

Establishing Model Credibility for Process-Microstructure-Property Relationships in Additive Manufacturing Using Exascale Computing

J Belak, M Rolchigo, M Titov, J Coleman, J Turner,
R Carson, S Reeve, G Knapp, M Bement, A
Plotkowski

January 2026

The International Journal of High Performance
Computing Applications

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Establishing Model Credibility for Process-Microstructure-Property Relationships in Additive Manufacturing Using Exascale Computing

Matt Rolchigo^{1*}, John Coleman¹, Robert Carson², Mikhail Titov³, Samuel Temple Reeve¹, Gerry Knapp¹, Alex Plotkowski¹, James Belak², Matt Bement¹, and John Turner¹

¹ Oak Ridge National Laboratory, Oak Ridge, TN

² Lawrence Livermore National Laboratory, Livermore, CA

³ Brookhaven National Laboratory, Upton, NY

January 12, 2026

Abstract

Additive Manufacturing (AM) of alloys holds significant promise as a disruptive technology in various industries, yet its adoption is often hindered by challenges in achieving consistent part quality. These issues are primarily due to the complex process-microstructure-property (PSP) relationships inherent to AM. Computational models can greatly aid in understanding these relationships, but their widespread impact and adoption has been limited by a lack of validated, open-source, and computationally efficient PSP modeling frameworks and hardware limitations. This study leverages the ExaAM software suite and data from the AMBench-2018 series of laser powder bed fusion (LPBF) benchmark experiments to perform a comprehensive model assessment, including verification, validation, sensitivity analysis, and uncertainty quantification. The RADICAL-EnTK workflow manager was used to perform an ensemble of heat transport, solidification, and mechanical response simulations on the exascale computer Frontier, considering uncertainties in critical model inputs such as laser spot size and nucleation parameters, and consisting of 125 explicit grain structure simulations and 7,875 crystal plasticity simulations. For a selected location within the Inconel 625 AMBench-2018 test artifact, sensitivity analysis and uncertainty quantification were performed using the predicted distributions of grain structure and mechanical properties. Qualitative agreement was found between the predicted grain size and texture and the observed AMBench-2018 microstructure, the mean predicted yield stress was within 5% of the experimental measurement mean, and the mean predicted engineering stress at 5% strain was within 10% of the experimental measurement mean. The insights gained from development and validation of the ExaAM PSP modeling framework will help guide future directions for enhancing the credibility and reliability of PSP models in AM, thereby accelerating the adoption of AM technologies in various industries.

1 Introduction

Additive Manufacturing (AM) of alloys is a transformative processing technology that enables the production of parts with unique geometries, microstructures, and mechanical properties unattainable through conven-

* Corresponding author: 1 Bethel Valley Rd., Oak Ridge, TN 37831-2008, USA, rolchigomr@ornl.gov. Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/doe-public-access-plan>).

tional processing methods [Herzog²⁰¹⁶, Askari²⁰²⁰, Yan²⁰²⁰]. These capabilities, coupled with flexible part design, make AM highly attractive in sectors such as biomedical, aerospace, automotive, and nuclear energy [Gisario²⁰¹⁹, Sabahi²⁰²⁰, Blakey-Milner²⁰²¹, Salmi²⁰²¹, Madhavadas²⁰²²]. However, widespread adoption in industries requiring high reliability is hindered by challenges in achieving consistent part quality and reproducibility [Mostafaei²⁰²²]. These challenges often stem from the complex multiscale physical phenomena that occur during the build process [Sames²⁰¹⁶]. Understanding relationships between processing, structure, and properties (PSP) is essential for addressing these challenges.

A comprehensive understanding of PSP relationships, supported by computational models and experimental data, can potentially be used to optimize the microstructure and mechanical performance of AM parts for a desired application. While empirical relationships and simple process models offer key insights into certain PSP relationships, they can be limited by the complex nature of coupled PSP phenomena across processing space, unclear transferability across different AM systems and alloys, and uncertainty to both model form and model input values. To address these limitations, various multiscale, multiphysics modeling workflows for AM have been developed [Yan²⁰¹⁸, Herriott²⁰¹⁹, Markl²⁰²⁰, Heo²⁰²¹, Mitchell²⁰²³, Saunders²⁰²³].

Prediction of PSP relationships for AM is a complicated, multidisciplinary problem that requires the following features model and software features for a PSP modeling workflow:

- **Validated physics-based models:** It is well known that the processing conditions affect the grain structure via the conditions in the wake of the melt pool produced by the heat source, and that the grain size, shape distribution, and texture affect the constitutive properties and their anisotropy [Nadammal²⁰¹⁷, Marattukalam²⁰²⁰, Gokcekaya²⁰²¹, Wei²⁰²¹]. Even when utilizing performant and scalable codes, simulation of all physical phenomena associated with the PSP relationships is intractable, and models must make reasonable approximations regarding any unresolved features of processes such as melt pool evolution and grain structure formation. To justify such approximations, the components of the AM modeling workflow should be validated against melt pool dimensions, grain structure, and constitutive properties from experiments. Existing process-microstructure modeling workflows have successfully reproduced the grain size distribution and texture of AM parts [Mohebbi²⁰²⁰, Pauza²⁰²¹·MSMSE, Zinovieva²⁰²², Camus²⁰²³, Koepf²⁰²³] and many workflows include models of constitutive properties to predict how quantities such as Young’s modulus and yield strength vary as a function of processing conditions or local microstructure [Zinovieva²⁰²², Pauza²⁰²¹·JMPEP, Kergabner²⁰²¹, Heo²⁰²¹]. However, validation against experimental data for each stage of the PSP workflow – prediction of melt pool, microstructure, and constitutive properties – is rarely performed. For full confidence in a PSP modeling framework this multi-stage calibration is needed; furthermore, the models should be able to reproduce not just a single data point, but also the variability and trends observed in experimental data to be useful as a predictive tool for future AM experiment design. This includes, for example, melt pool dimension variation with spot size and texture variation with melt pool dimensions.
- **Software and workflow flexibility:** The PSP modeling framework and associated workflow should be adaptive and customizable for various user needs. Examples of modeling workflows of interest for AM problems include individual simulations of process, microstructure, and properties; UQ studies on one or more component models; and sensitivity studies to examine specific model linkages or the effects of subsets of input parameters. The majority of the existing PSP software frameworks for AM (with the notable exception of JAX-AM [jax^{am}]) are not publicly developed, limiting their widespread use across academia and industry as well as their adaptability to specific user needs. Open source software and workflow tools are both more accessible and more easily allow individual users to customize the workflow to suit their individual simulation needs, including changing microstructure models or adding a material deformation model, for example.
- **Code performance and scaling:** Due to the complex and multiscale nature of the phenomena governing heat transport, phase transformation, and mechanical response during AM processing, the software used to accurately and explicitly simulate PSP relationships tend to require long runtimes and significant compute resources. Additionally, many model inputs, such as the absorptivity and nucleation site density, have significant uncertainties that lead to variability in model predictions at each

stage of the PSP workflow further exacerbating computational requirements with ensembles of simulations. A PSP modeling framework should account for these uncertainties by predicting distributions of melt pool shapes, grain structures, and constitutive properties. Multiple modeling workflows have investigated the sensitivity of model output to particular inputs; for example, Herriott et al. investigated the effect of nucleation parameter variation on the predicted stress-strain behavior of stainless steel 316L [Herriott’2019], and Saunders et al. constructed a Gaussian process-based surrogate from model data to predict PSP linkages [Saunders’2023]. To perform individual PSP simulations, let alone the more complex PSP workflows necessary to perform ensembles of simulations, computationally efficient models that can leverage high performance computing resources are necessary. Given that nearly all of the existing PSP modeling frameworks heavily utilize CPU-only codes, their ability to perform large ensembles of simulations will be limited by compute resource availability.

While these requirements are not exhaustive for understanding all aspects of PSP relationships (e.g., modeling porosity and residual stress is also crucial for improving part reproducibility), they represent a minimum threshold for an accessible workflow tool that can accurately and efficiently obtain distributions of predicted properties across a range of part locations and processing conditions.

The ExaAM software developed as part of the Exascale Computing Project (ECP) was designed to uniquely satisfy the aforementioned requirements, allowing not only for location-specific simulation of melt pool conditions, grain structure, and constitutive properties during an AM build but also uniquely examine aspects of the PSP relationships through large ensembles of simulations [Turner’2022]. The software’s physics capabilities and its ability to leverage Frontier, the exascale computer at the Oak Ridge Leadership Computing Facility (OLCF), are shown in this study to yield key insights into AM PSP relationships that would not have been possible with existing CPU-only software and smaller compute systems. Distributions of predictions from each model, accounting for input parameter uncertainty and variability in upstream model inputs, are verified and validated against data from the extensive Additive Manufacturing Benchmark (AMBench) suite produced by the National Institute of Standards and Technology (NIST). The following section describes the relevant experimental data used for software validation along with three core models in the ExaAM software suite and the workflow manager. We subsequently define the problem of interest - the calibration of the models and selection of input parameter uncertainties followed by a demonstration of the ExaAM software suite’s capability to execute ensembles of AM simulations. The example results focus on model trends with respect to uncertain inputs, highlighting comparisons to experiments and model inputs to which predicted properties are the most sensitive. Finally, we discuss areas for future improvement in establishing model credibility for predicting AM PSP relationships.

2 Methods

Data from the Additive Manufacturing Benchmark (AMBench) series of experiments and measurements conducted by the National Institute of Standards and Technology (NIST) [Levine’2020] was used to calibrate and assess the process, microstructure, and property model credibility. Details regarding specific measurements and datasets used for model calibration and validation are given in Subsection 2.1. The main components of the workflow are as follows, with the software stack diagram given in Figure 1, and more information about each modeling component in Subsection 2.2.

- **AdditiveFOAM** (<https://github.com/ORNL/AdditiveFOAM>) is a continuum heat and mass transfer code built on the OpenFOAM finite volume library. AdditiveFOAM addresses the challenge of disparate length scales between the heat source and the simulated part through (static) local mesh refinement and MPI for parallel execution on CPU. AdditiveFOAM will be detailed further in Subsubsection 2.2.1.
- **ExaCA** (<https://github.com/LLNL/ExaCA>) is a cellular automata (CA) code used to predict explicit representations of the grain structure in as-solidified alloys. ExaCA uses the Kokkos library [kokkos] and (GPU-aware) MPI for performance portability on both CPUs and GPUs and has simulated among the largest (in terms of number of computational cells) grain structures reported in the existing literature to date [Rolchigo’2022]. ExaCA will be detailed further in Subsubsection 2.2.2.

- **ExaConstit** (<https://github.com/LLNL/ExaConstit>) is a quasi-static crystal plasticity code built on the MFEM finite element library [**mfem-library**] utilizing the ExaCMech library for the crystal plasticity constitutive response [**ecmech**] and the RAJA library for performance portability on both CPUs and GPUs [**raja**]. At the time of this writing, ExaConstit is the only known open-source crystal plasticity finite element code capable of simulation on multiple GPU. ExaConstit will be detailed further in Subsubsection 2.2.3.
- **RADICAL-EnTK** (<https://github.com/radical-cybertools/radical.entk>) is a python library developed by the RADICAL Research Group [**Balasubramanian'2016**] to develop and execute large-scale ensemble-based workflows across computing systems. RADICAL-EnTK will be detailed further in Subsubsection 2.2.4.
- **TASMANIAN** (<https://github.com/ORNL/TASMANIAN>) is a collection of robust libraries for high dimensional integration and interpolation as well as parameter calibration. TASMANIAN will be detailed further in Subsubsection 2.2.4.

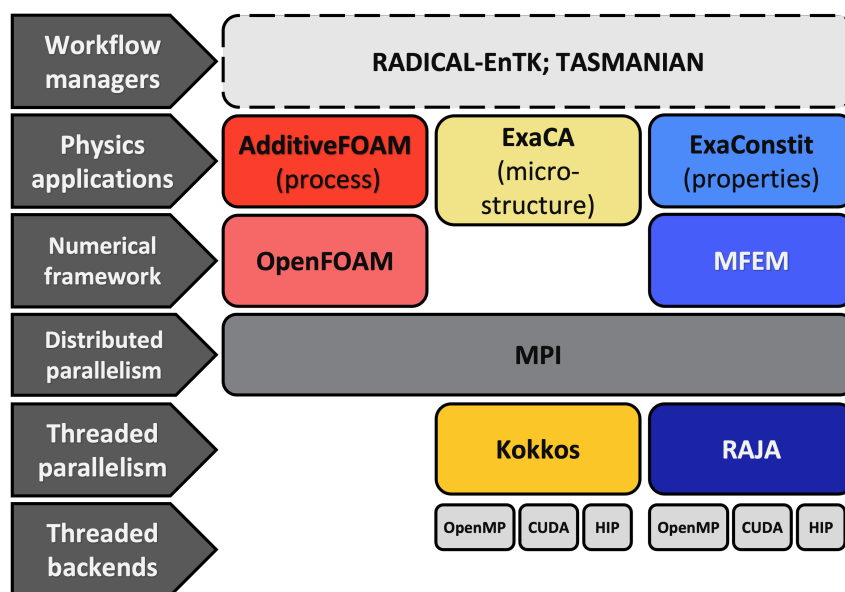


Figure 1: Software stack including each component of the ExaAM PSP workflow, with performance portability layers, framework libraries, and workflow managers included.

2.1 Experiments

Data and artifacts from the AMBench-2018-01 and AMBench-2018-02 experiments were used to calibrate and/or validate the models in the ExaAM workflow. The AMBench-2018-01 experiments built test artifacts with dimensions as shown in Figure 2a, where a pattern of support legs (“thick legs” of length 5 mm, “medium” legs of length 2.5 mm, and “thin” legs of length 0.5 mm) is repeated with a layer height of 20 μm to support a bridge-like structure. The test artifacts were built using a bidirectional scan strategy, with +/- X direction scans on odd numbered layers and +/-Y direction scans on even numbered layers. These artifacts were sectioned for Electron Backscatter Diffraction (EBSD) measurements of grain size and crystallographic orientation distributions. EBSD data reproduced from Stoudt et al [**Stoudt'2020**] for an unspecified location in the interior of L4 (see label in Figure 2a) along with corresponding inverse pole figures are shown in Figure 2b. The microstructure shows a bimodal distribution of grain areas, consisting of small area grains with relatively narrow widths (5-25 μm) often found between larger area grains (50-100 μm wide) with near-square cross-sections. These small grains had a preferential cube texture $\langle 100 \rangle$ aligned with

the specimen directions X, Y, and Z) while the large grains had a texture favoring $\langle 110 \rangle$ alignment with Z, $\langle 100 \rangle$ or similar crystallographic directions aligned with Y, and a weak alignment of $\langle 100 \rangle$ and $\langle 111 \rangle$ with X. Some of the large grains also exhibited a cube texture, or intermediate orientations such as the alignment of $\langle 210 \rangle$ with Z (appearing as yellow on the IPF-Z map). Thermography data from portions of the AMBench-2018-01 builds from Heigel et al. [Heigel'2020] was also considered in this study, and showed discrepancies in melt pool size and cooling conditions between odd and even numbered layers of the builds.

Additional grain structure measurements were taken at Oak Ridge National Laboratory (ORNL) from various locations within one of the AMBench-2018-01 test artifacts. Sectioning L7 (see label in Figure 2a) at a distance of 1 mm above the baseplate top, a stitched XY cross-section was prepared and imaged at the ORNL Manufacturing Demonstration Facility (MDF) metallography lab and is shown in Figure 2c. The surface of these samples was polished to a finishing step using 0.05 μm colloidal silica and EBSD was performed on a Gemini 450 FESEM with an Oxford detector using a 5 μm step size. Here, it is shown that the observed $\langle 110 \rangle // Z$ $\langle 100 \rangle // Y$ texture in the L4 leg interior from Figure 2b is more or less consistent with that observed in the XY plane of L7, and more large area grains with the cube texture appear in the 0.5 mm near each edge. Other XY sections at various distances from the baseplate top from L7, L8, L9, and the "bridge" supported by the legs showed that the grain structures and textures from various Z locations in the part interior were qualitatively the same as observed in Figure 2b [Levine'2023]. From these observations, it can be assumed that the grain structure in the interior of the supporting legs (at least 0.5 mm from the part edges and at least 1 mm above the baseplate top) are reasonably consistent.

