ASME Journal of Manufacturing Science and Engineering
Online journal at:
https://asmedigitalcollection.asme.org/manufacturingscience

Aaron Cornelius
Department of Mechanical, Aerospace, and Biomedical Engineering,
University of Tennessee,
Nathan W. Dougherty Engineering Building,
1512 Middle Dr,
Knoxville, TN 37916
e-mail: acornel5@vols.utk.edu

Jaydeep Karandikar[1]
Mem. ASME
Manufacturing Science Division,
Oak Ridge National Laboratory,
2350 Cherahala Boulevard,
Knoxville, TN 37932
e-mail: karandikarjm@ornl.gov

Chris Tyler
Manufacturing Science Division,
Oak Ridge National Laboratory,
2350 Cherahala Boulevard,
Knoxville, TN 37932
e-mail: tylerct@ornl.gov

Tony Schmitz
Fellow ASME
Department of Mechanical, Aerospace, and Biomedical Engineering,
University of Tennessee,
Nathan W. Dougherty Engineering Building,
1512 Middle Dr,
Knoxville, TN 37916;
Manufacturing Science Division,
Oak Ridge National Laboratory,
2350 Cherahala Boulevard,
Knoxville, TN 37932
e-mail: tony.schmitz@utk.edu

Process Damping Identification Using Bayesian Learning and Time Domain Simulation

Process damping can provide improved machining productivity by increasing the stability limit at low spindle speeds. While the phenomenon is well known, experimental identification of process damping model parameters can limit pre-process parameter selection that leverages the potential increases in material removal rates. This paper proposes a physics-informed Bayesian method that can identify the cutting force and process damping model coefficients from a limited set of test cuts without requiring direct measurements of cutting force or vibration. The method uses time-domain simulation to incorporate process damping and provide a basis for test selection. New strategies for efficient sampling and dimensionality reduction are applied to lower computation time and minimize the effect of model error. The proposed method is demonstrated, and the identified cutting and damping force coefficients are compared to values obtained using machining tests and least-squares fitting. [DOI: 10.1115/1.4064832]

Keywords: milling, dynamics, process damping, Bayesian learning, machine learning, machine tool dynamics, machining processes

Introduction

Process damping is a phenomenon where a machining process dissipates vibrational energy through the impact of the tool and workpiece [1]. This energy dissipation can increase the stability limit and improve machining productivity, particularly at low spindle speeds. Various models have been proposed to predict the effect of process damping on the machining process [2–7]. These methods apply process damping coefficients which depend on the workpiece material and the tool's cutting-edge geometry and enable the selection of increased stable depths of cut at low cutting speeds.

Several approaches for identifying the process damping coefficients have been proposed. Altintas et al. directly measured the coefficients using a piezoelectric fast tool servo to oscillate the tool relative to the workpiece at selected oscillating frequencies and amplitudes, and measured the cutting force components with a dynamometer [2]. Tyler et al. determined the process damping coefficient by performing least-squares fitting to a set of stable and unstable cutting tests at low spindle speeds [3]. Budak et al. determined the process damping model by recording the tool's vibration with a laser vibrometer and performing energy analysis [4,5]. Sellmeier et al. performed planning cuts by fixing the spindle and oscillating the crossfeed axis. They measured the resulting indentation force to establish the model [6]. As an alternative, Tuysuz et al. described an algorithm for analytically predicting the process damping based on the cutting-edge geometry [7].

Bayesian inference has recently been proposed as a method to predict milling stability and the underlying system parameters, including the frequency response function (FRF) and cutting force model. These machine learning approaches begin with a probabilistic guess about the shape of the stability map and iteratively update it based on the results of cutting tests. This method was

[1]Corresponding author.

Manuscript received July 7, 2023; final manuscript received February 6, 2024; published online April 24, 2024. Assoc. Editor: Y.B. Guo.

first implemented by Karandikar et al. using a non-physics-informed method which updated the probabilistic stability map based on proximity to stable/unstable test results [8,9]. While this approach can find the cutting parameters with the highest material removal rate (MRR), it does not provide new information about the underlying system parameters. Several authors have recently proposed physics-informed Bayesian learning models which update not only the probabilistic stability map but also the underlying uncertainties about the FRF and cutting force models [10–14]. These approaches can converge in a smaller number of cutting tests compared to the non-physics-informed method [12]. However, the methods presented thus far have three significant limitations.

First, the previous methods have used approximate stability models which linearize the process dynamics. References [11–13] utilized the zero-order approximation (ZOA) frequency domain stability model [15]. This method is computationally efficient, but offers reduced accuracy for low radial immersions and does not include process damping or variable flute spacing. Chen et al. employed a semi-discretization method, which utilizes Floquet theory to determine stability and can incorporate additional model complexities [10]. Because semi-discretization is more computationally expensive than the frequency domain stability solution, Chen trained a surrogate machine learning model to approximate the semi-discretization solution and reduce the computation time.

Second, previous work has utilized Markov Chain Monte Carlo (MCMC) sampling algorithms to draw samples from the posterior distribution. These methods create a chain of correlated samples by randomly walking around the posterior distribution. MCMC methods are well-suited for exploring complex high-dimensional probability distributions since they naturally concentrate sampling in high-probability regions of the target distributions [16]. However, they have several disadvantages. MCMC implementation can be highly complex, and it is difficult to design optimal samplers. As of this writing, every physics-based Bayesian stability paper has utilized a different MCMC sampler [9–12], and there has been no direct comparison to determine which is the most efficient. The most advanced MCMC samplers exploit knowledge about the target distribution to achieve efficient sampling, but these methods are generally difficult to apply to stability algorithms. For example, the Hamiltonian Monte Carlo algorithm utilizes the gradient of the target distribution to improve efficiency, but there is no closed-form solution for its application to stability maps [16]. MCMC methods are also difficult to parallelize efficiently. While there are a variety of methods for parallelizing the samplers [17], these methods still impose parallelization overhead which makes it difficult to scale the MCMC approach to utilize high-performance computing resources.

Third, no prior work has evaluated the ability of the Bayesian methods to converge to the true cutting force model and system FRF solely from stability test data. Instead, these efforts have focused on stability map convergence [10,11] and the identification of high MRR cutting parameters [12,13]. Ahmadi et al. compared the posterior distribution to tap testing data during their study on Bayesian identification of turning dynamics [14]; however, their updating method required tool vibration measurement during machining, which is not feasible for production environments. Since stability maps may not be one-to-one (i.e., multiple sets of input parameters may produce similar stability maps), stability map convergence does not guarantee that the cutting force model and FRF have converged to their true values. This limits the generalizability of the posterior solution. For example, the cutting force model learned in one setup could ideally be used to generate accurate stability maps for other setups, such as changing the tool extension length or increasing the radial depth of the cut. If an incorrect force model or FRF is identified, this may not be true.

This work presents a Bayesian updating algorithm which can identify both the process damping and cutting force model from a set of cutting experiments. The algorithm is highly flexible and can be adapted to different types of cuts and process damping models. The work has four novel contributions.

- The Bayesian stability updating method is extended to incorporate time-domain simulation (TDS). TDS utilizes numerical integration to solve the instantaneous cutting forces and vibrations in discrete time-steps [1]. TDS is well-suited to include complex effects such as process damping [18], accurately incorporates process nonlinearities, and can provide additional information about the machining process. However, TDS is computationally expensive compared to the ZOA, requiring computation times that are at least 1000× higher.
- A grid-based sampling method and updating strategy is proposed instead of the Markov Chain Monte Carlo (MCMC) methods used in prior work. Grid-based sampling is computationally less efficient since the sample is not concentrated in high-probability regions, but is trivially parallelizable and enables the stability maps to be precalculated so that the Bayesian updating can be performed quickly. As such, it is well-suited for iterative Bayesian updating and can easily leverage high-performance parallel computing resources.
- A divided training method is proposed which uses subsets of the data to learn different system parameters. This approach reduces the effect of model errors by training each variable using the most reliable data and can provide dimensionality reduction for decreased computing time.
- A comparison between the posterior for the system parameters and measured values is presented to evaluate how well the Bayesian approach was able to converge to the underlying values.

The paper is organized as follows. The Proposed Method section presents the TDS algorithm, cutting force and process damping models, and the Bayesian updating method. The Results section demonstrates the approach using a set of milling tests. The Discussion section analyzes the results and discusses the advantages and limitations of the new method. Finally, the Conclusion section summarizes the study.

Proposed Method

This section describes the proposed method. First, the TDS and stability classification methods are described. Next, the proposed iterative Bayesian updating scheme is detailed. The flowchart shown in Fig. 1 provides an overview of the process flow.

Time-Domain Simulation. TDS combines numerical integration with a mechanistic force model to estimate the time-dependent tool (and workpiece) vibration and cutting force components. After the simulation is complete, the vibration and force signals can be analyzed to classify the cut as stable or unstable. TDS methods have been widely studied and are easily extended to include complex effects such as process damping [18], variable flute spacing tools [19], and low radial immersion milling [20].

The TDS algorithm is outlined in Fig. 2. As TDS methods have been well-documented in the literature, this paper does not present detailed step-by-step instructions (for further details and example code, see [1]). Instead, this section examines the model choices necessary to use TDS for Bayesian learning. Specifically, the cutting force and process damping models are presented, followed by the method for classifying simulations as stable or unstable and calculating the limiting axial depth.

