Inference for Deformation and Interference in 3D Printing

Arman Sabbaghi* Tirthankar Dasgupta† Qiang Huang‡ Jizhe Zhang§

Abstract

Additive manufacturing, or 3D printing, is a promising manufacturing technique. Material solidification in the printing process leads to product deformation, which reduces the utility of printed products. Control of product deformations can be achieved by compensation plans. However, little attention has been paid to interference in compensation, which is thought to result from the inevitable discretization of a compensation plan. We investigate interference with an experiment involving the application of discretized compensations to cylinders. Our treatment illustrates a principled framework for detecting interference, and for addressing complications in modeling deformations under discretized compensations, which facilitates the study of printing processes. Properly defining experimental units and understanding interference are critical for quality control in complex engineering processes. Our work ultimately provides a step in that direction for 3D printing.

Key Words: 3D printing, additive manufacturing, posterior predictive checks

1. Interference in Shrinkage Compensation

Additive manufacturing, or 3D printing, refers to a class of technologies for the direct fabrication of physical products from 3D Computer-Aided Design (CAD) models. In contrast to material removal processes in traditional machining, the printing process adds material layer by layer. This technique enables direct printing of products with complex geometry (e.g., complicated gear structures), without affecting building efficiency. No extra effort is necessary for molding construction or fixture tooling design, making 3D printing a promising manufacturing technique (Hilton and Jacobs, 2000; Gibson, Rosen, and Stucker, 2009; Melchels, Feijen, and Grijpma, 2010; Campbell, Williams, Ivanova, and Garrett, 2011). Despite these promising features, control of a product's printed dimensions (i.e., dimensional accuracy control) remains a major bottleneck. Material solidification during layer formation leads to product deformation, or shrinkage (Wang, Cheah, Fuh, and Lu, 1996). Shrinkage control is crucial to overcome the accuracy barrier in 3D printing.

To control detailed features along the boundary of a printed product, Tong, Lehtihet, and Joshi (2003) and Tong et al. (2008) used polynomial regression models to analyze shrinkage in different directions separately, and compensate for product deformation by changing the original CAD design accordingly. Unfortunately, their predictions of shrinkage are independent of the product's geometry, which is not consistent with the manufacturing process. Huang, Zhang, Sabbaghi, and Dasgupta (2012) built on this work, using polar coordinates to develop a physically consistent approach to model and compensate shrinkage. Validation experiments suggest that this approach achieves greater accuracy.

An important issue not yet addressed is how the application of compensation to one section of a product will affect the deformation of neighboring sections. Technically, compensation plans are always discretized (according to the tolerance of the 3D printer), in the

*Harvard University, Department of Statistics, 1 Oxford Street, 7th Fl., Cambridge, MA 02138

†Harvard University, Department of Statistics, 1 Oxford Street, 7th Fl., Cambridge, MA 02138

‡University of Southern California, Daniel J. Epstein Department of Industrial and Systems Engineering, 3715 McClintock Avenue, Los Angeles, CA 90089

§University of Southern California, Daniel J. Epstein Department of Industrial and Systems Engineering, 3715 McClintock Avenue, Los Angeles, CA 90089

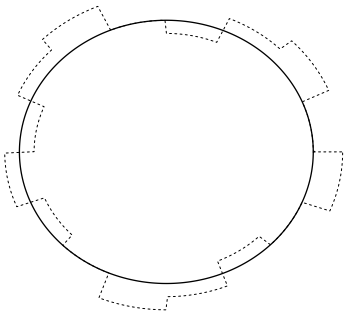

Figure 1: A discretized compensation (dashed-line) to the nominal boundary (solid line).

sense that sections of the CAD design are altered by a single amount, as, for example, in Figure 1. Furthermore, when planning an experiment to assess the effect of compensations on product deformation, it is natural to discretize the quantitative "compensation" factor into a finite number of levels, which also leads to a product having a more complex boundary. Ultimately, such changes to a product's boundary may introduce interference between different sections of the printed product. Interference is defined to occur when one section's deformation depends not only on its assigned compensation, but also on compensations applied to its neighbors (Rubin, 1980). For example, in Figure 1, the deformation for points that lie at the boundary of two neighboring sections should depend on compensations applied to both sections. By the same logic, interference becomes a practical issue when printing products with complex geometry. Therefore, to better understand dimensional accuracy control of 3D printed products, it is of broad importance to formally investigate complications introduced by interference, which result from the inevitable discretization in compensation plans. We take the first step with an experiment involving a discretized compensation plan for a simple geometric shape, or profile.

We begin in Section 2 with a review of interference, models for shape shrinkage, and the effect of compensation. Adoption of the Rubin Causal Model (RCM, Holland, 1986) is a significant and novel feature of our investigation in direct 3D printing, and greatly facilitates discussions on interference. Section 3.1 summarizes the basic model for shape shrinkage of cylinders given by Huang et al. (2012). Our analyses are in Sections 3.2 - 3.5: we first describe an experimental design hypothesized to generate interference, then proceed with posterior predictive checks to demonstrate the existence of interference, and finally conclude with a model that captures interference. A statistically significant idea arising in Section 3.3 is that, in experiments explicitly making a distinction between *units of analysis* and *units of interpretation* (Cox and Donnelly, 2011, p. 18–19), the posterior distribution of model parameters, constructed using "benchmark" data, naturally leads to a simple assessment and inference for interference similar to that suggested by Sobel (2006) and Rosenbaum (2007). Analyses in Sections 3.4 and 3.5 demonstrate how discretized compensations complicate dimensional accuracy control along a product's boundary through the introduction of interference. This illustrates the fact that in complex engineering processes, a proper definition of experimental units and understanding of how units may interfere with each other are critical to quality control.

2. Potential Outcomes and Interference

2.1 Definition of Experimental Units and Potential Outcomes

Suppose a product has intended shape ψ_0, and observed shape ψ. Shape deformation is informally described as the difference between ψ and ψ_0, where we can represent both either in the Cartesian coordinate system (x, y, z) or cylindrical coordinate system (r, θ, z). Cylindrical coordinates greatly facilitate modeling of shrinkage, and are used throughout.

