

The Potential of Least Squares Vector Field Fitting in the Improvement of DESI Fiber Positioner Performance

by

Brandon Freudenstein

A senior thesis submitted in partial fulfillment
of the requirements for the degree of
Honors Bachelor of Science
(Physics)
in The University of Michigan
2018

Advisers:

Gregory Tarlé
Michael Schubnell

ACKNOWLEDGEMENTS

I would like to thank my advisers Gregory Tarlé and Michael Schubnell, both for aiding me in the development of this thesis project, and for guiding me as I begin my scientific career. I would also like to thank the entire DESI team both at the University of Michigan and at Lawrence Berkeley National Laboratory, who's dedication to DESI made this thesis and in fact the entire experiment possible. In particular, my colleague Kevin Fanning has provided constant guidance on the techniques of testing positioners. Finally, I would like to thank all of my friends and family who have supported me throughout my time as an undergraduate.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
ABSTRACT	iv
CHAPTER	
I. Introduction	1
II. Fiber Positioner Quality Control Testing	4
2.1 The Positioner	4
2.2 The Test Stand	6
2.3 The Testing Software	8
III. Least Squares Vector Field Fitting	11
IV. Analysis of Lifetime Data	14
4.1 Overview and Procedure	14
4.2 Results	15
V. Analysis of Random Point Testing	21
5.1 Overview and Procedure	21
5.2 Results	23
VI. Conclusion	25

ABSTRACT

The Potential of Least Squares Vector Field Fitting in the Improvement of DESI
Fiber Positioner Performance

by

Brandon Freudenstein

The Dark Energy Spectroscopic Instrument (DESI) is currently being constructed to collect cosmological redshifts from a large portion of our sky. It will accomplish this task by aiming optical fibers at a predetermined list of targets using 5,000 robotic positioners. The devices currently being produced operate within specification, but there is much to gain from decreasing the number of movements it takes for these positioners to achieve the desired accuracy; they are allowed a blind move without feedback, and then four correction moves using feedback from a camera. Decreasing the needed submoves could be achieved by performing a least squares fit on the positioners' performance history, and harnessing any systematic error present to subtract their past mistakes from their future movement attempts. By performing a hypothetical correction of this nature, using data from a test meant to operate positioners far past their expected lifetime, it was discovered that in the regular testing sequence significant systematic error is present. By using only a fit on a 192-point test dispersed across the positioner's range and subtracting this from the following test for 59 tests in a row, on average 7 of 11 positioners would have reached specification after

only the blind move. While initially encouraging, the same target sequence was used for all of these tests. After repeating this technique on tests with a scrambled target order, the success rate decreased and no positioners reached specification this quickly. There was still uniform improvement on the blind move however, with many positioners likely a single submove away from the required accuracy; therefore, a physical implementation of this correction should be designed and tested for potential use on the actual instrument. Producing results as good as those found in this study has the potential to significantly increase the speed of the survey, improving DESI as an experiment.

CHAPTER I

Introduction

The Dark Energy Spectroscopic Instrument (DESI) is being constructed to conduct an unprecedentedly vast survey of redshifts, aiming to collect the spectra of 30 million receding galaxies and quasars. This data, effectively a 3-dimensional map of a significant portion of our universe, will be used after the conclusion of the survey to find novel constraints on various problems in cosmology and to provide reference for future astronomers. The primary motivation for DESI lies in exploring the nature of dark energy, and in particular to discover what redshift-space distortions suggest about the evolution of large-scale structure. In addition to this, DESI will provide constraints on modified gravity and inflation, allow for the first measurement of the sum of neutrino masses, and include a survey of the Milky Way that among other benefits will allow precise measurements of this galaxy's dark matter halo. A more technical review of the experiment's science goals can be found in the first half of its final design report [1].

The experiment will be conducted using the Mayall telescope, outfitted with new collection optics, at the National Optical Astronomy Observatory at Kitt Peak, Arizona. A particularly striking aspect of DESI is the 5,000 robotic positioners that will be used to individually move the fiber-optic cables to a list of positions corresponding to astronomical targets. These positioners must be able to, on average,

bring a fiber to within $5\ \mu\text{m}$ of a specified location. This accuracy is not achieved in a single move; the fibers will be backlit so that a Fiber View Camera can track their actual positions, and with this information the positioners can be instructed to perform correction moves. A positioner can perform four correction moves after the initial “blind” move within the specified time constraint of 45 seconds per target, so a positioner that can on average achieve the desired accuracy within that number of moves satisfies the requirement. In addition to this specification, there is an upper bound on the maximum error a positioner can have on its blind move. This is because the positioners are packed close enough that their 6 mm patrol radii overlap one another; an algorithm is needed to avoid collisions, and such an algorithm could never be designed without a reasonable upper limit on movement error.

Currently, positioners are being consistently produced that meet these basic requirements. Further improvement of positioner performance would not greatly add to the data’s potential for scientific discovery, as even if a fraction of the positioners were performing sub-optimally, they could be arranged to collect the desired portion of the sky. On the other hand, if the positioners were improved so that they reached the desired accuracy in less submoves than they are currently achieving, it could greatly increase the speed of the survey. The work presented here is a proposal to achieve this. Visual inspection of repetitive patterns from quality control test data has suggested the presence of significant systematic error in a standard positioner’s performance; a least squares fit performed on a positioner’s recent history is a way to harness this error by actively correcting its future performance. After summarizing the test performed on the positioners and deriving the least squares method to be used, this paper presents data collected over a long period and the result of performing such a fit on this data. For both the traditional test and a followup test meant to more closely simulate the effective randomness of the target positions in the actual experiment, the results are very promising. Although there are some further tests

that could be done, the data taken so far suggests that implementing this method could decrease the needed submoves for most positioners. Thus, as the commissioning of the actual instrument begins in the near future, the implementation of this method for tracking positioner history to predict future error is strongly recommended.