The cutting force required to shear a chip from the workpiece can be projected into the tangential and normal directions on each flute. These force components F_t, F_n (illustrated in Fig. 3) are defined in Eqs. (1) and (2), where K_s is the specific cutting force coefficient (N/m^2), which describes how much force is required to cut away a chip of unit area, β is the cutting force angle (rad), which describes how that cutting force is oriented relative to the surface normal direction, h is the chip thickness (m), and db is the chip width (m) for a differential element of the flute's cutting edge. For comparison, the alternate cutting force component representations

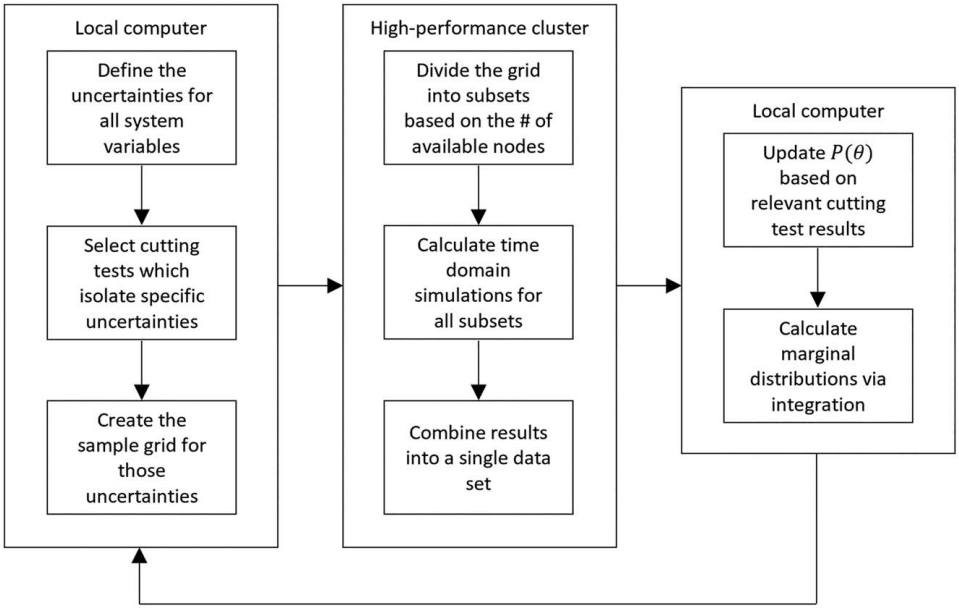

Fig. 1 Flowchart showing the proposed method for Bayesian updating and variable separation

using the tangential, k_t, and normal, k_n, direction coefficients are also shown. Edge coefficients are not included in this study since they have a secondary effect on process stability (i.e., they are modeled as proportional to chip width only and not the time-dependent chip thickness). Note that the cutting force is set equal to zero when $h < 0$ since the tool is not engaged with the workpiece. TDS is well-suited for modeling this nonlinearity, which cannot be captured in linear approximations such as the frequency domain stability algorithms.

$$F_t = \begin{cases} K_s \sin(\beta) h\, db = k_t h\, db & h \geq 0 \\ 0 & h < 0 \end{cases} \tag{1}$$

$$F_n = \begin{cases} K_s \cos(\beta) h\, db = k_n h\, db & h \geq 0 \\ 0 & h < 0 \end{cases} \tag{2}$$

The process damping force, on the other hand, is a dissipative force caused by interference between the machined workpiece surface and the cutting tool's clearance face. The process damping effect is modeled here using the nonlinear wavelength-dependent model introduced in Ref. [2]. Let $V = \frac{\Omega d \cdot 2\pi}{60}$ (m/s) be the tangential velocity of the cutting tool's surface due to spindle rotation (where Ω is spindle speed in rpm and d is the tool diameter in m), $\dot{x}$ and $\dot{y}$ be the relative vibrational velocities (m/s) between the tool and workpiece in the X and Y directions, and c_n (N/m) be a process damping coefficient. The normal direction process damping force F_D(N) on a single flute element can then be calculated using Eq. (3), where ϕ is the tool rotation angle (rad). Similar to the cutting force model, the process damping force is zero when the tool is not engaged with the workpiece. Note that there have been several other proposed damping models and any of them can be implemented in the proposed Bayesian framework (e.g., the indentation volume model [18]).

$$F_D = \begin{cases} -c_n\, db \dfrac{\dot{x}\sin(\phi) - \dot{y}\cos(\phi)}{V} & h \geq 0 \\ 0 & h < 0 \end{cases} \tag{3}$$

Stability classification is completed using once-per-flute sampling from the simulated displacement signal $x(t)$ [21]. Let x_i be the ith sample according to Eq. (4), where τ is the tool rotation (rad), N_t is the number of flutes on the cutting tool, and $i = 0, 1, 2, \ldots$ The N samples are then used to calculate the stability metric m using Eqs. (4) and (5). In an ideally stable cut, the displacement repeats with each new flute passage and $m = 0$. This is repeated for all possible different cutter rotations lags $\tau = \{0{:}2\pi/N_t\}$ since the peak m value may not occur at the same rotation for every set of cutting parameters (see Fig. 4). If any of the calculated values exceed a preset stability cutoff of $m = 1\ \mu$m, then the cut is classified as unstable. Otherwise, it is classified as stable. In summary, the approach confirms that the milling operation exhibits forced vibration, which is labeled as a stable milling condition. In forced vibration, the force and corresponding vibration repeat with each tooth passage, so sampling the displacement once-per-flute and calculating the difference between adjacent samples confirms stable milling when the value is (ideally) zero. When self-excited vibration occurs during regenerative chatter, a new frequency (the chatter frequency) is introduced and the samples do not repeat with each tooth passage. The metric defined in Eq. (5) is sensitive to this variation from one tooth to the next. The cutoff of $m = 1\ \mu$m is selected to avoid numerical noise, while establishing a quantitative value for separating stable and unstable (chatter) conditions automatically.

$$x_i = x\left(\frac{60\ i}{\Omega N_t} + \frac{60\ \tau}{2\pi\Omega}\right), \quad i = 0, 1, 2, \ldots \tag{4}$$

$$m = \frac{1}{N}\sum_{i=1}^{N} |x_i - x_{i-1}| \tag{5}$$

Two different stability functions are necessary for the Bayesian updating. The first function $S(\boldsymbol{\theta}, \Omega, b)$ tests stability at a single spindle speed and axial depth of cut b for a given set of system parameters $\boldsymbol{\theta}$, returning 1 if the simulation is stable and 0 if the simulation is unstable. This is determined by checking the result of a single time-domain simulation. The second function $b_{lim}(\boldsymbol{\theta}, \Omega)$ calculates the limiting axial depth for a given spindle speed and set of system parameters. Because TDS does not calculate an explicit limiting axial depth, as with the frequency domain algorithm, b_{lim} is defined as the midpoint between the two simulations where the cut transitions from stable to unstable (Fig. 5(*a*)). There are two cases to consider. First, if there were no unstable simulations at a given spindle speed, then $b_{lim}(\theta, \Omega)$ is undefined. As such, it is important to run simulations at greater axial depths than test cuts will be performed. Second, due to period-n

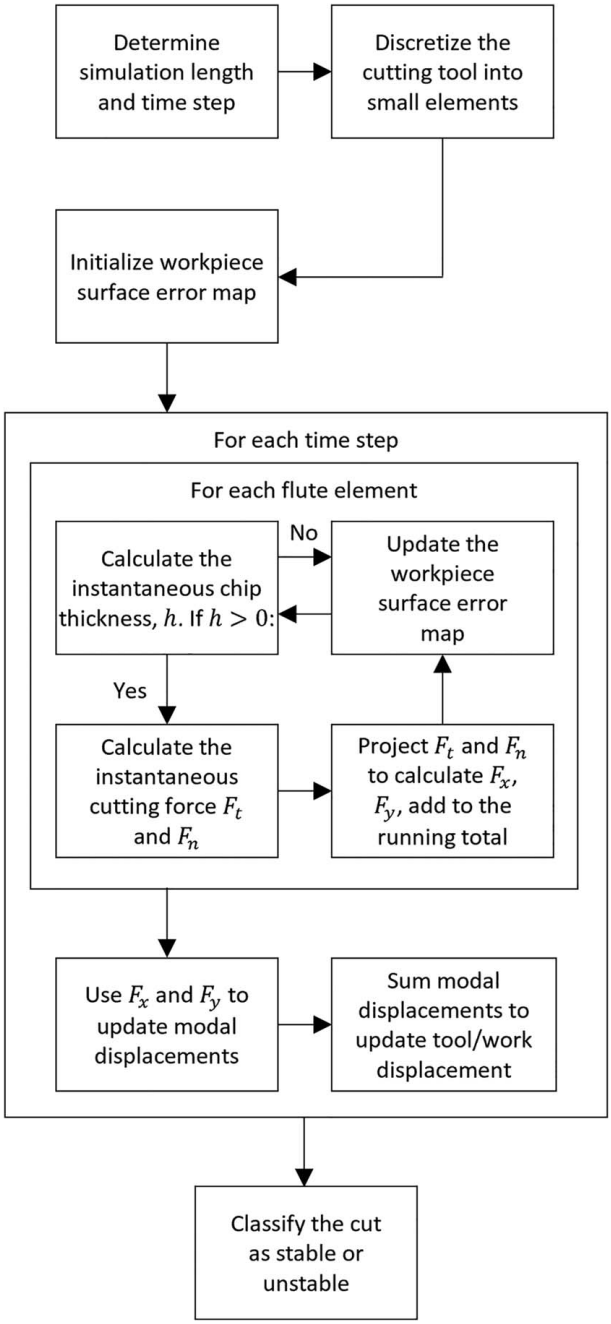

Fig. 2 Flowchart for time-domain simulation. This simulation process is repeated for each set of cutting parameters to be tested.

bifurcations that can occur in milling [22], there can sometimes be multiple stable/unstable transitions at a given spindle speed which result in multiple b_{lim} values, illustrated in Fig. 5(*b*). These bifurcations are susceptible to small changes in system parameters and can occur for individual sampled stability maps, even if the experimental test cuts do not demonstrate bifurcations. Therefore, all transitions must be returned by the b_{lim} function so the Bayesian updating can incorporate this behavior.

