



Sequential coronagraphic low-order wavefront control

Michael Bottom^a, Samuel A. U. Walker^{a*}, Ian Cunyningham^{a*}, Charlotte Guthery^b, and Jacques-Robert Delorme^b

^aInstitute for Astronomy, University of Hawai'i, 640 N. Aohoku Place, Hilo, HI, 96720, United States

^bW. M. Keck Observatory, 65-1120 Mamalahoa Hwy, Waimea, HI 96743, United States

ABSTRACT

Coronagraphs are highly sensitive to wavefront errors, with performance degrading rapidly in the presence of low-order aberrations. Correcting these aberrations at the coronagraphic focal plane is key to optimal performance. We present two new methods based on the sequential phase diversity approach of the “Fast and Furious” algorithm that can correct low-order aberrations through a coronagraph. The first, called “2 Fast 2 Furious,” is an extension of Fast and Furious to all coronagraphs with even symmetry. The second, “Tokyo Drift,” uses a deep learning approach and works with general coronagraphic systems, including those with complex phase masks. Both algorithms have 100% science uptime and require effectively no diversity frames or additional hardware beyond the deformable mirror and science camera, making them suitable for many high contrast imaging systems. We present theory, simulations, and preliminary lab results demonstrating their performance.

Keywords: Wavefront control, coronagraphs, high-contrast imaging

1. INTRODUCTION

Coronagraph designs have advanced greatly in recent years, driven by the needs of “extreme” adaptive optics systems (ExAO, [9]) operating on large ground-based telescopes (SCEAO, GPI, SPHERE, MagAO-X, etc). Modern coronagraphs like the PIAA [8] or vortex [21] deliver close to optimal performance within the limits imposed by wave optics. However, there are fundamental trades between coronagraph contrast and inner working angle, and sensitivity to low-order aberrations such as tip/tilt, focus, astigmatism, etc. Even in coronagraphs with modest performance, when aiming for high contrast, the sensitivity to low-order aberrations is important.

Low-order aberrations are not all equal, however. Aberrations occurring before the coronagraphic focal plane are more destructive than those occurring after. For example, Figure 1 shows a small (0.2λ) defocus term applied to a unocculted PSF (such as a planet), which barely affects the peak brightness (eg, Strehl ratio). However




Figure 1. Simulation of Keck’s NIRC2 PSF with no errors, with .2 (wave / D) defocus, then repeated with a vortex coronagraph, demonstrating the much higher sensitivity of the vortex to small aberrations

applying the same error before a coronagraphic focal plane (in this case a vortex) leads to degradation in contrast by almost an order of magnitude.

Given this sensitivity to low-order aberrations, there are two complementary approaches to sensing and controlling these aberrations at the coronagraph plane.

One way to solve such a problem is to use extra hardware to detect reflected or refracted light from the coronagraph optics and use this to determine the level of aberration in the system. Guyon 2009 [10] proposed such a scheme using a defocused image of light reflected from a coronagraphic spot. The high contrast imaging system SPHERE on the VLT and the Stellar Double Coronagraph at Palomar used similar approaches to control pointing only on the coronagraph, with a tip/tilt sensor hardware mounted very close to the coronagraphic focal plane [2, 6]. Beyond tip and tilt, Singh et al. 2014 [31] demonstrated control of many low-order modes using light reflected from Lyot-style coronagraphic masks. In these proceedings Haffert et al. also presents a near-optimal wavefront sensor based on beam apodizing optics and a Zernike phase mask. The main drawbacks of hardware-based coronagraphic low-order wavefront sensors is that they require extra hardware.

An alternative method is to use the the science camera focal plane as the wavefront sensor, imaging the light transmitted through the coronagraph. This approach is similar to PSF sharpening algorithms in non-coronagraphic imaging. For example, The Gerchberg-Saxton algorithm uses known injected phase diversity (usually defocus) to solve for the residual phase aberrations in the pupil plane, which can then be corrected resulting in Strehl improvement. The need for the phase diversity is derived from the well-known even sign degeneracy from symmetric pupils such as those on most telescopes. Asymmetric pupils allow for direct sensing of all modes [20] without phase diversity, as can other specialized (asymmetric) optical designs like vector apodizing phase plates [4] or self-coherent cameras [1]. These may be seen as incorporating wavefront sensing directly into the coronagraph design, but their principles and operation are not applicable to general coronagraphic systems.

Focal-plane wavefront control through a general coronagraphic system was first described in Paul et al. with the COFFEE algorithm [27], which used a phase-diversity approach (using large defocus and astigmatism terms) to derive aberrations upstream and downstream of the coronagraph. This approach was shown to work well in simulation and was demonstrated using the internal source at SPHERE [26]. While it can be used with any coronagraphic system, it has the disadvantage of requiring two diversity frames per correction step, reducing science time.

