Summer 7-29-2022

Run-to-Run Control via Constrained Optimization of a Mechanical Serial-Sectioning System

Damian L. Gallegos-Patterson

"Run-to-Run Control via Constrained Optimization of a Mechanical Serial-Sectioning System" a thesis prepared by Damian L. Gallegos-Patterson in partial fulfillment of the requirements for the degree, Masters of Science, has been approved and accepted by the following:

Dr. Claus Danielson

July 28, 2022

Date

Committee in charge:

Dr. Claus Danielson

Dr. Jonathan Madison

Dr. Meeko Oishi

Run-to-Run Control via Constrained Optimization of a Mechanical Serial-Sectioning System

BY

Damian L. Gallegos-Patterson,

B.S. New Mexico Institute of Mining and Technology

A thesis submitted to the Department of Mechanical Engineering

in partial fulfillment of the requirements

for the degree

Masters of Science

Mechanical Engineering

University of New Mexico

Albuquerque, New Mexico

December 2022

DEDICATION

I dedicate this work to my wife, mother, father, brother, and all of the amazing friends I have met along my path.

ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Claus Danielson, for his patience, interest, and knowledge to help me complete this research. I would like to thank him for investing his time in helping me become a successful graduate student. I would also like to thank Dr. Jonathan Madison and Dr. Andrew Polonsky for their support, encouragement, and belief in my abilities. I can not emphasize enough how much they have contributed to helping me become successful on my academic journey. Lastly, I would like to thank the University of New Mexico Department of Mechanical Engineering for providing me a strong foundation on my path to becoming a successful engineer.

Run-to-Run Control via Constrained Optimization of a Mechanical Serial-Sectioning System

by

Damian L. Gallegos-Patterson

B.S., Mechanical Engineering, New Mexico Institute of Mining and Technology, 2019

M.S., Mechanical Engineering, The University of New Mexico, 2022

ABSTRACT

This thesis develops a methodology for run-to-run (R2R) control of a mechanical serial sectioning (MSS) system for microstructural investigations. MSS is a destructive material characterization process which repeatedly removes a thin layer of material and images the exposed surface. The images are then used to gain insight into the internal structure and arrangement of a material and are often used to generate a 3-dimensional (3D) reconstruction of the sample. Currently, an experienced human operator selects the parameters for MSS to achieve the desired per slice removal rate. The proposed R2R control methodology automates this process while improving the precision and repeatability of material removal. The proposed methodology does this by solving an optimization problem designed to

minimize the variance of the material removal subject to achieving the expected target removal rate. This optimization problem was embedded in an R2R framework to provide iterative feedback for disturbance rejection and convergence to the target removal amount. Since an analytic model of the MSS system is unavailable, a data-driven approach to synthesize our R2R controller from historical data was used.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 Introduction

1.1 Materials Characterization

Essentially everything around us is composed of materials. From the soles of shoes to the components that make up the James Webb telescope, a selection process for what materials are to be used has been undertaken. In most industries, components are expected to become smaller, lighter, more durable, and more responsive. To achieve this, a significant amount of research and development is invested in understanding the linkage between how a material is arranged at a micron scale and how that arrangement impacts the material's properties and performance. This type of study is known as "materials characterization". Materials characterization is a broad field and is used to investigate the microstructural arrangement of materials, provide insight into properties such as hardness, tensile or compressive strength, elasticity, toughness and others. These properties ultimately translate to material performance. Characterization tools currently available are plentiful and encompass both destructive and non-destructive techniques [1, 2]. Among non-destructive tools, Scanning Electron Microscopy (SEM) [3], X-Ray Diffraction (XRD) [4], and X-ray Computed Tomography [5] are often employed. However, non-destructive techniques may not provide the insight

needed for all cases, at which point destructive analysis may be employed. In destructive analysis, the sample is consumed during the process of data collection. Among destructive and non-destructive techniques, some provide an ability to generate a three-dimensional volume of material. Figure 1 indicates the approximate length-scales for experimental observation across a variety of these destructive and non-destructive techniques. The method selection depends on the resolution needed as well as the material properties of the sample itself. This thesis will demonstrate a methodology for closed-loop control in MSS which, should be stated, is a destructive technique.

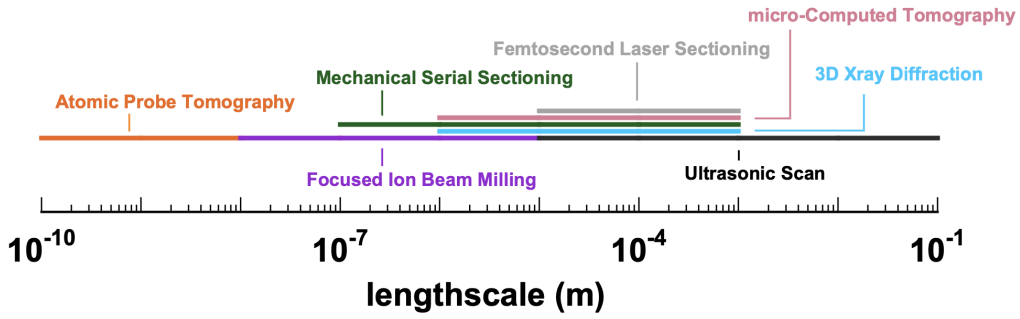

Figure 1: Length scales for three dimensional data collection tools [6].

1.2 Closed-Loop Control

Closed-loop control, also known as feedback control, leverages the output from a system to produce feedback which is then used to determine or adjust further inputs to the system thereby achieving a desired output. Feedback control has

been documented as far back as ancient Greece with items like the water clock of Ktesibios and self regulating oil lamps [7]. Use of feedback control saw a surge during the industrial revolution [8] with applications in machinery and continues to be either applied or investigated across a significant amount of industries today [9–11]. Closed-loop control methods can be applied almost anywhere that a system can be modeled and controlled. Various feedback control methods have been developed and deployed for systems from which information can be collected or produced by the system that can then be exploited to manipulate further inputs and produce a desired output. Some of the advantages of a closed-loop system as compared to an open-loop one is the ability to self-correct over time and reduce the influence of external disturbances on the system. Advances in the area of data-driven, adaptive closed-loop control have improved system performance while decreasing the amount of user intervention required.

Iterative learning control (ILC) is a data driven adaptive control method commonly used for reference tracking of a system in which the system is measured and compared to a reference. This method requires knowledge only of the inputs and outputs of the system and can be done without an explicit model of the system [12]. ILC uses knowledge of the previous inputs and outputs to determine subsequent inputs to the system. ILC is applied to systems which display repetitive and repeatable characteristics. ILC is mainly used for batch processes such as the control of robotic arms [13] which perform the same movements repeatedly. ILC

also allows for in-situ changes to be implemented.

Repetitive control (RC) is a disturbance rejection method primarily used in continuous systems [14]. For most systems using an RC controller, the period of disturbance is known. Similar to ICL, RC has a reference trajectory it uses for tracking, but the reference is checked periodically instead of at the conclusion of a run or cycle. A common application of RC is tracking of a periodic reference signal and a disturbance rejection signal [15]. RC is implemented in continuous processes where in-situ measurements of input signals can be made and disturbances can be rejected.

The methods listed here along with Run-to-Run control are all methods employed to control repetitive systems which use information about the system at prior iterations to make decisions about future iterations. In section 2 we further investigate Run-to-Run control as applied to the challenge at hand. These control methods are classified as learning-type controls.

2 Literature Review

In this section, an overview of Mechanical Serial-Sectioning (MSS) and Run-to-Run (R2R) control will be covered.

2.1 Mechanical Serial Sectioning

Mechanical serial-sectioning in metallugry was first performed by Otto Forsman in 1918 as a process for investigating the microstructure of metals [16]. It was a labor intensive process involving meticulous grinding and polishing of samples using a manual grinding wheel to remove a near uniform amount of material over and over again. This process was revisited by Hull et al. in 1991[17], and by 1998 the process became automated with the introduction of the Focused Ion Beam (FIB) "slice and view" process [18]. In 2001 Voorhees and Alkemper further improved upon the serial sectioning process using a diamond miller and applied their technique to softer metals [19]. To date, a collection of researchers have employed this technique in metallographic settings and reported their findings in the open literature [20–22]. Each of these serial-sectioning methods can be combined with various imaging techniques such as optical microscopy [23] or Scanning Electron Microscopy [24]. In this work, we will focus our attention specifically on MSS with optical imaging.

Mechanical serial-sectioning uses a three phase, repetitive process of grinding, polishing, and optical imagining to collect data from a sample. Here we adopt a common convention of the literature and denote each cycle of polishing and optical imaging as a *slice* [21, 25–27]. In this work, these *slices* are composed of a montage of *tiles* which can be stitched together to produce a large-field cross-sectional view

of the sample or specific region of interest at a given height within the sample. For any set of conditions, the operator inputs the number of slices the system will perform. For our purposes, each set of slices will be referred to as a *run*. For each run, the operator selects a sequence of grinding pads, polishing pads, wheel speed in rotations per minute (RPM), polishing times, polishing solutions, and solution dispensing time which serve as inputs to the system. For the purposes of this thesis, the collection of all these items shall be called a *recipe* which is the sum total of all arguments and inputs provided to the system to yield a desired removal amount per slice over the user determined number of slices. Once the sample has been completely sectioned, or the run has been terminated due to some other criteria, all runs performed are considered in aggregate as the *experiment*. The parameters of the recipe are the inputs to the MSS system.

This thesis develops an autonomous controller for iteratively selecting the appropriate recipe to achieve a target material removal amount. The baseline method for achieving the target removal per slice is for an experienced human operator to run a series of test slices. The material removal is then measured using the average focal height of the revealed sections used for the imaging montage. Based on the calculated material removal and their experience, the human operator adjusts the recipe, typically by adjusting the polishing times, polishing pads or number of polishing steps for each pad. This process is repeated with another run of test slices until the target removal amount is achieved. This process

is not ideal for several reasons. First, it requires significant human intervention from a skilled operator whose valuable experience could be better used for other pursuits. Second, the "calibration-phase" of the recipe can require multiple test runs, removing a large amount of material. This is inappropriate for small samples as a significant portion of the sample will not be sectioned at a consistent rate and calibration steps to optimize material removal rate have to be minimized to preserve the sample. Third, even with the achievement of an optimal recipe, pad wear over the course of a long run can cause the slice thickness to drop causing inconsistent slice thickness during reconstruction. An automated MSS controller could both reduce human intervention and improve the performance of the system by mitigating and reducing these three aforementioned scenarios.

