

MIT Open Access Articles

Generative pretrained autoregressive transformer graph neural network applied to the analysis and discovery of novel proteins

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Markus J. Buehler; Generative pretrained autoregressive transformer graph neural network applied to the analysis and discovery of novel proteins. J. Appl. Phys. 28 August 2023; 134 (8): 084902.

As Published: 10.1063/5.0157367

Publisher: AIP Publishing

Persistent URL: <https://hdl.handle.net/1721.1/156895>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context



Terms of use: Creative Commons Attribution



RESEARCH ARTICLE | AUGUST 29 2023

Generative pretrained autoregressive transformer graph neural network applied to the analysis and discovery of novel proteins

Special Collection: [Multiscale Modeling and Analysis in Biophysics](#)

Markus J. Buehler  



J. Appl. Phys. 134, 084902 (2023)
<https://doi.org/10.1063/5.0157367>



Nanotechnology & Materials Science


Optics & Photonics

Impedance Analysis

Scanning Probe Microscopy


Sensors

Failure Analysis & Semiconductors



Unlock the Full Spectrum.
From DC to 8.5 GHz.
Your Application. Measured.

[Find out more](#)



Generative pretrained autoregressive transformer graph neural network applied to the analysis and discovery of novel proteins

Cite as: J. Appl. Phys. **134**, 084902 (2023); doi: [10.1063/5.0157367](https://doi.org/10.1063/5.0157367)

Submitted: 7 May 2023 · Accepted: 27 July 2023 ·

Published Online: 29 August 2023



View Online



Export Citation



CrossMark

Markus J. Buehler^{1,2,a)} 

AFFILIATIONS

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

²Center for Computational Science and Engineering, Schwarzman College of Computing, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, Massachusetts 02139, USA

Note: This paper is part of the Special Topic on Multiscale Modeling and Analysis in Biophysics.

a) Author to whom correspondence should be addressed: mbuehler@mit.edu

ABSTRACT

We report a flexible language-model-based deep learning strategy, applied here to solve complex forward and inverse problems in protein modeling, based on an attention neural network that integrates transformer and graph convolutional architectures in a causal multi-headed graph mechanism, to realize a generative pretrained model. The model is applied to predict the secondary structure content (per-residue level and overall content), protein solubility, and sequencing tasks. Further trained on inverse tasks, the model is rendered capable of designing proteins with these properties as target features. The model is formulated as a general framework, completely prompt-based, and can be adapted for a variety of downstream tasks. We find that adding additional tasks yields emergent synergies that the model exploits in improving overall performance, beyond what would be possible by training a model on each dataset alone. Case studies are presented to validate the method, yielding protein designs specifically focused on structural materials, but also exploring the applicability in the design of soluble, antimicrobial biomaterials. While our model is trained to ultimately perform eight distinct tasks, with available datasets, it can be extended to solve additional problems. In a broader sense, this study illustrates a form of multiscale modeling that relates a set of ultimate building blocks (here, byte-level utf8 characters that define the nature of the physical system at hand) to complex output. This materiomorphic scheme captures complex emergent relationships between universal building block and resulting properties, via a synergizing learning capacity, to express a set of potentialities embedded in the knowledge used in training via the interplay of universality and diversity.

Significance statement: Predicting the properties of materials based on a flexible description of their structure, environment, or process, is a long-standing challenge in multiscale modeling. Our MaterioFormer language model, trained to solve forward and inverse tasks, incorporates a deep learning capacity through attention and graph strategies to yield a multimodal approach to model and design materials. Since our model is prompt-based and information is encoded consistently via byte-level utf8 tokenization, it can process diverse modalities of information, such as sequence data, description of tasks, and numbers, and offers a flexible workflow that integrates human intelligence and artificial intelligence. Autoregressive training, using pre-training against a large unlabeled dataset, allows for straightforward adjustment of specific objectives.

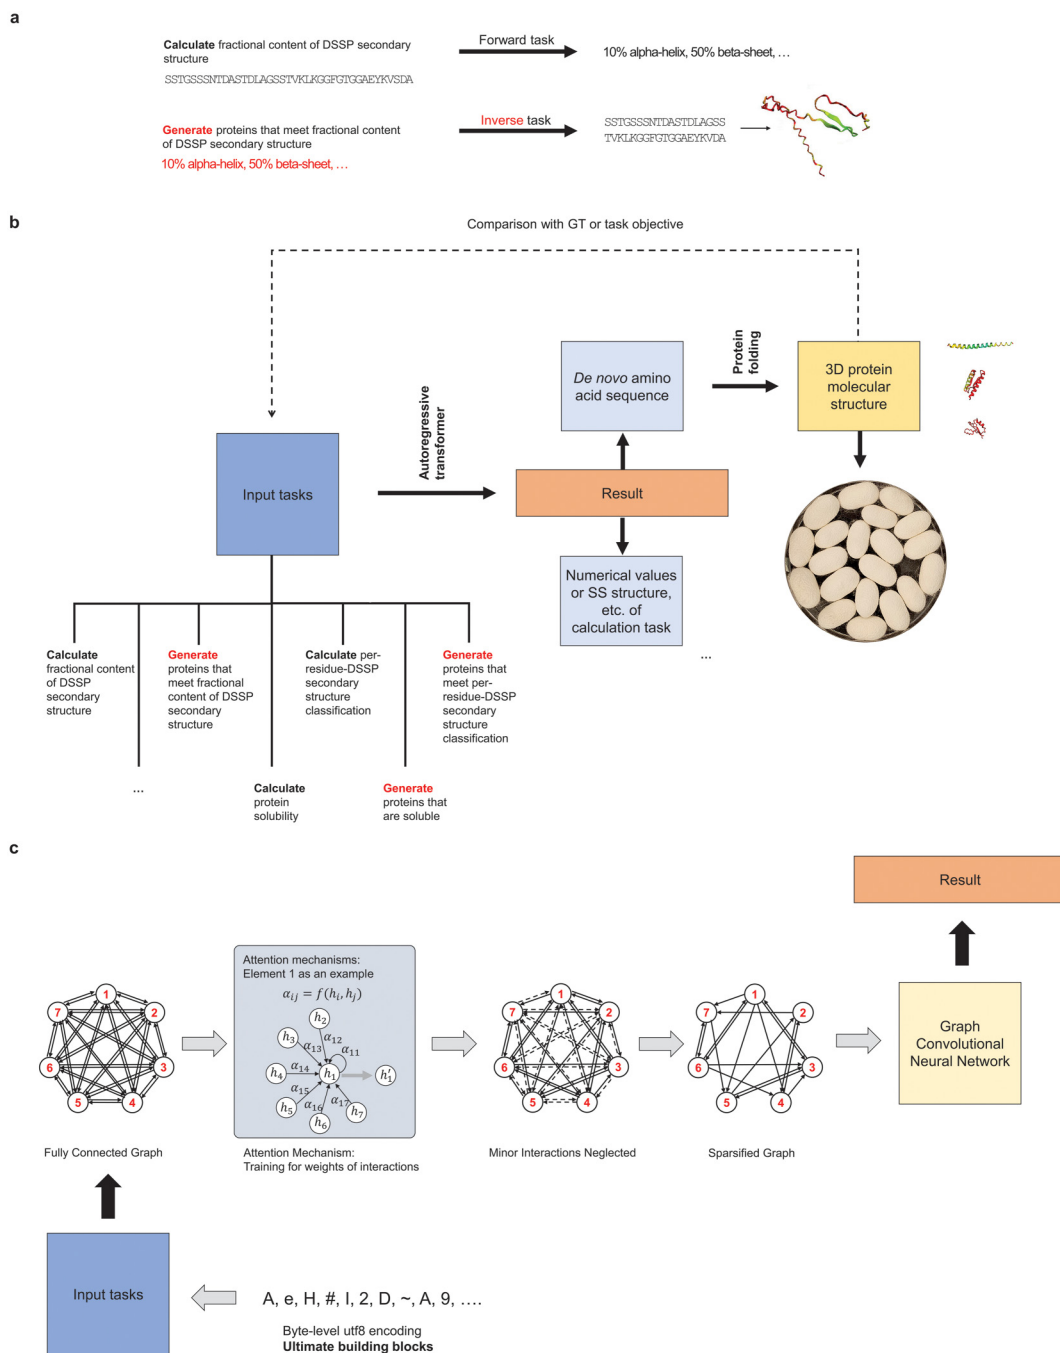
© 2023 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0157367>