Ideally, the focal plane wavefront control algorithm will require minimum or no diversity frames, and like COFFEE, work with any coronagraph. Additionally, it should work reliably in an operational environment, and converge rapidly. For example, an algorithm called QACITS [12] to control just tip/tilt has been in regular on-sky use with the Keck observatory vortex coronagraph [30, 11] effectively correcting slow pointing drifts off the vortex and dramatically improving performance. At its core, QACITS uses focal-plane images of the tip/tilt response to build an linearized model. While QACITS was developed specifically for a vortex coronagraph, the main ideas

Send correspondence to MB, mbottom@hawaii.edu

*contributed equally to this work

could be applied to almost any coronagraphic system to control tip/tilt. Orban de Xivry et al. [25] extended this work to derive many low-order aberrations through a vortex coronagraph using the weak phase solution (ie small aberrations applied), and with two phase diversities applied. This technique was shown to work very well in simulation, but relied on an analytic formulation only applicable to a vortex coronagraph with an unobstructed pupil.

In summary, we submit that an ideal low-order focal-plane wavefront control algorithm would possess the following properties:

1. Work reliably on-sky and converge quickly
2. Require no extra hardware beyond a science camera and deformable mirror
3. Require no diversity frames
4. Work with arbitrary pupils
5. Work with any coronagraph

It would seem that points 2, 3, and 4 are in mutual contradiction, as a symmetric pupil would naturally inherit the even phase degeneracy ultimately due to Fourier symmetries, and one would need to use extra hardware or diversity frames to determine the sign. However, there is a practical workaround to this, which is used in the (non-coronagraphic) sequential phase diversity algorithm Fast and Furious.

Fast and Furious is a high-performance focal-plane PSF optimization algorithm. First developed in 2013 by Keller et al. [13], it is based on the sequential phase-diversity approach of Gonsalves [7] in the weak-aberration limit. An on-sky astronomical demonstration occurred at the Subaru telescope in 2019, generating excitement by showing potential for correcting the notorious “low-wind effect” [5], where thermal discontinuities create large piston errors between telescope spiders, destroying PSF quality. Fast and Furious uses the following steps to sense the aberrations:

1. The PSF image is divided into even (symmetric) and odd (antisymmetric) parts.
2. The odd part can directly solve for the odd electric field.
3. The even part can find the amplitude of the even electric field but not its sign.
4. Using the previous PSF and correction, the sign of the even component is determined.
5. The sign and amplitude components provide the full even electric field.
6. By Fourier transforming the odd and even electric field, the complete wavefront error is reconstructed.
7. A correction is made ($\Phi_i \rightarrow \Phi_d$), and the process repeats.

Fast and Furious has also been used on-sky at Keck on many occasions (Bos et al. [3] also see Guthery et al., these proceedings), and has demonstrated reliable and rapid convergence. One key idea of the algorithm is to use the previous correction as the diversity phase, which means the PSFs stays very close to optimal at every iteration, and no science time is lost to large diversity phases. Indeed, the algorithm can be run concurrently with observations, which also tends to stabilize the PSFs, leading to better systematics control. The other aspect which is helpful is that it only requires a model of the pupil and uses Fraunhofer propagators, so calculation speed and complexity is minimal. In its native implementation, F&F has several requirements which preclude operation with arbitrary optical systems: the exit pupil must be real and symmetric, the PSF must be distorted only by phase aberrations, and the PSF must be shift-invariant. The only coronagraph design that fulfils these three conditions is a symmetric apodized pupil coronagraph.

Therefore, the goal of this work is to try to modify F&F to work with arbitrary coronagraphic systems. The first method we present, nominally called “2 Fast 2 Furious,” (2F2F) is a linear method similar to the regular F&F algorithm but including some extra coronagraphic operators. 2F2F nearly fulfils the ideal algorithm criteria presented in the list above, but is restricted to work on symmetric coronagraphs, such as classic Lyots, PIAAs, or several classes of shaped pupil masks. It does not work on hybrid designs or with complex phase masks like the vortex coronagraph. The second method, informally referred to as “Tokyo Drift,” (TD) is F&F in spirit only, with the same trick of using the previous correction to break the sign degeneracy, but uses a neural network trained on the optical model to derive the aberrations. The benefit of all this extra work is flexibility, as TD can

which describes the final focal plane image p as a Fourier transform ($\mathcal{F}$) of an electric field with amplitude A and phase ϕ from the pupil plane (where the incoming wavefront is defined) to the focal plane (where the observed image is defined). The case of $\phi = 0$, corresponds to an unaberrated image with no phase error. When using a coronagraph, one typically places a focal plane mask (denoted by C) and a pupil plane Lyot stop (denoted by M) into this optical path, thus:

$$p = |\mathcal{F}\{M\mathcal{F}^{-1}\{C\mathcal{F}\{A \exp(i\phi)\}\}\}|^2. \quad (2)$$

Here, M , C and A are all properties of the telescope system in use, as A is defined by the telescope aperture. One defining assumption of the Fast and Furious approach is that the aperture A is a real and even function of pupil plane position. Without this assumption the expansion out of the modulus in equation 1 involves several terms that would otherwise cancel, significantly complicating the math. If both M and C are also real and even functions, as is the case for many coronagraphs, then the three system properties in equation 2 can be combined into one real and even system parameter, hereinafter denoted by X and given by

$$X = M\mathcal{F}^{-1}\{C\mathcal{F}\{A\}\}. \quad (3)$$

Substituting this back into equation 2 gives that the image is described by the equation

$$p = |\mathcal{F}\{X \exp(i\phi)\}|^2, \quad (4)$$

which is entirely analogous to equation 1 with the substitution $A \rightarrow X$.

