

IMECE2024-144919

PREDICTING MICROSTRUCTURE-PROPERTY OF SILICA AEROGEL MATERIALS VIA BAYESIAN CONVOLUTIONAL NEURAL NETWORKS SURROGATE MODEL

Md Azharul Islam

Department of Mechanical
and Aerospace Engineering
University at Buffalo
Buffalo, New York, 14260
Email: mdazharu@buffalo.edu

Dwyer Scout Deighan

Department of Computational
and Data Enabled Sciences
University at Buffalo
Buffalo, New York, 14260
Email: dwyerdei@buffalo.edu

Danial Faghihi

Department of Mechanical
and Aerospace Engineering
University at Buffalo
Buffalo, New York, 14260
Email: danialfa@buffalo.edu

ABSTRACT

Deep neural networks have become essential for developing data-driven surrogate models of complex multiscale and multi-physics simulations. Trained with high-fidelity simulation data, these surrogate models enable computational predictions with significantly reduced time and resources compared to physics-based simulations. Surrogate models based on Convolutional Neural Networks (CNNs) are emerging as powerful tools for learning the complex relationships between microstructural images and the corresponding macroscopic properties of materials, facilitating tasks once computationally prohibitive, such as optimal design and synthesizing materials with target properties. However, the common neural network training method, relying on maximum likelihood parameter estimation, limits CNNs' ability to handle uncertainty due to sparse and limited high-fidelity data generated by physics-based simulations. This often leads to overfitting and overly confident predictions, compromising the reliability of CNNs, especially in high-consequence tasks such as model-based material design. This contribution proposes a Bayesian CNN for surrogate modeling by treating the network's training as statistical inference to overcome the formidable challenge of uncertainty assessment in predictions provided by neural network-based models. We employ Variational Inference to introduce probability distributions over the CNN's weights, ensuring accurate uncertainty estimation. The proposed Bayesian CNN surrogate model is applied to learn microstructure image-mechanical property relations in silica aerogel porous materials, known for its superior insulation properties but suffer from low

mechanical strength. Training data is obtained from elastic deformation simulations of the solid phase in the porous materials governed by stochastic partial differential equations. Results demonstrate the effectiveness of the Bayesian CNN in predicting the strain energy corresponding to a given microstructure image while considering confidence levels in predictions. The impact of training data points on prediction accuracy and reliability is also investigated using Bayesian CNN.

INTRODUCTION

In recent years, the simulation-based design of microstructures with the aim of achieving enhanced material properties has gained significant attention [13]. Despite notable advancements in image processing techniques and the development of physics-based models to predict microstructure-property relationships [3, 4, 6, 14, 22], the utilization of these tools often results in computational bottlenecks in microstructure design, primarily due to the necessity for iterative evaluations of the forward model. To overcome the computationally prohibitive microstructure design, the construction of surrogate models based on physics-based simulations has emerged as a pivotal solution. After training with data obtained from physics-based simulations, these surrogate models allow for computational predictions with significantly reduced time.

The emergence of deep learning methodologies has led to a paradigm shift in surrogate modeling of high-fidelity physical simulations, facilitating numerous computationally prohibitive

tasks such as optimization, design, and uncertainty quantification. Among these methodologies, Convolutional Neural Networks (CNNs) have emerged as a prominent class, offering robust capabilities in capturing the complex relationships between microstructural images and the corresponding macroscopic properties of materials. CNNs are particularly favored for microstructure-property surrogate modeling due to their ability to autonomously extract features from image-like inputs using convolution operations, thereby facilitating the linkage of local features within an image to associated properties. For instance, Yang et al. [25] developed a CNN-based model to establish structure-property linkages for high-contrast elastic 3-D composites, demonstrating superior predictive performance compared to simplified physics-based approaches. Shishir et al. [16] employed a CNN-based surrogate model of molecular dynamics to predict fracture thoroughness of polycrystalline graphene from microstructure images. In a similar vein, Mann et al. [10] introduced a novel CNN architecture for capturing complex microstructure-property relationships in high-contrast composite materials that significantly reduces the number of trainable parameters by eliminating fully connected layers and exclusively leveraging 2-point spatial correlations of the microstructure as input.

