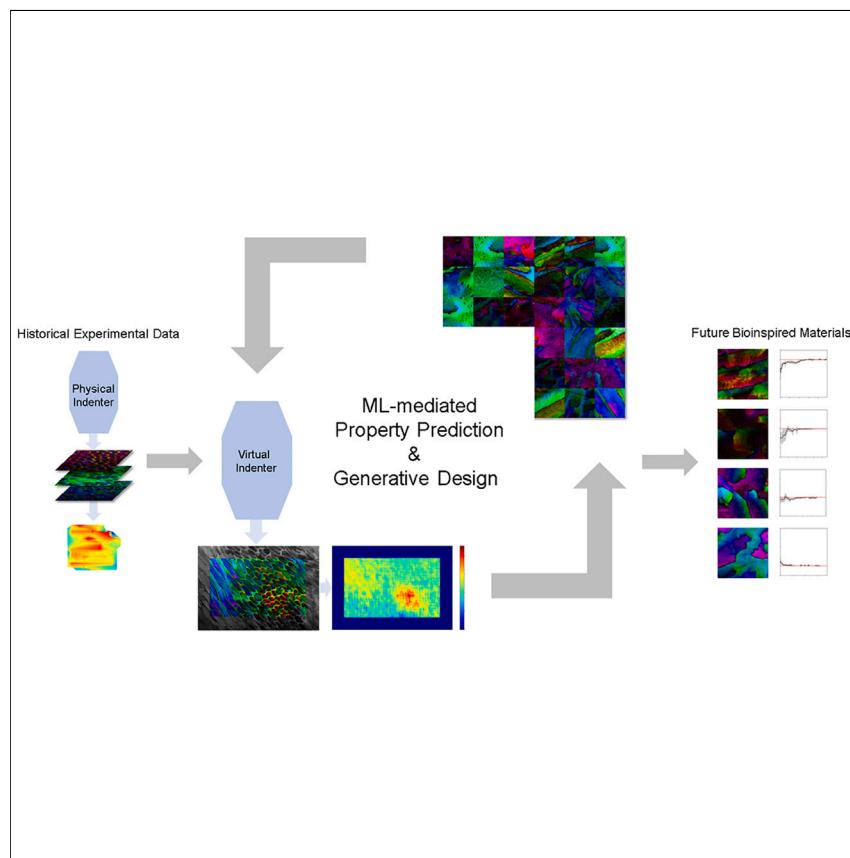


Article

Deep learning virtual indenter maps nanoscale hardness rapidly and non-destructively, revealing mechanism and enhancing bioinspired design



Andrew J. Lew, Cayla A. Stifler,
Astrid Cantamessa, Alexandra
Tits, Davide Ruffoni, Pupa U.P.A.
Gilbert, Markus J. Buehler

mbuehler@mit.edu

Highlights

Traditional ways of measuring hardness can be augmented with ML models

Expanded insight can be derived from known data, without experimental advancement

Adding generative models and genetic algorithms enables targeted bioinspired design

Hardness measurement is destructive and time consuming, but it is a very common process with much historical study. Thus, we train an ML model that can quickly predict hardness given a material structure. We also train a generative model that can generate bioinspired structures. Using the predictive and generative models together allows us to emulate the progress of nature and “evolve” structures that possess desired hardness values. This process is an iterative approach for targeted, bioinspired materials design.



Benchmark

First qualification/assessment of material properties and/or performance

Article

Deep learning virtual indenter maps nanoscale hardness rapidly and non-destructively, revealing mechanism and enhancing bioinspired design

Andrew J. Lew,^{1,2} Cayla A. Stifler,³ Astrid Cantamessa,⁴ Alexandra Tits,⁴ Davide Ruffoni,⁴ Pupa U.P.A. Gilbert,^{3,5,6,7} and Markus J. Buehler^{1,8,*}

SUMMARY

Over evolution, organisms develop complex material structures fit to their environments. Based on these time-tested designs, human-engineered bioinspired structures offer exciting possible materials configurations. However, navigating diverse structure spaces for attaining desired properties remains non-trivial. We focus on the hardest biological tissue in humans, tooth enamel, to examine the structure-property relationship. While typical hardness measurements are time consuming and destructive, we propose that artificial intelligence models can predict properties directly and enable high-throughput, non-destructive characterization. We train a deep image regression neural network as a surrogate model and visualize with gradient ascent and saliency maps to identify structural features contributing most to hardness. This model demonstrates improved spatial resolution and sensitivity compared with experimental hardness maps. Using this rapid hardness testing model, a generative adversarial model, and a genetic algorithm that operates in latent space, allows for guided materials design, yielding proposed designs for bioinspired structures with precisely controlled hardness.

INTRODUCTION

Hardness measurement is a ubiquitous materials characterization technique that is performed more frequently than any other mechanical test.¹ Despite this, there are numerous limitations on traditional indentation experiments. To name a few, indentation is a destructive method that changes material properties, meaning it cannot be used in production settings. Furthermore, this destruction limits how closely spaced test points can be,² otherwise measurements will suffer from deformation artifacts. Additionally, imperfect indentation quality may lead to unclear results³ that depend on interpretation by the observer. Different operators of varying skill levels may yield inconsistent measurements of similar samples. Even when these elements are not of concern, the best-quality results are obtained through low-throughput experiments with a large experimental time per sample.⁴

Of course, such issues are only exacerbated when attempting high-throughput characterization. Completing several thousands of indents in a reasonable time frame requires shortening the typical testing rate of tens of measurements per hour.⁵ When doing so, one must take care to choose a correct indentation depth⁶ that avoids deformation overlap. This is because faster loading rates make consideration of

PROGRESS AND POTENTIAL

Bioinspired structures offer exciting possible materials configurations, but navigating diverse structure spaces for attaining desired properties remains non-trivial. We focus on the hardest biological tissue in humans, tooth enamel, to examine the structure-property relationship. While typical hardness measurements are time consuming and destructive, we propose that AI models can predict properties directly and enable high-throughput, non-destructive characterization. We train a deep image regression neural network as a surrogate model and visualize with gradient ascent and saliency maps to identify structural features contributing most to hardness. Using this rapid hardness testing model, a generative adversarial model, and a genetic algorithm together allows for guided materials design, yielding proposed bioinspired structures with precisely controlled hardness. There is potential for the approach to impact other material systems and properties given an appropriate training dataset.

strain rate sensitivity particularly important.⁵ In short, rapid, consistent, and scalable hardness testing remains an issue of great import and active research.

Artificial intelligence (AI)-based models have been greatly successful in accelerating and scaling materials characterization and prediction tasks. In the realm of mechanical properties, previous models have successfully predicted fracture,^{7–9} compliance optimization,¹⁰ and buckling¹¹ using a combination of neural networks, latent space representations, and genetic algorithms.

Thus, here we propose using a combination of three AI methods—regression, adversarial generation, and genetic algorithm optimization—to not only tackle the problem of hardness characterization, but also to gain a deeper understanding of the factors contributing to hardness and facilitate the design of structures with controlled hardness. Specifically, we focus here on the material system of human enamel, the hardest tissue in the human body and a classic subject of hardness characterization,¹² where recent polarization-dependent imaging contrast (PIC) mapping¹³ has provided detailed surface structure information in various biominerals. We hypothesize that this rich surface information may be suitable for hardness evaluation and provide a unique opportunity for identifying structure-property relationships.

RESULTS

Obtaining a hardness model

A primary challenge of utilizing machine learning models is the acquisition of suitable training data, ideally of both large volume and high accuracy. However, there is often a practical trade-off between large volume and high accuracy experimentation. Thus, here we employ a strategy of combining datasets with different levels of volume and accuracy to obtain a more complete view of enamel hardness.

We first utilize a hardness dataset with accurate labeling, shown in *Figure 1A*, comprised of 94 hardness measurements across a $125 \times 100 \mu\text{m}$ region of human enamel. Nanoindentation details are provided in the *experimental procedures* section. The second dataset we use is generated via estimated labeling from cross referencing literature, shown in *Figure 1B*. Specifically, we use enamel images from the literature of different regions of a human molar,¹⁴ subdivide each image into square 200×200 pixel segments, and estimate hardness labels for each square by comparing the square location with human enamel hardness maps in the literature,^{15,16} for a total of 333 image-label pairs. It should be noted that this estimation runs a risk of including error between hardness values and microstructures, due to natural variation of biological samples from tooth age, oral environment, tooth position, etc. The third dataset is comprised of 17 large enamel maps from the literature¹⁷ with noisy labeling, shown in *Figure 1C*. Despite the large volume of image area in this third dataset, only the average hardness values and standard deviations for each image are known. There is low granularity in the hardness labels of any exact location. To summarize, the first dataset has small volume but accurate and precise labels; the second dataset is larger but has approximate labels with fewer significant digits and the possibility for error; and the third dataset is the largest, but any local image section is likely to have some unknown deviation from the average, resulting in label errors.