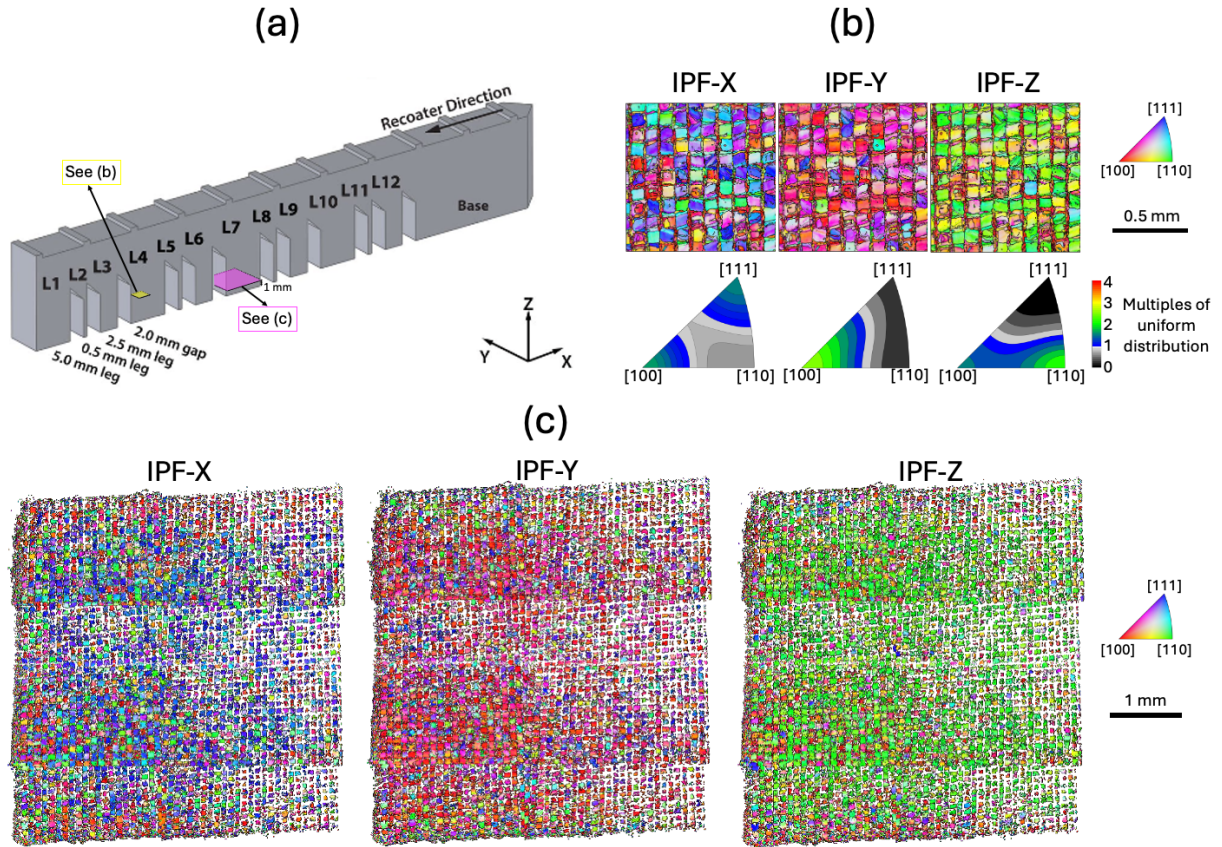


Figure 2: (a) Schematic of the AMB2018-01 test artifact, adapted from [Stoudt'2020]. (b) EBSD images of an XY plane within L4, reproduced from [Stoudt'2020]. (c) EBSD images of a full XY cross-section of L7, sectioned 1 mm above the baseplate top.

The AMBench-2018-01 test artifacts were also sectioned for macroscopic uniaxial compression experiments performed at both Lawrence Livermore National Laboratory (LLNL) and the Cornell High Energy

Synchrotron Source (CHESS). These experiments were conducted at a quasi-static strain rate of 0.001 s^{-1} to measure the yield strength (YS) and stress at 5% strain (S5) for compression loading in the build direction (Z) and one of the transverse directions (Y). Results are summarized in Table 1.

Table 1: Summary of experimental macroscopic stress-strain data at a strain rate of 0.001 s^{-1} where BD is the build direction (Z), TD is one of the transverse directions (Y), YS is the yield strength, and S5 is the stress at -5% strain.

Experiment Name	YS BD (MPa)	YS TD (MPa)	S5 BD (MPa)	S5 TD (MPa)
Experiment 1 / LLNL Exp 1	-750.0	-740.0	-951.0	-915.0
Experiment 2 / CHESS Exp 1	-741.26	-726.19	-967.0	-921.0
Experiment 3 / CHESS Exp 2	-741.24	-720.87	-951.0	-908.0
Experiment 4 / CHESS Exp 3	-686.13	-675.80	-916.0	-858.2

The AMBench-2018-02 experiments consisted of detailed measurements of melt pool dimensions, taken from individual laser tracks on an Inconel 625 baseplate and laser spot sizes ranging from 76 to 311 μm [Weaver²⁰²²]. The reported spot size used to build the AMBench-2018-01 test artifact was 100.0 μm , which was shown to be on the transition between conduction and keyhole mode for the AMBench-2018-02 single tracks [Weaver²⁰²²]. In conduction mode the melt pools are hemispherical, whereas in keyhole mode the melt pools become deeper with elongated tails due to multiple internal reflections in the vapor cavity increasing the laser absorption.

2.2 Models

2.2.1 Process

Simulating AM parts starts with a process model linking the moving heat source (in this case, a laser) to the temperature field evolution and phase transformations in the material. The conservation of energy in AdditiveFOAM is defined as:

$$\rho c_p \left[\frac{\partial T}{\partial t} + \nabla \cdot (\mathbf{u}T) \right] = \nabla \cdot (k\nabla T) + \rho L_f \frac{\partial f_s}{\partial t} + \dot{Q}(\mathbf{x}, t) \quad (1)$$

where ρ is density, c_p is specific heat at constant pressure, \mathbf{u} is velocity, T is temperature, t is time, k is thermal conductivity, L_f is latent heat of fusion, f_s is the solid mass fraction, and $\dot{Q}(\mathbf{x}, t)$ is a volumetric heating rate which is a function of position \mathbf{x} and time.

The first source term in Eq.(1) accounts for latent heat evolution during solidification and melting, which is modeled using a linear relationship of fraction solid between the liquid temperature, T_l , and the eutectic temperature, T_e . The second source term in Eq.(1) accounts for the heat input to the system by the moving laser, which is modeled using the two-parameter transient heat source derived in Coleman et al [Coleman²⁰²⁴]:

$$\dot{Q}(\mathbf{x}, t) = \frac{\eta P}{V_0} \exp \left(- \left[\frac{\Delta x^2 + \Delta y^2}{r^2} \right]^{k/2} \right) \quad (2)$$

where η is the effective absorption, P is the laser power, Δx , and Δy are the relative distances from the laser center, and the coefficient k is the radial distribution parameter for a super-Gaussian profile. The radial width of the volumetric heating profile is a function of depth, such that:

$$r = \frac{2\sigma}{\sqrt[k/2]{}} \left[1 - \left| \frac{\Delta z}{d} \right|^{m-1} \right]^{1/m} \quad (3)$$

where 2σ is half the International Organization for Standards (ISO) definition for beam width ($D4\sigma$), Δz is the relative distance from the laser center, d is the heat source depth, and m is the volumetric shape factor that controls the final shape of heat input. Finally, the volume integral of the 3D distribution is:

$$V_0 = \pi \Gamma \left(1 + \frac{2}{k} \right) \left(\frac{2\sigma}{\sqrt[2]{2}} \right)^2 d \left[\frac{\Gamma \left(1 + \frac{1}{m} \right) \Gamma \left(1 + \frac{2}{m} \right)}{\Gamma \left(1 + \frac{3}{m} \right)} \right] \quad (4)$$

where $\Gamma(x)$ is the gamma function, which can be numerically computed following DiDonato et al. [DiDonato'1986].

This volumetric source term accounts the effects of changing melt pool shape on transient absorption of the laser energy using the geometric expression for a conical cavity:

$$\eta(a) = \eta_0 \frac{1 + (1 - \eta_0)(G - F)}{1 - (1 - \eta_0)(1 - G)} \quad (5)$$

where η_0 is the Fresnel absorption, $G = 1/(1 + \sqrt{1 + a^2})$ is the ratio between the area of the cavity opening to the surface area of the cavity, $F = \sin^2(\theta)$ is the view factor, and $a = d/(2\sigma)$ is the aspect ratio of the cavity. The depth of the liquidus isotherm is used to calculate both the transient heat source depth in Eq.(3) and the dynamic absorption in Eq.(5).

As a consequence of AdditiveFOAM's heat source formulation, it accurately reproduces a wide range of melt pool shapes relevant to metal AM [Knapp'2023, Coleman'2024]. As surface tension-driven forces at the vapor-liquid interface have been observed to significantly affect melt pool evolution in the AMBench experiments (see Weaver et al. [Weaver'2022]), AdditiveFOAM serves a need for a process model that incorporates these necessary fluid flow effects without the computational cost of simulating higher-fidelity physics such as vapor cavity formation [Coleman'2024].

2.2.2 Microstructure

ExaCA is a cellular automata (CA)-based model for the microstructure at the scale of grain envelopes ($\sim 10^{-5}$ - 10^{-6} m), approximating the solute field at and around the solidification front. This is the more commonly-used length scale for modeling AM problems as it enables comparison of larger-scale microstructural features such as grain size, shape and texture to EBSD data from AM builds [Zinovieva'2023, Koepf'2023, Xie'2023, Chen'2023, Teferra'2021, Mohebbi'2020]. It should also be noted that Phase field (PF) and CA methods utilizing GPU hardware have been used to explicitly predict the solute field and the cellular and/or dendritic nature of the solidification front at a more refined length scale ($\sim 10^{-6}$ - 10^{-7} m) [Takaki'2018, Ohno'2020, Sakane'2020, Ghosh'2022, Fattebert'2023, Yuan'2024]. However, accurately modeling coupled solute transport and solidification at the part scale (multiple locations and volumes on the order of 1 mm^3) with morphology-scale models remains intractable even with exascale computers. As a grain-scale model, ExaCA therefore provides a balance of the necessary physical fidelity for accurate predictions and the need for sufficiently large volumes from which grain and texture statistics can be extracted.

ExaCA simulates a grain structure using a series of simple, geometry-based rules that approximate the physics of dendritic solidification for alloys with a primary cubic phase (such as Inconel 625) [Gandin'1997]. The advance of the solidification front from existing grains is tracked via octahedral grain envelopes associated with active cells. The growth rate of a grain envelope per time step is governed by the local undercooling ΔT in the associated cell and an interfacial response function $V(\Delta T)$. This function in turn relates the growth velocity of the $\langle 100 \rangle$ dendrite tips (which are not tracked explicitly) at the grain edges to the undercooling at these dendrite tips as calculated a priori from either experimental data, analytical model approximations, or information from higher resolution dendritic solidification models. In the present work, the interfacial response function is of the form

$$V(\Delta T) = \frac{\Delta t}{\Delta x} \cdot (A(\Delta T)^3 + B(\Delta T)^2 + C(\Delta T)), \quad (6)$$

where A , B , and C are fitting parameters for a given material, Δx is the cell size, and Δt is the time step.

In addition to the solidification of existing grains, ExaCA simulates the nucleation of grains in the undercooled liquid as the competition between these phenomena plays a significant role in the grain structure evolution [Hunt'1984]. The columnar to equiaxed grain structure transition (CET) is typically suppressed in LPBF due to the large thermal gradients in the undercooled liquid. As a result, the grain structure of a given layer of an AM part typically consists of epitaxial growth from the grains of the previously deposited layer. However, the large undercooling at the solidification front will promote the nucleation of new grains

that also play a key role in the development of texture and competition between grains. In ExaCA, a heterogeneous nucleation density N_0 is considered to approximate the heterogeneity in the undercooled liquid. The undercooling needed to trigger a given nucleation event is assigned from a Gaussian distribution described by

$$f(\Delta T) = \frac{N_0}{\Delta T_\sigma \sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{\Delta T_N - \Delta T}{\Delta T_\sigma}\right)^2\right). \quad (7)$$

where ΔT_N is the mean nucleation undercooling and ΔT_σ is the standard deviation of the distribution.

As a GPU-capable grain scale model, ExaCA fulfills the need for a microstructure model that can efficiently simulate sufficiently large microstructures to obtain representative statistics on grain size and texture, in turn enabling prediction of trends in process-microstructure relationships. The time-temperature history data produced by AdditiveFOAM is supplied through an input file as was first described in Rolchigo et al. [Rolchigo'2020], and was recently extended to include more details of the remelting and resolidification behavior common to AM problems [Rolchigo'2024]. This file-based approach decouples the ExaCA and AdditiveFOAM resource utilization and performance for the large multilayer simulations performed in this work while retaining sufficient time-temperature history data to accurately predict microstructure.

2.2.3 Properties

ExaConstit is a crystal plasticity model used to calculate anisotropic material deformation due to slip under various loading conditions. The elastic stress-strain behavior of the material can be written as

$$\sigma = \mathbf{C} : \varepsilon, \quad (8)$$

Where σ is the Cauchy stress, \mathbf{C} is the 4th order single crystal elasticity tensor, and ε is the elastic strain tensor. There are three independent constants represented in \mathbf{C} for the cubic crystal symmetry observed in solidified Inconel 625: C_{11} , C_{12} , and C_{44} [Mouhat'2014].

ExaConstit solves for the balance of linear momentum as given in Eq.(9) using typical solid mechanics FEM discretizations of the weak form for nonlinear materials. Since ExaConstit is an updated Lagrangian FEM code, the primal variable being solved for during this solution is the velocity field at the mesh nodes.

$$\nabla \cdot \sigma = 0 \quad (9)$$

While ExaConstit solves for the balance of linear momentum, it relies on external libraries to provide the end of time step Cauchy stress at each quadrature point throughout the mesh. Typically, a material model might rely on either the deformation gradient or velocity gradient as the kinematic quantity of interest to update the model. For the crystal mechanic models used here, ExaConstit relies on ExaCMech to update the material state variables and provide an updated Cauchy stress based on the velocity gradient and values of the beginning of time step state variables.

A basic description of the crystal mechanics kinematics and models used in ExaCMech will now be covered. From the velocity gradient, \mathbf{L} , we can decompose the kinematic response as

$$\mathbf{L} = \dot{\mathbf{F}}\mathbf{F}^{-1} = \dot{\mathbf{V}}^e \mathbf{V}^{e-1} + \mathbf{V}^e \hat{\mathbf{L}}^p \mathbf{V}^{e-1}, \quad (10)$$

where \mathbf{F} is the deformation gradient, \mathbf{V}^e is the elastic left stretch tensor, and $\hat{\mathbf{L}}^p$ is defined as

$$\hat{\mathbf{L}}^p = \dot{\mathbf{R}}^* \mathbf{R}^{*T} + \mathbf{R}^* \dot{\mathbf{F}}^p \mathbf{F}^{p-1} \mathbf{R}^{*T}, \quad (11)$$

where \mathbf{R}^* is the lattice rotation and \mathbf{F}^p is the plastic deformation gradient. The plastic velocity gradient, $\bar{\mathbf{L}}^p = \dot{\mathbf{F}}^p \mathbf{F}^{p-1}$ is related to the combined macroscopic shearing of the slip systems via

$$\bar{\mathbf{L}}^p = \sum_{\alpha=1}^n \dot{\gamma}^\alpha \bar{\mathbf{s}}^\alpha \otimes \bar{\mathbf{m}}^\alpha \quad (12)$$

where $\dot{\gamma}^\alpha$ is the plastic shearing rate on the α slip system and $\bar{\mathbf{s}}^\alpha \otimes \bar{\mathbf{m}}^\alpha$ is the Schmid tensor defined in the crystal reference frame. The plastic shearing rate is described by the chosen slip kinetics and hardening

models. While ExaCMech has support for a number of different slip kinetics and hardening models, the present work uses a power law slip kinetics formulation and a Voce hardening model form. The slip kinetics take the form

$$\dot{\gamma}^\alpha = \dot{\gamma}_0 \left(\frac{\tau^\alpha}{g^\alpha} \right)^{\frac{1}{m}} \text{sgn}(\tau^\alpha) \quad (13)$$

where $\dot{\gamma}_0$ is a reference shearing rate, τ^α is a resolved shear stress on the α slip system, g^α is the critically resolved shear strength on the α slip system, and m is the rate sensitivity of slip. The resolved shear stress can be computed using

$$\tau^\alpha = \sigma : (\hat{\mathbf{s}}^\alpha \otimes \hat{\mathbf{m}}^\alpha)_{sym}, \quad (14)$$

where σ is the Cauchy stress. g^α is described using a Voce hardening model through

$$\dot{g}^\alpha = h_0 \left(\frac{g_{sat} - g^\alpha}{g_{sat} - g_0^\alpha} \right)^{n_{slip}} \sum_{\beta=1}^{n_{slip}} \dot{\gamma}_\beta, \quad (15)$$

where h_0 is the strength hardening coefficient, g_0 is the initial critically resolved shear strength, and g_{sat} is the saturation critically resolved shear strength.