2.2 Run-to-Run Control

Run-to-Run control (R2R) is a model-based adaptive closed-loop control method for discrete systems in which parameters are changed from one run or cycle to the next. This method allows for compensation of variability in the system over time. R2R is commonly used in semiconductor manufacturing where in-situ measurements can not be performed [28]. W.J. Campbell an expert in semiconductor manufacturing and control stated, "Typically, in semiconductor manufacturing, the goal is to control qualities such as film thickness or electrical properties which are difficult, if not impossible to measure in real-time in the process environment.

Most semiconductor products must be moved from the processing chamber to a metrology tool before an accurate measurement of the control variable value can be taken." [29]. Because of the lack on in-situ measurements, adjustments to parameters are made at the conclusion of a run. R2R is able to compensate for drift over time by employing a commonly used technique called exponentially weighed moving average which tunes the system based upon previous runs [30]. Due to the lack of in-situ measurements and an inability for in-run changes to be made to the recipe, we employ an R2R algorithm within our closed-loop control methodology.

3 Design and Methodology

In this section the current limitations of the open-loop approach for a Robo-Met.3D™ mechanical serial-sectioning machine will be presented. We then introduce a deterministic model of the MSS plant, operational constraints of the system, and control objectives. Using historical data we are able to create a model of system dynamics from which we can form a stochastic model of the system. Finally, we create a framework for finding optimal inputs to the system for a given target removal, which can be inserted into a R2R algorithm framework.

3.1 Limitations of Open-Loop Control

Many of the experiments performed on the Robo-Met.3D™ system require a specific resolution of data to adequately resolve features of interest within a given

Figure 2: Open loop nature of the Robo-Met.3D™ system, where a set of fixed user inputs, u_0, form a polishing routine (recipe) that generates an output removal amount y_i which may vary due to system dynamics.

sample. In-plane dataset resolution is ultimately limited by the resolution of the microscope, while the out-of-plane resolution is determined by the slice thickness, thereby explicitly linking at least one dimension of dataset resolution with the removal rate achieved for a given recipe. While optical microscopy can routinely achieve spatial resolutions on the order of 1 to 3 μm, typical slice thicknesses are generally on the order of several microns, and are therefore the limiting factor in dataset resolution. Consistent slice thicknesses are critical for ensuring accurate measurements of features of interest that span multiple slices. Consistent slice thickness also reduces the amount of post processing needed during dataset reconstruction, which involves combining each individual slice into a coherent 3D volume. The open loop nature of the Robo-Met.3D™ system requires a priori

knowledge of suitable polishing parameters for a given sample, making examination of novel or unexplored material systems a challenge, as the operator must rely on institutional knowledge or otherwise be forced to make a "best guess" estimate of starting parameters and iterate toward an acceptable recipe as discussed previously in the introduction.

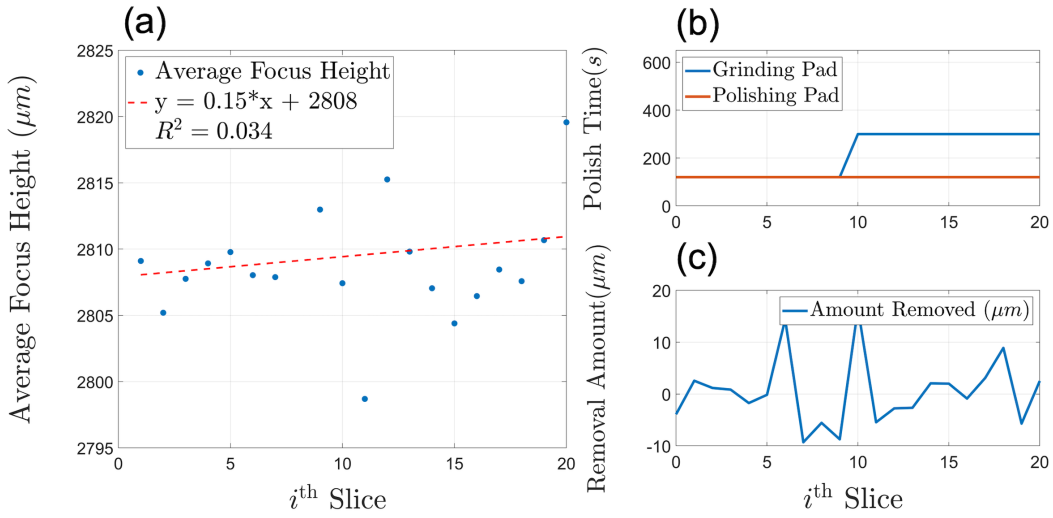

Figure 3: First 20 slices of a historical data set from the Robo-Met.3D™ system. (a) Removal rate as a function of slice number for the two 10-slice runs. Despite a consistent set of system inputs for the first 10-slice run, the measured removal rate is highly variable with a very poor linear correlation ($R^2 = 0.034$). (b) Polish times for both pads as defined by user input. (c) Slice-to-slice removal rate as measured by average focus height.

Figure 3(a) shows the average focus height of the microscope as a function of slice number for the first 20 slices of an experiment performed on the Robo-Met.3D™ system during an initial recipe development process. These initial slices

are comprised of two sets of input parameters, with the grinding step time increasing for slices 10-20. The output of the system is noisy from slice to slice, with a highly variable removal rate. A best-fit line of the focus height data has a very low correlation, and the square of the correlation coefficient (R^2) is close to zero, indicating that the data is poorly fit with a linear model. This inconsistent removal rate can be caused by error in the auto-focus routine of the microscope, or inconsistencies in the slice thickness due to changing material hardness, changing cross-sectional area, pad wear, and other factors[23]. Figure 3(b) shows the polish times for the grinding step and the polishing step in this two step recipe, as a function of slice number. The discontinuity in polishing time for the grinding pad occurring on the tenth slice shows that the operator increased the polish time in an attempt to increase the removal rate, which does not appear to have a significant effect on removal rate as shown in Figure 3(a) despite more than doubling the time on this pad.

The measured removal amounts for each slice in Figure 3(c) demonstrate that the measurement of the removal amount can be highly variable from slice-to-slice, with more than a quarter of slices reporting a negative removal rate. This non-physical result is caused in part by the inherent limitations of the microscope auto-focus, which can be particularly challenging for thin slices only several microns thick, as the depth of field of many optical lenses (the range of focus heights over which an image appears sharp and well-focused) can be greater than the slice

thickness. The total depth of field can be calculated using the following equation:

$$d_{tot} = \frac{\lambda \cdot n}{NA^2} + \frac{n}{M \cdot NA} e \tag{1}$$

where d_{tot} is the total depth of field, λ is the wavelength of the light source, n is the refractive index of the medium, NA is the numerical aperture of the lens, a dimensionless number that quantifies the angular range over which the lens can accept or emit light, M is the magnifying power, and e is the pixel resolution of the detector. The Zeiss AxioObserver microscope used in the Robo-Met.3D™ system has a sensor pixel resolution of 3.4 µm, and an LED light source with a color temperature of 5700K. The wavelength of the light source can then be approximated using the color temperature, T by taking the maximum wavelength of the light source's spectrum according to Wien's displacement law, given by:

$$\lambda_{max} = \frac{b}{T} \tag{2}$$

where b is Wien's displacement constant, and is equal to $2.898 \times 10^{-3} \mathrm{m} \cdot \mathrm{K}$, yielding a wavelength of 508 nm which can be used in Equation 1. Given a refractive index of 1.000 in air for the microscope, the depth of field for various objective lenses used in the Robo-Met.3D™ system can now be determined using the numerical aperture for each lens, yielding depths of field of 35.3 µm for the 5× objective ($NA = 0.13$), 6.78 µm for the 10× objective ($NA = 0.3$), and 1.81 µm,

for the 20× objective ($NA = 0.6$). These large ranges of foci depth, particularly for the 5× and 10× lenses, mean that for sectioning thicknesses on the order of several micron, a successful auto-focusing to produce a sharp image could vary at a fidelity greater than the removal rate. Negative measures of removal rate for an individual slice are therefore not necessarily indicative of a poorly focused image, but more a reflection of the large depths of acceptable foci relative to the slice thickness. This inaccuracy is most clearly observed over a small number of slices, as shown in Figure 3a, and represents the short-range disorder arising from the auto-focus measurement routine due to large depths of field. The standard deviation of slice thickness for these data is 6.8 μm, well within the optical depth of field for the 5× objective used to make these measurements. Despite the inherent limitations of measuring slice thickness via optical methods, a more consistent average slice thickness is observed over longer runs and larger numbers of total slices sectioned (10s to 100s of slices) as the variability from the auto-focus routine is expected to be random. Therefore, although the auto-focus routine may over- or under-predict the actual slice thickness for an individual slice, over many slices this effect averages out to yield a consistent global removal rate. An example of this longer-range consistency is shown in Figure 4, in which the average removal rate is stable on the length scale of nearly 150 slices, as evidenced by an R^2 value (square of the correlation coefficient) of 0.92, which had an average removal rate of 0.63 μm per slice. Around slice 400 of this experiment, the operator increased

the polish time of the grinding pad by roughly a factor of two over the course of 50 slices, resulting in another consistent, measurable regime of material removal rate ($R^2 = 0.95$), with an average material removal rate roughly double the previous rate. Although the linear approximation of average removal rate is quite good over 100 or more slices, these data still display large local variance as the standard deviation for both these regimes is equal to 8.1 μm, which is on a similar order of variance as observed in the data in Figure 3(a). These focus height measurements were also collected using a 5× objective, so the variance is still on the order of one-quarter the calculated depth of field of the lens. These data demonstrate that consistent material removal is achieved globally over many slices, but there are local deviations over smaller runs that necessitate increased operator intervention and additional post-processing routines to uniformly sample the collected data.

