

# The future looks dark: improving high contrast imaging with hyper-parameter optimization for data-driven predictive wavefront control

J. Fowler<sup>a</sup>, Rebecca Jensen-Clem<sup>a</sup>, Maaike A. M. van Kooten<sup>b</sup>, Vincent Chambouleyron<sup>a</sup>, and Sylvain Cetre<sup>c</sup>

<sup>a</sup>Department of Astronomy & Astrophysics, University of California, Santa Cruz, CA, USA

<sup>b</sup>National Research Council Canada, Herzberg Astronomy and Astrophysics Research Center, Victoria, Canada

<sup>c</sup>W. M. Keck Observatory, Waimea, Hawaii, USA

## ABSTRACT

The direct imaging and characterization of exoplanets requires extreme adaptive optics (XAO), achieving exquisite wavefront correction (upwards of 90% Strehl) over a narrow field of view (a few arcseconds). For these XAO systems the temporal error is often a leading term in the error budget, wherein the wavefront evolves faster than the lag between wavefront sensing and control. For atmospheres with high-velocity wind layers, this can result in a wind-driven halo in the coronagraphic dark-zone, limiting sensitivity to faint, close-in companions. The AO system's lag-time is often limited by the wavefront sensor exposure time, especially in the case of fainter guidestars. Predictive control mitigates the temporal error by predicting the shape of the wavefront by time the system correction is applied. One such method of prediction is empirical orthogonal functions (EOF), wherein previous states in the wavefront sensor history are used to learn linear correlations with a minimization problem. This method has been demonstrated on-sky at Subaru/SCExAO and Keck/NIRC2, but has yet to be optimized. With this work as a starting point, we explore the optimal filter hyper-parameter space for implementing EOF on-sky, study its stability under varying atmospheric parameters, and discuss future paths for facilitization of predictive control. This work not only offers a pathway to optimize Keck and Subaru observing, but also acts as a pathfinder for predictive control methods with extremely large telescopes.

**Keywords:** predictive control, temporal error, empirical orthogonal functions (EOF), hyper-parameter optimization, simulated annealing, Keck, Subaru

## 1. INTRODUCTION

Predictive wavefront control, first suggested by Dessenne in 1997,<sup>1</sup> has been explored through many implementations over the years, ranging from classic and predictive Linear Quadratic Gaussian controllers,<sup>2,3</sup> neural network approaches,<sup>4,5</sup> through reinforcement learning,<sup>6–9</sup> and through classic minimization problems;<sup>10–13</sup> for a more complete review of predictive methods, see (Fowler, 2023).<sup>14</sup> Predictive controllers consistently show promising performance with on-sky testing,<sup>15–17</sup> and in particular, empirical orthogonal functions (EOF) at Keck has shown up to 3x improvement in contrast at separations of 3–7  $\lambda/D$ , and 2x for 3  $\lambda/D$ .<sup>18</sup>

However, no facilitated implementations of predictive control are running on-sky and little has been published in the way of optimizing its performance for regular usage under varying atmospheric conditions. The implementation at Keck<sup>18</sup> explored the optimal mixing between an integrator and predictor and optimal gain for the high order correction, but built a predictive filter with a consistent set of hyper-parameters. Similarly, one work<sup>4</sup> explored the optimal amount of data for an EOF filter to train on with Subaru/SCExAO telemetry, but predictive control at SCExAO is run on-sky with a static set of filter hyper-parameters (private communication, SCExAO team).

---

Further author information: (Send correspondence to J. F.)

J. F.: E-mail: [jumfowle@ucsc.edu](mailto:jumfowle@ucsc.edu)

In the effort to move predictive control from an on-sky demonstration to a facilitated mode, we must address how to consistently set up the predictive controller for varying nightly conditions and make recommendations for a hyper-parameter optimization strategy during observing nights. Here, we address the tuning of individual hyper-parameters in the predictive filter as the first step towards developing such an on-sky strategy.

## 2. TUNEABLE HYPER-PARAMETERS IN THE EOF FILTER

Empirical orthogonal functions (EOF)<sup>10</sup> builds a linear filter  $\mathbf{F}$  by assembling  $\mathbf{l}$  previous states of wavefront sensor information (or history vectors  $\mathbf{h}$ ) into a training matrix  $\mathbf{D}$ , and comparing those to known future states  $\mathbf{P}$ . The history vector contains  $\mathbf{n}$  frames of wavefront sensor data (each with  $\mathbf{m}$  points for each mode) flattened into a single vector; this is shown visually in Figure 1. The filter matrix encodes the wind velocities and other linear system errors.