For illustrative purposes, we define terms for two-dimensional products (notation for three dimensions follow immediately). Improved product accuracy requires an understanding of what happens when different regions of the CAD model receive different amounts of compensation. We therefore define a finite number N points on the boundary of the product, corresponding to specific angles $\theta_1, \ldots, \theta_N$, as the experimental units. The desired boundary of a product, dictated by the CAD model, is represented by a function $r_0(\theta)$ denoting the nominal radius at angle θ. We consider only one (quantitative) treatment factor, compensation to the CAD model, defined as changing the CAD model to either increase or decrease the nominal radius by x_i units at θ_i. The potential radius for θ_i under the application of treatments $\mathbf{x} = (x_1, \ldots, x_N)$ to units $\theta_1, \ldots, \theta_N$ is a function of θ_i, CAD model $r_0(\cdot)$, and $\mathbf{x}$, denoted by $r(\theta_i, r_0(\cdot), \mathbf{x})$. The difference between the potential and nominal radius at θ_i defines deformation at θ_i. Therefore, we define

$$\Delta r(\theta_i, r_0(\cdot), \mathbf{x}) = r(\theta_i, r_0(\cdot), \mathbf{x}) - r_0(\theta_i) \tag{1}$$

as our *potential outcome* for θ_i. For now, potential outcomes are fixed numbers. Randomness is introduced in Section 2.3 in our general model for the potential outcomes.

This definition of the potential outcome is convenient for visualizing shrinkage. For example, suppose the desired shape of the product is the solid-line, and the manufactured product when $\mathbf{x} = \mathbf{0} = (0, \ldots, 0)$ is the dashed-line, in Figure 2(a)). Plotting the shrinkage at each angle leads to a visualization (Figure 2(b)) amenable to analysis. Orientation is fixed: we match the coordinate axes of the printed product with those of the CAD model.

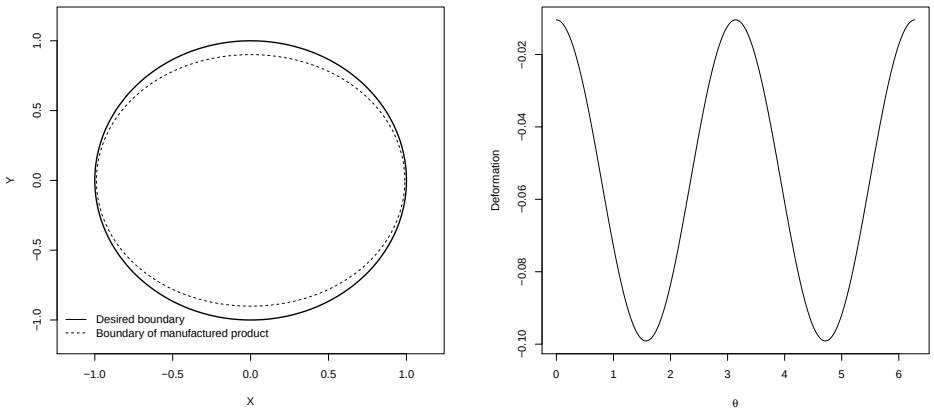

(a) (b)

Figure 2: **(a)**: Ideal shape vs. printed shape. **(b)**: Shrinkage visualization.

2.2 Interference

A unit θ_i is said to be affected by interference if

$$\Delta r(\theta_i, r_0(\cdot), \mathbf{x}) \neq \Delta r(\theta_i, r_0(\cdot), \mathbf{x}')$$

for at least one pair of distinct treatment vectors $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^N$ with $x_i = x_i'$ (Rubin, 1980). If there is no interference, then

$$\Delta r(\theta_i, r_0(\cdot), \mathbf{x}) = \Delta r(\theta_i, r_0(\cdot), x_i).$$

As the experimental units reside on a connected boundary, the deformation of one unit may depend on compensations assigned to its neighbors when the compensation plan is discretized (according to the tolerance of the 3D printer). Perhaps less plausible, but equally serious, is the leakage of assigned compensations across units. All these considerations explain the presence of the vector $\mathbf{x}$, containing treatment assignments for all units, in the potential outcome notation (1). Practically, accommodations made for interference should reduce bias in compensation plans for complex products.

2.3 The General Deformation Model

Potential outcomes under treatment assignment $\mathbf{x} = \mathbf{0}$ are decomposed into

$$\Delta r(\theta_i, r_0(\cdot), \mathbf{0}) = f_1(r_0(\cdot)) + f_2(\theta_i, r_0(\cdot), \mathbf{0}) + \epsilon_i. \tag{2}$$

Function $f_1(r_0(\cdot))$ represents average deformation of a given nominal shape $r_0(\cdot)$ independent of location θ_i, and $f_2(\theta_i, r_0(\cdot), \mathbf{0})$ is the additional location-dependent shrinkage, geometrically and physically related to the CAD model $r_0(\theta)$. We can also interpret $f_1(\cdot)$ as a lower order component and $f_2(\cdot, \cdot, \mathbf{0})$ as a higher order component of shrinkage. The ϵ_i are random variables representing high frequency components that add on to the main trend, with expectation $\mathbb{E}(\epsilon_i) = 0$ and $\mathrm{Var}(\epsilon_i) < \infty$ for all $i = 1, \ldots, N$. In all that follows, the ϵ_i are assumed to be independent and identically distributed. Correlation among the ϵ_i can be addressed in various ways (e.g., (Colosimo, Semeraro, and Pacella, 2008)), and is the focus of our future work.

Figure 2 demonstrates model (2), with $r_0(\theta) = r_0$, so $f_1(\cdot)$ is a function of r_0, and $f_2(0, r_0, \mathbf{0}) = f_2(2\pi, r_0, \mathbf{0})$. Decomposition of shrinkage reduces (2) to

$$\Delta r(\theta_i, r_0, \mathbf{0}) = c_{r_0} + \sum_k \{a_{r_0,k} \cos(k\theta_i) + b_{r_0,k} \sin(k\theta_i)\} + \epsilon_i, \tag{3}$$

where $f_1(r_0) = c_{r_0}$, and $\{a_{r_0,k},\ b_{r_0,k}\}$ are coefficients of a Fourier expansion of $f_2(\cdot, \cdot, \mathbf{0})$. The $\{a_{r_0,k},\ b_{r_0,k}\}$ with large k represent surface roughness, and is not of primary interest.

2.4 General Compensation and Interference Models

Under the polar coordinate system, a compensation of x_i units at θ_i can be thought of as an extension of the product's radius by x_i units in that specific direction. Bearing this in mind, we now extend (2) to accommodate compensations.