Table 1: Microscope lens magnification and depths of focus

Lens	Depth of Focus
5x	35.3 μm
10x	6.78 μm
20x	1.81 μm

In order to address the increased labor associated with "guess-and-check" data collection and subsequent post-processing issues, the Robo-Met.3D™ can be con-

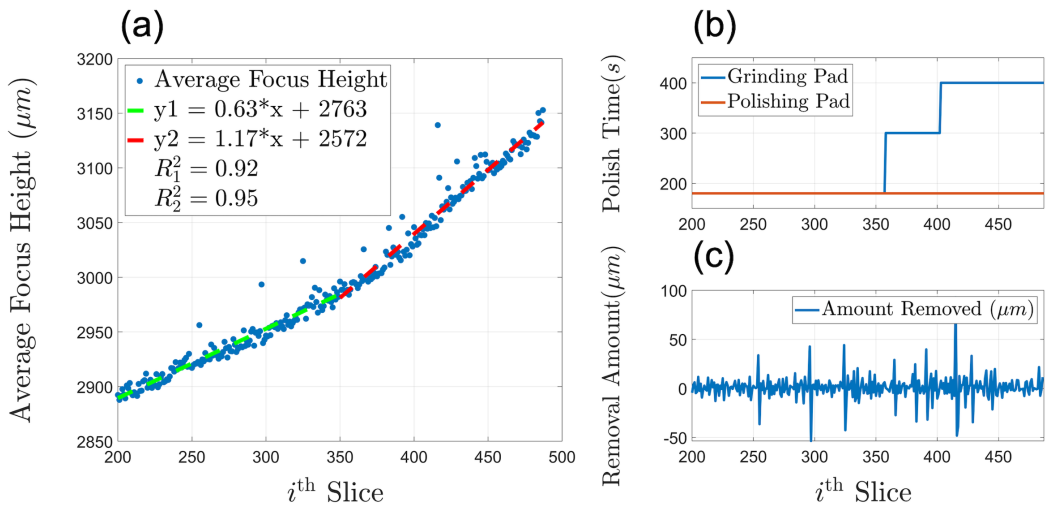

Figure 4: Last 300 slices of a historical data set from the Robo-Met.3D™ system. (a) Average focus height as a function of slice number. (b) Polish time for both Pads as a function of slice number. (c) Slice-to-slice removal rate as measured by average focus height. Manual changes by the operator result in increases in average removal rate, which over many slices, fit well to a linear trendline.

verted into a closed-loop feedback controlled system for more automated operation. This can be achieved by defining a constrained optimization problem within an R2R algorithm framework with the goal of reducing operator intervention during optimal recipe identification. This allows for slice-to-slice recipe changes, and convergence of the removal amount towards an operator specified target amount. R2R has been identified as a useful control algorithm for systems where in-situ changes are not possible and must be made at the conclusion of a cycle[31–34]. In order to implement this solution, we must first model the system dynamics of the Robo-Met.3D™.

3.2 The MSS Plant

Due to the operational nature of MSS systems, we do not have direct access to in-situ measurements, nor can we change the recipe during a run. Thus, we model the MSS system plant as a discrete time algebraic map of the recipe u_i to removal amount y_i

$$y_i = f(u_i, d_i) \tag{3}$$

where $u_i \in \mathbb{R}^{n_u}$ and $y_i \in \mathbb{R}^1$ are the MSS i^{th} system inputs and outputs respectively, and the 'time'-index i represents the *slice number* for the plant. We consider the plant (3) to be deterministic, but unknown. We note that the plant (3) 'dynamics' do not depend on the system 'state' y_i, Thus, the plant is a static non-linearity mapping $f(u_i, d_i) : \mathbb{R}^{n_u} \times \mathbb{R}^{n_d} \to \mathbb{R}^1$ that maps recipes u_i and additional hidden variables d_i to the removal amount y_i. The hidden variables $d_i \in \mathbb{R}^{n_d}$ characterize all unknown factors within the system, such as material hardness, thermal effects, sensor measurement error, grinding and polishing pad wear, cavities in the material, etc. Since we do not have an ability to reliably and consistently provide a universal value for all hidden variables d_i, it is not possible to produce a model of the static nonlinearity (3) from historical data. Furthermore, even if the plant model were known, it cannot be utilized since the hidden variables d_i are unmeasured and time-varying. Instead, we use real-time feedback to adjust

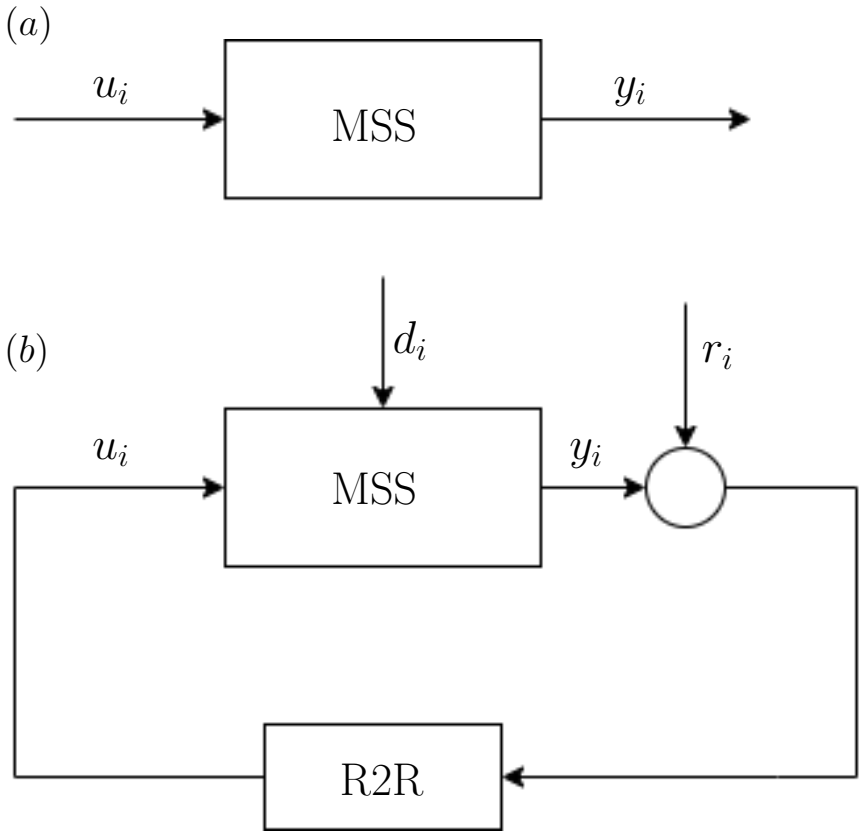

Figure 5: (a) Depicts the MSS System plant with user full user intervention input ui_i. (b) Depicts the Feedback controllable system block diagram, where the plant is based solely off of plant observations of the MSS and the feed-back R2R controller is defined in section 3.6.

the recipe u_i online to achieve the desired removal amount $y_i \to r$ where $r \in \mathbb{R}^1$ is the target removal amount. Feedback control is ideally suited to the problem of rejecting unknown and varying disturbances d_i. Although the plant (3) is static, the closed-loop system will be dynamic due to the dynamics of the controller.

The MSS plant (3) is an over-actuated system; there are multiple $n_u > n_y = 1$ inputs $u_i \in \mathbb{R}^{n_u}$ that can be manipulated to drive $y_i \to r$ one output y_i to the desired removal amount $r \in \mathbb{R}^1$. The manipulated inputs are the recipe

parameters summarized in Table 2. For the automated system, we envision that the MSS will operate with a fixed sequence of multiple grinding and polishing pad types, as well as different solution types for each pad. The automated system will then select the polishing speed and polishing time for each of the pads in the sequence. The vector $u_i \in \mathbb{R}^{n_n}$ of polishing-speeds and polishing-times is the control input (recipe) for the MSS system. For this preliminary work, we will only vary the polishing times for one grinding pad followed by one polishing pad. Thus, we restrict our problem to a two-input system. Importantly for controller development, the removal amount (3) is monotonically non-decreasing with respect to the polish times i.e. polishing for a longer time will not result in less material removed.

The over-actuation of the MSS renders human-in-the-loop operation difficult, requiring the operator to have significant experience and expertise to choose the appropriate recipe $u \in \mathbb{R}^{n_u}$ to achieve the desired removal amount $y_i \in \mathbb{R}^1$. Therefore an automated system also provides the benefit of improving the user-friendliness of the MSS. Furthermore, an automated system could harness this over-actuation to improve system performance, shorten the calibration-phase, and reduce human error. Performance objectives will be described in Section 3.4.

Table 2: Inputs (Recipe Parameters) for the MSS

Input Variable	Constraints
u_{rt} cutting pad type	$1 \leq \gamma \leq 8$
u_{pt} polishing pad type	$1 \leq \gamma \leq 8$
u_s polishing solution type	$1 \leq \gamma \leq 4$
u_ω polishing speed (RPM) for each pad	$1 \leq \gamma \leq 300$
u polishing time for each pad	$5 \leq \gamma$

3.3 Operational Constraints

Our controller must produce recipes u that are physically implementable by the MSS (i.e. non-negative, real numbers). The operational constraints shown in Table 2 describe the physical limitations of the MSS system. The maximum number of pads the system is able to hold at a time is eight, therefore $u_{rt} + u_{pt} \leq 8$ for any recipe. For polishing solutions u_s we are able to select only one of the following particle sizes $1\mu m$, $3\mu m$, $6\mu m$, or $9\mu m$ which can be used in any combination with the pad types. The speed at which the pads can be rotated u_ω is limited to 300 RPM. The minimum amount of time a pad can polish is five seconds. To ensure a favorable imaging surface while avoiding over use of polishing pads, polishing pad use is constrained to no less than sixty seconds and no more than 200 seconds per

step. The constraints for our system form a n_u-dimensional polytope

$$\mathcal{U} = \{u \in \mathbb{R}^{n_u} : Hu \leq h\}. \tag{4}$$

We will analyze our R2R controller with and without these constraints (4).

3.4 Control Objectives

Beyond automating the operation of the MSS, our control objectives include improving its performance. Our control objectives can be summarized by the following conceptual stochastic optimization problem

$$u_i = \arg\min_{u} \quad \mathbb{V}[y - r|u] \tag{5a}$$

$$\text{s.t.} \quad \mathbb{E}[y|u] = r \tag{5b}$$

$$u \in \mathcal{U}. \tag{5c}$$

The desired controller should compute a recipe u such that the expected material removal $\mathbb{E}[y|u]$ matches the target removal amount r i.e. the recipe u should satisfy the equality constraint (5b). Since the MSS is over-actuated $n_u > 1$, there are potentially an infinite-number of recipes $u \in \mathbb{R}^{n_u}$ that can achieve the desired removal (5b). Among these recipes u, we would like to select the recipe that produces the lowest variance (5a) so that the slices have uniform thickness. Finally, the recipe u must be implementable (5c) given the input constraints described in Section 3.3.