As input, ExaConstit takes representative volume elements (RVE) provided by ExaCA (in this case, consisting of 400^3 cells) [Rolchigo’2022b]. Each RVE is coarsened to a 200^3 voxel microstructure, taking the most populated grain ID in a 2^3 neighborhood and using that grain ID for 1 voxel in the new microstructure. This voxelized microstructure was then meshed into a 200^3 linear hexahedron element mesh with each grain associated with a unique crystallographic orientation. ExaConstit’s model formulation and GPU-capabilities enable scalable, efficient, and accurate representation of the complex physics and model requirements associated with material deformation (such as remeshing a finite element grid after a material has undergone large deformation), satisfying the needs of the PSP modeling workflow.

2.2.4 Process-Structure-Properties Workflow

RADICAL-EnTK is the workflow manager used to couple and launch ensembles of AdditiveFOAM, ExaCA, and ExaConstit simulations. An example of such a workflow run on Frontier is available at <https://github.com/ExascaleAM/Workflows>. One use case for this workflow is the propagation of uncertainties through a chain of coupled models to understand sources of potential variability in the overall model predictions. Because these codes are computationally intensive and designed for large computational resource sets, direct sampling of specific model outputs to obtain statistically significant distributions of the outputs of interest is not possible. Instead, we used the TASMANIAN library (<https://github.com/ORNL/TASMANIAN>) [Stoyanov’2016] to both construct sparse grids for sampling the model input space and constructing surrogate models through linear interpolation after obtaining the simulated outputs for each point on the model input space. Because multiple codes must be connected for this purpose, each containing its own inputs and associated uncertainties, the overall workflow to properly perform the uncertainty quantification (UQ) analysis is complex. The workflow manager was designed to manage this complexity by automatically scheduling and launching parallel jobs and communicating necessary data between models as the data become available.

A schematic of the workflow executed by RADICAL-EnTK in the present study is shown in Figure 3. First, AdditiveFOAM and ExaCA input files are generated using TASMANIAN sparse grids along with bounds for chosen uncertain input parameters. The workflow manager launches a series of AdditiveFOAM simulations for the chosen location within the AMBench-2018-01 part to generate thermal histories for ExaCA. In this case, the “thick leg” L7 (which has an area of 5 mm by 5 mm) from the AMBench-2018-01 artifact from Figure 2a is the area of interest. For each set of AdditiveFOAM inputs, one odd numbered layer (with scans in the +/- X direction) and one even numbered layer (scans in the +/-Y direction) are simulated using a refined mesh near the center of L7 and a coarser mesh for the outer edges of the leg and the baseplate. Once these jobs are complete, the workflow manager launches a series of ExaCA simulations for the 1 mm by 1 mm area at the center of L7 using AdditiveFOAM data. From the explicit microstructures generated

by ExaCA, each spanning at least 1 mm in the Z direction and each representing a unique permutation of AdditiveFOAM and ExaCA input parameters, representative volume elements (RVEs) of size 0.5 mm^3 are extracted for use by ExaConstit. These RVEs were taken from the central 0.5 by 0.5 mm area spanning $Z = 0.5$ to 1.0 mm. The workflow manager then launches a series of ExaConstit simulations using the RVEs to generate yield surfaces. At each stage, relevant model predictions can be used to generate surrogate models and response surfaces to evaluate uncertainty in model outputs given the uncertain model inputs.

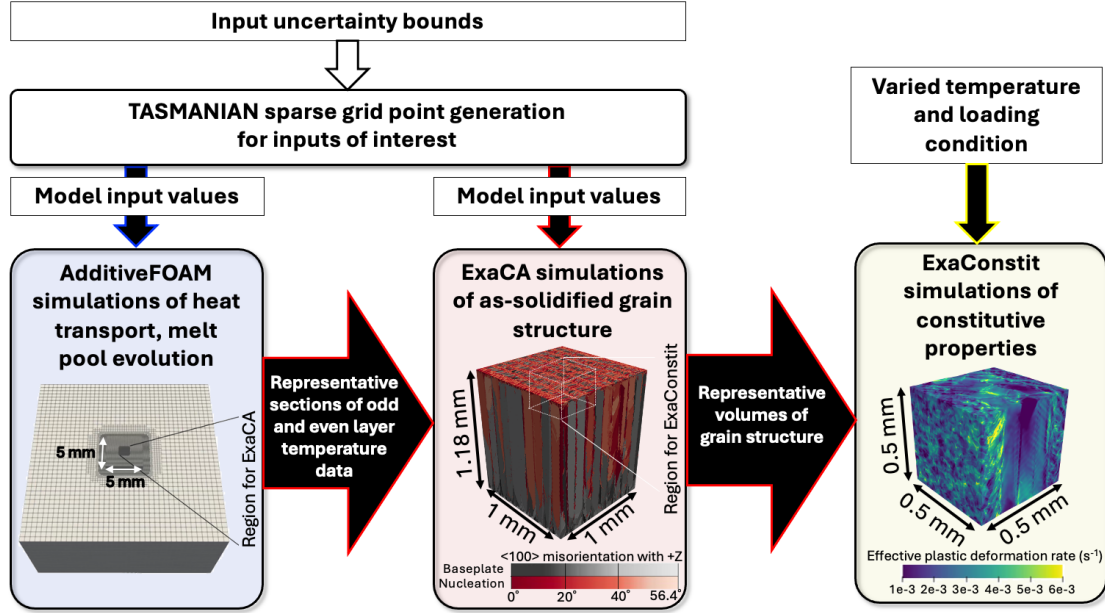


Figure 3: The process-microstructure-properties workflow considered in the present work for a representative region within an AMBench-2018-01 “thick leg” (L1, L4, L7, etc from Figure 2a). The heat transport for the full area of the leg and the surrounding baseplate is simulated using AdditiveFOAM, with representative time-temperature history data for one odd and one even numbered layer passed to ExaCA. ExaCA in turn stacks these representative time-temperature histories to produce explicit representations of the as-solidified grain structure, representative volumes of which are used by ExaConstit to predict distributions of material responses under various loading conditions. TASMANIAN is used to generate distributions of input parameters for AdditiveFOAM and ExaCA.

2.3 Problem Statement

This goal of this study is to use the ExaAM software suite on Frontier to execute a large ensemble of AdditiveFOAM, ExaCA, and ExaConstit simulations, enabling new insights into the complex PSP relationships inherent to AM. As part of this work, we demonstrate the ability of ExaAM software to accurately predict location-specific microstructure and properties in the interior of the AM-Bench-2018-01 test artifact and perform a comprehensive model assessment for the workflow components. To accomplish this goal, appropriate model inputs should be calibrated, uncertain inputs to which simulation results are the most sensitive must be identified, and the inputs to vary in the simulation ensemble should be selected. These topics and preliminary simulation results used to support our choices for input parameter selection and variation are given in Subsection 3.1. With the ranges of inputs parameters and simulation ensemble size identified, the ExaAM workflow was executed on the exascale computer Frontier, the results of which are shown in Subsection 3.2.

Table 2: AdditiveFOAM model input parameters for Inconel 625.

Fixed Parameters			
Name	Symbol	Value	Units
Density	ρ	7569.92	kg/m ³
Specific heat capacity	c_p	750.65	J/(kg · K)
Thermal conductivity (solid)	k_s	8.441 + 0.0147 · T	J/(m · s · K)
Thermal conductivity (liquid)	k_l	5.1182 + 0.0148 · T	J/(m · s · K)
Latent heat of fusion	L_f	2.9 × 10 ⁵	J/(kg)
Liquidus temperature	T_L	1620	K
Eutectic temperature	T_E	1410	K
Layer preheat temperature	T_0	353	K
Radial distribution parameter	k	7.95	—
Volumetric shape parameter	m	2.72	—
Fresnel absorption	η_0	0.3	—
Variable Parameters			
Name	Symbol	Values	Units
Spot size	$D4\sigma$	80 —120	μm

3 Results

3.1 Model Input Parameter Selection

3.1.1 AdditiveFOAM

AdditiveFOAM was calibrated against AMBench-2018 data as part of a previous study [Coleman’2024], and the calibrated values for the radial distribution parameter k , the volumetric shape parameter m , and the Fresnel absorption η_0 are given in Table 2. Given that the laser spot size $D4\sigma$ of 100 μm was shown to be at a transition between conduction and keyhole mode heat transport regimes for single laser tracks on a bare plate, and that there were significant differences in the melt pool behavior for odd and even numbered layers of the AMBench-2018-01 build due to asymmetry in the part dimensions and scan paths [Heigel’2020], there is significant uncertainty in the exact melt pool dimensions and shapes during the build of the AMBench-2018-01 artifacts. Rolchigo et al. showed that these changes in melt pool orientation can significantly change the grain texture by influencing the volume of remelted material and the misorientation angle of thermal gradients in this remelted region [Rolchigo’2024]. As shown in Table 2, the spot size $D4\sigma$ in the present work was varied over a range that reasonably captures the uncertainty in experimental spot size and spans both sides of the conduction-keyhole mode transition.

3.1.2 ExaCA

The nucleation input parameters have been shown to be the most significant source of model uncertainty and sensitivity for both CA models [Li’2018, Andersson’2023, Camus’2023] and analytical models of microstructure development under AM conditions [Haines’2018]. The most significant of these input parameters, the heterogeneous nucleation density N_0 and the mean nucleation undercooling ΔT_N , are selected for variation in the ensemble of ExaCA simulations. Reasonable ranges for these N_0 and ΔT_N must be selected, and values for other fixed input parameters must be chosen as well. For the fixed input parameters, a rough calibration for the mean substrate grain size S_0 can be found in the Appendix along with rationale for the chosen cell size Δx , time step Δt , and interfacial response function $V(\Delta T)$.

ΔT_N controls the relative distance between nucleated grains and the epitaxial solidification front in the undercooled liquid trailing the melt pool. For a fixed $V(\Delta T)$, varying ΔT_N will affect grain competition and selection. An athermal nucleation model is used here, where

$$\Delta T_N = 3\Delta T_\sigma, \quad (16)$$

ensuring that some finite nucleation is possible at zero undercooling and the distribution of nucleation undercooling gets wider as the mean undercooling increases. This follows the assumptions of Queded and Greer for deterministic nucleation of potent nucleant particles during rapid cooling [Quested'2005]. Despite the large thermal gradients present in the melt pool, new grains appear in the microstructure at different points along the build direction as shown in the auxiliary EBSD image in the Appendix. Given that the solidification front velocity in the majority of the melt pool is near or less than 0.15 m/s based on prior thermal gradient and solidification velocity data from AdditiveFOAM, the upper bound for ΔT_N is set to the undercooling at which $V(\Delta T_N) = 0.15$ m/s. Using the chosen interfacial response function $V(\Delta T)$, this upper bound was around 30 K. The lower bound of ΔT_N is set to 3 K to ensure that only a negligible portion of nucleation events occur near or at the liquidus temperature.

N_0 values used in the literature span several orders of magnitude from 7×10^8 [Zinoviev'2016] to 10^{15} m^{-3} [Liu'2020], with some simulations not considering nucleation at all. As the balance of growth from the substrate and the nucleation and growth of new grains is expected to be critical to grain size and texture prediction, preliminary simulations were performed to provide reasonable bounds for N_0 for the present study. These simulations were performed at $\Delta T_N = 16.5$ K (midway through the 3 to 30 K range selected for study) and XY cross-sections colored using the IPF-Z map are shown in Figure 4 alongside a reproduction of the EBSD data from Figure 2b. Qualitatively, the best matches for the distribution of grain areas observed in the EBSD data from Figure 2b occur at $N_0 = 10^{14}$ and 10^{15} m^{-3} . Additionally, the simulation at $N_0 = 10^{14}$ m^{-3} showed the best match with the EBSD's strength of the $\langle 110 \rangle$ alignment with the Z direction. As a result of these observations, the range of N_0 selected for this study spanned 10^{14} through 10^{15} m^{-3} .

The microstructure model inputs are summarized in Table 3. The layer height H of 20 μm was chosen to match that reported for the AMBench-2018-01 build, and simulations consisted of 56 layers to obtain representative grain structures 1 mm above the baseplate top while also avoiding sampling grain structure from the region melted on successive layers.

Table 3: ExaCA model input parameters for Inconel 625.

Constant Parameters			
Name	Symbol	Value	Units
Cell size	Δx	1.25	μm
Time step	Δt	0.0625	μs
Layer height	H	20	μm
Number of layers	N_{layers}	56	-
Baseplate grain size	S_0	12.3	μm
IRF third order	A	-1.03×10^{-7}	$m/K^3 \cdot s$
IRF second order	B	1.05×10^{-4}	$m/(K^2 \cdot s)$
IRF first order	C	2.22×10^{-3}	$m/(K \cdot s)$
Varied Parameters			
Name	Symbol	Values	Units
Heterogeneous nucleation density	N_0	10^{14} to 10^{15}	m^{-3}
Mean nucleation undercooling	ΔT_N	3 to 30	K

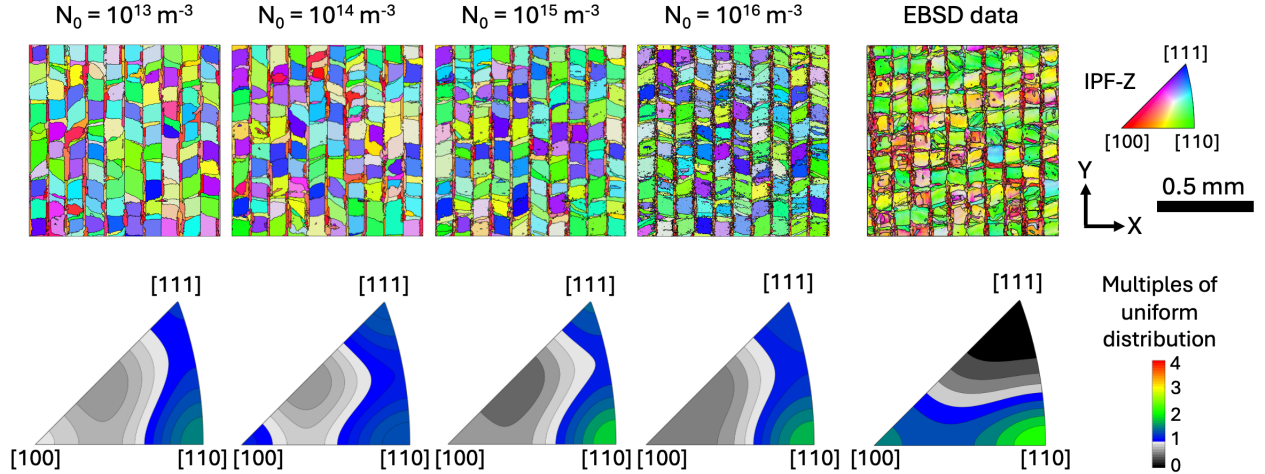


Figure 4: ExaCA simulation results perpendicular to the Z direction and 1 mm above the baseplate top from simulations using four heterogeneous nucleation density (N_0) values, colored using the IPF-Z map. Also shown is EBSD data from the AMBench-2018-01 build [Stoudt’2020] and corresponding inverse pole figures from the ExaCA and EBSD data.

