Supporting Information

Pessimistic asynchronous sampling in high-cost Bayesian optimization

Amanda A. Volk[1,2]*, Kristofer G. Reyes[3,4], Jeffrey G. Ethier[2], Luke A. Baldwin[2]*

[1]*National Research Council, Washington, District of Columbia 20001, USA*

[2]*Materials and Manufacturing Directorate, Air Force Research Laboratory, Wright–Patterson Air Force Base, Dayton, OH 45433, USA*

[3]*Department of Materials Design and Innovation, University at Buffalo, Buffalo, NY 14260, USA*

[4]*Brookhaven National Lab, Computational and Data Science Directorate, Upton, NY 11973, USA*

Table of Contents

Section S.1 – Additional Simulation Results

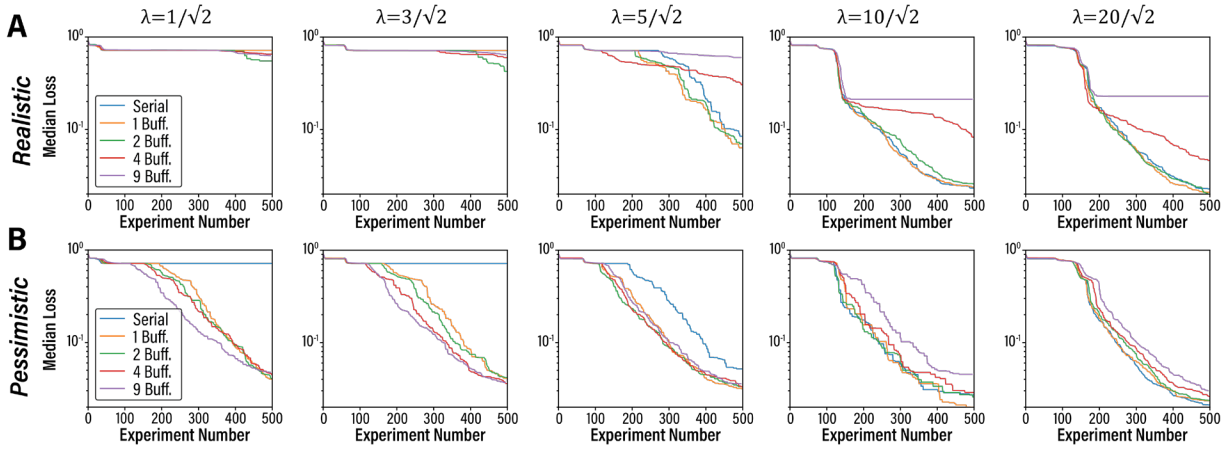

Figure S.1. Simulation results of pessimistic and realistic asynchronous decision policies on five-dimensional *TriPeak*. The median loss across all 200 randomized simulated campaigns as a function of the number of experiments for the (A) realistic and (B) pessimistic prediction policies, with (columns) increasing exploration constants, lambda values, in the upper confidence bounds decision policy. The no buffer replicates were repeated for each of the policies and lambda values.

