Data-driven subspace predictive control of adaptive optics for high-contrast imaging

Sebastiaan Haffert

NASA Postdoctoral Hubble Fellow
University of Arizona

Wavefront sensing in the VLT/ELT era 2020
15 October 2020

NASA Hubble Fellowship Program

“Data-driven subspace predictive control of adaptive optics for high-contrast imaging.”,
Sebastiaan Y. Haffert, Jared R. Males, Laird M. Close, Kyle Van Gorkom, Joseph D. Long, Alexander D. Hedglen, Olivier Guyon, Lauren Schatz, Maggie Kautz, Jennifer Lumbres, Alex Rodack, Justin M. Knight, He Sun, Kevin Fogarty, under review

THE UNIVERSITY OF ARIZONA

Temporal control errors

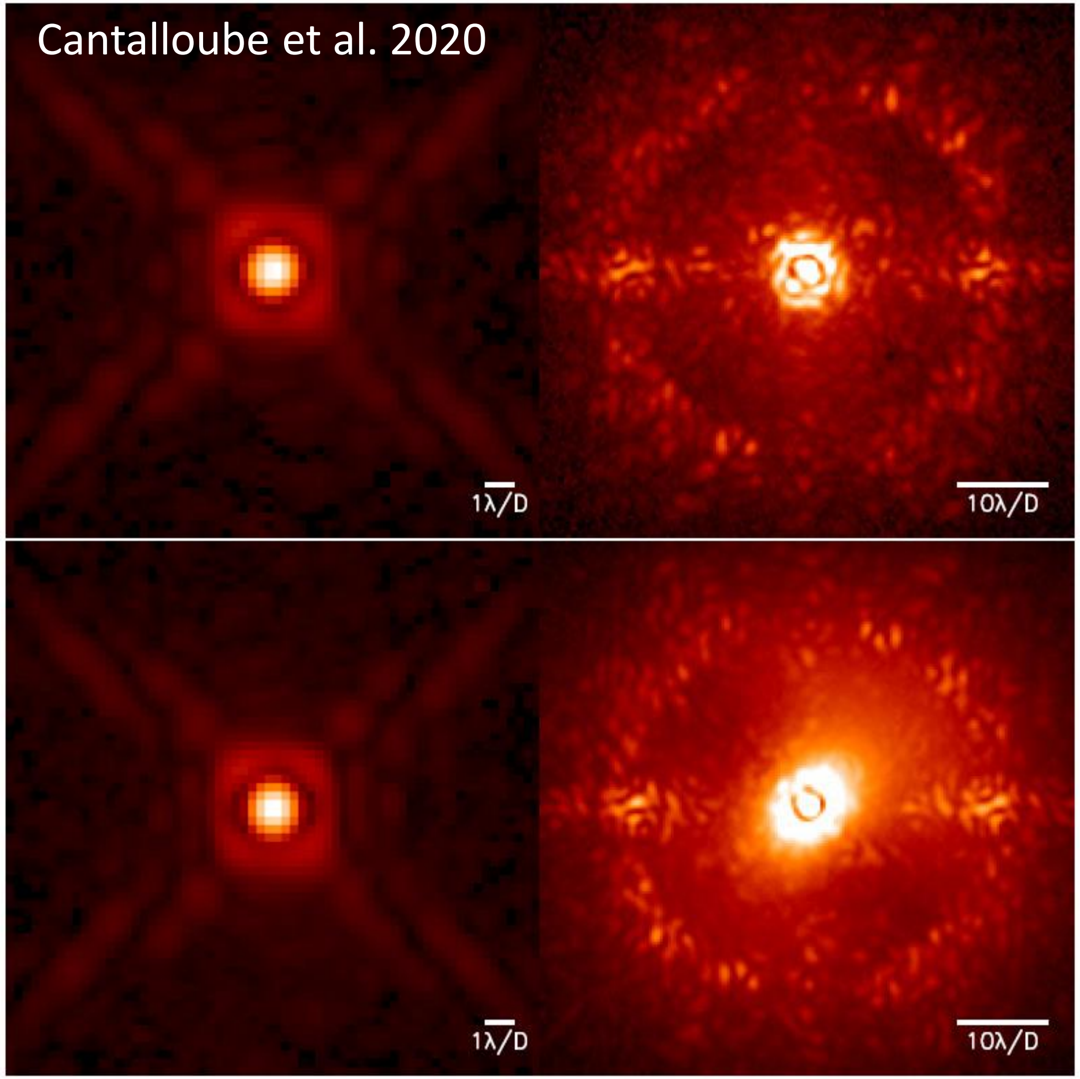

$\tau_0 = 9\text{ms}$

$\tau_0 < 4\text{ms}$

- The impact on the spatial-resolution of instruments is not very strong.
- But there is a strong impact on the post-coronagraphic contrast.

Temporal control errors

shaffert@arizona.edu

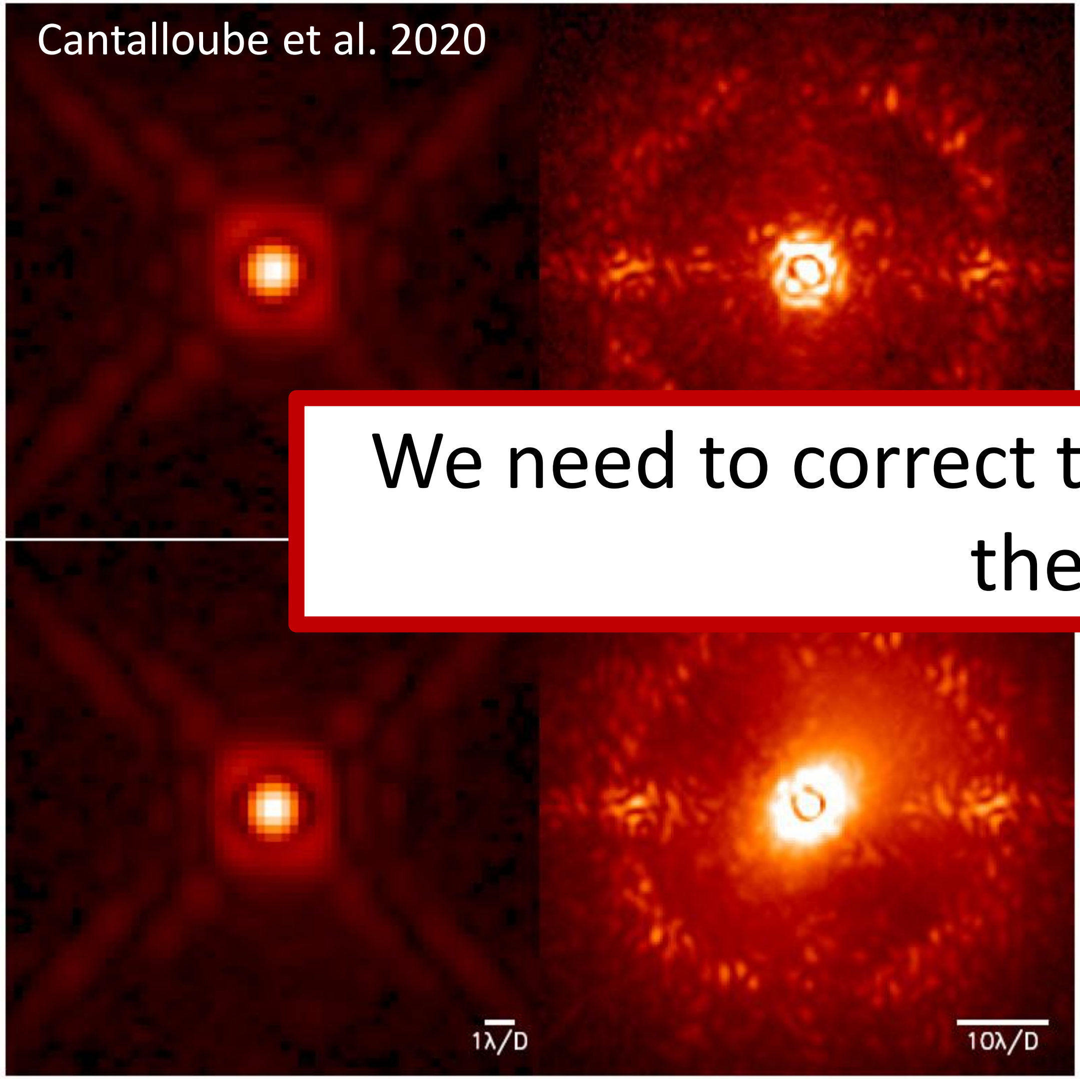

$\tau_0 = 9\text{ms}$

$\tau_0 < 4\text{ms}$

- The impact on the spatial-resolution of instruments is
- g impact on the post-coronagraph contrast.

We need to correct the wavefront errors before they appear!

Why data-driven predictive control?

shaffert@arizona.edu

- I would like a linear model because those are easy to analyze.
- Therefore, we need to work in closed-loop because we may have non-linear behavior in open-loop control.
- The control law needs to be adaptive to be robust against changes in the atmospheric conditions.
- A model free approach is preferred, because we do not know how well we know the actual the system and the disturbances.

shaffert@arizona.edu

Time series of measurements

$e_0, e_1 \cdots e_N$

Time series of commands

$\Delta u_0, \Delta u_1 \cdots \Delta u_N$

shaffert@arizona.edu

Time series of measurements

$e_0, e_1 \cdots e_N$

Time series of commands

$\Delta u_0, \Delta u_1 \cdots \Delta u_N$

$$e_p = \begin{bmatrix} e_i \\ \vdots \\ e_{i-N} \end{bmatrix} \quad e_f = \begin{bmatrix} e_{i+M} \\ \vdots \\ e_{i+1} \end{bmatrix}$$

Now choose how much of the past and future data you want to connect.

shaffert@arizona.edu

Time series of measurements

$e_0, e_1 \cdots e_N$

Time series of commands

$\Delta u_0, \Delta u_1 \cdots \Delta u_N$

$$e_p = \begin{bmatrix} e_i \\ \vdots \\ e_{i-N} \end{bmatrix} \quad e_f = \begin{bmatrix} e_{i+M} \\ \vdots \\ e_{i+1} \end{bmatrix}$$

Now choose how much of the past and future data you want to connect.
And just add a matrix in between!

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

shaffert@arizona.edu

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

shaffert@arizona.edu

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

The past errors are necessary because the errors are what drive the system.

Data-driven subspace predictive control

shaffert@arizona.edu

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

The past commands are required because the deformable mirror does not react instantaneously. And it could have a complicated temporal response!

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

The future commands are necessary because the deformable mirror will be commanded in the future and that will have an influence on the wavefront errors.

Data-driven subspace predictive control

shaffert@arizona.edu

$$e_f = \begin{bmatrix} A & B & C \end{bmatrix} \begin{bmatrix} e_p \\ \Delta u_p \\ \Delta u_f \end{bmatrix}$$

Haffert et al. 2020 under review
Favoreel et al. 1999
B. Huang and R. Kadali 2008

The data-flow

- In principle we could feed the algorithm the slopes (or whatever measurement the wfs spits out) and let it learn the complete spatial-temporal interaction matrix directly.
- However, due to the large number of modes (>1000) it is computationally expensive.
- Therefore, we use a distributive approach where we control each individual spatial mode. This assumes little spatial-temporal cross-coupling between the modes.

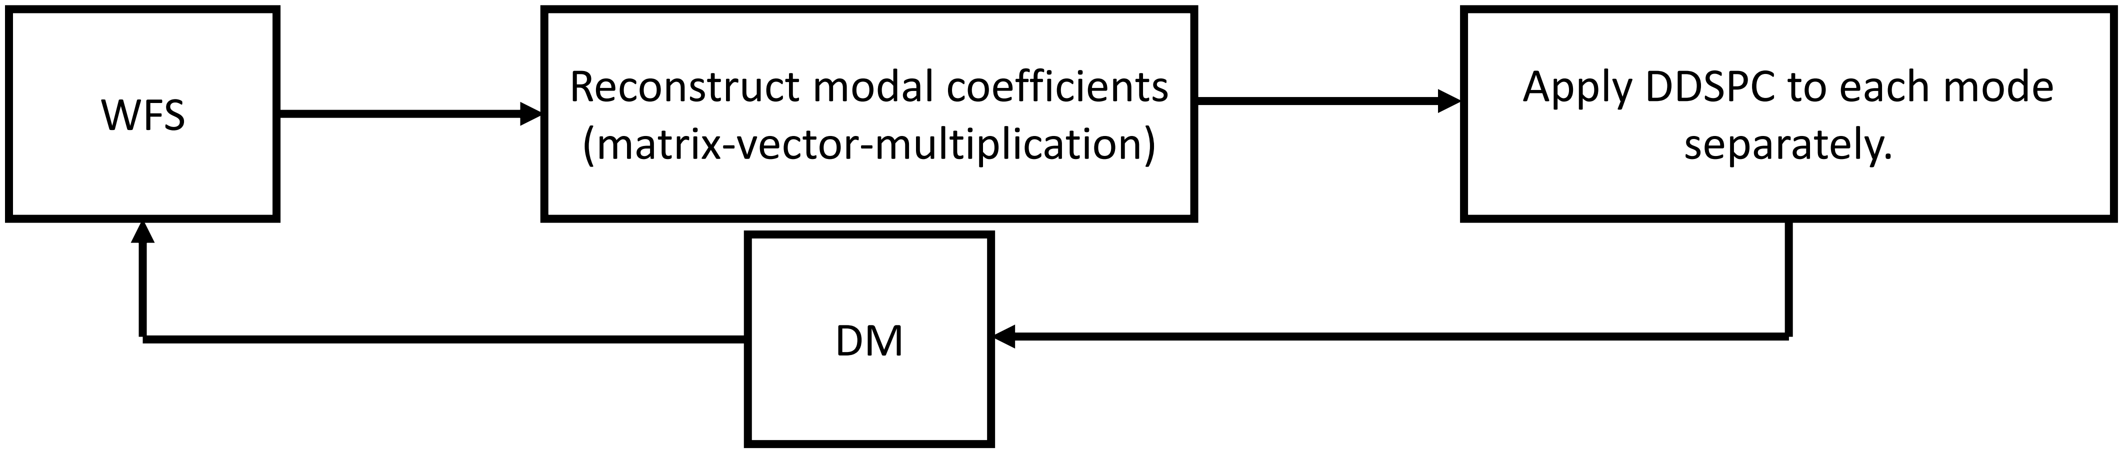

Simulation results – system overview

shaffert@arizona.edu

AO system parameters	
Telescope diameter	6.5 m
Number of actuators across the pupil	50
Sensing wavelength	0.8 um
Science wavelength	1 um
Loop speed	1000 Hz
System delay	2 frames
Wavefront sensor	Direct phase sensor