I. INTRODUCTION

Multiscale modeling provides a powerful foundation for the analysis and design of hierarchical biological materials.^{1–4} Special attention is given to protein materials that form the basis for

numerous biological and biologically derived materials.^{5–7} In that realm of analysis, data-driven modeling using machine learning and related approaches has emerged as a powerful strategy^{8–14} that includes both analysis tasks (such as predicting properties from

18 September 2024, 18:20:51



18 September 2024, 18:20:51

FIG. 1. A deep language model is developed that can solve forward and inverse protein modeling problems. Panel (a) shows two sample tasks, forward (e.g., calculate secondary structure content of a protein given its sequence) and inverse (design a protein to meet a specified secondary structure content). Overview of the approach implemented, generating molecular structures from amino acid sequences (b). The model realizes a variety of calculate and generate tasks to solve multiple protein analysis and design problems. At the heart of the algorithm used here is a text-based transformer architecture that builds interaction graphs using deep multi-headed attention, which serve as the input for a deep graph convolutional neural network to form a nested transformer-graph architecture (c). In a broader sense, the modeling conducted here relates an ultimate set of building blocks—here, byte-level utf8 encoded characters—to complex output, which can take many forms. This multiscale scheme captures complex emergent relationships between the basic building block of matter and resulting properties. DSSP is the acronym that refers to the define secondary structure of proteins (DSSP) algorithm.

sequences) and inverse design tasks (designing proteins or other biomaterials to meet a set of target properties).¹⁵ Specifically, generative biomaterials science is an emerging frontier in materials discovery and has been applied to proteins,¹⁶ organic molecules, inorganics including drug design,¹⁷ bioactive materials,¹⁸ and architected materials,^{19–22} among numerous others, recently facilitated by the use of language models.²³ With the advent of attention-based transformer models in a variety of realizations,^{24–32} we are beginning to see emergent behaviors of such models³³ with important questions that should be explored specific to applications in science and engineering and, as explored here, multiscale modeling of biological protein materials.

Figure 1 shows an overview of the problem tackled in this paper, focused on solving forward and inverse problems [Fig. 1(a)]. The model features a capacity to both analyze protein sequences within the scope of end-to-end sequence-to-property predictions as well as generate molecular protein structures to meet a variety of target properties, all within a single model [Fig. 1(b)]. Our example applications include secondary structure targets (Table I) that are critical for structural, functional, and assembly properties of proteins. At the heart of the algorithm used here is a multimodal text-based autoregressive transformer architecture that builds a set of interaction graphs using deep multi-headed attention, which serve as the input for a deep graph convolutional neural network to form a nested transformer-graph architecture [Figs. 2(a) and 2(b)]. This transformer-graph architecture combines an autoregressive, causal self-attention model with a deep graph convolutional neural network [Fig. 2(b)]. Inputs to the model are completely text-based and allow, through the use of byte-level tokenization [Fig. 2(c)], flexible inputs (for details, see Materials and Methods). Due to this formulation, the model can easily be trained, and new tasks can be added. By implementing a pretraining strategy, we can endow the model with knowledge derived from unlabeled data derived from a variety of diverse protein sequences from across species (for details on datasets and training strategy, see Materials and Methods).

TABLE I. Summary of DSSP secondary structure codes used in the modeling (the tables show both, DSSP 8 and DSSP 3 codes).

DSSP 8 code	Description
h	Alpha-helix (AH)
e	Extended parallel and/or anti-parallel beta-sheet (BS) conformation
t	Hydrogen bonded turn (3, 4, or 5 turn)
~	Unstructured
b	Beta-bridge (single pair beta-sheet hydrogen bond formation)
g	3/3 ₁₀ helix
i	pi-helix
s	Bend
DSSP 3 code	Description
h	Alpha-helix (AH) (h, g, i from DSSP 8)
e	Beta-sheet (BS) (b and e from DSSP 8)
~	Unstructured (~, t, s from DSSP 8)

The plan of this paper is as follows. First, we introduce the overall approach, model development, and validation. We then cover a series of application studies focused on using the model to design new proteins with targeted properties. We focus on structural proteins and discuss how this interactive tool can be used to evolve existing proteins (here, silk protein) into new designs and how it can be used to develop proteins that incorporate antimicrobial motifs into a new protein that has high solubility, and which can also serve as a structural material to achieve multifunctionality. We conclude with a discussion and outlook to future opportunities.