The logic to build a filter matrix is shown below; the matrix inversion is done with a least-squares inverse, and uses  $\alpha$  as a regularization parameter.

$$\min \|\mathbf{D}^T \mathbf{F}^T - \mathbf{P}^T\|^2 \quad (1)$$

$$\mathbf{F} = ((\mathbf{D}^T)^\dagger \mathbf{P}^T)^T \quad (2)$$

$$\mathbf{F} = \mathbf{P} \mathbf{D}^T (\mathbf{D} \mathbf{D}^T + \alpha \mathbf{I})^{-1} \quad (3)$$

$$\text{prediction} = \mathbf{F} \mathbf{h} \quad (4)$$

At every iteration of the real-time-controller, we multiply the filter matrix  $\mathbf{F}$ , by the history vector  $\mathbf{h}$  to predict an iteration one lag-time in the future. This full logic is visualized in Figure 2.

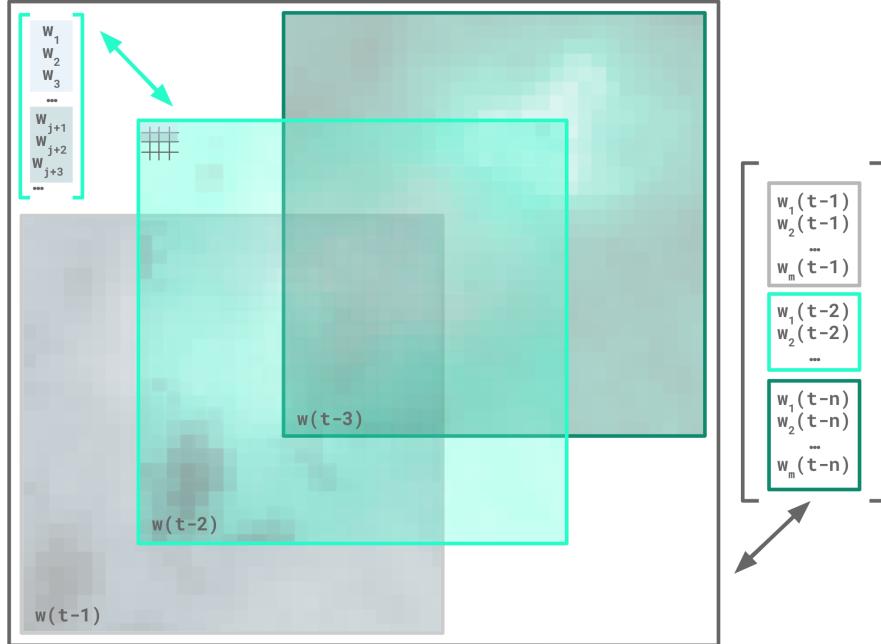


Figure 1. The history vector contains  $\mathbf{n}$  frames of wavefront sensor data flattened into a single vector.

The filter matrix provides a good correction while the atmosphere is driven by the same linear forces on which it was trained, therefore we must plan to update the filter matrix every few minutes or have a rolling update of the filter matrix for on-sky operations. To generate a predictive filter, we consider the number of frames in our training data  $\mathbf{l}$ , the number of frames in our history vector  $\mathbf{n}$ , and the regularization constant  $\alpha$  as tuneable hyper-parameters of the EOF filter. In simulation, our matrix is well conditioned, and we can run with an  $\alpha = 1$ , so the simulations presented in the rest of this work aim to explore optimal values for  $\mathbf{l}$  and  $\mathbf{n}$ .

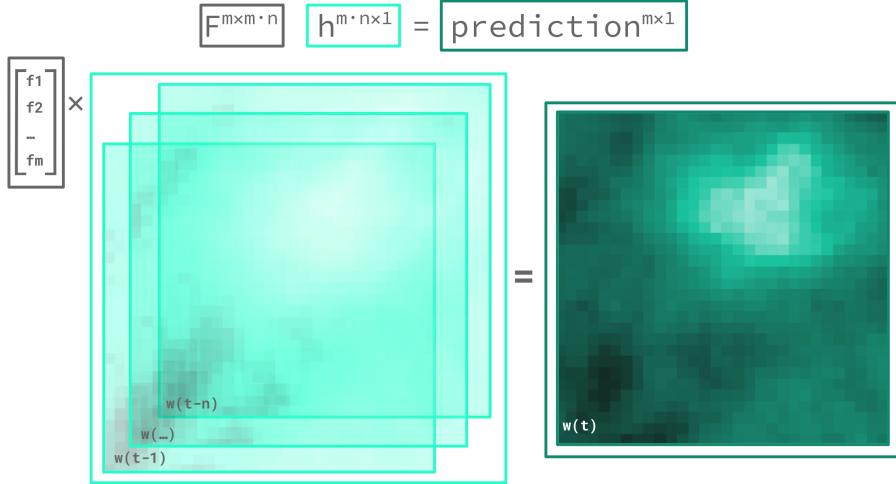


Figure 2. To predict a future iteration with  $\mathbf{m}$  modes, we multiply a filter matrix that is  $\mathbf{m}$  by  $\mathbf{m} \cdot \mathbf{n}$  (with  $\mathbf{m}$  modes and  $\mathbf{n}$  points in time) by the current history vector  $\mathbf{h}$ .