Simulation results – learning speed

shaffert@arizona.edu

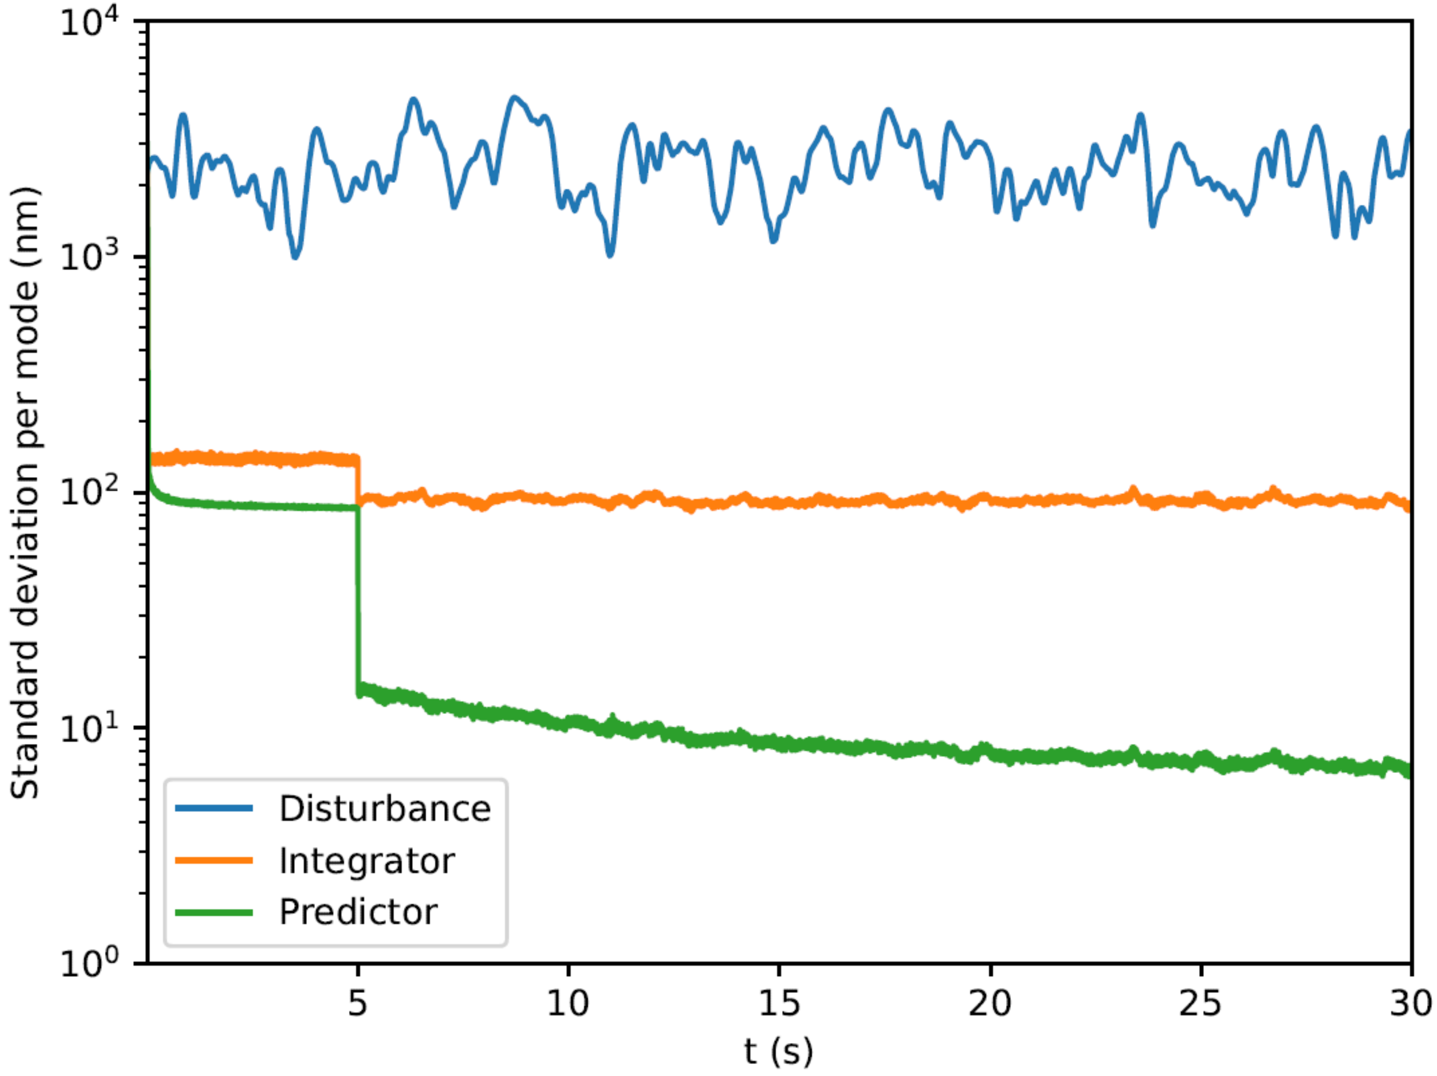

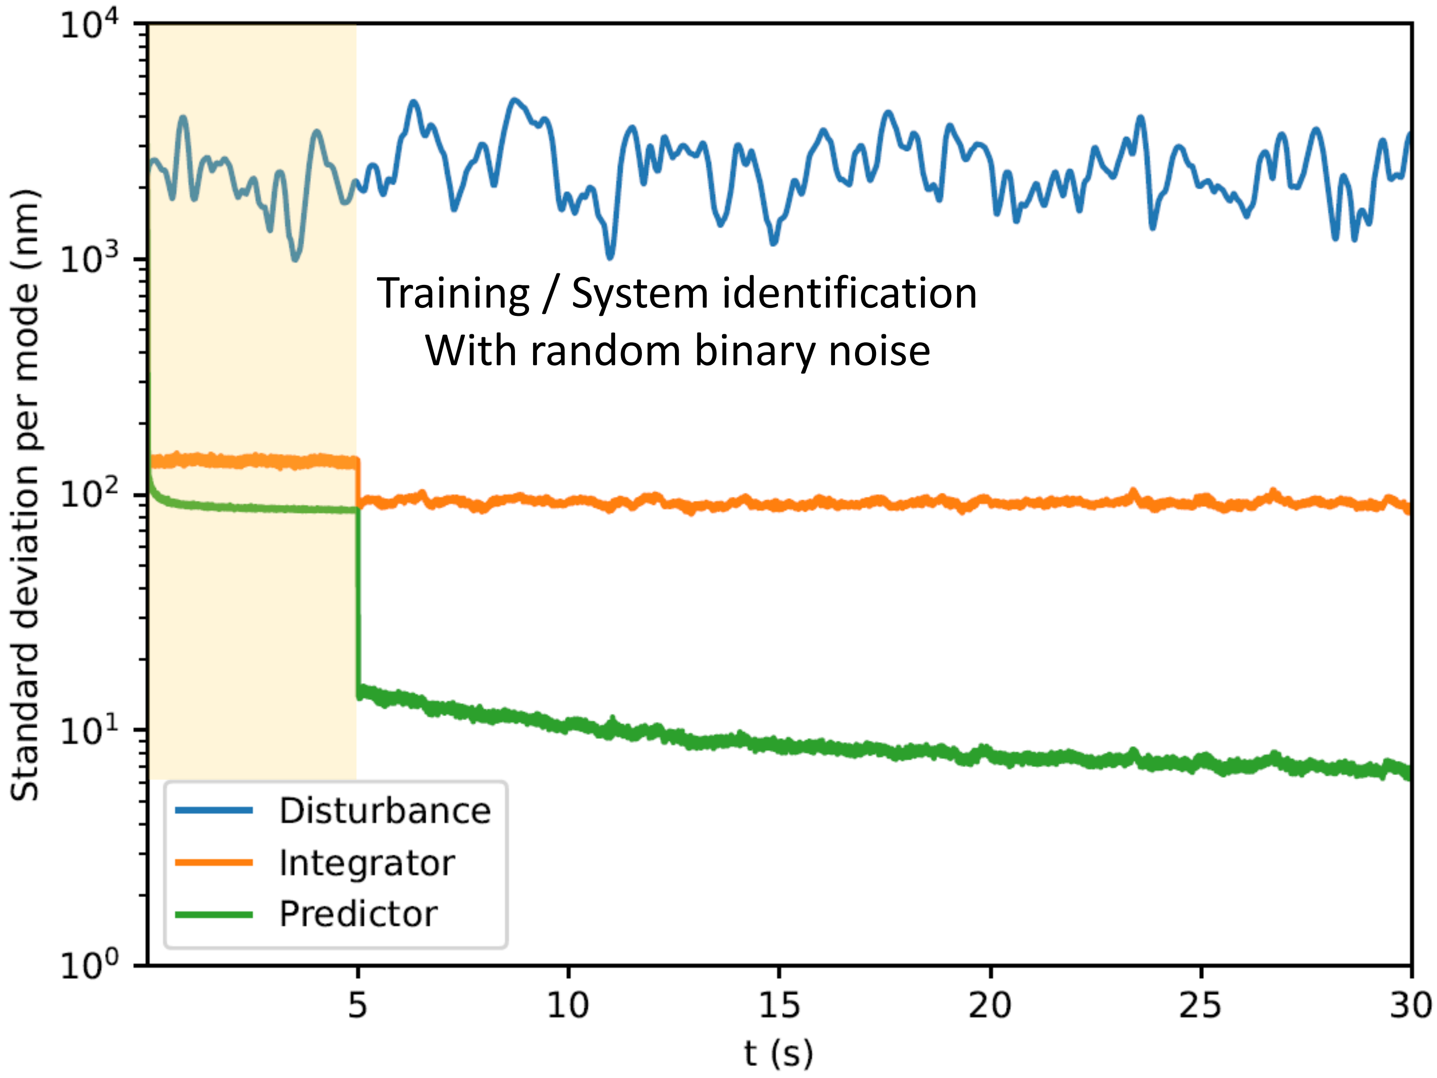

Model training / System identification

We start by adding random binary noise on the actuators to let the system learn about itself.

Simulation results – learning speed

shaffert@arizona.edu

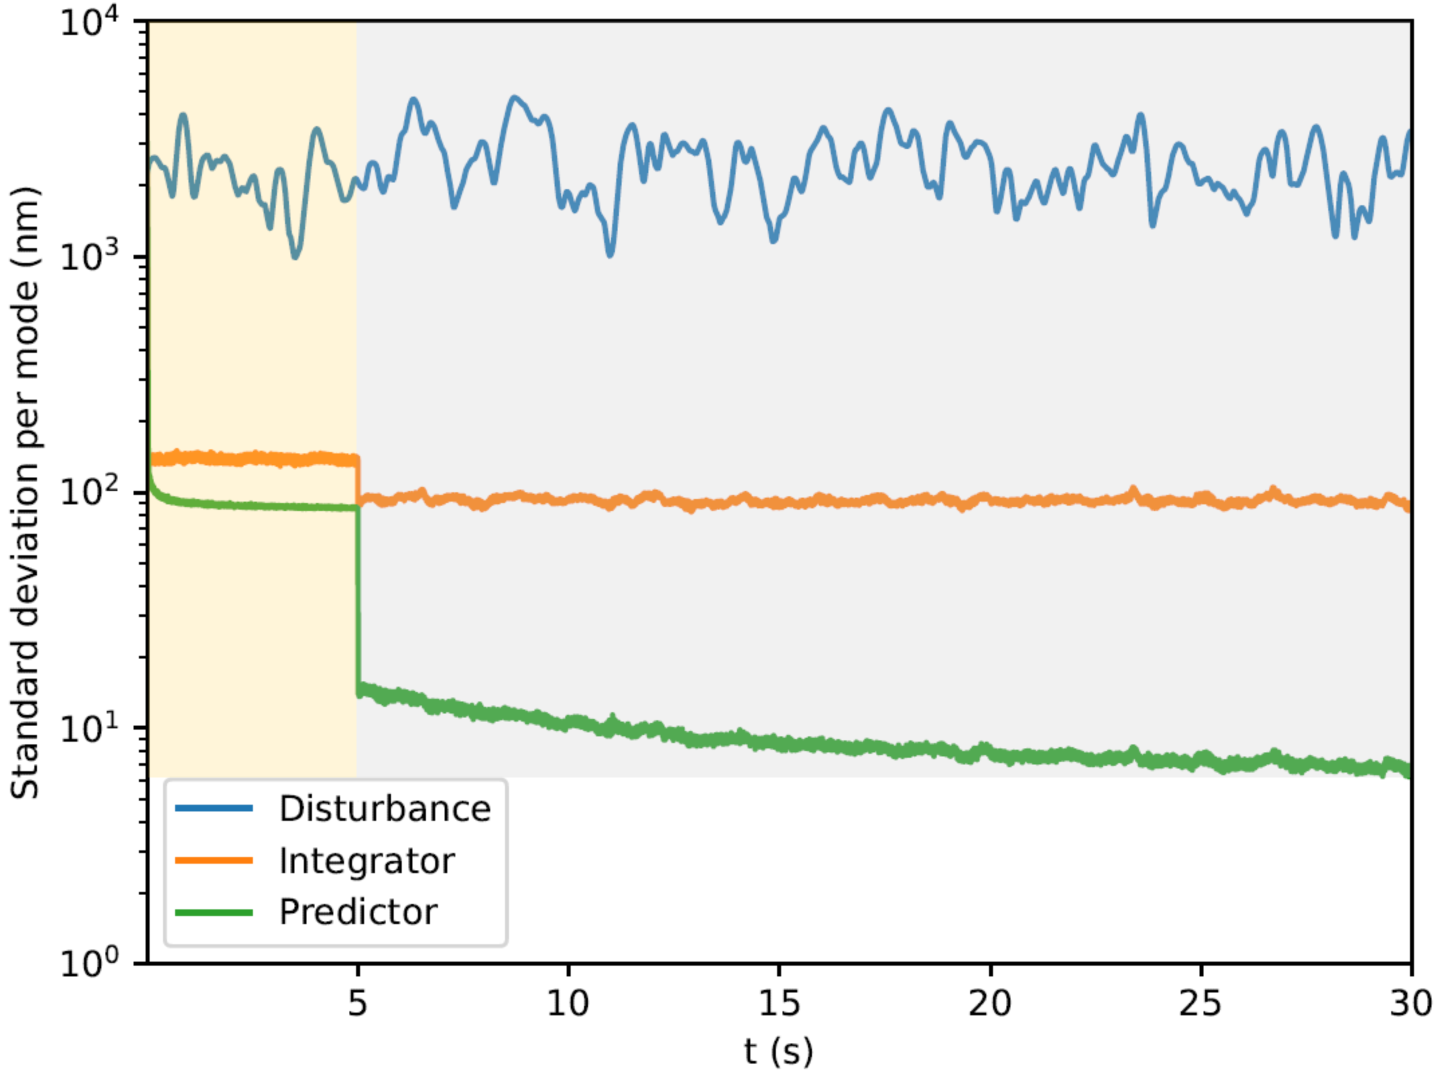

Model training / System identification
We start by adding random binary noise on the actuators to let the system learn about itself.