CHAPTER II

Fiber Positioner Quality Control Testing

In order to understand the importance of this paper’s proposal, it is essential to understand how a positioner actually works and what qualifies as a satisfactory test of its performance.

2.1 The Positioner

An example of one of the robotic positioners is presented by Figure 2.1. Commands are sent to the positioner through the pigtailed connected at the end of the electronics board. CAN commands are used to control the positioner’s two motors by Pulse Width Modulation. These two motors are referred to as “theta” and “phi”, and each control a 3 mm arm that reaches perpendicularly from the axis of the positioner. The theta arm rotates around this positioner axis, and the phi arm rotates around the end of the theta arm. Thus, with over 360° of rotation allowed in theta and over 180° in phi, the end of the phi arm can theoretically reach any point within a 6 mm radius circle centered on and within the plane perpendicular to the positioner axis. A cylindrical metal idler with a wedge-shaped gap surrounds the central axle of the positioner to define the edges of the theta arm’s movement range. These boundaries are called hard-stops, and are needed both to avoid tangling the attached optical fiber and to define the angular coordinates used to track the position of the theta arm.



Figure 2.1: One of the positioners that has been produced for use on DESI. Both a full depiction for scale and a closer view of both sides of the upper housing are provided.

The phi hard-stops are defined by the shape of the positioner’s upper housing. A central shaft of the positioner is open for the electronics wiring and for a fiber optic cable to be passed through, with its ferrule to be held within the holder at the end of the phi arm by a set screw. Each positioner is visually identifiable by the serial number printed on its upper housing, and electronically by its unique CAN ID.

2.2 The Test Stand

The goal of the test stand is to measure the ability of every produced positioner to move fiber optics with the required accuracy using at most four correction moves. After various calibrations of the positioners and test stand that will be described in more detail in the following section, this is achieved by instructing each positioner to move through an array of evenly spaced target points across the expected movement range and logging the error after each move. Figure 2.2 presents the physical apparatus used to perform this test. Custom made blocks with notches to hold positioners and clamps to keep them secure are attached to the optical table. Across from this is the Fiber View Camera, specifically the SBIG STF-8300, held secure by a custom built mount. The heights of these elements assure that up to forty positioners can be placed in the block with their tips in full view of the camera. Optical fibers are fed through the positioners, with their ferrules securely fastened within each ferrule holder. Their other ends are collected in an LED-lit receptacle, which provides the backlight that renders the position of the fiber tips measurable with the camera. One of the forty allowed spaces in the block contains an electronic fiducial light source that is used to define a global coordinate system immune to vibrations of the table or the camera. Each positioner, the fiducial, and the LED illuminator are all wired to a custom-designed breakout board. Through this, which is additionally connected to a power supply, each element is provided with 7.5V. The breakout board also provides the pathway through which all of the elements can be controlled by CAN commands



Figure 2.2: One of the test stands used for the quality control tests of DESI positioners sent from a BeagleBone Black running with a DESI-specific operating system. This is powered externally and communicates through the university’s network with a PC neighboring the test stand. This PC is also connected by USB to the camera for readout. The camera, block, illuminator, and breakout board are enclosed by panneling in an aluminum frame, which provides the darkness needed for the camera to accurately capture the fiber positions.

The standard test run on all positioners has four fundamental differences in conditions to the actual instrument that either have been or will be separately tested. First, the positioners on the instrument will be held vertically, with their fibers pointing toward the ground at the collection optics, as compared to their horizontal orientation in the test stand. A more compact stand with all of the same elements has

been constructed attached to a rotating frame that allows tests to be performed at a series of angles between horizontal and vertical orientations. Initial results of such tests do not suggest a loss of positioning accuracy. Second, the actual experiment will require positioners that are able to perform to the specified accuracy for 372,000 moves [2], while a standard test only contains 5,000 moves. This has been tested in a series of lifetime tests, including one on twenty positioners that reached 700,000 moves and left most of their performance intact. The data of this maximal lifetime test has been used for this least squares project. Third, the block in the standard test stand holds the positioners far enough apart that there is no chance of collision and therefore no anti-collision algorithm is used. This algorithm has been designed and a tight-array testing block produced to test it; tests of this nature will be performed soon. Fourth, the positioners in the test stand move through an evenly spaced grid of points, moving along each vertical line of points in the same direction. This could hide performance issues only apparent during tests with random targets. To ensure the chance of effectively random collisions, the anti-collision tests will require a scattered collection of targets, so this will also be tested during anti-collision tests. The latter two aspects of the instrument will be further tested with an engineering model “petal” (the positioners on the actual instrument will be arranged into large petal-shaped pieces) before the actual experiment is performed.

2.3 The Testing Software

The testing software is made up of an extensive hierarchy of scripts written in Python. In general, movement of the positioners can be done using functions varying in complexity from direct movement commands, only called directly for visual performance checks, to a full set of potentially distinct positioner-specific targets organized into a table, which is the standard method used by the test script. For a positioner to move to a target defined in the positioner-centric Cartesian coordinate system,

physical aspects of the positioner determined by calibration are used to convert to the needed change in theta and phi. In every case, the moves are all sent to the petal controller (the BeagleBone) in this simple form where they can be converted to CAN instructions for the number of motor rotations. Measurement of positioner location is done by fitting 2D Gaussians to the output of the camera's CCD. Every positioner has an accompanying configuration file that tracks its last measured global position, which is used to match each positioner to the dots seen on the CCD at each new measurement. The distance between this and the last sent target gives an error on the positioner's move. The fiducial and illuminator also have configuration files, allowing the user to set and track the brightness of every element in the test stand to ensure accuracy in finding the dots' centers.