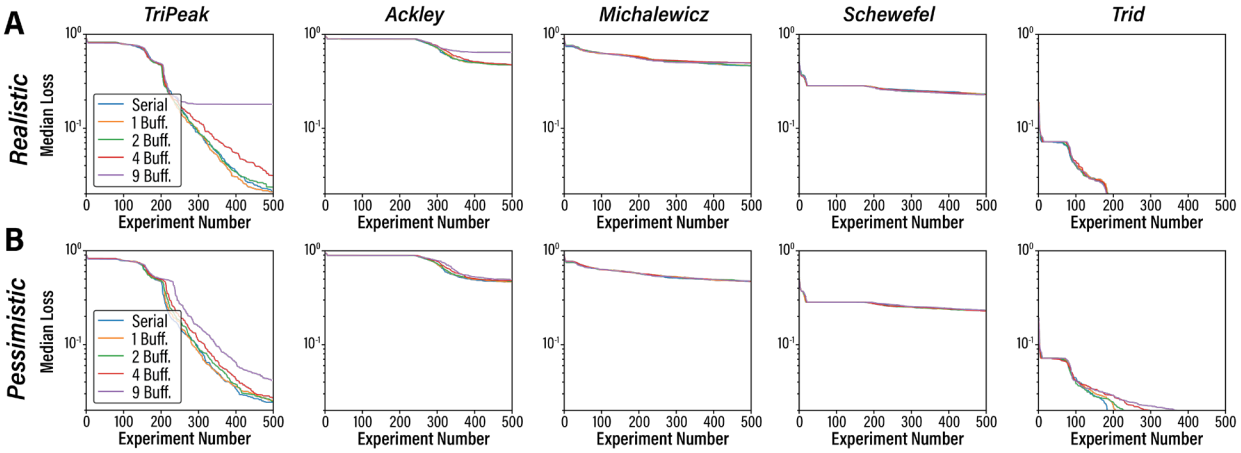

Figure S.2. Logarithmic lambda simulation results of pessimistic and realistic asynchronous decision policies on five-dimensional *TriPeak*, *Ackley*, *Michalewicz*, *Schewefel*, and *Trid*. The median loss across all 200 randomized simulated campaigns as a function of the number of experiments for the (A) realistic and (B) pessimistic prediction policies, with a logarithmic upper confidence bounds decision policy. The exploration constant, lambda, is defined as $\lambda = 2\ln\left(\frac{Dt^2\pi^2}{3d}\right)$, where D is the dimensionality of the input space, t is the current experiment number, and d is equal to 1. The no buffer replicates were repeated for each of the policies.

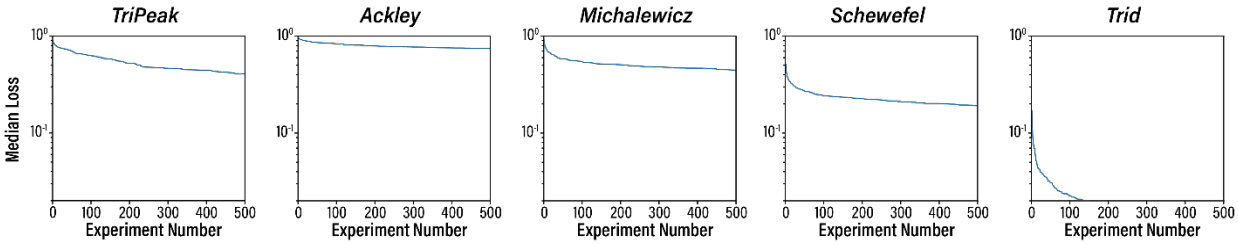

Figure S.3. Random serial sampling control group on five-dimensional *TriPeak*, *Ackley*, *Michalewicz*, *Schewefel*, and *Trid*. The median loss across all 200 randomized simulated campaigns as a function of the number of experiments.