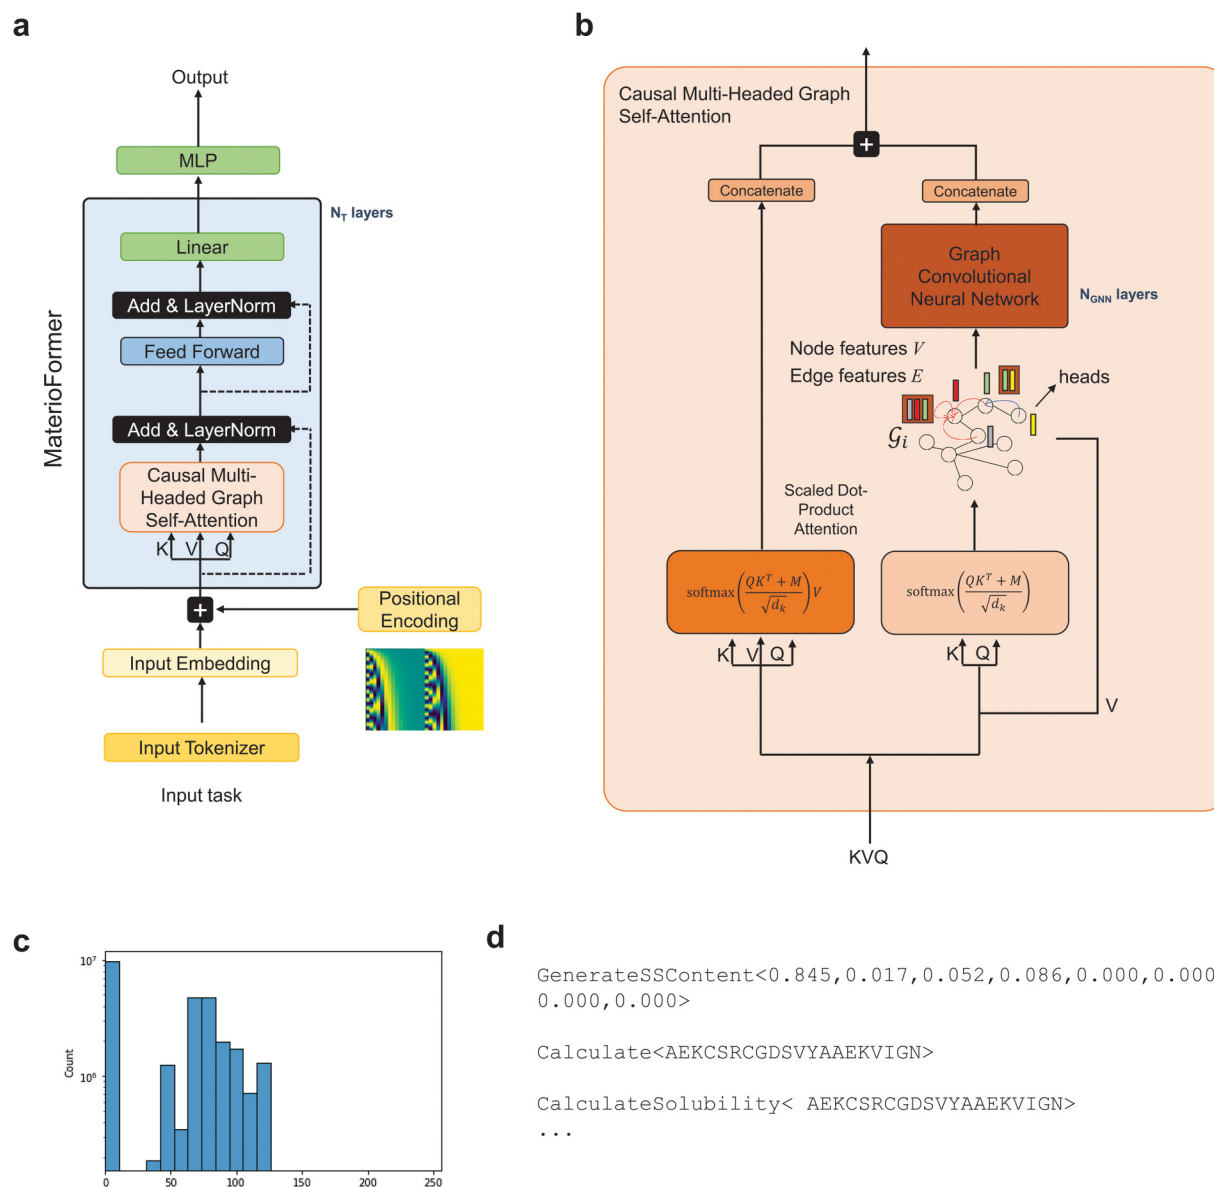
II. RESULTS AND DISCUSSION

Multi-headed attention mechanisms are used in many existing transformer models in both sequence data and graph data.³⁴ The concept to complement the attention mechanism by incorporating a graph convolutional neural network within the causal multi-headed self-attention block yields multi-headed graph-forming convolutional self-attention approach, akin to category theoretical olog models.^{31,35,36} This idea, based on viewing the attention mechanism as a graph-forming framework, allows us to exploit the complex knowledge graphs generated by the attention mechanism, one per each attention head h . We do this by performing deep graph convolutional operations on the discovered graphs. Our strategy provides a powerful framework for many other materiomorphic transformer applications, in which, we want to exploit attention and graph generation in an integrated, synergistic manner [Fig. 2(b)].

With this model, we now proceed to discuss the training strategy and the results obtained. We use a multi-stage training strategy to develop a generalizable model that progressively learns first general and then increasingly targeted and complex tasks, as shown in Fig. 3. Stage I consists of pre-training the model against unlabeled sequences (we explore various pretraining strategies, including one where 15% of the tokens are masked using a corruption strategy, inspired by the training strategy used in BERT).²⁹ The use of masking tokens adds additional complexity to the problem by teaching the model not only how to predict the correct next token, but to accomplish this task with missing information about previous sequence elements. Unlike in BERT, here, we use causal masking so that the model can only attend to tokens to the left; hence, training for both regressive predictions of next tokens and addressing the task under masked circumstances. Since the masking is randomized in each batch, pretraining is less prone to overfitting. We did not, however, find an improvement in the performance while using such masking and, ultimately, used the no masking strategy shown at the bottom of Fig. 3(b), for the results shown in the paper.