Let $r(\theta_i, r_0(\cdot), (x_i, \ldots, x_i)) = r(\theta_i, r_0(\cdot), x_i)$ denote the potential radius for θ_i under application of x_i units of compensation to all units. Assuming that the manufacturing and deformation dynamics remain the same under compensation, we note that

$$r(\theta_i, r_0(\cdot), x_i) - \{r_0(\theta_i) + x_i\} = \Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}).$$

Consequently, the potential outcome for θ_i is

$$\begin{aligned}\Delta r(\theta_i, r_0(\cdot), x_i) &= r(\theta_i, r_0(\cdot), x_i) - r_0(\theta_i) \\ &= \Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}) + x_i. \end{aligned} \tag{4}$$

If x_i is small relative to $r_0(\theta_i)$, then the expectation of (4) can be approximated using the first and second terms of the Taylor expansion of $\mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0})\}$ at $r_0(\theta_i)$:

$$\begin{aligned}\Delta r(\theta_i, r_0(\cdot), x_i) &\approx \mathbb{E}\{\Delta r(\theta_i, r_0(\cdot), \mathbf{0})\} \\ &\quad + (x_i - 0)\left[\frac{d}{dx}\mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x, \mathbf{0})\}\right]_{x=0} + x_i + \epsilon_i \\ &= \Delta r(\theta_i, r_0(\cdot), \mathbf{0}) + \{1 + h(\theta_i, r_0(\cdot), \mathbf{0})\}x_i, \end{aligned} \tag{5}$$

where $h(\theta_i, r_0(\cdot), \mathbf{0}) = \left[\frac{d}{dx}\mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x, \mathbf{0})\}\right]_{x=0}$. When a parametric model is specified, this expansion is performed conditional on the model parameters.

Interference can then be incorporated into (5) in a simple manner for a treatment assignment $\mathbf{x}$. Specifically, although unit θ_i is assigned compensation x_i, we do not use only x_i to model the potential outcome for θ_i. Instead, we alter (5) to

$$\Delta r(\theta_i, r_0(\cdot), \mathbf{x}) \approx \Delta r(\theta_i, r_0(\cdot), \mathbf{0}) + \{1 + h(\theta_i, r_0(\cdot), \mathbf{0})\}g_i(\mathbf{x}), \tag{6}$$

where the *effective treatment* $g_i(\mathbf{x})$ is a function of x_i and compensations applied to neighbors of θ_i (with the definition of neighboring units naturally dependent on the specific product), hence potentially a function of the entire vector $\mathbf{x}$. This change from x_i to $g_i(\mathbf{x})$ in (5) and (6) is based on the belief that the physical logic ultimately leading to (5) holds true on some level for general compensation plans. The simplest way to introduce interference in the potential outcomes model is to allow the treatment effect for θ_i to not be a function solely of x_i, but to depend on treatments assigned to neighboring units as well. As we shall see in the analysis of our experiment, this modification of the potential outcomes model effectively incorporates interference in a meaningful manner.

3. Experimental Design and Analysis for Interference

3.1 Compensation Model for Cylinders

Huang et al. (2012) constructed four cylinders with $r_0 = 0.5, 1, 2$, and 3 inches, and used $N_{0.5} = 749, N_1 = 707, N_2 = 700$, and $N_3 = 721$ equally-spaced units from each, displayed in Figure 3(a). They fitted

$$\Delta r(\theta_i, r_0, \mathbf{0}) = x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i) + \epsilon_i \tag{7}$$

to the data, according to the reasoning in Section 2.3, with $\epsilon_i \sim \mathrm{N}(0, \sigma^2)$ independently, and parameters α, β, a, b, x_0, and σ independent of r_0. Independent errors were used because the correlation of the ϵ_i was not substantial enough to merit being modeled.

They specified

$$a \sim \mathrm{N}(1, 2^2),\ b \sim \mathrm{N}(1, 1^2),\ \log(x_0) \sim \mathrm{N}(0, 1^2),$$

and placed flat priors on α, β, and $\log(\sigma)$, with all parameters independent *a priori*. Posterior draws of the parameters were obtained by Hamiltonian Monte Carlo (HMC, Duane, Kennedy, Pendleton, and Roweth, 1987), with 1000 draws after a burn-in of 500 (summarized in Table 1. A simple comparison of the posterior predictive distribution of product

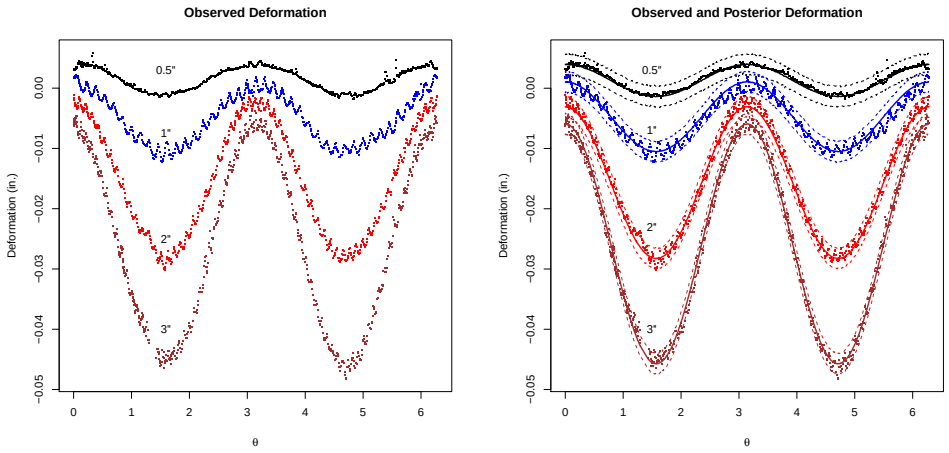

(a) (b)

Figure 3: **(a)**: Observed deformation for $r_0 = 0.5, 1, 2$, and 3 inches, when no compensation is applied. **(b)**: Comparison of observed data with posterior predictions of deformation when no compensation is applied. Bold solid lines denote posterior means, and dashed lines denote the 2.5% and 97.5% posterior quantiles of shrinkage for each angle.

Table 1: Summary of posterior draws of parameters when no compensation is applied.

	Mean	SD	Median	95% Credible Interval	ESS
α	-1.34×10^{-2}	1.6×10^{-4}	-1.34×10^{-2}	$(-1.37, -1.31) \times 10^{-2}$	8198
β	5.7×10^{-3}	3.1×10^{-5}	5.71×10^{-3}	$(5.65, 5.8) \times 10^{-3}$	9522
a	8.61×10^{-1}	7.33×10^{-3}	8.61×10^{-1}	$(8.47, 8.75) \times 10^{-1}$	8223
b	1.13	5.46×10^{-3}	1.13	(1.12, 1.14)	9424
x_0	8.79×10^{-3}	1.5×10^{-4}	8.79×10^{-3}	$(8.5, 9.07) \times 10^{-3}$	8211
σ	8.7×10^{-4}	1.18×10^{-5}	8.7×10^{-4}	$(8.5, 8.9) \times 10^{-4}$	9384

deformation to the observed data (Figure 3(b)) demonstrates the good fit, and so we proceed with this specification.