### 3. OPTIMIZING FILTER HYPER-PARAMETERS IN SIMULATION

We simulate the performance of EOF with varying hyper-parameters  $\mathbf{l}$  and  $\mathbf{n}$  as applied to atmospheres with varying windspeeds and Fried parameters. We simulate a classic single-layer Kolmogorov<sup>19</sup> phase screen with the Taylor frozen flow approximation<sup>20</sup> with a given Fried parameter  $r_0$  and windspeed for a 10 meter telescope at a resolution of 48x48 pixels. We subtract piston, tip, and tilt modes, and apply a circular aperture. We sample and correct the atmosphere at 1kHz, with a lag of two time steps (2 ms), with perfect wavefront sensing and perfect correction – so the error is indicative only of error from the control itself.

For each hyper-parameter grid search, we build a predictive EOF filter with varying numbers of frames of training data ( $\mathbf{l}$ ) and varying history vector lengths ( $\mathbf{n}$ ), and build a heat map of the the median performance in root mean square (RMS) error. We compute a residual by subtracting the prediction of the full phase-screen at each iteration from the known injected phase at each iteration, and take an RMS of that difference at each point in time, leaving us with the RMS for each of 10000 iterations. We then take the median of that time-series distribution for the final value plotted and used for optimization. Figures 3 and 4 show examples of differing  $r_0$  with the same windspeed and differing windspeed with the same  $r_0$  respectively. We note that the optimal history vector length changes based on different turbulence conditions.

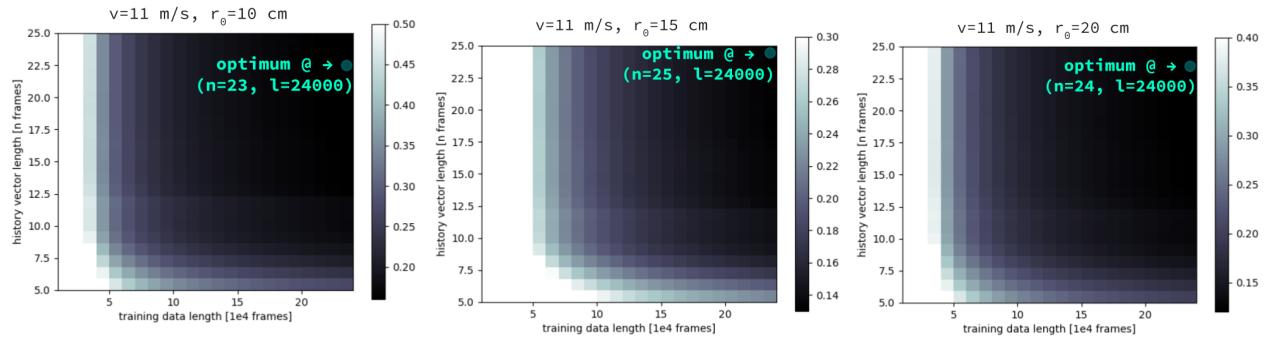


Figure 3. Heatmap of EOF predictor performance (in median RMS error in nm) as training data length  $\mathbf{l}$  and history vector length  $\mathbf{n}$  vary. Left:  $r_0$  of 10 cm. Middle:  $r_0$  of 15 cm. Right:  $r_0$  of 20 cm. All of the above simulations have a windspeed of 11 m/s. Note that the optimal history vector length changes based on  $r_0$ .

While these simulations show that the optimal history vector length changes, simulations consistently show that the optimal amount of training data is the maximum amount available. We expect that predictor perfor-

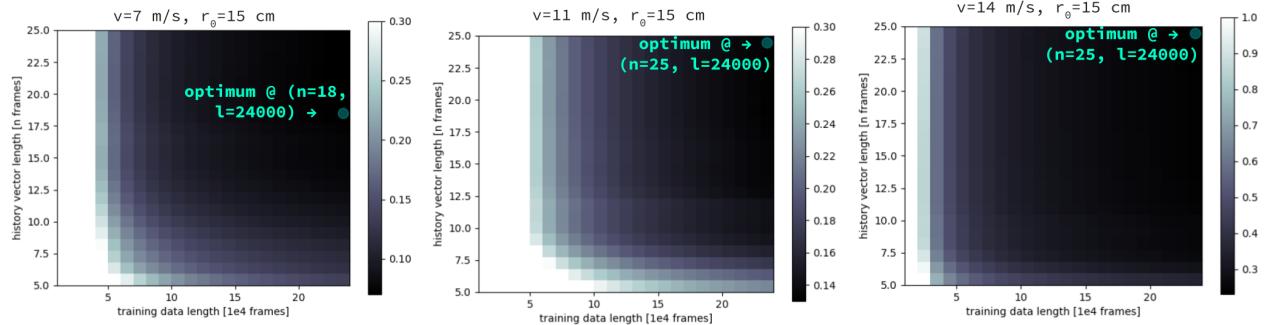


Figure 4. Heatmap of EOF predictor performance (in median RMS error in nm) as training data length  $l$  and history vector length  $n$  vary. Left: windspeed of 7 m/s. Middle: windspeed of 11 m/s. Right: windspeed of 14 m/s. All of the above simulations have an  $r_0$  of 15 cm. Note that the optimal history vector length is different for the slowest atmosphere. Given the two right-most plots find an optimal in the corner of our grid, they may also find different optimal values given an expanded search space.