This methodology is summarised below. We adopt the notation that the Fourier transform of a quantity denoted by an upper-case letter is denoted by a lower-case letter, i.e. $x = \mathcal{F}\{X\}$.

Firstly, the exponential in equation 4 is expanded to first order in ϕ , and then the phase is expressed as a sum of its even and odd terms, $\phi = \phi_e + \phi_o$. These terms are carried through to obtain expressions for the even and odd parts of the final image. The odd part of the final image is then given by

$$p_o = 2ix\mathcal{F}\{X\phi_o\}, \quad (5)$$

where $i = \sqrt{-1}$. Similarly, the even part of the final image is given by

$$p_e = x^2 + (\mathcal{F}\{X\phi_e\})^2 - (\mathcal{F}\{X\phi_o\})^2, \quad (6)$$

where it can be seen that if ϕ_o is already known from equation 5, the only unknown becomes ϕ_e . The F&F approach breaks the even aberration sign degeneracy evident in this equation by using the deformable mirror correction calculated in the previous iteration as a diversity phase. If p_n is the image at the most recent iteration, and p_{n-1} is the image at the iteration prior to that with a phase offset of ϕ_d from p_n (given by the deformable mirror correction applied from the previous iteration's calculations), then the difference between the even parts of these two images can be written as

$$\begin{aligned} p_{e,n} - p_{e,n-1} &= (\mathcal{F}\{X\phi_{o,d}\})^2 - (\mathcal{F}\{X\phi_{e,d}\})^2 \\ &\quad + 2\mathcal{F}\{X\phi_o\}\mathcal{F}\{X\phi_{o,d}\} - 2\mathcal{F}\{X\phi_e\}\mathcal{F}\{X\phi_{e,d}\}. \end{aligned} \quad (7)$$

[14] find their version of this equation to be unstable due to the subtraction of two similar images, using it purely to find the sign of ϕ_e in combination with the absolute value of the ϕ_e obtained from equation 6. We continue to follow this approach going forward.




Figure 3. Average final RMS WFE as a function of focal plane mask size in units of λ/D for various combinations of focal plane mask and filter wavelength available on Keck/NIRC2. The algorithm performs well for coronagraphs with small spot sizes $\lesssim 4\lambda/D$, which is expected given that we are only correcting the first 20 Zernike modes. Errorbars represent the standard error on 100 measurements of each datapoint, with differing initial phases for each measurement.

2.2 Testing and results

Despite the mathematical substitutions above, it was not clear that algorithm would work in practice due to the violation of conservation of energy in coronagraphic systems from pupil to focal plane, an implicit assumption in Fast and Furious [13]. We implemented the algorithm coupled with simulations of the Keck/NIRC2 optical path using the Python package HClPy [28]. In contrast to Fast and Furious, we find that 2 Fast 2 Furious performs poorly under direct inversion of equations (5), (6) and (7), necessitating a numerical approach for solving these equations. To increase computation speed, we fit for the coefficients of the Zernike modes that we are aiming to correct, which we find works well and performs at speeds on the order of 2-3 seconds, which is efficient enough for eventual on-sky deployment. Throughout our analysis we focus on correcting for the first 20 Zernike modes, under the assumption that the majority of the power of any on-sky aberrations would be in these modes.

Using these simulations, we characterised the regime in which 2 Fast 2 Furious converges. The first of our tests focused on the coronagraph sizes. For this, we simulated PSF images with a range of focal plane mask sizes and observing bands based on those currently available for science use on Keck/NIRC2. For each image, we then applied 0.5 radians total root-mean-square wavefront error (RMS WFE) injected randomly across the 20 Zernike modes that we aim to correct, before running the image through 20 iterations of 2 Fast 2 Furious and recording the RMS WFE of the final image. We repeated this process 100 times for each instrument configuration, randomly altering the shape of the wavefront error each time. We found that the algorithm converges best when operated with coronagraphs that have spot sizes $\leq 4\lambda/D$, as shown in Figure 3. Beyond this limit, the phase of the image either diverges from the initial injected value or else converges weakly. This is likely due to the low signal from Zernikes of index 20 or less at those spatial frequencies.

Our second test concerned the maximum amplitude of aberrations that 2 Fast 2 Furious could reliably correct. We took simulated images of the smallest 100 mas focal plane mask in the *K*-band and performed 100 tests of 2 Fast 2 Furious for various initial input aberrations RMS WFE amplitudes from 0.5 radians to 1.5 radians, again allowing 20 iterations for convergence. This experiment showed that 2 Fast 2 Furious reliably converges when dealing with aberrations below ~ 1.1 radians (1/6 of a wave), which would correspond to approximately 33% Strehl ratio without a coronagraph in the optical path (see Fig. 4). This value is key as the AO system at Keck is typically able to correct to this level or better [18], so the fact that 2 Fast 2 Furious can correct images




Figure 4. Average final RMS WFE after 20 iterations as a function of initial injected RMS WFE for a 100 mas coronagraph in the K -band on Keck/NIRC2. 2 Fast 2 Furious converges reliably every time for initial aberrations with an RMS of less than 1.1 radians. Errorbars represent the standard error on 100 measurements of each datapoint, with differing initial phases for each measurement.

at these Strehl levels demonstrates that the algorithm operates reliably at the point where non-common-path aberrations begin to dominate the wavefront error budget. It also is the same as the correction range of the non-coronagraphic F&F algorithm.