Before a test sequence is run, the software locates each positioner given in an additional hardware configuration file by using its known CAN ID (stored in the positioners' configuration file) to nudge it by a small angle in phi. The location of any dots that do not move when all are instructed to are stored as fiducial dot positions. After this, calibration of each positioner entails moving each of theta and phi in points evenly spread along an arc spanning each arms' movement range. With this data, the software determines the theta and phi arm lengths, the theta and phi offsets (used to define each positioner-centric angular coordinate system), and the theta and phi movement range. All of these attributes are updated in the positioner configuration files to be referenced during any movement command.

Another configuration file is used to define the attributes of the desired test sequence. This customizability allows for the design of alternate performance tests, including those that were described in the previous section and the tests used in this project. The standard test sequence performs two 72-point tests each followed by 500 unmeasured moves to burn in the positioners. Then a sequence of 24-point tests at varying fractions of their maximum supplied current are run to collect data on the

relationship between positioner performance and provided power. Finally, the actual performance test consists of three sequential 192-point tests. This latter sequence is done using four correction moves after each blind move. Every test loop is preceded by an additional calibration sequence.

For a given test sequence, every test loop generates a separate data file for each positioner. This file contains the coordinates of each target, the actual measured position of the positioner trying to reach it during each of its submoves, and the error represented by the difference between these. Various visualizations of this data are typically generated, as will be done for the tests in this study. Patterns in these plots can be used to propose the mechanical defect or test stand malfunction that led to the failure.

CHAPTER III

Least Squares Vector Field Fitting

A least squares fit is a continuous mathematical function meant to approximate some discrete collection of points. This is done by minimizing the sum of squared errors, or residuals, between the fit and the points. The data that is output by a positioner's quality assurance test is a set of vectors, so in this case the points will be fit to a polynomial vector field of arbitrarily chosen degree. The mathematics of least squares fitting is fairly straightforward, but will be developed in detail here. This has been adapted from the method presented by Lage, et al. [3].

A mapping $\mathbf{F} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is called a vector field and can be written

$$\mathbf{F}(x, y) = (u(x, y), v(x, y)) \tag{3.1}$$

for some $u : \mathbb{R}^2 \rightarrow \mathbb{R}$ and $v : \mathbb{R}^2 \rightarrow \mathbb{R}$. A vector field is a bivariate polynomial of degree d when

$$u(x, y) = \sum_{0 \leq i+j \leq d} a_{ij} x^i y^j, \quad v(x, y) = \sum_{0 \leq i+j \leq d} b_{ij} x^i y^j \tag{3.2}$$

for some coefficients $a_{ij}, b_{ij} \in \mathbb{R}$. These can be written more conveniently by defining

$$\begin{aligned}\mathbf{a} &= \begin{bmatrix} a_{00} & a_{10} & \cdots & a_{d0} & a_{01} & \cdots & a_{(d-1)1} & a_{02} & \cdots & a_{0d} \end{bmatrix}^T \\ \mathbf{b} &= \begin{bmatrix} b_{00} & b_{10} & \cdots & b_{d0} & b_{01} & \cdots & b_{(d-1)1} & b_{02} & \cdots & b_{0d} \end{bmatrix}^T \\ \mathbf{P}(x, y) &= \begin{bmatrix} 1 & x & \cdots & x^d & y & \cdots & x^{d-1}y & y^2 & \cdots & y^d \end{bmatrix}^T\end{aligned}$$

so that we can write (3.2) as

$$u(x, y) = \mathbf{P} \bullet \mathbf{a} = \mathbf{P}^T \mathbf{a}, \quad v(x, y) = \mathbf{P} \bullet \mathbf{b} = \mathbf{P}^T \mathbf{b}. \quad (3.3)$$

Our goal is to take some sample of n vectors, composed say of $\mathbf{v}_k = (v_{x_k}, v_{y_k})$ defined at a series of points (x_k, y_k) , and best approximate the set by a continuous function (3.1) of the form (3.3). In other words, we want to minimize the sum of squared errors between \mathbf{F} and our sample by choosing all of the components of vectors \mathbf{a} and \mathbf{b} . The sum of squared errors is

$$\begin{aligned}\sum_{k=0}^n \|\mathbf{F}(x_k, y_k) - \mathbf{v}_k\|^2 &= \sum_{k=0}^n (\mathbf{F}(x_k, y_k) - \mathbf{v}_k)^T (\mathbf{F}(x_k, y_k) - \mathbf{v}_k) \\ &= \sum_{k=0}^n \|\mathbf{F}(x_k, y_k)\|^2 - 2\mathbf{F}(x_k, y_k)^T \mathbf{v}_k + \|\mathbf{v}_k\|^2 \\ &= \sum_{k=0}^n \|\mathbf{P}(x_k, y_k)^T \mathbf{a}\|^2 + \|\mathbf{P}(x_k, y_k)^T \mathbf{b}\|^2 \\ &\quad - 2(\mathbf{P}(x_k, y_k)^T \mathbf{a} v_{x_k} + \mathbf{P}(x_k, y_k)^T \mathbf{b} v_{y_k}) + \|\mathbf{v}_k\|^2 \\ &= \sum_{k=0}^n \mathbf{a}^T \mathbf{P}(x_k, y_k) \mathbf{P}(x_k, y_k)^T \mathbf{a} + \mathbf{b}^T \mathbf{P}(x_k, y_k) \mathbf{P}(x_k, y_k)^T \mathbf{b} \\ &\quad - 2(\mathbf{a}^T \mathbf{P}(x_k, y_k) v_{x_k} + \mathbf{b}^T \mathbf{P}(x_k, y_k) v_{y_k}) + \|\mathbf{v}_k\|^2.\end{aligned}$$