Despite advancements in deep learning, a fundamental challenge in surrogate modeling stems from the scarcity and uncertainty of training data derived from high-fidelity physical simulations. Conventional training methods, relying on maximum likelihood estimation of parameters, are susceptible to overfitting, wherein the model fits to noise in the training data rather than capturing underlying patterns, thereby resulting in overly confident predictions. In scientific domains such as simulation-based material design, which informs critical decision-making processes, the compromised reliability and trustworthiness of surrogate model predictions lead to significant implications. Addressing this challenge involves framing deep learning training as Bayesian inference [2, 11], which enables the representation of parameter uncertainty through probability distributions resilient to overfitting, facilitating learning from small datasets and enabling the characterization of surrogate model predictions, e.g., [12, 19]. However, the application of Bayesian inference to CNNs has been limited due to the substantial number of parameters inherent in this class of deep learning models. Zhu et al. [26] introduced a novel approach for uncertainty quantification and propagation in systems governed by stochastic partial differential equations (PDEs) using deep convolutional encoder-decoder networks. They employed a variational gradient descent algorithm based on Stein's method to extend approximate Bayesian inference to deep convolutional networks with millions of uncertain parameters. Their surrogate model effectively quantifies flow in heterogeneous media, despite the absence of shared underlying structures between input (permeability) and output (flow/pressure) fields. Furthermore, Shridhar et al. [17] utilized

a variational inference method to learn the posterior weight distribution of CNN parameters and propagate this uncertainty in classification tasks across various datasets, including the MNIST [8], CIFAR-10, and CIFAR-100 datasets.

This contribution introduces a Bayesian CNN for surrogate modeling under limited and uncertain training data to predict material properties given a microstructure image. The Bayesian CNN model is applied to learn microstructure-property relations in silica aerogel, a class of porous materials with superinsulation properties but low mechanical strength. The training data are obtained from the elastic deformation of the solid phase in the silica aerogel materials, where the samples of the stochastic microstructural patterns are derived from another generative model trained using the lattice Boltzmann simulation of the foaming process for synthesizing the aerogel. The results demonstrate the effectiveness of the Bayesian CNN framework in predicting strain energy for a given microstructure image of the silica aerogel while accounting for the level of confidence in the surrogate model prediction. Using Bayesian CNN, the effect of the number of training data points (ranging from 250 to 1000) on the accuracy and reliability of the microstructure-property prediction is investigated.

THEORY AND METHODOLOGY

High-fidelity Simulation

We aim to develop a surrogate model of microstructure image to mechanical property relation in silica aerogel materials. Silica aerogel involves a highly porous microstructure composed of interconnected amorphous ceramic nanoparticles at the nanometer scale. These materials have exceptional thermal properties, making them highly desirable for insulation applications in extreme environments, including civil and military aircraft, fire-resistant wearables, and high-performance insulation material for net-zero buildings, e.g., [1, 15, 21]. The high-fidelity simulation of mechanical deformation in silica aerogel consists of the stochastic partial differential equation (PDE) to determine a stochastic displacement field and compute the strain energy as the Quantity of Interest (QoI), given the microstructural images. Let Ω be a bounded domain in $\mathbb{R}^d$, $d = 1, 2, 3$, with Lipschitz boundary denoted as $\partial\Omega$. The problem is to determine a stochastic displacement field $\mathbf{u}(\omega, \mathbf{y})$ and compute strain energy as the QoI from it. Here, $\mathbf{y} \in \Omega$ is the spatial points and ω belongs to the sample set of possible outcomes describing realizations of microstructural patterns of silica aerogel. The governing equation for the high-fidelity model is expressed as follows,

$$\begin{aligned} \nabla \cdot \mathbf{T}(\omega, \mathbf{u}) &= \mathbf{f}(\mathbf{y}), \quad \mathbf{y} \in \Omega, \\ \mathbf{T}(\omega, \mathbf{u})\mathbf{n} &= \mathbf{t}(\mathbf{y}), \quad \mathbf{y} \in \Gamma_N, \\ \mathbf{u}(\omega, \mathbf{y}) &= \mathbf{u}^*, \quad \mathbf{y} \in \Gamma_D, \end{aligned} \quad (1)$$

where ∇ represents the spatial gradient operator, $\mathbf{f}$ and $\mathbf{t}$ are prescribed source and traction terms, Γ_N is a subset of $\partial\Omega$ on which Neumann boundary condition is prescribed, and $\Gamma_D = \partial\Omega \setminus \Gamma_N$ denotes the domain boundary subjected to Dirichlet condition. The Cauchy stress tensor, $\mathbf{T}$ is defined by