mance due to different lengths of training data in the filter matrix would be more impactful on real data, where including more frames of data introduces more noise into the training (our simulations are lacking measurement noise), and when the velocity layer will vary over some timescale.

#### 4. ON-THE-FLY OPTIMIZATION WITH SIMULATED ANNEALING

If the desired strategy for EOF implementation is to do on-sky hyper-parameter tuning, building a full grid of filter hyper-parameter runs is likely not feasible. For these initial grids, we used 10 seconds of data and built a filter based on up to 24 seconds of training data; even for a system optimized to run in real time, to do this 600 times would take  $\sim 6$  hours. We present an alternative way to explore this hyper-parameter space more optimally with simulated annealing, a probabilistic method for finding the global optimum of a non-differentiable function.

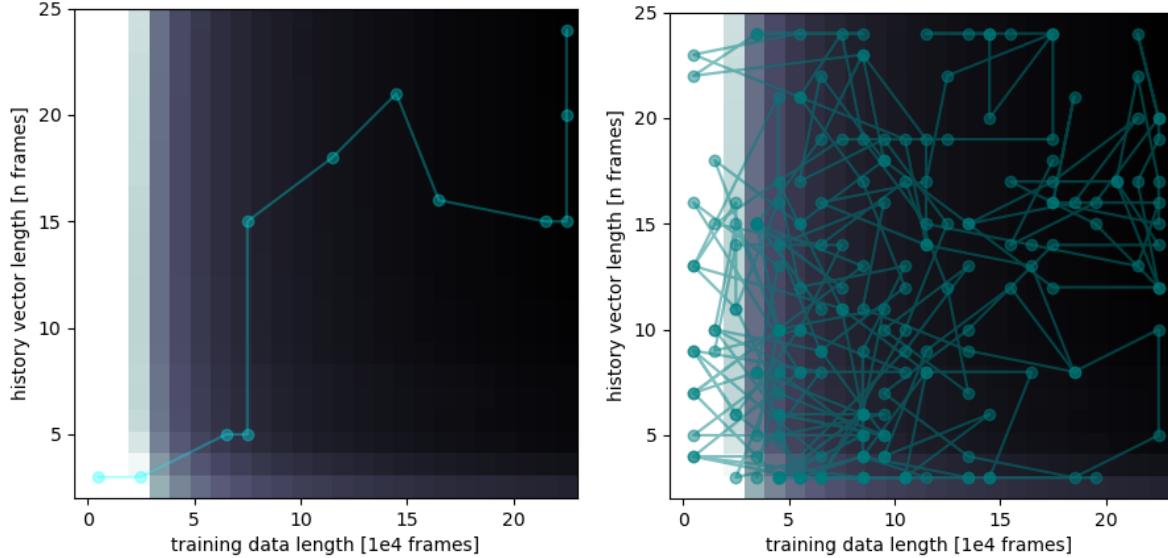


Figure 5. Simulated annealing samples many values, but it does not accept every step. This is the run of a simulated annealing algorithm applied to a grid with a windspeed of 14 m/s and an  $r_0$  of 15 cm, chosen because the grid had the most variation. Left: Entire walk of the simulated annealing to find the optimal solution. Right: Every point sampled (whether or not it was accepted) over the 400 steps.

Simulated annealing<sup>21</sup> is a probabilistic way to explore a discrete parameter space. The algorithm picks a random point of a given step size away from the current point and accepts it if the point provides a better solution, or if a Metropolis criterion is met with a given temperature. This process is named after metal annealing, where metal is heated and cooled strategically for optimal performance. Pseudo-code of this algorithm and more discussion of the Metropolis criterion are presented in Appendix A. Figure 5 shows how the algorithm walks through one of the hyper-parameter search grids shown above.