Thus if we define

$$\begin{aligned}
 S &= \sum_{k=0}^n \mathbf{P}(x_k, y_k) \mathbf{P}(x_k, y_k)^T \\
 S_x &= \sum_{k=0}^n \mathbf{P}(x_k, y_k) v_{x_k} \\
 S_y &= \sum_{k=0}^n \mathbf{P}(x_k, y_k) v_{y_k} \\
 S_{xy} &= \sum_{k=0}^n \|\mathbf{v}_k\|^2,
 \end{aligned}$$

we can write the sum of squared errors in the following form:

$$\sum_{k=0}^n \|\mathbf{F}(x_k, y_k) - \mathbf{v}_k\|^2 = \mathbf{a}^T \mathbf{S} \mathbf{a} + \mathbf{b}^T \mathbf{S} \mathbf{b} - 2(\mathbf{a}^T \mathbf{S}_x + \mathbf{b}^T \mathbf{S}_y) + S_{xy}. \quad (3.4)$$

This is the equation we want to minimize with a choice of coefficients. If we consider (3.4) as a function $f(\mathbf{a}, \mathbf{b})$, then only varying \mathbf{a} , the critical point of f occurs when $\nabla f = 0$, with the gradient taken in respect to the components of \mathbf{a} . We have, using basic vector derivative rules,

$$\nabla f = \mathbf{S} \mathbf{a} + \mathbf{S}^T \mathbf{a} - 2\mathbf{S}_x = 2\mathbf{S} \mathbf{a} - 2\mathbf{S}_x$$

where in the second equality the symmetry of \mathbf{S} was used. The gradient with respect to \mathbf{b} follows similarly, and thus the critical point of f occurs when

$$\mathbf{S} \mathbf{a} = \mathbf{S}_x, \quad \mathbf{S} \mathbf{b} = \mathbf{S}_y. \quad (3.5)$$

This provides easily calculable equations for the form of the best fit multivariate polynomial for any given collection of vectors.

CHAPTER IV

Analysis of Lifetime Data

4.1 Overview and Procedure

To use the result derived in the previous chapter to improve positioner performance, a script was written in Python that solves (3.5) to find the coefficients of (3.2) given a sample set of vectors. Thus, given some volume of errors from a positioner’s performance history, a fit can be performed on them to be compared to its future performance. For such a fit to be useful, the coordinate system used to define the errors has to be consistent. The output file of the test stand gives targets, measured positions, and errors in a Cartesian coordinate system with the origin set at the positioner but the axes defined perpendicular to the camera frame. These coordinates would render past data useless after significant rotations of the positioner around its central axis, and at the accuracy required by the instrument, “significant” rotation can occur during normal operation. Thus, before the fitting procedure, the recorded theta offset measured by the calibration preceding each test loop was used to rotate the Cartesian coordinates, placing the x-axis at the positioner’s theta hard-stop.

To demonstrate the efficacy of such a script in improving positioner performance, the data of the lifetime test described in section 2.2 was used. This test included standard 192-point tests separated by various amounts of unmeasured moves (from 2,000 to 10,000). 2 months worth of data taken in this way for fourteen positioners

was analyzed, with the result that a span of approximately 240,000 moves was used for this study. By fitting the data taken over various time periods and subtracting the resulting vector field from the offset normalized output of the next test, the potential improvement from a physical implementation of this correction was revealed.

4.2 Results

Before analyzing the data, any positioner that exhibited a period of tests significantly higher in error than the specifications required of a good positioner was removed from the data set. This meant removing three of the fourteen positioners. Because in all three cases these issues were short-lived, this is most likely an issue caused by the test stand. Of course, fitting based on bad performance data will lead to an unhelpful fit, and ignoring data from unusually bad performance should be implemented if the correction technique proposed here is used in the actual experiment. In addition, a couple other positioners exhibited a single test with nonphysically high error. In these cases, the offending test was removed and the rest of the data analyzed.

To demonstrate both the effectiveness of the fit and an appropriate polynomial degree, Figure 4.1 presents a series of fits to a single 192-point data set. The progression to higher degrees clearly shows an approach toward the actual test result. The difference between the 5 and 10 degree fits is subtle, but nuances of the actual test result are clearly mimicked better by the highest degree fit. Because there is no specifically desired accuracy, this analysis will be kept qualitative, and a 10 degree fit will be used for the study. It should be noted that degrees much higher than this often led to wildly inaccurate fits that were most likely the result of some error in Python's numerical methods.

Before analyzing all of the positioners at once, it is illuminating to view the results of the fit directly for one representative example. Figure 4.2 gives the fit of one full month of data set next to the data of the following test. One month of data is




Figure 4.1: The result of performing the least squares fit to three different polynomial degrees on a typical test output. In each plot, the base of the arrows are set at the target, and the point at the measured position.

equivalent to 30 different measurements at each target, giving a total of 5760 moves contributing to the fit. Although on its own anecdotal, this is a very promising visual suggestion of systematic error; effectively averaging a large data set closely resembles the following test. To check whether this is stable, Figure 4.3 shows effectively this same comparison done over an extended time period; the max and rms error for every test over the course of a month is plotted both for the original data set and the set resulting from subtracting a fit on the previous month. These plots show, at least for one positioner, that the promising visual result present in Figure 4.2 is stable over