$$\begin{aligned}\mathbf{T}(\omega, \mathbf{u}) &= 2\mu_s \varphi(\omega, \mathbf{y}) \mathbf{E}(\mathbf{u}) + \lambda_s \varphi(\omega, \mathbf{y}) \text{tr}(\mathbf{E}(\mathbf{u})) \mathbf{I}, \\ \mathbf{E}(\mathbf{u}) &= \frac{1}{2} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T),\end{aligned}\quad (2)$$

where λ_s and μ_s are the Lamé constants of the solid aerogel phase (equivalent to Young's modulus E_s and Poisson's ratio ν_s), $\varphi(\omega, \mathbf{u})$ is the microstructure indicator function, taking $\varphi = 0$ for the spatial points inside pores and $\varphi = 1$ at the aerogel solid skeleton, and $\mathbf{E}(\mathbf{u})$ is the strain tensor. Samples of the stochastic microstructure indicator function are derived from a generative model, trained using microstructural images of silica aerogel obtained from a lattice Boltzmann simulation of the foaming process. It efficiently generates φ while preserving the same morphological properties across arbitrary domain sizes. Thus, the high-fidelity simulation for the elasticity problems involves the finite element solution of the PDE (1), employing a uniformly fine mesh to resolve the resolution of microstructure patterns as dictated by $\varphi(\omega, \mathbf{y})$. The model output is defined as the scalar strain energy of the material system, determined through the solution of the stochastic high-fidelity simulation as,

$$u_D = \int_{\Omega} \mathbf{T}(\omega, \mathbf{u}) : \mathbf{E}(\mathbf{u}) \, d\mathbf{y}. \quad (3)$$

The high-fidelity simulation is utilized to generate training datasets for constructing the surrogate models, as depicted in Figure 1. Specifically, strain energy is assessed for 1000 microstructure images of silica aerogel. As illustrated in the kernel density estimate presented in this figure, the strain energies range from $8.80\text{e-}05$ mJ to $9.80\text{e-}05$ mJ, with the mean strain energy of the dataset approximately $9.36\text{e-}05$ mJ and a standard deviation of approximately $14\text{e-}07$.

Convolutional Neural Network (CNN)

Convolutional neural networks are specially designed neural networks for image processing-related applications. It can deal with images in a specific way by using filters (or kernels) on images and storing the convoluted values as weights. There are four main layers of a CNN [5], i.e. Input layer, convolutional layer, pooling layer and fully-connected layer. The input layer serves as the initial entry point for data into the CNN, storing the raw pixel values of the input images. These pixel values are then processed through subsequent layers. The convolutional layer




FIGURE 1. Training data from high-fidelity simulation: (left) Four sample microstructure images (left). (right) Kernel density estimate of the strain energies of two different datasets with sizes of 1000 denoted in blue and 250 denoted in red.

follows the input layer. It contains the filters that extract features such as patterns, edges, and color depth from the input images. These filters are adjustable during training, enabling the network to learn relevant features. The convolutional layer outputs feature maps, which are then passed to the pooling layer. The pooling layer reduces the spatial dimensions of the feature maps, aiding computational efficiency and memory usage. This layer doesn't perform feature extraction but rather aggregates information, commonly using max or average pooling methods. Finally, the fully connected layer integrates the features learned from previous layers to make predictions. Each neuron in this layer is connected to every neuron in the preceding layer, facilitating comprehensive feature utilization for accurate predictions.

The architecture of CNN used in this study is shown in Figure 2. It contains four convolutional layers, three average pooling layers, and two fully-connected layers. The last layer of CNN (linear layer) is also a fully connected layer that performs the regression task. Each convolutional layer comprises three components: the number of filters, batch normalization, and activation function. The number of filters determines the variety of features extracted from the input data, with each filter capturing different aspects. Batch normalization ensures stable and accelerated training by normalizing the inputs of each layer, reducing internal covariate shifts. The Rectified Linear Unit (ReLU) activation function introduces non-linearity, allowing the network to learn complex patterns and gradients efficiently. Table 1 shows the number of parameters and output shape of each layer. This architecture has a total of 14,721,121 trainable parameters.

Bayesian CNN surrogate model

We define the CNN surrogate model for microstructure-property linkage as a mapping between the input vector $\mathbf{x}$ (representing a microstructure image) and the output $\mathbf{u}$ (representing material properties, such as strain energy). Using the CNN model, this map is parameterized by a set of weights $\mathbf{w}$ within the filters




FIGURE 2. CNN architecture breakdown of the full model and a sample convolutional layer.