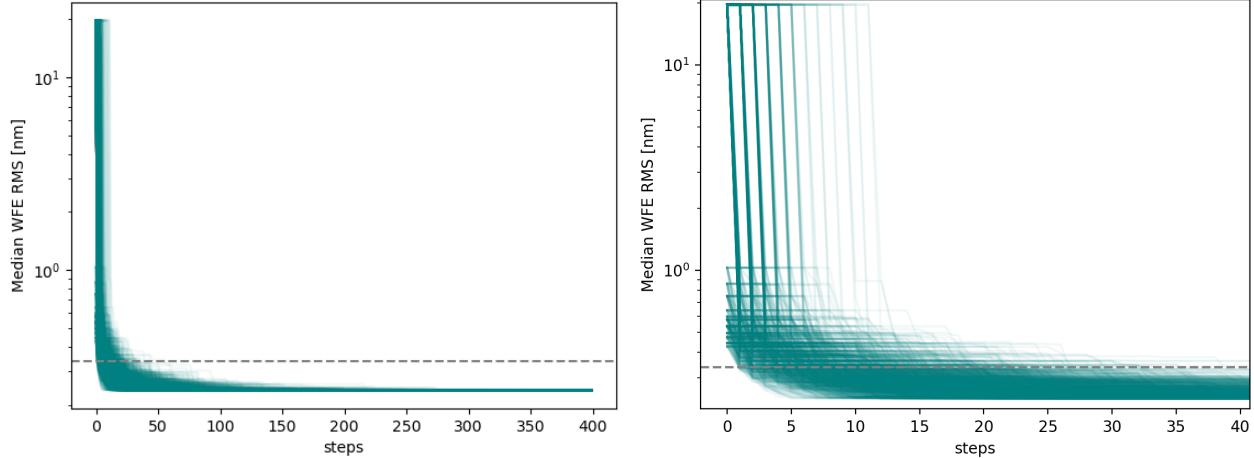


Figure 6. Simulated annealing is applied to the previously shown dataset, with a windspeed of 14 m/s and an  $r_0$  of 15 cm. The input temperature is 119 and the step size is 5. The dashed line indicates when the algorithm reaches within 0.1 nm of the optimal solution. Left: full run. Right: closer detail on the first few iterations.

We applied simulated annealing to the previously shown filter hyper-parameter grid for a windspeed of 14 m/s and an  $r_0$  of 15 cm, and found that the algorithm could repeatably find within 0.1 nm of the optimal value within 10 steps and repeatably recover the exact solution within 400 steps. For an on-sky strategy, we expect we could use closer to 10 steps, to find a very good if not exactly optimal solution. However, how quickly the simulated annealing algorithm converges to near its minimum will be dependent on the performance of each filter over a set of conditions, so testing with lab-data and on-sky telemetry will be needed before specific strategies can be recommended.

Figure 6 shows 1000 runs of simulated annealing on the same data set. This example was run with a step size of 5 (which maps to a single history vector frame and 50000 training data frames), and a temperature of 119. Initial testing found that higher temperatures (100-200) worked better for this data set but no quantitative optimization was done for temperature and step size in this case, due to how well the algorithm performed with these parameters. A strength of simulated annealing is its effective application to problems without significant optimization or previous knowledge of a parameter space.

## 5. PRELIMINARY LAB RESULTS

To further this work we plan to examine  $\alpha$  as third hyper-parameter to optimize, as well as explore how more realistic conditions (e.g., imperfect noisy data) impact the optimal training data length. To do this, we plan to run similar grids of optimization on the SEAL (Santa Cruz Extreme AO Laboratory) testbed.<sup>22</sup> Along with more realistic conditions, SEAL has a coronagraphic branch with a vector vortex coronagraph;<sup>23</sup> moving to the bench will allow us to optimize hyper-parameters based on coronagraphic contrast as well as wavefront error.

Figures 7 and 8 show a preliminary implementation of EOF predictive control on the SEAL testbed, comparing the performance of EOF (with a filter built with a 5 frame history vector and 60000 frames of training data) to an integrator with a gain of 0.2. We simulate a single layer atmosphere with a  $D/r_0$  of 100 (e.g., a 10 meter telescope with an  $r_0$  of 10 cm) and a pupil crossing time of 1s (e.g., 10 m/s wind layer) and sample its phase every millisecond. We apply turbulence with the same  $D/r_0$  in the lab, but with the amplitude scaled down

to the linear regime of the Boston MEMS DM (within  $2\pi$  of phase), and then applied and corrected with the MEMS DM for 10000 iterations.

We use a high speed Thorlabs Shack-Hartmann Wavefront Sensor with 640 slopes and a Boston Micromachines MEMS deformable mirror with 24 actuators across the pupil. Initial results show improvement over a classic integrator both in reduction of wavefront error (Figure 8), and in the the focal plane PSF images (Figure 7).