After pretraining (stage I) is complete (for around 70 000 steps), we proceed to stage II (training on forward tasks, see Table II for an overview). Figure 4 shows the performance of the forward model for the CalculateSS task [predicting overall content of the secondary structure, as defined in the define secondary structure of proteins (DSSP) algorithm, in DSSP 8]. Results are shown in Fig. 4(a), and sample secondary structure predictions are shown in Fig. 4(b). The model shows strong forward capacity for a variety of tasks, notably all integrated in a single model. The good performance suggests that there are likely synergy between the training

18 September 2024, 18:20:51



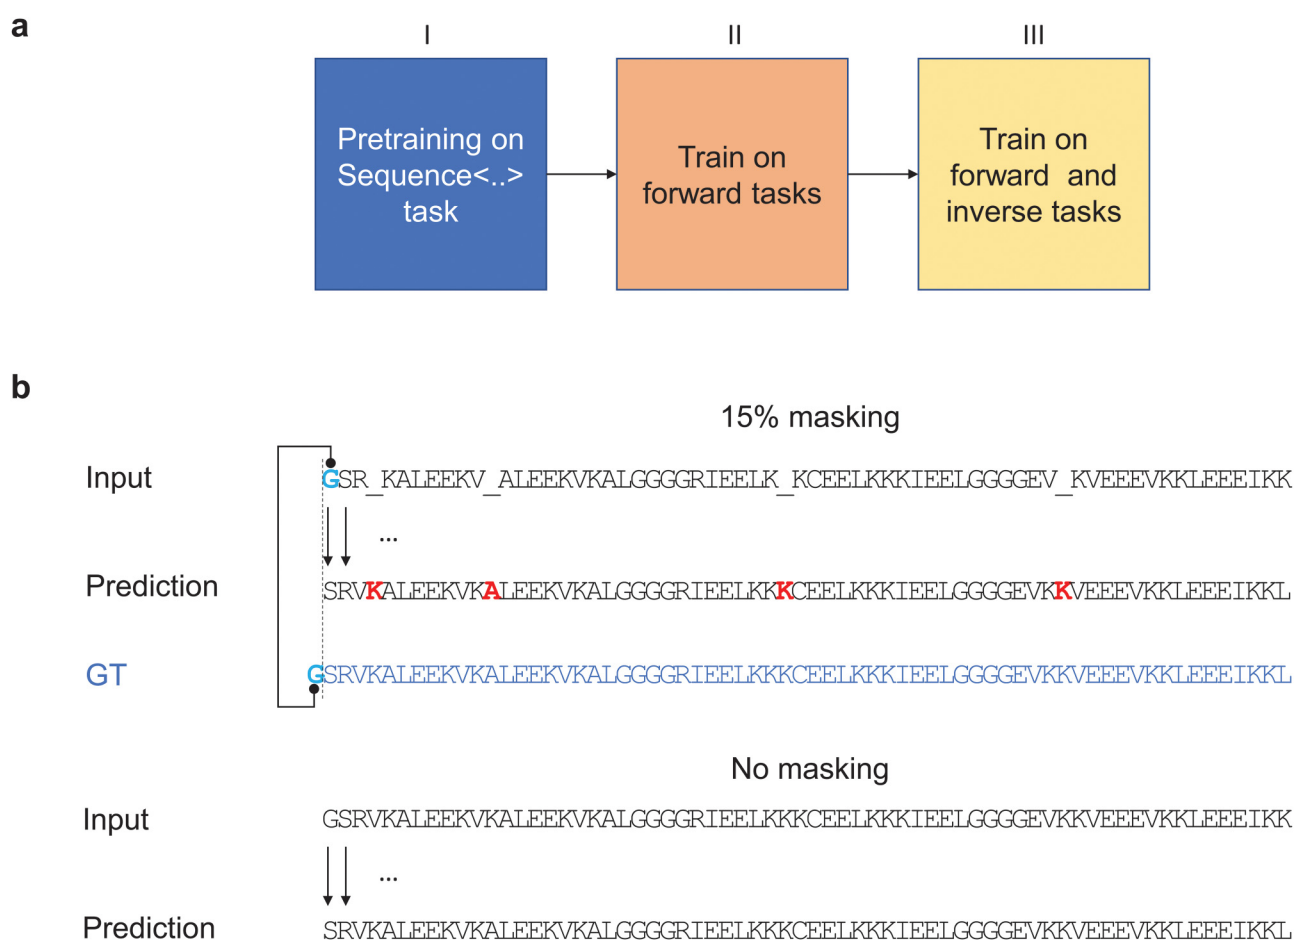
18 September 2024, 18:20:51

FIG. 2. Overview of the MaterioFormer model, an autoregressive transformer-graph convolutional model built on text-based prompt input for diverse tasks. Panel (a) depicts details of the implementation of the model, with (b) showing the causal multi-headed graph self-attention strategy used. The model features a conventional scaled dot-product attention mechanism, using causal self-attention via the triangular mask M , complemented by a graph convolutional neural network. Based on the concept schematically shown in Fig. 1(b), $\text{softmax}\left(\frac{QK^T + M}{\sqrt{d_k}}\right)$ is used to define the edge features of a set of N_{heads} nodes (N being the length of the input sequence), with node features defined by the corresponding part of V . Graph convolutional operators are applied by creating a deep graph neural network with N_{GNN} layers (hidden dimension is equal to each head dimension, generally $d_{\text{head}} = d/N_{\text{heads}}$). The message passing approach is schematically illustrated in the image, where edge information is used to scale aggregation. Panel (c) shows the statistics of the byte tokenizer, which encodes generic utf8 string data into 256 distinct tokens. The padding token is 0, most commonly seen in the padded sequence data fed to the model. Panel (d) shows sample prompts used (a complete list, see Table II).

tasks, exploited by a collective and emergent behavior captured by the model.

Next, we explore whether generative tasks can be added to the model, proceeding to training stage III. Figure 5 shows generative

tasks solved, showing examples for generating new proteins based on given ratios of the secondary structure content. The designed sequences are shown on the left, images of the folded proteins in the center, and a comparison of the design objective (labeled as



18 September 2024, 18:20:51

FIG. 3. Training strategy, featuring three stages (a). The first stage represents general-purpose masked pretraining [as shown in (b) we explore both, a strategy where we corrupt 15% of the input tokens randomly (randomized in every training step) with a masking token “_”, and a pretraining strategy without masking]. We use ~333 000 unlabeled sequences as training data. The second stage focuses on training forward tasks (calculating various protein properties), and the third stage trains on both forward and inverse tasks (designing sequences to meet a certain target). Fewer training epochs are needed from left to right, as the model learns complex relationships and, ultimately, synergistically builds on knowledge from forward and inverse tasks.