Establishing and Sampling the Prior. Let $\boldsymbol{\theta}$ be a variable containing all relevant system parameters necessary to define the workpiece/machine/spindle/tool holder/cutting tool system, including the tool tip FRF, process damping coefficient, and cutting force model parameters. $\boldsymbol{\theta}$ can take various forms depending on what

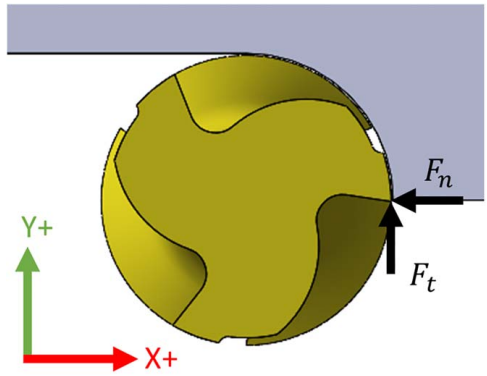

Fig. 3 Normal and tangential force directions on the cutting tool

variables are known and what physics models are used. TDS requires a set of modal parameters, so to use TDS $\boldsymbol{\theta}$ should be a vector containing a cutting force model and a set of modal parameters, e.g., $\boldsymbol{\theta} = \{K_s, \beta, f_n, k, \zeta\}$. Alternatively, the FRF could be calculated using receptance coupling substructure analysis [23], where $\boldsymbol{\theta}$ would contain information about the coupling stiffness and damping, tool and holder models, and spindle/machine FRFs.

The prior distribution $P_{prior}(\boldsymbol{\theta})$ is an initial probabilistic guess about those system parameters based on the information that is already known, describing how likely it is that a specific set of values for $\boldsymbol{\theta}$ represents the true system. The prior can be established using a variety of techniques. Here, the prior is constructed using the author experience with similar work materials and setups.

Previous work has drawn samples from the prior using Monte Carlo sampling and then used MCMC sampling to draw samples from the posterior [10–14]. Both these methods create samples with equal weight and the density of the target distribution is denoted by the concentration of samples. However, these samples are difficult to update to represent the posterior, requiring the MCMC sampler to be run again. Since this requires calculating new samples and, thus, new stability maps, it is not suitable for iterative Bayesian updating with time-consuming stability algorithms such as TDS.

Instead, samples are selected from a fixed grid pattern. Each sample contains both the sampled values $\boldsymbol{\theta}$ and the precalculated TDS stability maps. This method may require computing more samples than the MCMC approach since samples are not concentrated in regions of high probability, but still has several significant advantages versus MCMC. First, it is straightforward to parallelize stability map calculations across multiple cores or computers since each sample is calculated independently. Second, the Bayesian updating is performed on the probability associated with each sample, rather than by drawing new samples and calculating a new stability map. As such, the grid integration method is well-suited for iterative Bayesian updating with computationally expensive TDS. Specifically, the TDS results can be precalculated using a high-performance multi-node computing cluster, and the Bayesian updating is performed by looping through all the samples to update the probabilities and calculated the updated distributions. Calculation times for this process are presented in the Results section.

After all the stability maps have been calculated, they are used to generate a probabilistic stability map. Previous work using MCMC sampling generated samples with equal weight and could calculate the probability of stability as the percentage of samples which predict that a given set of cutting parameters would be stable using the stability prediction function (i.e., for N samples, the probability of stability was $P_S(\Omega, b) = \frac{1}{N}\sum_{i=1}^{N} S(\boldsymbol{\theta}_i, \Omega, b)$). In this study, each sample has a different weight based on its probability, and the probability of stability for a given point $P_S(\Omega, b)$ is calculated by integrating $S(\boldsymbol{\theta}, \Omega, b)$ using Eq. (6). An example probabilistic stability map is displayed in Fig. 6. The grayscale indicates the

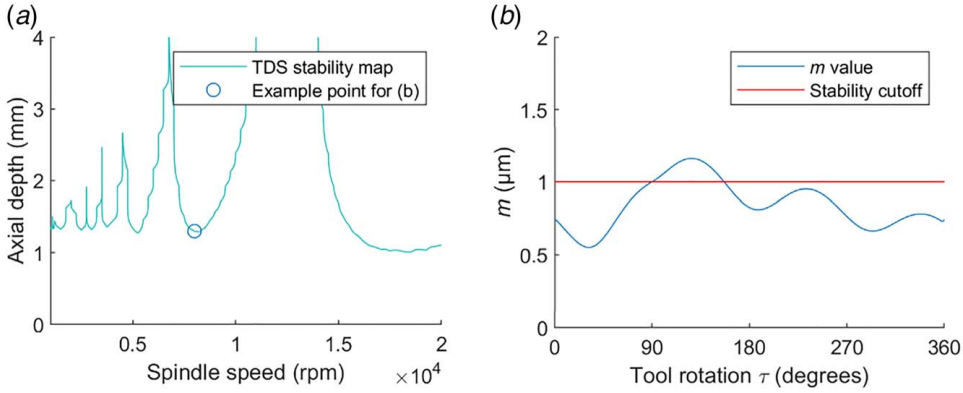

Fig. 4 Stability classification example: (*a*) stability map calculated using time-domain simulation, and (*b*) once-per-flute standard deviations for a TDS at 8000 rpm, 1.3 mm axial depth. This example exceeds the m = 1 μm threshold and is classified as unstable.

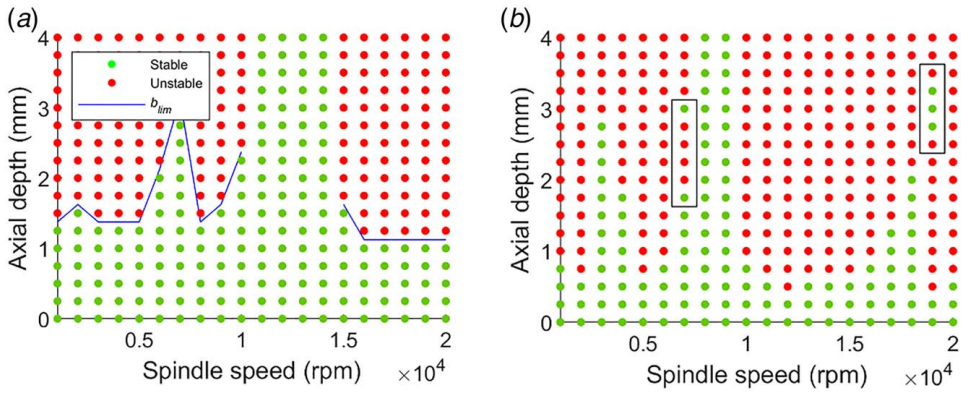

Fig. 5 (*a*) b_{lim} is determined from a set of time-domain simulations. b_{lim} is undefined between 11,000 and 14,000 rpm since there were no unstable simulations at those spindle speeds. (*b*) Multiple values for b_{lim} due to system nonlinearities, with specific nonlinear regions boxed.

probability of stability, where brighter regions are more likely to be stable and darker regions are more likely to chatter.

$$P_S(\Omega, b) = \int_{\boldsymbol{\theta}} P_{prior}(\boldsymbol{\theta})\, S(\boldsymbol{\theta}, \Omega, b)\, d\boldsymbol{\theta} \tag{6}$$

Bayesian Updating. Once the tests have been selected, performed, and classified, their results are used to define the posterior distribution. Let α be a variable containing the results of the ith test, including the tested spindle speed α_Ω and axial depth α_b, whether the cut was stable or unstable, and any other observed information about the cutting test (e.g., chatter frequency). The posterior distribution $P(\boldsymbol{\theta}|\alpha)$ is calculated using Bayes' theorem, as shown in Eq. (7). In this equation, $P(\alpha|\boldsymbol{\theta})$ is the likelihood function which states how likely the test result α is for a given set of parameters $\boldsymbol{\theta}$ and c is a normalizing constant that ensures $\int_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\alpha)\, d\boldsymbol{\theta} = 1$. This constant does not have an analytical solution for this case. The non-normalized probability distribution $p(\boldsymbol{\theta}|\alpha)$ is therefore used (Eq. (8)) and normalized numerically during the updating process. If more than one cutting test has been run, then the posterior for the first test serves as the prior for the next, giving the general form shown in Eq. (9).