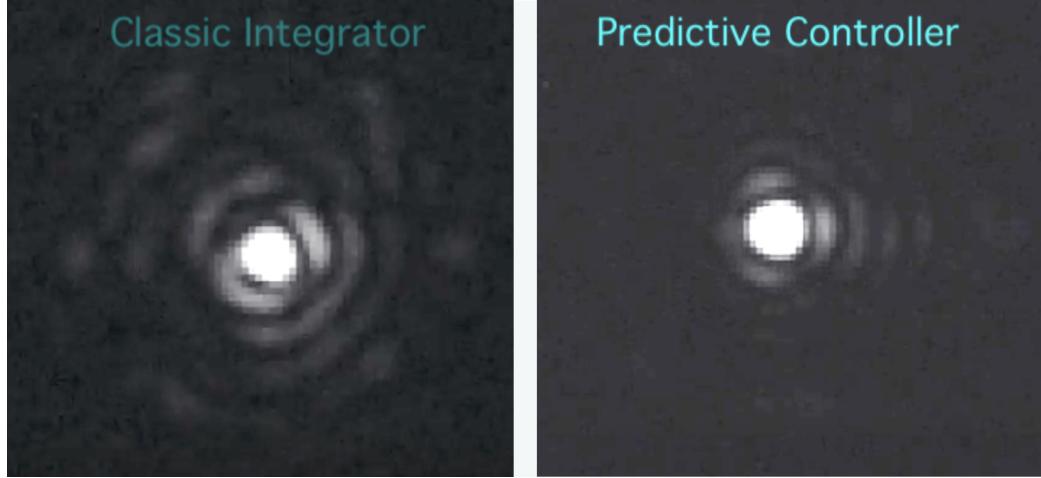


Figure 7. PSFs of predictive control and a classic integrator side-by-side on the same color-scale. Note: these are *preliminary* results without robust image calibration and serve only as an initial visual comparison.

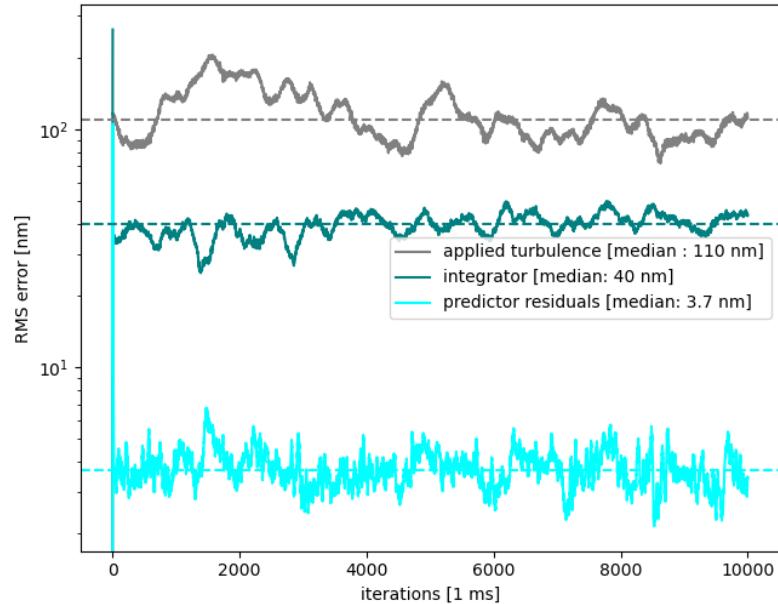


Figure 8. Bench results of an EOF controller as compared to a classic integrator. RMS wavefront is estimated from the SHWFS residuals as converted to DM space. Input turbulence is a single wind layer with a pupil crossing time of 1s and an atmosphere with a  $D/r_0$  of 100, with amplitude scaled down the linear  $2\pi$  dynamic range of the Boston MEMS.

Plans for future lab testing include recreating these hyper-parameter grid searches in an environment more indicative of a true AO system; we will apply turbulence with a Meadowlark spatial light modulator of 1116 pixels across, and correct with a woofer/tweeter system of a 97 actuator ALPAO deformable mirror and the aforementioned Boston MEMS. This will give us the ability to mimic more realistic effects of measurement error (with turbulence applied as a much finer spatial scale than we can correct.)

## 6. CONCLUSIONS

In conclusion, we explore future facilitization of predictive wavefront control by looking at optimal hyper-parameters for the training data and history vector length of the predictive filter. We find that in simulation, optimal filter hyper-parameters vary based on the atmospheric turbulence, and expect that we will need to optimize on-sky. We explore simulated annealing as a way to probe filter hyper-parameters more efficiently, and find it can repeatably converge to within 0.1 nm RMS of the optimal solution within 10 iterations.

We present preliminary lab results, that compare EOF to a classic integrator, and see promising performance from EOF. This lab implementation is a precursor to future exploration of  $\alpha$  in the EOF matrix inversion, as well as an opportunity to study the impact of these corrections on the coronagraphic dark hole. We also hope to test this optimization on-sky with Keck/NIRC2 or Subaru/SCExAO.

Facilitating predictive control for regular on-sky use will not only improve current high contrast imaging for large telescopes, but will show technology maturation for its use in upcoming extremely large telescopes. When predictive control is in use, the future of our coronagraphic dark hole is looking dark.