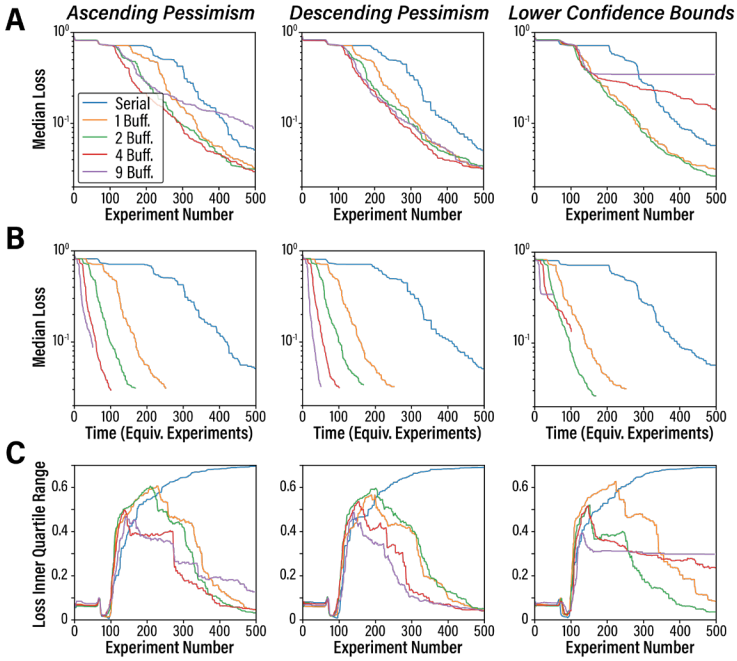

Figure S.4. Simulation results of three asynchronous decision policies utilizing different prediction strategies on five-dimensional *TriPeak*. The median loss across all 200 randomized simulated campaigns as a function of (A) the number of experiments and (B) the effective optimization time relative to a single experiment and (C) the inner quartile range of the loss as a function of experiment number across the four decision policies, (first column) ascending pessimism, (second column) descending pessimism, and (third column) lower confidence bound based prediction. Serial replicates were repeated for each of the four policies.