3.1.3 ExaConstit

The crystal plasticity models in ExaConstit were fit to experimental macroscopic uniaxial compression data [AMBench22’Comp, AMBench22’exaam], with single crystal elastic moduli C_{11} , C_{12} , and C_{44} determined by Wang et al [Wang’2016] and given in Table 4. ExaConstit used data from “Experiment 1” in Table 1 (see [AMBench22’Comp]) at room temperature with loading in the build direction through roughly 30% strain to calibrate the strength hardening coefficient h_0 , initial critical resolved shear stress g_0 , saturation critical resolved shear stress g_{sat} , and rate sensitivity of slip m . ExaConstit calibration did not consider the transverse loading direction data from “Experiment 1” as the model inherently accounts for anisotropic effects through the single crystal elasticity tensor and non-uniformity of the crystallographic texture. Additionally, given the isotropic nature of the hardening model, including the transverse loading direction data in the calibration would not necessarily have lead to an improved model fit. The other three experimental datasets were not used in model calibration, and were instead used to validate the distribution of predicted properties from the PSP workflow.

A “first guess” microstructure was used to calibrate the viscoplastic input parameters in ExaConstit. This microstructure was generated using AdditiveFOAM and ExaCA with estimated model inputs. The RVE from ExaCA used for ExaConstit calibration is depicted in Figure 5a showing the expected bimodal distribution of grain sizes and the presence of columnar grains penetrating multiple layers. As shown in Figure 5b, the model calibration RVE had a very weak texture in the build direction compared to the EBSD data from Figure 2b despite the similar grain size and shape distribution. Finally, a comparison of the experimental data to the calibrated macroscopic compression prediction from ExaConstit is shown in Figure 5c, showing that the model successfully reproduced both the yield strength and large strain behavior of the material under loading in the build direction. We note that use of the weakly textured “first guess” microstructure introduces a potential source of error for the constitutive response of RVEs along directions with strong fiber textures - this point will be returned to in the discussion on ExaConstit prediction sensitivities (Subsubsection 4.1.2).

The calibrated ExaConstit model parameters are shown in Table 4. Due to the complexity of the ExaConstit model formulation compared to AdditiveFOAM and ExaCA, variation in the input parameters in Table 4 was not considered. ExaConstit used fixed input parameters with varying temperatures and loading conditions to obtain the temperature-dependent yield surface for each RVE. It was found that 18 varying degrees of biaxial and 3 simple shear loading conditions across the XY, XZ, and YZ planes were the minimum requirement to fit an anisotropic yield surface using ExaConstit. While not explored in the present work, the yield surfaces in turn could be used in macroscopic part-scale thermomechanic simulations

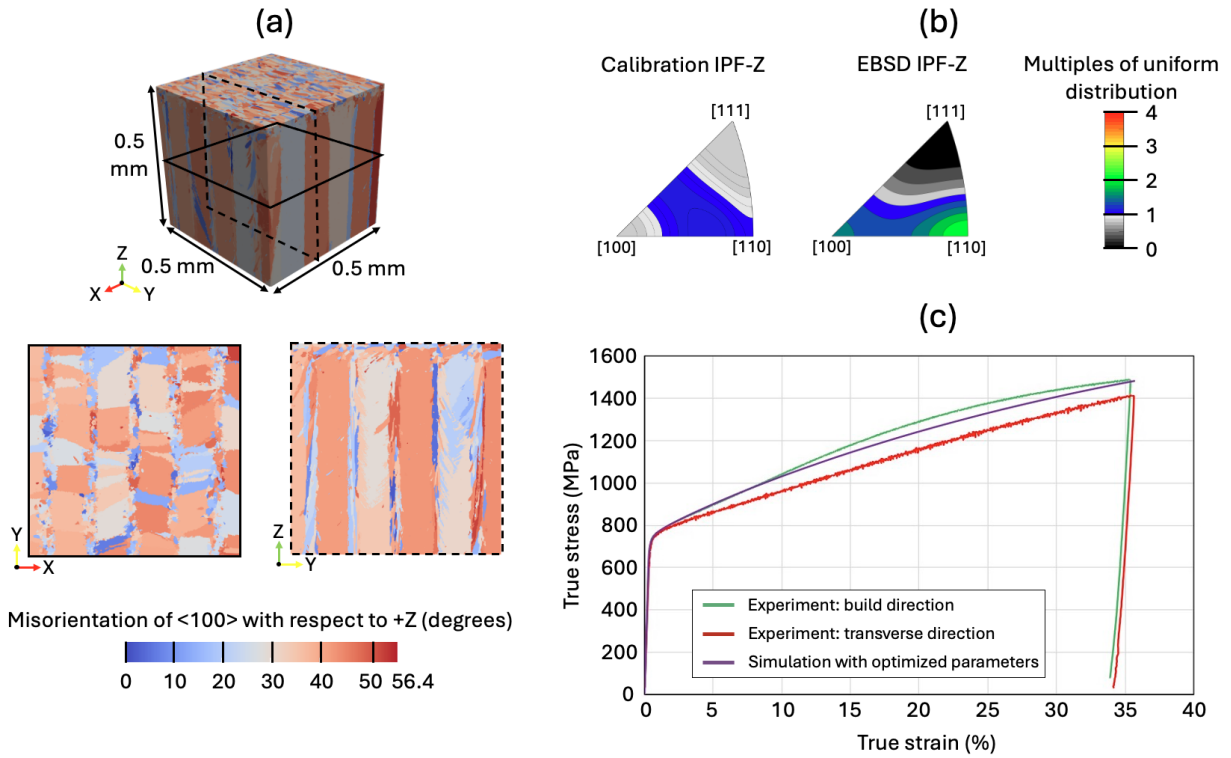


Figure 5: (a) ExaCA-generated RVE used to calibrate ExaConstit. (b) Inverse pole figures with respect to the Z (build) direction for the calibration RVE with a comparison to the EBSD data reproduced from Figure 2b. (c) Macroscopic true stress-strain curves measured from the relevant locations within the AMBench-2018-01 test artifact for the transverse (Y) and build (Z) loading directions, and ExaConstit-predicted true stress-strain behavior for the Z loading direction using the optimized simulation parameters.

of the AMBench-2018-01 part as described in Turner et al [Turner'2022].

Table 4: ExaConstit model input parameters for Inconel 625.

Constant Parameters			
Name	Symbol	Value	Units
Single crystal elasticity matrix coefficient 1, 1	C_{11}	243.3	GPa
Single crystal elasticity matrix coefficient 1, 2	C_{12}	156.7	GPa
Single crystal elasticity matrix coefficient 4, 4	C_{44}	117.8	GPa
Strength hardening coefficient	h_0	485.0	MPa
Initial critical resolved shear stress	g_0	312.75	MPa
Saturation critical resolved shear stress	g_{sat}	745.0	MPa
Rate sensitivity of slip	m	0.015	-
Reference shear rate	$\dot{\gamma}_0$	1.0	s^{-1}
Varied Parameters			
Name	Symbol	Values	Units
Temperature	T	293, 518, and 773	K
Loading condition	σ_{ij}	Various uniaxial, biaxial, and simple shear stress states and magnitudes	-

3.2 ExaAM PSP workflow demonstration

Using the calibration data and preliminary results from Subsection 3.1, the ExaAM PSP workflow was demonstrated for the interior of L7 of the AMBench-2018-01 test artifact. This calculation was run on Frontier, the exascale computer at Oak Ridge National Laboratory. First, the input ranges for $D4\sigma$ in AdditiveFOAM (see Table 2) along with the N_0 , ΔT_N , and ΔT_σ ranges in ExaCA (see Table 3) were used by a Python script that incorporated TASMANIAN’s makeLocalPolynomialGrid function to generate input files for both models. The AdditiveFOAM input space was sampled uniformly using a 1D sparse grid with a depth of 1 to obtain 5 $D4\sigma$ values, while the ExaCA input space was sampled uniformly using a 2D sparse grid with a depth of 2 to obtain 25 permutations of N_0 , ΔT_N , and ΔT_σ (though as mentioned in Subsubsection 3.1.2, ΔT_σ was not independent of ΔT_N). Following heat transport and microstructure simulation, 125 RVEs were taken from the 125 ExaCA simulated microstructures and provided to ExaConstit. For each RVE, the material response to varying stress magnitudes for 21 loading conditions and at each of 3 temperatures were simulated, yielding 7875 total ExaConstit simulation results for this location in the test artifact. As most of the experiments were conducted at room temperature, only the ExaConstit results from $T = 298$ K simulations will be analyzed here. These model inputs and outputs are summarized in Table 5. Code performance, resource utilization, and areas for further workflow optimization will be discussed in Subsection 4.2.

3.2.1 Process

AdditiveFOAM simulations with spot sizes $D4\sigma$ of 80, 90, 100, 110, and 120 μm were performed each using both the odd numbered and even numbered layer scan pattern for L7 of the AMBench-2018-01 build, with data extracted from the central 1 mm by 1 mm region for use by ExaCA. Figure 6 shows outlines of the melt pool cross-sections perpendicular to the scan directions, taken from the center of the 1 mm by 1 mm region for each $D4\sigma$. The bottoms of these melt pool profiles show a series of valleys corresponding to melt pool centerlines and peaks corresponding to the bottom of melt pool overlap regions, outlined in white. Figure

Table 5: Model inputs and outputs for the PSP demonstration calculation ensemble

Model	AdditiveFOAM	ExaCA		ExaConstit	
Input data	N/A	5 odd/even time-temperature histories		125 representative grain structures	
Varied input parameter(s)	Spot size ($D4\sigma$)	Nucleation density (N_0)	Mean nucleation undercooling (ΔT_N)	Temperature (T)	Loading condition (σ_{ij})
Varied input value(s)	80, 90, 100, 110, 120 μm	$1.00 \cdot 10^{14}$, $3.25 \cdot 10^{14}$, $5.50 \cdot 10^{14}$, $7.75 \cdot 10^{14}$, $1.00 \cdot 10^{15}$ m^{-3}	3.00, 9.75, 16.50, 23.25, 30.00 K	293, 523, 773 K	Varied uniaxial, biaxial stress state
Varied input permutations	5	5	5	3	21
Output	5 odd/even time-temperature histories	125 representative grain structures		7875 mechanical response datasets	

6a shows that the even layer melt pools formed for each $D4\sigma$ are larger than the corresponding odd layer melt pools due to the effect of heat accumulation between laser passes. As $D4\sigma$ is increased, the melt pools generally became wider and more shallow for both the odd and even layers. However, the extent of the effect of $D4\sigma$ on the melt pool dimensions differed for the odd and even layers. This is most notable in the sharp transition in melt pool depth for the odd layer melt pools between keyhole mode heat transport at $D4\sigma = 100$ and conduction mode heat transport at $D4\sigma = 110 \mu\text{m}$, which does not appear in the even layer melt pools. Additionally, the melt pool widths for the odd layers are much less sensitive to $D4\sigma$ than for the even layers.

As a consequence of the nonlinear relationships between $D4\sigma$ and melt pool dimensions for the odd and even numbered layers, the extent of remelting when stacking an even numbered layer “n+1” atop an odd numbered layer “n” will depend on $D4\sigma$. This is shown in Figure 6b, where it can be seen that the melting and remelting behavior can be classified into three distinct “zones”. In the top zone (hatched spacing), with height equivalent to the layer height H of 20 μm , the even layer “n+1” will not remelt any of the material that solidified during the odd layer “n”. In the middle zone (dark shading), all locations that solidified in the odd layer will remelt as part of the even layer. In the bottom zone (light shading), only some locations that solidified in the odd layer will remelt as part of the even layer; locations at Y coordinates aligned with even layer melt pool centerlines will be more likely to remelt than locations at Y coordinates aligned with even layer melt pool overlaps. Notably, while all three zones are present for $D4\sigma = 80, 90,$ and $100 \mu\text{m}$, the even layer melt pools for $D4\sigma = 110$ and $120 \mu\text{m}$ completely remelt the material that solidified as part of the odd layer. This transition from partial to complete remelting of the material that solidified during odd numbered layers has significant implications for the microstructure simulations, as the grain structure for simulations with $D4\sigma = 110$ and $120 \mu\text{m}$ should predominantly be influenced by the even layer melt pool geometry, and those with $D4\sigma = 80, 90,$ and $100 \mu\text{m}$ should be more significantly influenced by the odd layer melt pool geometry.

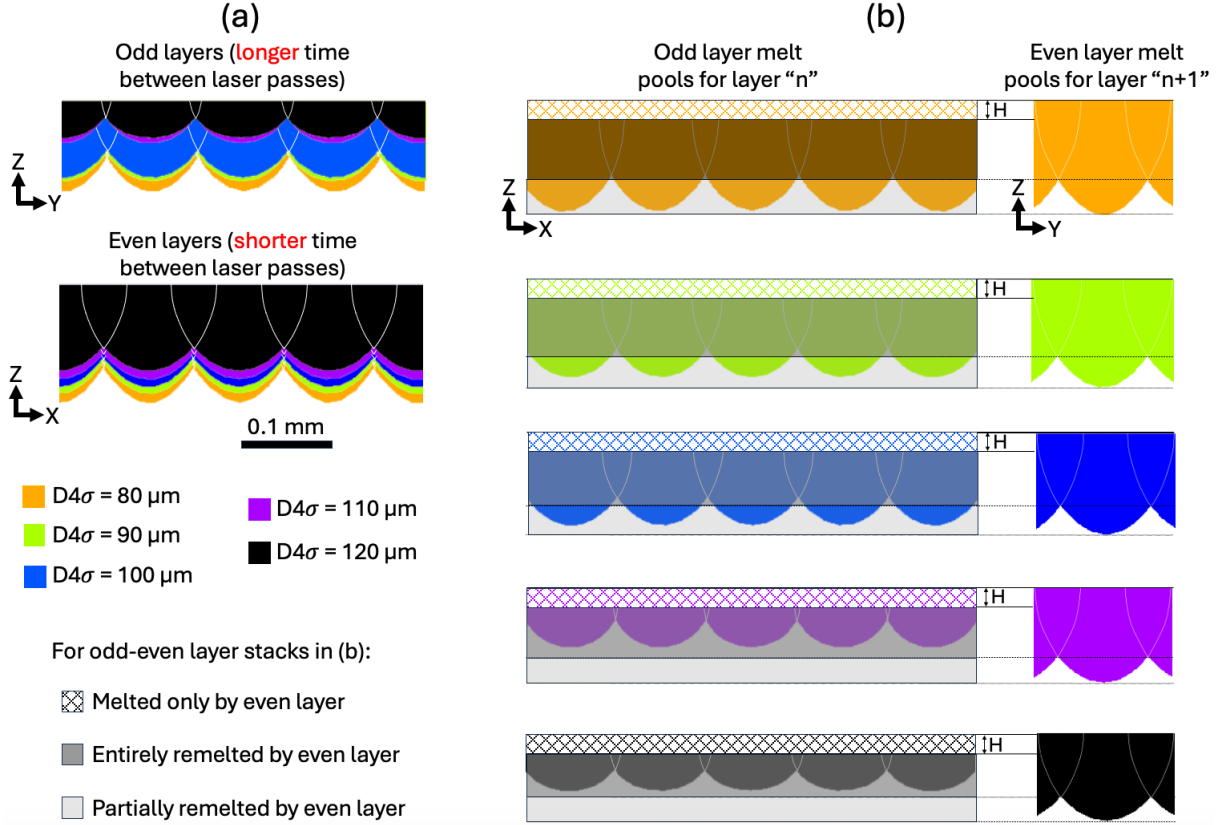


Figure 6: (a) Melt pool cross-sections perpendicular to the scan direction for each $D4\sigma$ and a subset of the scans in the odd and even layers. (b) The extent of remelting of the odd layer's melt pool footprint during the next simulated even layer.