Once we had proved that 2 Fast 2 Furious worked in simulations, we sought lab validation. As such, we carried out several tests on the Keck/NIRC2 bench, and these tests are still ongoing. Our main preliminary result is that we have successfully used 2 Fast 2 Furious to correct for injected tip/tilt errors using the 100 mas coronagraph in the K -band on the bench. Evidence of this is presented in Figure 5, wherein we use the 5σ contrast at $0.1''$ as an indicator of image quality. The significant improvements in limiting contrast, as well as the improvement in PSF quality that is evident from the raw NIRC2 images, has us satisfied that 2 Fast 2 Furious can correct tip/tilt well in a real-world setting. We are now working to extend this to more Zernike modes, to mirror the work we have done in simulation.

3. TOKYO DRIFT

Korkiakoski et al. [14] modified the F&F algorithm to accommodate asymmetric pupils using a Gerchberg-Saxton like approach where the pupil amplitudes are iteratively calculated. However, this approach was found to have instability issues in the small-phase limit (eg, as the algorithm converged) which caused wildly divergent behavior due to subtraction of similar images leading to “divide by zero” errors, requiring further filtering and regularization. It is likely this issue would be exacerbated with coronagraphs. We instead opted to investigate a conceptually different approach using deep neural networks.

The deep neural network (DNN) analog of F&F was trained using supervised techniques where large datasets were simulated comprising PSFs of coronagraphic systems and the random wavefront errors which generated them. The network then learned to directly estimate the Zernikes that would best correct the observed PSF aberrations via regression. In operation, this would be used as any other WFS with some gain in an integrator loop.




Figure 5. Results from Keck/NIRC2 bench testing when correcting for tip-tilt with the 100mas coronagraph in the optical path. Images are presented in log scale. The 5σ $0''.1$ contrast improves with each iteration as the PSF becomes sharper and more well-defined, denoting improved PSF/coronagraph alignment.

3.1 Model training

For training Tokyo Drift, simple DNN architectures were used consisting of a stack of 2D convolutional layers [16] followed by multiple dense layers before the output. ReLU activation [22] was used in-between layers, with no additional activation after the output layers. While more complex architectures may yield some improvements to performance, experience on similar problems has suggested that the gains from this strategy appear to be minor compared to the importance of carefully specifying the problem and crafting the datasets.

Simulations were generated via HCIPy [28], using a custom module: `TelescopeSim`, built for use in machine learning tasks, with the ability to simulate filter bandwidth, detector and photon noise, deformable mirrors, Zernike aberrations in the pupil-plane, etc. 5. Specifics of the NIRC2 instrument from work on 2F2F were used, including precise telescope and instrument pupil functions, filter specifics, plate scales, and details of the vortex coronagraph which were validated against on-sky operation. Later iterations used similar details of the SCExAO VAMPIRES instrument 5.

For a supervised approach to training a DNN FPWFS using a single observation (i.e. with an asymmetric pupil [15]), one simple approach is using the focal-plane image as the input, and as the target for regression, using the specific, randomly drawn Zernike errors imposed on the pupil-plane to generate them. The loss function which proved most effective for these models is mean-absolute error (MAE). Inputs and outputs were then stacked into batches (with typical batch size of 128). Therefore, for a 64×64 pixel focal plane image, predicting 35 Zernikes, the input batch would have the shape (128, 64, 64, 1) (where the first dimension is batch size, and the final dimension indicates a single convolutional input channel), and the target batch shape would be (128, 35). However, to utilize the phase diversity offered by the F&F technique, a more complex model was needed.




Figure 6. Sample Tokyo Drift model architecture utilizing the previous actuation as a separate input from the images. This input is concatenated with the outputs of the convolution stack. Afterwards, a set of dense layers are utilized to fuse this information. Finally, the output is broken up into the magnitude and sign of the errors with corresponding loss functions.




Figure 7. Schematic of the multi-input, multi-output Tokyo Drift architecture. The model is trained on simulated data of the optical system, with inputs of two-image stacks, along with their intervening actuations and output of the predicted, best Zernike modes for correction. The outputs are split into the magnitude of the Zernikes using MAE loss, and their sign using BCE.

The F&F algorithm requires two input images: the most recent and the one preceding it, as well as the differential deformable mirror actuation used between these frames. As 2D convolutional layers were used for the initial layers of simpler models, adding the second image was simply accomplished by stacking it as an additional convolutional channel. Adding the qualitatively different input vector of the intervening actuation meant this new model needed transition to multiple-inputs. A schematic of how this was accomplished is illustrated in Figure 6 where the actuation vector input is fused with the output of the image convolution stack. Therefore, with a batchsize of 128, the input now had a structure of (128, 64, 64, 2) for the input images, and (128, 35 for a 35 Zernike deformable mirror actuation between frames.