We then combine these three datasets into one single dataset of 22,730 image-label pairs, with a distribution of hardness values between 2 and 7 GPa as shown in *Figure 1D*. By uniting these three streams of data, we obtain a large volume dataset

¹Laboratory for Atomistic and Molecular Mechanics (LAMM), Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

²Department of Chemistry, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

³Department of Physics, University of Wisconsin, Madison, WI 53706, USA

⁴Mechanics of Biological and Bioinspired Materials Laboratory, Department of Aerospace and Mechanical Engineering, University of Liège, 4000 Liège, Belgium

⁵Department of Chemistry, University of Wisconsin, Madison, WI 53706, USA

⁶Departments of Materials Science and Engineering, Geoscience, University of Wisconsin, Madison, WI 53706, USA

⁷Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

⁸Lead contact

*Correspondence: mbuehler@mit.edu

<https://doi.org/10.1016/j.matt.2023.03.031>

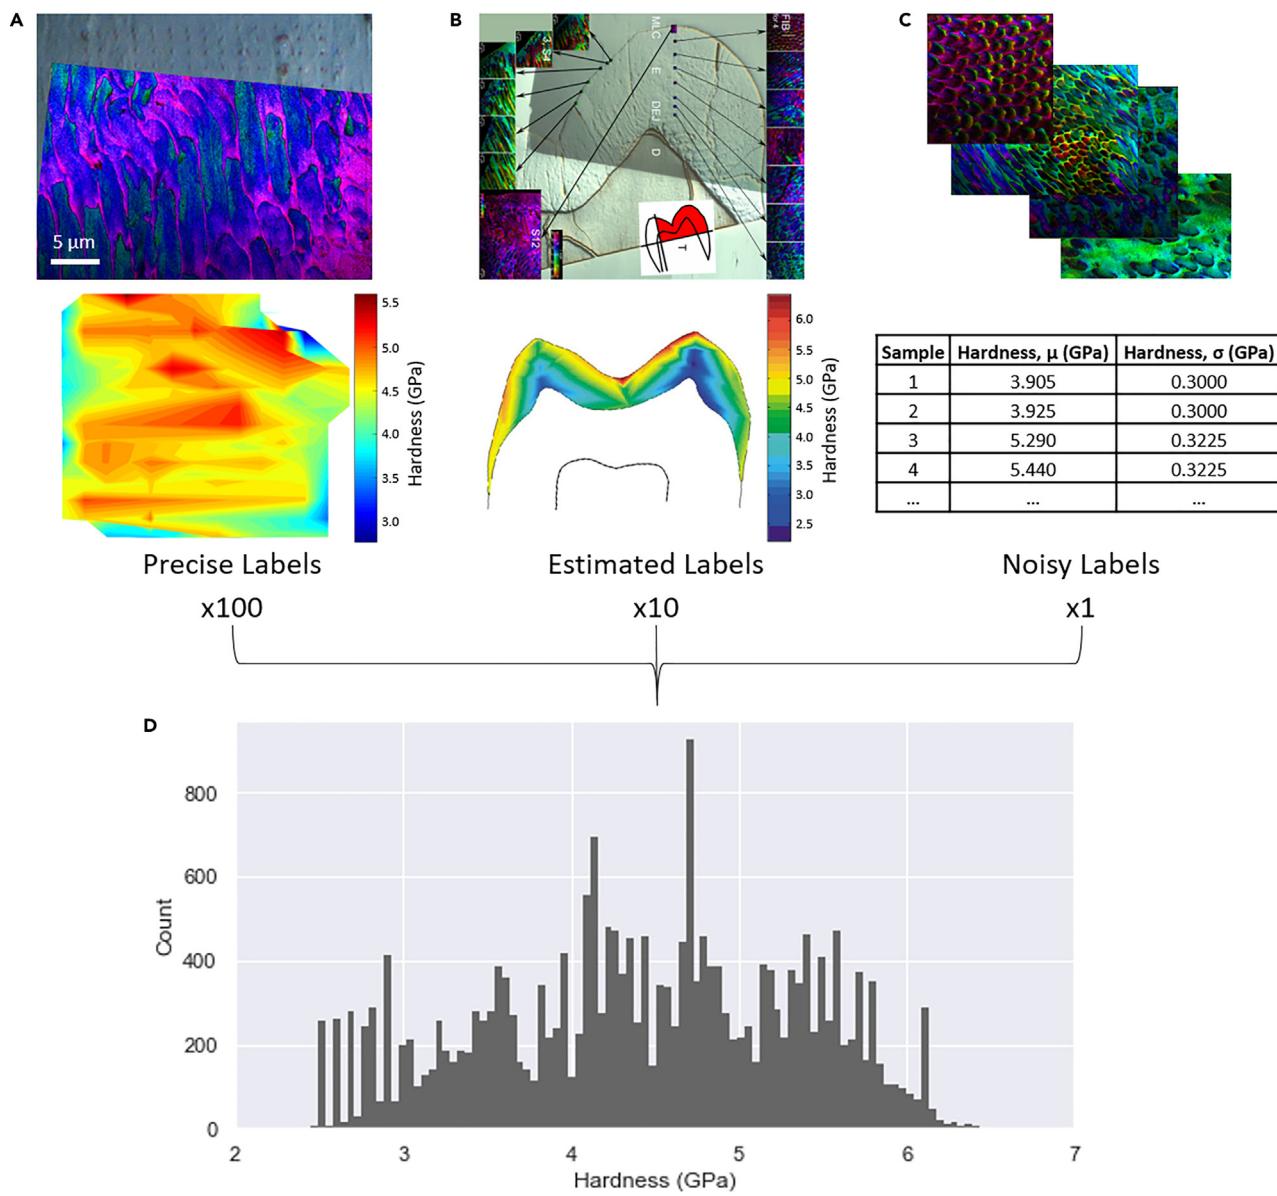


Figure 1. Dataset curation

Performant ML models require training data of both high accuracy and large volume.

(A–C) To achieve this comprehensive representation of hardness, we combine three imperfect datasets together: (A) a small dataset of accurate locally measured labels, (B) a mid-size dataset with estimated labels from known literature trends (Beniash et al.¹⁴ and Cuy et al.¹⁵), and (C) a large dataset of structures where only the average and standard deviation of each image is known (adapted from Stifler et al.¹⁷), resulting in some noise associated with any local label.

(D) We obtain the united dataset by weighting the representation of the more reliable image-label pairs accordingly, for a total of 22,730 image-label pairs spanning hardness values from 2 to 7 GPa.

that also incorporates a representation of high accuracy measurement. Details of how these three datasets are combined are provided in the [experimental procedures](#) section.