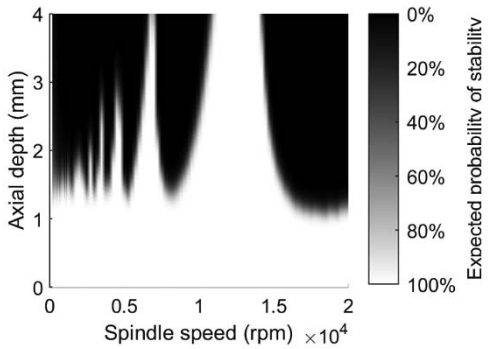

Fig. 6 Probabilistic stability map generated through grid sampling

$$P(\boldsymbol{\theta}|\alpha) = \frac{P_{prior}(\boldsymbol{\theta})\, P(\alpha|\boldsymbol{\theta})}{c} \tag{7}$$

$$p(\boldsymbol{\theta}|\alpha) = P(\boldsymbol{\theta}|\alpha) \cdot c = P_{prior}(\boldsymbol{\theta})\, P(\alpha|\theta) \tag{8}$$

$$p(\boldsymbol{\theta}|\alpha_{1:n}) = P_{prior}(\boldsymbol{\theta}) \prod_{i=1}^{n} P(\alpha_i|\boldsymbol{\theta}) \tag{9}$$

The likelihood function here is the same as used in Refs. [12,13], shown in Eqs. (10) and (11), where Δ_b is the error in the stability limit and σ_b is a stability error hyperparameter (larger values for σ_b allow for larger errors in stability limit). If $\boldsymbol{\theta}$ correctly predicts the cut stability, then $\Delta_b = 0$ and $P(\alpha|\boldsymbol{\theta}) = 1$. If the $\boldsymbol{\theta}$ prediction is incorrect, then $0 \le P(\alpha|\boldsymbol{\theta}) < 1$, where larger errors give smaller likelihoods according to a Gaussian dropoff (Fig. 6(*b*)). For the specific case where there is an unstable cut and $b_{lim}(\boldsymbol{\theta}, \alpha_\Omega)$ is undefined

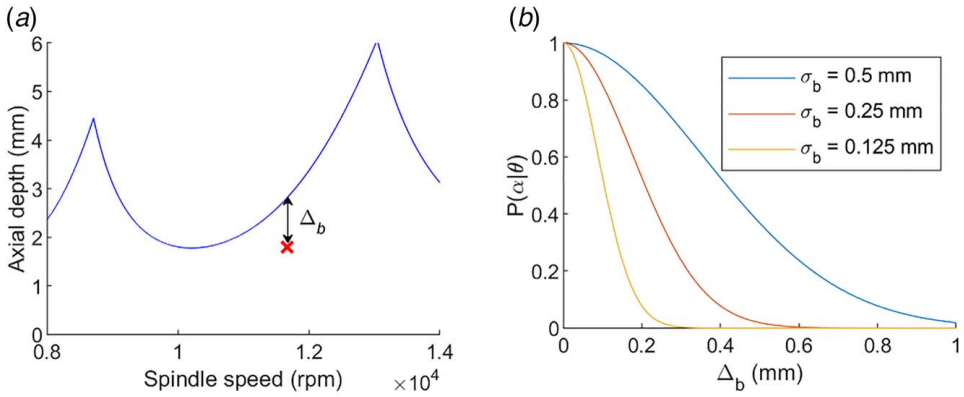

Fig. 7 (*a*) Demonstration of how Δ_b is calculated for an incorrectly predicted stability test. (*b*) Likelihood functions using the Gaussian drop-offs for different stability error hyperparameters.

since no simulations at that spindle speed were unstable, then $P(\alpha|\boldsymbol{\theta})$ is assumed to be 0. This is not an issue with stable cuts since the time-domain simulation at $b = 0$ will have zero cutting force and is always stable. If there are multiple b_{lim} values for the given spindle speed, then the b_{lim} value closest to α_b is selected. Figure 7 displays Δ_b and the Gaussian likelihood function for three representative σ_b values.

$$P(\alpha|\boldsymbol{\theta}) = \exp\left(-\left(\frac{\Delta_b}{\sigma_b}\right)^2\right) \tag{10}$$

$$\Delta_b = \begin{cases} 0 & S(\boldsymbol{\theta},\ \alpha_\Omega,\ \alpha_b) = \alpha_{Stable} \\ \min\left(|b_{lim}(\boldsymbol{\theta},\ \alpha_\Omega) - \alpha_b|\right) & S(\boldsymbol{\theta},\ \alpha_\Omega,\ \alpha_b) \neq \alpha_{Stable} \\ \infty & b_{lim}(\boldsymbol{\theta},\ \alpha_\Omega) \text{ is undefined} \end{cases} \tag{11}$$

The likelihood function can be extended to include additional terms, such as chatter frequency [12]. For this paper, chatter frequency is omitted since it is not significantly affected by the cutting force model or process damping, and thus does not significantly accelerate the learning process.

Once the posterior probability function has been defined, it is evaluated for each calculated sample. To maintain numerical accuracy, this should be done using the logarithm of the non-normalized probability $\log\left(p(\theta|\alpha_{1:n})\right)$ (see Eq. (12)). The logarithmic probabilities are offset to ensure that the highest probability samples are centered in the highest resolution area for floating point calculations, forming a temporary non-normalized distribution $\tilde{p}(\theta|\alpha_{1:n})$ (Eq. (13)). The normalizing constant c is calculated (Eq. (14)), and the final normalized distribution $P(\boldsymbol{\theta}|\alpha_{1:n})$ is calculated as follows (Eq. (15)):

$$\log\left(p(\boldsymbol{\theta}|\boldsymbol{\alpha}_{1:n})\right) = \log\left(P_{prior}(\boldsymbol{\theta})\right) + \sum_{i=0}^{n} -0.5\left(\frac{\Delta_{b\,i}}{\sigma_b}\right)^2 \tag{12}$$

$$\tilde{p}(\boldsymbol{\theta}|\alpha_{1:n}) = \exp\left(\log\left(p(\boldsymbol{\theta}|\alpha_{1:n})\right) - \max\left(\log\left(p(\boldsymbol{\theta}|\alpha_{1:n})\right)\right)\right) \tag{13}$$

$$c = \int_{\boldsymbol{\theta}} \exp(\log\left(\tilde{p}(\boldsymbol{\theta}|\boldsymbol{\alpha}_{1:n})\right))\, d\boldsymbol{\theta} \tag{14}$$

$$P(\boldsymbol{\theta}|\alpha_{1:n}) = \frac{1}{c}\ \tilde{p}(\boldsymbol{\theta}|\alpha_{1:n}) \tag{15}$$

The normalized distribution is used to calculate an updated probabilistic stability map using Eq. (6). Additionally, the marginal distribution for each subvariable is calculated and used to determine the posterior mean and standard deviation for each variable. Let θ_i be the ith element of the $\boldsymbol{\theta}$ vector. The marginal distribution $P(\boldsymbol{\theta}_i|\alpha_{1:n})$ can be calculated by integrating across all the other dimensions of $\boldsymbol{\theta}$ (Eq. (16)) and the mean and standard deviation calculated by integration (Eqs. (17) and (18)).

$$P(\theta_i|\alpha_{1:n}) = \int_{\boldsymbol{\theta}_j,\, j\neq i} P(\boldsymbol{\theta}|\alpha_{1:n})\, d\theta_j \tag{16}$$

$$\mu(\theta_i|\alpha_{1:n}) = \int_{\boldsymbol{\theta}_i} \theta_i \cdot P(\theta_i|\alpha_{1:n})\, d\theta_i \tag{17}$$

$$\sigma(P(\boldsymbol{\theta}_i|\alpha_{1:n})) = \sqrt{\int_{\boldsymbol{\theta}_i} (\boldsymbol{\theta}_i - \mu(\boldsymbol{\theta}_i|\alpha_{1:n}))^2 \cdot P(\boldsymbol{\theta}_i|\alpha_{1:n})\, d\boldsymbol{\theta}_i} \tag{18}$$

Separated Learning. Prior Bayesian stability models treated all test cuts identically and attempted to fit a global stability map to the data simultaneously. However, cutting tests in different areas of the stability map may provide information on different aspects of the system. For example, high spindle cutting tests provide very little information on process damping since the increased stability is generally only significant at low spindle speeds. This is naturally accounted for in the Bayesian updating model: if a variable has little effect on the stability map close to a cutting test, then that test will not significantly alter the posterior distribution. However, the inverse is not true: low spindle-speed cutting tests are affected by both the process damping and cutting force models. Specifically, modifying either can affect the local stability limit. Increasing process damping increases the stability limit, while decreasing cutting force model coefficients likewise increases the stability limit. This is problematic since process damping models can have local errors even if they are broadly accurate [3,18]. This error affects the posterior distributions for the stability map, process damping model, and cutting force model.

Rather than learning using the entire data/parameter domain simultaneously, Bayesian updating is separated here into discrete

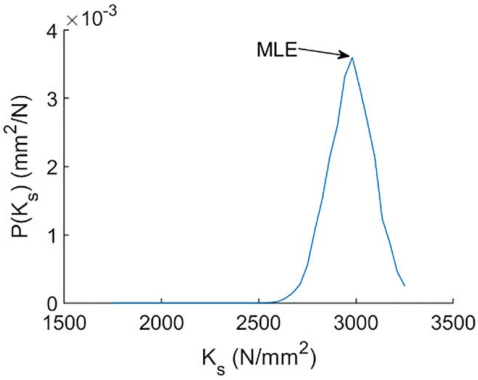

Fig. 8 Maximum-likelihood estimator

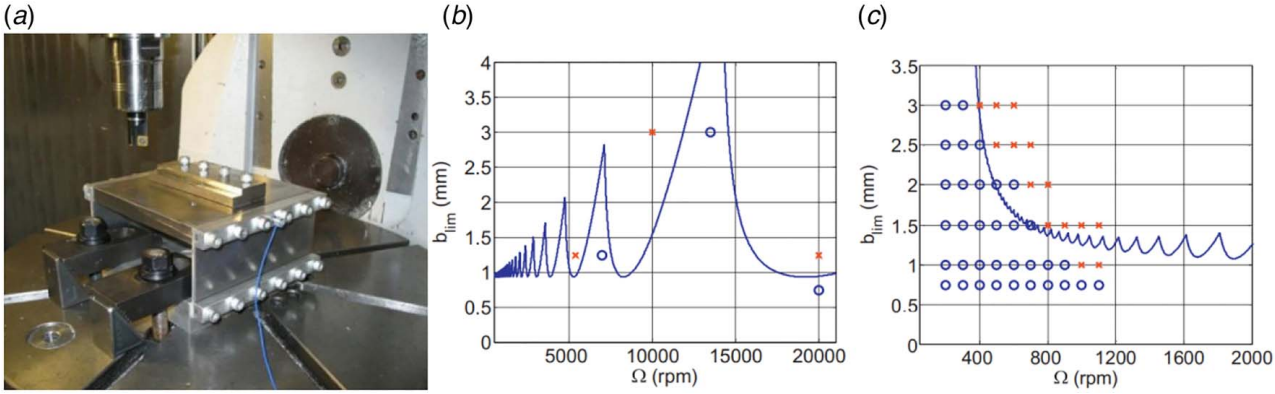

Fig. 9 Cutting test setup and results from Ref. [3]: (*a*) flexure test setup; (*b*) high spindle-speed cutting tests; (*c*) low spindle-speed cutting tests with fit process damping stability map

subdomains which each learn specific variables from a subset of the cutting tests. This provides the most reliable information for the selected subdomain. In this study, the cutting force model is learned using only cutting tests conducted at high spindle speeds, with process damping fixed at zero. The posterior maximum-likelihood estimator (MLE) is established for the cutting force model variables K_s and β, where the MLE is the value in the posterior with the highest probability using Eq. (19) with the reduced model $\theta = \{K_s,\ \beta\}$. Figure 8 shows a univariate example for the specific cutting force coefficient where the MLE for $K_s \approx 3000\,\mathrm{N/mm^2}$. These values for K_s and β are then used to calculate new stability maps which incorporate the process damping uncertainty for c_n. Finally, the c_n distribution is updated using the low-speed cutting tests.