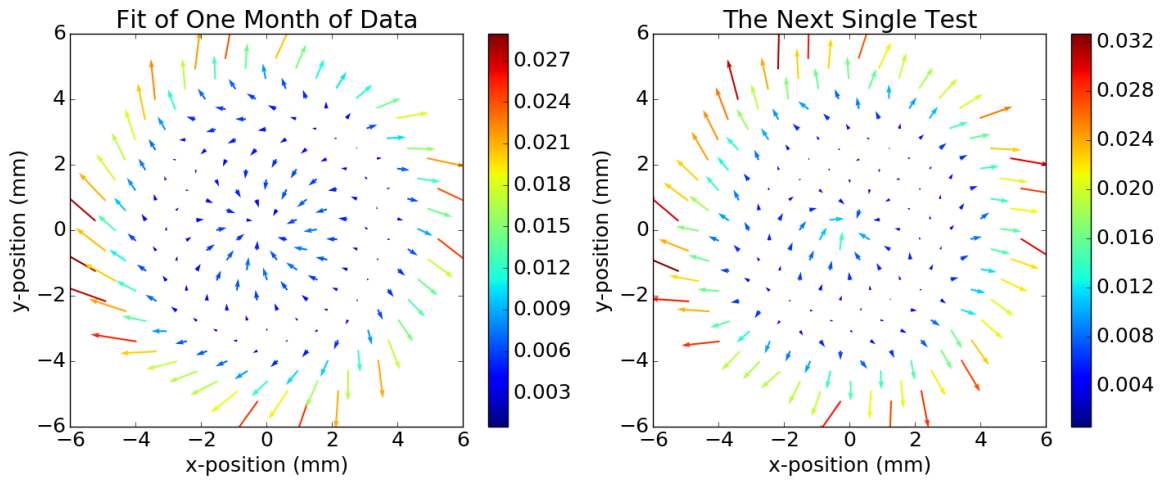


Figure 4.2: The result of a fit on one month of blind move data from positioner M00989, and the regular blind move data from the next performed test loop. Note that every data set before contributing to the fit, and the data of the next test, have been rotated according to the position of the theta offset to make comparison reasonable.




Figure 4.3: The max and rms errors of the blind moves in every test loop taken over the course of a month of the lifetime test, compared to these same values after subtracting a fit on the previous 30 tests.

many moves. Subtracting the fit greatly improves the results, and does so consistently over the entire testing period.

The next natural question to answer with this model positioner's test results is whether or not using a shorter history retains such a promising trend. Figure 4.4 presents the same plot as 4.3, but after shortening the length of time fitted over to about a week and then to only the single most recent test. The former case specifically means 7 different measurements at each target and therefore a fit generated from 1344 points. The latter is of course a fit on only 1 measurement at each target, meaning a fit generated from 192 points. The results are surprising: although it is more prone to errant tests, using only the previous test seems to retain the quality of the hypothetical correction. Although not included here, a similar trend can be seen in other positioners. There is probably no need to make this comparison more precise: accumulation of performance data will happen slowly on the actual experiment, so if a similar magnitude of correction occurs for less data, this should be the focus of further analysis.

Motivated by this idea, Figure 4.4 presents the averaged result of performing the same technique on all eleven positioners. All positioner data sees improvement after subtracting the fit, with most exhibiting significant improvement. Also plotted in Figure 4.5 is the line representing the rms error required for an instrument-ready positioner; seven of the positioners would be within specification after only a blind move, while most remaining may require only a single correction move. This data strongly supports the use of this method; if all of the positioners on DESI performed similarly well in the blind move, the experiment could be performed in significantly less time. Of course, the data up to this point does not necessarily reflect the result of using this method on the actual experiment. The next chapter takes a step closer toward the conditions that will be seen on the instrument.