Substituting $\Delta r(\theta_i, r_0, \mathbf{0})$ from (7) into the general model (4), we have

$$\Delta r(\theta_i, r_0, x_i) = x_0 + x_i + \alpha(r_0 + x_0 + x_i)^a + \beta(r_0 + x_0 + x_i)^b \cos(2\theta_i) + \epsilon_i. \quad (8)$$

Further approximation by the Taylor expansion at $r_0 + x_0$, as in (5), yields

$$\begin{aligned}\Delta r(\theta_i, r_0, x_i) = {} & x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i) \\ & + \left\{1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1} \cos(2\theta_i)\right\} x_i + \epsilon_i. \quad (9)\end{aligned}$$

We incorporate interference by changing x_i in (9) to $g_i(\mathbf{x})$, with the functional form of $g_i(\mathbf{x})$ derived by exploratory means in Section 3.3.

3.2 Experimental Design for Interference

Under a discretized compensation plan, the boundary of a product is divided into sections, with all units in one section assigned the same compensation and neighboring sections having different compensations. In the terminology of Cox and Donnelly (2011, p. 18–19), these sections constitute units of analysis, and individual angles are units of interpretation.

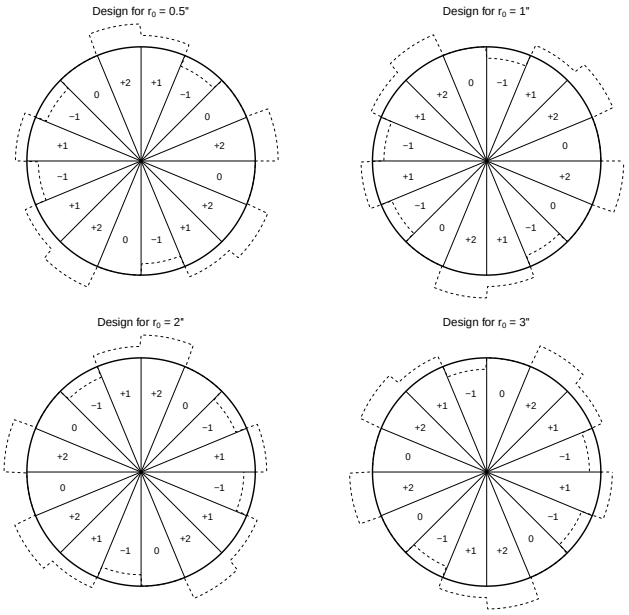

Figure 4: Experimental designs: dashed lines represent assigned compensations.

We then expect interference for angles near neighboring sections. Interference should be substantial for large differences in neighboring compensations, and should become negligible as the differences in compensations decrease.

This reasoning led to the following restricted Latin square design to study interference. We apply compensations to four cylinders of radius $0.5, 1, 2$, and 3 inches, with each cylinder divided into 16 equal-sized sections of $\pi/8$ radians. One unit of compensation is $0.004, 0.008, 0.016$, and 0.03 inch for each respective cylinder, and there are only four possible levels of compensation, $-1, 0, +1$, and $+2$ units. Two blocking factors are considered. The first is the quadrant, and the second is the "symmetry group" consisting of units in $\pi/8$-radian sections that are reflections about the coordinate axes from each other. Symmetric sections form a meaningful block: if compensation x is applied to all units, then we have from (9) that for $0 \leq \theta \leq \pi/2$,

$$\begin{aligned}\mathbb{E}\left\{\Delta r(\theta, r_0, x) \mid \alpha, \beta, a, b, x_0, \sigma\right\} &= \mathbb{E}\left\{\Delta r(\pi - \theta, r_0, x) \mid \alpha, \beta, a, b, x_0, \sigma\right\} \\ &= \mathbb{E}\left\{\Delta r(\pi + \theta, r_0, x) \mid \alpha, \beta, a, b, x_0, \sigma\right\} \\ &= \mathbb{E}\left\{\Delta r(2\pi - \theta, r_0, x) \mid \alpha, \beta, a, b, x_0, \sigma\right\},\end{aligned}$$

suggesting a need to control for this symmetry in the experiment. Finally, based on prior concerns about the possible severity of interference and resulting scope of inference from our Taylor-expansion model (5), treatments were randomly assigned to the 16 sections under a further constraint that the absolute difference in assigned treatments between two sections cannot exceed two levels of compensation.

Our restricted Latin square design blocks on two factors suggested by the previous deformaton model (and yet is model robust to a certain extent) to form a discretized compensation plan, as desired. The chosen experimental designs are in Figure 4, and results of this experiment are in Figure 5. There are $N_{0.5} = 6159, N_1 = 6022, N_2 = 6206$, and $N_3 = 6056$ equally-spaced observations used for the four cylinders of nominal radius $r_0 = 0.5, 1, 2$, and 3 inches throughout.