In Section 3.6, we translate the conceptual stochastic optimization problem (5) into an implementable deterministic optimization. We will use a data-driven approach to formulate this deterministic optimization problem from historical data. This deterministic optimization problem will be embedded in a R2R control framework to iteratively ensure that the removal amount converges $y_i \to r$ to the target removal amount r.

3.5 Historic Operational-Data

We will use a data-driven approach to translate the conceptual stochastic optimization problem (5) into an implementable deterministic form [35]. We use historic operational data to estimate the mean and variance of the stochastic linear model (6) to approximate a deterministic model for (3). Historical data was collected over a ten year span and include over 150 runs containing up to 500 slices per run. The data files contain the inputs of the system (polish times, RPM, pads used) and outputs (microscope focus height). We utilize the observations from the system plant which come in the form of multiple data files, which contain the inputs (polish times, RPM, pads used) and outputs (microscope focus height). We developed a script that extracts this data from thousands of separate text-files and collects the data into a unified data set $\mathcal{D} = (\mathcal{Y}, \mathcal{U})$ where $\mathcal{U} \in \mathbb{R}^{n \times 2}$ are the recipes and $\mathcal{Y} \in \mathbb{R}^{n \times 1}$ are the resulting amounts of material removed. When a slice is imaged, the focal height of the microscope is recorded for each image in

the montage. The average focal height is then calculated and used to estimate the amount of material removed for each slice. Microscope auto focus errors can occur which will cause incorrect average focal height readings which can lead to readings of negative or minimal removal amounts. Therefore, the need to pre-process the data by removing outliers is required. This includes removing all data points associated with negative values. Also, we compute a preliminary estimate of the mean and variance of the parameters c and b. Any data-points $y_i = c + b^\top u_i$ outside of three standard deviations of the estimated mean value $y_i = \mu_c + \mu_b^\top u_i$ are removed. Once pre-processing is complete, we use the historical data to estimate the mean and variance of the parameters of the stochastic model (6) using the generalized method of moments (GMM) method [36]. The curve-fit of the mean $y_i = \mu_c + \mu_b^\top u_i$ of (6) is shown in Fig. 6.

Fig. 6(a) shows that the available historical data is not very exciting (in the sense of persistency of excitation). The expert human operators tend to use a few different recipes and the polish times are round numbers, typically multiples of 60 seconds. Indeed, the excitation of this input data u_i is

$$\underline{\sigma}\left(\frac{1}{N}\sum_{i=1}^{N}\frac{u_i u_i^\top}{u_i^\top u_i}\right) = 0.0678$$

where $\underline{\sigma}(\,\cdot\,)$ is the smallest singular-value of a matrix. While the low-level of excitation makes it difficult to accurately estimate the parameters, it demonstrates the room for improvement through automation. The proposed R2R control algorithm

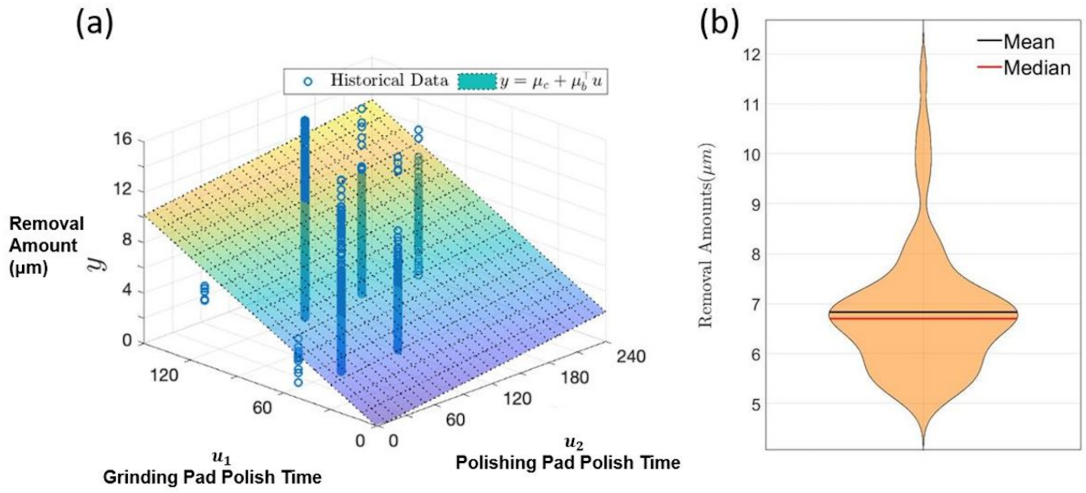

Figure 6: (a) Least squares regression of historical Robo-Met.3D™ data containing over 1000 measurements of removal rate for a variety of Polish times. (b) Violin plot of a single set of Polish time input parameters showing the vast majority of measured removal rates clusters around the mean and median values of the distribution.

will not artificially restrict itself to a small number of recipes with round numbers. This greater flexibility can potentially lead to improved performance. In future work, we will consider active-learning/dual-control to produce more exciting data for further improvements of our data-driven R2R control design.

Fig. 6 shows that material removal y_i is highly variable. Even when the same recipe $u_i = \bar{u}$ is used, the resulting removal y_i varies greatly. This is partially due to measurement noise, but the hidden parameters d_i play a significant role. Polishing a soft material will remove far more material than polishing a hard material for the same amount of time. Likewise, a fresh pad will remove material more quickly than an old pad even if both pads are assigned the same polishing

times. This high variance of the material removal shown in Fig. 6 motivates our decision to model the material removal as stochastic (6). This also motivates our objective of finding recipes that minimize the variance in the removal of material.

3.6 Optimal Run-to-Run Controller

In this section, we describe the proposed R2R controller for automating the MSS system. R2R control is the appropriate paradigm for this problem due to the lack of in-situ measurements and our inability to alter the recipe during a slice. Our algorithm embeds a deterministic formulation of the stochastic optimization problem (5) into a R2R framework in order to compute optimal recipes $u_{i+1}^{\star}$. The R2R framework provides a feedback mechanism for adjusting the recipe $u_{i+1}^{\star}$ based on the material removal y_i, which is measured after each slice. This feedback is used to reject the hidden variables d_i, which we consider as disturbances.

3.7 Optimal Recipe

The main challenge for MSS controller synthesis is that the plant model (3) mapping recipes u_i to removal amount y_i is an *unknown* static non-linearity. However, since the removal amount (3) is monotonic, we can use a linear approximation

$$y_i = c + b^{\top} u_i \tag{6}$$

where $c \approx f(u_i, d_i)$ is the drift coefficient and $b \approx \nabla f(u_i, d_i)$ is the slope coefficient around the operating point u_i. The parameters $c \in \mathbb{R}^1$ and $b \in \mathbb{R}^{n_u}$ are

uncertain and time-varying due to both the changing linearization point u_i and the hidden variables d_i. To capture this uncertainty, we will model these parameters as stochastic. Although the probability density function of these stochastic variables is unknown, we will use historical-data to quantify our uncertainty using their empirical moments. We will use these empirical moments to translate the conceptual stochastic optimization problem (5) into a deterministic optimization problem [37]. Since we consider stochastic *parameters*, estimating the parameters is non-trivial.

First, we translate the stochastic equality constraint (5b) into a deterministic constraint based on empirical moments. Substituting the stochastic linear model (6) into the equality constraint (5b) yields $\mathbb{E}[y_i|u_i] = \mathbb{E}[c + b^\top u_i|u_i] = r_i$. Exploiting the linearity of the expectation, we obtain $\mathbb{E}[c] + \mathbb{E}[b]^\top u_i = r_i$ where r_i and u_i are deterministic. This equality becomes the deterministic constraints (8b) when the expectations $\mathbb{E}[c]$ and $\mathbb{E}[b]$ are replaced by their empirical estimates $\mu_c \approx \mathbb{E}[c]$ and $\mu_b \approx \mathbb{E}[b]$.

Next, we translate the conceptual stochastic cost (5a) into a deterministic cost function. Substituting the stochastic linear model (6) into the cost (5a) yields

$$\mathbb{V}[y_i|u_i] = \mathbb{E}[(y_i - r_i)^2|u_i] = \mathbb{E}[(c + b^\top u_i - r_i)^2|u_i] \tag{7}$$

where the mean value of y_i is r_i due to the equality constraint (5b). Expanding the cost, yields $\mathbb{V}[y_i|u_i] \propto u_i^\top \mathbb{E}[bb^\top] u_i + 2u_i^\top \mathbb{E}[bc] - 2u_i^\top \mathbb{E}[b] r_i$ where the terms from

$\mathbb{E}[(c+r_i)^2]$ were omitted since they do not depend on the decision variables u_i. Substituting the empirical estimates $\mathbb{E}[bb^\top] \approx \Sigma_{bb}+\mu_b\mu_b^\top$ and $\mathbb{E}[bc] \approx \Sigma_{bc}+\mu_b\mu_c$, we obtain $\mathbb{V}[y_i|u_i] \propto u_i^\top(\Sigma_{bb}+\mu_b\mu_b^\top)u_i+2u_i^\top\Sigma_{bc}-2u_i^\top\mu_b(\mu_c-r_i)$. Since $\mu_c-r_i=-\mu_b^\top u_i$ according to (8b), we obtain $\mathbb{V}[y_i|u_i] \propto u_i^\top(\Sigma_{bb}-\mu_b\mu_b^\top)u_i+2u_i^\top\Sigma_{bc}$ Finally, noting that $\mu_b^\top u_i$ is constant, we obtain the deterministic cost (8a). Thus, the conceptual stochastic optimization problem (5) can now be approximated by the following deterministic optimization problem

$$u_i = \arg\min_u \quad u^\top\Sigma_{bb}u + 2u^\top\Sigma_{bc} \tag{8a}$$

$$\text{s.t.} \quad \mu_c + \mu_b^\top u_i = r_i \tag{8b}$$

$$u_i \in \mathcal{U} \tag{8c}$$

where the approximation is due to the use of empirical estimates of the means μ_c, μ_b and variances Σ_{bb}, Σ_{bc} of the parameters. Conveniently, this problem formulation (8) only requires second-order statics for the model (6) parameters. Solving (8) will produce the optimal recipe $u_i^\star$.