TABLE 1. Summary of the CNN architecture used for developing the microstructure-property surrogate model.

Layer (type)	Output Shape	Param #
Conv2d-1	[1, 16, 168, 168]	416
BatchNorm2d-1	[1, 16, 168, 168]	32
ReLU-1	[1, 16, 168, 168]	0
Conv2d-2	[1, 32, 168, 168]	12,832
BatchNorm2d-2	[1, 32, 168, 168]	64
ReLU-2	[1, 32, 168, 168]	0
AvgPool2d-2	[1, 32, 84, 84]	0
Conv2d-3	[1, 64, 84, 84]	51,264
BatchNorm2d-3	[1, 64, 84, 84]	128
ReLU-3	[1, 64, 84, 84]	0
AvgPool2d-1	[1, 64, 42, 42]	0
Conv2d-4	[1, 128, 42, 42]	204,928
BatchNorm2d-4	[1, 128, 42, 42]	256
ReLU-4	[1, 128, 42, 42]	0
AvgPool2d-2	[1, 128, 21, 21]	0
Flatten-1	[1, 56448]	0
Linear-1	[1, 256]	14,450,944
ReLU-5	[1, 256]	0
Linear-2	[1, 1]	257

of convolutional layers and fully-connected layers. The training process of the CNN involves adjusting the values of these weights using the training dataset $\mathbf{D} = \{(\mathbf{x}_i, \mathbf{u}_i)\}_{i=1}^{N_D}$. Traditional training methods in regression problems aim to minimize an error function between the output and the training data without explicitly considering the uncertainty. In a Bayesian framework,

a prior probability distribution function (PDF) is assigned to the weights $\pi_{\text{prior}}(\mathbf{w})$ before observing data. Subsequently, the prior is updated based on the data using Bayes' rule, e.g., [9, 18, 20],

$$\pi_{\text{post}}(\mathbf{w} | \mathbf{D}) = \frac{\pi_{\text{like}}(\mathbf{D} | \mathbf{w}) \pi_{\text{prior}}(\mathbf{w})}{\pi_{\text{evid}}(\mathbf{D})}, \quad (4)$$

where $\pi_{\text{like}}(\mathbf{D} | \mathbf{w})$ represents the likelihood of observing the data $\mathbf{D}$ given the parameter $\mathbf{w}$, and $\pi_{\text{evid}}(\mathbf{D})$ serves as the evidence PDF, acting as a normalization factor. Given the posterior PDF $\pi_{\text{post}}(\mathbf{w} | \mathbf{D})$, the output can be predicted for any new input $\mathbf{x}^*$ from the predictive distribution such that,

$$\pi(\mathbf{u}^* | \mathbf{x}^*, \mathbf{D}) = \int \pi(\mathbf{u}^* | \mathbf{x}^*, \mathbf{w}) \pi_{\text{post}}(\mathbf{w} | \mathbf{D}) d\mathbf{w}. \quad (5)$$

Computing the posterior PDF is computationally intractable due to the high-dimensional parameter space inherent in CNNs. To tackle this challenge, some inference methods attempt to approximate the posterior distribution using alternative distributions such as variational inference described in the following section or generate samples from it using Markov chain Monte Carlo (MCMC) algorithms. After the inference phase, a collection of parameter samples denoted as $\{\hat{\mathbf{w}}_j\}_{j=1}^M$ are obtained, and the prediction distribution can be approximated using Monte Carlo estimation as follows,

$$\pi(\mathbf{u}^* | \mathbf{x}^*, \mathbf{D}) = \mathbb{E}_{\mathbf{w} | \mathbf{D}}[\pi(\mathbf{u}^* | \mathbf{x}^*, \mathbf{w})] \approx \frac{1}{M} \sum_{j=1}^M f(\mathbf{u}^* | \mathbf{x}^*, \hat{\mathbf{w}}_j), \quad (6)$$

where, f is a deterministic map between $\mathbf{x}^*$ and $\mathbf{u}^*$ given $\hat{\mathbf{w}}_j$.