Long-term stability test
We run the system for 30000 iteration.
And even after 30000 iterations we still see improvement in the rejection.

The online learning improves as $\frac{1}{\sqrt{t}}$. This is the statistically optimal learning rate if data samples are uncorrelated.

The learning rate puts a limit on how fast the model can adapt to changing atmospheric conditions.

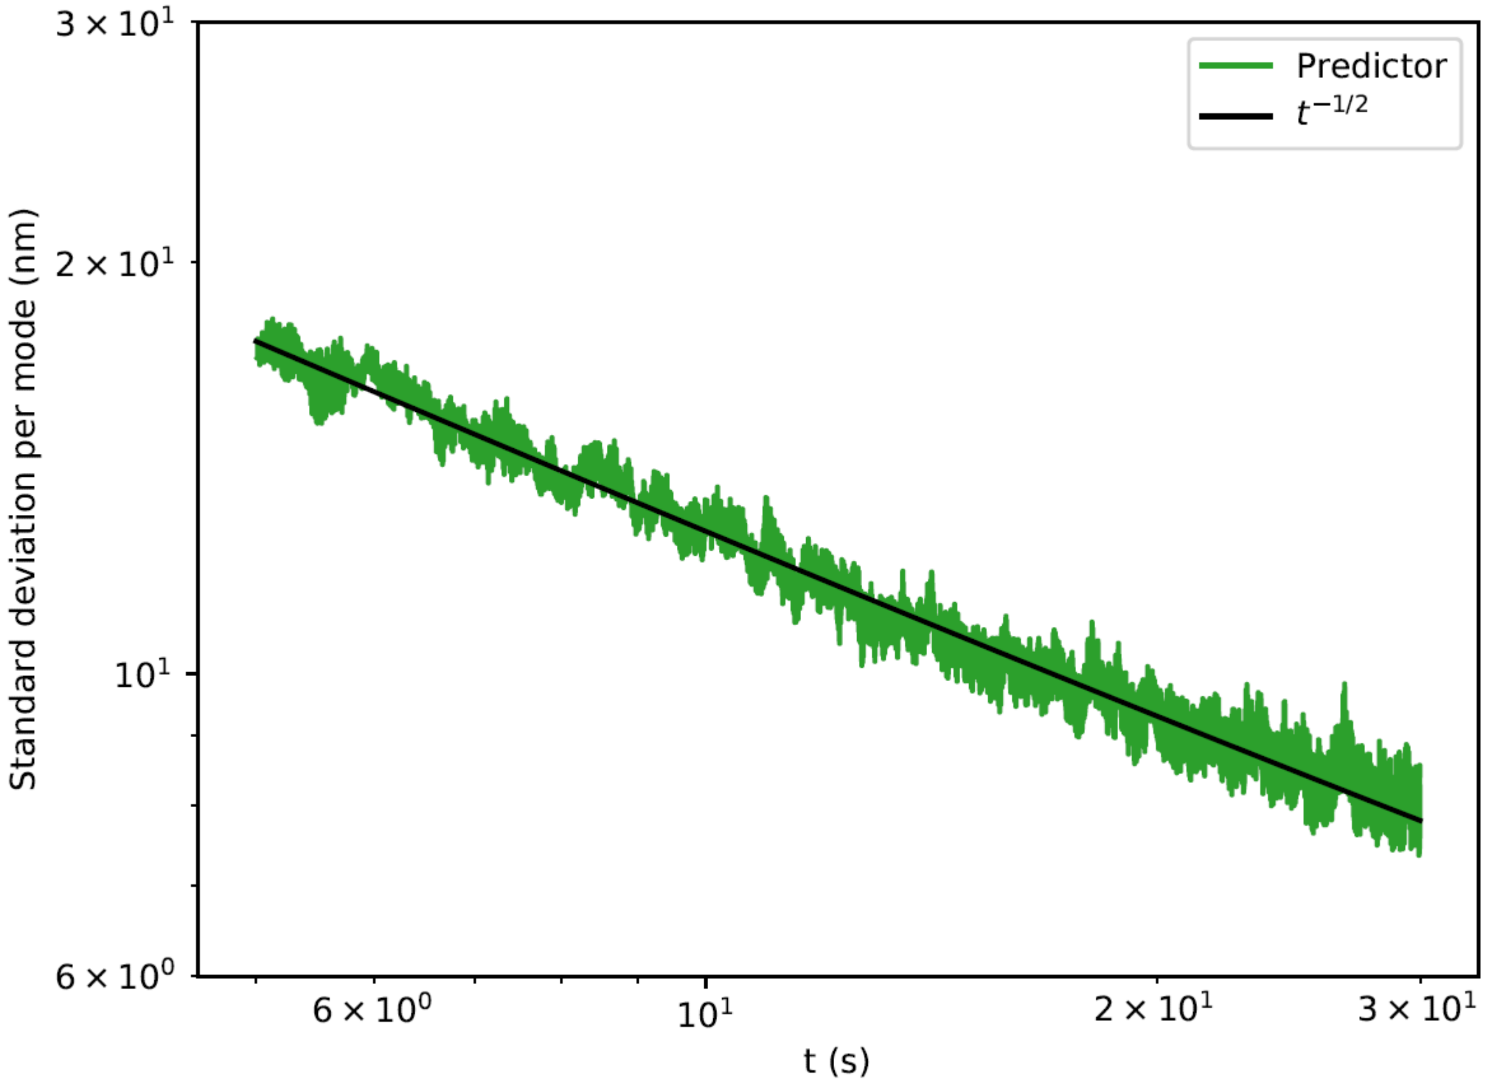

Simulation results – full atmosphere

shaffert@arizona.edu

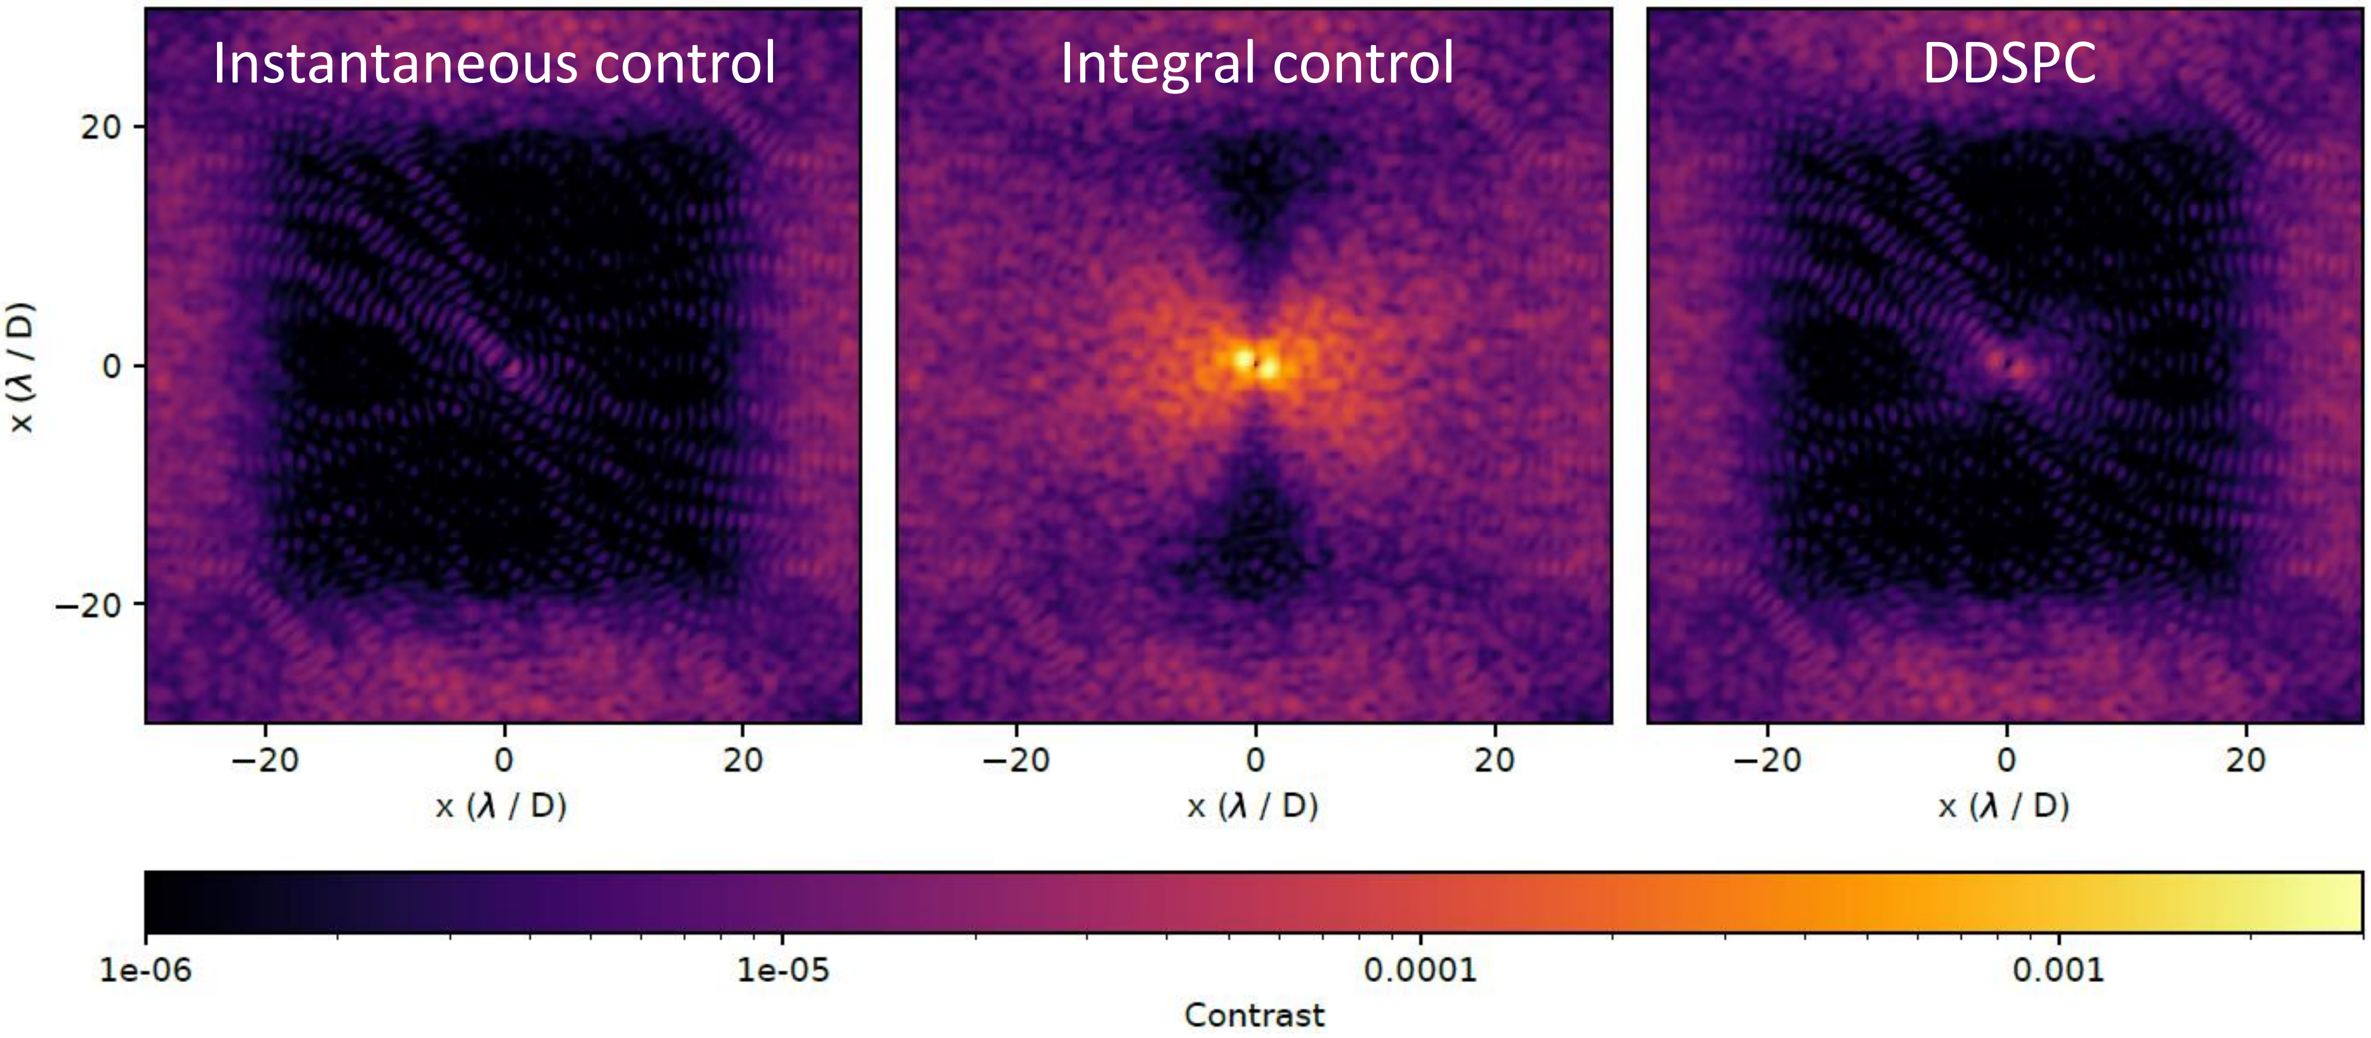

Simulation results – full atmosphere

shaffert@arizona.edu

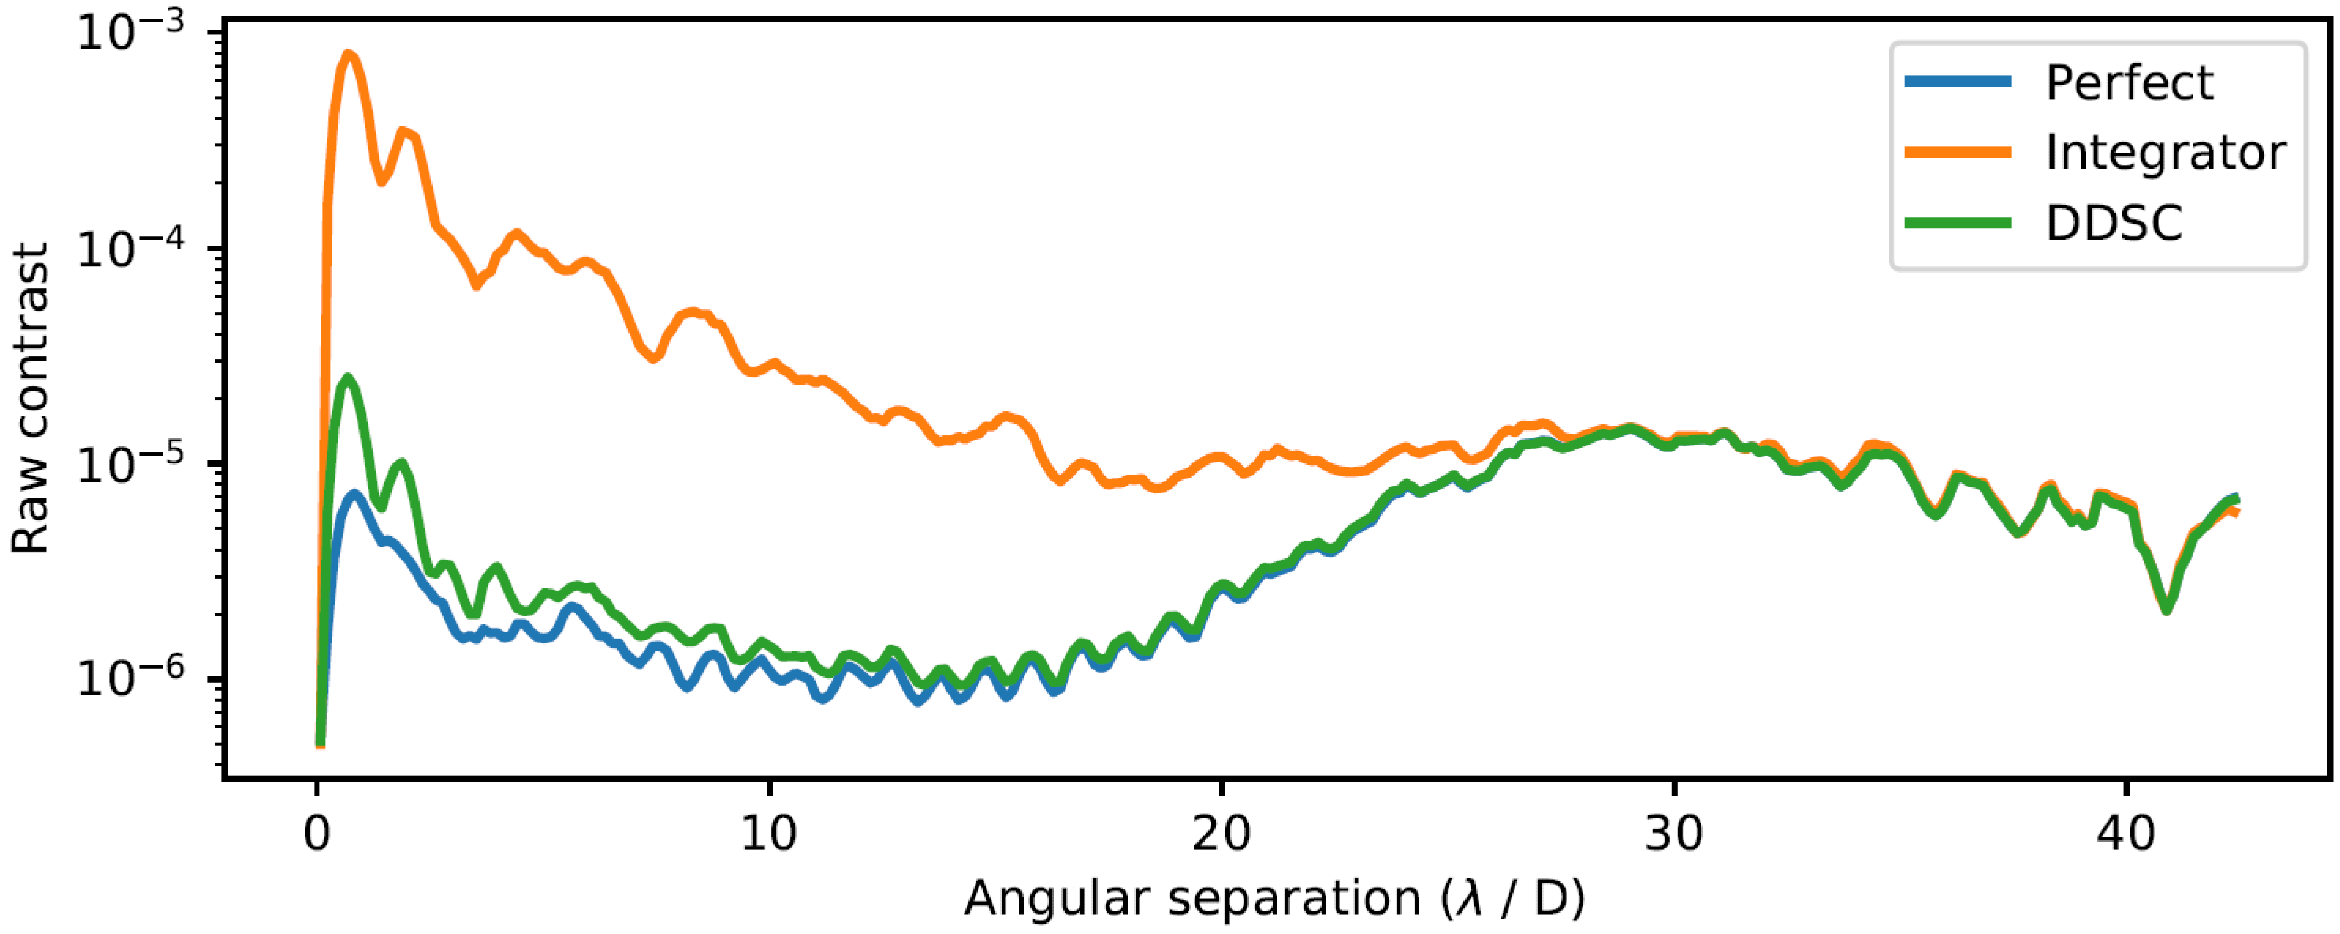

Simulation results – full atmosphere

shaffert@arizona.edu