3.2.2 Microstructure

ExaCA used the representative odd and even layer time-temperature history data at each of the five $D4\sigma$ values to generate an ensemble of 125 ExaCA grain structures, each using a unique permutation of $D4\sigma$, N_0 , and ΔT_N (in turn linked to ΔT_σ). Data at the XY cross-section located 1 mm above the baseplate top was post-processed using the MTEX toolbox in Matlab to obtain grain size, shape, and texture statistics. The grain structures colored using each of the inverse pole figure X, Y, and Z colormaps are shown in Figure 7 for the 125 cross-sections, organized by the AdditiveFOAM $D4\sigma$ and ExaCA N_0 and ΔT_N . The cross-sections for various input parameters were generally similar, with $\langle 110 \rangle$ and $\langle 111 \rangle$ crystallographic directions predominantly aligned with X and Z and $\langle 100 \rangle$ crystallographic directions predominantly aligned with Y. The effects of the odd-even scan pattern with fixed hatch spacing are visible in the symmetrical nature of the grain shapes and sizes. The most notable difference among these microstructures is the transition between $D4\sigma = 100$ and 110 μm from grains with square cross-sections and $\langle 110 \rangle$ and $\langle 111 \rangle$ crystallographic directions aligned with Z, to grains with rectangular cross-sections (elongated in Y, the scan direction for the even layers) and primarily $\langle 110 \rangle$ directions aligned with Z. A similar but more subtle trend towards increased $\langle 110 \rangle$ alignment with X occurred between $D4\sigma = 100$ and 110 μm . In the IPF-Y maps, a subtle shift in the proportions of $\langle 100 \rangle$ and similar ($\langle 210 \rangle$ and $\langle 211 \rangle$) grain alignments between $D4\sigma = 100$ and 110 μm occurred as well. These texture transition correspond to the transition in odd layer melt pool sizes and odd-even remelting behavior from Figure 6. Additionally, for any given $D4\sigma$, the grain size was reduced with increasing N_0 and decreasing ΔT_N and the texture slightly strengthened. The microstructure trends will be examined further in Section 4.

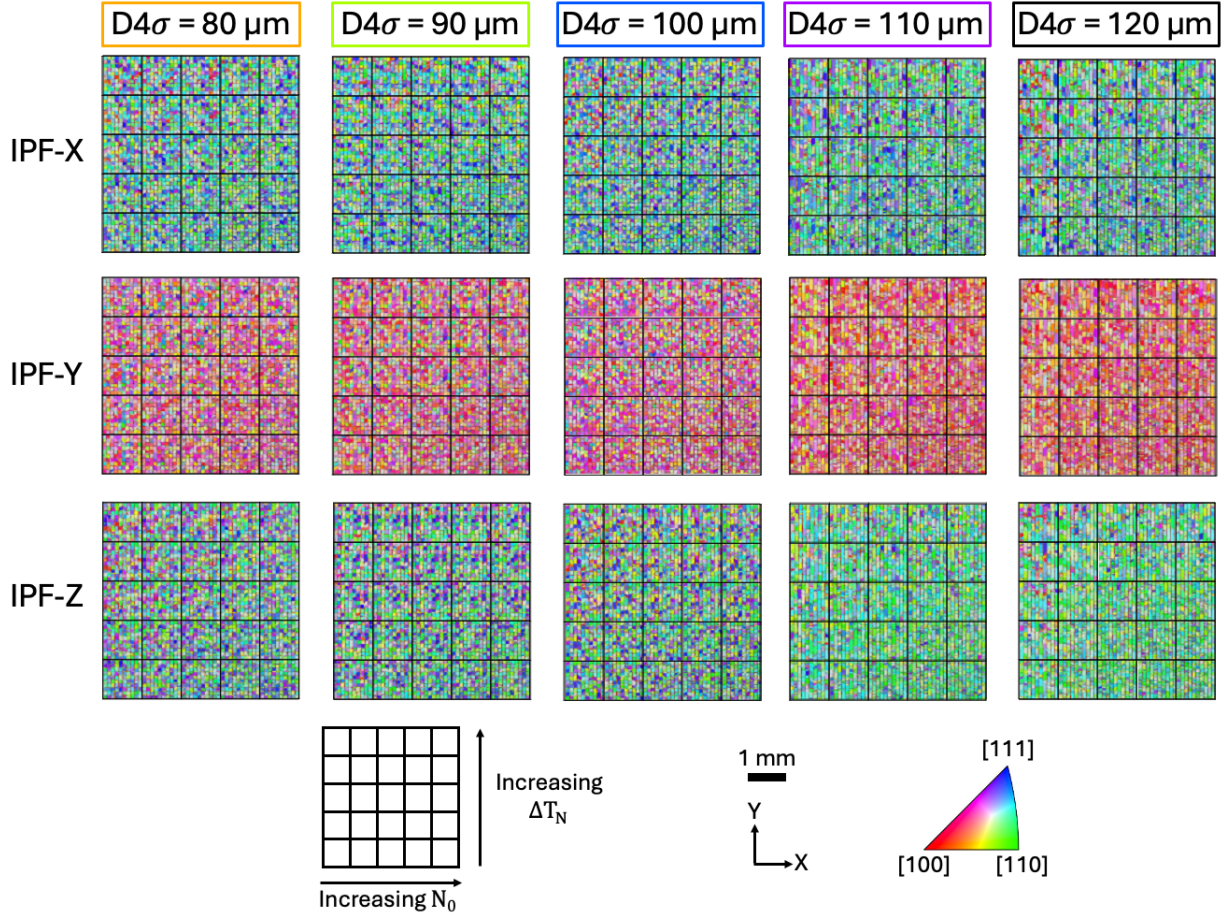


Figure 7: XY cross-sections from each of the 125 ExaCA microstructures, taken at 1 mm above the baseplate top, and organized by the input parameters $D4\sigma$, N_0 , and ΔT_N . The microstructures are colored using each of the inverse pole figure X, Y, and Z maps.

Features of the predicted grain structures and the transition between grain structures from small spot size simulations (where the odd and even layer melt pools contribute towards the predicted the solidification behavior) to large spot size simulations (where the even layer melt pools dominate the predicted solidification behavior) can be rationalized using the quantity θ_s to describe the texture favored by a given melt pool overlap geometry [Rolchigo'2024]. While the transient thermal conditions in the AdditiveFOAM data result in a range of θ_s for a given $D4\sigma$ rather than the fixed θ_s in the prior work, these ranges tended to span only around ± 2.5 degrees from the mean value. For the even layer melt pool overlaps, the mean θ_s was estimated to be around 27 degrees for $D4\sigma = 80 \mu\text{m}$ and increasing to around 35 degrees for $D4\sigma = 120 \mu\text{m}$. The odd layer melt pool overlaps consistently had smaller θ_s than the corresponding even layer melt pool overlaps, but even for $D4\sigma = 80 \mu\text{m}$, θ_s appeared to nearly always be 22.5 degrees or larger. All overlap geometries therefore correspond to Regime II (as defined in [Rolchigo'2024]), with odd layers favoring a $\langle 110 \rangle // Z$ and $\langle 100 \rangle // X$ texture and even layers favoring a $\langle 110 \rangle // Z$ and $\langle 100 \rangle // Y$ texture. For the $D4\sigma = 110$ and $120 \mu\text{m}$ simulations, the resulting $\langle 110 \rangle // Z$ and $\langle 100 \rangle // Y$ texture is due to the fact that most solidification occurs during the even layers. Further, the grain areas for these simulations are elongated in the Y direction due to the effects of the $\pm Y$ direction scans almost exclusively shaping the microstructure development. For the $D4\sigma = 80, 90,$ and $100 \mu\text{m}$ simulations, the $\langle 110 \rangle // Z$ and $\langle 100 \rangle // Y$ texture is still observed as the majority of solidification is dominated by the larger melt pools of the even layers. However, these textures are weaker than those for $D4\sigma = 110$ and $120 \mu\text{m}$ due to destructive contributions from the odd layer melt pools favoring $\langle 110 \rangle // Z$ and $\langle 100 \rangle // X$.

3.2.3 Properties

Out of the 7875 ExaConstit simulation results, the focus of the analysis in the present work is the temperature and loading conditions that corresponded to the measurements in Table 1. These simulations were the 375 conducted at room temperature (298 K) and considered uniaxial monotonic compression loading in each of the X, Y, and Z directions. For loading in the Y and Z directions, histograms of the predicted yield stress (YS) and engineering stress at 5% strain (S5) are plotted in Figure 8a and Figure 8b, respectively. It can be noticed that for both YS and S5 that the material was stronger in the Z loading direction compared to the Y loading direction. This feature is also present in the model predictions, as the distributions of YS and S5 for each loading direction are distinct and exhibit little overlap. The distributions of YS and S5 for the Z loading direction are similar between experiment and simulation as 3 of the 4 experimental measurements are within the bounds of the predicted distributions (and the fourth measurement within 10 MPa of the predicted distributions), and all predictions are within 40 MPa of at least one experimental measurement. Furthermore, for the Z loading direction, the mean of the YS measurements was underestimated by just 1.5%, and the mean of the S5 measurements by just 3.3%. For the Y loading direction, the model predictions tended to underestimate the measured YS and S5 by somewhat more significant margins (8.2% for YS and 6.7% for S5), and just 1 of the 4 experimental measurements were within the ranges of YS and S5 predicted by ExaConstit. Despite the discrepancies for the Y loading direction, the model reasonably reproduced the constitutive mechanical behavior and all experimental measurements were within 15% of all model predictions for a given loading direction and stress metric.

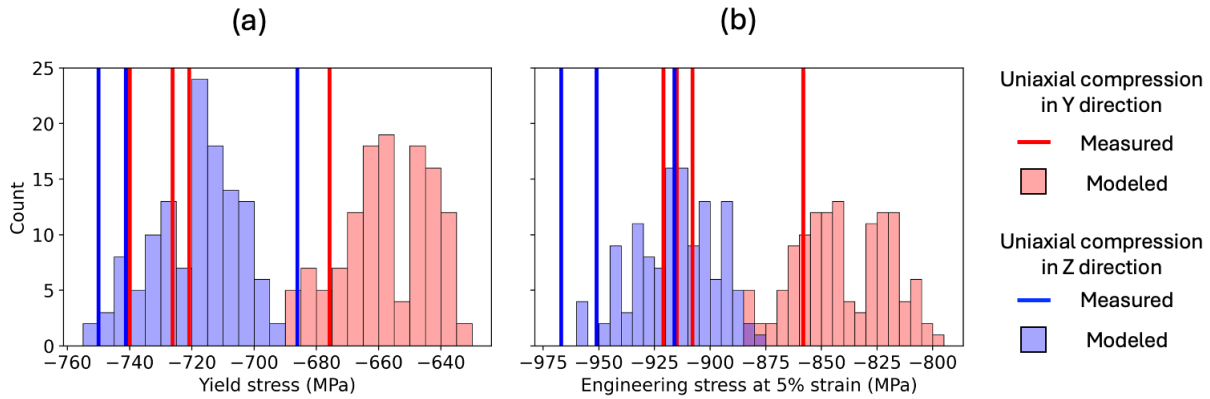


Figure 8: Histograms of ExaConstit predictions for (a) yield stress and (b) engineering stress at 5% strain across all 125 RVEs, alongside measured values (represented as vertical lines) from Table 1. Model predictions and experimental measurements for Y and Z direction loading are represented by red and blue coloring, respectively.

4 Discussion

4.1 Prediction accuracy and input sensitivity

4.1.1 Microstructure

The microstructure data can be used to understand trends regarding individual model outputs, such as the mean grain cross-sectional area as a function of the process and microstructure model inputs. To compare the predicted mean grain area to the measured value from the EBSD data in Figure 2b, TASMANIAN was used to plot the percent difference between simulated and measured values as a function of N_0 and ΔT_N for each of the five $D4\sigma$. These plots are shown in Figure 9, where shades of blue indicate ExaCA overestimating the experimental mean grain area and shades of red indicate ExaCA underestimating the mean grain area. For each $D4\sigma$, increasing N_0 and decreasing ΔT_N results in a decreasing mean grain area. There is also no unique combination of N_0 and ΔT_N that yield the best match with the experimental data; a “valley” of

equally accurate N_0 and ΔT_N exists for all spot sizes, though the precise location of the valley in nucleation parameter space varies depending on $D4\sigma$. Highlighted in Figure 9 for $D4\sigma = 100 \mu\text{m}$ are the locations in nucleation input space that produced the smallest (i), largest (iv), and most accurate compared to EBSD (ii-iii) mean grain areas. From the IPF-Z sections, it is observed that (i) results in a relatively weak texture while (ii-iv) feature textures of similar strengths, featuring large grains with $\langle 110 \rangle // Z$, $\langle 111 \rangle // Z$, and adjacent orientations along with smaller grains with $\langle 100 \rangle // Z$ orientations. This suggests that the development of strong texture in these simulations requires a certain amount of nucleation to force competition with the epitaxially growing grains from previous layers, though there is a limited amount of texture strengthening possible through this effect alone in the current model. Though N_0 and ΔT_N did have a modest effect on texture, Figure 7 showed that the predicted texture was most strongly controlled by melt pool geometry through variation in $D4\sigma$. We also note that with prior information on the ranges of inputs that produce the best match with the observed mean grain size from experiment, improved TASMANIAN sampling of the input space could be performed to produce a smoother representation of the output surfaces in Figure 9.

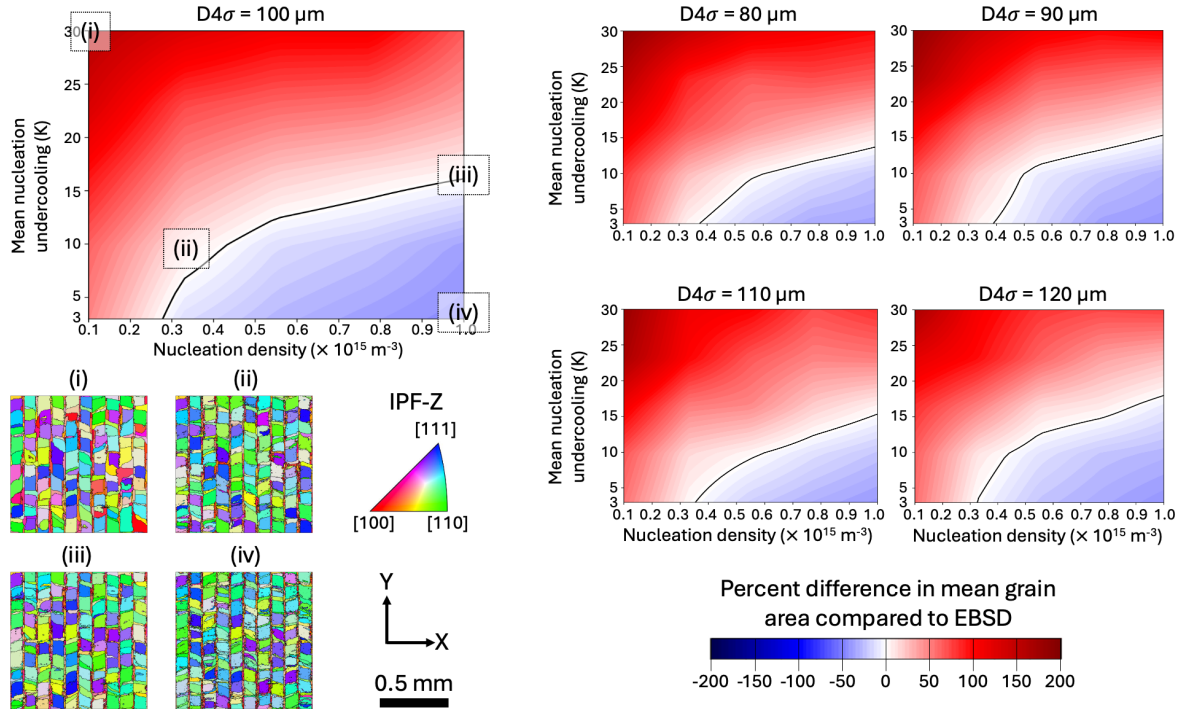


Figure 9: Response surfaces for the percent difference between the modeled and experimentally observed mean grain area as a function of nucleation parameter inputs N_0 and ΔT_N . These response surfaces are given for each of the five $D4\sigma$ values examined in this study, with IPF-Z colored XY cross-sections of the microstructure annotated for the $D4\sigma = 100 \mu\text{m}$ case. The solid contours represent locations in nucleation input space predicted by TASMANIAN to have the best agreement between simulations and measurements.