This united dataset is used to train an image regression model that maps enamel structure images to hardness values. It should be noted that, along with the known potential for label errors described above, the PIC maps are only a proxy for the

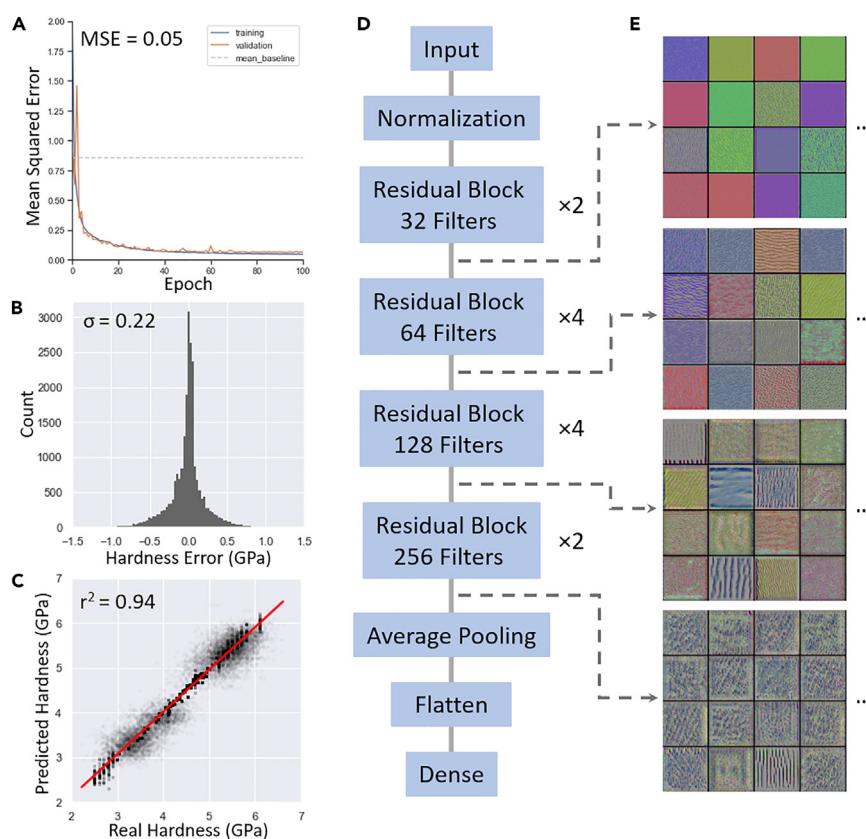


Figure 2. Deep residual network regression

Our model achieves (A) mean squared error loss <0.1 , with good agreement between training and validation sets. (B) The error in predictions is centered and narrow, with $\mu = -0.01$ and $\sigma = 0.22$, and (C) the linear correlation between predicted and labeled hardness values has $r^2 > 0.9$. Furthermore, the fit between predicted and labeled hardness values with 99% confidence interval visually appears as a singular line, indicating highly accurate predictions. (D) The architecture of our model incorporates shortcut connections in residual blocks to map structure images to hardness labels. Using filter activation maximization reveals (E) that initial layers of the model learn crystal orientation in the form of image colors, while deeper layers focus on orientation agnostic structure in the form of patterns and textures.

structural details probed by indentation, as PIC mapping only penetrates approximately 3 nm into the sample, while nanoindentation will probe approximately 150 nm of material. Despite this difference in scale, tracking the mean squared error of predicted hardness values shows good agreement between training and validation sets over 100 epochs of training, shown in Figure 2A. This indicates that the 3 nm surface probed by PIC maps encapsulates sufficient crucial structural features for determining hardness of this material system. There is very low error compared with the average mean squared error baseline (calculated by simply predicting each image as having the mean hardness value of the entire input dataset), demonstrating that the model has learned far beyond trivial performance. Specifically, the error in hardness predictions is narrow ($\sigma = 0.22$ GPa) with a balanced distribution only slightly underpredicting values ($\mu = -0.01$ GPa), as shown in Figure 2B. Predicted hardness values map linearly to hardness labels with an r^2 value of 0.94, shown in Figure 2C, demonstrating that the model has successfully learned to quantitatively map enamel structure images to hardness values. The architecture of our deep residual neural network incorporating skip connections¹⁸ is shown in Figure 2D.

Details of model architecture and training are provided in the [experimental procedures](#) section.

Understanding structure-property relationship: Misorientation hardness

Beyond simply generating a black-box model able to rapidly and non-destructively predict hardness values directly from nanoscale structure, it is important to understand what the model does. Identifying what factors contribute to hardness predictions may help us understand deeper mechanisms of hardness as a material property and increase our confidence that such a machine learning (ML)-driven approach is not simply a fluke.¹⁹ Thus, we employ filter activation maximization to visualize what our model "sees." To do this, we feed in an initial randomized image of gray static and repeatedly apply gradient ascent to maximize the activation of a specific filter in a target layer.^{20,21} We visualize 16 maximized filters after each section of residual blocks in [Figure 2E](#) to extract the sequential process of how the model comes to a hardness prediction. We observe a progression from bright single-color images at earlier layers to monochrome textures at deeper layers. Physically, this indicates that the model first pays attention to crystal orientations, which are encoded by pixel color in PIC maps, with different orientations handled by different filters. Conversely, at the deep end of the model, filters are attuned to pick up different structural patterns and crystal sizes regardless of crystal orientation.

After visualizing the general approach of how the model processes hardness prediction, we subsequently delve further into how specific values are determined. To do this, we use a saliency map visualization across a selection of enamel structures shown in [Figure 3A](#). Saliency maps compute the spatial support of a prediction, identifying local regions in an image most important to the determination of the output.²² We then convert the hue and brightness values of each image (which represent c-axis orientations with respect to the polarization plane) to orientations ϕ and θ with respect to the image plane (since the polarization plane is 60° from the sample surface²³). Subsequently, in [Figure 3B](#), we plot saliency as a function of the in-image plane orientation ϕ and the out-of-image plane orientation θ , where $\phi = 0^\circ$ corresponds to the vertical in-image direction and $\theta = 0^\circ$ corresponds to the perpendicular direction coming straight out of the image. We find that the saliency distribution roughly corresponds to histograms of the frequency of orientations in [Figure 3C](#), as one might expect: when certain crystal orientations are more plentiful than others, there is a higher chance that highly salient regions would be comprised of these orientations. Additionally, it may also indicate that the model has attributed higher importance to the orientations that occur more frequently. There is an important deviation to this correlation, however, when one scales saliency values by orientation frequency in [Figure 3D](#). Specifically, we note that the least common orientations have larger relative saliencies. In other words, there appears to be an additional factor maintaining a level of importance to rare orientations—the model does not simply ignore them. To interpret this physically, we note that rare orientations are less likely to be similar to other surrounding orientations. In other words, there are likely to be local misorientations in their vicinity. From the literature, it has been reported that misorientations in tooth enamels, measured as the angular distance in 3D space of c-axis unit vectors via dot product, are known to be positively correlated with hardness.¹⁷ The model's relative fixation on rare orientations may indicate an understanding of this misorientation hardness effect.

By applying domain knowledge to visualization tools, we are able to better understand both what the model is predicting and the important factors behind hardness as a property. The identification of intelligible factors behind the model indicate that