$$MLE(\theta) = \arg\max_{\theta}\left(P(\theta)\right) \tag{19}$$

This method has two advantages compared to learning with all variables simultaneously. First, since each variable is learned using those cutting tests which provide the most accurate information, it reduces the impact of model errors and bias on the learning. Second, it enables the number of variables to be reduced significantly. Without this approach, the computational requirements scale as $N = m^n$, where N is the total number of required samples, m is the number of samples per variable, and n is the number of variables used. The demonstration shown in the Results section utilizes three variables (K_s, β, and c_n) and 40 samples per variable. When learning all variables simultaneously (hereafter referred to as "combined learning"), a total of $40^3 = 64{,}000$ samples are required. Using the separated learning approach, only $40^2 + 40 = 1640$ samples must be calculated.

Results

Setup Description. The cutting tests reported by Tyler and Schmitz are used to demonstrate the proposed approach [3]. They ran a series of cutting tests on a 1018 steel workpiece mounted on a single-degree-of-freedom flexure setup and used least-mean squares fitting to fit a frequency domain process damping model (Fig. 9). These cutting tests demonstrate process damping, showing an exponential increase in stability limit as the spindle speed decreases. The cutting tool was a 19.05 mm single insert endmill with zero helix angle (Cutting Tool Technologies model DRM-030), and all up-milling cuts were performed using 50% radial immersion and a feed per tooth of 0.05 mm. The flexure dynamics were measured by tap testing and are summarized in Table 1.

These results were used for this study for three reasons. First, the system is computationally efficient to model with TDS since a straight flute tool with one insert does not need to be discretized into multiple segments. Second, the flexure dynamics were measured by tap testing and are thus known with low uncertainty. While tool tip FRFs may shift at high spindle speeds due to changes in bearing dynamics [24], workpiece dynamics are not subject to this variation. Instead, the flexure dynamics will change as the material is removed from the workpiece, but this effect is easier to predict and quantify than the spindle dynamics shifts. The largest cut conducted in the reference paper had an axial depth of 3 mm, a radial depth of 9.525 mm, and a total cut length of approximately 125 mm, removing roughly 30 grams of material. This is much lower than the modal masses (Table 1) so the FRF does not significantly shift and can be treated as a constant to reduce the number of updating variables. Third, the paper measured

Table 1 Flexure dynamics determined by tap testing [3]

Direction	f_n (Hz)	k (N/μm)	ζ	m (kg)	c (N-s/m)
X (feed)	228	2.77	0.063	1.35	244
Y	1482	174	0.037	2.01	1380

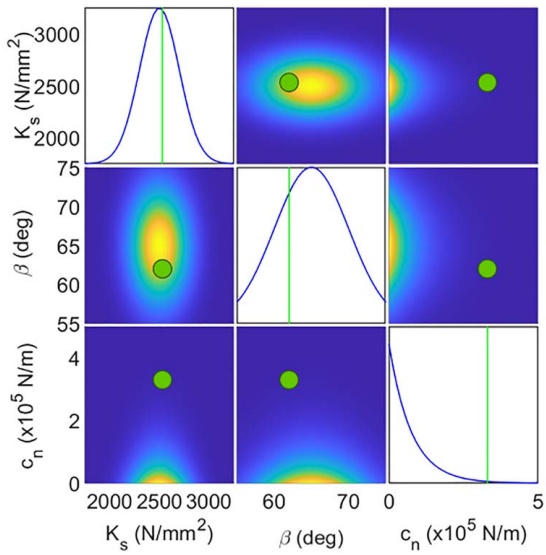

Fig. 10 Prior distributions for the three variables K_s, β, and c_n. The on-diagonal frames show the marginal distribution for each variable, with the vertical line representing the reference value from [3] and other curve representing the prior. The off-diagonal frames show the conditional distribution between each pair of variables, with brighter colors representing high-probability densities.

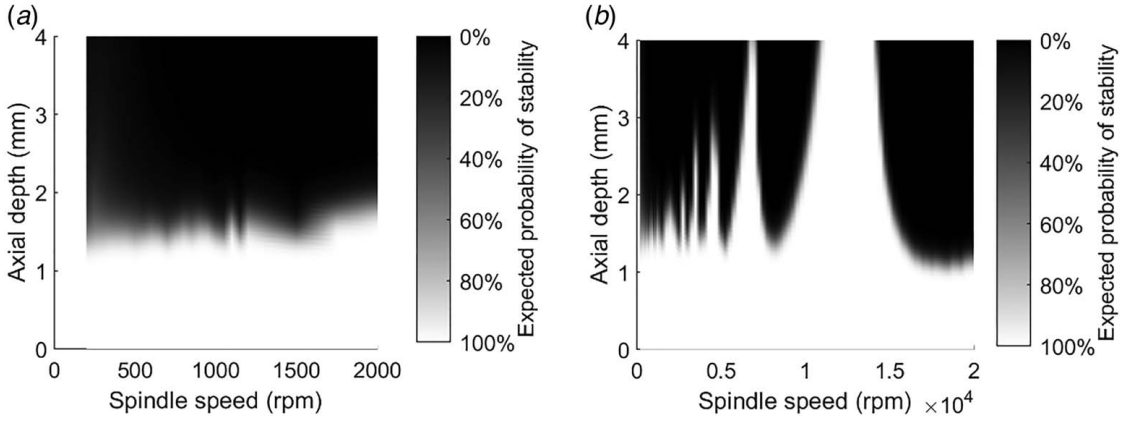

Fig. 11 Prior probabilistic stability map for the example setup from Ref. [3]: (*a*) low spindle-speed range, and (*b*) high spindle-speed stability map

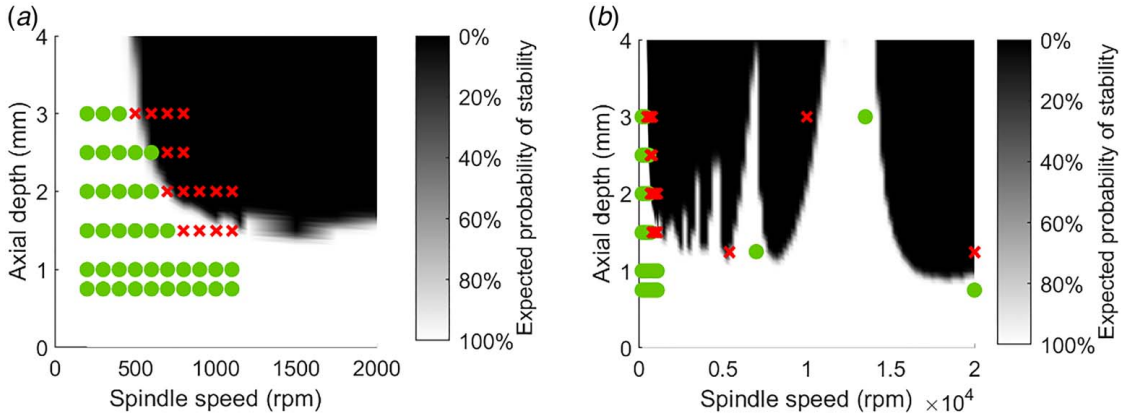

Fig. 12 Posterior stability maps for the combined updating: (*a*) low spindle-speed range, and (*b*) high spindle-speed stability map

and reported values for the cutting force and process damping models which can be used as a benchmark to compare with the posterior values. The cutting force model was quantified using a piezoelectric cutting force dynamometer to measure the force for known cutting conditions, while the process damping coefficient was calibrated by least-squares fitting to stability tests.

The prior was established based on the author's experience. K_s and β both used normal distributions, with $P_{prior}(K_s) = \mathcal{N}(\mu = 2500,\ \sigma = 200)\text{N/mm}^2$ and $P_{prior}(\beta) = \mathcal{N}(\mu = 65, \sigma = 5)\text{deg}$. The normal process damping term c_n had an exponential prior $P_{prior}(c_n) = \mathcal{E}(\lambda = 75000)\text{N/m}$ using Eq. (20).

$$E(\lambda) = \begin{cases} \lambda \exp(-\lambda x) & x \geq 0 \\ 0 & x < 0 \end{cases} \tag{20}$$

This is a long-tailed distribution which assigns a high probability that c_n is close to 0 (Fig. 10). The exponential distribution was chosen for two reasons. First, it is difficult to predict the process damping coefficient and, as such, it is not possible to select a meaningful mean value for a normal distribution. Second, process damping often does not have a significant effect on the stability map, either because the cutter has a high clearance angle or because milling operations are conducted at high spindle speeds, so it is appropriate to select a prior which assumes that process damping most likely does not have a significant effect on the stability map. The prior assigns only a ~1% probability that the process damping is as high as the best-fit value from [3]. The three uncertainties are combined to form a single multivariate distribution $\boldsymbol{\theta} = \{K_s, \beta, c_n\}$, where $P_{prior}(\boldsymbol{\theta}) = P_{prior}(K_s) \cdot P_{prior}(\beta) \cdot P_{prior}(c_n)$.