Variational inference: This approach involves approximating the true posterior $\pi_{\text{post}}(\mathbf{w} | \mathbf{D})$ using a variational distribution $q_{\boldsymbol{\theta}}(\mathbf{w})$ that is computationally tractable and parameterized by $\boldsymbol{\theta}$. Typically, the variational distribution is expressed as a product of independent Gaussian distributions, known as the mean-field approximation $q_{\boldsymbol{\theta}}(\mathbf{w}) = \prod_i q_{\theta_i}(w_i)$, where $\boldsymbol{\theta}$ encompasses the means μ and standard deviations σ of each Gaussian component. The Bayesian training is then framed as learning the variational parameters $\boldsymbol{\theta}$ by minimizing the Kullback-Leibler (KL) divergence between the density $q_{\boldsymbol{\theta}}(\mathbf{w})$ and the posterior PDF $\pi_{\text{post}}(\mathbf{w} | \mathbf{D})$, yielding,

$$\boldsymbol{\theta}^{\text{opt}} = \arg \min_{\boldsymbol{\theta}} \mathcal{F}(\boldsymbol{\theta}), \quad (7)$$

where $\mathcal{F}$ is the loss function and denoted as [2],

$$\begin{aligned}\mathcal{F} &= \text{KL} [q_{\boldsymbol{\theta}}(\mathbf{w}) \| \pi_{\text{post}}(\mathbf{w} | \mathbf{D})] \\ &= \text{KL} [q_{\boldsymbol{\theta}}(\mathbf{w}) \| \pi_{\text{prior}}(\mathbf{w})] \\ &\quad - \mathbb{E}_{q_{\boldsymbol{\theta}}(\mathbf{w})} [\ln \pi_{\text{like}}(\mathbf{D} | \mathbf{w})],\end{aligned}\quad (8)$$

and the KL divergence term in (8) can be expressed as,

$$\text{KL} [q_{\boldsymbol{\theta}}(\mathbf{w}) \| \pi_{\text{post}}(\mathbf{w} | \mathbf{D})] = \int q_{\boldsymbol{\theta}}(\mathbf{w}) \ln \frac{q_{\boldsymbol{\theta}}(\mathbf{w})}{\pi_{\text{post}}(\mathbf{w} | \mathbf{D})} d\mathbf{w}, \quad (9)$$

representing the dissimilarity between the posterior and its variational approximation. Minimizing the loss function in (8) ensures that the variational distribution captures a good representation of the data while imposing regularization through the prior distribution to prevent overfitting. The data dependent term in (8) is known as “data misfit” and usually computed with sampling methods,

$$\mathbb{E}_{q_{\boldsymbol{\theta}}(\mathbf{w})} [\ln \pi_{\text{like}}(\mathbf{D} | \mathbf{w})] = \frac{1}{n} \sum_{i=1}^n \ln \pi_{\text{like}}(\mathbf{D} | \hat{\mathbf{w}}^{(i)}), \quad (10)$$

where $\hat{\mathbf{w}}^{(i)}$ are samples drawn from the variational distribution $q_{\boldsymbol{\theta}}(\mathbf{w})$ and n is the number of draws. To this end, the loss function to update the parameters, known as the evidence lower bound (ELBO), can be expressed as,

$$\mathcal{F} \approx \beta \text{KL} (q_{\boldsymbol{\theta}}(\mathbf{w}) \| \pi_{\text{prior}}(\mathbf{w})) - \frac{1}{n} \left[\sum_{i=1}^n \ln \pi_{\text{like}}(\mathbf{D} | \hat{\mathbf{w}}^{(i)}) \right], \quad (11)$$

where β is a weighting factor for the regularization term in comparison to the data misfit term. Use of β weighting is not a standard component of the variational inference algorithm and is instead empirically motivated rather than theoretically justified. It’s inclusion is motivated by the Cold Posterior Effect [23] which suggests that the standard Bayesian posterior can be sub-optimal in certain deep learning models and requires artificial sharpening (i.e., temperature scaling) to achieve optimal performance. It can be shown that sharpening the posterior is equivalent to reducing the prior weight. Consequently, we empirically adjusted the prior weight to enhance performance. Details on the determination of this weighting factor are provided in the following section.