GT) with the actually obtained secondary structure content (labeled as Prediction) shown on the right (for DSSP8 and DSSP3, see Table I for definitions of the secondary structure codes).

In stage III, the model has also been trained to solve sequence, i.e., amino acid residue level, design tasks. Figure 6 shows the results of such sequence-level generative tasks, where the residue-level secondary structure is provided as an input, and proteins are designed. The result shows experimentation with design objectives of alpha-helical proteins with varying lengths. A sample task (regular font) and an output (in bold) are given as follows:

~Generate<~hhhhhhhhh~> [MSEVAALGVGALDWGKIK]\$.

Figures 6(a) and 6(b) show results for two different sampling temperatures (a, $T = 0.1$; b, $T = 0.5$). For higher sampling temperatures, proteins tend to be more diverse and novel. However, if the

temperature increases >1 , the design objectives may be less rigorously met. It is noted that the sampling temperature does not refer to real temperature units, rather it is a measure of how much Gaussian noise is added during sampling. A temperature of 1 refers to added noise with a standard deviation of 1 and, hence, indicates the point where significant effects are expected in terms of influencing the probability distributions of the predictions and, hence, the output of the model.

Figure 7 explores the effect of sampling temperature T and sampling threshold (defined as the fraction of highest rated logit candidates from which it is sampled from). The higher the temperature, the more diverse the designs become and the less they tend to adhere to the objective. Increasing the sampling threshold and the temperature provides a mechanism to yield highly diverse outcomes.

Figure 8(a) shows the design of a beta-sheet-rich protein structure, using the following prompt:

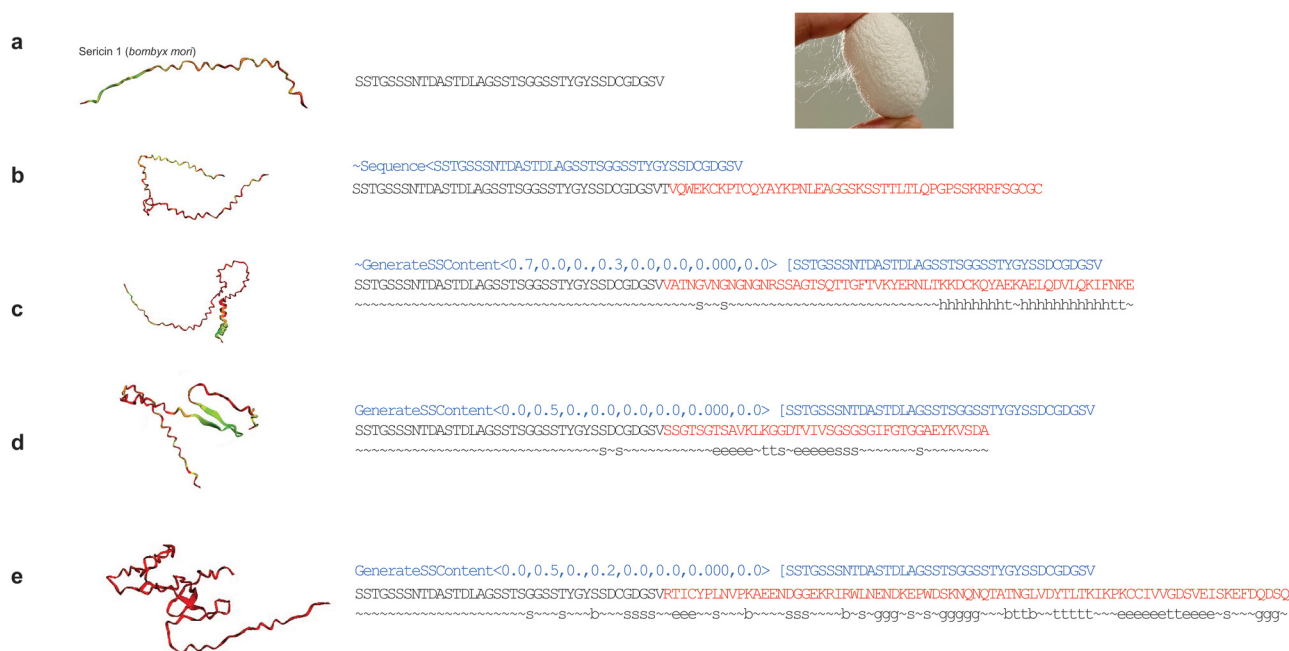


FIG. 10. Using an amino acid sequence extracted from an existing protein, Sericin 1 (*bombyx mori*, P07856, SER11_BOMMO, ser1 gene), and re-engineering the natural protein toward particular design objectives. Herein, panel (a) shows the original proteins' structure and sequence of sericin. Panel (b) shows a sequence completion task, where the initial sequence is continued in an unconstrained manner. Panel (c) shows a design task where the design objective is provided alongside the original sequence and then continued to meet the design task. The design task, in this case, is to generate an alpha-helical protein, which is indeed found toward the end of the protein. Panels (d) shows a similar example, however, with the design task to generate a beta-sheet-rich protein. This task is more difficult, but after a few trials, a solution that meets the design target is obtained. Finally, panel (e) shows that another example where the design task is given is a target with 50% beta-sheet, 20 random coils. This results in a more complex overall protein structure.

conducted using layer-wise learning rate decay⁴¹ (LLRD, a method in which we use higher learning rates for top layers and lower learning rates for bottom layers).

III. CONCLUSIONS

We have shown that generative language methods provide a flexible platform for protein materials' discovery and design. We can easily incorporate these models into a wide range of applications and solve multiple, complex tasks, as summarized in Table II. While we have considered a total of eight tasks in this study, these can be easily extended to feature additional tasks, which provide more data for the model to learn on. While our model solves these tasks overall well, there are certain advantages of using dedicated models that focus on one task at a time (e.g., sequence-to-property predictions or generative tasks using diffusion models).¹⁶ For instance, in the design task to create protein sequences that meet a certain per-residue secondary structure, the model reported in this paper sometimes fails to accurately reflect the desired length in the prediction. A similar aspect is seen when secondary structure predictions are made from an input protein sequence. In contrast, a diffusion model trained against solely one generative task¹⁶ solves it more accurately when it comes to the sequence length. However, it is noted that the model in Ref. 16 that generated sequences from

overall secondary structure contents struggled to identify novel protein designs. The model reported here can solve this task exceptionally well, with a very high degree of novel protein sequence designs.