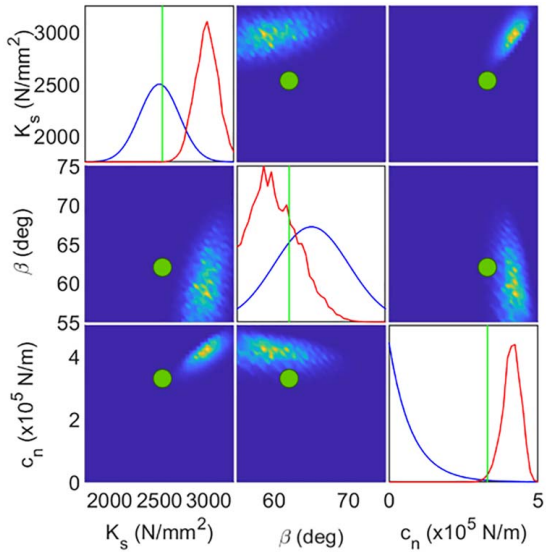

Fig. 13 Posterior distribution for K_s, β, and c_n after updating using the single-stage strategy. These follow the same notation as Fig. 10 with the new lines in the diagonal plots representing the posterior.

Combined Learning. The combined learning approach was used to learn both the process damping and cutting force models from all test points simultaneously. Stability maps were calculated

Table 2 Comparison of prior and posterior uncertainties for each marginal distribution

Variable	Prior	Combined learning posterior	Separated learning stage 1 posterior	Separated learning stage 2 posterior	Reference value
K_s (N/mm^2)	$\mathcal{N}(2500,\ 200)$	$\mathcal{N}(2970,\ 115)$	$\mathcal{N}(2671,\ 156)$	2671	2531.0
β (deg)	$\mathcal{N}(65,\ 5)$	$\mathcal{N}(60.2,\ 3.1)$	$\mathcal{N}(62.8,\ 3.9)$	62.8	62.0
c_r (N/m)	$\mathcal{E}(75{,}000)$	$\mathcal{N}(411000,\ 30{,}500)$	N/A	$\mathcal{N}(309000,\ 10{,}400)$	330000

Note: $\mathcal{N}(\mu, \sigma)$ is a fit normal distribution with mean μ and standard deviation σ, $\mathcal{E}(\lambda)$ is an exponential distribution with rate parameter λ.

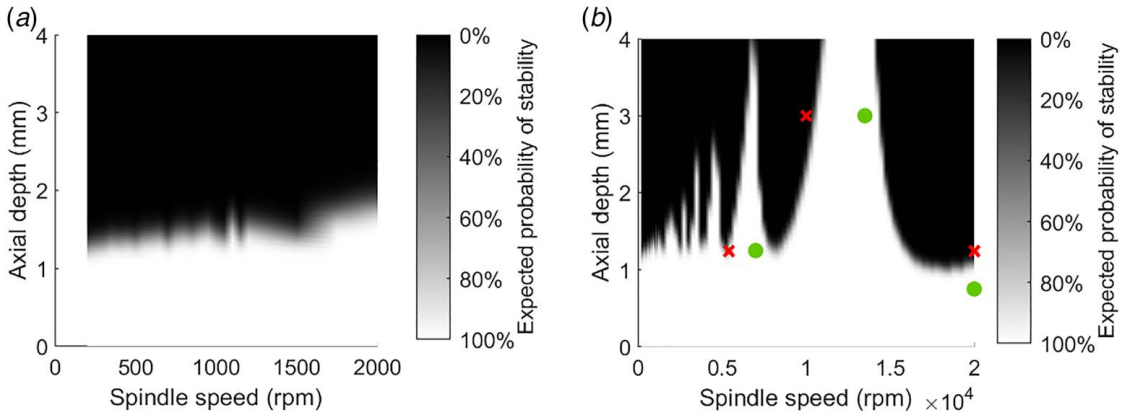

Fig. 14 Posterior distributions after learning the cutting force model using the high spindle-speed cutting tests: (*a*) closeup of the low spindle-speed range, and (*b*) high spindle-speed range

for each sample using the ISAAC-NG high-performance computing cluster at the University of Tennessee, Knoxville. Six nodes were applied with Intel Xeon Gold 6248R processors [25]. The calculations were parallelized using the Slurm Workload Manager's array calculation functions, dividing the job into 24 equally sized subjobs which each utilized twelve threads. Since each subjob utilized only a fraction of a node, this allowed the Slurm scheduler to backfill the jobs into idle time on the nodes, rather than having to wait for entire nodes to be available for use. After all the subjobs had been run, the results were concatenated to produce the final set of stability maps. Forty samples were used for each variable to give a total of 64,000 samples. Each stability map was calculated at 96 different spindle speeds (increments of 50 rpm from 200 rpm to 1200 rpm and increments of 250 rpm from 1500 rpm to 20,000 rpm) and 65 axial depths from zero to 4 mm, for a total of 4896 simulations per map. Each simulation was classified as stable or unstable using a stability metric cutoff of $m = 1\,\mu$m. These calculations took approximately 4 h. This resulted in the probabilistic stability map shown in Fig. 11.

The Bayesian updating was performed on a standard laptop and took approximately one minute (Intel i7-10750H processor, parallelized on six threads). Updating was performed using all 56 test cuts reported in Ref. [3], resulting in the posterior probabilistic stability map displayed in Fig. 12 and the parameter distributions shown in Fig. 13. Table 2 in the Comparison section compares the prior and posterior distributions. The posterior successfully identified the presence of process damping. However, it failed to converge to the experimental values for the underlying system variables due to model error on the low-speed cutting tests, overestimating both c_n and K_s as explained in the Discussion section.

Separated Learning. The separated learning strategy using the maximum-likelihood estimator was completed in two steps. Figure 14 shows the posterior stability maps for the first step, where process damping is assumed to be negligible so the cutting force model can be learned from the high spindle-speed cutting tests. Figure 15 shows the corresponding posterior distributions for K_s and β. This step required calculating 1600 stability maps and took less than ten minutes on the ISAAC computing cluster. It is feasible that this number of stability maps could be calculated using a high-power personal computer. Both K_s and β shift toward their reference values. Since only six test cuts were performed at high spindle speeds and because those test cuts were not selected to maximize information gain, these values still had significant uncertainty. K_s had an error of 5.5% between the posterior mean and experimental value, while β had an error of 1.3%.

Figures 16 and 17 show the results from the second updating step where the maximum-likelihood estimates for K_s and β are set as fixed and new cutting tests are used to update the prior for c_n.

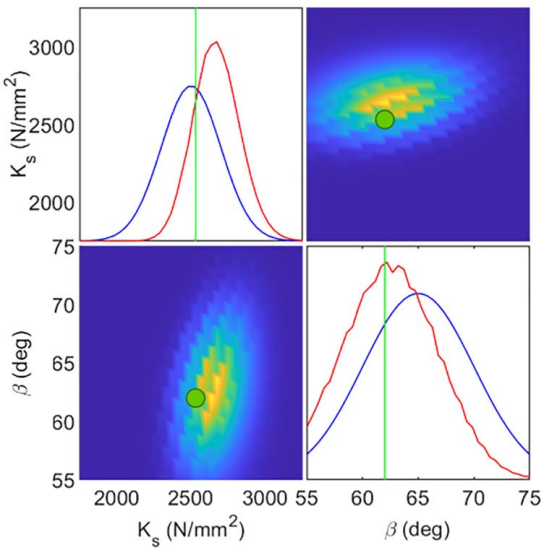

Fig. 15 Posterior distribution for K_s and β after the first separated learning updating step

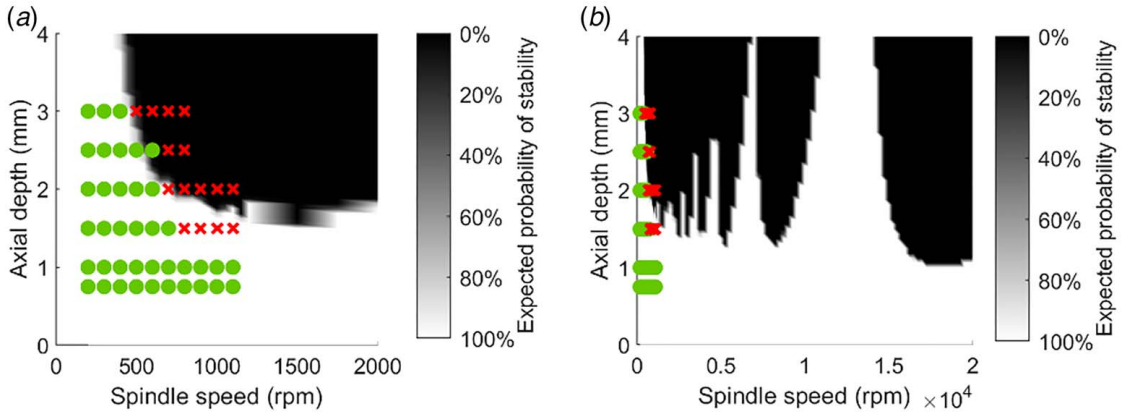

Fig. 16 Posterior distributions after learning K_s and β after the second separated learning step using the low spindle-speed cutting tests: (*a*) closeup of the low spindle-speed range, and (*b*) high spindle-speed range

This required calculating 40 stability maps, which was done on a personal computer (laptop). This posterior converged to the experimental value for c_n, with an error of 6.4% and low uncertainty.