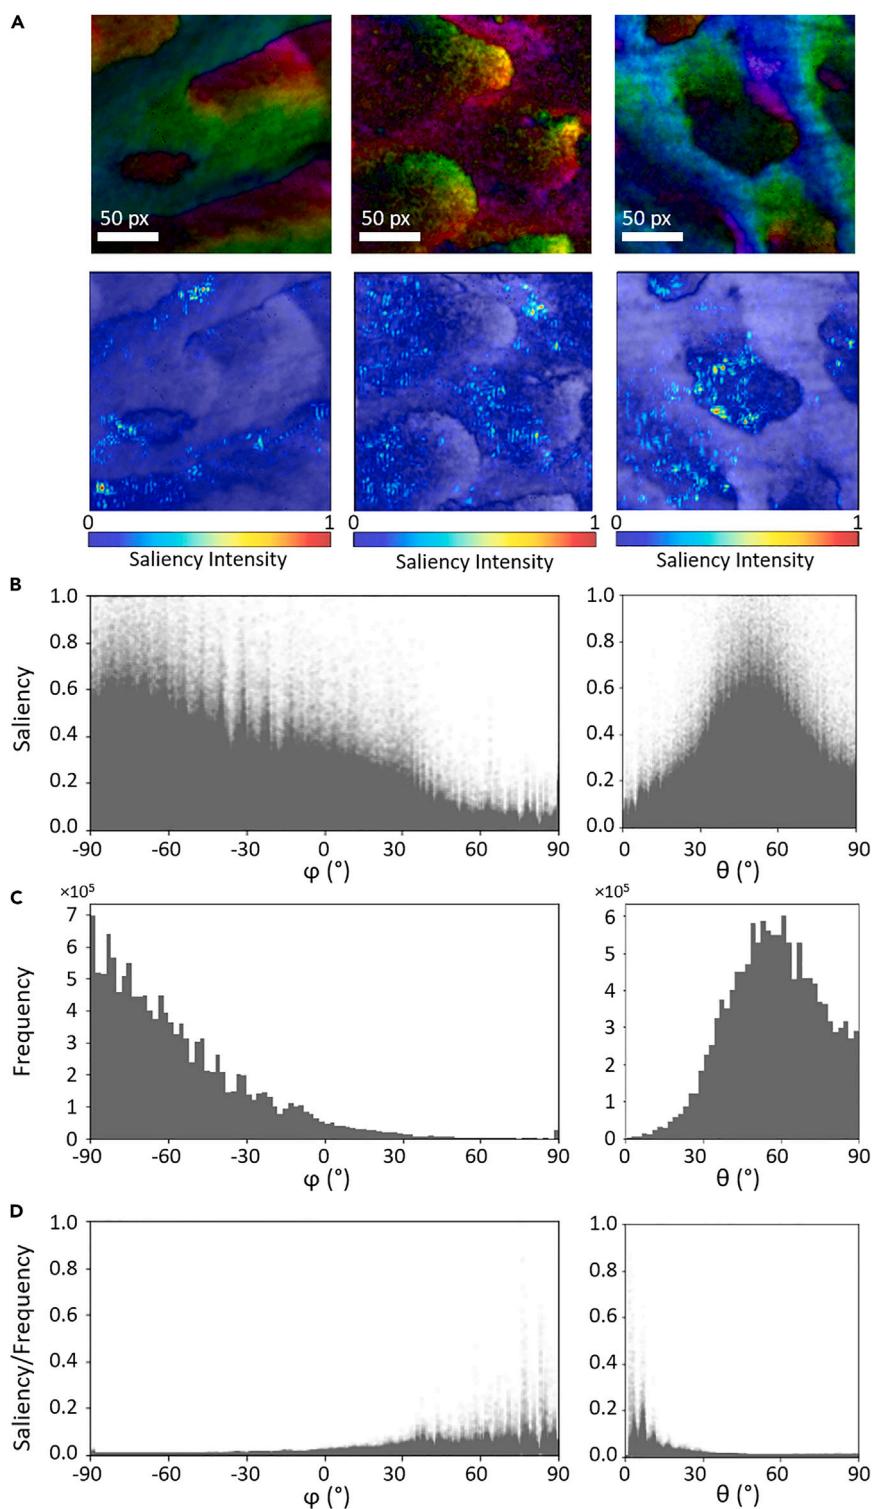


Figure 3. Understanding model predictions

To better understand how the model assigns hardness values to structures, we subject enamel images to (A) saliency mapping, which highlights regions of each image according to how much it contributes to hardness predictions. We (B) plot the saliency for each pixel of each image in our total dataset as a function of its ϕ and θ , the in-image plane and out-of-image plane orientation of the crystalline c -axis, respectively. These saliency distributions roughly correspond with

Figure 3. Continued

(C) orientation frequency histograms. A notable exception to this correspondence, however, is highlighted by (D) scaling saliency values by orientation frequency. Specifically, we find that the least-common orientations have larger relative saliences.

it may encapsulate a degree of physical understanding of hardness, as opposed to simply being a black box that memorizes random distributions of pixels that happen to give low-loss predictions without grasping real-world structure-property relationships. Beyond the post-hoc explainability techniques employed here, however, transparently understanding exactly what functions or algorithms an ML model implements in its predictions remains a non-trivial problem and a field of increasingly active research²⁴ for future work.

Applications of AI-based hardness modeling

Immediately, our hardness regression model has powerful consequences in the realm of materials characterization. The model can be utilized as a virtual indentation tool to rapidly and non-destructively probe the hardness value of another human enamel structure. This residual net (res-net) virtual indenter has the added benefit of outputting consistent results independent of which the operator runs the algorithm. Furthermore, the rapidity of res-net virtual indenter predictions allows one to map out the hardness of regions much larger than are experimentally feasible. Specifically, we can extract a sample window from an image, predict a hardness value at that location, and raster the sample window across the full image for repeated hardness predictions, as shown in Figure 4A. This model allows us to obtain high-fidelity hardness maps as in Figure 4B, where hardness values can be assigned with pixel-level resolution. This aspect of hardness mapping again surpasses experiment, as the destructive nature of experimental indentations that deform material and change local stress conditions puts a physical limit on the smallest spacing between adjacent indentations.²

Furthermore, rapid hardness mapping provides accelerated structure feedback that opens the door for non-traditional avenues of design. One can imagine manually applying knowledge of enamel structures and hardness to design bioinspired materials through a traditionally arduous process of trial and error. We take this paradigm a step further, taking inspiration from natural evolution and fueling the automation of the process with AI-based models. As a result, we are able to employ an iterative genetic algorithm^{25,26} inverse design process,^{9,27} shown in Figure 5A, to rapidly attain structures with specific hardness values.

First, we train a secondary generative model with our dataset of enamel images. Specifically, we apply the StyleGAN2 model^{28,29} on the enamel image dataset with image augmentations, which learns a 512D latent design space of enamel-inspired structures from the original images. Details of the training are provided in the [experimental procedures](#) section. By inputting an arbitrary 512D code to the trained generative model, a “fake” enamel-inspired structure is generated as output. Using this generative model, we generate a randomized population of 10 enamel-inspired structures as our starting point for design.

In addressing tasks such as engineering design, there are two broad categories of search techniques: local and global optimizers.³⁰ Local techniques, such as conjugate gradient, are highly dependent on the initial guess and tightly coupled to the solution space, often relying on smooth differentiability and continuity to efficiently reach optima.³¹ However, the 512D design space we explore here in this

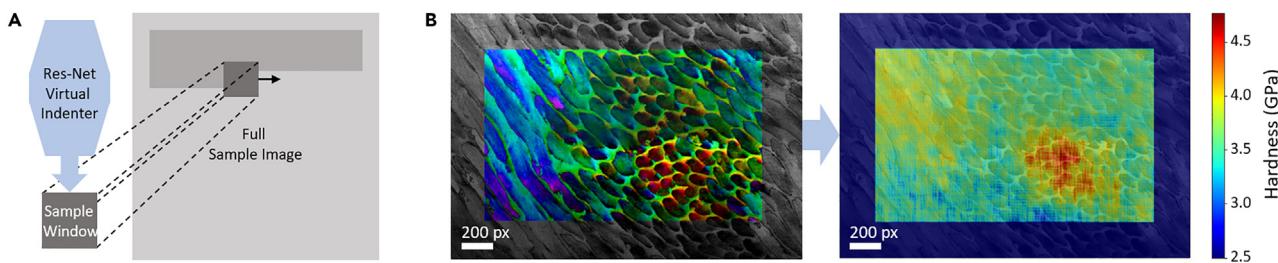


Figure 4. ML-driven hardness characterization

We can utilize the deep residual network regression model as a high-throughput and high-resolution virtual indenter by (A) rastering an observation window across a sample image and repeatedly predicting hardness values for each window. In doing so, we can identify regions of high vs. low hardness and provide (B) a non-destructive method for rapid, high-resolution hardness mapping at the nanoscale.

work has no guarantee on differentiability, with many potential local optima to fall into. Hence, we turn toward global optimization techniques that excel in situations where the topology of the solution space is unknown. Specifically, genetic algorithms are suited to constrained optimization problems with complex geometries and offer more efficient convergence than other global optimizers like random-walk searches.³²

Thus, the 512D latent codes of each structure are used as the chromosomes for a genetic algorithm. Here, a pair of structures are selected to act as parents for the subsequent generation; a process of genetic crossover is employed in order to generate two child structures; a round of specialization, migration, and mutation is employed to maintain genetic diversity in the population for combating premature convergence; and a desired hardness value is used in concert with our virtual indenter model to evaluate each structure's fitness in competition. The remaining structures after competition are fed back into the loop as the starting population for a subsequent generation. Details of implementation and each step of the genetic algorithm are provided in the [experimental procedures](#) section.

Over tens of generations, the structure population gradually drifts to have hardness values closer to the desired target. Approximately, it takes about 2 h of computation time on a workstation with a 10th Generation Intel Core i7 Processor and an NVIDIA GeForce RTX 2060 GPU to run the genetic algorithm for 250 generations. One is compelled to recall that traditional experimental hardness measurements typically occur at a rate of tens per hour⁵ compared with our res-net virtual indenter's capability to make tens of measurements every few seconds. If we assume 30 experimental hardness measurements an hour, the evaluation process for 250 generations each with 15 structures would take over 5 days of constant work. Instead, our res-net virtual indenter can accomplish these measurements (as well as automated structure generation with directed design) in less than 2% of the time.