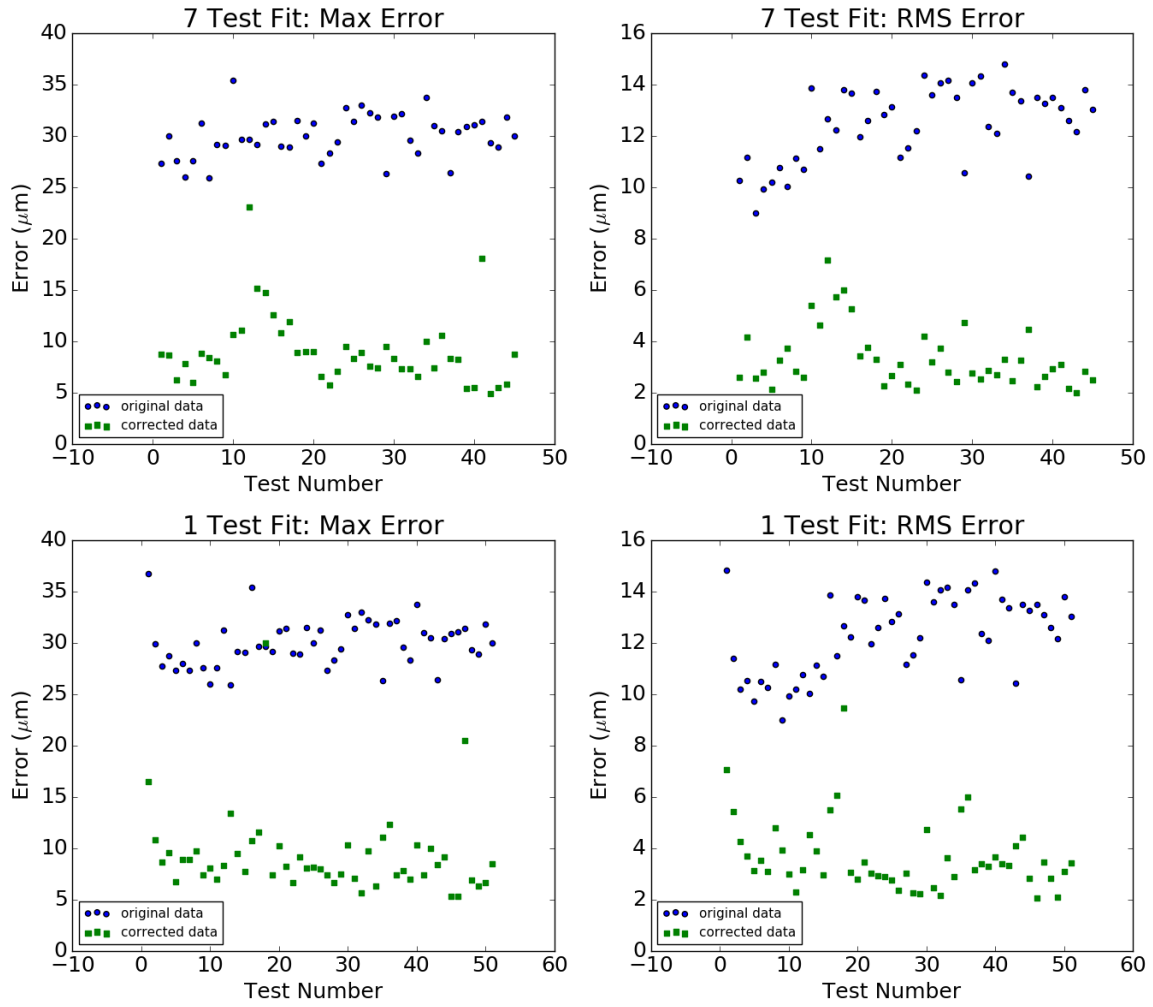


Figure 4.4: The max and rms errors of the blind moves in every test loop taken over the two months of analyzed lifetime data, compared to these same values after subtracting both a fit on the previous 7 tests and a fit on the single previous test.

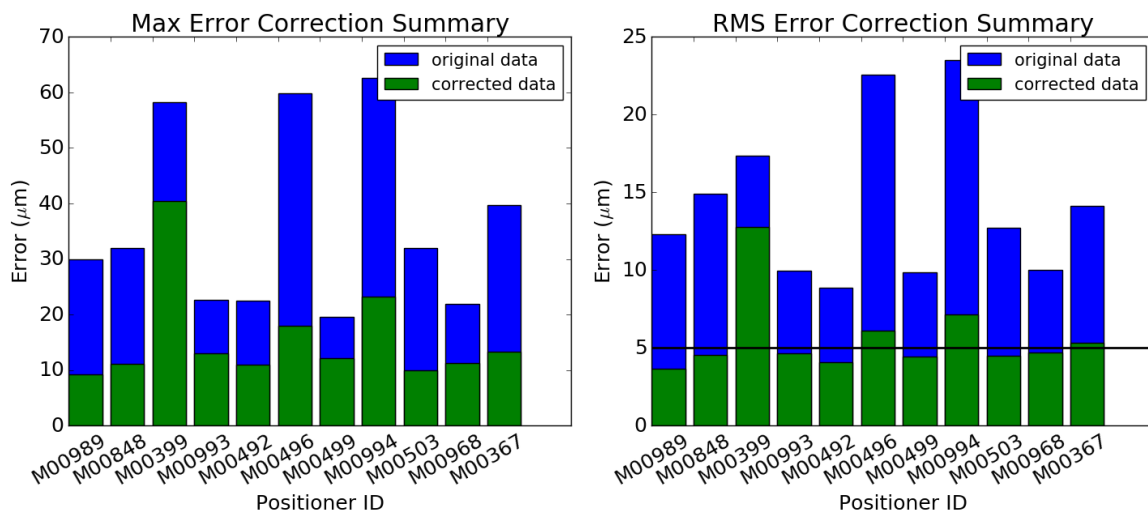


Figure 4.5: The average max and rms error achieved by each well-performing lifetime positioner's blind move over the extent of the analyzed lifetime data, and the average result of subtracting from each test a fit on the previous test. The rms plot also includes a line marking the rms required for the instrument (to be reached within at most four submoves).

CHAPTER V

Analysis of Random Point Testing

5.1 Overview and Procedure

Although the results of the previous chapter are promising, there is an important factor being overlooked by studying just the lifetime data. On the instrument, the points will be randomly distributed around each positioner's range, and so the positioner will approach points from seemingly random directions. It is not understood what the main source of positioner error is, but it is reasonable to think that it could depend on distance traveled before a target, or the direction from which the target is approached.

To test this, six 192-point tests were run in a row on the same positioners used in the lifetime test, aside from a couple that had stopped operating within specification; these were replaced by new positioners. The test was performed on fifteen positioners that all produced viable performance data. This time, the same array of points were used but in a scrambled order. The same scrambled order was used for the first three tests, and then three distinct orders used for the following ones. Figure 5.1 shows the actual pattern used for each of these target arrays, demonstrating that the utilized patterns did effectively randomize the points as desired.