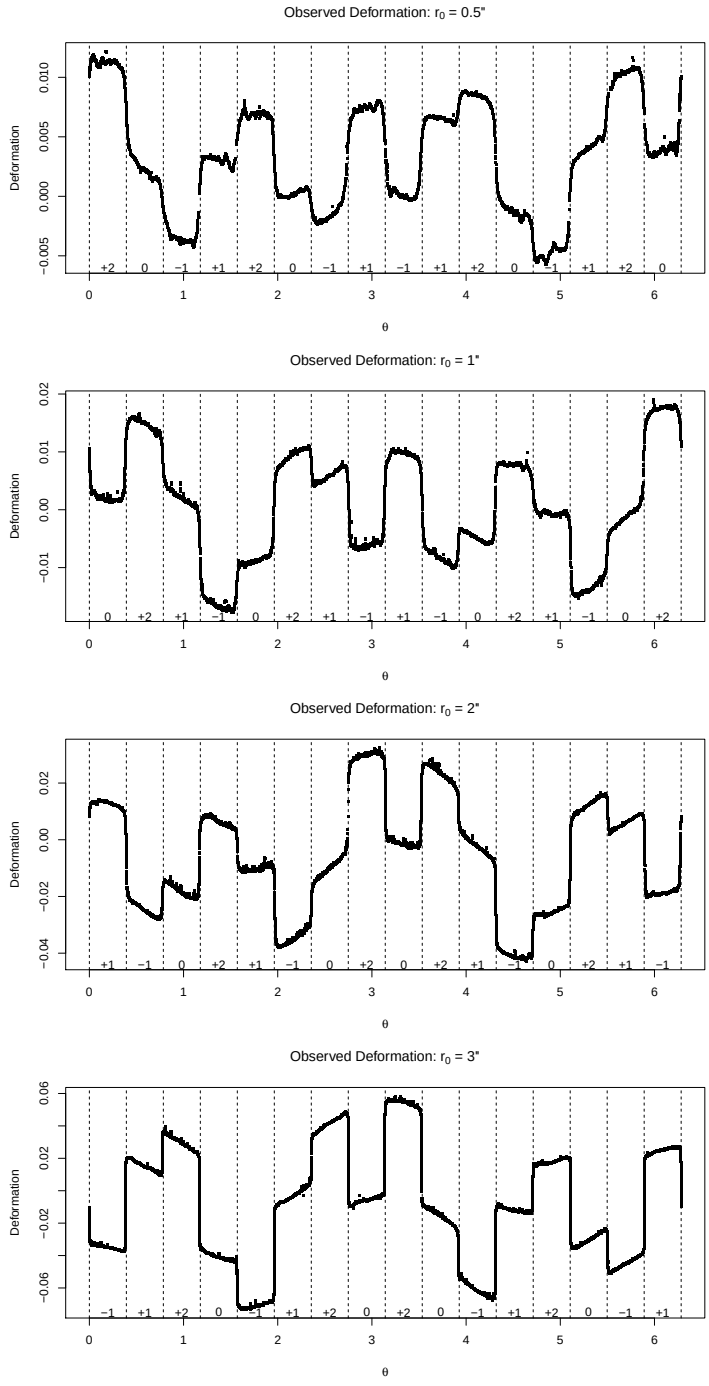

Figure 5: Observed deformations in the experiment. Dashed lines represent sections, and numbers at the bottom of each represent applied compensations.

3.3 Assessing the Structure of Interference

If a unit θ_i has negligible interference under treatment assignment $\mathbf{x} = (x_1, \ldots, x_N)$, then

$$\Delta r(\theta_i, r_0, \mathbf{x}) = \Delta r(\theta_i, r_0, (x_i, \ldots, x_i)) = \Delta r(\theta_i, r_0, x_i).$$

Our first task is to assess which units have negligible interference. To do so, we use the suggestions of Sobel (2006) and Rosenbaum (2007), who describe when interest exists in comparing a treatment assignment $\mathbf{x}$ to a baseline.

We have in Section 3.1 data on cylinders that receive no compensation, and a model (7) that provides a good fit. Furthermore, we have a hypothesized model (9) when interference is negligible, which is a function of parameters in (7). If the manufacturing process is in control, posterior inferences made conditional on the no-compensation data in Section 3.1 (denoted by $\mathbf{D}_n$) then yield, by (9), predictions for the experiment. In the absence of any other information, units in the experiment with observed outcomes deviating strongly from their predictions can be argued to have substantial interference. This suggests the following procedure to assess interference:

1. Calculate the posterior distribution of the parameters conditional on $\mathbf{D}_n$, which we denote by $\pi(\alpha, \beta, a, b, x_0, \sigma \mid \mathbf{D}_n)$.

2. For every angle in the four cylinders, form the posterior predictive distribution of the potential outcome corresponding to the observed treatment assignment (Figure 4) using model (9) and $\pi(\alpha, \beta, a, b, x_0, \sigma \mid \mathbf{D}_n)$.

3. Compare the posterior predictive distributions to the experiment.

    - If a unit's observed outcome falls within the 99% central posterior predictive interval and follows the posterior predictive mean trend, it is deemed to have negligible interference.

    - Otherwise, we conclude that the unit has substantial interference.

This is similar to the construction of control charts (Box et al., 2009). When an observed outcome lies outside the 99% posterior predictive interval, we suspect existence of a special cause. As the product is constructed simultaneously, the assignable cause is interference.

We implemented this procedure and observed that approximately $70\% - 80\%$ of units have negligible interference, primarily in the central regions of sections. This is clearly seen with another graph that assesses effective treatments, which we proceed to describe. Taking expectations in (9), the treatment received by θ_i is

$$\frac{\mathbb{E}\left\{\Delta r(\theta_i, r_0, \mathbf{x}) \mid \alpha, \beta, a, b, x_0, \sigma\right\} - x_0 - \alpha(r_0 + x_0)^a - \beta(r_0 + x_0)^b \cos(2\theta_i)}{1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1}\cos(2\theta_i)}. \tag{10}$$

We then use (10) to gauge $g_i(\mathbf{x})$, i.e., the effective treatment for θ_i, by plugging observed data from the experiment and posterior draws of the parameters based on $\mathbf{D}_n$ into this equation. These calculations, summarized in Figure 6, again suggest that central angles in each section have negligible interference: estimates of their effective treatments correspond to their assigned treatments. There is a slight discrepancy between assigned treatments nd posterior predictive quantities for some central angles, but this is likely due to different parameter values for the two data sets. Of more importance is the observation that the effective treatment of a boundary angle θ_i is a weighted average of treatments assigned to its section and its nearest neighboring section, with the weights a function of the distances (in radians) between θ_i and the midpoint angle of its section, $\theta_{i,M}$, and the midpoint angle of its nearest neighboring section, $\theta_{i,NM}$. All these observations correspond to the intuition that interference should be substantial near section boundaries, and negligible in the middle.

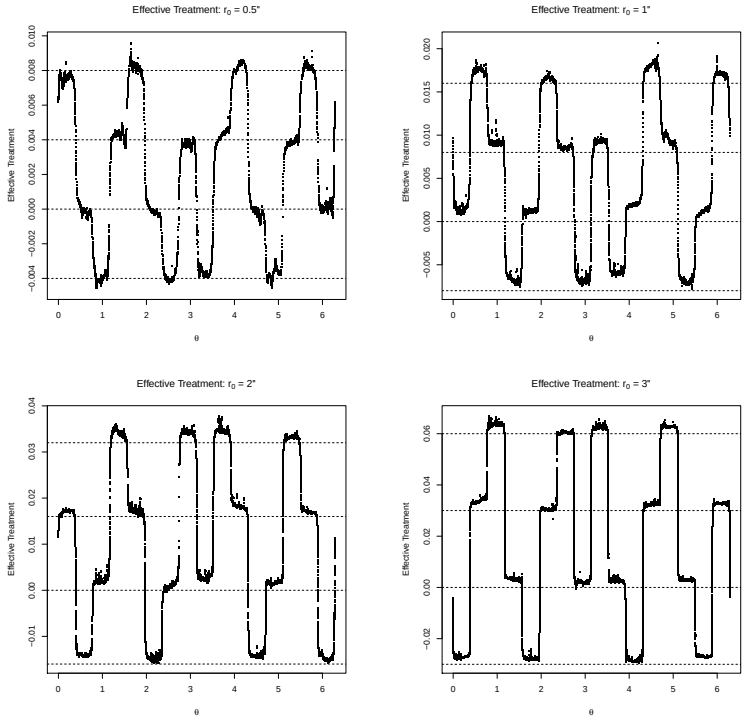

Figure 6: Gauging effective treatments $g_i(\mathbf{x})$ using (10). Four horizontal lines in each subfigure denote the possible compensations for the cylinders, and the dots denote estimates of treatments that these units effectively received in the experiment.