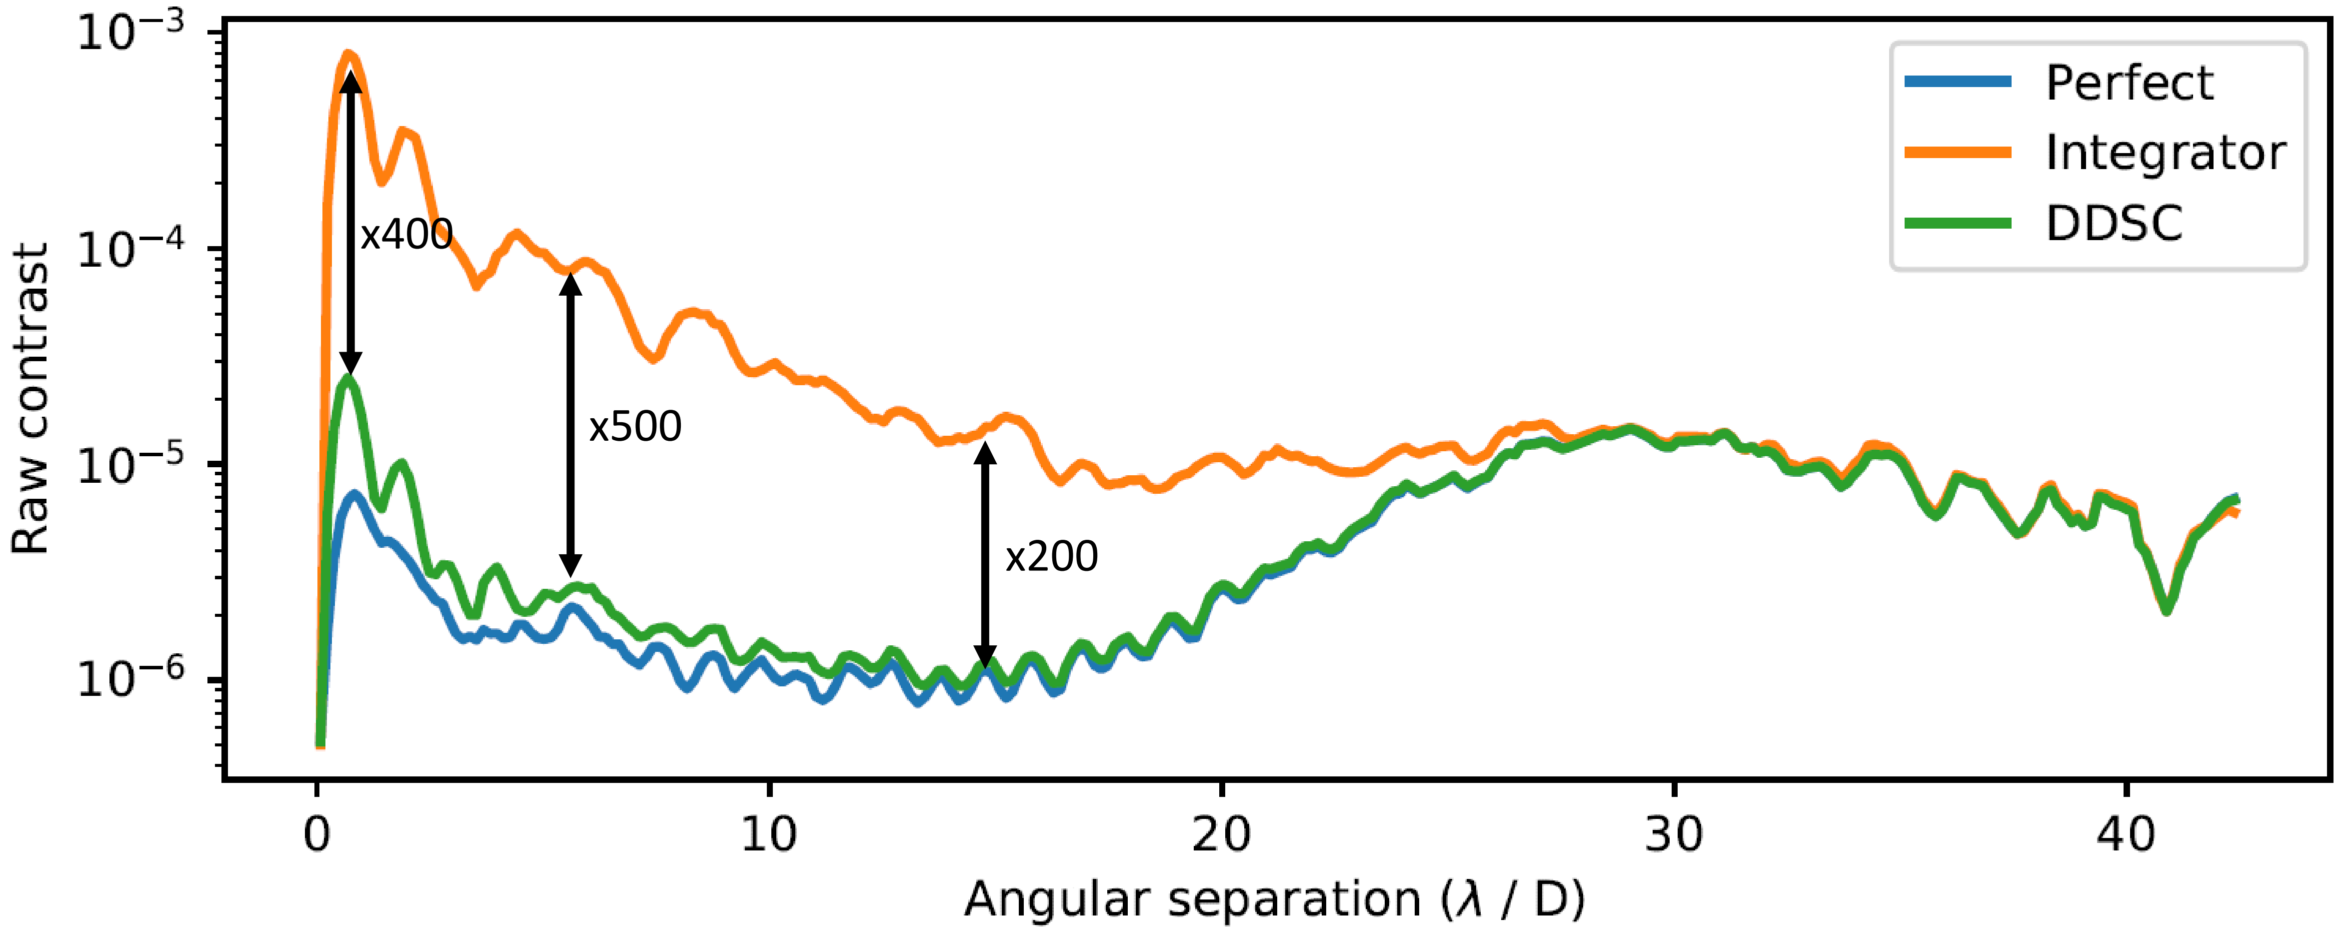

Simulation results – full atmosphere

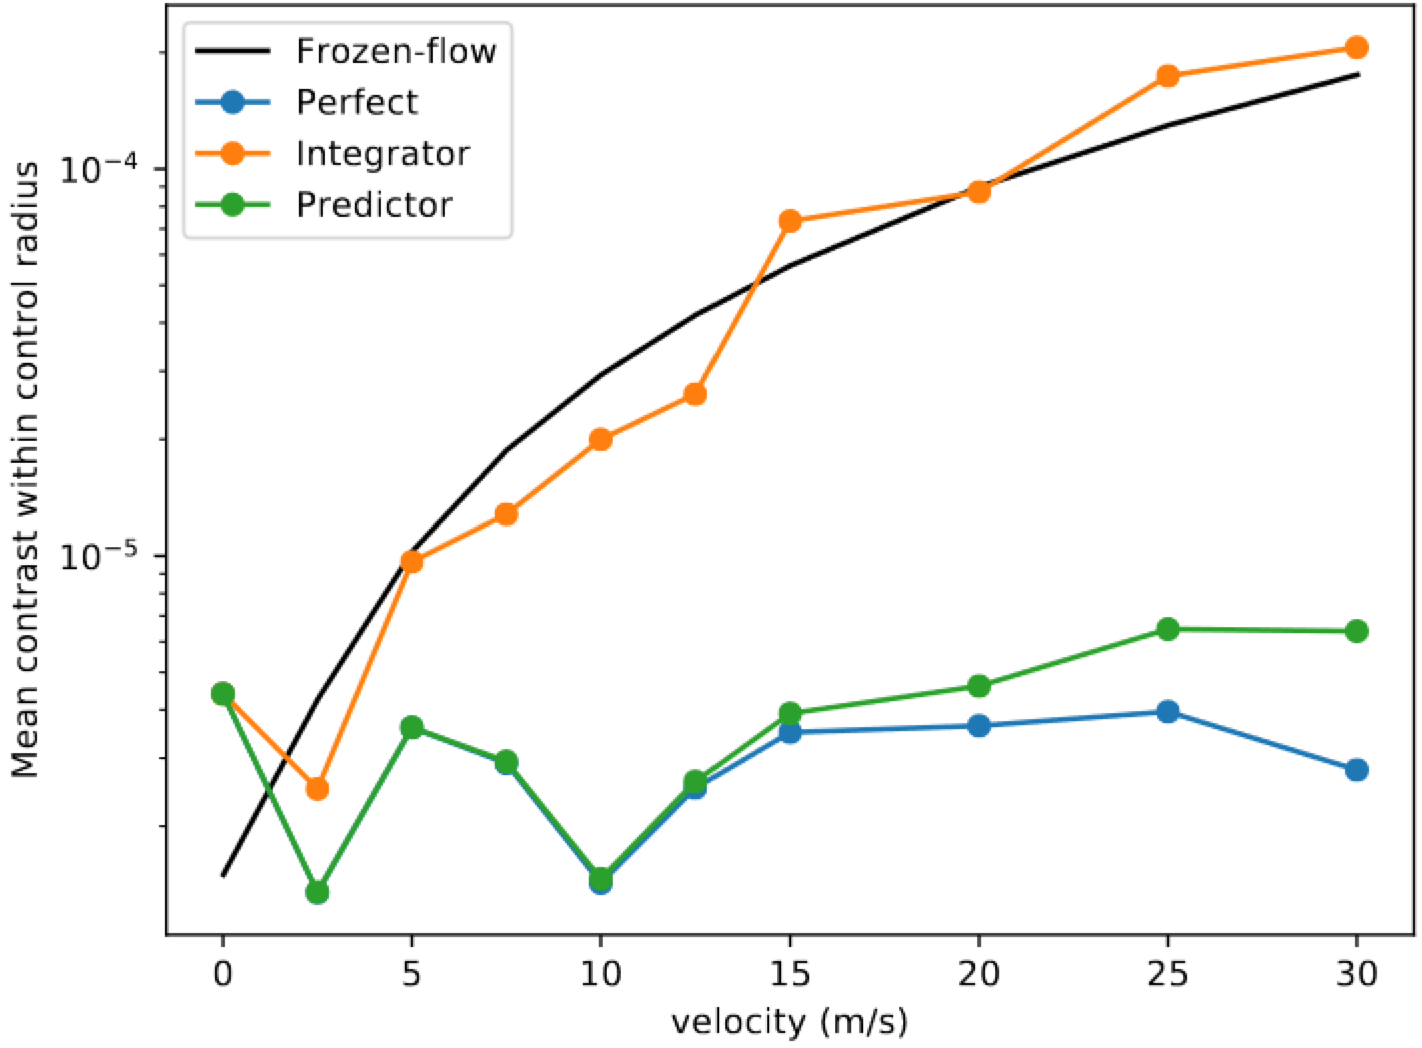

shaffert@arizona.edu

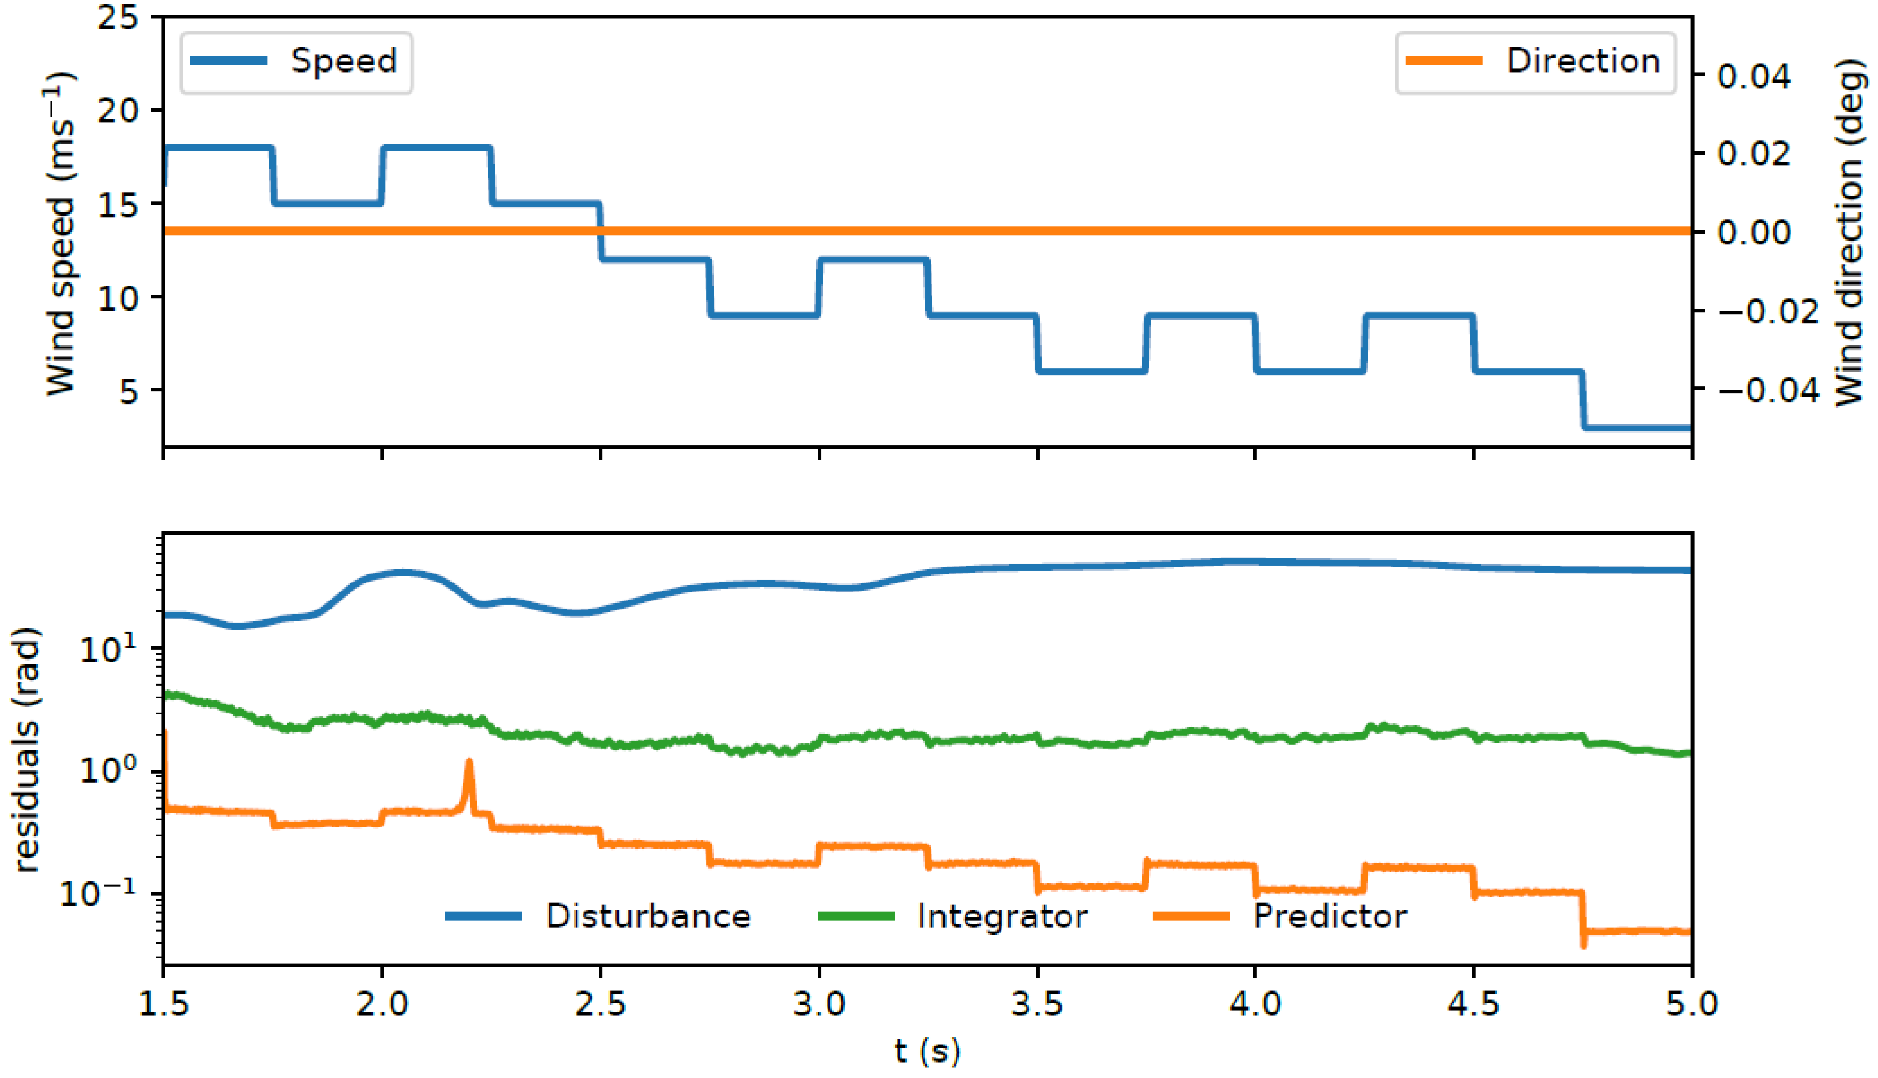

Simulation results – non-stationarity

shaffert@arizona.edu

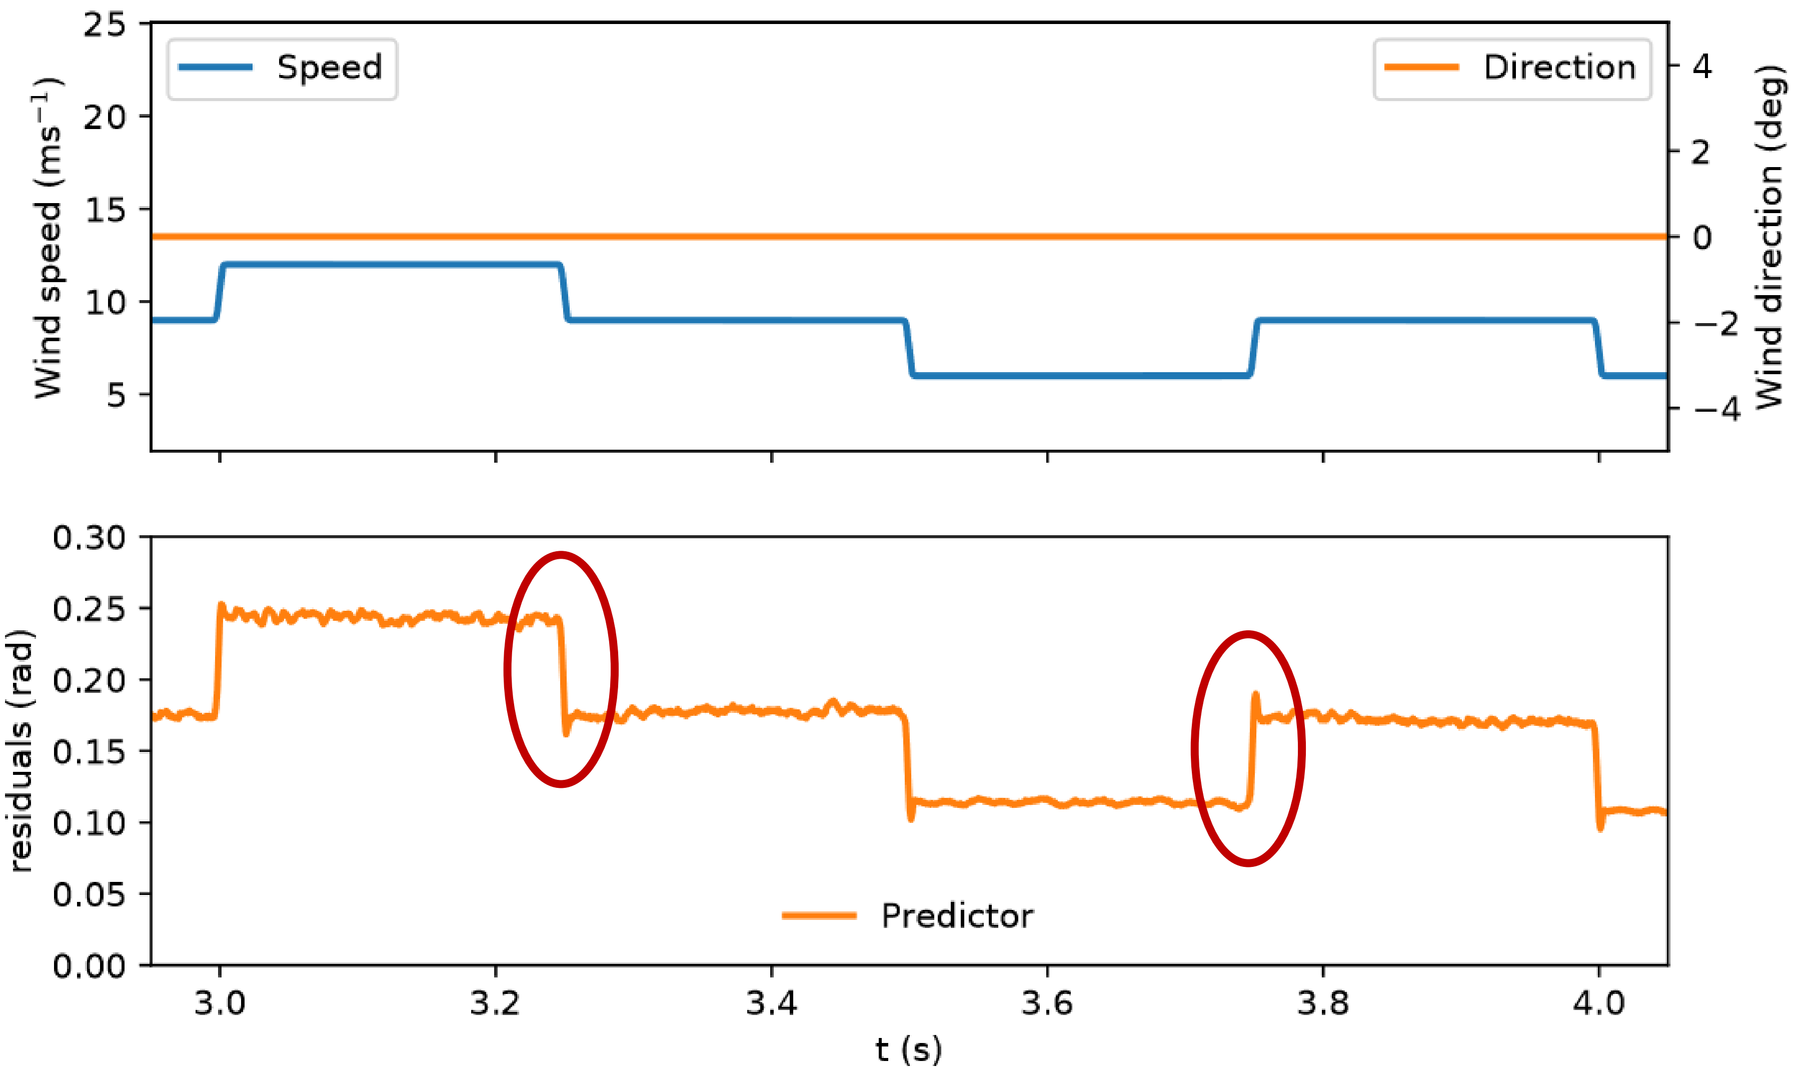

Simulation results – non-stationarity

shaffert@arizona.edu