Here, in [Figure 5B](#), we provide an example selection of 7 enamel-inspired structures generated from our genetic algorithm process with target hardness values from 3 to 6 GPa. Below each structure, we plot the average and standard deviation of population hardness over generations as the black line and gray envelope, respectively. These hardness traces are overlaid with the target hardness, plotted as the dashed red lines. Due to initial randomization of the starting population, the hardness average and standard deviation at generation 0 varies across each run. Nevertheless, in all trials, the population hardness successfully hones in on the target.

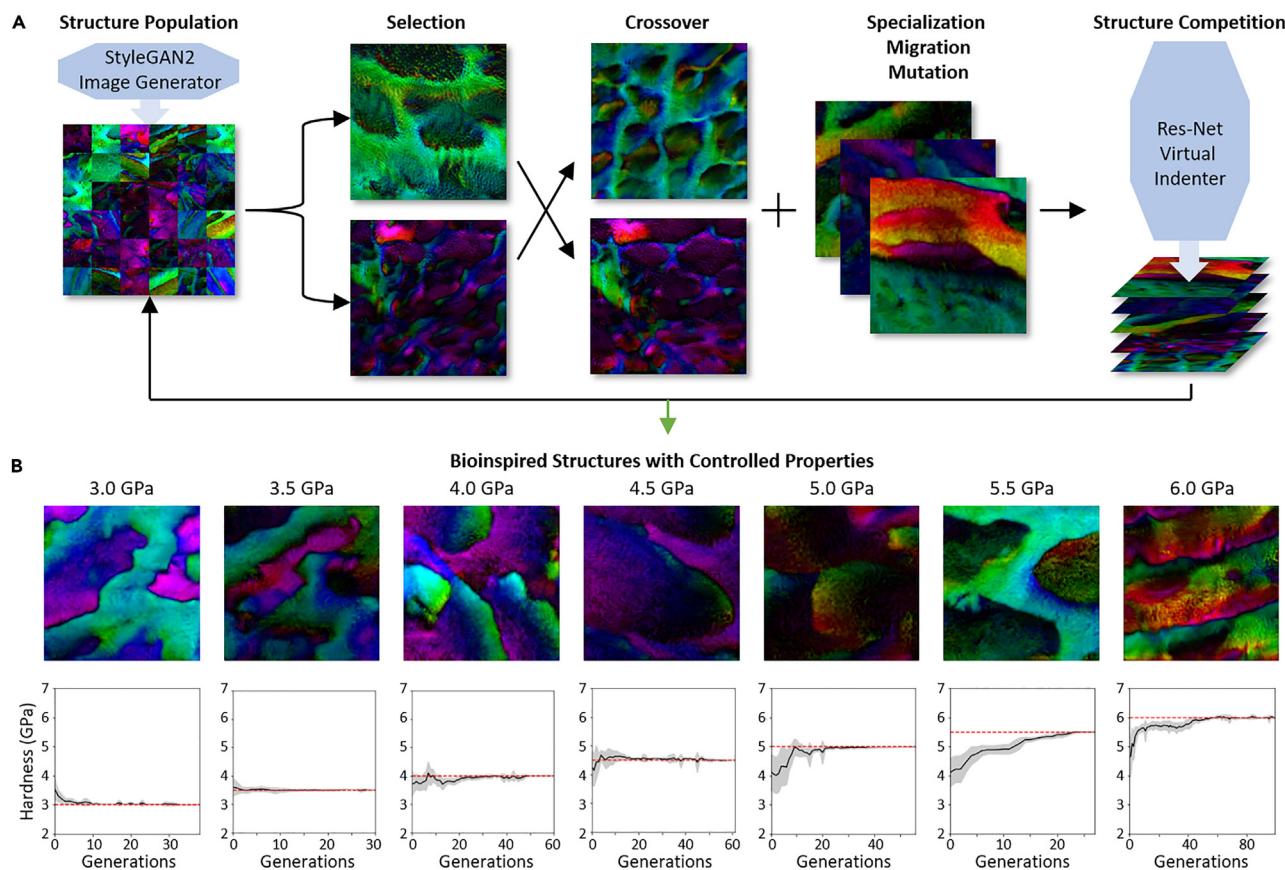


Figure 5. Targeted evolutionary design

Inspiration from the genetic processes of nature allows us to design structures with desired properties. Specifically, we (A) generate an initial population via StyleGAN2 and feed this ensemble of structures through an iterative process of selection, crossover, specialization/migration/mutation, and competition. Our regression-based virtual indenter can rapidly screen structure hardness and (B) provide the necessary feedback to drift structure populations toward desired hardness values. We run the process multiple times to obtain an example series of bioinspired “fake” enamel-like structures with predicted hardness values from 3 to 6 GPa, in increments of 0.5 GPa.

The 7 generated enamel-inspired structures vary widely in their average crystal orientation. To demonstrate that the design approach is indeed controlling the generation of specific microstructures to meet desired hardness values, and not simply assuming a simpler relationship between the distribution or average of crystal orientations with hardness,¹⁷ we set up a comparison study in Figure 6. Specifically, we take the generated enamel-inspired structures, shown in Figure 6A, and make a separate series of structures that are homogeneously comprised of each sample’s average crystal orientation in Figure 6B. Similarly, we make a second series of altered structures that maintain the exact same distribution of crystal orientations per sample, but with the rows and columns of each image randomly scrambled to remove ordered spatial arrangement of enamel microstructures in Figure 6C. When evaluated by the res-net virtual indenter in Figure 6D, we find that these averaged and scrambled structures no longer have hardness values tightly corresponding to the desired target values. This indicates that our design process indeed acts on the microstructure and arrangement of crystals, and not simply the orientation of the c-axis, to produce desired properties.

DISCUSSION

AI can be a tool for enhancing the value of information and uncovering hidden insights from historical datasets. One might have expected that previously impossible

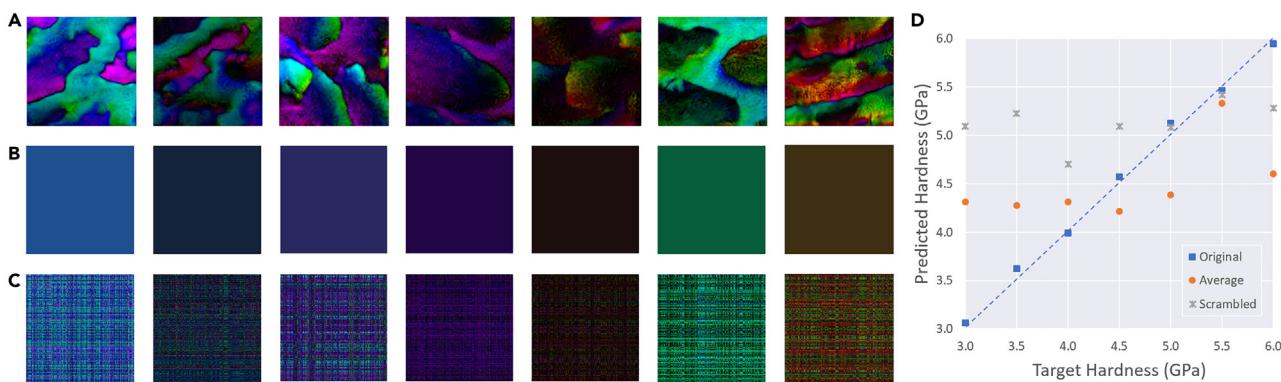


Figure 6. Design operates on a spatial microstructural level

The (A) candidate enamel-inspired structures vary in (B) average crystal orientations. To demonstrate that the approach controls the generation of specific microstructures to meet desired hardness values, we compare the predicted hardness values of the original enamel-inspired structures with structures homogeneously comprised of only the average orientation. Furthermore, we also compare the hardness values of (C) scrambled images, which contain the exact same distribution of crystal orientations as the original structures in (A) but with random spatial ordering of the rows and columns of each structure image. We find that (D) the altered structures from (B) and (C) do not achieve the tightly controlled target hardness values of (A), indicating that our inverse design approach indeed acts on spatial microstructure rather than on a simpler mapping of either crystal orientation averages or distributions.

large-scale hardness mapping and inverse design tasks would necessitate some revolution in experimental indentation methods. However, we demonstrate here that simply looking at the data we already have through alternate perspectives (i.e., via ML-driven processing workflows) is sufficient to make these tasks possible. While initial experimentation to obtain hardness measures and PIC maps may be destructive and time consuming, our approach of leveraging information from the literature to train a residual neural network for regression expands capabilities for hardness evaluation without commensurate addition of destructive characterization. After the one-time cost of training the virtual indenter model, it can be applied for rapid, non-destructive characterization without limit. In other words, to address the initial drawbacks stated about hardness testing, we have demonstrated a hardness characterization pathway that is non-destructive, high-fidelity, repeatedly consistent regardless of operator, and rapidly scalable to high-throughput testing.