The closed form solution that was used to compute the KL divergence function analytically is as below,

$$\begin{aligned}\text{KL} [q_{\boldsymbol{\theta}}(\mathbf{w}) \| \pi_{\text{prior}}(\mathbf{w})] &= \\ \frac{1}{2} \sum_{i=1}^d \left[2 \log \left(\frac{\sigma_p}{\sigma_q} \right) - 1 + \left(\frac{\sigma_q}{\sigma_p} \right)^2 + \left(\frac{\mu_p - \mu_q}{\sigma_p} \right)^2 \right].\end{aligned}\quad (12)$$

In (12), d is the dimension of the random variables $\mathbf{w}$ and μ_p, σ_p, μ_q , and σ_q are the mean and standard deviation of the prior and posterior distributions respectively.

To solve the minimization problem given by (7) with the loss function in (11) using gradient-based optimization algorithms, it is necessary to compute the gradient of the ELBO with respect to the variational parameters $\boldsymbol{\theta}$. However, directly optimizing over random distributions is not feasible. Therefore, we employ the “Reparameterization Trick” [7] to facilitate the optimization of the distribution’s parameters.

Local Reparametrization Trick: The core idea of this approach is to apply an affine transformation to a standard normal distribution for each weight w_i such that $w_i = \mu_i + \sigma_i * x_i$, where $x_i \sim \mathcal{N}(0, 1)$. During training, x_i is treated as a constant, and the parameters of the affine transformation (μ_i and σ_i) are optimized. To ensure that $\sigma_i > 0$, we use a common practice of parameterizing it with the softplus function, $\sigma_i(\rho_i) = e^{1+\log(\rho_i)} > 0$, where ρ_i is optimized instead of σ_i . This allows for the variational parameters μ and σ to be updated through standard backpropagation. Using this method for updating the variational parameters is known as “Bayes by Backprop” [2].

This work uses stochastic gradient descent (SGD) to solve the optimization problem. The performance of the SGD largely depends on the variance of the gradients. So, for keeping the variance low and computational acceleration, the local reparameterization trick was used, where instead of sampling $\mathbf{w}$ (weights) directly, samples are taken from the neuron pre-activation sums from their implied Gaussian distribution. This approach essentially involves treating the mean μ and standard deviation σ as deterministic, while introducing stochasticity via noise (typically standard Gaussian) in the pre-activation sums.

RESULTS AND DISCUSSION

Hyperparameters Tuning

The hyperparameters of the Bayesian CNN were determined through a validation process. One such hyperparameter is the learning rate utilized within the Adam optimizer during variational inference. To this end, the entire high-fidelity dataset was partitioned into training and validation sets, comprising 80% and 20% of the data, respectively. The validation loss, defined as the mean squared error between the validation data and the corresponding Bayesian CNN mean prediction trained on the training data, serves as an indicator of the model’s performance on new data at each epoch. Figure 3 (a) depicts the progressive reduction in validation loss over epochs for four different learning rates: 0.001, 0.002, 0.003, and 0.004. These plots indicate that a learning rate of 0.003 corresponds to the plot displaying minimal validation losses. Moreover, it is observed that the losses plateau approximately after 300 epochs, coinciding with the epoch exhibiting the lowest validation loss. Consequently, based on this study, a learning rate of 0.003 was selected for the model train-




FIGURE 3. Tuning the hyperparameters of the Bayesian CNN surrogate model: (a) Validation loss decreases over the number of epochs for four distinct learning rates. A learning rate of 0.003 was selected due to its association with the minimum validation loss. (b) The standard deviation of a randomly chosen weight in the first convolutional layer declines as the number of epochs increases. This trend is observed across all weights, prompting the selection of a regularization weight of $\beta = 0.0001$.

ing.

The second hyperparameter is β in (11), which represents the weighting factor of the regularization term relative to the likelihood function. The efficacy of Bayesian inference crucially hinges on the accurate determination of this parameter to strike a balance between data misfit in the likelihood function and the prior distribution of parameters. An overestimated β may lead to the neglect of valuable information contained in the data, while underestimating this weight can result in overfitting, leading to overly confident and biased parameter estimation. To avoid bias in Bayesian inference, we determine this parameter through iterative experimentation, monitoring the variation of the standard




FIGURE 4. Prediction of 4 images, labeled respectively 1,2,3, and 4 (left to right) in Table 2. The blue distribution represents prediction trained with 1000 data points, the red distribution represents prediction with 250 data points and the black dashed line is the true strain energy.

deviation of inferred weights across epochs. Figure 3 (b) illustrates an example of the standard deviation of a kernel within the first convolutional layer. In this study, $\beta = 0.0001$ is chosen as it corresponds to a decreasing trend in the standard deviation of all kernel parameters, aligning with the expected increase in confidence levels as the model is informed by additional data at each epoch, in accordance with the notion of Bayesian learning. Moreover, the determined value of the regularization weight falls within the range of β values reported in prior studies [24].