Further, it was recognized that while even modes were encoded ambiguously in the focal-plane, it was only the signs of these features which were unknown. A DNN that is trained to regress on an single image of an even Zernike-mode ends up learning to predict zeros in order to minimize error while predicting odd modes correctly. However, the structure in the image-plane which encodes the magnitude of these even aberrations should still be present to learn. It was intuited that splitting up the prediction target into magnitude and sign should be useful, so each could be separately back-propagated with their own loss functions. This meant the convolutional stack could focus on extracting the features encoding magnitudes, and the dense stack responsible for fusing in the actuation information could learn to correctly predict the signs as a separate task. Moving to a multi-output model, the target batches were shaped (128, 35) for the magnitudes encoded as floating point numbers utilizing MAE loss, and (128, 35) for the signs of those targets, encoded as floats representing Boolean values and utilizing Binary Cross Entropy (BCE) loss. A schematic of the full flow of inputs, outputs, and losses are shown in Figure 7.

Constructing the distributions of errors for training this model was quite a bit more complex than previous, single-image models where it is only necessary to select a reasonable distribution of initial errors for each sample. With this F&F analog, a representative actuation for the consequent image also had to be constructed. The final strategy employed was as follows:

1. For the initial frame, draw random errors as before




Figure 8. Example of iterative correction in an AO loop on a simulation of NIRC2 vortex coronagraph. The first two frames are created via the training set generator. These PSFs, plus the actuation between them, are fed into the trained DNN to produce the first prediction (third frame). The process then repeats several more times until it converges to minimal wavefront error on the coronagraph.

2. To construct a reasonable actuation, assume the FPWFS from first step is only somewhat accurate, and is running in a loop with some variable gain. Therefore:
 - (a) Take the true error of the first frame that would provide perfect correction and scale it by a random gain (for instance $[.01 - .99]$). In fact, it proved important to scale beyond an ideal correction to simulate over-correction and subsequent errors flipping sign, so in the end, a simulated gain of $[0.01 - 1.99]$ was used.
 - (b) Add some extra, independent, random errors to simulate inaccurate sensing
3. Finally, in order for the model to avoid simply propagating the signs from the input, and to provide a balanced training set for the BCE loss, the signs of the actuation used for the final frame were randomly flipped. This simulated the FPWFS's blindness to sign in its first estimation.

Several generations of models were produced in this fashion, necessitating large hyper-parameter sweeps over DNN architecture details, learning-rate, etc. Bayesian hyperparameter optimization [32] and Hyperband early stopping [17] were used to search the large parameter spaces efficiently. After which, the most successful models were trained to convergence. Along the way, these tests verified that the multi-output, multi-loss formulation did indeed provide a significant advantage, with a single output regression failing to learn much at all during training. Additionally, after the initial architecture search was complete, further hyperparameter sweeps proved unnecessary for training TD models on completely different optical systems. Simply generating a new dataset and retraining a model specific to those systems appears to be effective, allowing this technique to be quickly applied to different telescopes and their specific instrument configurations.

3.2 Simulation results

The performance of this model was evaluated by simulating an adaptive-optics loop, where the initial state was created with the same dataset generation function as training. Using the TD model then next predicted correction was applied with a gain of .6. This process continued from there, iteratively feeding the predicted actuations and resulting observations back into the model (Figure 8).

The resulting model seems to be capable of correcting aberrations with surprisingly high initial RMS errors (Figure 9). Typically, focal-plane sensing techniques are limited to around 1 radian RMS error before being




Figure 9. 1000 random initial states with a wide range of initial errors were selected, then each was corrected over 16 iterations. The runs were grouped by initial RMS error and the results averaged here.

unable to converge. While errors this large should rarely be encountered in a stable coronagraphic instrument with high-quality traditional WFS, it is still promising to see that the technique has capabilities beyond what has previously been shown.

3.3 Bench results

Testing with real optical systems is currently underway on two fronts. First, a simple adaptive-optics test-bed was built in the Institute for Astronomy - Hilo’s optical lab. This consists of a fiber-fed laser, which is collimated, the pupil then re-imaged on a 12x12 Boston Micromachines DM, and then imaged on a detector (5). Aberrations are then introduced via the DM and sensed with a TD model trained on simulated data. While there have been many difficulties with this setup – being unstable optically, a stuck actuator and other issues with the DM, and the model not being trained on a very precise simulation of the system as it changes – the Tokyo Drift model still performs well, converging with large initial input errors.

More recently, we have been doing daytime testing using the VAMPIRES instrument [23, 19] on Subaru. The first model was trained without a coronagraph and with only 10 Zernikes (Figure 10), with the intention of scaling up the difficulty by adding higher order aberrations and utilizing the recently installed vector vortex coronagraph. Again, with many small sim-to-real barriers to overcome, the Tokyo Drift model has shown successful convergence. Both of these optical systems demonstrate that the success of the TD approach is viable for real optical systems.