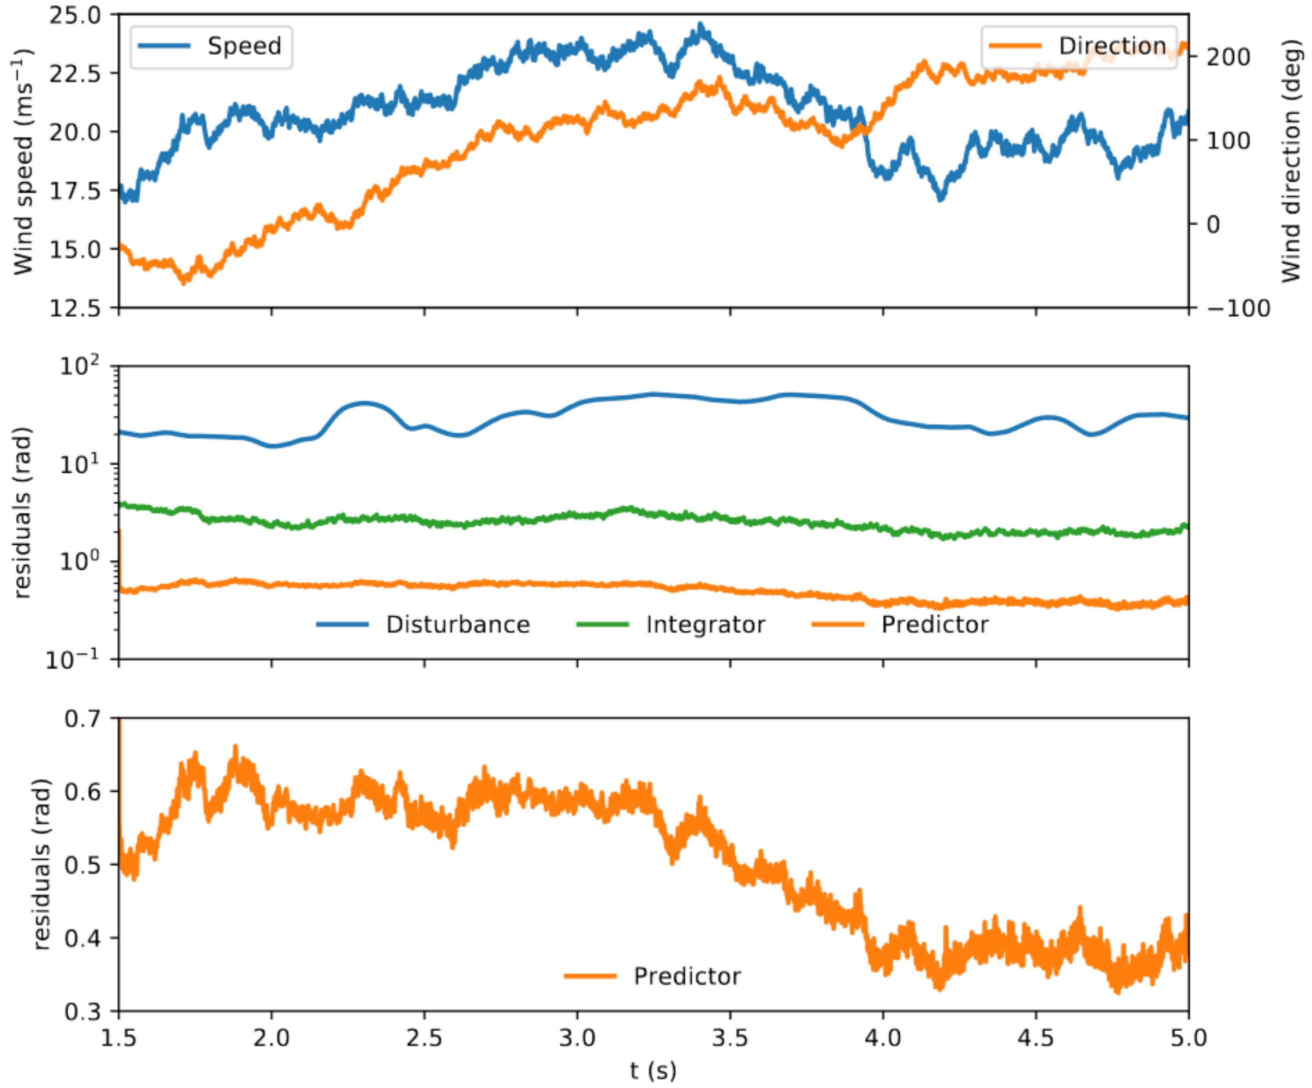

Simulation results – non-stationarity

shaffert@arizona.edu

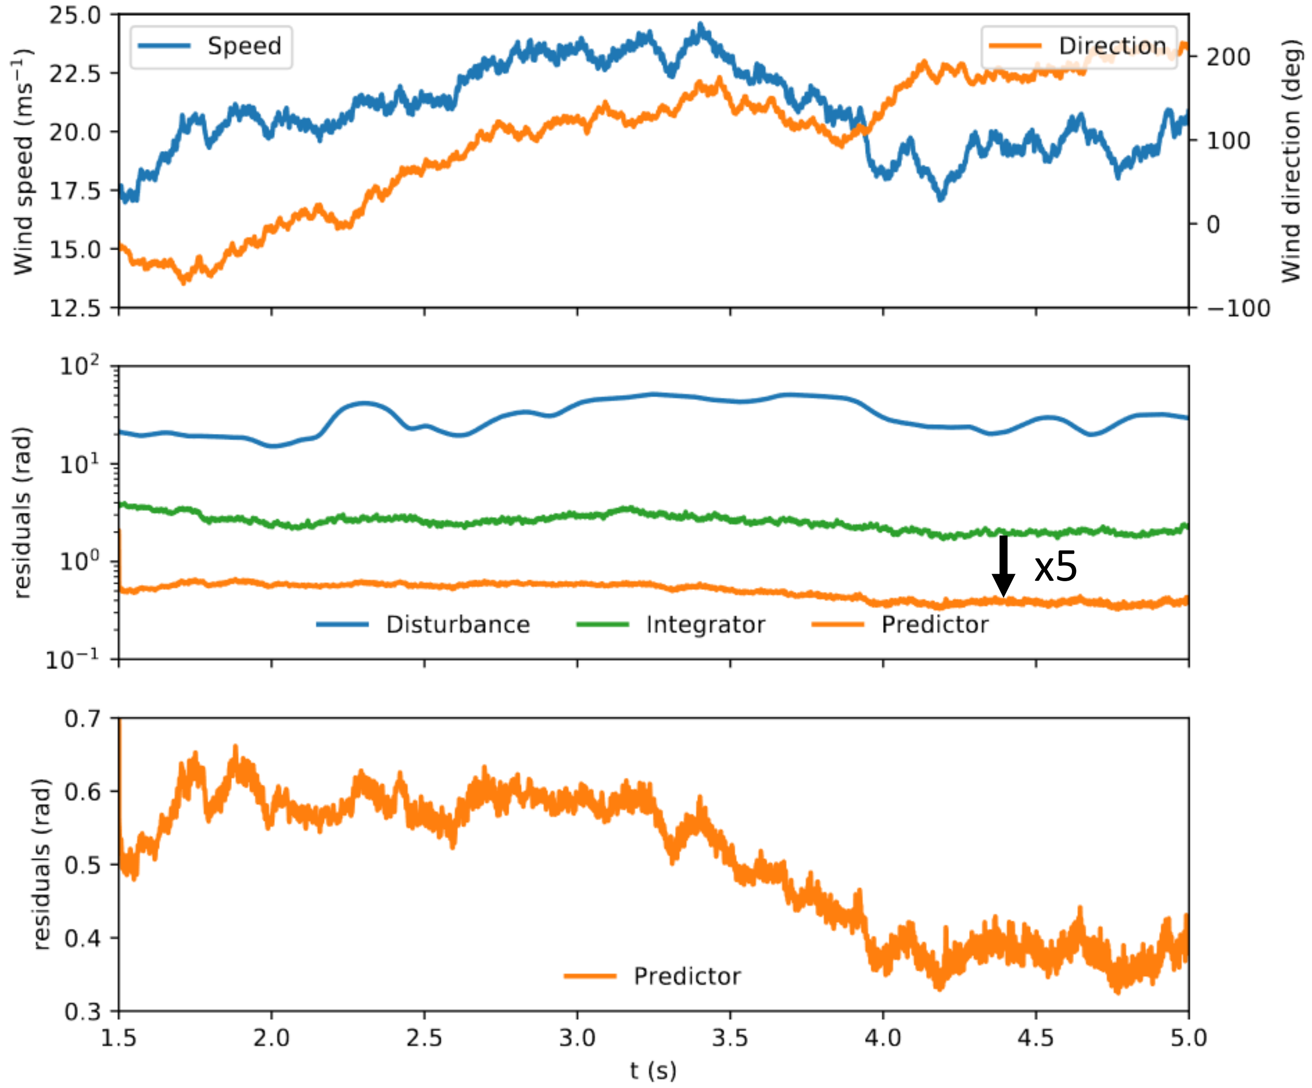

- Improvement even with a completely non-stationary atmosphere. But, with less additional rejection.
- Starlight rejection improvement is roughly $5^2 = 25$ in the very strong non-stationary case.