3.4 A Simple Interference Model

We first alter (9) to

$$\Delta r(\theta_i, r_0, \mathbf{x}) = x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i) + \left\{1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1} \cos(2\theta_i)\right\} g_i(\mathbf{x}) + \epsilon_i, \quad (11)$$

where

$$g_i(\mathbf{x}) = \{1 + \exp(-\lambda_{r_0}|\theta_i - \theta_{i,NM}| + \lambda_{r_0}|\theta_i - \theta_{i,M}|)\}^{-1} x_{i,M} + \{1 + \exp(\lambda_{r_0}|\theta_i - \theta_{i,NM}| - \lambda_{r_0}|\theta_i - \theta_{i,M}|)\}^{-1} x_{i,NM}, \quad (12)$$

with $\theta_{i,M}, \theta_{i,NM}$ denoting midpoint angles for the $\pi/8$-radian sections containing and neighboring nearest to θ_i, respectively, and $x_{i,M}, x_{i,NM}$ treatments assigned to these sections. Effective treatment $g_i(\mathbf{x})$ is a weighted average of the unit's assigned treatment $x_i = x_{i,M}$ and the treatment $x_{i,NM}$ assigned to its nearest neighboring section. Although the functional form of the weights is chosen for computational convenience, we recognize that (12) belongs to a class of models agreeing with prior subject-matter knowledge that interference may be negligible if the implemented compensation is sufficiently "continuous", in the sense that the compensation plan is a continuous function of θ and the tolerance of the 3D printer is sufficiently fine so that discretization of compensation is negligible.

We fit the model in (11), (12), having 10 total parameters, to the experiment data. The prior specification remains the same, with $\log(\lambda_{r_0}) \sim \mathrm{N}(0, 4^2)$ independently *a priori* for $r_0 = 0.5, 1, 2$, and 3 inches. A HMC algorithm was used to obtain 1000 draws from the joint posterior distribution after a burn-in of 500, and these are summarized in Table 2.

Table 2: Summary of posterior draws for simple interference model.

	Mean	SD	Median	95% Credible Interval	ESS
α	-1.06×10^{-2}	1.53×10^{-4}	-1.06×10^{-2}	$(-1.09, -1.03) \times 10^{-2}$	8078
β	5.79×10^{-3}	3.69×10^{-5}	5.79×10^{-3}	$(5.72, 5.86) \times 10^{-3}$	8237
a	9.5×10^{-1}	9.46×10^{-3}	9.5×10^{-1}	$(9.31, 9.69) \times 10^{-1}$	8150
b	1.12	6.64×10^{-3}	1.12	(1.0, 1.13)	8504
x_0	7.1×10^{-3}	1.43×10^{-4}	7.1×10^{-3}	$(6.82, 7.39) \times 10^{-3}$	8404
σ	3.14×10^{-3}	1.36×10^{-5}	3.14×10^{-3}	$(3.11, 3.17) \times 10^{-3}$	8924
$\lambda_{0.5}$	32.66	2.05	32.62	(28.69, 36.76)	8686
λ_1	48.24	2	48.12	(44.5, 52.6)	8666
λ_2	76.83	1.78	76.78	(73.42, 80.44)	8770
λ_3	86.08	0.83	86.06	(84.49, 87.68)	8385

This model provides a good fit to the observed data for 0.5 and 1 inch cylinders. However, predictions for the other cylinders are not as good. As an example, in Figure 7(a) the posterior mean trend does not capture the observed transition across sections.

The problem appears to reside in (12). This specification for the weights implies that effective treatments for units $\theta_i = k\pi/8$ ($k \in \mathbb{Z}_{>0}$) are averages of compensations applied to units $k\pi/8 \pm \pi/16$. To assess the validity of this implication, we use the posterior distribution of the parameters to calculate, for each θ_i, the inferred effective treatment in (10). An example of these calculations, Figure 7(b), shows that the inferred effective treatment for $\theta_i = \pi$ is nearly 0.06 inch, the compensation applied to its neighboring section. Thus, specification (12) is invalidated by the experiment.

Another posterior predictive check helps clarify the problem. From (12),

$$g_i(\mathbf{x}) = w_i x_{i,M} + (1 - w_i) x_{i,NM}.$$

We solve for w_i to obtain

$$w_i = \frac{g_i(\mathbf{x}) - x_{i,NM}}{x_{i,M} - x_{i,NM}}. \tag{13}$$

Plugging in the inferred effective treatments, calculated from (10), into (13), we then diagnose how to modify (12) to better model interference in the experiment.

This calculation was made for all cylinders, and the results for $r_0 = 3''$ are summarized in Figure 8 as an example. Rows in this figure show the weights for each of the four quadrants, and we focus on the behavior of the weights in neighborhoods of integral multiples of $\pi/8$. Neither the decay in the weights (represented by λ_{r_0} in (12)) nor the weight for integral multiples of $\pi/8$ remain constant across sections. In fact, these figures suggest that λ_{r_0} is a function of $\theta_{i,M}, \theta_{i,NM}$, and that a location term is required.

3.5 A Refined Interference Model

Our refined effective treatment model is of the same functional form as (12), with λ_{r_0} replaced by $\lambda_{r_0}(\theta_{i,M}, \theta_{i,NM})$, and $|\theta_i - \theta_{i,M}|$ and $|\theta_i - \theta_{i,NM}|$ replaced by $|\theta_i - \theta_{i,M} - \delta_{r_0}(\theta_{i,M}, \theta_{i,NM})|$ and $|\theta_i - \theta_{i,NM} - \delta_{r_0}(\theta_{i,M}, \theta_{i,NM})|$, respectively. Here, $\delta_{r_0}(\theta_{i,M}, \theta_{i,NM})$ represent location shifts across sections suggested by our posterior predictive checks.