4. CONCLUSIONS

We have presented two approaches to coronagraphic low-order wavefront control using a sequential method based on the Fast and Furious algorithm of Keller and Korkiakoski.

The first, “2 Fast 2 Furious” is an extension of the linear methodology of F&F to Lyot-style coronagraphs. The simple implementation and advantageous properties of F&F are preserved, though as before, a symmetric optical system is required, precluding usage with coronagraphic architectures like vortex phase masks. Simulations show similar convergence properties of the original F&F. Bench testing indicates the algorithm works in principle, though further work is needed to demonstrate correction of multiple low-order modes.




Figure 10. Data from a recent Tokyo Drift bench-test run on SCExAO’s VAMPIRES instrument in the F760 filter without a coronagraph. A model was trained to correct 10 Zernike-modes, random initial errors were chosen and set on the DM and iteratively corrected using only this technique.

The second approach, “Tokyo Drift,” uses a deep neural network trained on simulated images of the coronagraphic system with aberrations, but employs the same inputs and outputs as F&F. This gives it the ability to correct asymmetric optical systems, as well as outside the weak-aberration limit. The drawback of this approach is the time it takes to train the network, though this is a one-time sunk cost. Simulations show excellent performance with a vortex coronagraph, and lab testing shows very good convergence, albeit in a system without a coronagraph. Bench tests with coronagraphs in place are now commencing.

It is intriguing to consider extending TD to dig dark holes in order to further enhance contrast. One possibility for achieving this would involve pre-calculating Zernike patterns that create dark holes in a perfectly corrected system, and then generating new simulated datasets with perturbations, with the target of converging on these patterns. In this way, the problem remains a supervised one, and so trivial to adapt the current training process. Another possibility, is to reframe the problem using reinforcement learning (using the TD neural network as the “policy,” in reinforcement learning nomenclature). The advantage is that the target for prediction does not need to be differentiable. Instead, one could directly use measurements of the quality of the dark hole, e.g. $\text{minimizing} (\text{flux in the desired dark spot location}) / (\text{central PSF brightness})$. This has the advantages of not needing to guess at useful distributions of errors in the training set, learning better the model’s own prediction errors, and the ability to use with real optical systems instead of simulations.

Future work will focus on validation on coronagraphic testbeds. Keck/NIRC2 finished a major detector upgrade and VAMPIRES on Subaru recently had both Lyot coronagraphs [19] and a vortex coronagraph (Doelman et al. in prep) installed, which will facilitate further testing and validation.

5. CODE, DATA, AND MATERIALS

Framework for simulating telescope systems with HCIPy, formulated for running machine learning experiments <https://github.com/icunnyngham/TelescopeSim>

ACKNOWLEDGMENTS

Thanks to Miles Lucas for providing precise optical system details for SCExAO’s VAMPIRES instrument as well as extensive help running tests on-bench.

The technical support and advanced computing resources from University of Hawaii Information Technology Services – Cyberinfrastructure, funded in part by the National Science Foundation CC* awards # 2201428 and # 2232862 are gratefully acknowledged.

Thanks to Christoph Baranec for the use of his deformable mirror and other important optical components for the IFA AO bench. Thanks to James Ou for help with interfacing with the DM. Thanks to Charles-Antoine Claveau and Aidan Walk for help interfacing with the detector.