## APPENDIX A. SIMULATED ANNEALING ALGORITHM

Here we further describe the Metropolis criterion, and present Python pseudo-code of the simulated annealing algorithm to demonstrate the application of this algorithm to our data. For each iteration, the algorithm picks a new point, evaluates if that point is a better solution, and if not decides to keep the new point only if it meets the Metropolis criterion, which is met by comparing a randomly generated number to a function built from a cooling input temperature and how far the evaluation of the new point is from the previous. We note that simulated annealing and the Metropolis criterion make this algorithm vulnerable to getting trapped in a deep local minimum, and plan to study this impact with more realistic lab data. For a given iteration  $i$ :

$$\text{cooled temperature} = \frac{\text{temperature}}{i + 1} \quad (5)$$

$$\text{difference} = \text{evaluation}(i - 1) - \text{evaluation}(i) \quad (6)$$

$$\text{metropolis} = \exp\left(\frac{\text{difference}}{\text{cooled temperature}}\right) \quad (7)$$

This logic runs for a given number of steps, but on-sky this could also be a cutoff based on how much the algorithm is still improving the solution or some minimum allowable wavefront error.

```

def simulated_annealing(objective, bounds, n_iterations, step_size, temp):
    # start with an initial guess
    best = preset or random guess
    # evaluate it and set it as a starting point
    best_eval = objective(best)
    current, current_eval = best, best_eval
    # for every iteration
    for i in range(n_iterations):
        # take a step
        candidate = move from current
        # evaluate the new candidate
        candidate_eval = objective(candidate)
        # check for new best solution
        if candidate_eval < best_eval:
            # store new best point
            best, best_eval = candidate, candidate_eval
        # otherwise, use metropolis to see if we keep the new point
        # or stay with the previous guess
        # this determines where the guess next iteration starts from
        diff = candidate_eval - current_eval
        t = temp/(i + 1)
        # calculate metropolis acceptance criterion
        metropolis = np.exp(-diff / t)

        # check if we should keep the new point
        if diff < 0 or random number < metropolis:
            # store the new current point
            current, current_eval = candidate, candidate_eval

    # after however many iterations return the best solution
    return best, best_eval

```

Figure 9. Psuedo-code of simulated annealing algorithm, heavily influenced by [Machine Learning Mastery](#). The details of making consistent integer steps over a two dimensional grid make this algorithm more complex, so we have chosen to present representative code as opposed to the literal implementation.

## ACKNOWLEDGMENTS

This work was supported by NSF ATI Grant 2008822. Thank you to the folks at Keck Observatory and members of the SCExAO team for discussing their implementations of EOF.

## REFERENCES

- [1] Dessenne, C., Madec, P. Y., and Rousset, G., “Modal prediction forclosed-loop adaptive optics,” *Optics Letters* **22**, 1535–1537 (Oct. 1997).
- [2] Sivo, G., Kulcsár, C., Conan, J.-M., Raynaud, H.-F., Gendron, É., Basden, A., Vidal, F., Morris, T., Meimon, S., Petit, C., Gratadour, D., Martin, O., Hubert, Z., Sevin, A., Perret, D., Chemla, F., Rousset, G., Dipper, N., Talbot, G., Younger, E., Myers, R., Henry, D., Todd, S., Atkinson, D., Dickson, C., and Longmore, A., “First on-sky SCAO validation of full LQG control with vibration mitigation on the CANARY pathfinder,” *Optics Express* **22**, 23565 (Sept. 2014).
- [3] Poyneer, L. A., Ammons, S. M., Kim, M. K., Bauman, B., Terrel-Perez, J., Lemmer, A. J., and Nguyen, J., “Laboratory demonstration of the prediction of wind-blown turbulence by adaptive optics at 8 kHz with use of LQG control,” *62*, 1871 (Mar. 2023).
- [4] Wong, A. P., Norris, B. R. M., Tuthill, P. G., Scalzo, R., Lozi, J., Vievard, S., and Guyon, O., “Predictive control for adaptive optics using neural networks,” *Journal of Astronomical Telescopes, Instruments, and Systems* **7**, 019001 (Jan. 2021).

[5] Swanson, R., Lamb, M., Correia, C. M., Sivanandam, S., and Kutulakos, K., “Closed loop predictive control of adaptive optics systems with convolutional neural networks,” **503**, 2944–2954 (May 2021).

[6] Haffert, S. Y., Males, J. R., Close, L. M., Van Gorkom, K., Long, J. D., Hedglen, A. D., Guyon, O., Schatz, L., Kautz, M., Lumbres, J., Rodack, A., Knight, J. M., Sun, H., and Fogarty, K., “Data-driven subspace predictive control of adaptive optics for high-contrast imaging,” *arXiv e-prints* , arXiv:2103.07566 (Mar. 2021).