Prediction Uncertainty

Following the validation of the Bayesian inference process, the entire dataset is utilized to determine the probability distribution of the CNN parameters. Subsequently, the resulting predictive CNN surrogate model can be leveraged to predict strain energy with quantified uncertainty for any given microstructural image input. Figure 4 presents the Bayesian CNN predictions on the testing set, represented through kernel density estimates of strain energies across four microstructural images outside the training set. These prediction distributions, as per (6), incorporate $M = 50$ samples of the posterior PDF of the weight parameters, each characterized by two means and variances. Furthermore, this figure compares the prediction results obtained from two sets of training data: one comprising 1000 images and the other 250 images. The mean and standard deviation of the predictions, compared with the true strain energy values obtained from high-fidelity simulations, are presented in Table 2. The

findings suggest that, as anticipated, employing a larger dataset of 1000 images for training yields more precise and reliable predictions from the surrogate model, as evidenced by the closer alignment of the mean with the true strain energy and the sharper probability distributions of strain energy.

TABLE 2. The mean and standard deviation of strain energies (in 10^{-7} mJ) predicted by the Bayesian CNN surrogate model (trained with 1000 and 250 data points) in the testing set, contrasted with the true values acquired from high-fidelity simulation.

Testing image	True value	1000 training data		250 training data	
		Mean	StD	Mean	StD
1	937	936	0.65	929	1.60
2	940	948	0.85	935	1.66
3	950	948	0.82	931	1.88
4	919	931	0.78	932	1.11

CONCLUSION

This paper introduces a Bayesian CNN surrogate model aimed at capturing the relationship between microstructural images and their corresponding mechanical strain energy in silica aerogel materials. To achieve this, variational inference, along with Bayes by Backprop, has been utilized to determine the probability distributions of weights within the filters of convolutional layers and the weights in fully-connected layers. The training data are obtained from a pore-scale elasticity model of aerogel governed by a stochastic partial differential equation. Hyperparameters, including the learning rate and the weighting factor of the regularization term, were determined using a validation dataset. Following hyperparameter tuning, the entire dataset is employed within Bayesian inference to determine the distributions of parameters for the Bayesian CNN surrogate model. The accuracy and reliability of the surrogate model's predictive capability are assessed using a testing dataset. The results demonstrate that the proposed Bayesian CNN is capable of characterizing uncertainty in training data obtained from high-fidelity simulations, thus translating into uncertainty in surrogate model predictions.

Future work will concentrate on expanding the application of the Bayesian CNN surrogate model to predict aerogel properties across a wide range of microstructural features, encompassing different pore sizes and morphologies. Additionally, thorough investigations will be conducted on the effect of CNN architecture on its predictive ability, and leveraging Bayesian model plausibility to guide the determination of optimal architecture and hyperparameters.

ACKNOWLEDGMENT

MAI and DF extend their sincere gratitude for the financial support received from the U.S. National Science Foundation (NSF) CAREER Award CMMI-2143662. The authors acknowledge the Center for Computational Research (<http://www.buffalo.edu/ccr.html>) at University at Buffalo for providing HPC resources that have contributed to the research results reported here.

REFERENCES

- [1] L. An, M. Di Luigi, J. Tan, D. Faghihi, and S. Ren. Flexible percolation fibrous thermal insulating composite membranes for thermal management. *Materials Advances*, 4(1):284–290, 2023.
- [2] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.
- [3] D. Brands, D. Balzani, L. Scheunemann, J. Schröder, H. Richter, and D. Raabe. Computational modeling of dual-phase steels based on representative three-dimensional microstructures obtained from ebsd data. *Archive of Applied Mechanics*, 86:575–598, 2016.
- [4] S. Ghosh, K. Lee, and S. Moorthy. Multiple scale analysis of heterogeneous elastic structures using homogenization theory and voronoi cell finite element method. *International journal of solids and structures*, 32(1):27–62, 1995.
- [5] M. A. Islam, S. I. Rashid, N. U. I. Hossain, R. Fleming, and A. Sokolov. An integrated convolutional neural network and sorting algorithm for image classification for efficient flood disaster management. *Decision Analytics Journal*, 7:100225, 2023.
- [6] S. Kalidindi and S. Schoenfeld. On the prediction of yield surfaces by the crystal plasticity models for fcc polycrystals. *Materials Science and Engineering: A*, 293(1-2):120–129, 2000.
- [7] D. P. Kingma, T. Salimans, and M. Welling. Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28, 2015.
- [8] Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [9] B. Liang, J. Tan, L. Lozenski, D. A. Hormuth, T. E. Yankeelov, U. Villa, and D. Faghihi. Bayesian inference of tissue heterogeneity for individualized prediction of glioma growth. *IEEE Transactions on Medical Imaging*, 42(10):2865–2875, 2023.
- [10] A. Mann and S. R. Kalidindi. Development of a robust cnn model for capturing microstructure-property linkages and building property closures supporting material design. *Frontiers in materials*, 9:851085, 2022.