Our specific model is

$$\delta_{r_0}(\theta_{i,M}, \theta_{i,NM}) = \delta_{r_0,0} + \sum_{k=1}^{3} \left\{ \delta^c_{r_0,k} \cos(k\theta_{i,B}) + \delta^s_{r_0,k} \sin(k\theta_{i,B}) \right\}, \tag{14}$$

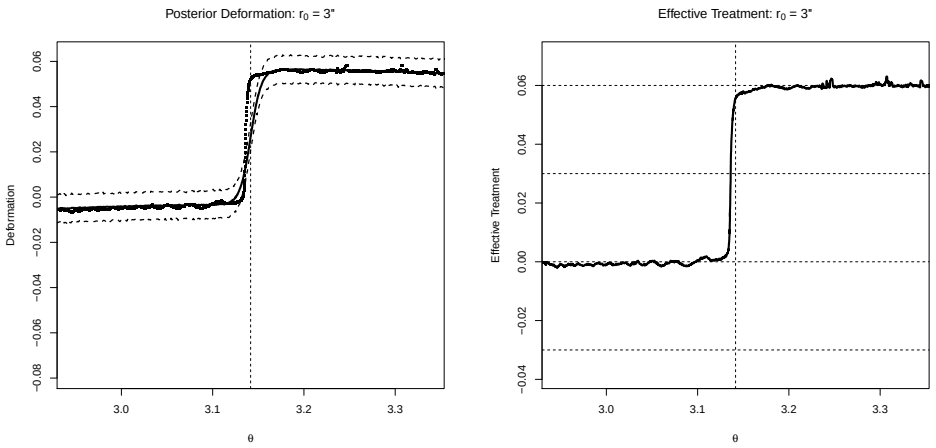

(a)

(b)

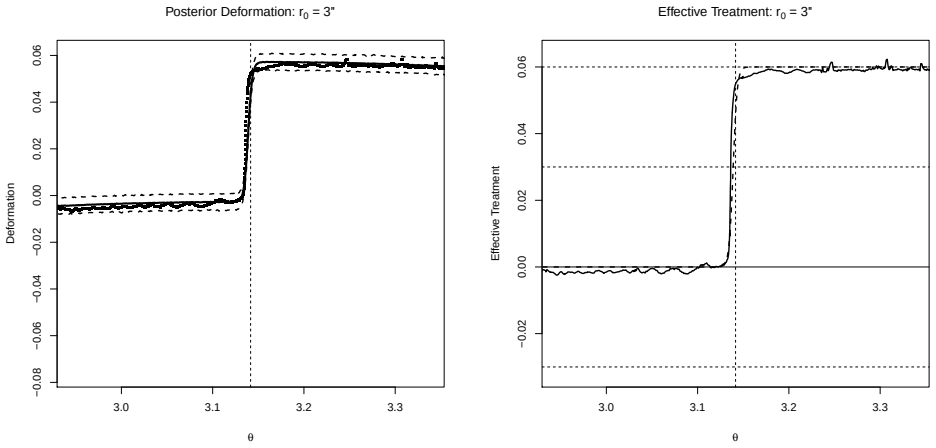

(c)

(d)

Figure 7: **(a)**: An example of the type of erroneous predictions made by model (11), (12) for the 3 inch cylinder. The vertical line is drawn at $\theta = \pi$, marking the boundary between two sections, with units to the left given $+1$ compensation, and units to the right given 0.
(b): Corresponding inferred effective treatment for $15\pi/16 \leq \theta \leq 17\pi/16$.
(c): Refined posterior predictions for $r_0 = 3''$, $15\pi/16 \leq \theta \leq 17\pi/16$.
(d): Comparing inferred effective treatments (solid line) with refined effective treatment model (dashed line) for the 3 inch cylinder.

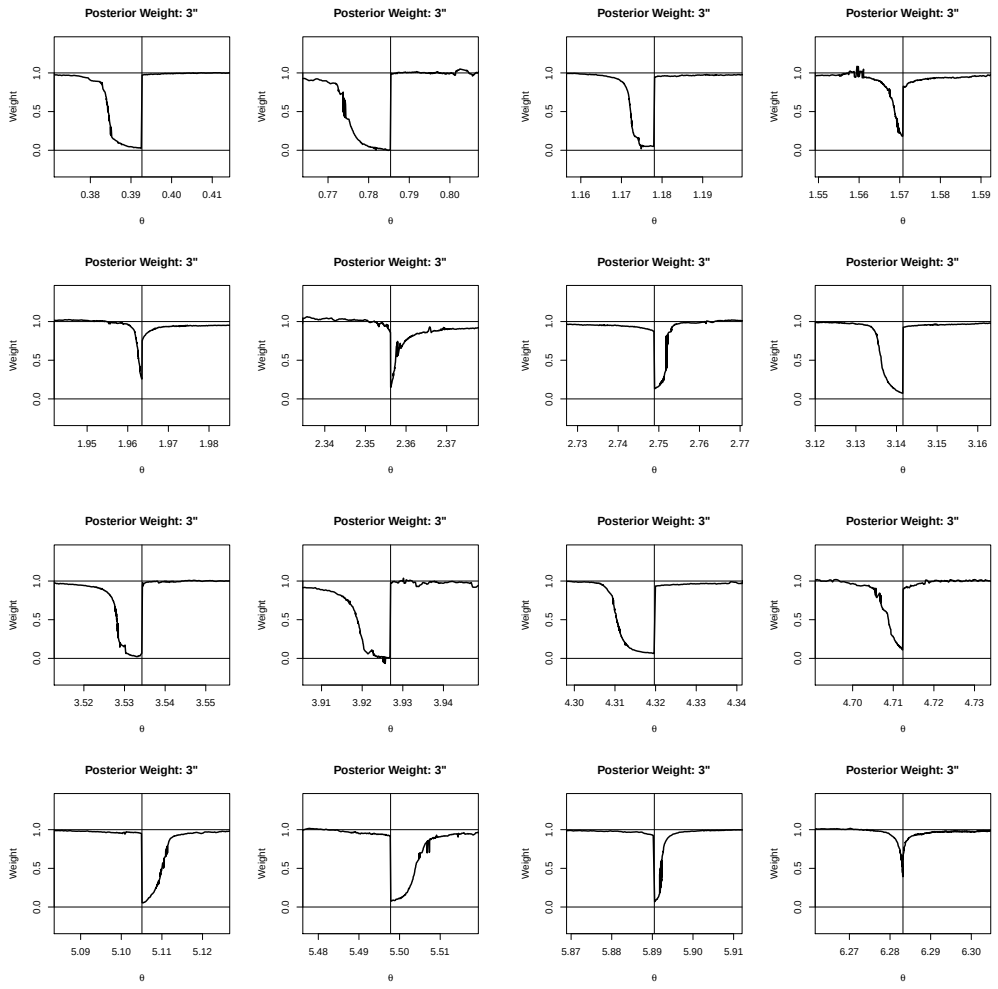

Figure 8: Inferring weights in the interference model for the $r_0 = 3''$ cylinder, using effective treatments calculated from equation (10) and the posterior distribution of parameters from Section 3.4. Vertical lines are drawn at $\theta = k\pi/8$ for $k = 1, \ldots, 16$.