We demonstrate our work focusing on a particular structure and material property of interest—human enamel and hardness. Of course, every stage of our method from data curation to image regression and from image generation to genetic algorithm optimization is portable to other datasets. There is nothing theoretically preventing our approach from treating, for example, the grain structures of some metal alloy labeled with electronic conductivities. One simply needs access to an appropriate dataset, at which point one can retrain the model with that dataset. For specifically expanding the capabilities of hardness evaluation to other similar systems, future work may be interested in mapping scanning electron microscopy (SEM) micrographs to final properties, as they are much more ubiquitous than specialized PIC mapping.

Though more sophisticated search techniques for inverse design were not necessary for this work, future work may additionally be interested in exploring variants such as adaptive genetic algorithms, in which the fitness function itself updates upon certain criteria, in order to tackle multi-objective optimization.³⁰ Such an optimizer may be useful for a design dataset incorporating multiple material properties of interest, such as both hardness and stiffness. Furthermore, other complex directed evolution

schemes³³ can incorporate the benefits of local optimization with global optimization, by iteratively cycling through periods of global structure diversification and local structure screening, to achieve rapid convergence while avoiding local minima. Algorithms such as these may be of further interest as more complex design landscapes and parameters of interest continue to be explored in the future.

With respect to the particular structures obtained here by inverse design, average rod sizes in the final enamel-inspired structures appear smaller than those in typical human enamel. This may be due to the relationship between grain boundaries and hardness—the generative model may have found it easier to tune for specific hardness by incorporating more grain boundaries in the structure that it can perturb. However, determining exactly what the model may or may not be doing points to deeper questions of model interpretability that remain an active field of research.

Currently, the success of our inverse design is heavily dependent on the genetic algorithm parameters and characteristics of the generative latent space. Namely, the genetic algorithm drives the population toward an optimized solution, but a lack of genetic diversity in structure population may still lead to premature convergence to an off-target suboptimal structure. A balance must be tuned between genetic diversity and selective pressure³⁴ in how structures are selected for crossover, mutation rates are defined, etc.—the exact nature of which may vary widely across handling different datasets. Furthermore, if the generative latent space is trained on too narrow a set of structures, a desired target may be impossible to reach no matter how tuned the genetic algorithm parameters are. Future work may also be interested in (1) devising an automated, scalable solution to parameter optimization for treating larger libraries of material structures and properties and (2) ensuring the extrapolability of generative latent space.

Additionally, as more advancements in language models come to pass, more accessible and natural design interfaces can be made. Language models such as GPT-3³⁵ have been applied to generate images from text descriptions³⁶ and to translate multi-clause sentences into intelligible images. One could imagine a situation in which one could query “Computer, I need a replacement part for this instrument, but with 20% higher hardness,” and a language model maps those words to structures, whereupon a genetic algorithm coordinates a dance between regressive and generative models to automatically design the desired part for immediate 3D printing. A chain of events like this may have been fiction not long ago, but now it is a plausible future that may yet be actualized.

Indeed, recent successful demonstrations of translating words to matter³⁷ have established the proof of concept. The technical challenges remaining are many, however, and include gathering the large datasets required to expand generality of the approach. A model that knows how to automatically design for one property in one environmental condition may be insufficient in only slightly varied conditions, and the amount of training data to account for each possible design context may be insurmountably large. Furthermore, completeness and consistency across large-volume data are difficult to achieve but important for a large-scale general model to achieve valid results. Additionally, the requests a human operator may input could be informationally incomplete. If someone requests high strength, for example, would the model optimize for yield strength, ultimate tensile strength, or some combination? Additionally, terminology across fields and domains may not be consistent, so some form of common language may be necessary to ensure that no

information is lost in translation or misinterpreted. The future may require a more interdisciplinary perspective and corpus of research before a truly general model can be constructed.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Markus J. Buehler (mbuehler@mit.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The data that support the findings of this study are available from the [lead contact](#) upon reasonable request. All original code and data has been deposited at <https://github.com/lamm-mit/HardnessMapDesign>.

Human enamel images

We use PIC maps of human enamel that provide detailed crystal c-axis orientation information¹³ labeled with hardness values obtained by nanoindentation. The color of each pixel in a PIC map displays the in-plane and out-of-plane angles of the crystalline c-axis direction, with the hue representing the in-plane orientation angle from -90° to 90° and the brightness representing the out-of-plane orientation angle from 0° to 90° . (In-plane cyan = 0° . Out-of-plane angle = 0° is full brightness of each hue, and 90° out of plane is black. All angles are measured with respect to the polarization plane, which is rotated by 60° from the image plane.) [Figure 1A](#) enamel PIC maps and nanoindentation data are debuted here, while [Figures 1B](#) and [1C](#) enamel PIC maps were previously published in Beniash et al.¹⁴ The PIC map presented was acquired before nanoindentation. All PIC maps were acquired at the Advanced Light Source, Berkeley, CA, USA, at the Ca L-edge. We acquired images at 352.8 and 352.4 eV at each linear polarization of the illuminating X-rays, repeated the acquisitions after rotating the linear polarization in 5° increments from 0° – 90° . We then did image ratios of 352.8 and 352.4 eV images and stacked the ratio images for PIC mapping analysis.

The human tooth in [Figure 1A](#), a third molar from a young adult patient, was analyzed with PIC mapping and nanoindentation to generate the accurately labeled dataset. We acquired PIC maps before and after nanoindentation to assign a precise position to each indent in the pre-indentation PIC maps. The visible light microscopy (VLM) image of indents in [Figure 1A](#) was obtained using a Zeiss AxioImager.A1m, which was used in reflected light and bright field.

Nanoindentation of human enamel

Nanoindentation tests were performed with a Triboindenter TI-950 (Bruker, Billerica, MA, USA) on the polished sample using a Berkovich diamond probe calibrated on standard fused quartz. Indentation was done on the same locations characterized by PIC mapping by exploiting specific anatomical features visible in the optical microscope of the nanoindenter. A displacement-controlled trapezoidal load function (10 s loading, 4 s holding, and 10 s unloading) with a maximum depth of 100 nm (corresponding to a peak load of approximately 1,200 μ N) was used. Areas of the same dimensions as those studied with PIC mapping were probed, with a later spacing

between indents of 2 μm . Indentation force-displacement curves were analyzed with the Oliver-Pharr method³⁸ to extract hardness H.

Pre- and post-indentation indexing of PIC maps and hardness data

Nanoindentation data were arranged in table form, including indent number, spatial coordinates, and measured H. PIC maps acquired before and after indentation, and the VLM images were imported into a single file using Adobe Photoshop 2023, with one image per layer. First, the post-indentation PIC and VLM images were aligned so the indents coincided in position. Then, the pre-indentation PIC map was aligned to the post-indentation one until all enamel rods in the image coincided with the post-indentation ones. Each indent was evaluated carefully such that it could be unambiguously located on the PIC map and therefore assigned the proper measured H value. Due to image distortions present in PIC maps, as in all microscopy methods, out of many hundreds of total indents in the area utilized, here it was only possible to identify 94 measurements with certainty.

Hardness dataset curation

To obtain one dataset that incorporates both large volume and high accuracy, we begin by treating the largest dataset with “noisy” labeling. Specifically, we use the known hardness averages and standard deviations to assign randomly sampled hardness values to 200 \times 200 pixel image regions. Any particular image is thus assigned a value within a ballpark of its real hardness value, with some unknown error or “noise.” To treat this uncertainty, we loop through each image in this set 5 times and assign a different randomly sampled hardness value each time, drawn from a Gaussian distribution defined by the known mean and standard deviation values of each image. In doing so, while we do not know the accurate hardness value of any particular image in this set, we can provide a set of likely values for each structure while neither over representing nor completely ignoring the possibility for lower probability tail cases. As a result, we obtain a set of 10,000 image-label pairs from the original noisy data.