[7] Landman, R., Haffert, S. Y., Radhakrishnan, V. M., and Keller, C. U., “Self-optimizing adaptive optics control with reinforcement learning for high-contrast imaging,” *Journal of Astronomical Telescopes, Instruments, and Systems* **7**, 039002 (July 2021).

[8] Nousiainen, J., Engler, B., Kasper, M., Rajani, C., Helin, T., Heritier, C. T., Quanz, S. P., and Glauser, A. M., “Laboratory experiments of model-based reinforcement learning for adaptive optics control,” *Journal of Astronomical Telescopes, Instruments, and Systems* **10**, 019001 (Jan. 2024).

[9] Pou, B., Smith, J., Quinones, E., Martin, M., and Gratadour, D., “Model-free reinforcement learning with a non-linear reconstructor for closed-loop adaptive optics control with a pyramid wavefront sensor,” in [*Adaptive Optics Systems VIII*], Schreiber, L., Schmidt, D., and Vernet, E., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **12185**, 121852U (Aug. 2022).

[10] Guyon, O. and Males, J., “Adaptive Optics Predictive Control with Empirical Orthogonal Functions (EOFs),” *arXiv e-prints* , arXiv:1707.00570 (July 2017).

[11] Jensen-Clem, R., Bond, C. Z., Cetre, S., McEwen, E., Wizinowich, P., Ragland, S., Mawet, D., and Graham, J., “Demonstrating predictive wavefront control with the Keck II near-infrared pyramid wavefront sensor,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **11117**, 111170W (Sept. 2019).

[12] van Kooten, M. A. M., Doelman, N., and Kenworthy, M., “Robustness of prediction for extreme adaptive optics systems under various observing conditions. An analysis using VLT/SPHERE adaptive optics data,” **636**, A81 (Apr. 2020).

[13] Fowler, J., Van Kooten, M. A. M., and Jensen-Clem, R., “Battle of the predictive wavefront controls: comparing data and model-driven predictive control for high contrast imaging,” in [*Adaptive Optics Systems VIII*], Schreiber, L., Schmidt, D., and Vernet, E., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **12185**, 1218582 (Aug. 2022).

[14] Fowler, J. and Landman, R., “Tempestas ex machina: A review of machine learning methods for wavefront control,” *arXiv e-prints* , arXiv:2309.00730 (Sept. 2023).

[15] Guyon, O., Sevin, A., Ferreira, F., Ltaief, H., Males, J., Deo, V., Gratadour, D., Cetre, S., Martinache, F., Lozi, J., Vieuard, S., Fruitwala, N., Bos, S., and Skaf, N., “Adaptive optics real-time control with the compute and control for adaptive optics (Cacao) software framework,” in [*Adaptive Optics Systems VII*], Schreiber, L., Schmidt, D., and Vernet, E., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **11448**, 114482N (Dec. 2020).

[16] Dessenne, C., Madec, P.-Y., and Rousset, G., “Optimization of a Predictive Controller for Closed-Loop Adaptive Optics,” **37**, 4623–4633 (July 1998).

[17] Marquis, L., Raynaud, H., Galland, N., De La Rosa, J., Garcia, I., Tubio, O., Garcia-Talavera, M., Lombardi, G., Lopez, M., Reverte, D., and Kulcsar, C., “First on-sky tests of lqg control for a 10m-class telescope: prelude on the gran telescopio canarias adaptive optics system,” *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **in same proceedings** (2024).

[18] van Kooten, M. A. M., Jensen-Clem, R., Cetre, S., Ragland, S., Bond, C. Z., Fowler, J., and Wizinowich, P., “Predictive wavefront control on Keck II adaptive optics bench: on-sky coronagraphic results,” *Journal of Astronomical Telescopes, Instruments, and Systems* **8**, 029006 (Apr. 2022).

[19] Kolmogorov, A., “The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds’ Numbers,” *Akademii Nauk SSSR Doklady* **30**, 301–305 (Jan. 1941).

[20] Taylor, G. I., “The Spectrum of Turbulence,” *Proceedings of the Royal Society of London Series A* **164**, 476–490 (Feb. 1938).

[21] Laarhoven, P. J. M., [Simulated annealing: Theory and applications] (1987).

- [22] Jensen-Clem, R., Dillon, D., Gerard, B., van Kooten, M. A. M., Fowler, J., Kupke, R., Cetre, S., Sanchez, D., Hinz, P., Laguna, C., Doelman, D., and Snik, F., “The Santa Cruz Extreme AO Lab (SEAL): design and first light,” in [*Techniques and Instrumentation for Detection of Exoplanets X*], Shaklan, S. B. and Ruane, G. J., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **11823**, 118231D (Sept. 2021).
- [23] Moreno, A., Chambouleyron, R., Jensen-Clem, R., Dillon, D., Hinz, P., and Macintosh, B., “Implementation and characterization of the vector vortex coronagraph on the seal testbed,” *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **in same proceedings** (2024).