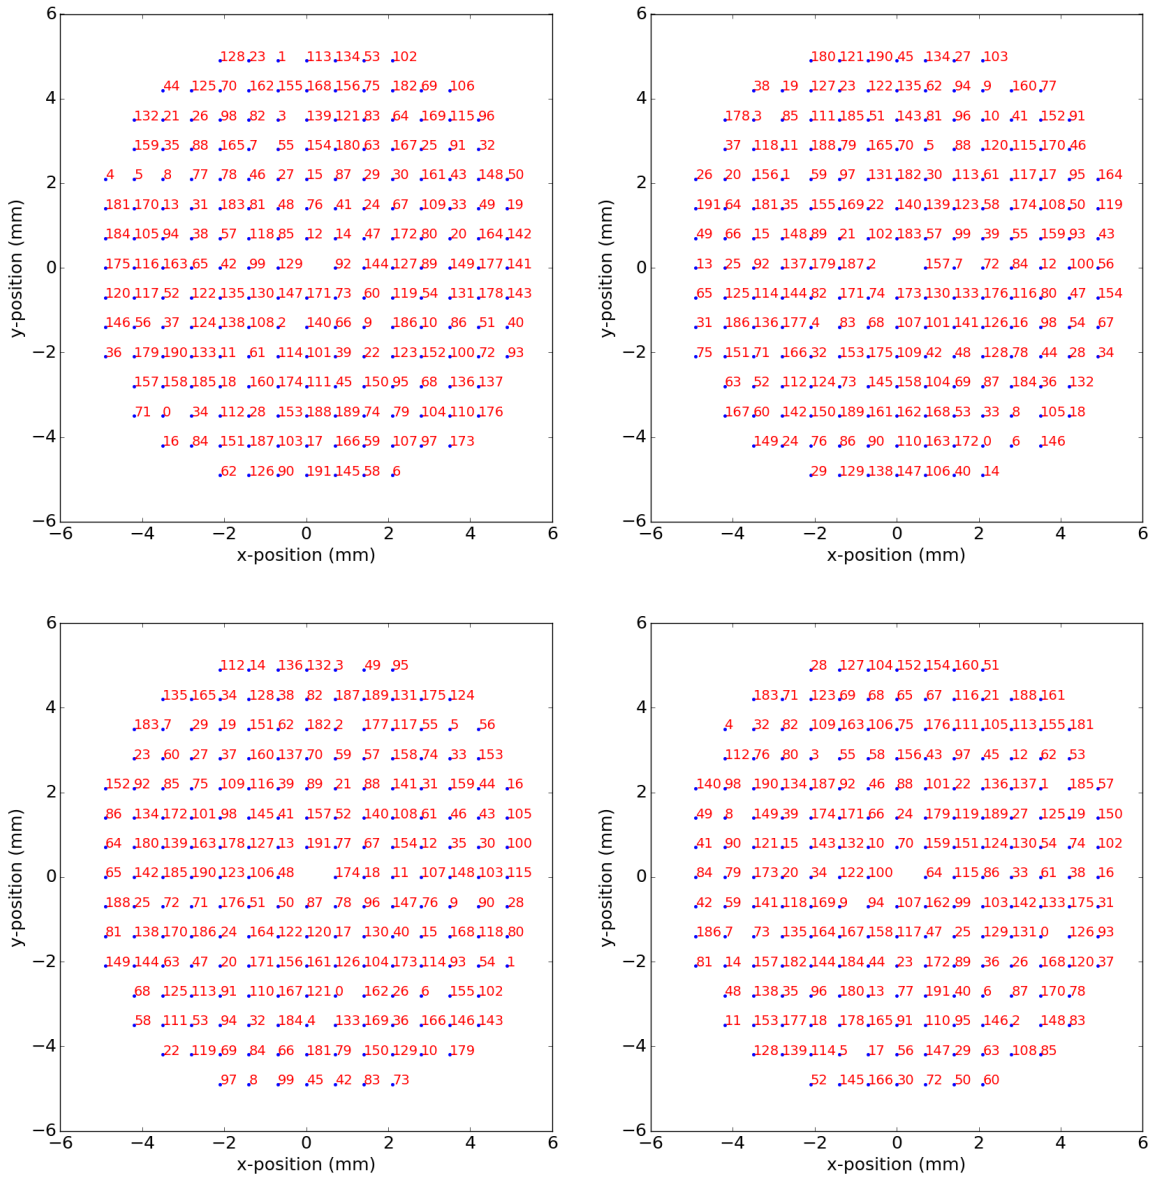


Figure 5.1: The four “random” target sequences used to test the least squares fitting method on unpredictable positioner movements. The shapes of these arrays are the same as the standard 192-point tests; only the order of the targets has been changed.

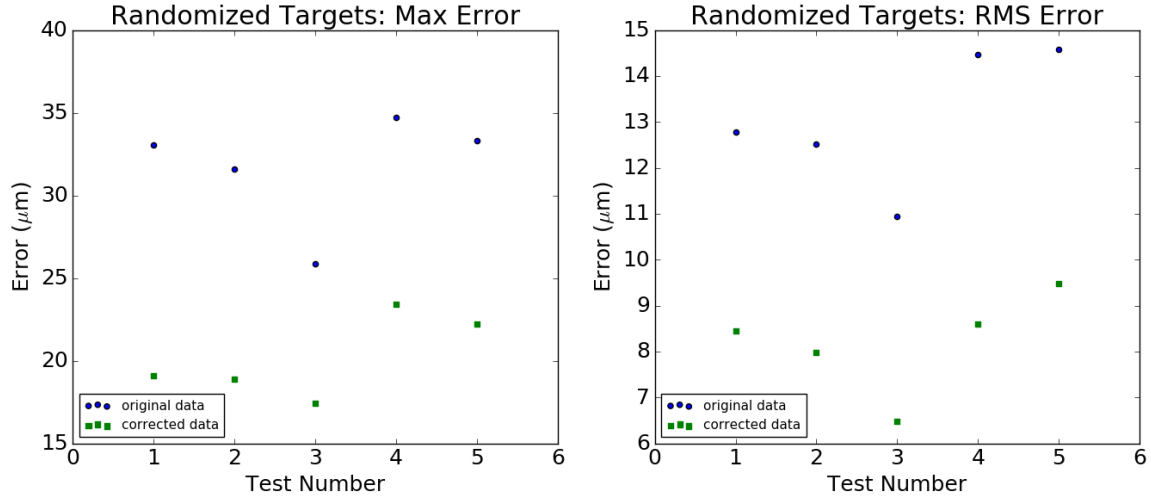


Figure 5.2: The actual max and rms error of positioner M00989’s blind moves resulting from the performance test with scrambled target order, along with the result of subtracting a least squares fit on the previous test. The first three tests were run with the same pattern, so that the first two tests plotted here were corrected by results from the same movement sequence. On the other hand, the other three plotted tests each had distinct target sequences, and so were each corrected by data taken using a different target order than their own.

5.2 Results

Figure 5.2 presents the same model positioner as the previous chapter, again used to demonstrate the potential of the fitting method. The average max and rms error is plotted for each test result along with the result of subtracting the fit of the previous test. Note that the first two corrected tests in each plot of Figure 5.2 have used a fit on data from the same target pattern as the test being corrected, while the following three have been corrected by different patterns. The data is encouraging; for this positioner, scrambling the points has not ruined the efficacy of the fit method.

