol. 160). https://doi.org/10.1007/b138659

Li, S. E. (2023). Reinforcement Learning for Sequential Decision and Optimal Control. In Reinforcement Learning for Sequential Decision and Optimal Control. https://doi.org/10.1007/978-981-19-7784-8

Rainforth, T., Foster, A., Ivanova, D. R., & Smith, F. B. (2024). Modern Bayesian Experimental Design. Statistical Science, 39(1). https://doi.org/10.1214/23-STS915

Ruthotto, L., Chung, J., & Chung, M. (2018). Optimal experimental design for inverse problems with state constraints. SIAM Journal on Scientific Computing, 40(4). https://doi.org/10.1137/17M1143733

Ryan, E. G., Drovandi, C. C., Mcgree, J. M., & Pettitt, A. N. (2016). A Review of Modern Computational Algorithms for Bayesian Optimal Design. International Statistical Review, 84(1). https://doi.org/10.1111/insr.12107

Shen, W., & Huan, X. (2023). Bayesian sequential optimal experimental design for nonlinear models using policy gradient reinforcement learning. Computer Methods in Applied Mechanics and Engineering, 416. https://doi.org/10.1016/j.cma.2023.116304

Shen, Z., Wang, Y., Wu, D., Yang, X., & Dong, B. (2022). Learning to scan: A deep reinforcement learning approach for personalized scanning in CT imaging. Inverse Problems and Imaging, 16(1). https://doi.org/10.3934/ipi.2021045

Stuart, A. M. (2010). Inverse problems: A Bayesian perspective. Acta Numerica, 19, 451–559. https://doi.org/10.1017/S0962492910000061

Wang, T., Florian, V., Schielein, R., Kretzer, C., Kasperl, S., Lucka, F., & van Leeuwen, T. van. (2024). Task-Adaptive Angle Selection for Computed Tomography-Based Defect Detection. Journal of Imaging, 10(9), 208. https://doi.org/10.3390/jimaging10090208

Wang, T., Lucka, F., & van Leeuwen, T. (2024). Sequential Experimental Design for X-Ray CT Using Deep Reinforcement Learning. IEEE Transactions on Computational Imaging, 10, 953–968. https://doi.org/10.1109/TCI.2024.3414273

Zhao, Z., Ye, J. C., & Bresler, Y. (2023). Generative Models for Inverse Imaging Problems: From mathematical foundations to physics-driven applications. IEEE Signal Processing Magazine, 40(1). https://doi.org/10.1109/MSP.2022.3215282