Comparison. This section compares the combined and separated learning methods. Figure 18 shows the 50% probability boundaries for each case. The prior is generally a poor match, incorrectly predicting stability results at low speeds as well as the test cut at 5400 rpm. The posteriors for the combined and separate learning algorithms are similar with only minor differences in stability boundary. Both have successfully identified the presence of process damping at low spindle speeds. However, the posterior distributions in Table 2 clearly show that the separated learning method converged more accurately to the reference values from [3].

Comparison to MCMC Sampling. This section compares the computational efficiency and accuracy of the proposed numerical integration method to the parallel adaptive MCMC sampling method used in Ref. [12]. The new sampler is designed to approximate a target distribution in a small number of samples. This approach first samples the prior distribution using Monte Carlo sampling and then calculates the updated posterior distribution using MCMC. While the framework does allow for parallelization, it is not infinitely scalable and cannot practically be used with multimode computing. As such, the MCMC sampling was evaluated using a single node on the ISAAC cluster with 48 cores. Performing the initial Monte Carlo sampling with 4000 samples took 3 h and 20 min. Based on this, calculating the posterior using the MCMC sampler with a typical acceptance rate of 30–40% would take between 8 h and 11 h.

Due to the computational expense, the accuracy is evaluated using two trivariate distributions which resemble the prior and posterior distributions obtained from the combined learning updating and can be evaluated without the cost of calculating the stability map. For ease of comparison, both MCMC and the integration method use a total of 64,000 samples, representing the best case for the MCMC sampler. The case shown in Fig. 19(*a*) and Table 3 is based on the prior distribution. Here, the MCMC sampler had an acceptance rate of ~30%. Figure 19(*b*), and Table 4 shows the case based on the combined learning posterior distribution and had an acceptance rate of ~40%. In general, the MCMC output is somewhat noisy and slightly underestimates the probability at the tails. Including more samples improves accuracy but increases calculation time. However, there are a wide variety of MCMC samplers available in the literature, which may perform better or worse for specific applications. In contrast, the integration-based method result is nearly indistinguishable from the nominal. The slight divergence on the x_2 distribution is due to the range chosen for the integration: x_2 was only evaluated between 55 and 75, resulting in a truncated normal distribution which has slightly increased probability density near the mean.

Discussion

Results show that the system parameters were able to converge close to the experimental values using the Bayesian approach. This confirms the viability of the Bayesian approach for identifying system parameters and suggests the potential for transfer learning. For example, a cutting force model for a tool/workpiece combination can be learned on one machine tool and then re-used later, e.g., if the cutting tool is transferred to a different machine. These results also highlight the importance of intelligent test selection. Even with six high spindle-speed cutting tests, there was still significant uncertainty about the true values of K_s and β since the cutting tests were not selected to maximize information gain. Specifically, only one of the high-speed tests significantly altered the posterior distribution. One advantage of the Bayesian approach is that test cuts can be iteratively selected based on the current uncertainties about the system in order to maximize information gain. Prior Bayesian-based approaches have either performed test cuts in a grid pattern [10,11] or selected test cuts designed to identify high-productivity machining parameters [8,9,12,13]. Future work will investigate test selection strategies which focus on reducing uncertainty in the underlying system variables.

Even after updating, the posterior stability map showed some errors at low spindle speeds, with the stability boundary located slightly above the 1.5 mm axial depth test cuts from 800 rpm to

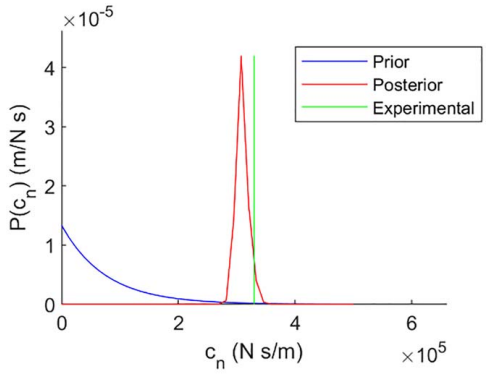

Fig. 17 Distribution for c_n after updating using the low spindle-speed data. Distributions are not presented for K_s and β since the maximum-likelihood estimators were used as fixed values after the first updating step.

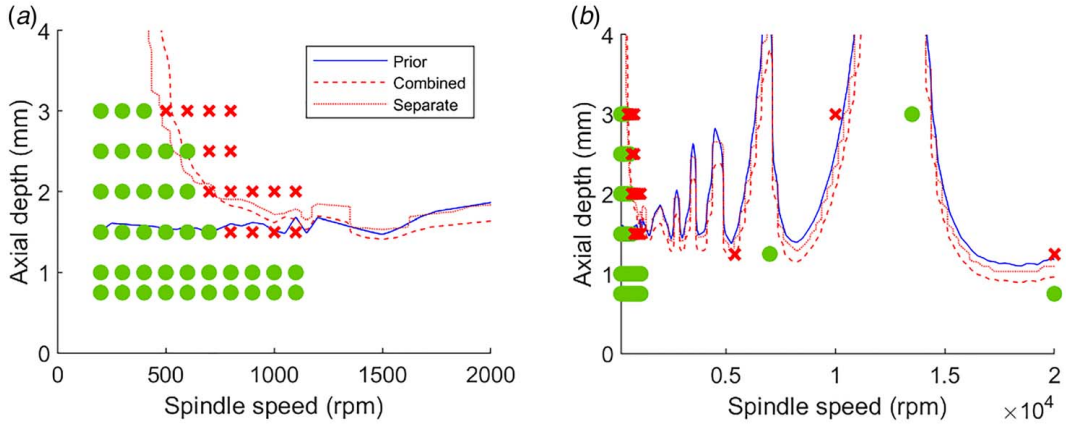

Fig. 18 Posterior stability maps, comparing the 50% probability of stability boundary for the different updating methods: (*a*) low spindle-speed region, and (*b*) high spindle-speed region

1100 rpm (Fig. 16). This error is attributed to three sources. First, while they were treated as fixed for this study, there is still some uncertainty in the flexure dynamics obtained by tap testing, both due to uncertainty in the tap testing [26] and the changes in dynamics as the material is removed from the workpiece. Second, while the Bayesian method applied here treats the stability boundary as a sharp cutoff, there is usually some uncertainty about what exactly qualifies as a stable or unstable cut. The in-process data for the cutting tests in Ref. [3] are not available, and the results were classified as stable or unstable without additional qualifying information. Finally, inaccuracies in the process damping model affect the fit. With only the single variable c_n, it is not possible for a stability map to exactly fit all experimental cutting tests. Note that, even in the paper that published the original dataset [3], there were several cutting tests which were not predicted by the best-fit stability map (see Fig. 6(*c*)).

These model errors are currently a major limitation of the physics-informed Bayesian learning approach. Using the physics-based stability models to inform learning can enable the models to efficiently converge and discover information about the underlying system that cannot be identified by naïve approaches. However, errors in the stability models can cause the physics model to fail to converge to the ground truth for underlying system parameters or, in extreme cases, cause the Bayesian updating to fail to identify the correct stability map [13]. This can occur even when the models are broadly accurate and experience only local errors. In contrast, a learning method which is not constrained to the process physics will necessarily be able to converge to the true stability boundary given sufficient cutting tests [8]. However, this does not necessarily provide new insight into the process behavior.

The integration-based method introduced here has both advantages and disadvantages compared to the MCMC sampling

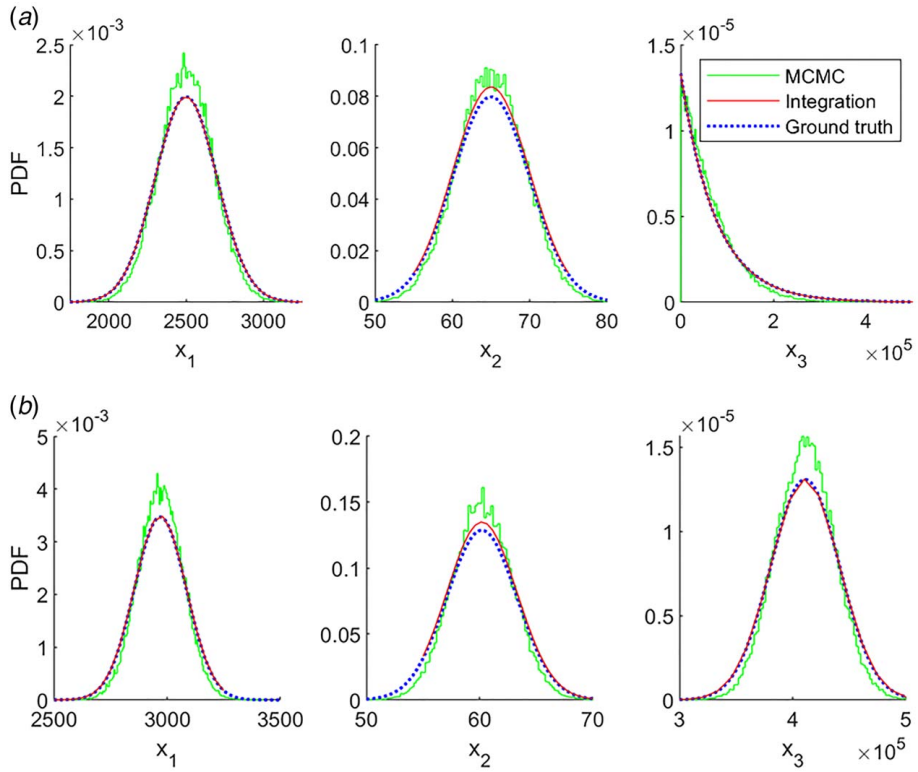

Fig. 19 Accuracy comparison between MCMC and numerical integration methods: (*a*) case shown in Table 3, and (*b*) case shown in Table 4