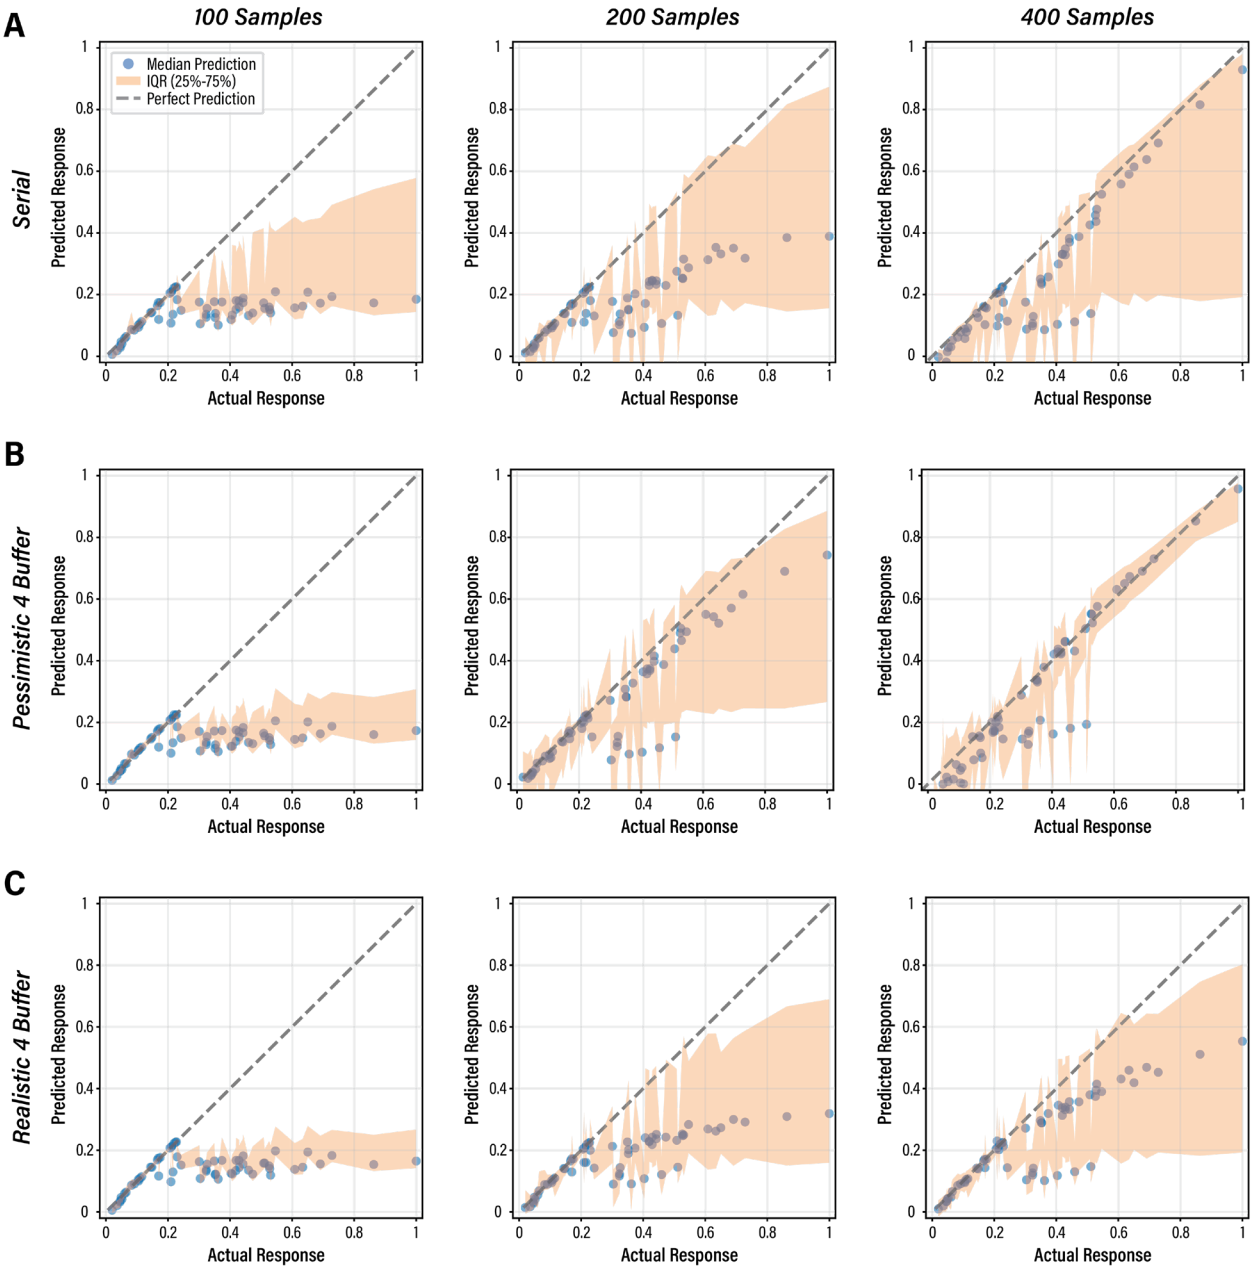

Figure S.5. Test set prediction regression on five-dimensional *TriPeak* at different sample numbers in the campaign. The median and inner quartile range of the model predicted values for (A) serial, (B) four buffer pessimistic asynchronous, and (C) four buffer realistic sampling after (left) 100, (middle) 200, and (right) 400 samples in the optimization campaign, as a function of the actual ground truth value. The test set is the same for all panels and comprised of the optimal condition plus 100 randomly selected conditions with grouped stratification across the ground truth response range. It should be noted that most purely randomly selected conditions would fall below 0.3. All asynchronous models included the prediction values in the model training. All figures are the result of 200 campaign replicates of each selection algorithm.