The primary components of the texture from the ExaCA predictions and EBSD data were the $\langle 100 \rangle$ alignment with Y and the $\langle 110 \rangle$ alignment with Z. To quantitatively compare the ExaCA predictions to both each other and to the EBSD data, Figure 10 plots the percent difference in MTEX-estimated multiples of uniform distribution (M.U.D) between the ExaCA predictions and EBSD measurements for the primary texture components as functions of N_0 and ΔT_N . This is performed at $D4\sigma = 100 \mu\text{m}$ and $110 \mu\text{m}$ to consider both sides of the observed texture transition as a function of $D4\sigma$ from Figure 7. While noise is present in the response surfaces, the strength of the $\langle 110 \rangle // Z$ and $\langle 100 \rangle // Y$ texture components generally increased with increasing N_0 and decreasing ΔT_N . For $D4\sigma = 100 \mu\text{m}$, ExaCA underestimates the strength of both $\langle 110 \rangle // Z$ and $\langle 100 \rangle // Y$ texture components across the nucleation input space. This underestimation is both larger (around 50% of the measured value) and more uniform across input parameter space for $D4\sigma =$

100 μm and $\langle 100 \rangle // Y$. For $D4\sigma = 110 \mu\text{m}$, a “valley” of predicted M.U.D $\langle 110 \rangle // Z$ values match the EBSD observation, and predicted M.U.D $\langle 100 \rangle // Y$ values at the largest N_0 approached the EBSD observation as well. Notably, while the $D4\sigma = 100 \mu\text{m}$ predictions underestimate the two texture components by 25-50% for nearly all permutations of N_0 and ΔT_N , nearly all of the $D4\sigma = 110 \mu\text{m}$ predictions fall within $\pm 25\%$ of the measured values.

Selecting simulation results that exhibited the best match for either M.U.D. $\langle 110 \rangle // Z$ or $\langle 100 \rangle // Y$ at both $D4\sigma = 100$ and $110 \mu\text{m}$ (labeled (i-iv) in Figure 10), Figure 11 shows the IPF-Y and IPF-Z colored grain structure maps and inverse pole figures for comparison to the EBSD data reproduced from Figure 2b. Qualitatively, despite underestimating texture strength, the square grain cross-sections in the $D4\sigma = 100 \mu\text{m}$ data (i, iii) better matched the EBSD data than the rectangular grain cross-sections at $D4\sigma = 110 \mu\text{m}$ (ii, iv). While not a perfect match for the texture observed experimentally (there is some deviation in features such as the $\langle 211 \rangle$ alignment with Y and the $\langle 210 \rangle$ alignment with Z), the $D4\sigma = 110 \mu\text{m}$ predictions well-reproduced the strength of both of the primary texture components.

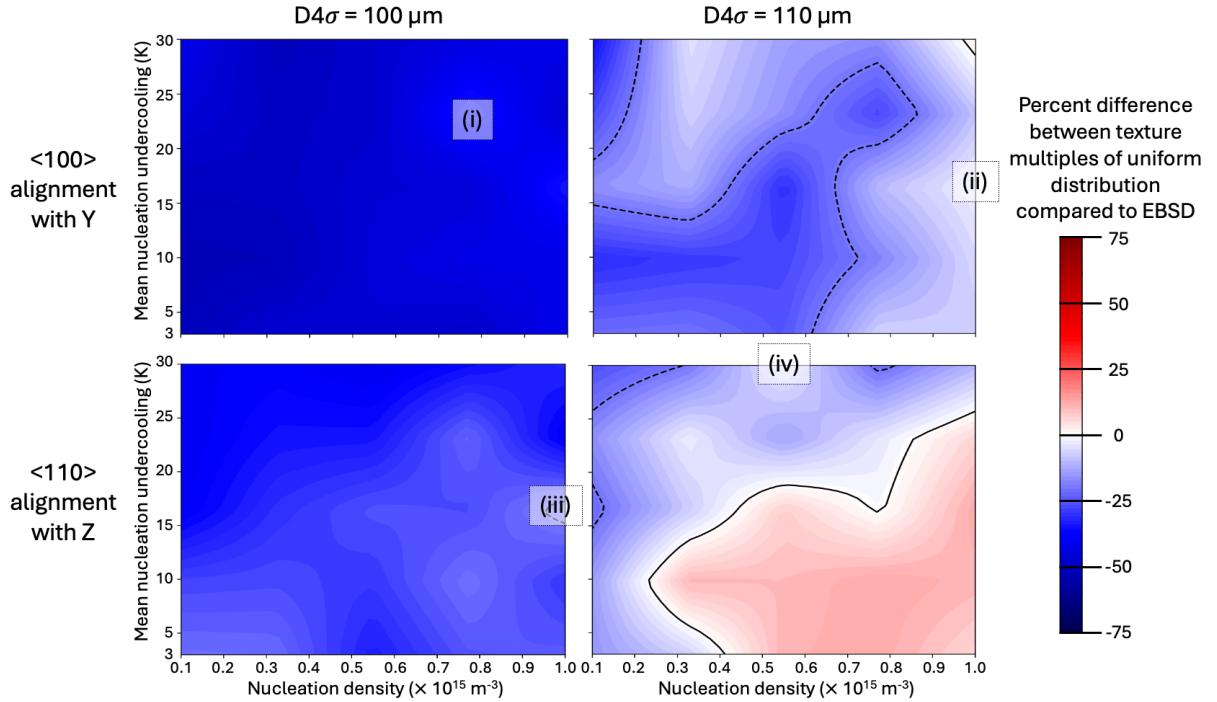


Figure 10: Response surfaces in nucleation parameter input space (N_0 , ΔT_N) for the percent difference between modeled and EBSD observed multiples of uniform distribution (M.U.D.) for the $\langle 100 \rangle // Y$ (top) and $\langle 110 \rangle // Z$ (bottom) texture components. EBSD values were taken from the XY cross-sections from Figure 2b, while the simulated values were from the XY cross-sections from Figure 7. Comparisons are performed for $D4\sigma = 100 \mu\text{m}$ (left) and $D4\sigma = 110 \mu\text{m}$ (right) simulations. Markers (i-iv) denote the nucleation inputs that produced the prediction with the best agreement with the experimental data for a given $D4\sigma$ and texture component. The solid contours represent locations in nucleation input space predicted by TASMANIAN to have the best agreement between simulations and measurements, while the dashed contours represent $\pm 20\%$ error between simulations and measurements.

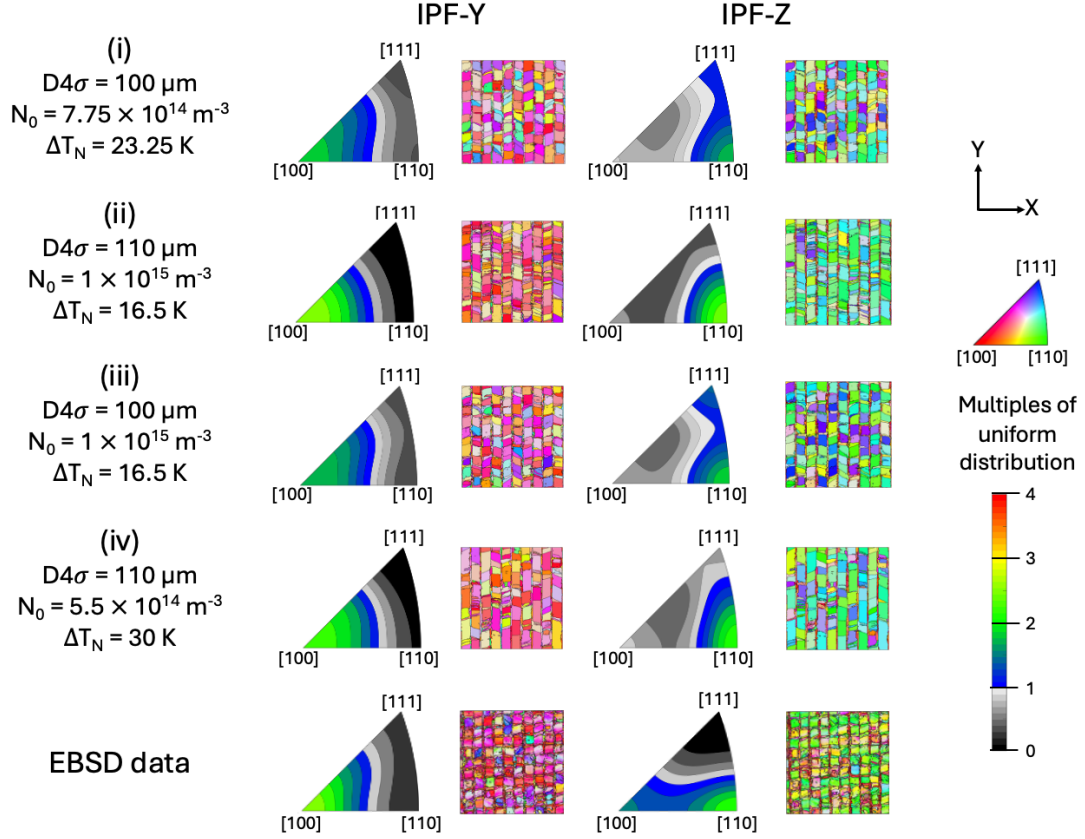


Figure 11: XY cross-sections of selected inverse pole figure-colored grain structure predictions and the associated inverse pole figures in the Y and Z directions, where markers (i-iv) are consistent with those from Figure 10 and the EBSD data is reproduced from Figure 2b for purposes of model comparison.

4.1.2 Properties

The ExaConstit simulation results shown and compared to the observed values in Figure 8 can be broken down by RVE, where each RVE was generated using different AdditiveFOAM and ExaCA input parameters. First, we investigate the effect of AdditiveFOAM spot size $D4\sigma$ on yield stress (YS) and engineering stress at 5% strain (S5) predictions for monotonic uniaxial compression loading in the X, Y, and Z directions. These results are shown in Figure 12, alongside measured values from Table 1 shown as dashed lines for the Y and Z loading directions. The larger predicted YS and S5 values for the Z loading direction compared to Y are immediately apparent, as is the similar mechanical response for the X and Z loading directions. From Figure 7, this observation can be rationalized as the predicted grain structures show similar crystallographic orientations aligned with the X and Z directions (primarily $\langle 110 \rangle$, $\langle 111 \rangle$, and other similar directions). The smaller YS and S5 for the Y loading direction, consisting primarily of $\langle 100 \rangle$ and similar crystallographic directions, is consistent with these directions being the most compliant orientation for FCC materials. Similar results have been shown in powder diffraction experiments done on AM Inconel 625 samples which elucidate how grains oriented in the $\langle 100 \rangle$ crystallographic fiber direction carry the most elastic strain during the sample deformation [Knox'2025]. For both loading directions, the predicted properties at $D4\sigma = 100 \mu\text{m}$ (the reported spot size for the AMBench-2018-01 builds) tended to be the closest to the measured properties, though underestimation of YS and S5 for the Y loading direction was still present. Figure 12 also shows a significant change in the predicted YS and S5 distributions between $D4\sigma = 100$ and $110 \mu\text{m}$, a change that is again reflected in the predicted textures from Figure 7.

Next, we examine the effect of the ExaCA input parameters - nucleation density N_0 and mean nucleation undercooling ΔT_N on the ExaConstit predictions for YS and S5. The $D4\sigma = 100 \mu\text{m}$ response surfaces in

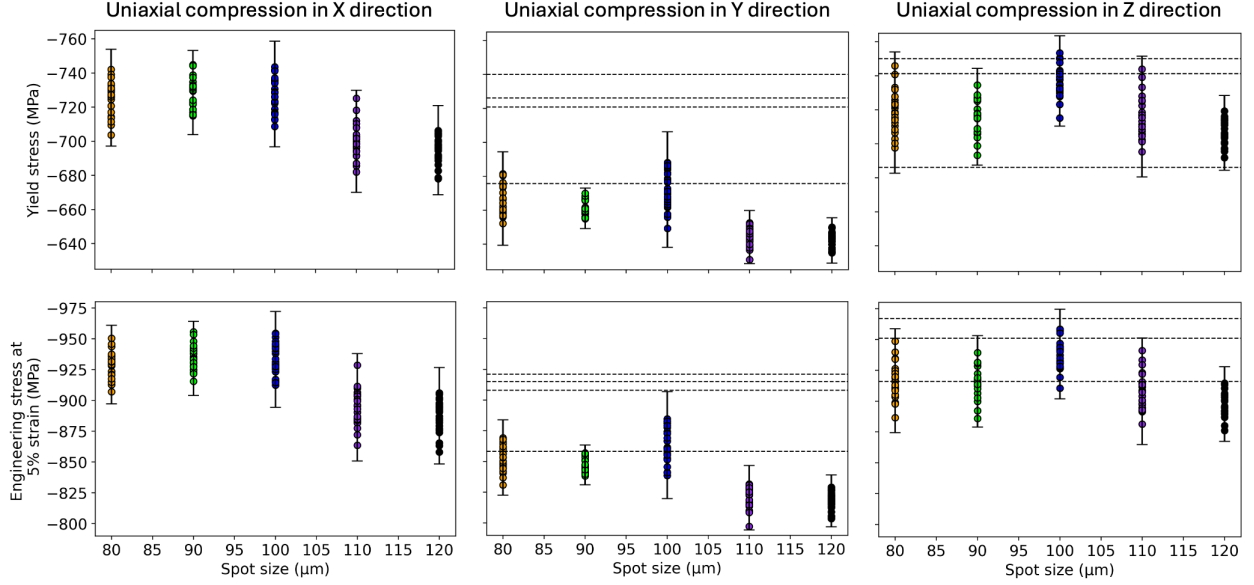


Figure 12: ExaConstit predictions of yield stress (top) and engineering stress at 5% strain (bottom) for monotonic uniaxial compression loading along each cardinal direction, broken down by the AdditiveFOAM spot size $D4\sigma$ used to generate each RVE. Error bars represent three standard deviations from the mean for each $D4\sigma$, and dashed lines for the Y and Z loading directions represent the experimental measurements from Table 1.

nucleation input parameter space for the percent error between the magnitude of the mean measured values from Table 1 and the magnitude of the simulated YS and S5 are plotted in Figure 13. While both YS and S5 vary within nucleation input parameter space for a given loading direction, trends with respect to the input parameters are both subtle and noisy. There was relatively little variation in percent error for the simulated YS and S5 as a function of nucleation input parameters, mainly between -5 and 5% for the Z loading direction quantities and between 5 and 10% for YS and S5 for the Y loading direction quantities. The magnitudes of the predicted YS and S5 generally appear to slightly increase as N_0 is increased and ΔT_N decreased - particularly for the Y loading direction, where large N_0 and small ΔT_N lead to more significant strength overestimation. It is possible that these subtle trends with N_0 and ΔT_N are related to the trends towards smaller grain size and slightly stronger texture with increasing N_0 and decreasing ΔT_N shown in Figure 9 and Figure 10. Given the noisy nature of the texture strength predictions with respect to the EBSD as shown in Figure 10, it is difficult to draw definitive conclusions on the effect of nucleation input parameters on properties. However, the results from Figure 12 and Figure 13 suggest that the predicted properties and their anisotropy are more sensitive to texture variability than grain size variability for the range of uncertain inputs examined in the present work.