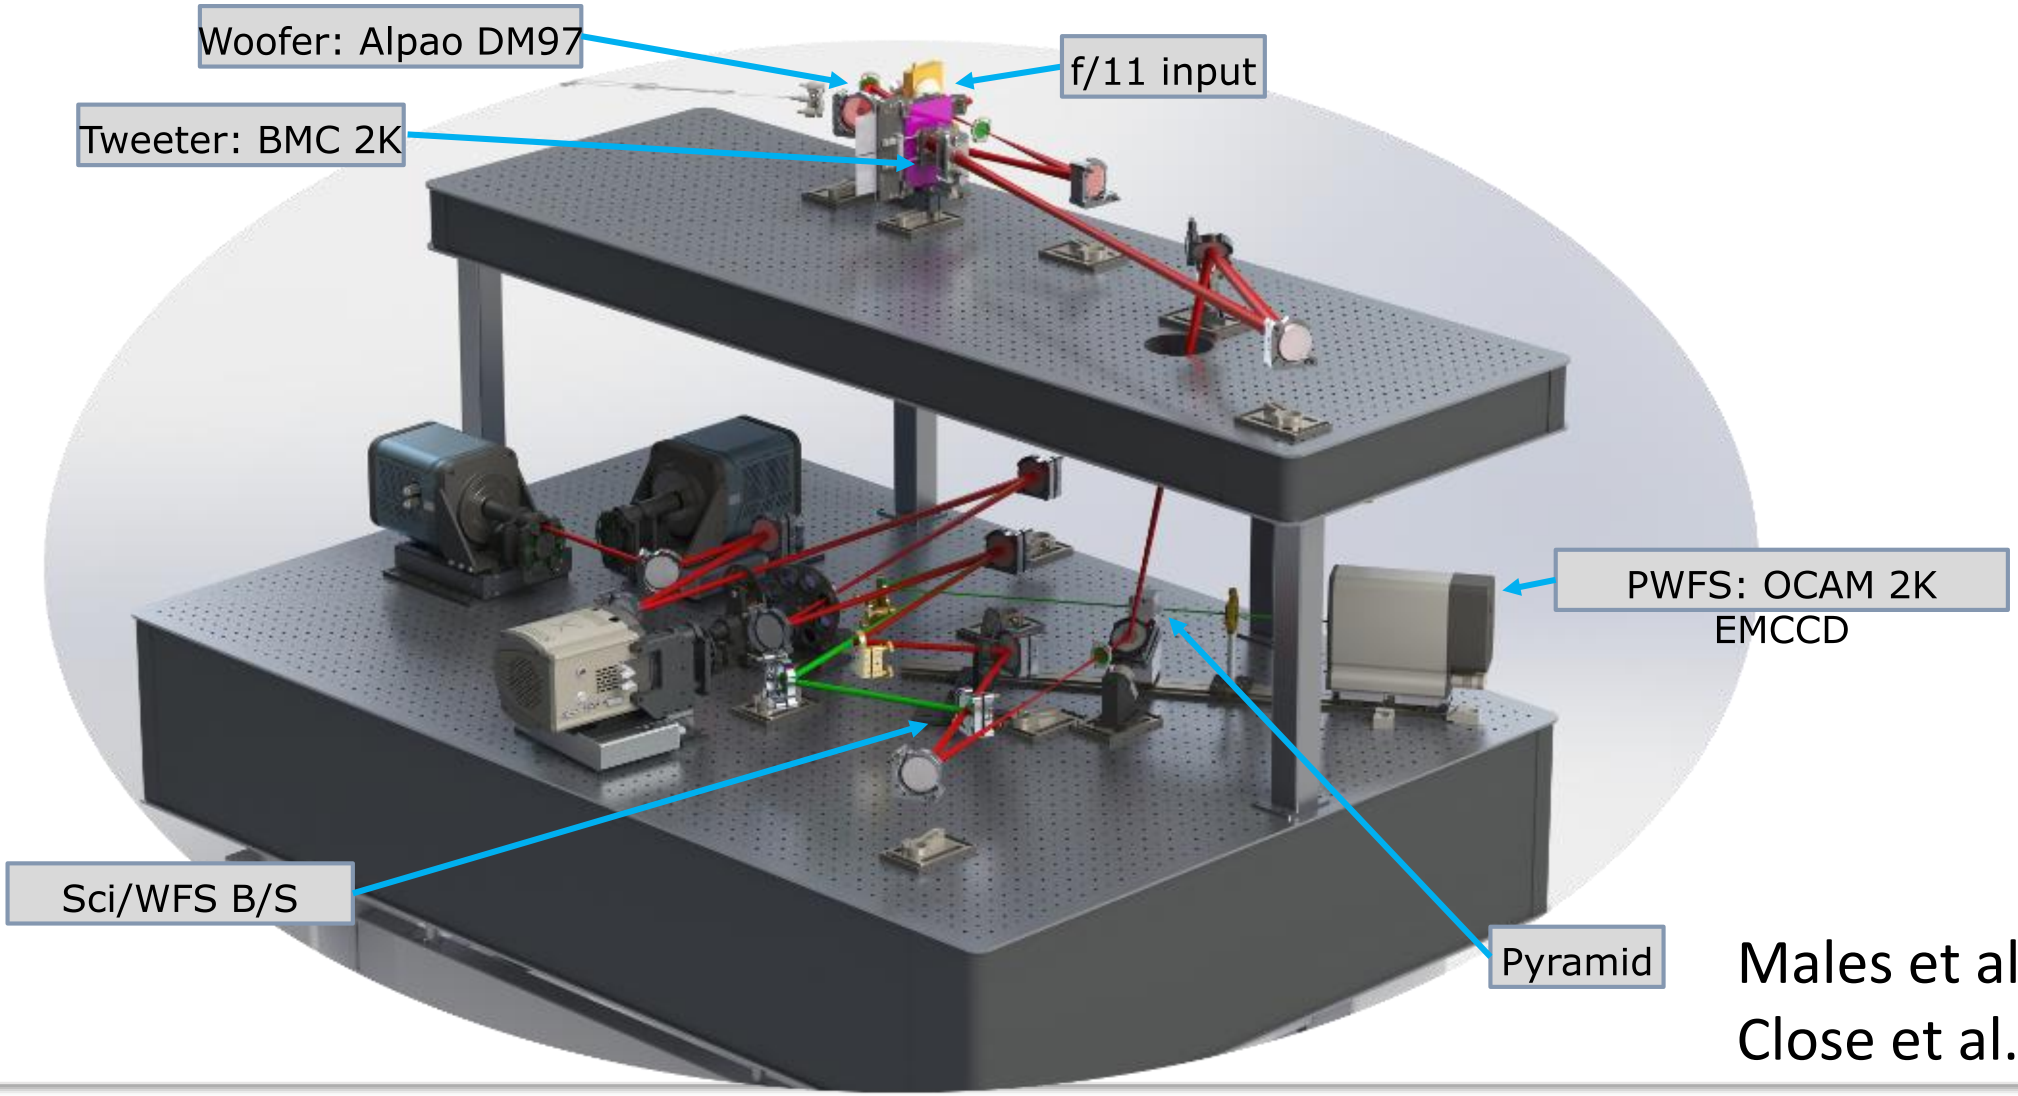

Males et al. 2018

Close et al. 2018

MagAO-X – testing the controller

shaffert@arizona.edu

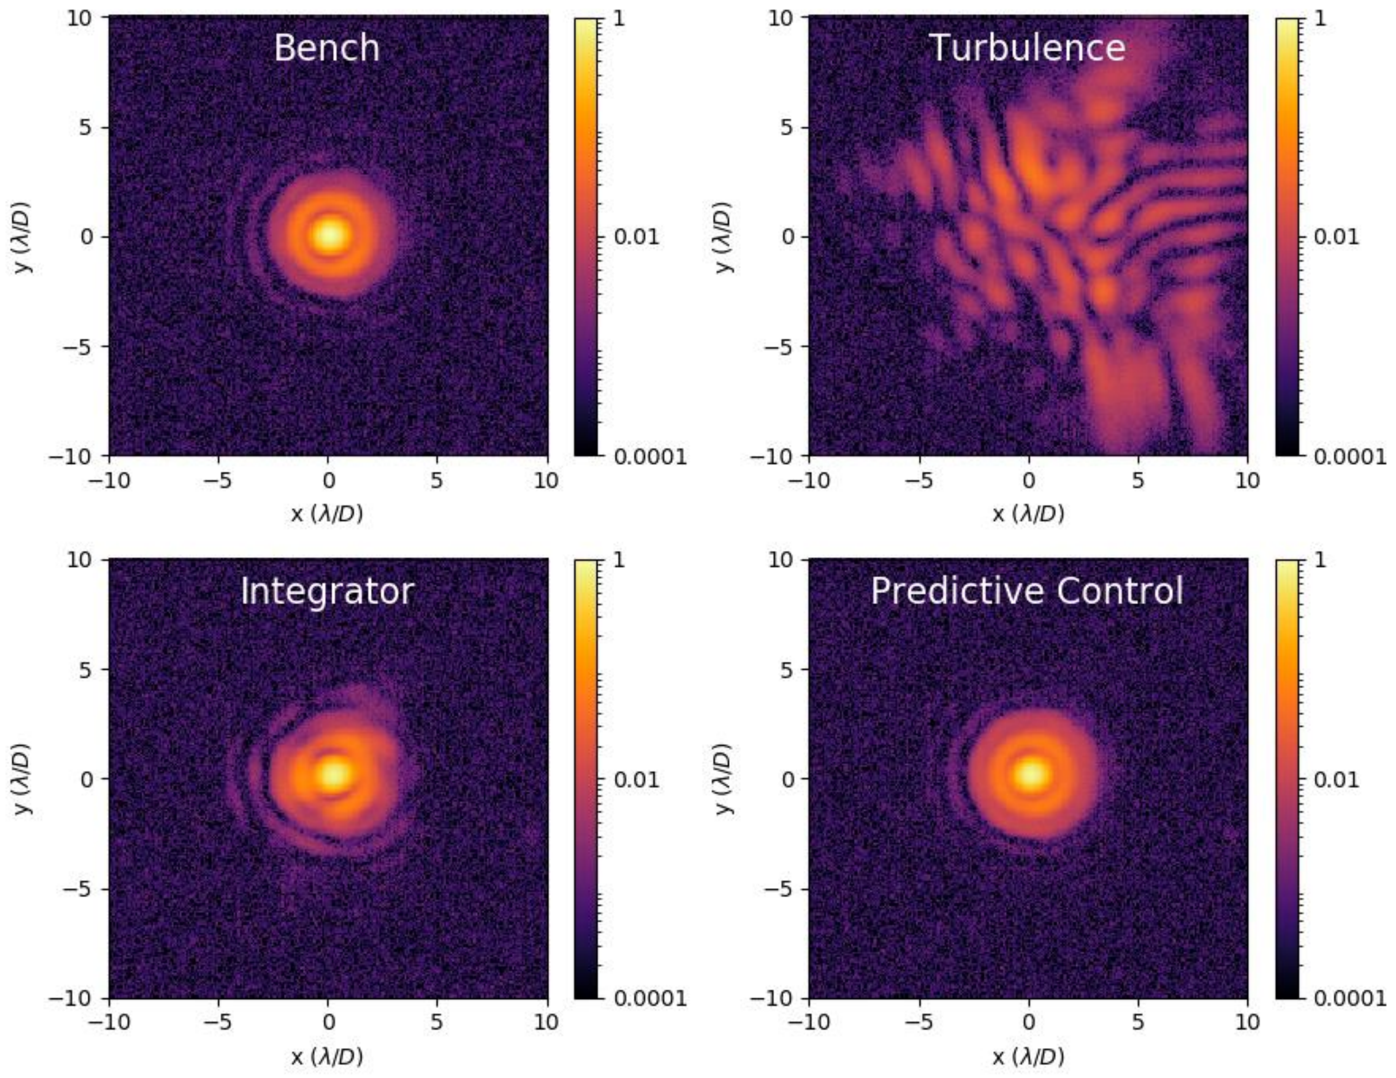

- Using the Alpao 97 DM
- 86 modes could be controlled
- We are using the DM to create and control the turbulence.

MagAO-X – contrast gain estimation

shaffert@email.arizona.edu

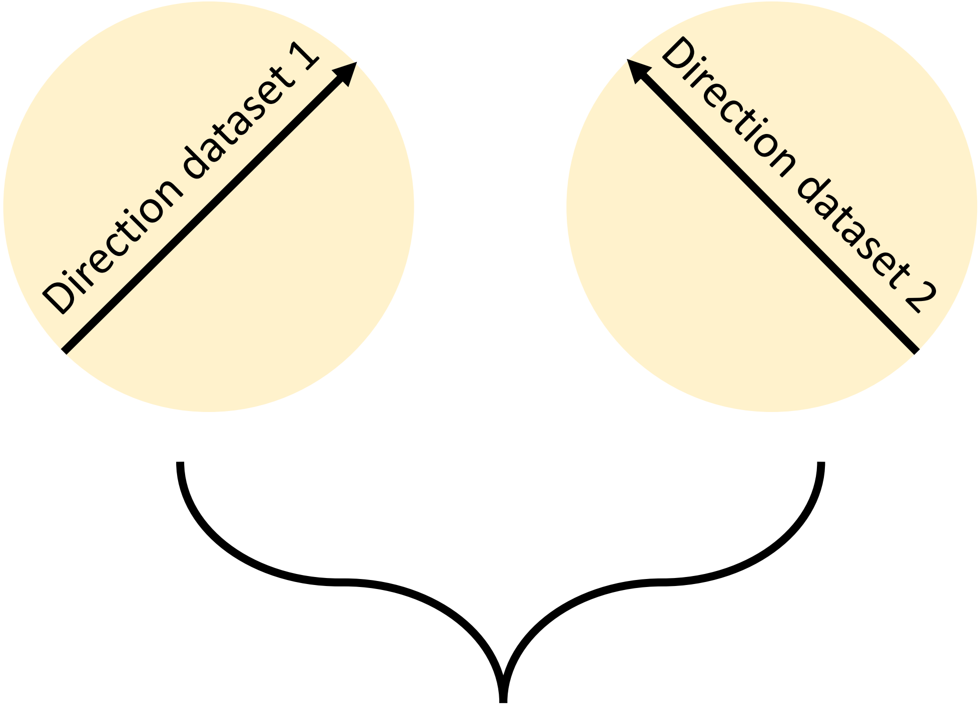

Difference shows temporal stability.