$$\lambda_{r_0}(\theta_{i,M}, \theta_{i,NM}) = \mathbb{I}(|x_{i,M} - x_{i,NM}| = 1)\lambda_{r_0,1} + \mathbb{I}(|x_{i,M} - x_{i,NM}| = 2)\lambda_{r_0,2}, \quad (15)$$

where $\theta_{i,B} = (\theta_{i,M} + \theta_{i,NM})/2$ and $|x_{i,M} - x_{i,NM}|$ is measured in absolute units of compensation here, not inches. From the location shift pattern in Figure 8, and the fact that

$$\delta_{r_0}(\theta_{i,M}, \theta_{i,NM}) = \delta_{r_0}(\theta_{i,M} + 2\pi, \theta_{i,NM} + 2\pi),$$

we believe location shifts should be modeled using a harmonic function, as in (14). Parameters $\lambda_{r_0,1}$ and $\lambda_{r_0,2}$ correspond to absolute differences of 1 and 2 assigned compensations.

This model provides a better fit. Comparing Figure 7(c), which displays posterior predictions from the refined model (based on one chain of posterior draws using a standard random walk Metropolis algorithm), with the previous model's predictions in Figure 7(a), we immediately see that the refined model better captures the posterior mean trend and interference. Similar improvements in prediction exist in the other sections and cylinders.

3.6 Summary of the Experimental Design and Analysis

Three key ingredients relating to the data, model, and experimental design have made our series of analyses possible. First is the availability of benchmark data, e.g., every unit on the cylinder receiving zero compensation. Second is the potential outcomes model (9) for the effect of compensation, defined in terms of a fixed number of parameters that do not depend on the vector of treatment assignments $\mathbf{x}$. These two ingredients enable one to form the posterior predictive distribution of potential outcomes under the assumption of negligible interference. The final ingredient is the distinction between units of analysis and units of interpretation, made explicit in our choice of design, which provides the means to assess and model interference in the experiment. Comparing observed outcomes from the experiment to posterior predictions allows one to infer the structure of interference, which can ultimately be validated by further experimentation. These considerations suggest that our approach and methodology can be generalized and applied to other experimental situations in the physical sciences having units residing on connected surfaces. In general, when experimenting with units on a connected surface, a principled and step-by-step analysis using the three ingredients above can ultimately shed more light on the question of interest.

4. Conclusion: Interference Inhibits Improvements

To construct 3D printed products satisfying manufacturing demands on dimensional accuracy, it is important to consider the problem of interference and address it in a principled manner. Huang et al. (2012) recognized that continuous compensation plans implemented on printers with a sufficiently fine tolerance can effectively control a product's printed dimensions, without inducing additional complications through interference. Their models for product deformation motivated our experimental design that introduces interference through the application of a discretized compensation plan to the boundary of a cylinder. Combining this experiment's data with inferences based on data for which every unit received no compensation led to an assessment of the overall structure of interference in terms of how treatments across units effectively differed from that physically assigned.

It is important to note that the refined interference model's location and scale parameters (14), (15) are also a function of the compensation plan. For example, reflecting the assigned compensations across the y axis would accordingly change the location shifts. The implication of this and all our previous observations for manufacturing is that *severely discretized compensation plans inhibit improved dimensional accuracy control for 3D printed products, and so should not be used for manufacturing*.

References

Box, G., A. Luceño, and M. Paniagua-Quiñones (2009). *Statistical Control by Monitoring and Adjustment* (2 ed.). Wiley.

Campbell, T., C. Williams, O. Ivanova, and B. Garrett (2011). Could 3d printing change the world? technologies, potential, and implications of additive manufacturing.

Colosimo, B. M., Q. Semeraro, and M. Pacella (2008). Statistical process control for geometric specifications: on the monitoring of roundness profiles. *Journal of Quality Technology 40*(1), 1–18.

Cox, D. and C. Donnelly (2011). *Principles of Applied Statistics* (1 ed.). Cambridge University Press.

Duane, S., A. Kennedy, B. J. Pendleton, and D. Roweth (1987, September). Hybrid monte carlo. *Physics Letters B 195*(2), 216–222.

Gibson, I., D. Rosen, and B. Stucker (2009). *Additive manufacturing technologies: rapid prototyping to direct digital manufacturing*. Springer Verlag.

Hilton, P. and P. Jacobs (2000). *Rapid tooling: technologies and industrial applications*. CRC.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association 81*(396), 945–960.

Huang, Q., J. Zhang, A. Sabbaghi, and T. Dasgupta (2012). Optimal offline compensation of shape shrinkage for 3D printing processes. Technical report, submitted to IIE Transactions 12-21-2012.

Melchels, F., J. Feijen, and D. Grijpma (2010). A review on stereolithography and its applications in biomedical engineering. *Biomaterials 31*(24), 6121–6130.

Rosenbaum, P. (2007). Interference between units in randomized experiments. *Journal of the American Statistical Association 102*(477), 192–200.

Rubin, D. (1980). Comment on "Randomization analysis of experimental data: the Fisher randomization test," by D. Basu. *Journal of the American Statistical Association 75*, 591–593.

Sobel, M. (2006). What do randomized studies of housing mobility demonstrate?: causal inference in the face of interference. *Journal of the American Statistical Association 101*(476), 1396–1407.

Tong, K., S. Joshi, and E. Lehtihet (2008). Error compensation for fused deposition modeling (fdm) machine by correcting slice files. *Rapid Prototyping Journal 14*(1), 4–14.

Tong, K., E. Lehtihet, and S. Joshi (2003). Parametric error modeling and software error compensation for rapid prototyping. *Rapid Prototyping Journal 9*(5), 301–313.

Wang, W., C. Cheah, J. Fuh, and L. Lu (1996). Influence of process parameters on stereolithography part shrinkage. *Materials & Design 17*(4), 205–213.