Table 3 Comparison for the case shown in Fig. 19(a)

Variable	Nominal	MCMC	Integration
$\boldsymbol{x}_1$	$\mathcal{N}(2500,\ 200)$	$\mathcal{N}(2498,\ 176.7)$	$\mathcal{N}(2500,\ 199.7)$
$\boldsymbol{x}_2$	$\mathcal{N}(65,\ 5)$	$\mathcal{N}(64.9,\ 4.4)$	$\mathcal{N}(65.2,\ 4.4)$
$\boldsymbol{x}_3$	$\mathcal{E}(75000)$	$\mathcal{E}(71980)$	$\mathcal{E}(75086)$

Table 4 Comparison for the case shown in Fig. 19(b)

Variable	Nominal	MCMC	Integration
$\boldsymbol{x}_1$	$\mathcal{N}(2970,\ 115)$	$\mathcal{N}(2970,\ 99.0)$	$\mathcal{N}(2968,\ 111.9)$
$\boldsymbol{x}_2$	$\mathcal{N}(60.2,\ 3.1)$	$\mathcal{N}(60.2,\ 2.66)$	$\mathcal{N}(60.5,\ 2.8)$
$\boldsymbol{x}_3$	$\mathcal{N}(411000,\ 30500)$	$\mathcal{N}(410995,\ 26360)$	$\mathcal{N}(410810,\ 30226)$

approaches that have been applied in previous work. The grid sampling approach enables stability maps to be precalculated, removing the heavy computational load so that the Bayesian updating itself can be performed in close to real time. This is particularly important if iterative test cut selection is to be implemented, since MCMC would require new sampling after each update in order to select new test cuts. This is not computationally feasible when algorithms such as TDS are used. In contrast, the integration approach updates the probability associated with precalculated samples and can therefore be performed quickly, regardless of the computation time for the selected algorithm, although this approach scales poorly as the dimensionality of $\boldsymbol{\theta}$ increases. Even with access to high-performance computing resources, this becomes infeasible for complex systems with many uncertainties.

This issue can be partly addressed by separating the Bayesian updating into smaller subproblems with fewer variables. Here, this separation was achieved using high and low-speed cutting tests, but it can be accomplished in other ways with various data sources. For example, if the goal is to learn the system's natural frequencies, cutting force model, and process damping model, then the natural frequencies could be learned from the chatter frequencies obtained from a microphone that records the machining sound [12], the cutting force model could be learned from spindle torque or dynamometer data, and the process damping learned from binary stable/unstable classifications at low spindle-speed cutting tests. This approach can significantly reduce the number of required samples and the overall computation time. However, since it does not consider the uncertainty for all variables during each updating step, this approach must be used cautiously, especially in cases where cutting tests have not been selected to optimize information gain, and there is still significant uncertainty remaining.

Conclusion

This study demonstrated for the first time that Bayesian learning can identify both the stability map and underlying system parameters from a set of cutting tests. Novel sampling and updating strategies were proposed to enable the use of computationally expensive stability models. These methods were tested using experimental results from the literature, and the updating results were found to closely agree with measured values. It was demonstrated that the method was able to learn a process damping model represented by a single coefficient with no direct force or vibration measurements. It is anticipated that the method can be extended to other learning applications in milling and, potentially, other manufacturing processes.

Conflict of Interest

There are no conflicts of interest.

Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request.

Nomenclature

b = axial depth of cut
h = instantaneous chip thickness
m = once-per-flute stability metric
c_n = normal direction process damping coefficient
F_D = process damping force
K_s = specific cutting force coefficient
$b_{lim}(\boldsymbol{\theta}, \Omega)$ = limiting axial depths at a spindle speed for a set of system parameters
$x,\ y$ = tool/workpiece relative displacement in the X and Y axes
$\dot{x}, \dot{y}$ = tool/workpiece relative velocity in the X and Y axes
$S(\boldsymbol{\theta}, \Omega, b)$ = stability prediction for a given set of system parameters, spindle speed, and axial depth
F_t, F_n = cutting force components in the tangential and normal directions
α = a set of cutting test results
β = cutting force angle
Δ_b = stability error used to calculate the likelihood function $P(\alpha|\boldsymbol{\theta})$
$\boldsymbol{\theta}$ = a set of system parameters which can be used to calculate a stability map
σ_b = stability error tolerance hyperparameter used in the likelihood function
Ω = spindle speed

References

[1] Schmitz, T., and Smith, K. S., 2009, *Machining Dynamics: Frequency Response to Improved Productivity*, Springer, New York.
[2] Altintas, Y., Eynian, M., and Onozuka, H., 2008, "Identification of Dynamic Cutting Force Coefficients and Chatter Stability With Process Damping," CIRP Ann., **57**(1), pp. 371–374.
[3] Tyler, C., and Schmitz, T., 2013, "Analytical Process Damping Stability Prediction," J. Manuf. Processes, **15**(1), pp. 69–76.
[4] Budak, E., and Tunç, L. T., 2010, "Identification and Modeling of Process Damping in Turning and Milling Using a New Approach," CIRP Ann., **59**(1), pp. 403–408.
[5] Budak, E., and Tunc, L. T., 2009, "A New Method for Identification and Modeling of Process Damping in Machining," ASME J. Manuf. Sci. Eng., **131**(5), p. 051019.
[6] Sellmeier, V., and Denkena, B., 2012, "High Speed Process Damping in Milling," CIRP J. Manuf. Sci. Technol., **5**(1), pp. 8–19.
[7] Tuysuz, O., and Altintas, Y., 2019, "Analytical Modeling of Process Damping in Machining," ASME J. Manuf. Sci. Eng., **141**(6), p. 061006.
[8] Karandikar, J., Traverso, M., Abbas, A., and Schmitz, T., 2014, "Bayesian Inference for Milling Stability Using a Random Walk Approach," ASME J. Manuf. Sci. Eng., **136**(3), p. 031015.
[9] Karandikar, J., Honeycutt, A., Smith, S., and Schmitz, T., 2020, "Milling Stability Identification Using Bayesian Machine Learning," Procedia CIRP, **93**, pp. 1423–1428.
[10] Chen, G., Li, Y., Liu, X., and Yang, B., 2021, "Physics-Informed Bayesian Inference for Milling Stability Analysis," Int. J. Mach. Tools Manuf., **167**, p. 103767.
[11] Li, K., He, S., Liu, H., Mao, X., Li, B., and Luo, B., 2020, "Bayesian Uncertainty Quantification and Propagation for Prediction of Milling Stability Lobe," Mech. Syst. Signal Process, **138**, p. 106532.
[12] Schmitz, T., Cornelius, A., Karandikar, J., Tyler, C., and Smith, S., 2022, "Receptance Coupling Substructure Analysis and Chatter Frequency-Informed Machine Learning for Milling Stability," CIRP Ann., **71**(1), pp. 321–324.
[13] Cornelius, A., Karandikar, J., Gomez, M., and Schmitz, T., 2021, "A Bayesian Framework for Milling Stability Prediction and Reverse Parameter Identification," Procedia Manuf., **53**, pp. 760–772.
[14] Ahmadi, K., 2022, "Bayesian Updating of Modal Parameters for Modeling Chatter in Turning," CIRP J. Manuf. Sci. Technol., **38**, pp. 724–736.
[15] Altintaş, Y., and Budak, E., 1995, "Analytical Prediction of Stability Lobes in Milling," CIRP Ann., **44**(1), pp. 357–362.
[16] Brooks, S., Gelman, A., Jones, G., and Meng, X. L., 2011, *Handbook of Markov Chain Monte Carlo*, CRC Press, New York.

[17] Calderhead, B., 2014, "A General Construction for Parallelizing Metropolis–Hastings Algorithms," Proc. Natl. Acad. Sci., **111**(49), pp. 17408–17413.
[18] Denkena, B., Grabowski, R., Krödel, A., and Ellersiek, L., 2020, "Time-Domain Simulation of Milling Processes Including Process Damping," CIRP J. Manuf. Sci. Technol., **30**, pp. 149–156.
[19] Altıntas, Y., Engin, S., and Budak, E., 1999, "Analytical Stability Prediction and Design of Variable Pitch Cutters," J. Manuf. Sci. Eng., **121**(2), pp. 173–178.
[20] Campomanes, M. L., and Altintas, Y., 2003, "An Improved Time Domain Simulation for Dynamic Milling at Small Radial Immersions," ASME J. Manuf. Sci. Eng., **125**(3), pp. 416–422.
[21] Honeycutt, A., and Schmitz, T., 2017, "Milling Stability Interrogation by Subharmonic Sampling," ASME J. Manuf. Sci. Eng., **139**(4), p. 041009.
[22] Honeycutt, A., and Schmitz, T., 2018, "Milling Bifurcations: A Review of Literature and Experiment," ASME J. Manuf. Sci. Eng., **140**(12), p. 120801.
[23] Schmitz, T., Betters, E., Budak, E., Yüksel, E., Park, S., and Altintas, Y., 2022, "Review and Status of Tool Tip Frequency Response Function Prediction Using Receptance Coupling," Precis. Eng.., **79**, pp. 60–77.
[24] Cao, H., Li, B., and He, Z., 2012, "Chatter Stability of Milling With Speed-Varying Dynamics of Spindles," Int. J. Mach. Tools Manuf., **52**(1), pp. 50–58.
[25] UTK Office of Innovative Technologies, "ISAAC-NG," UTK, https://oit.utk.edu/hpsc/isaac-open-enclave-new-kpb/. Accessed June 12, 2023
[26] Kim, H. S., and Schmitz, T., 2007, "Bivariate Uncertainty Analysis for Impact Testing," Meas. Sci. Technol., **18**(11), pp. 3565–3571.