3.8 Run-to-Run Controller

The deterministic optimization problem (8) is static. Thus in this section, we embed this optimization problem (8) into an R2R framework to provide feedback. Our R2R algorithm iteratively adjusts the recipe $u_{i+1}^\star$ based on the measured material removed y_i during the previous slice i. The R2R feedback allows the ma-

Algorithm 1 Optimal Run-to-Run Control

1: Implement initial recipe u_0

2: **repeat**

3: Measure material remove y_i for i-th slice

4: Update (9) drift coefficient μ_c

5: Solve (8) for optimal recipe u_i

6: Implement recipe u_i

7: **until** All slices complete

terial removal to converge $y_i \to r$ to the desired removal amount r while rejecting the unmeasured disturbances d_i. Our R2R controller is described by Algorithm 1. The R2R Algorithm 1 indirectly adjusts the recipe u_i by updating the drift coefficient μ_c. After each slice, the R2R measures the resulting material removal y_i. The difference $y_i - r$ between the actual y_i and desired r removal amount is used to update the drift coefficient μ_c. We update the drift coefficient using exponentially weighted moving average (EWMA) dynamics [38]

$$\mu_{c,i+1} = \mu_{c,i} + \lambda(y_i - r) \tag{9}$$

where $\mu_{c,i+1}$ is the updated drift coefficient and $\lambda \in [0, 1]$ is a tuning parameter. The EWMA update-law has many beneficial properties [38]. The optimization problem (8) is solved with the updated drift coefficient μ_c to obtain a new recipe u_i which is then implemented. Thus, we can interpret the drift coefficient μ_c as a

state and the equality constraint (8b) as dynamics. The R2R controller continues to refine the recipe u_i until all slices have been completed.

3.9 Comparison with Existing R2R Controllers

In this section, we compare our R2R controller algorithm with existing R2R controllers from the literature to assess its value as a viable means of closed loop optimization. As noted in the survey [14], most R2R controllers have the following integral dynamics

$$u_{i+1} = u_i + \lambda \mu_b^+ (r - y_i) \tag{10}$$

where $\mu_b^+ = \mu_b / \mu_b^\top \mu_b$ is the Moore-Penrose pseudo-inverse and $\lambda \in [0, 1]$ is a tuning parameter. See equation (16) from [14] for details. We will show that our R2R controller has integral dynamics (10) when the input constraints (8c) are ignored, although with a novel pseudo-inverse. In contrast, when the input constraints are included the integral dynamics no longer apply. Without the input constraints (8c), the Karush-Kuhn-Tucker (KKT) optimality conditions for (8) are

$$\begin{bmatrix} 2\Sigma_{bb} & \mu_b \\ \mu_b^\top & 0 \end{bmatrix} \begin{bmatrix} u \\ \nu \end{bmatrix} = \begin{bmatrix} -2\Sigma_{bc} \\ r - \mu_c \end{bmatrix} \tag{11}$$

where $\nu \in \mathbb{R}$ is the dual variable associated with the equality constraint (8b). Solving (11) for u, we obtain the control-law

$$u = \mu_b^\dagger (r - \mu_c) + \tfrac{1}{2} \big(I - \mu_b^\dagger \mu_b^\top \big) \Sigma_{bb}^{-1} \Sigma_{bc} \tag{12}$$

where $\mu_b^\dagger = \Sigma_{bb}^{-1}\mu_b/(\mu_b^\top\Sigma_{bb}^{-1}\mu_b)$ is an alternative pseudo-inverse of μ_b i.e. $\mu_b^\top\mu_b^\dagger = \mu_b^\top\Sigma_{bb}^{-1}\mu_b/(\mu_b^\top\Sigma_{bb}^{-1}\mu_b) = 1$. This pseudo-inverse was derived from (8) to minimizes the variance $\mathbb{E}[(y-r)^2|u]$ of the material removal. Combining (12) with the EWMA dynamics (9), we obtain the following integral dynamics

$$\begin{aligned} u_{i+1} &= \mu_b^\dagger(r-\mu_{c,i}) + \tfrac{1}{2}\big(I-\mu_b^\dagger\mu_b^\top\big)\Sigma_{bb}^{-1}\Sigma_{bc} + \lambda\mu_b^\dagger(r-y_i) \\ &= u_i + \lambda\mu_b^\dagger(r-y_i) \end{aligned} \tag{13}$$

with the specific initial condition $u_0 = \mu_b^\dagger(r-\mu_{c0}) + \frac{1}{2}\big(I-\mu_b^\dagger\mu_b^\top\big)\Sigma_{bb}^{-1}\Sigma_{bc}$. This initial condition is important since otherwise the dynamic controller (13) would not include the second-term which compensates for possible cross-correlation between the model (6) parameters c and b. Note that our integral dynamics (13) match the literature dynamics (10).

Next, we show our R2R Algorithm does not necessarily have literature dynamics (10) when the input constraints (8c) are included. With the input constraints (8c), the KKT optimality conditions for (8) are

$$\begin{bmatrix} 2\Sigma_{bb} & \mu_b & H_{\mathcal{A}}^\top \\ \mu_b^\top & 0 & 0 \\ H_{\mathcal{A}} & 0 & 0 \end{bmatrix} \begin{bmatrix} u \\ \nu \\ \lambda_{\mathcal{A}} \end{bmatrix} = \begin{bmatrix} -2\Sigma_{bc} \\ r-\mu_c \\ h_{\mathcal{A}} \end{bmatrix} \tag{14}$$

where H and h are the half-space parameters of the input constraints (4) and $H_{\mathcal{A}}$ and $h_{\mathcal{A}}$ are the rows corresponding to the subset $\mathcal{A}$ of constraints that are active at the optimal. The active dual variables are denoted by $\lambda_{\mathcal{A}} \geq 0$ where the dual variable corresponding to inactive constraints are zero. Through brute-force

computation, we obtain

$$u = \mu_b^{\ddagger}(r - \mu_c) + \tfrac{1}{2}\big(I - \mu_b^{\ddagger}\mu_b^{\top}\big)(\Sigma_{bb}^{-1}\Sigma_{bc} + \Gamma h_{\mathcal{A}}) \tag{15}$$

where $\Gamma = \Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top}(H_{\mathcal{A}}\Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top})^{-1}$ and $\mu_b^{\ddagger}$ is yet another pseudo-inverse of μ_b given by

$$\mu_b^{\ddagger} = \frac{(\Sigma_{bb}^{-1} - \Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top}(H_{\mathcal{A}}\Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top})^{-1}H_{\mathcal{A}}\Sigma_{bb}^{-1})\mu_b}{\mu_b^{\top}(\Sigma_{bb}^{-1} - \Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top}(H_{\mathcal{A}}\Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top})^{-1}H_{\mathcal{A}}\Sigma_{bb}^{-1})\mu_b}. \tag{16}$$

Note that if (8) is feasible then μ_b does not lie in the null-space of: $\Sigma_{bb}^{-1} - \Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top}(H_{\mathcal{A}}\Sigma_{bb}^{-1}H_{\mathcal{A}}^{\top})^{-1}H_{\mathcal{A}}\Sigma_{bb}^{-1}$. This follows from the fact that the active inequality constraints cannot bind the equality constraint (8b).

Although (15) has a similar structure as (12), it cannot necessarily be transformed into the integral-form (10). As μ_c changes (9), the optimal active-set $\mathcal{A}$ can change. Thus, the pseudo-inverse $\mu_b^{\ddagger}$ and matrix Γ are time-varying. Thus, the nonlinear map provided by the optimization problem (8) replaces rather than integrates (10) the control inputs. Note that, although our R2R controller does not have the integral dynamics (10), it is still dynamic due to the EWMA dynamics (9). Finally, note that our R2R controller can be trivially put in the general form $u_{i+1} = \alpha u_i + \delta u_i$ given by equation (17) in [14] since any arbitrary feedback controller $\kappa(x)$ can be written in this form by defining $\delta u_i = \kappa(x) - \alpha u_i$.

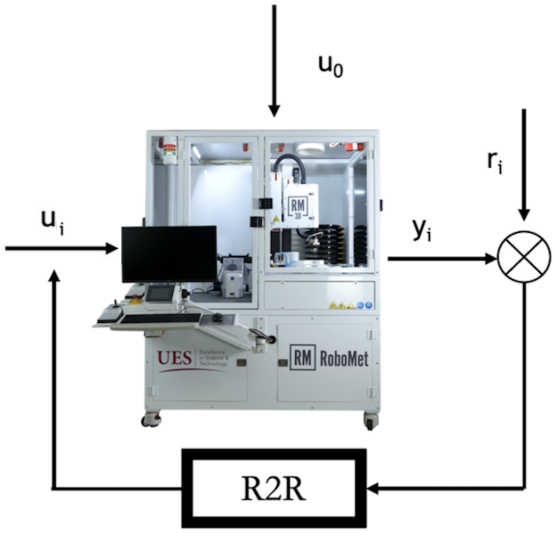

Figure 7: Closed-loop Robo-Met.3D™ system. u_0 denotes the initial recipe, y_i is removal amount produced which is then compared to the target removal amount r ($\otimes$ operation). The difference is acted upon by the R2R controller which calculates the input changes for the next slice u_i.

4 Results

In this section we cover initial simulation results and discuss a simulation comparison to a state-of-art R2R controller from literature. Finally, we cover physical system tests for both a previously executed and never before run recipe.

4.1 Simulation Results

Prior to implementation on the physical system, the R2R algorithm was tested initially via simulated runs with a targeted removal amount of 10 µm. Simulated runs were performed without the use of the Robo-Met.3D™ in order to explicitly test the efficacy of the R2R algorithm. This involves synthetic generation of

removal rates by assigning an amount of material removed per unit time to both the grinding pad and the polishing pad. This allows for direct calculation of a removal rate for a given set of input parameters to test the output behavior of the R2R algorithm.

The results of two simulated runs are shown in Figure 8, with Figure 8(a,b) showing a run with conservative initial conditions to determine if the R2R algorithm could achieve a targeted removal rate, and Figure 8(c,d) showing a run with simulated pad wear. An additional constraint of a maximum polish time of 200 s was imposed on the polishing pad to ensure the majority of material removal would occur on the grinding pad. Given initial conditions of 100 s for both pads in the first simulation, the R2R algorithm quickly increases the polish time for both pads, leading to a sharp increase in slice thickness from 5 μm to 10 μm after only a few slices. As the R2R algorithm favors adjustment of input parameters to reach the targeted removal in a minimal amount of time, there is a slight overshoot of the targeted removal rate on slice 4, and so for the remainder of the simulation, the R2R algorithm slowly decreases the amount of polishing time on the grinding pad to settle on the targeted removal rate.