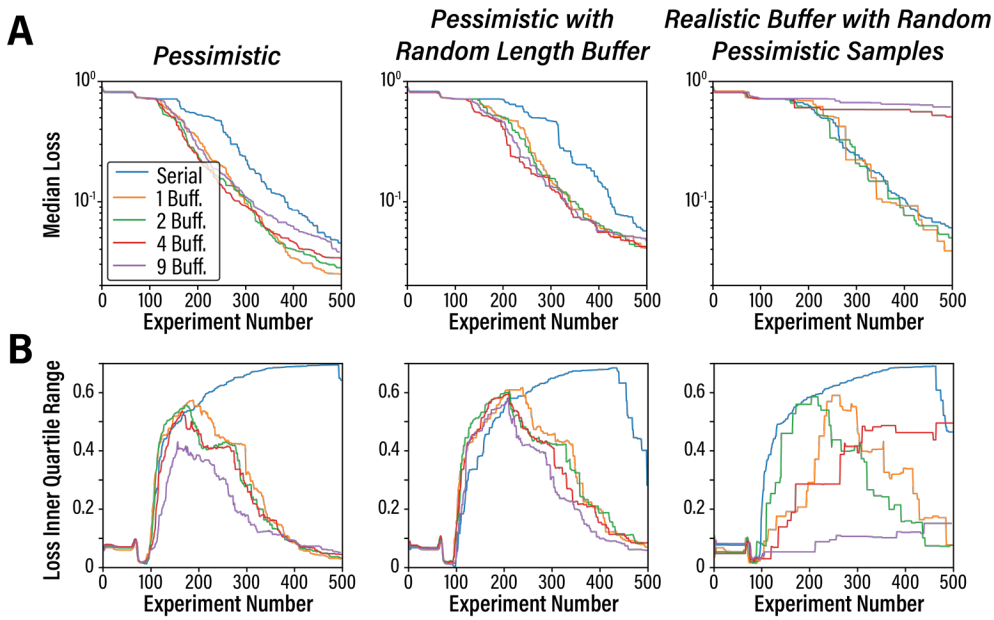

Figure S.6. Simulation results of three asynchronous decision policies utilizing different pessimism-based prediction strategies on five-dimensional *TriPeak*. The median loss across all 140 to 200 randomized simulated campaigns as a function of (A) the number of experiments and (B) the inner quartile range of the loss as a function of experiment number across the four decision policies, (first column) pessimistic, (second column) pessimistic with a random length buffer, and (third column) realistic buffer with randomly sampled pessimism. The random length buffer was selected uniformly between 0 and the specified buffer length. The pessimistic samples were formed by randomly assigning a pessimistic value to a number of samples equal to the current buffer length. Serial replicates were repeated for each of the four policies.