Figure 5.3 gives a summary of this analysis done for every positioner. The uncorrected tests maintained similar errors. There is still uniform improvement, suggesting there is a large presence of error that depends only on target position. On the other hand, the rms line reveals a quantitative degradation in the utility of this method:

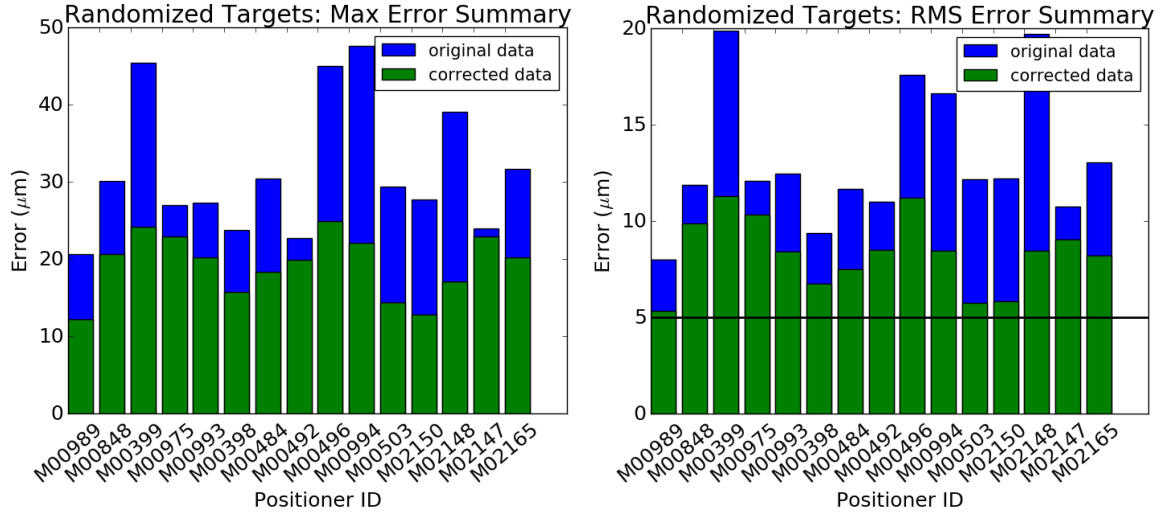


Figure 5.3: The average max and rms error achieved on the blind move by each positioner over the six tests with randomized target orders, and the average result of subtracting from each test a fit on the previous test. The rms plot also includes a line marking the rms required for the instrument (to be reached within at most four submoves).

no positioners would theoretically be performing to the specified accuracy after the blind move in this case, compared to the majority of the lifetime positioners that did demonstrate this. Thus, there is also some amount of error that is dependent on the way a target is approached. Regardless of the loss of the more striking improvement suggested in correcting the lifetime data, the random point data still suggests that this method could decrease the average number of submoves needed to fall below the desired rms error (while significantly lowering the max as well).

CHAPTER VI

Conclusion

Using the hardware and software already in place for regular performance testing, the aim of this project was to demonstrate the potential of performing a least squares fit of a positioner's recent error history and using this to predict its future errors and correct for them. This was simulated by subtracting a fit on performance data from the result of a following test, revealing the approximate effect of such a correction. The result of subtracting the fit of a single 192-point test from the following one for 51 tests in a row was seven of eleven positioners hypothetically reaching rms error specification after only the blind move; in addition, all eleven positioners saw significant improvement in both the blind max and blind rms errors. The result of randomizing the order of target points was to diminish this success rate; no positioners from this test reached specification in only the blind move. Uniform improvement was still seen, though; most positioners would have most likely reached the required rms error in less submoves than they would have without correction. This data strongly supports the implementation of an actual active correction based on this fit method.

There are multiple avenues of further study available. First, it may be worthwhile to repeat the random point test with randomized target positions instead of simply randomized order. The tests presented here do not explore the full capability of this method in that errors at the same target points that were used to produce

the fit are predicted by it. It seems likely that a fit of the same density of points covering a comparable portion of the positioner’s movement range would have similar results to the randomized target order tests; as can be seen in the actual test data of Figure 2.1, the tests of satisfactory positioners very often produce what appear to be discretized versions of continuous vector fields. With such recognizable patterns, large “discontinuities” between the tested points seem unlikely. Nevertheless, unexpected behavior could still arise. Second, as mentioned in Section 2.2, while initial tests of positioners held at varying angles from the horizontal have been promising in that the orientations seem to maintain average positioner accuracy, there has not yet been any check that the actual error patterns are maintained. This would be a instructive test before implementing this correction method on the actual instrument. Finally, and most importantly, this method needs to actually be performed in the test stand. The results presented here all represent an estimation of what an actual correction would do; only by actually performing this correction, and physically reaching specification with less submoves, will the viability of this method truly be proven. The data found by this project should make a convincing case for this implementation. In fact, the random target and varying orientation tests could be performed with this implemented correction. An additional benefit of this test could be the use of new positioners; by the time the random point test was run, many of its positioners had seen lifetimes of moves.

The combination of the encouraging data presented by this study and the potential gains for DESI should hopefully provide a convincing argument for the reader: Further work on the design and improvement of this method is worth the collaboration’s time.

BIBLIOGRAPHY

- [1] Amir Aghamousa et al. The DESI Experiment Part I: Science, Targeting, and Survey Design. December 2016.
- [2] Amir Aghamousa et al. The DESI Experiment Part II: Instrument Design. December 2016.
- [3] Marcos Lage et al. Vector field reconstruction from sparse samples with applications. pages 297–306, October 2006.