A second simulation which incorporated an explicit modelling of pad wear was also performed. In this simulation, the prescribed removal rate was explicitly reduced over time. Pad wear was modelled using a simple linear decay function,

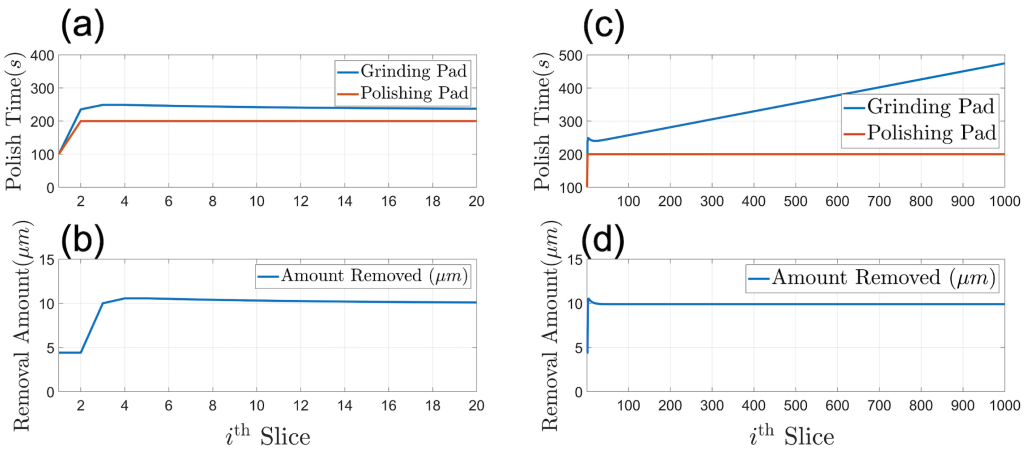

Figure 8: Simulated runs of the R2R algorithm. (a) Simulated inputs and (b) simulated removal amounts as a function of slice number for a recipe with a targeted removal amount of 10 µm for a 20-slice run. The system quickly finds the targeted removal rate, with a slow decay of the time on the Grinding Pad over the next 15 slices. (c) Simulated inputs and (b) simulated removal amounts as a function of slice number for a recipe with simulated pad wear. Pad wear was modelled using a linear decay function as given by Equation 17

.

given by the following:

$$y_i = y_0 \left(1 - x \cdot i\right) \tag{17}$$

where y_i is the removal rate on slice i, y_0 is the initial removal rate of the pad, x is the decay factor, and i is the slice number. For values of $x > 0$, the pad will remove less material with each subsequent slice. For the simulation shown in Figure 8(c,d), pad wear was modelled at 1% decay per slice, ($x = 0.01$). Although actual pad wear is a complex process, depending on the hardness of material being sectioned, the weight applied to the robotic polishing arm, and the amount

of polishing solution dispensed, a simple linear approximation of this process is instructive as it simulates one of the dynamic processes at play during MSS experiments. The R2R algorithm is able to successfully accommodate pad wear by continually increasing the amount of polishing times on both pads. Given the maximum time constraint on the polishing pad, this pad quickly reaches its maximum polish time. The grinding pad, on the other hand, has an initial large jump similar to the first R2R simulation which then leads to a slowly reducing removal rate in the first 20 slices. As the pad wear becomes more significant with increased slice number, the polishing time steadily increases at a linear rate, matching the linear decay of the pad's effectiveness with simulated pad wear. The ability of the R2R algorithm to adequately account for simulated pad wear demonstrates the efficacy of the closed loop approach in accommodating the complex dynamics of the Robo-Met.3D™ system. This demonstrates the framework's ability to handle any variety of real-life performance degradations or reduction in consumable effectiveness over time as is common in metallographic preparation. Therefore, the framework can be seen to be capable of adjusting for any removal rate changes over time and still optimize polishing parameters to achieve the desired outcome.

We then compared our R2R controller in Algorithm 1 with the literature controller (10) using the Moore-Penrose pseudo-inverse $\mu_b^+ = \mu_b/\mu_b^\top\mu_b$ of μ_b. We also compare our controller without the input constraints (8c). As we showed in

Section 3.9, without (8c) our controller has the form (10) with the pseudo-inverse

$$\mu_b^\dagger = \frac{\Sigma_{bb}^{-1}\mu_b}{\mu_b^\top \Sigma_{bb}^{-1}\mu_b}. \tag{18}$$

For these simulation results, we model the removal function (3) as

$$y_i = \mu_c + \mu_b^\top u_i$$

where y_i is the removal amount, μ_c is estimated variance of the output, μ_b is the estimated variance of the inputs, and u_i is the inputs.

The simulation results are shown in Fig. 9. For each of the three R2R algorithms we show the removal amount y_i and the recipe u_i versus slice i. The desired removal amount is $r = 10.5$ μm.

Each of the R2R algorithms converged to the desired removal amount $y_i \to r$ after 14 slices. This is fast convergence considering an experiment is typically comprised of hundreds of slices. Furthermore, a human operator can require up to 40 slices or more to find an appropriate recipe for a unique sample with adjustments needed during the course of an experiment. However, the linear R2R controller (10) produced different recipes to achieve the desired removal amount. In the next section, we examine the benefits of the pseudo-inverse (18).

Figure 9 and 10 show both the linear R2R controller and the unconstrained controller (10) produce non-implementable recipes [37]. In Figure 9, the literature R2R controller (10) with the Moore-Penrose pseudo-inverse produces a negative polishing time for one of the pads, which is obviously unimplementable. Figure

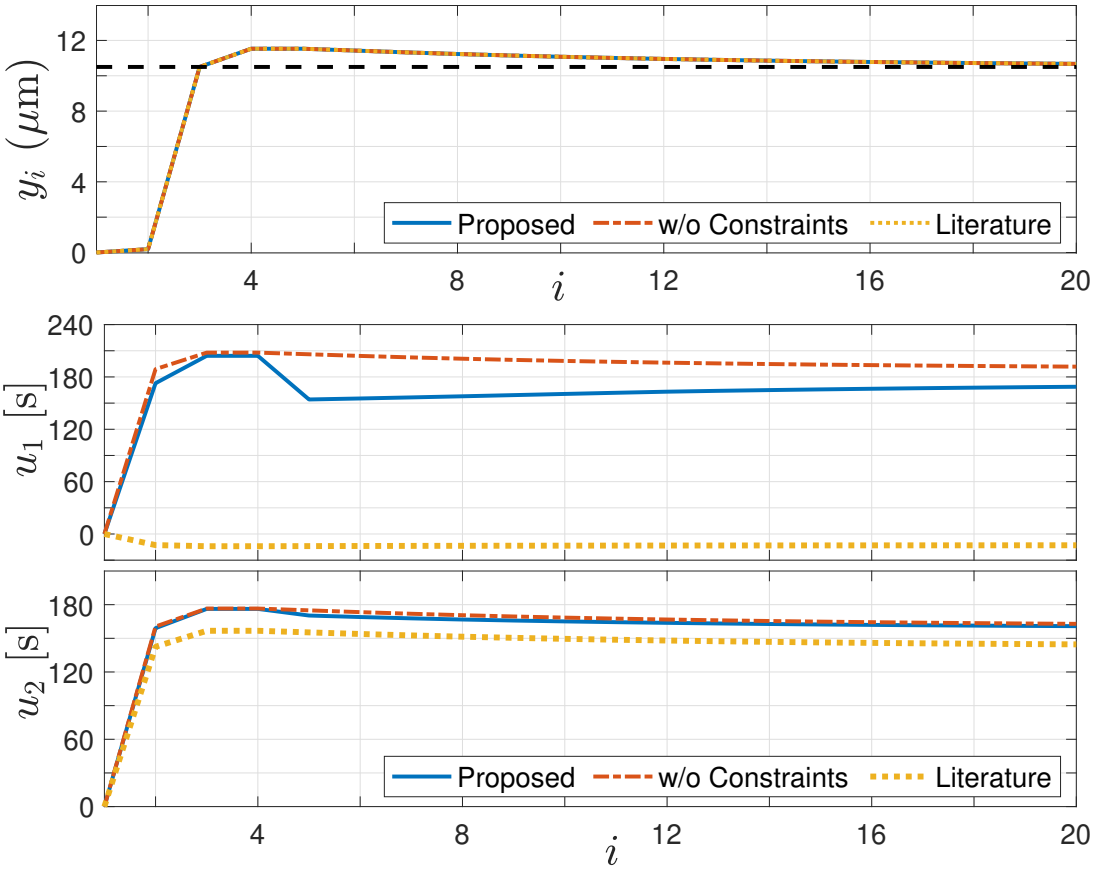

Figure 9: Simulation results comparing the proposed R2R control with and without constraints with the literature R2R controller using the Moore-Penrose pseudo-inverse.

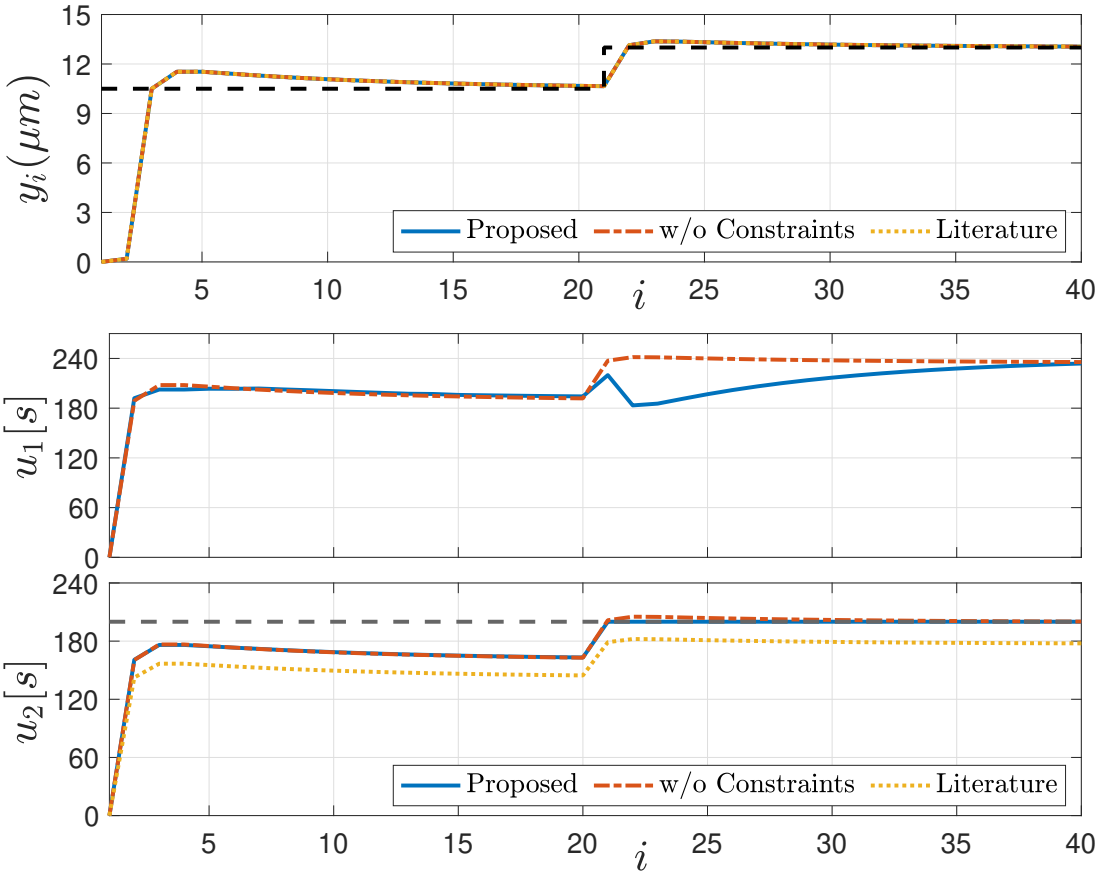

Figure 10: Simulation results comparing the proposed R2R control with and without constraints with the literature R2R controller (10) using the Moore-Penrose pseudo-inverse.