We then treat the medium-size dataset with “estimated” labels. Specifically, we leverage the fact that these labels are more reliable than those from the noisy dataset. In other words, when training on an image from this set compared with an image from the previously described set, we want the model to consider estimated hardness values with greater weight than noisy hardness values. Thus, we resample each image-label pair in this dataset by a factor of 10 to increase their proportional representation in the total dataset. As a result, we obtain a set of 3,330 image-label pairs from the original estimated data.

Finally, we treat the small-size dataset with “accurate” labels. Similar to the logic of the previous dataset, we use the fact that these measurements are more reliable than both the noisy and estimated labels and resample each image-label pair by a factor of 10². As a result, we obtain a set of 9,400 image-label pairs from the original accurate data.

The united dataset of 22,730 image-label pairs is thus obtained by simply appending these three treated datasets together into one data frame. By processing the original data with sampling and resampling, we obtain a dataset that is both large volume and incorporates a representation of high accuracy.

Hardness prediction via regression model

We use a residual neural network to learn the regression task of mapping enamel structure images to a numerical hardness value. The architecture of the model is as [Figure 2D](#), with the residual blocks incorporating a skip connection. Specifically,

an input x is passed through a transformation $F(x)$ comprised of a 2D convolutional layer, a rectified linear unit (ReLU) batch normalization, and a second 2D convolutional layer. Then, the two $F(x)$ and x are added together with another ReLU batch normalization. In this way, the model has the ability to simply pass the identity mapping for these layers by setting $F(x)$ to zero and to combat the degradation problem of falling accuracy with excess layers.¹⁸

Training is done with a learning rate of 0.001 and batch size of 32 using the Ranger optimizer,³⁹ which combines the Rectified Adam optimizer⁴⁰ with the LookAhead optimizer⁴¹ to benefit from both a dynamic rectifier that provides a solid start to training and a secondary copy of weights that stabilizes exploration during the entirety of training. Image augmentations include a zoom range of 0.5–1, horizontal flips, and vertical flips. 70% of the images are used immediately for training, 20% used for validation, and 10% reserved for testing. Implementation is in Python with the TensorFlow⁴² 2.3.0 package.

Structure generation via StyleGAN2 model

The original curated dataset used to train the virtual indenter contained 22,730 image-label pairs, in which resampling and repetition were used to combat imperfect labeling. Now, for a model learning to generate enamel-inspired structures, only unique structures matter—not property labels. Thus, we do not include the $\times 5$ resampling of noisy labels, the $\times 10$ resampling of estimated labels, or the $\times 100$ resampling of accurate labels. As a result, there are 2,427 unique enamel images we use as a base. From this image set, we incorporate image augmentations including brightness range of 0.5–2, zoom range of 0.5–1, horizontal flips, and vertical flips. We then resized the images to 256 \times 256 pixels for compatibility with the StyleGAN2 model²⁹ and trained for 2,000 kimg. Implementation is done in Python with the TensorFlow⁴² 1.14 package.

Directed design via genetic algorithm optimization

We begin by randomly selecting two members of the structure population to serve as parents for the next generation. We can average the values of the parent chromosomes together in order to form a child structure with intermediary qualities. However, if we simply average together every structure, we quickly converge on the average hardness of the starting population. Thus, we introduce some stochasticity by using a weighted average with a randomly sampled weight. Furthermore, we generate not just one child interpolated between the parents but also a second child extrapolated beyond one of the parents. Specifically, we use two randomly sampled weights from Gaussian distributions with $\mu_1 = 0.5$, $\sigma_1 = 0.2$ and $\mu_2 = 1.5$, $\sigma_2 = 0.2$, respectively.

Furthermore, we implement three additional ways of introducing genetic diversity into our population. The first is specialization, in which we magnify the properties of a random chromosome in the population by multiplying its code by some specialization factor (SF). In this case, we use a randomly sampled SF from a Gaussian distribution with $\mu_{SF} = 1.5$, $\sigma_{SF} = 0.2$. The second method we use to introduce genetic diversity is migration, in which we introduce an entirely separate, randomly generated chromosome to the population. In this case, we pull 512 values from a uniform distribution with range –5 to 5 to create an immigrant chromosome code. The third method is mutation, in which we alter the properties of a random chromosome in the population by selecting random entries in its code R_i and replacing them with random values, MF_i , which are sampled from Gaussian distributions with $\mu_i = R_i$ and variable σ_i . In this case, we tune the chance of an entry being chosen as well as the magnitude of σ as inverse functions of the population convergence, which

we take as the fitness difference between the most and least fit structures. In this way, greater genetic diversity is introduced as the optimization progress, combating the tendency to become stuck in premature local optima.

From the initial population size of 10 structures, 5 more have been generated: interpolated child, extrapolated child, specialized structure, immigrant structure, and mutated structure. Now, we subject the entire set of 15 structures to a round of competition, in which we define a target hardness as our fitness metric and discard 5 structures to maintain a constant population size for the next generation. The trivial implementation would be to use the hardness regression model to predict the hardness of each structure and discard the 5 farthest from the target hardness value. However, this strategy risks falling into local optima, as it disincentivizes intermediary structures that may be on the right track toward a better solution. Instead, we discard only 3 unfit individuals with hardness values farthest away from the target. Then, 2 random individuals from the population that are not the current best individual are discarded. Thus, even if a proposed structure is temporarily unfit, we allow opportunities to incorporate diverse information into later generations and achieve results closer to our target.

Importantly, the hardness regression model was crafted using TensorFlow 2.3.0, whereas the StyleGAN2 model runs on the older TensorFlow 1.14. These TensorFlow versions "follow fundamentally different programming paradigm[s]"⁴³ and thus cannot trivially be run together. Furthermore, the hardness regression model also runs a more recent version of NumPy (1.19.2 vs. 1.14.3) and Python (3.8.12 vs. 3.6). While one may convert one or two of the package dependencies into one unified environment, the simpler solution we implemented was to create two separate virtual environments for the two models, respectively. Then, the genetic algorithm starts up in the Python 3.6 StyleGAN2 environment, and the hardness regression model is repeatedly called as a subprocess with its own Python 3.8.12 environment for the competition step at the end of every generation. This order was chosen because it was computationally quicker to reload the hardness regression model than the StyleGAN2 model. In this way, the hardness regression model can work together with StyleGAN2 while retaining compatibility with more recent software packages.

ACKNOWLEDGMENTS

This material is based upon work supported by the NSF GRFP under grant no. 1122374. We acknowledge support by the Office of Naval Research (N000141612333 and N000141912375), AFOSRMURI (FA9550-15-1-0514), and the Army Research Office (W911NF1920098). Related support from the MIT-IBM Watson AI Lab, MIT Quest, and Google Cloud Computing is acknowledged. P.U.P.A.G. received 40% support from the Department of Energy, Basic Energy Science, Chemical Sciences, Geosciences, Biosciences, Geosciences (DOE-BES-CSGB-Geosciences) grant DE-FG02-07ER15899 at University of Wisconsin, 40% support from award FWP-FP00011135 also from DOE-BES-CSGB-Geosciences at Lawrence Berkeley National Laboratory, and 20% support from the National Science Foundation (NSF), Biomaterials grant DMR-2220274. All PIC maps were acquired at the Advanced Light Source, a US DOE Office of Science User Facility under contract no. DE-AC02-05CH11231.

AUTHOR CONTRIBUTIONS

A.J.L and M.J.B. conceived the idea and designed the deep learning regression and generative model together. A.J.L. devised the data curation procedure, the structure and implementation of the genetic algorithm, and subsequent analyses.