References

- [1] P. Baudoz et al. “The Self-Coherent Camera: a new tool for planet detection”. en. In: *Proceedings of the International Astronomical Union* 1.C200 (Oct. 2005), pp. 553–558. ISSN: 1743-9213, 1743-9221. DOI: [10.1017/S174392130600994X](https://doi.org/10.1017/S174392130600994X). URL: https://www.cambridge.org/core/product/identifier/S174392130600994X/type/journal_article (visited on 11/07/2023).
- [2] Pierre Baudoz et al. “The differential tip-tilt sensor of SPHERE”. en. In: ed. by Ian S. McLean, Suzanne K. Ramsay, and Hideki Takami. San Diego, California, USA, July 2010, 77355B. DOI: [10.1117/12.858274](https://doi.org/10.1117/12.858274). URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.858274> (visited on 11/16/2023).
- [3] Steven P. Bos et al. “Fast” and Furious focal-plane wavefront sensing at W. M. Keck Observatory. en. arXiv:2107.07601 [astro-ph]. July 2021. URL: <http://arxiv.org/abs/2107.07601> (visited on 12/06/2023).
- [4] Steven P. Bos et al. “Focal-plane wavefront sensing with the vector Apodizing Phase Plate”. en. In: *Astronomy & Astrophysics* 632 (Dec. 2019). arXiv:1909.08317 [astro-ph], A48. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201936062](https://doi.org/10.1051/0004-6361/201936062). URL: <http://arxiv.org/abs/1909.08317> (visited on 11/29/2023).
- [5] Steven P. Bos et al. “On-sky verification of Fast and Furious focal-plane wavefront sensing: Moving forward toward controlling the island effect at Subaru/SCEXAO”. en. In: *Astronomy & Astrophysics* 639 (July 2020). arXiv:2005.12097 [astro-ph], A52. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202037910](https://doi.org/10.1051/0004-6361/202037910). URL: <http://arxiv.org/abs/2005.12097> (visited on 11/29/2023).
- [6] Michael Bottom et al. “Stellar Double Coronagraph: A Multistage Coronagraphic Platform at Palomar Observatory”. en. In: *Publications of the Astronomical Society of the Pacific* 128.965 (July 2016), p. 075003. ISSN: 0004-6280, 1538-3873. DOI: [10.1088/1538-3873/128/965/075003](https://doi.org/10.1088/1538-3873/128/965/075003). URL: <https://iopscience.iop.org/article/10.1088/1538-3873/128/965/075003> (visited on 11/29/2023).
- [7] Robert A Gonsalves. “Phase retrieval and diversity in adaptive optics”. In: *Optical Engineering* 21.5 (1982), pp. 829–832.
- [8] O. Guyon. “Phase-induced amplitude apodization of telescope pupils for extrasolar terrestrial planet imaging”. en. In: *Astronomy & Astrophysics* 404.1 (June 2003), pp. 379–387. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:20030457](https://doi.org/10.1051/0004-6361:20030457). URL: <http://www.aanda.org/10.1051/0004-6361:20030457> (visited on 11/29/2023).
- [9] Olivier Guyon. “Extreme Adaptive Optics”. en. In: (2018).
- [10] Olivier Guyon, Taro Matsuo, and Roger Angel. “CORONAGRAPHIC LOW-ORDER WAVE-FRONT SENSOR: PRINCIPLE AND APPLICATION TO A PHASE-INDUCED AMPLITUDE CORONAGRAPH”. en. In: *The Astrophysical Journal* 693.1 (Mar. 2009), pp. 75–84. ISSN: 0004-637X, 1538-4357. DOI: [10.1088/0004-637X/693/1/75](https://doi.org/10.1088/0004-637X/693/1/75). URL: <https://iopscience.iop.org/article/10.1088/0004-637X/693/1/75> (visited on 11/18/2023).
- [11] Elsa Huby et al. “On-sky performance of the QACITS pointing control technique with the Keck/NIRC2 vortex coronagraph”. en. In: *Astronomy & Astrophysics* 600 (Apr. 2017). arXiv:1701.06397 [astro-ph], A46. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201630232](https://doi.org/10.1051/0004-6361/201630232). URL: <http://arxiv.org/abs/1701.06397> (visited on 11/16/2023).
- [12] Elsa Huby et al. “The QACITS pointing sensor: from theory to on-sky operation on Keck/NIRC2”. en. In: arXiv:1607.05497 [astro-ph]. July 2016, p. 990920. DOI: [10.1117/12.2233274](https://doi.org/10.1117/12.2233274). URL: <http://arxiv.org/abs/1607.05497> (visited on 11/29/2023).