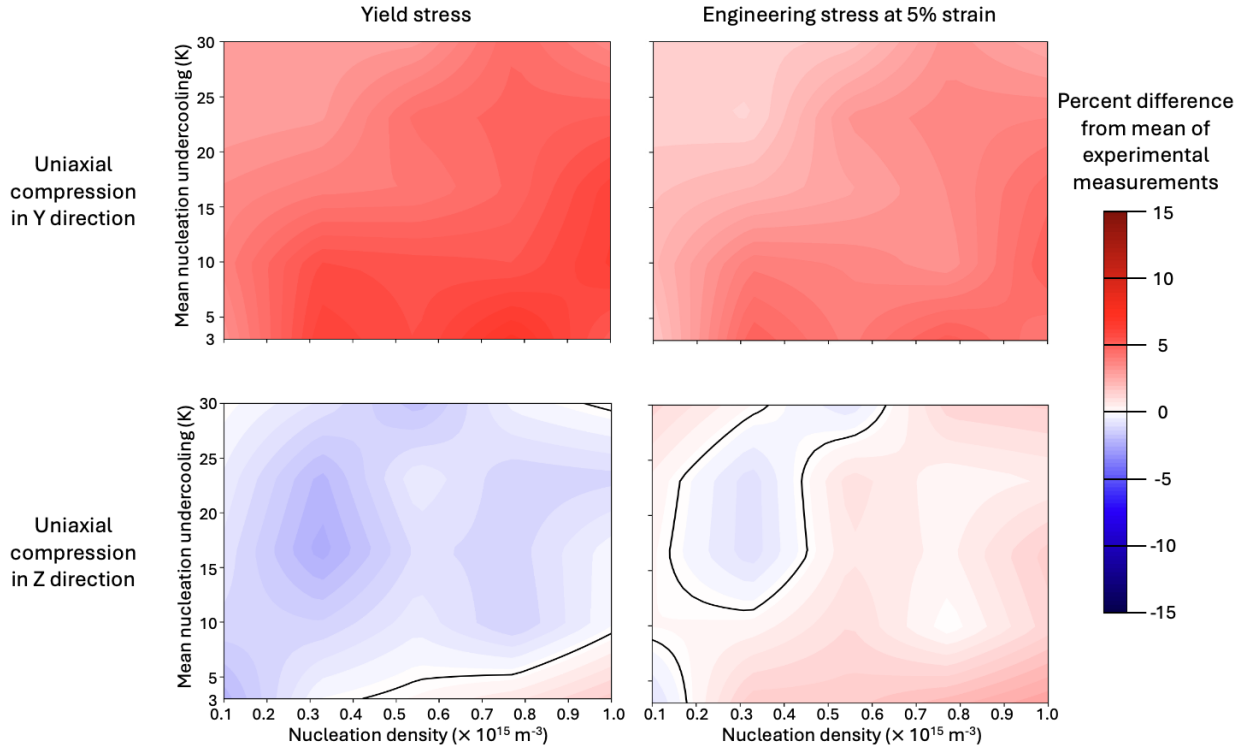


Figure 13: Percent error between the magnitude of the mean measured and ExaConstit predicted strength values in nucleation parameter input space for a fixed spot size $D4\sigma$ of $100\mu\text{m}$. Results shown are for yield stress (left) and engineering stress at 5% strain (right) for monotonic uniaxial compression loading along the Y (top) and Z (bottom) loading directions. The solid contours represent locations in nucleation input space predicted by TASMANIAN to have the best agreement between simulations and measurements.

4.2 Compute resource utilization

The calculations performed for the ExaAM Frontier run represent what is among the largest published PSP AM simulations, both in terms of resources utilized and output dataset size. The AdditiveFOAM, ExaCA, and ExaConstit resource utilization output from RADICAL-EnTK is shown in Figure 14. Here, resource utilization measures the fraction of allocated resources that RADICAL components or computing tasks actively employ during execution. The “bootstrap” state denotes the startup period for components of RADICAL-EnTK (workflow management system) and RADICAL-Pilot (runtime system). The “exec setup” state refers to the task preparation for the execution, including sandbox creation and launching scripts generation; this duration may vary depending on the underlying file system. The “running” state captures the period during which tasks execute on assigned resources, and the “idle” state corresponds to periods when resources are allocated but not actively used for computation. This discussion is centered around the time spent in the “running” state, which occupied the vast majority of the simulation time.

Each of the 10 AdditiveFOAM runs (consisting of simulations of an odd numbered and an even numbered layer raster pattern simulation for each of the 5 $D4\sigma$ values) used 4 nodes (224 MPI ranks per run) of Frontier. As shown in Figure 14a, AdditiveFOAM utilized only the CPUs on these 40 nodes, completing the simulations in around 40 minutes. The runtime for these simulations varied with some runs taking as few as 20 minutes depending on the melt pool size and whether the odd or even layer scan path was used. The bulk of the runtime was during execution of ExaCA and ExaConstit. It is shown in Figure 14b that the ExaCA runtimes were heavily dependent on the number of solidification events in the AdditiveFOAM data, with the runs corresponding to the smaller melt pools ($D4\sigma = 110$ and $120 \mu\text{m}$) completing in under 2 hours and the runs corresponding to the largest melt pools ($D4\sigma = 80 \mu\text{m}$) taking up to 2 hours and 40

minutes. ExaCA fully utilized all GPUs per node, where each ExaCA simulation was executed on a single node (8 MPI ranks per run, 1 GPU per MPI rank). The vast majority of the node hours were used during the ExaConstit simulations - the CPU and GPU on each of the 8000 nodes were heavily utilized for around 2 hours to complete the ensemble of runs for various loading conditions on each RVE. Job management was performed such that 1000 concurrent ExaConstit simulations, each utilizing 8 nodes of Frontier (64 MPI ranks per simulation, 1 GPU per MPI rank), were performed at a time to maximize utilization of the available resources. Overall, the entirety of the demonstration run was performed in 5 to 6 hours, with the first stage (AdditiveFOAM and ExaCA) using smaller portions of Frontier and the second stage (ExaConstit) utilizing nearly the entire machine.

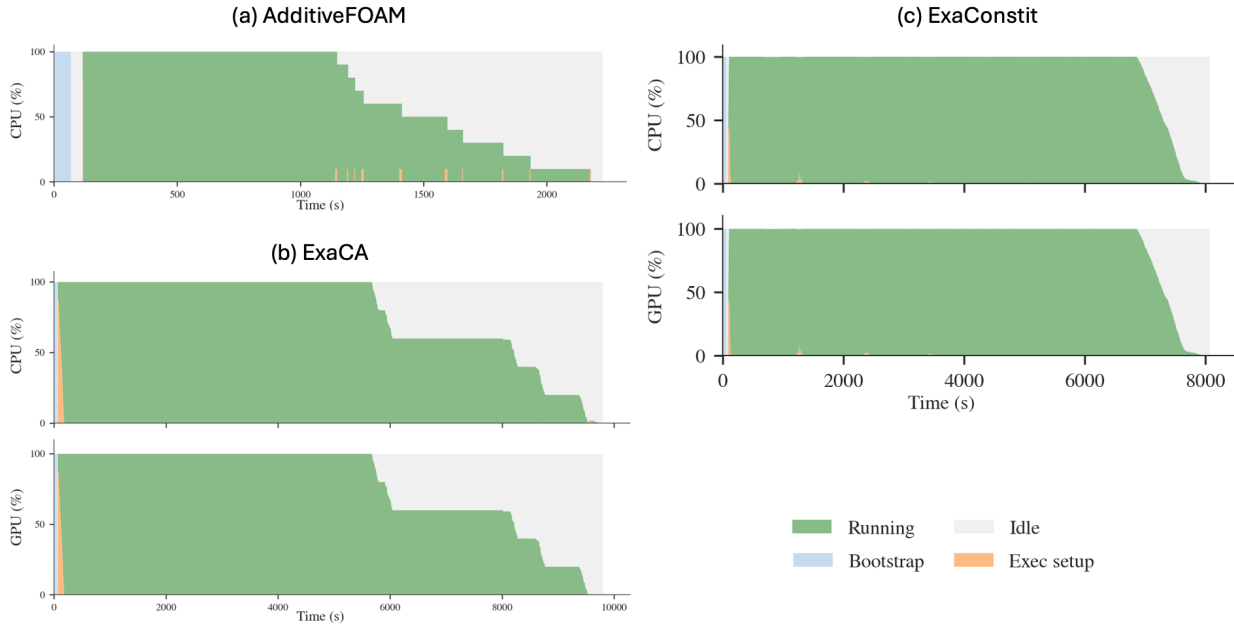


Figure 14: CPU and GPU resource utilization for each of the three PSP modeling components: (a) Additive-FOAM, (b) ExaCA, and (c) ExaConstit during the demonstration run on Frontier. AdditiveFOAM utilized the CPUs, while ExaCA and ExaConstit used both CPUs and GPUs. Resource utilization is broken down by job state, where “bootstrap” refers to the startup period for the RADICAL components, “exec setup” refers to task preparation, “running” refers to task execution, and “idle” refers to periods where resources were not actively used for computation.

Based on the resource utilization of the codes, there are several areas for potential improvement for running the coupled workflow and more efficiently utilizing Frontier resources. While AdditiveFOAM consumed the smallest amount of wall clock time of the three workflow components, it is also the only code in the ExaAM workflow that cannot use the GPU. Replacing this with a GPU-capable code, such as the finite difference-based Finch software currently under development, would allow more efficient resource utilization and shorter run times [finch]. Due to the fine grid size required for the CA model compared to that required to resolve the temperature fields, implementation of a process model with adaptive mesh refinement and/or in-memory coupling to ExaCA for temperature interpolation could reduce both computational and memory costs. ExaCA utilizes the GPUs and is highly performant for large ensemble jobs, but individual ExaCA AM runs do not strong scale well. This is a known issue with domain decomposition-based parallelization of AM solidification models; as most of the calculations to be done on a given time step are located in and around the melt pool, which in turn represents only a small fraction of the overall domain [Lian’2018, Rolchigo’2022]. One potential solution for this scaling bottleneck is a discrete event approach, where as an alternative to applying rules and updating cell states on a time step by time step basis, the system is reformulated as a series of cell transition events to be ordered in time [Nutaro’2023]. Such optimizations for tighter process-microstructure model coupling, improved GPU utilization, and improved strong scaling

would allow future simulations that use considerably fewer node hours and allow an increase in scope for the ranges and sampling of inputs considered.

ExaConstit, the most computationally expensive code used in the workflow, used over an order of magnitude more resources than the others, though a similar amount of wall clock time as ExaCA. One way of reducing this amount of work is to determine the minimum number of elements required to produce an RVE. The work performed in [Carson’2023] showed that various coarsening levels were able to produce less than 5% root mean square relative error (RMSRE) compared to the experiments across all simulated microstructures. However, the coarser microstructures led to distributions more skewed towards larger errors. Determining the optimal coarsening level for the desired level of precision for the properties of interest will be critical to determining the number of resources required to run these large ExaConstit simulation ensembles. Furthermore, streamlining the RVE data transfer between ExaCA and ExaConstit to minimize file size may further optimize the process of taking the regular grid representation from ExaCA and remeshing it into the coarser finite element representation used by ExaConstit.

4.3 Re-evaluation of modeling approximations and uncertainty sources

Each model in the ExaAM workflow, and each measurement used for calibration, has some inherent uncertainty due to aleatoric and epistemic sources. AdditiveFOAM showed excellent agreement with the AM-Bench-2018-02 calibration data [Coleman’2024], but additional complications arise when applying the calibrated model to the AMBench-2018-01 simulations. At the time of the simulations, melt pool shapes and dimensions from the odd and even numbered layers of the AMBench-2018-01 test artifacts were not available and therefore AdditiveFOAM could not be directly validated for the workflow demonstrated in this study. However, given the high sensitivity of microstructure predictions to melt pool shape and dimensions [Rolchigo’2024], sources of error and uncertainty in the melt pool predictions merit further discussion. One such notable source is the use of bare plate laser track dimensions for calibration, which will lead to some error in the prediction of melt pools with powder present. Additionally, while the dynamic heat source in AdditiveFOAM accurately captured melt pool dimensions on both sides of the conduction to keyhole mode transition without explicitly modeling fluid flow, the predictions consisted of artificially ellipsoidal melt pool shapes. AdditiveFOAM was also calibrated to mean observations of melt pool dimensions for single track experiments on bare plates, despite the fact that a range of dimensions were observed for a given $D4\sigma$. The fact that the small $D4\sigma$ results exhibited a better agreement with grain shape, and the larger $D4\sigma$ results exhibited a better agreement with texture, suggests that the melt pool dimensions within the AMBench-2018-01 leg interiors may fall between those predicted for $D4\sigma = 100$ and $110 \mu\text{m}$ or oscillate over a range of dimensions to produce microstructures with features of both predictions.

In addition to error in the input data from AdditiveFOAM, several ExaCA model approximations may have also lead to the observed discrepancies with the experimental microstructure data. The most notable of these discrepancies was the overprediction of the frequency of grains with $\langle 111 \rangle$ and $\langle 221 \rangle // Z$ crystallographic orientations and the underprediction of the overall texture strength and the number of grains with $\langle 100 \rangle // Z$ crystallographic orientations. Though the cell size of $1.25 \mu\text{m}$ was sufficient for the prediction of some grain competition, the competition between grains of similar orientations (such as $\langle 111 \rangle$, $\langle 221 \rangle$ and $\langle 111 \rangle$ alignments in a given direction) is highly sensitive to cell size when large thermal gradients are present [Rolchigo’2025]. As smaller cell sizes are not necessarily feasible given compute and memory limitations, algorithmic improvements to better predict the grain competition may be required for improved prediction accuracy. Some inputs such as the fitting parameters used in the interfacial response function were not varied in the present work and calibration of these may affect the microstructure predictions. Other fixed inputs have inherent variability not accounted for in this work as well; one such example is the layer height, which is expected to have some statistical variation during a build due to the heterogeneous nature of the powder feedstock and may potentially affect grain structure evolution. While the single XY cross-section used for validation (Figure 2b) appeared to be a reasonable representation of the bulk microstructure based on the larger area cross-sections in Figure 2c and data from other Z locations from Levine et al. [Levine’2023], calibration or validation in the future should include other cross-sections and ideally other cross-section orientations (i.e., XZ and YZ) to avoid overfitting model inputs.

Despite the observed discrepancies between the various CA simulation results and the EBSD data, the constitutive properties modeled using the 125 microstructures were relatively consistent with the experimen-

tally measured values, particularly for loading in the Z direction. Some discrepancies between the model and experiment may be due to the fact that a relatively untextured RVE was used for ExaConstit calibration; redoing this calibration using microstructures more representative of AM textures (stronger $\langle 100 \rangle$ and/or $\langle 110 \rangle$ fiber alignment with the loading direction) may improve model accuracy. As the simulations predicted a stronger texture with respect to Y than Z, error induced by using a weakly textured calibration RVE may be more significant for Y loading direction predictions. Additionally, while the ExaConstit RVE size of 0.5^3 mm was selected to ensure the calculations were tractable given the number of finite elements per RVE, the number of grains necessary to produce a fully representative portion of the predicted microstructure is not known. This may be particularly true for Z direction loading as Figure 5a shows that many grains may extend through the entirety of the RVE in the Z direction; inability to capture the true grain size distribution in the Z direction may have affected model prediction trends with respect to $D4\sigma$ for Z direction loading from Figure 12. This sample size error, in conjunction with the noise in ExaCA texture predictions as a function of nucleation input parameters and ExaCA’s tendency to underpredict the expected texture, may have also caused the lack of property sensitivity to ExaCA nucleation input parameters. Future studies should consider larger representative volumes to better represent the grain structure statistics and reduce noise in property prediction trends, potentially at higher coarsening levels to keep the crystal plasticity simulations tractable. A comparison between RVE constitutive behavior and experimental measurements that includes microstructure coarsening effects can be found in Carson et al [Carson’2023]. Additionally, some error can be attributed to the use of a single experimental dataset for calibration, and use of additional experimental datasets could enable a proper Bayesian calibration of ExaConstit and tighten the posterior distribution of the fitted model inputs.