P.U.P.A.G. devised the process of PIC mapping. P.U.P.A.G. and C.A.S. acquired the PIC mapping dataset. D.R., A.C., and A.T. acquired the nanoindentation hardness dataset. C.A.S., A.C., and A.T. indexed the two datasets. P.U.P.A.G. participated in all discussions of this model. M.J.B. supervised the overall project. A.J.L. wrote the manuscript, and all other authors edited it.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 30, 2022

Revised: February 11, 2023

Accepted: March 23, 2023

Published: April 17, 2023

REFERENCES

1. Callister, W.D., and Rethwisch, D.G. (2013). *Materials Science and Engineering: An Introduction*, 7th ed. (John Wiley & Sons, Inc.).
2. Sudharshan Phani, P., and Oliver, W.C. (2019). A critical assessment of the effect of indentation spacing on the measurement of hardness and modulus using instrumented indentation testing. *Mater. Des.* 164, 107563. <https://doi.org/10.1016/j.matdes.2018.107563>.
3. Tanaka, Y., Seino, Y., and Hattori, K. (2020). Automated Vickers hardness measurement using convolutional neural networks. *Int. J. Adv. Manuf. Technol.* 109, 1345–1355. <https://doi.org/10.1007/s00170-020-05746-4>. Published.
4. Kranenburg, J.M., Tweedie, C.A., Van Vliet, K.J., and Schubert, U.S. (2009). Challenges and progress in high-throughput screening of polymer mechanical properties by indentation. *Adv. Mater.* 21, 3551–3561. <https://doi.org/10.1002/adma.200803538>.
5. Hintsala, E.D., Hangen, U., Stauffer, D.D., -Bruker, N., Surfaces, E., and Prairie, U. (2018). High-throughput nanoindentation for statistical and spatial property determination. *J. Miner. Met. Mater. Soc.* 70, 494–503. <https://doi.org/10.1007/s11837-018-2752-0>.
6. Constantinides, G., Ravi Chandran, K., Ulm, F.-J., and Van Vliet, K.J. (2006). Grid indentation analysis of composite microstructure and mechanics: principles and validation. *Mater. Sci. Eng., A* 430, 189–202. <https://doi.org/10.1016/j.msea.2006.05.125>.
7. Hsu, Y.C., Yu, C.H., and Buehler, M.J. (2020). Using deep learning to predict fracture patterns in crystalline solids. *Matter* 3, 197–211. <https://doi.org/10.1016/j.matt.2020.04.019>.
8. Lew, A.J., Yu, C.-H., Hsu, Y.-C., and Buehler, M.J. (2021). Deep learning model to predict fracture mechanisms of graphene. *npj 2D Mater. Appl.* 5, 48. <https://doi.org/10.1038/s41699-021-00228-x>.
9. Lew, A.J., and Buehler, M.J. (2021). A deep learning augmented genetic algorithm approach to polycrystalline 2D material fracture discovery and design. *Appl. Phys. Rev.* 8, 041414. <https://doi.org/10.1063/5.0057162>.
10. Lew, A.J., and Buehler, M.J. (2021). Encoding and exploring latent design space of optimal material structures via a VAE-LSTM model. *Forces in Mechanics* 5, 100054. <https://doi.org/10.1016/j.finmech.2021.100054>.
11. Lew, A.J., and Buehler, M.J. (2022). DeepBuckle: extracting physical behavior directly from empirical observation for a material agnostic approach to analyze and predict buckling. *J. Mech. Phys. Solid.* 164, 104909. <https://doi.org/10.1016/j.jmps.2022.104909>.
12. Craig, R.G., and Peyton, F.A. (1958). The microhardness of enamel and dentin. *J. Dent. Res.* 37, 661–668. <https://doi.org/10.1177/00220345580370041301>.
13. Gilbert, P.U.P.A., Young, A., and Coppersmith, S.N. (2011). Measurement of c-axis angular orientation in calcite (CaCO₃) nanocrystals using X-ray absorption spectroscopy. *Proc. Natl. Acad. Sci. USA* 108, 11350–11355. <https://doi.org/10.1073/pnas.1107917108>.
14. Beniash, E., Stifler, C.A., Sun, C.Y., Jung, G.S., Qin, Z., Buehler, M.J., and Gilbert, P.U.P.A. (2019). The hidden structure of human enamel. *Nat. Commun.* 10, 4383–4413. <https://doi.org/10.1038/s41467-019-12185-7>.
15. Cuy, J.L., Mann, A.B., Livi, K.J., Teaford, M.F., and Weihis, T.P. (2002). Nanoindentation mapping of the mechanical properties of human molar tooth enamel. *Arch. Oral Biol.* 47, 281–291. [https://doi.org/10.1016/S0003-9969\(02\)00006-7](https://doi.org/10.1016/S0003-9969(02)00006-7).
16. Habelitz, S. (2015). Materials engineering by ameloblasts. *J. Dent. Res.* 94, 759–767. <https://doi.org/10.1177/0022034515577963>.
17. Stifler, C.A., Jakes, J.E., North, J.D., Green, D.R., Weaver, J.C., and Gilbert, P.U.P.A. (2021). Crystal misorientation correlates with hardness in tooth enamels. *Acta Biomater.* 120, 124–134. <https://doi.org/10.1016/j.actbio.2020.07.037>.
18. He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1512.03385>.
19. Adadi, A., and Berrada, M. (2018). Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access* 6,
20. Erhan, D., Bengio, Y., Courville, A., and Vincent, P. (2009). Visualizing Higher-Layer Features of a Deep Network, 1341 (University of Montreal), pp. 1–13.
21. Chollet, F. (2020). Visualizing what Convnets Learn (GitHub).
22. Simonyan, K., Vedaldi, A., and Zisserman, A. (2014). Deep inside convolutional networks: visualising image classification models and saliency maps. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1312.6034>.
23. Gilbert, P.U.P.A. (2018). Polarization-dependent imaging contrast (PIC) mapping in 2018. *Microsc. Microanal.* 24, 454–457. <https://doi.org/10.1017/S1431927618014514>.
24. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al. (2020). Explainable Artificial Intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>.
25. Holland, J. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (University of Michigan Press).
26. Whitley, D. (1994). A genetic algorithm tutorial. *Stat. Comput.* 4, 65–85. <https://doi.org/10.1007/BF00175354>.
27. Chakraborti, N. (2013). Genetic algorithms in materials design and processing. *Int. Mater. Rev.* 49, 246–260. <https://doi.org/10.1179/095066004225021909>.
28. Karras, T., Laine, S., and Aila, T. (2019). A style-based generator architecture for generative adversarial networks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1812.04948>.
29. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. (2020). Analyzing and improving the image quality of StyleGAN. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1912.04958>.
30. Jafar-Zanjani, S., Inampudi, S., and Mosallaei, H. (2018). Adaptive genetic algorithm for

optical metasurfaces design. *Sci. Rep.* 8, 11040–11116. <https://doi.org/10.1038/s41598-018-29275-z>.

31. Nazareth, J.L. (2009). Conjugate gradient method. *Wiley Interdiscip. Rev. Comput. Stat.* 1, 348–353. <https://doi.org/10.1002/WICS.13>.

32. Johnson, J.M., and Rahmat-Samii, Y. (1997). Genetic algorithms in engineering electromagnetics. *IEEE Antenn. Propag. Mag.* 39, 7–21.

33. Yang, K.K., Wu, Z., and Arnold, F.H. (2019). Machine-learning-guided directed evolution for protein engineering. *Nat. Methods* 16, 687–694. <https://doi.org/10.1038/s41592-019-0496-6>.

34. Gupta, D., and Ghafir, S. (2012). An overview of methods maintaining diversity in genetic algorithms. *International Journal of Emerging Technology and Advanced Engineering* 2, 56.

35. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). language models are few-shot learners. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2005.14165>.

36. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. (2021). Zero-shot text-to-image generation. Preprint at arXiv. <https://arxiv.org/abs/2102.12092>.

37. Yang, Z., and Buehler, M.J. (2021). Words to matter: de novo architected materials design using transformer neural networks. *Front. Mater.* 8, 417. <https://doi.org/10.3389/FMATS.2021.740754> [BIBTEX].

38. Oliver, W.C., and Pharr, G.M. (1992). An improved technique for determining hardness and elastic modulus using load and displacement sensing indentation experiments. *J. Mater. Res.* 7, 1564–1583. <https://doi.org/10.1557/JMR.1992.1564>.

39. Wright, L. (2019). Ranger - a Synergistic Optimizer (GitHub). <https://github.com/lessw2020/Ranger-Deep-Learning-Optimizer>.

40. Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. (2019). On the variance of the adaptive learning rate and beyond. Preprint at arXiv. <https://doi.org/10.48550/arxiv.1908.03265>.

41. Zhang, M.R., Lucas, J., Hinton, G., and Ba, J. (2019). Lookahead Optimizer: k steps forward, 1 step back. Preprint at arXiv. <https://doi.org/10.48550/arxiv.1907.08610>.

42. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). TensorFlow: a system for large-scale machine learning. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*, pp. 265–283.

43. (2022). TensorFlow 1.x vs TensorFlow 2 - behaviors and APIs. *TensorFlow Core*.