MagAO-X – contrast gain estimation

shaffert@arizona.edu

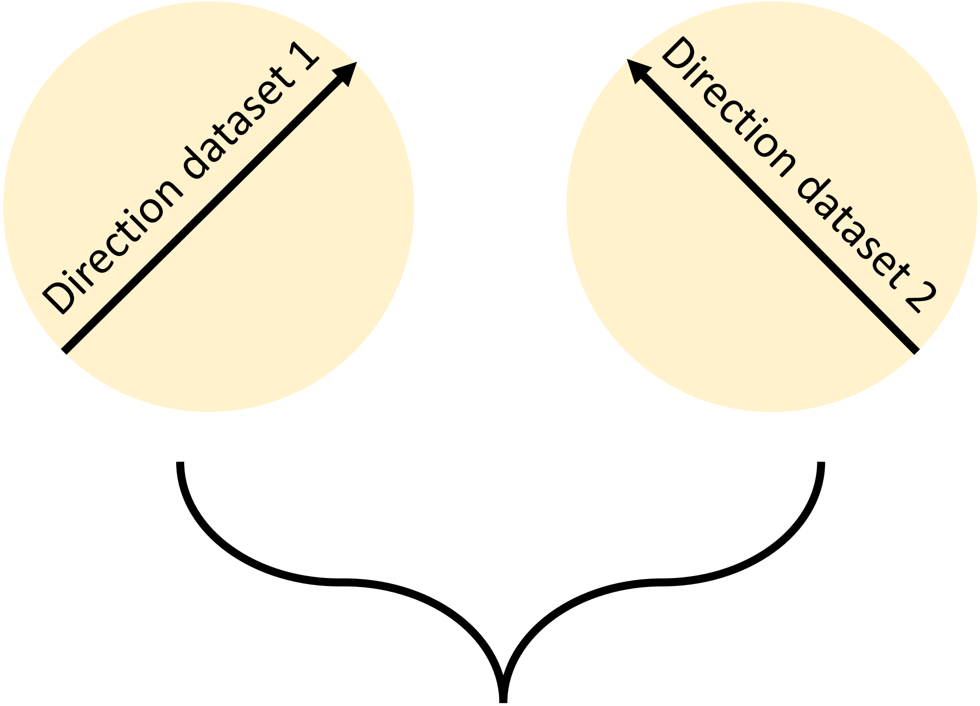

Difference shows temporal stability.

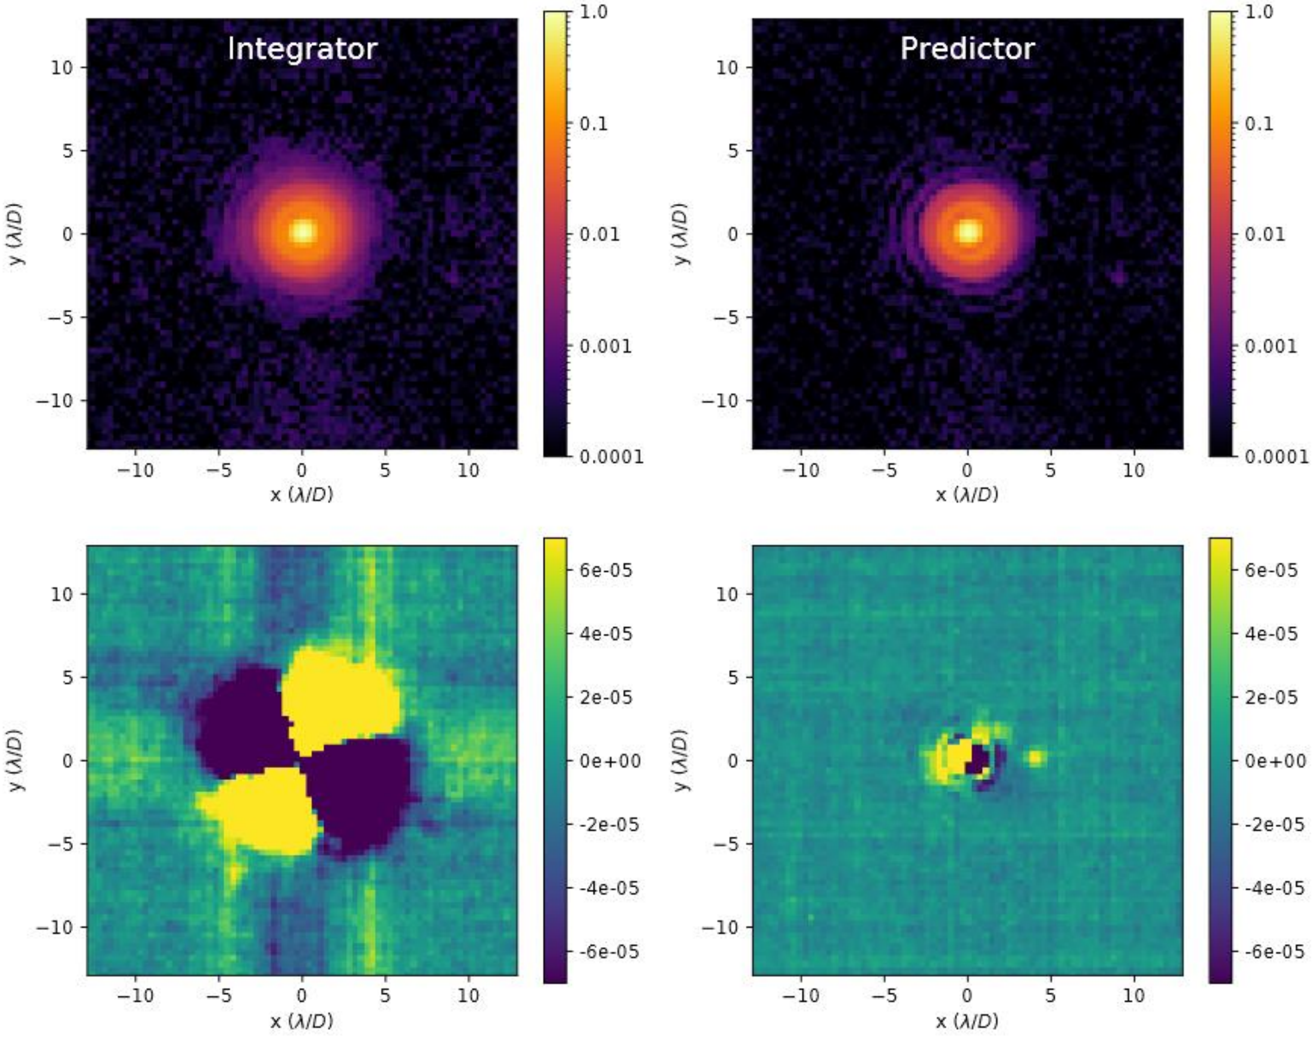

shaffert@arizona.edu

MagAO-X – contrast gain estimation

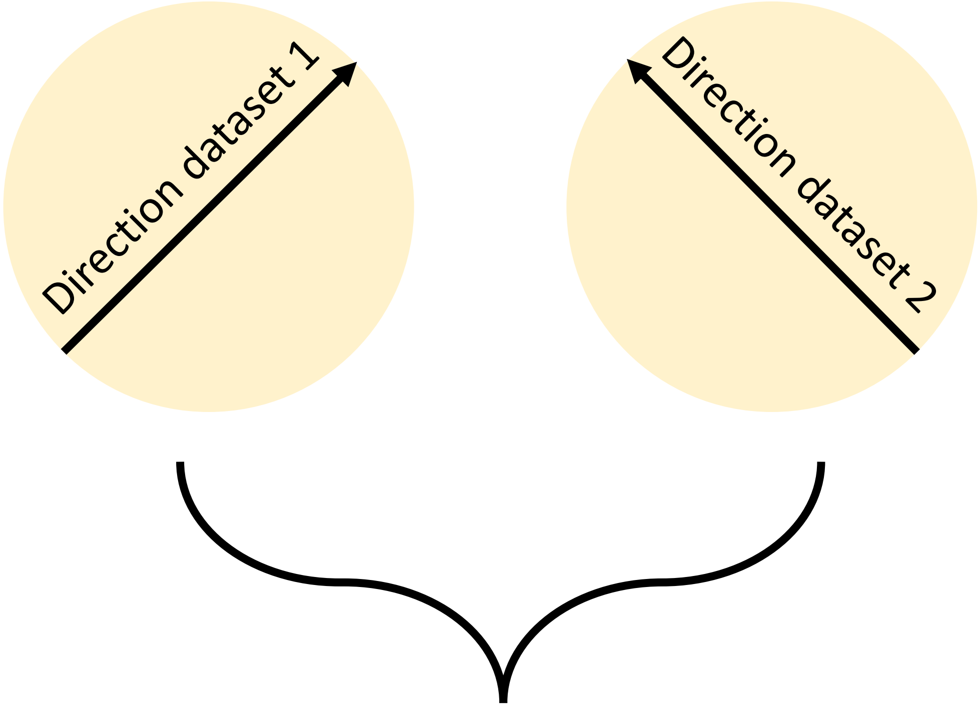

Difference shows temporal stability.

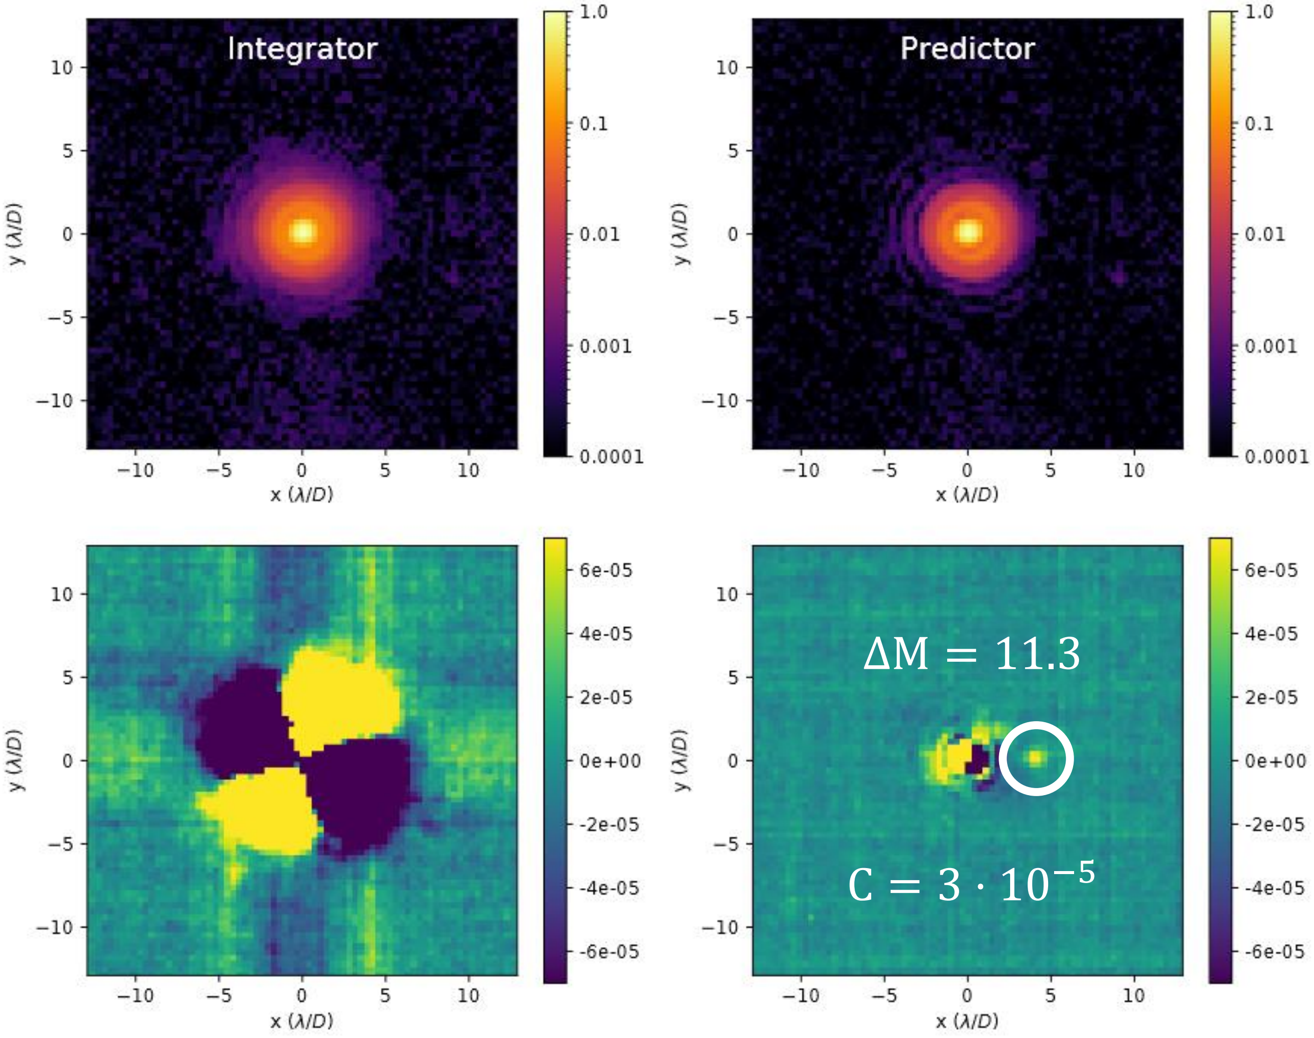

Conclusion

shaffert@arizona.edu

- We have derived and implemented a new distributed predictive controller for high-contrast imaging.
- Simulations show 2 orders of magnitude improvement for frozen flow, but only 1 order of magnitude for the non-stationary turbulence.
- We have tried the algorithm in the lab and have been able to successfully close the loop!

Outlook

shaffert@arizona.edu

- Current work is focused on implementing a high-speed version in the RTC of MagAO-X.

- Figure out how much additional effects such as photon noise and boiling of the atmosphere impact the predictability.

- We need to investigate what the optimal hyper parameters of the controller are.

Outlook

shaffert@arizona.edu

- Current work is focused on implementing a high-speed version in the RTC of MagAO-X.
- Figure out how much additional effects such as photon noise and boiling of the atmosphe

Questions?

- We need to investigate what the optimal hyper parameters of the controller are.