Another potential source of error is the approximation made by both ExaCA and ExaConstit of a microstructure without intragranular crystallographic misorientations; grains are assumed to be perfect crystals without defects and without residual elastic strain that might exist in the material’s unloaded state. This is a common assumption for crystal plasticity models and is typically reasonable for microstructures with little intragranular heterogeneity. However, such intragranular misorientations are present not only in the AMBench-2018-01 microstructure data [Stoudt’2020] but have also been observed in both electron and laser-based AM microstructures [Fernandez-Zelaia’2021, Zhang’2023] as a result of the rapid solidification of the material and the accumulation of elastic strain during the build process. These local crystallographic orientation variations are expected to lead to noticeable spread in the response of the material’s macro scale stress-strain curves. While intragranular heterogeneity was neglected in the present modeling workflow due to the complexity of including such effects, doing so in the future may enable more accurate property predictions.

4.4 Opportunities for future use

While the demonstration in this work focused on a single location in an LPBF-fabricated part using a single material, the workflow is easily extensible to LPBF of other parts and materials. Provided thermophysical properties and melt pool dimensions from single laser tracks, repetition of the AdditiveFOAM heat source calibration enables AdditiveFOAM-ExaCA simulation ensembles for alternative materials, scan paths, and part geometries. ExaCA showed that a range of nucleation input parameters may provide a match with experimental data, and future work should consider additional outputs of interest to further constrain the input calibration. For individual AdditiveFOAM-ExaCA simulations, nucleation parameter calibration is not necessarily required to obtain a reasonable qualitative match with microstructures from experiment and the nucleation inputs from this study may be transferable to others, particularly other builds using Inconel 625. While ExaCA’s algorithm limits the workflow to materials that solidify with a cubic crystal geometry, this captures the vast majority of materials printed using LPBF, such as ferritic and austenitic stainless steels and Ni-based alloys. ExaConstit support has been recently added for multiple materials as well, enabling property simulations of individual ExaCA RVEs as well as ensembles of simulations using a series of ExaCA RVEs. Based on the demonstrated ability to rapidly and accurately predict PSP relationships, future simulation ensembles using ExaAM software will likely play key roles in future experiment design, enabling downselection of process parameter input space to produce parts with optimal mechanical properties or microstructural features.

The PSP insights from these ExaAM simulations can also be used to guide future model development

and simulation campaigns. While the ExaAM workflow considered much of the key physics governing AM heat transport, solidification grain structure, and constitutive behavior, it does not include all aspects of the PSP relationships. However, the workflow has been structured such that existing components could be modified to account for additional physics, and new components can augment the current capabilities. One example of a potential workflow augmentation is a thermomechanical model, as thermomechanical stress during the build has been linked to the intragranular crystallographic orientation gradients in AM grain structures [Pistor’2021]. Addition of a thermomechanics model to the ExaAM workflow would enable prediction of these intragranular misorientation gradients and potentially lead to improved microstructure and property predictions. Adding features to the existing software may also improve prediction accuracy. For example, improvements to the nucleation portion of the CA algorithm have been shown to yield strong texture development in CA simulations in the literature [Xie’2023, Koepf’2023], and such updates along with other modifications to the grain envelope geometry will be considered in the future. Improved prediction accuracy is also possible, for example, through explicit consideration of fluid flow in AdditiveFOAM or refinement of the cell size used by ExaCA. However, such changes to the workflow should be weighted against the computational cost necessary to achieve such accuracy improvements as small improvements may not be worth a large increase in required resources or may not be necessary for a reasonable representation of the microstructural feature or mechanical property of interest.

5 Conclusions

Additive process-structure-property (PSP) relationships are complex and span multiple length and time scales, but modeling these relationships in a computationally efficient manner is necessary for part and process qualification and AM process optimization. The ExaAM software suite addresses this need with open source models of the various relevant physics during AM builds, coupled using autonomous workflows allowing critical insight into the PSP relationships. This work showcases an example of the ExaAM suite’s modeling capabilities with a validation of grain structure and constitutive property prediction against benchmark data, using appropriately calibrated input parameters to the heat transport model AdditiveFOAM, the microstructure model ExaCA, and the crystal plasticity model ExaConstit. The workflow showcased the ability to leverage exascale computing hardware, with the most computationally intensive model utilizing 8000 total nodes as part of an ensembles of simulations accounting for input uncertainty. The rich dataset of predicted microstructure and properties was obtained in under 6 hours of compute time, a fraction of the time that would be required with existing, CPU-only software. Furthermore, the models accurately reproduced distributions of grain structure features (texture, grain area, grain shape) and mechanical properties (yield stress and engineering stress at 5% strain) from the AMBench-2018 experiment and measurement suite. This demonstrated ability to rapidly and accurately simulate PSP relationships accounting for the uncertainty in input parameters is critical for utilizing modeling workflows to guide parameter selection for producing reproducible and defect-free parts with AM.

In addition to the successful benchmarking of model predictions against experimental data and the demonstration of the ExaAM workflow on a large portion of Frontier, the simulations performed in this study also provided insight into the fundamental relationships between processing and microstructure, as well as microstructure and properties. The predicted grain shapes and texture from ExaCA were found to be highly sensitive to melt pool size and shape, particularly the differences in melt pool size and shape between different layers of the AdditiveFOAM simulations. Simulations where the grain structure from the odd numbered layers was entirely remelted during even numbered layers showed a better match for the experimentally observed texture; however, simulations where both the odd and even layers contributed to the final grain structure showed a better match with the experimentally observed grain shape distribution. While larger nucleation density and smaller nucleation undercooling yielded smaller grain areas in the ExaCA simulation results, it was found that ranges of possible nucleation inputs could yield very similar mean grain areas and textures. Further, the predicted properties using ExaConstit crystal plasticity simulations and representative volume elements (RVEs) from the ExaCA simulations yielded a distribution of properties that well-matched the observed properties from various experiments. ExaConstit property predictions were found to be more sensitive to AdditiveFOAM spot size used to generate a given ExaCA RVE, with a weaker dependence on the ExaCA nucleation input parameters associated with a given RVE. This finding will require further study

to ensure that the 0.5^3 mm microstructures used by ExaConstit are truly representative of a region’s grain structure, as trends in ExaCA texture prediction as a function of nucleation input parameters were often noisy. These insights would not have been possible without such a large ensemble of PSP calculations, which in turn would not have been possible without the ExaAM software suite and utilization of Frontier. These results set the stage for future simulations accounting for insights gained from this study, addressing sources of error believed to be responsible for quantities for which the model predictions were less accurate (such as the transverse direction constitutive response of the model RVEs) and for further uncertainty quantification of results to understand the relative importance of other uncertain inputs in AM processing.

Acknowledgements

Work was performed in part by Lawrence Livermore National Laboratory, Oak Ridge National Laboratory, and Brookhaven National Laboratory under contract DE-AC52-07NA27344 and was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. DOE Office of Science and the NNSA. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

The authors would like to thank Naren Raghavan, Sarah Graham, Andrés Marquez Rossy for EBSD data collection; the ExaAM team for useful planning discussions; and Adam Creuziger and Lyle Levine for their guidance in AMBench data interpretation.

Appendix: CA input parameter estimation and preliminary simulations

The ExaCA simulations performed as part of the ExaAM workflow used fixed values for the mean substrate grain size S_0 , interfacial response function $V(\Delta T)$, the cell size Δx , and the time step Δt . Here, we provide the rationale for using fixed values for each of these parameters, as well as the selection of values for these fixed parameters.

- S_0 : From additional EBSD data of the AMBench-2018-01 baseplate shown in Figure A1a (in turn reproduced from Levine et al [Levine’2023]), the decision was made to approximate the baseplate as equiaxed using a Voronoi scheme and randomly assigning one of 1,000,000 unique crystallographic orientations to each grain. From a rough analysis of the EBSD grain sizes, an S_0 value of $12.3 \mu\text{m}$ was selected. However, the EBSD data taken from the bottom of the part shows a distribution of grain sizes. Preliminary simulations using S_0 values of 5.3 and $19.3 \mu\text{m}$, in addition to the $12.3 \mu\text{m}$ simulations, were performed using AdditiveFOAM data at three $D4\sigma$ values to gauge the effect of different substrate grain sizes on the microstructure at the location of interest 1 mm above the baseplate top. As shown in the IPF-Z-colored cross-sections in Figure A1b, negligible differences in grain area were found across S_0 , particularly with respect to the significant difference in grain cross-sectional area and shape with varied $D4\sigma$. The minimal impact of S_0 on grain area and shape evolution is further corroborated by Rolchigo et al. [Rolchigo’2022b], where it was found that the grain areas 1 mm above the baseplate top are comparable when initializing simulations with similar S_0 values. The inverse pole figures in Figure A1c show that while the texture does vary with the selected value for S_0 , the differences are relatively minor compared to the difference in texture due to $D4\sigma$ variation. Additionally, no clear texture trend with S_0 is visible. Due to the lack of sensitivity of predicted grain area and texture to S_0 , variations in S_0 were not considered for the ensemble of CA simulations performed in the present work.
- $V(\Delta T)$: The interfacial response function relates the solidification velocity V in a CA cell to the undercooling ΔT in the cell. In the absence of new grain nucleation, Baumard et al observed minor differences in grain area predictions with varied $V(\Delta T)$ [Baumard’2021]. When nucleation is included in the simulations, $V(\Delta T)$ has a more significant effect on the grain size by changing the competition between nucleation and solidification via slowing or accelerating the early stages of solidification near

the liquidus isotherm [Koepf2023]. $V(\Delta T)$ has also been observed to change the contour of the solidification front in the melt pool tail, leading to variation in the grain structure along melt pool centerlines and at the top surface of melt pools [Xue2022]. Given that varying ΔT_N already varies the competition between nucleation and solidification and that the grain structure near the top surface of melt pools is primarily remelted during additional layers (and likely has only subtle impact on the final as-solidified texture prediction, which was noted by Xue et al [Xue2022]), a static $V(\Delta T)$ taken from a prior CA study using a Ni-base superalloy is used in the present work [Rai2016]. The effect of $V(\Delta T)$ on texture selection from the substrate in multilayer builds in the absence of nucleation parameter variation will be investigated in future work.

- Δx : A wide range of values for Δx have been used in CA studies of AM solidification, ranging from 1 μm [Akram2018] to 10 μm [Koepf2023]. In a previous study using ExaCA, 1.25 μm was selected as the cell size based on grain area convergence for a small simulation [Rolchigo2020] but no comprehensive investigation of the effect of Δx on texture development for large multilayer AM simulations has been performed. As a result, the 1.25 μm from the previous study (which to the best of the authors' knowledge, is the finest cell size used in a 3D CA model of multilayer AM solidification) is again used here.
- Δt : Generally, Δt is either adaptive to ensure that solidification over a given time step is less than some fraction of the cell size, or conservatively selected in a way that such that solidification will always be less than some fraction of the cell size per time step. As was the case with selection of Δx , we select a fixed Δt on the conservative side of the values that have been used in the literature. As the region of the simulated part is not at the start or end of scan paths, it can be assumed that the fastest solidification velocity expected in the melt pools is equal to the scan speed of 0.8 m/s. With $\Delta x = 1.25 \mu\text{m}$, and ensuring that the growth rate of the decentered octahedra in the CA algorithm does not exceed 1/25th of the cell spacing, a Δt of 0.0625 μs was selected.

These were performed at $\Delta T_N = 16.5 \text{ K}$ (midway through the 3 to 30 K range selected for study).

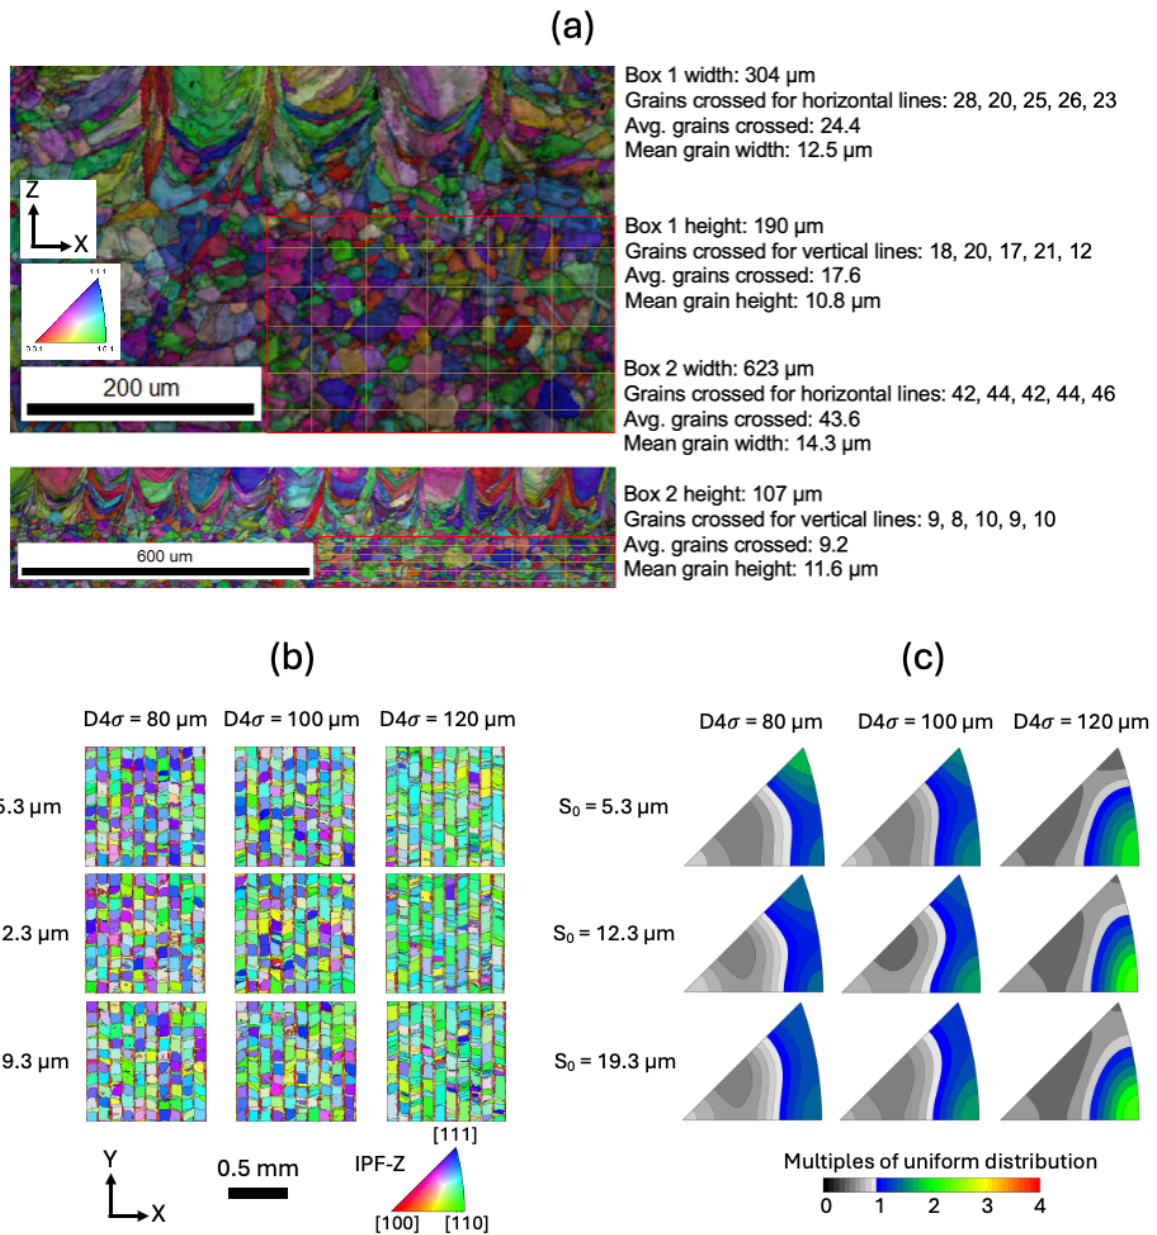


Figure A1: (a) EBSD data of the bottom of an AMBench-2018-01 specimen including the baseplate grain structure, colored using the IPF-Z map [Levine'2023]. (b) ExaCA simulation results perpendicular to Z and taken 1 mm above the baseplate top from simulations using three mean substrate grain size (S_0) values and three spot sizes ($D4\sigma$), colored using the IPF-Z map. (c) The inverse pole figures corresponding to the cross-sections from (b)