- [13] Christoph U. Keller et al. “Extremely fast focal-plane wavefront sensing for extreme adaptive optics”. en. In: arXiv:1207.3273 [astro-ph, physics:physics]. Sept. 2012, pp. 844721–844721–10. DOI: [10.1117/12.926725](https://doi.org/10.1117/12.926725). URL: <http://arxiv.org/abs/1207.3273> (visited on 11/06/2023).
- [14] Visa Korkiakoski et al. “Fast & Furious focal-plane wavefront sensing”. en. In: *Applied Optics* 53.20 (July 2014), p. 4565. ISSN: 1559-128X, 2155-3165. DOI: [10.1364/AO.53.004565](https://doi.org/10.1364/AO.53.004565). URL: <https://opg.optica.org/abstract.cfm?URI=ao-53-20-4565> (visited on 11/07/2023).
- [15] Jeff Kuhn et al. “The small-ELF project: toward an ultra-large coronagraphic optical receiver”. In: *Ground-based and Airborne Telescopes IX*. Vol. 12182. SPIE. 2022, pp. 161–184.
- [16] Yann LeCun et al. “Backpropagation applied to handwritten zip code recognition”. In: *Neural computation* 1.4 (1989), pp. 541–551.
- [17] Lisha Li et al. “Hyperband: A novel bandit-based approach to hyperparameter optimization”. In: *The journal of machine learning research* 18.1 (2017), pp. 6765–6816.
- [18] Michael C. Liu. “LGS AO science impact: present and future perspectives”. In: *Adaptive Optics Systems*. Ed. by Norbert Hubin, Claire E. Max, and Peter L. Wizinowich. Vol. 7015. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 2008, 701508, p. 701508. DOI: [10.1117/12.789766](https://doi.org/10.1117/12.789766). arXiv: [0807.0634](https://arxiv.org/abs/0807.0634) [astro-ph].
- [19] Miles Lucas et al. *A Visible-light Lyot Coronagraph for SCExAO/VAMPIRES*. en. arXiv:2208.02380 [astro-ph]. Aug. 2022. URL: <http://arxiv.org/abs/2208.02380> (visited on 12/06/2023).
- [20] Frantz Martinache. “The Asymmetric Pupil Fourier Wavefront Sensor”. en. In: *Publications of the Astronomical Society of the Pacific* 125.926 (Apr. 2013). arXiv:1303.6678 [astro-ph], pp. 422–430. ISSN: 00046280, 15383873. DOI: [10.1086/670670](https://doi.org/10.1086/670670). URL: <http://arxiv.org/abs/1303.6678> (visited on 11/16/2023).
- [21] D. Mawet et al. “Annular Groove Phase Mask Coronagraph”. en. In: *The Astrophysical Journal* 633.2 (Nov. 2005), pp. 1191–1200. ISSN: 0004-637X, 1538-4357. DOI: [10.1086/462409](https://doi.org/10.1086/462409). URL: <https://iopscience.iop.org/article/10.1086/462409> (visited on 11/29/2023).
- [22] Vinod Nair and Geoffrey E Hinton. “Rectified linear units improve restricted boltzmann machines”. In: *Proceedings of the 27th international conference on machine learning (ICML-10)*. 2010, pp. 807–814.
- [23] Barnaby Norris et al. “The VAMPIRES instrument: imaging the innermost regions of protoplanetary discs with polarimetric interferometry”. In: *Monthly Notices of the Royal Astronomical Society* 447.3 (2015), pp. 2894–2906.
- [24] G Orban de Xivry et al. “Focal plane wavefront sensing using machine learning: performance of convolutional neural networks compared to fundamental limits”. en. In: *Monthly Notices of the Royal Astronomical Society* 505.4 (July 2021), pp. 5702–5713. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stab1634](https://doi.org/10.1093/mnras/stab1634). URL: <https://academic.oup.com/mnras/article/505/4/5702/6295324> (visited on 11/16/2023).
- [25] Gilles Orban De Xivry et al. “Post-coronagraphic PSF sharpening with the vortex coronagraph”. en. In: *Proceedings of the Adaptive Optics for Extremely Large Telescopes 5*. Instituto de Astrofísica de Canarias (IAC), 2017. DOI: [10.26698/A04ELT5.0066](https://doi.org/10.26698/A04ELT5.0066). URL: <http://research.iac.es/congreso/A04ELT5/pages/proceeding0066.html> (visited on 11/17/2023).
- [26] B. Paul et al. “Compensation of high-order quasi-static aberrations on SPHERE with the coronagraphic phase diversity (COFFEE)”. en. In: *Astronomy & Astrophysics* 572 (Dec. 2014). arXiv:1409.2247 [astro-ph], A32. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201424133](https://doi.org/10.1051/0004-6361/201424133). URL: <http://arxiv.org/abs/1409.2247> (visited on 11/29/2023).
- [27] B. Paul et al. “High-order myopic coronagraphic phase diversity (COFFEE) for wave-front control in high-contrast imaging systems”. en. In: *Optics Express* 21.26 (Dec. 2013), p. 31751. ISSN: 1094-4087. DOI: [10.1364/OE.21.031751](https://doi.org/10.1364/OE.21.031751). URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-21-26-31751> (visited on 11/16/2023).
- [28] E. H. Por et al. “High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator”. In: *Adaptive Optics Systems VI*. Vol. 10703. Proc. SPIE. 2018. DOI: [10.1117/12.2314407](https://doi.org/10.1117/12.2314407). URL: <https://doi.org/10.1117/12.2314407>.

- [29] Maxime Quesnel et al. “Deep learning-based focal plane wavefront sensing for classical and coronagraphic imaging”. en. In: *Adaptive Optics Systems VII*. Ed. by Dirk Schmidt, Laura Schreiber, and Elise Vernet. Online Only, United States: SPIE, Dec. 2020, p. 56. ISBN: 978-1-5106-3683-5 978-1-5106-3684-2. DOI: [10.1117/12.2562456](https://doi.org/10.1117/12.2562456). URL: <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11448/2562456/Deep-learning-based-focal-plane-wavefront-sensing-for-classical-and/10.1117/12.2562456.full> (visited on 11/07/2023).
- [30] E. Serabyn et al. “THE W. M. KECK OBSERVATORY INFRARED VORTEX CORONAGRAPH AND A FIRST IMAGE OF HIP 79124 B”. en. In: *The Astronomical Journal* 153.1 (Jan. 2017), p. 43. ISSN: 1538-3881. DOI: [10.3847/1538-3881/153/1/43](https://doi.org/10.3847/1538-3881/153/1/43). URL: <https://iopscience.iop.org/article/10.3847/1538-3881/153/1/43> (visited on 11/16/2023).
- [31] Garima Singh et al. “Lyot-based Low Order Wavefront Sensor for Phase-mask Coronagraphs: Principle, Simulations and Laboratory Experiments”. en. In: *Publications of the Astronomical Society of the Pacific* 126.940 (June 2014). arXiv:1404.7201 [astro-ph], pp. 586–594. ISSN: 00046280, 15383873. DOI: [10.1086/677048](https://doi.org/10.1086/677048). URL: <http://arxiv.org/abs/1404.7201> (visited on 11/18/2023).
- [32] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. “Practical bayesian optimization of machine learning algorithms”. In: *Advances in neural information processing systems* 25 (2012).