An appealing aspect of the MaterioFormer model is the flexible, iterative workflow that can integrate human intelligence and artificial intelligence. As done in the various examples shown (Figs. 5–11), humans can enter a prompt, design a protein, and check whether it suits the design criteria (and if not, resample or adapt the design parameters), and then, use the output in a secondary task. This is demonstrated in Fig. 11 where we used an initial novel peptide design obtained in Fig. 5, as well as via the amalgamation of two naturally occurring sequences that, however, never occur jointly in a protein. Such iterative processes can also easily be combined with autonomous experimentation, providing an additional source of data generation, collection, and further training the model.

On a more theoretical side, the problem solved here is a complex building block assembly problem—building blocks are not just amino acid residues, secondary structures, but also numbers and various tasks by which these numerous combinatorial spaces are combined. Remarkably, the strategy used here learns foundational and transferrable insights. This results, as shown here, in a remarkable wealth of conditioned protein designs as well as

do not have to specially encode numerical values specifically (however, this can be done easily and would allow models to be trained to deal with very high-dimensional data, e.g., fields, images, time-series, which can be easily accomplished using vector-discretized encoding methods such as discrete variational auto-encoder as done in Ref. 42). This can be helpful for both task and prediction development and can allow for encapsulation of high-dimensional data within the architecture. There are also opportunities to introduce cross-attention mechanisms for more complex amalgamation of information processed in the attention and graph layers.

Other future explorations could incorporate additional prediction tasks in both forward and inverse directions and expand the training set to incorporate more sequences (e.g., during the pre-training stage). It would also be interesting to explore interactions with distinct biological molecules, such as mRNA or DNA, which can be added to the task training due to the flexible byte-level tokenizer. Such training tasks may also feature multiple-scale questions, such as coding not only the constituting proteins or biomolecules, but also other features such as relative concentrations, pH, or salt concentration, and others. This may ultimately be used to construct multi-modal multi-scale models that can incorporate knowledge developed from disparate simulation and experimental paradigms into all stages of training, from pretraining to tasks. A multiscale scheme, as used in this study, captures complex emergent relationships between the basic building block of matter and resulting properties. Hence, it offers a synergizing learning capacity to express a set of potentialities embedded in the foundational knowledge used to train the model that exploits unknown or little understood cross-fertilizing relationships. Mechanistically, this is facilitated by the elementary design of the approach to use a set of universal building blocks arranged in complex hierarchical patterns to create emergent functions.^{43–45}

IV. MATERIALS AND METHODS

A. Dataset construction and tokenizer

Pretraining (Fig. 3, stage I) is conducted with a dataset of ~333 000 sequences collected from the AlphaFold2 prediction database from a variety of organisms including *Humans*, *Methanocaldococcus jannaschii*, *Mouse*, *Maiz*, and many others (<https://alphafold.ebi.ac.uk/>; sequences up to a length of 256 amino acids are used, constructed from sequences for UP 000000805, UP000008816, UP000001584, UP000000625, UP000002485, UP000001450, UP000000559, UP000002311, UP000008153, UP000002195, UP000000803, UP000002296, UP000001940, UP000005640, UP000002494, UP000000589, UP000000437, UP000006548, UP000008827, UP000007305, UP000059680, and Swiss-Prot). For secondary structure tasks, we use the dataset reported in Ref. 46 that consists of 125 000 sequences, for which, overall secondary structure content and per-residue DSSP predictions have been calculated (of these, we select sequences with 128 or less amino acids or a total of ~14 500 sequences for training). Solubility is trained on the dataset reported in Ref. 37 with 40 000 sequences with associated solubility label (0 or 1) (we select sequences of 128 or less amino acids, resulting in ~4300 sequences for training and 293 test sequences; from “Test Set 1” in Ref. 37).

In the larger model, we use around 37 000 sequences of up to 512 length for training and 1792 sequences for testing solubility predictions.

We use byte-level tokenization to represent UCS Transformation Format 8 (utf8) codes in 256 tokens (in this scheme, each character is represented by one to four bytes). This strategy allows us to encode the nature of the physical system at hand, including a variety of tasks, and also opens the door to future training/fine-tuning of the model to include other sequences (e.g., amino acids, variants of natural amino acids, and DNA), tasks, or chemistries (e.g., SMILES). The distributions of tokens as obtained for the training set used in this study are shown in Fig. 2(b). Token sequences are encoded using trainable embedding layers. A physical interpretation of this strategy is that it defines all aspects of the physical including any operation applied to it, such as a design task or calculating properties. In a more abstract sense, it defines the system and tasks from an elementary building block perspective that can include numbers, characters, symbols, or other features. The model then learns to understand the relationship between these building blocks in order to solve tasks.

B. Attention based transformer and graph-convolutional models

Figures 2(a) and 2(b) depict a summary of the autoregressive transformer architecture, representing a decoder-only architecture that produces solutions iteratively from a start token during inference. The key mathematical operation is the masked attention mechanism^{23,32} defined as

$$\text{Attention}(Q, K, V; M) = \text{softmax}\left(\frac{QK^T + M}{\sqrt{d_k}}\right)V \quad (1)$$

with a triangular mask M (here, for a sequence of length 3),

$$M = \begin{pmatrix} 0 & -\infty & -\infty \\ 0 & 0 & -\infty \\ 0 & 0 & 0 \end{pmatrix}, \quad (2)$$

so that the model can only attend to tokens to the left (i.e., previous tokens to enforce causality). The causal attention calculation is implemented in the multi-headed form by using parallelly stacked attention layers, with a total dimension d . Instead of only computing the attention once, in the multi-head strategy (where h denotes the number of attention heads), we divide the input into segments (in the dimension of the hidden dimension, that is, $d_{v,i} = d/h$). We compute the scaled dot-product attention over each segment, allowing the model to jointly attend to information from different representation subspaces at different positions,

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O, \quad (3a)$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V). \quad (3b)$$

In self-attention as used in this study, all Q , K , and V come from either input or output embeddings (or other sources) only.