- [11] R. M. Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [12] A. Olivier, M. D. Shields, and L. Graham-Brady. Bayesian neural networks for uncertainty quantification in data-driven materials modeling. *Computer methods in applied mechanics and engineering*, 386:114079, 2021.
- [13] J. H. Panchal, S. R. Kalidindi, and D. L. McDowell. Key computational modeling issues in integrated computational materials engineering. *Computer-Aided Design*, 45(1):4–25, 2013.
- [14] F. Roters, P. Eisenlohr, L. Hantcherli, D. D. Tjahjanto, T. R. Bieler, and D. Raabe. Overview of constitutive laws, kinematics, homogenization and multiscale methods in crystal plasticity finite-element modeling: Theory, experiments, applications. *Acta materialia*, 58(4):1152–1211, 2010.
- [15] A. Sarkar, P. K. Singh, L. Zhu, D. Faghihi, and S. Ren. Carbon-sequestration straw cellulose-aerogel gradient thermal insulation material. *ACS Applied Engineering Materials*, 2024.
- [16] M. I. R. Shishir, M. S. R. Elapolu, and A. Tabarraei. A deep learning model for predicting mechanical properties of polycrystalline graphene. *Computational Materials Science*, 218:111924, 2023.
- [17] K. Shridhar, F. Laumann, and M. Liwicki. A comprehensive guide to bayesian convolutional neural network with variational inference. *arXiv preprint arXiv:1901.02731*, 2019.
- [18] P. K. Singh, K. A. Farrell-Maupin, and D. Faghihi. A framework for strategic discovery of credible neural network surrogate models under uncertainty. *arXiv preprint arXiv:2403.08901*, 2024.
- [19] P. K. Singh, K. A. Farrell-Maupin, and D. Faghihi. A framework for strategic discovery of credible neural network surrogate models under uncertainty. *Computer Methods in Applied Mechanics and Engineering*. (*arXiv preprint arXiv:2403.08901*), in press.
- [20] J. Tan, B. Liang, P. K. Singh, K. A. Farrell-Maupin, and D. Faghihi. Toward selecting optimal predictive multiscale models. *Computer Methods in Applied Mechanics and Engineering*, 402:115517, 2022.
- [21] J. Tan, P. Maleki, L. An, M. Di Luigi, U. Villa, C. Zhou, S. Ren, and D. Faghihi. A predictive multiphase model of silica aerogels for building envelope insulations. *Computational Mechanics*, 69(6):1457–1479, 2022.
- [22] E. Wargo, A. Hanna, A. Cecen, S. Kalidindi, and E. Kumbar. Selection of representative volume elements for pore-scale analysis of transport in fuel cell materials. *Journal of power sources*, 197:168–179, 2012.
- [23] F. Wenzel, K. Roth, B. S. Veeling, J. Światkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. How good is the bayes posterior in deep neural networks really? *arXiv preprint arXiv:2002.02405*, 2020.
- [24] F. Wenzel, K. Roth, B. S. Veeling, J. Światkowski, L. Tran, S. Mandt, J. Snoek, T. Salimans, R. Jenatton, and S. Nowozin. How good is the bayes posterior in deep neural networks really? *arXiv preprint arXiv:2002.02405*, 2020.
- [25] Z. Yang, Y. C. Yabansu, R. Al-Bahrani, W.-k. Liao, A. N. Choudhary, S. R. Kalidindi, and A. Agrawal. Deep learning approaches for mining structure-property linkages in high contrast composites from simulation datasets. *Computational Materials Science*, 151:278–287, 2018.
- [26] Y. Zhu and N. Zabarar. Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 366:415–447, 2018.