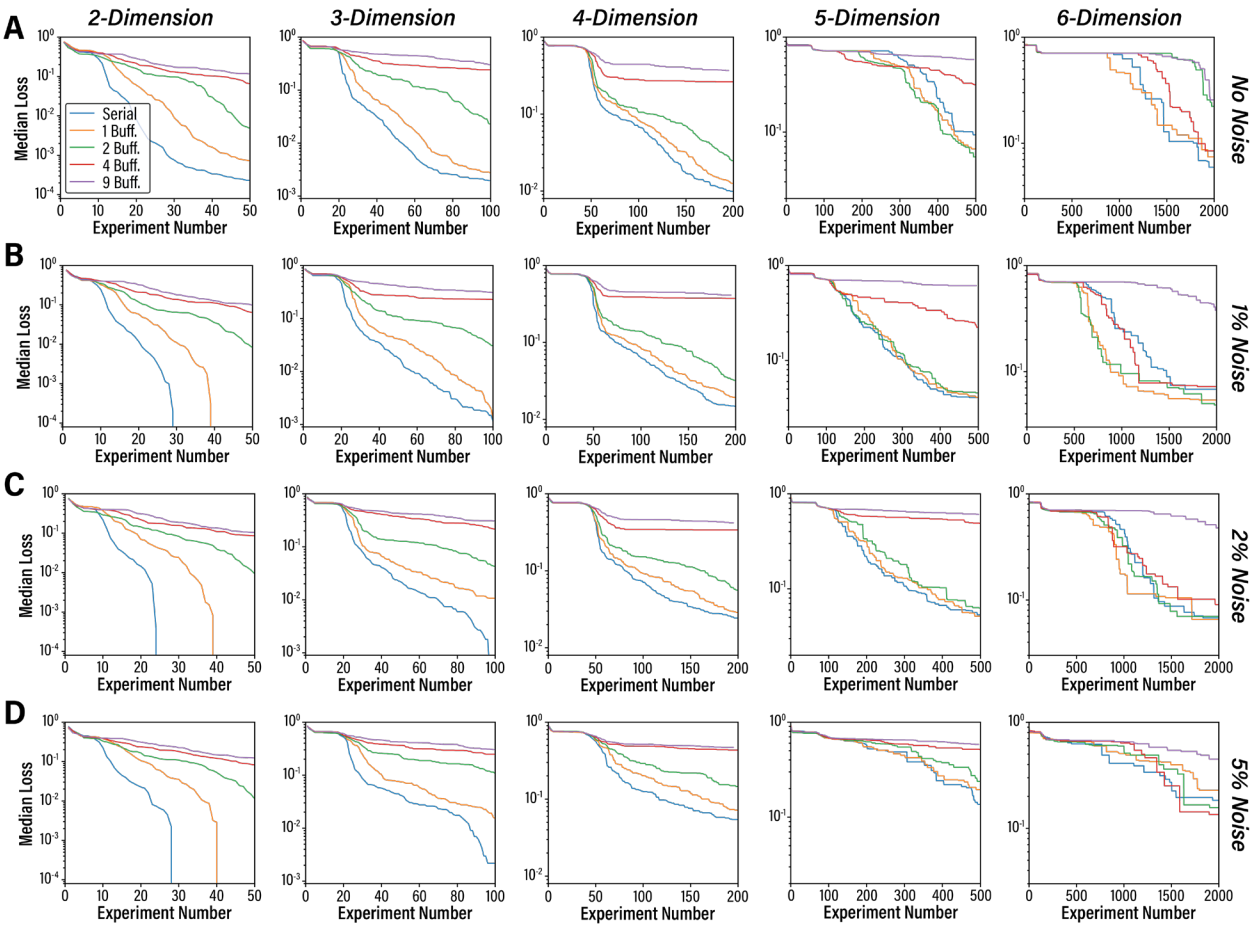

Figure S.7. Simulation results of realistic decision policies on *TriPeak* at different dimensionalities and noise levels. The median loss across all randomized simulated campaigns as a function of the number of experiments across (columns) two, three, four, five, and six-dimensional ground truth spaces with (A) 0%, (B) 1%, (C) 2%, and (D) 5% noise. Sampling noise is applied by randomly sampling from a normal distribution with a standard deviation equal to the specified noise value and adding the noise sample to the ground truth output. Reduced replicate counts were implemented due to computational resource constraints. The 2-, 3- and 4-dimensional plots are the result of 180 replicates. The 5-dimensional plots are the result of 68 replicates. The 6-dimensional plots are the result of 15 replicates. The loss is calculated from the noiseless ground truth and does not reflect the values sampled from the ground truth during each trials campaign

Section S.2 – Supplementary Asynchronous Function Descriptions

Pessimistic with Random Length Buffer - The pessimistic with random length buffer policy is designed to have a dynamically changing buffer of pessimistic predictions. The policy operates by randomly selecting a buffer length from the range specfied by the maximum buffer length. The policy then selects parameter sets corresponding to each buffer addition, continously padding pessimistic values with each parameter set. After completing the buffer padding, the policy resamples a new random buffer length. If the new buffer length is greater than the current buffer length, then additional conditions with pessimistic predictions are padded to the buffer. If the new buffer length is less than the current buffer length, then the sampled responses for each of the predictions are inserted into the buffer, thereby shrinking the buffer. This process continuously expands and contracts the buffer length to fit the selected value throughout the optimization.

Realistic Buffer with Random Pessimistic Samples - The realistic buffer with random pessimistic samples policy is designed to differentiate the performance of pessimism in the buffer versus general pessimism. The policy operates first using the realistic policy described in the main text. Then for each condition selection, a set of already sampled conditions in the training data set are forced to have pessimistic predictions. Experiments are then selected using the realistic buffer, the training data, and the pessimistic data. New conditions are sampled for pessimstim on each experiment selection.