We complement this standard transformer architecture by constructing a per-head graph neural network architecture based on the attention scores, providing edge features. Based on the approach shown in Fig. 1(b),

$$E = \text{softmax}\left(\frac{(QK^T + M)}{\sqrt{d_k}}\right) \quad (4)$$

is used to define directed edge features of a set of h graphs \mathcal{G}_i , each of which with N nodes (N being the length of the input sequence), where E_{ij} defines the edge feature of nodes i to j ($E \in \mathbb{R}^{N \times N}$). Node features of the graph are defined by the corresponding part of V , or

$$V_i = VW_i^V \in \mathbb{R}^{N \times d_{v,i}}. \quad (5)$$

A series of graph convolutional operators are applied by creating a deep graph neural network with N_{GNN} layers (in our implementation, the hidden dimension is equal to each head dimension, $d_{v,i}$, but this could be changed in principle to allow for additional learning capacity). We use message passing to neighbors defined by all non-zero elements in E and use mean aggregation weighted by all edge features to update node features. Graph processing is conducted for each of the h graphs, and the resulting node features V_i are then concatenated to form V . This way, the output of the graph convolutional processing and the scaled dot product attention have the same dimension, $\mathbb{R}^{N \times d}$. Since M is a triangulated causal mask used in the construction of E , causality is retained in graph convolutional operators via a directed graph.

The result of the deep convolutional graph neural network and the regular multi-headed attention operation are combined additively. Gaussian error linear unit (GELU) activation functions⁴⁷ are used in both the transformer and graph convolutional neural structures.

A start token $\mathcal{T} = \sim$ is added at the beginning of the prompt, so that

$$z = [\mathcal{T}, z_1, z_2, \dots, z_N]. \quad (6)$$

During generation, the start token \mathcal{T} followed by the task is fed into the model and the output is predicted from it. During sampling iterations, this process is repeated until the full output is produced, capped using an end token $\mathcal{E} = \$$. All conditioning and distinction of various tasks are provided by the input prompt. Additional sets of tokens are defined to encapsulate various tasks and input/output boundaries ($\langle \dots \rangle$ encapsulate task, $[..]$ to encapsulate prediction. Causal autoregressive training is performed using cross-entropy loss, where the next token in the input sequence z is the label for the current token [i.e., labels start with the second token of the input, and we remove the last logit since no label exists, see Fig. 3(b)]. The training data consist of \mathcal{T} , followed by the task and the corresponding prediction, ending with \mathcal{E} ,

$$z = [\mathcal{T}, z_1, z_2, \dots, z_N, \mathcal{E}]. \quad (7)$$

Gumbel softmax sampling^{48,49} is used during inference. This allows one to adapt the creativity of the model. This is achieved by

adding a defined level of noise controlled via the sampling temperature T to a fractional set of logit distributions (identified by a sampling threshold) predicted by the transformer model. We then sample the predicted token from this revised distribution. This helps to add expressivity to the generative tasks to achieve more variations in the predictions (T around or larger than 1). In forward prediction tasks, we find that they are best conducted using low sampling temperatures ($T = 0.1$ or lower). The model features a dimension of 256, eight heads (each with $d_{v,i} = 32$ dimension), depth = 12, feed forward multiplier = 4 (1024 channels), dropout = 0.1, embedding dimension = 32, 3 GNN layers nested within each of the 12 transformer layers (the hidden dimension of the GNN is 32, same as the head dimension $d_{v,i}$). Positional encoding is realized via Fourier encoding. The total depth of the model features 12 (transformer decoder layers) \times 3 (GNN layers) = 36 total layers.

The predicted sequences are folded into 3D protein structures using OmegaFold⁵⁰ and further analyzed using DSSP^{51,52} (to obtain secondary structure information). Additional analysis to assess the novelty of sequences generated is conducted using BLAST.³⁸

Table II summarizes all prompts used in the model.

C. Training process and other hyperparameters

All codes are developed in PyTorch.⁵³ All machine learning training is performed using an Adam optimizer,⁵⁴ with a learning rate of 0.0002. We use between 2000 and 4000 warmup training steps (during which the learning rate is ramped from 0 to the desired learning rate), followed by exponential decay.

Figure 3 depicts the training strategy, featuring a total of three stages. The first stage represents general-purpose pretraining. We use both masked (15% of the input tokens are randomly masked with a masking token “_”) and unmasked pretraining. We find that the unmasked pretraining strategy yielded better results overall, but it deserves further exploration and may be advantageous in certain scenarios. For instance, we found that masked pretraining yielded a better performance in certain forward tasks. The second stage focuses on training forward tasks (calculating various protein properties), and the third stage trains on both forward and inverse tasks (designing sequences to meet a certain target). Fewer training epochs are needed from left to right, as the model learns complex relationships and, ultimately, synergistically builds on knowledge from forward and inverse tasks.

SUPPLEMENTARY MATERIAL

See Fig. S1 for designing proteins using a generative solubility task.

ACKNOWLEDGMENTS

This work was supported by the MIT-IBM Watson AI Lab, the Army Research Office (Nos. W911NF1920098 and W911NF2220213), ONR (Nos. N00014-19-1-2375 and N00014-20-1-2189), as well as USDA (No. 2021-69012-35978).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

M.J.B. developed the overall concept and the algorithm, designed the ML model, developed the codes, oversaw the work, and drafted the paper.

Markus J. Buehler: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The code and data that support the findings of this study are openly available in Github at <https://github.com/lamm-mit/MateriomicTransformer>, Ref. 55.