10 shows that changes in the target removal amount leads to the unconstrained variant of our controller with the pseudo-inverse (18) produce an excessively long polishing time for one of the pads that violated the constraints.

4.2 Validation of Closed Loop Optimization

Given the success of the simulated application of the R2R algorithm, the closed loop controls were applied to a real-world sample in order to physically test the approach. Physical tests were performed on a mounted sample of copper that was cast within a steel crucible. This sample was selected as it was comprised of multiple metals, a common feature for many MSS samples, and also had a non-prismatic shape, so the cross-sectional area of the sample changed significantly along the serial-sectioning direction. The first physical test comprised a 20-slice run with a target removal rate of 10 μm, which is similar to those observed in the historical dataset, and therefore represents a test of the R2R algorithm in an input parameter space that has been well-sampled. The results of this physical test are shown in Figure 11 from a set of initial conditions of 100 s on each pad, and a maximum polish time of 400 s for the polishing pad. After an initial learning phase, the target slice thickness is achieved from the seventh slice onward, with an average removal rate of 9.47 μm per slice, which is approximately 95% of the targeted removal rate of 10 μm per slice. This local removal rate would be expected to continually approach the targeted removal rate over the next several slices.

Figure 11(b) shows that the R2R algorithm continually adjusts the polishing time on both pads as it accommodates the complex dynamics of the system, a feature not observed in the simulated experiments. Sample heights for this physical test were also measured using the 5× objective lens. Although the R2R algorithm does a good job of achieving the targeted average removal rate, it does not ameliorate the variance inherent to the auto-focus process, with standard deviations in slice thickness of 4.9 µm, which are similar to what was observed in the examples shown in Figure 3 and Figure 4.

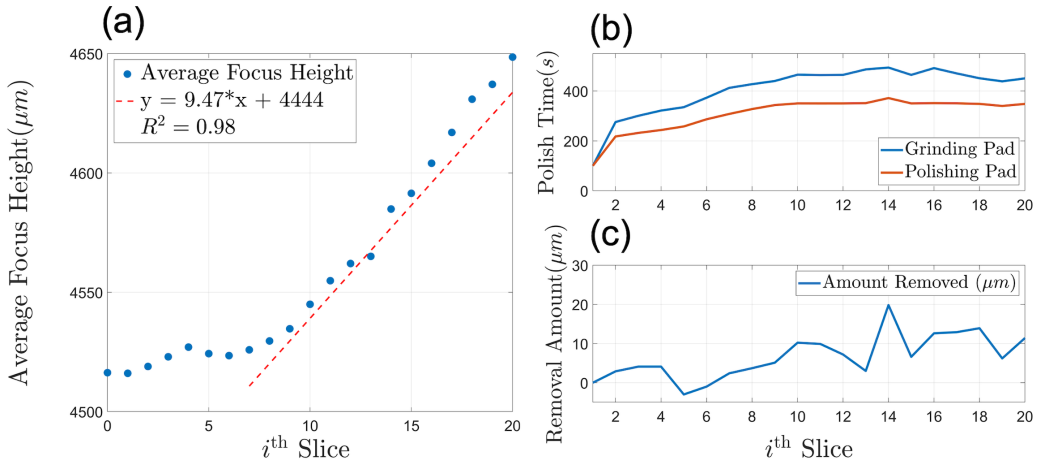

Figure 11: Physical test for targeted removal rate with extensive historical data. (a) Removal rate as a function of slice number for the 20 slice run. The R2R algorithm makes continual adjustments to system parameters throughout the run, with the targeted material removal rate of 10 µm achieved within 6 slices. (b) Polish times for both pads as determined by the R2R algorithm. (c) Slice-to-slice removal rate as measured by average focus height.

A second physical test was also performed on the same sample of cast copper

within a steel crucible, but with a targeted removal rate of 25 μm, which is far outside the bounds of the historical data. This test was designed to evaluate the flexibility of the R2R algorithm in exploring new parameter spaces, and also helps to assess its efficacy in handling a diverse range of samples and target removal rates which may or may not have the benefit of being previously explored. The targeted removal rate of 25 μm was coupled with a 200 s limit on polishing time for the polishing pad, and no initial starting conditions. The results of this physical test are shown in Figure 12. Given the large targeted removal rate and the lack of initial conditions, the R2R algorithm quickly increases the polish time on both pads. These large changes initially create an overshoot of the material removal rate for the first three slices, causing the R2R to over-correct with very small removal rates from slice 4 to slice 7 (Figure 12(b)). Despite the lack of initial conditions and large change in targeted removal rate, the R2R algorithm is able to generate an average removal rate of 21.4 μm per slice before the 10^{th} slice, which is within 84% of the target removal rate. As with the first physical test, this removal rate is expected to continually improve with additional slices in the run as the trend in data (see Figure 12a) shows. In both physical tests of the R2R algorithm, the average removal rate achieved by the system is less than the targeted removal rate [35]. This may be a result of the historic database having more values below the expected value determined by the linear regression, as evidenced by the large tail on the upper end of the distribution as shown in

Figure 6(b), or by the specific constraints of the optimization problem. However, this phenomena was not observed in the simulated experiments, in which the R2R algorithm initially exceeded the targeted removal rate before asymptotically approaching the target. Continued use of the R2R algorithm will reveal whether the trend of over- or under-shooting of the target removal rate is random or biased toward one behavior or the other. However, the authors suspect increasing historical data within recipe space previously unexplored may improve the R2R initial prediction and converge response toward the desired removal rate.

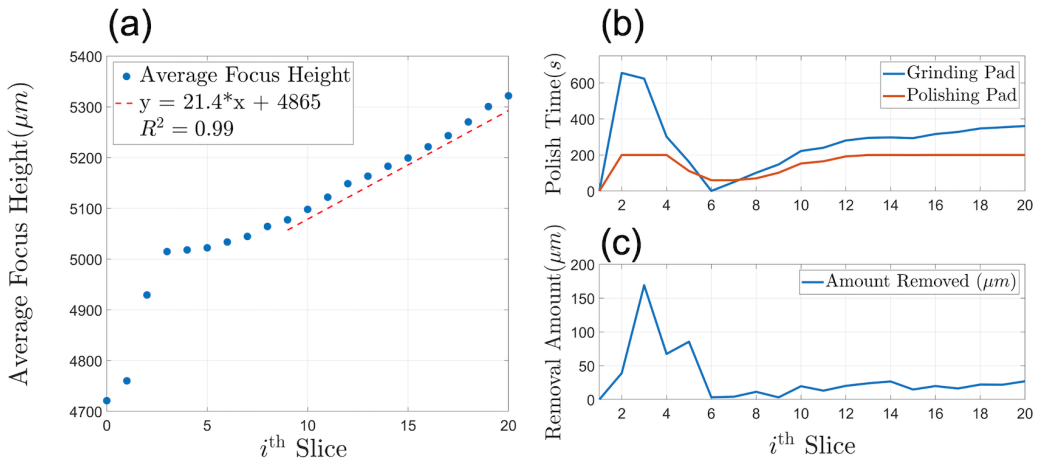

Figure 12: Physical system test for targeted removal rate without any historical data. (a) Removal rate as a function of slice number for the 20 slice run. The R2R algorithm makes large changes to system parameters in the first 7 slices, then makes more minor modifications to achieve a removal rate close to the target of 25 µm per slice. (b) Polish times for both pads as determined by the R2R algorithm. (c) Slice-to-slice removal rate as measured by average focus height.

5 Conclusion

5.1 Impact of Work

Taking advantage of over a decade of historical data, a data driven approach for modeling the dynamics of the Robo-Met.3D™ system using a linear least squares regression has been applied. Using a deterministic model, a constrained optimization problem was formulated, with the goal of minimizing the variance of removal amount, while also converging the removal amount towards an operator specified target amount. By inserting the constrained optimization problem into an R2R control framework, the Robo-Met.3D™ was deployed as a closed-loop, feedback controlled system. The following conclusions are offered:

- Using an engineering controls approach, we have developed an experimentally accurate estimate and system abstraction for a mathematical model of our Robo-Met.3D™ automated mechanical serial-sectioning system.

- Using a run-to-run control framework, we have developed and successfully demonstrated a means to transform an open-loop automated mechanical serial-sectioning system into a closed-loop operation that can iteratively revise inputs to produce an optimized experimental setup for a given criteria

- Using historical data from a decade of experiments, an optimization algorithm was trained which, when implemented, was shown to converge to

within 95% of a pre-determined target removal rate within 10 iterations or fewer for a previously executed experiment and within 84% of the target for a never-before run experiment

5.2 Future Work

As a recommendation for future, the authors suggests prioritizing full integration of the R2R controller into the physical MSS system to take advantage of the autonomous nature of the algorithm and the system. The authors also suggests improving the current model by creating a more complex model of the system through methods such as Gaussian Process Regression or "Kriging" to increase the accuracy of the stochastic model. Furthermore, the authors concur that exploration of various image processing techniques to identify sample surface quality to parameterize as a system constraint would be particularly impactful and provide an additional dimension of complexity and benefit to the R2R closed loop control approach. Finally, the authors suggest exploring active-learning/dual-control to improve the quality of the operational data used to synthesize the controller.

5.3 Acknowledgements

This material is based upon work supported by the National Science Foundation under NSF Grant Number CNS-2105631. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. [†]University of New Mexico. [‡]Sandia National Laboratories.

REFERENCES

[1] McLean P. Echlin, Timothy L. Burnett, Andrew T. Polonsky, Tresa M. Pollock, and Philip J. Withers. Serial sectioning in the sem for three dimensional materials science. *Current Opinion in Solid State and Materials Science*, 24 (2):100817, 2020. ISSN 1359-0286. doi: https://doi.org/10.1016/j.cossms.2020.100817. URL https://www.sciencedirect.com/science/article/pii/S1359028620300152.

[2] G Noguere, F Cserpak, C Ingelbrecht, AJM Plompen, CR Quetel, and P Schillebeeckx. Non-destructive analysis of materials by neutron resonance transmission. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 575(3): 476–488, 2007.

[3] K.D. Vernon-Parry. Scanning electron microscopy: an introduction. *III-Vs Review*, 13(4):40–44, 2000. ISSN 0961-1290. doi: https://doi.org/10.1016/S0961-1290(00)80006-X. URL https://www.sciencedirect.com/science/article/pii/S096112900080006X.

[4] B.E. Warren. *X-ray Diffraction*. Addison-Wesley series in metallurgy and materials engineering. Dover Publications, 1990. ISBN 9780486663173.

[5] Bruno De Man, Johan Nuyts, Patrick Dupont, Guy Marchal, and Paul Suetens. Metal streak artifacts in x-ray computed tomography: a simulation study. In *1998 IEEE Nuclear Science Symposium Conference Record. 1998 IEEE Nuclear Science Symposium and Medical Imaging Conference (Cat. No. 98CH36255)*, volume 3, pages 1860–1865. IEEE, 1998.

[6] Larry K Aagesen and Jonathan D Madison. Porosity in millimeter-scale welds of stainless steel : three-dimensional characterization. 5 2012. doi: 10.2172/1044948. URL https://www.osti.gov/biblio/1044948.

[7] D. Pidhayny. The origins of feedback control. *IEEE Transactions on Automatic Control*, 17(2):283–284, 1972. doi: 10.1109/TAC.1972.1099919.

[8] D.S. Bernstein. Feedback control: an invisible thread in the history of technology. *IEEE Control Systems Magazine*, 22(2):53–68, 2002. doi: 10.1109/37.993315.

[9] T. Nguyen, J. Wang, J. Hwang, and K. Alder. Force accommodation control of the space shuttle remote manipulator system: a unique problem. In *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, pages 2510,2511,2512,2513,2514,2515, Los Alamitos, CA, USA, apr 1991. IEEE Computer Society. doi: 10.1109/ROBOT.1991.132003. URL https://doi.ieeecomputersociety.org/10.1109/ROBOT.1991.132003.

[10] Xiaohui Li, Zhenping Sun, Qi Zhu, and Daxue Liu. A unified approach to local trajectory planning and control for autonomous driving along a reference path. In *2014 IEEE International Conference on Mechatronics and Automation*, pages 1716–1721, 2014. doi: 10.1109/ICMA.2014.6885959.

[11] Sandip Roy. Feedback control of soil moisture in precision-agriculture systems: Incorporating stochastic weather forecasts. In *2014 American Control Conference*, pages 2694–2698, 2014. doi: 10.1109/ACC.2014.6858834.

[12] Yunkai Lv and Ronghu Chi. Data-driven adaptive iterative learning predictive control. In *2017 6th Data Driven Control and Learning Systems (DDCLS)*, pages 374–377, 2017. doi: 10.1109/DDCLS.2017.8068100.

[13] Johanna Wallén, Mikael Norrlöf, and Svante Gunnarsson. Arm-side evaluation of ilc applied to a six-degrees-of-freedom industrial robot. *IFAC Proceedings Volumes*, 41(2):13450–13455, 2008. ISSN 1474-6670. doi: https://doi.org/10.3182/20080706-5-KR-1001.02278. URL https://www.sciencedirect.com/science/article/pii/S1474667016411444. 17th IFAC World Congress.

[14] Youqing Wang, Furong Gao, and Francis J. Doyle. Survey on iterative learning control, repetitive control, and run-to-run control. *Journal of Process Control*, 19(10):1589–1600, 2009. ISSN 0959-1524.

[15] Edi Kurniawan, Zhenwei Cao, and Zhihong Man. Design of robust repetitive control with time-varying sampling periods. *IEEE Transactions on Industrial Electronics*, 61(6):2834–2841, 2014. doi: 10.1109/TIE.2013.2276033.

[16] O Forsman. Undersökning av rymdstrukturen hos ett kolstå av hypereutectoid sammansättning. *Jernkontorets Ann*, 102:1–30, 1918.

[17] DA Hull, D McCammond, DW Hoeppner, and WG Hellier. Titanium prior-beta grain volume distribution by quantitative serial sectioning techniques. *Materials characterization*, 26(2):63–71, 1991.

[18] Tetsuo Sakamoto Tetsuo Sakamoto, Zhaohui Cheng Zhaohui Cheng, Masanori Takahashi Masanori Takahashi, Masanori Owari Masanori Owari,

and Yoshimasa Nihei Yoshimasa Nihei. Development of an ion and electron dual focused beam apparatus for three-dimensional microanalysis. *Japanese Journal of Applied Physics*, 37(4R):2051, 1998.

[19] J Alkemper and PW Voorhees. Quantitative serial sectioning analysis. *Journal of microscopy*, 201(3):388–394, 2001.

[20] Thomas A Ivanoff and Jonathan D Madison. Using 3d characterization and serial sectioning to improve materials analysis: Features, spatial locations, and morphological measures not readily attainable with 2d methods are now accessible using 3d characterization techniques. *Advanced Materials & Processes*, 178(1):16–22, 2020.

[21] Michael D Uchic, Michael A Groeber, and Anthony D Rollett. Automated serial sectioning methods for rapid collection of 3-d microstructure data. *Jom*, 63(3):25–29, 2011.

[22] G Spanos, DJ Rowenhorst, AC Lewis, and AB Geltmacher. Combining serial sectioning, ebsd analysis, and image-based finite element modeling. *Mrs Bulletin*, 33(6):597–602, 2008.

[23] J. D. Madison, O. D. Underwood, G. A. Poulter, and E. M. Huffman. Acquisition of real-time operation analytics for an automated serial sectioning system. *Integrating Materials and Manufacturing Innovation*, 6(2):135–146, 2017. doi: 10.1007/s40192-017-0091-6.

[24] David J Rowenhorst, Lily Nguyen, Aeriel D Murphy-Leonard, and Richard W Fonda. Characterization of microstructure in additively manufactured 316l using automated serial sectioning. *Current Opinion in Solid State and Materials Science*, 24(3):100819, 2020.

[25] McLean P Echlin, Alessandro Mottura, Christopher J Torbet, and Tresa M Pollock. A new tribeam system for three-dimensional multimodal materials analysis. *Review of Scientific Instruments*, 83(2):023701, 2012.

[26] McLean P Echlin, Naji S Husseini, John A Nees, and Tresa M Pollock. A new femtosecond laser-based tomography technique for multiphase materials. *Advanced Materials*, 23(20):2339–2342, 2011.

[27] Jonathan E Spowart. The 3-d analysis of discontinuously reinforced aluminum composite microstructures. *JOM*, 58(12):29–33, 2006.

[28] S. Adivikolanu and E. Zafiriou. Internal model control approach to run-to-run control for semiconductor manufacturing. In *Proceedings of the 1997*

American Control Conference (Cat. No.97CH36041), volume 1, pages 145–149 vol.1, 1997. doi: 10.1109/ACC.1997.611772.

[29] W.J. Campbell, S.K. Firth, A.J. Toprac, and T.F. Edgar. A comparison of run-to-run control algorithms. In *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, volume 3, pages 2150–2155 vol.3, 2002. doi: 10.1109/ACC.2002.1023955.

[30] Zheng Ying and Hou Mingyang. On variable weight ewma based run-to-run control in mixed product process. In *2008 27th Chinese Control Conference*, pages 6–10, 2008. doi: 10.1109/CHICC.2008.4604924.

[31] Jan Busch and Wolfgang Marquardt. Run-to-run control of membrane filtration in wastewater treatment - an experimental study 1 1the financial support by the dfg (german research foundation) in the project "optimization-based process control of chemical processes" and the excellent cooperation with koch membrane systems gmbh are gratefully acknowledged. *IFAC Proceedings Volumes*, 40(5):195–200, 2007. ISSN 1474-6670. doi: https://doi.org/10.3182/20070606-3-MX-2915.00080. URL https://www.sciencedirect.com/science/article/pii/S1474667015317237. 8th IFAC Symposium on Dynamics and Control of Process Systems.

[32] E. Del Castillo and Jinn-Yi Yeh. An adaptive run-to-run optimizing controller for linear and nonlinear semiconductor processes. *IEEE Transactions on Semiconductor Manufacturing*, 11(2):285–295, 1998. doi: 10.1109/66.670178.

[33] K. Miwa, T. Inokuchi, T. Takahashi, A. Oikawa, and K. Imaoka. Reduction of depth variation in an si etching process by applying an optimized run-to-run control system. *IEEE Transactions on Semiconductor Manufacturing*, 18 (4):517–521, 2005. doi: 10.1109/TSM.2005.858473.

[34] E.S. Hamby, P.T. Kabamba, and P.P. Khargonekar. A probabilistic approach to run-to-run control. *IEEE Transactions on Semiconductor Manufacturing*, 11(4):654–669, 1998. doi: 10.1109/66.728563.

[35] D. Gallegos-Patterson, K. Ortiz, C. Danielson, J. D. Madison, and A. T. Polonsky. A framework for closed-loop optimization of an automated mechanical serial-sectioning system via run-to-run control as applied to a robomet.3d. *JOM*, 74(8):2930–2940, 2022. doi: 10.1007/s11837-022-05372-3.

[36] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 1982.

[37] D. Gallegos-Patterson, K. Ortiz, C. Danielson, J. D. Madison, and A. T. Polonsky. Constrained run-to-run control for precision serial sectioning.

[38] James Moyne. Run-to-run control in semiconductor manufacturing. In *Encyclopedia of Systems and Control*. Springer London, 2014.