REFERENCES

- 1G. S. Jung and M. J. Buehler, “Multiscale modeling of muscular-skeletal systems,” *Annu. Rev. Biomed. Eng.* **19**, 435 (2017).
- 2D. L. Barreiro, J. Yeo, A. Tarakanova, F. J. Martin-Martinez, and M. J. Buehler, “Multiscale modeling of silk and silk-based biomaterials—A review,” *Macromol. Biosci.* **19**, 1800253 (2019).
- 3X. Chen and C. Drapaca, “On the dissipation of conforming and discontinuous Galerkin schemes for the incompressible Navier-Stokes equations,” *AIP Adv.* **12**(7), 075004 (2022).
- 4Y. Aboelkassem, J. D. Powers, K. J. McCabe, and A. D. McCulloch, “Multiscale models of cardiac muscle biophysics and tissue remodeling in hypertrophic cardiomyopathies,” *Curr. Opin. Biomed. Eng.* **11**, 35–44 (2019).
- 5G. Gronau, S. T. Krishnaji, M. E. Kinahan, T. Giesa, J. Y. Wong, D. L. Kaplan, and M. J. Buehler, “A review of combined experimental and computational procedures for assessing biopolymer structure-process-property relationships,” *Biomaterials* **33**(33), 8240–8255 (2012).
- 6S. Ling, D. L. Kaplan, and M. J. Buehler, “Nanofibrils in nature and materials engineering,” *Nat. Rev. Mater.* **3**(4), 18016 (2018).
- 7S. Ling, W. Chen, Y. Fan, K. Zheng, K. Jin, H. Yu, M. J. Buehler, and D. L. Kaplan, “Biopolymer nanofibrils: Structure, modeling, preparation, and applications,” *Prog. Pol. Sci.* **85**, 1–56 (2018).
- 8A. D. McCulloch, “How can AI accelerate advances in physiology?,” *J. Gen. Physiol.* **155**(6), e202313388 (2023).
- 9S. Wang, S. Sun, Z. Li, R. Zhang, and J. Xu, “Accurate de novo prediction of protein contact map by ultra-deep learning model,” *PLoS Comput. Biol.* **13**(1), e1005324 (2017).
- 10Z. Du, H. Su, W. Wang, L. Ye, H. Wei, Z. Peng, I. Anishchenko, D. Baker, and J. Yang, “The trRosetta server for fast and accurate protein structure prediction,” *Nat. Protoc.* **16**(12), 5634–5651 (2021).
- 11A. Suwardi, F. K. Wang, K. Xue, M. Y. Han, P. Teo, P. Wang, S. Wang, Y. Liu, E. Ye, Z. Li, and X. J. Loh, “Machine learning-driven biomaterials evolution,” *Adv. Mater.* **34**(1), 2102703 (2022).
- 12M. Alber, A. Buganza Tepole, W. R. Cannon, S. De, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W. W. Lytton, P. Perdikaris, L. Petzold, and E. Kuhl, “Integrating machine learning and multiscale modeling—Perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences,” *npj Digit. Med.* **2**(1), 1–11 (2019).

- 13F. Martínez-Martínez, M. J. Rupérez-Moreno, M. Martínez-Sober, J. A. Solves-Llorens, D. Lorente, A. J. Serrano-López, S. Martínez-Sanchis, C. Monserrat, and J. D. Martín-Guerrero, “A finite element-based machine learning approach for modeling the mechanical behavior of the breast tissues under compression in real-time,” *Comput. Biol. Med.* **90**, 116–124 (2017).
- 14Y. Hu and M. J. Buehler, “End-to-end protein normal mode frequency predictions using language and graph models and application to sonification,” *ACS Nano* **16**(12), 20656–20670 (2022).
- 15K. Xue, F. K. Wang, A. Suwardi, M. Y. Han, P. Teo, P. Wang, S. Wang, E. Ye, Z. Li, and X. J. Loh, “Biomaterials by design: Harnessing data for future development,” *Mater. Today Bio.* **12**, 100165 (2021).
- 16B. Ni, D. L. Kaplan, and M. J. Buehler, “Generative design of de novo proteins based on secondary structure constraints using an attention-based diffusion model,” *Chem* **9**, 1828 (2023).
- 17M. Popova, O. Isayev, and A. Tropsha, “Deep reinforcement learning for de Novo drug design,” *Science Advances* **4**(7), eaap7885 (2018).
- 18D. Merk, L. Friedrich, F. Grisoni, and G. Schneider, “De novo design of bioactive small molecules by artificial intelligence,” *Mol. Inform.* **37**(1), 1700153 (2018).
- 19A. J. Lew and M. J. Buehler, “Single-shot forward and inverse hierarchical architected materials design for nonlinear mechanical properties using an attention-diffusion model,” *Mater. Today* **64**, 10 (2023).
- 20Y.-C. Hsu, Z. Yang, and M. J. Buehler, “Generative design, manufacturing, and molecular modeling of 3D architected materials based on natural language input,” *APL Mater.* **10**(4), 041107 (2022).
- 21Z. Yang and M. J. Buehler, “Words to matter: De novo architected materials design using transformer neural networks,” *Front Mater.* **8**, 740754 (2021).
- 22K. Guo and M. J. Buehler, “A semi-supervised approach to architected materials design using graph neural networks,” *Extreme Mech. Lett.* **41**, 101029 (2020).
- 23Y. Hu and M. J. Buehler, “Deep language models for interpretative and predictive materials science,” *APL Mach. Learn.* **1**(1), 010901 (2023).
- 24Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, “Transformer-XL: Attentive language models beyond a fixed-length context,” in *ACL 2019—57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference* (Florence, Italy, 2019), pp. 2978–2988.
- 25P. Schwaller, T. Laino, T. Gaudin, P. Bolgar, C. A. Hunter, C. Bekas, and A. A. Lee, “Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction,” *ACS Cent. Sci.* **5**(9), 1572–1583 (2019).
- 26N. Kitaev, Ł. Kaiser, A. Levskaya, and G. Research, “Reformer: The efficient transformer,” *arXiv:abs/2001.04451* (2020).
- 27V. Micheli, E. Alonso, and F. Fleuret, see <https://openreview.net/forum?id=vhFu1Ac0xb> for “Transformers are sample-efficient world models” (2022).
- 28P. Esser, R. Rombach, and B. Ommer, “Taming transformers for high-resolution image synthesis,” *arXiv:abs/2012.09841* (2020).
- 29J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *NAACL HLT 2019—2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies—Proceedings of the Conference* (Florence, Italy, 2019), pp. 4171–4186.
- 30M. J. Buehler, “Modeling atomistic dynamic fracture mechanisms using a progressive transformer diffusion model,” *J. Appl. Mech.* **89**(12), 121009 (2022).
- 31M. J. Buehler, “Fieldperceiver: Domain agnostic transformer model to predict multiscale physical fields and nonlinear material properties through neural ologs,” *Mater. Today* **57**, 9–25 (2022).
- 32A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, *Advances in Neural Information Processing Systems* (Neural Information Processing Systems Foundation, 2017), pp. 5999–6009.
- 33S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M. T. Ribeiro, and Y. Zhang, “Sparks of Artificial General Intelligence: Early experiments with GPT-4,” *arXiv:2303.1271* (2023).
- 34P. Veličković, A. Casanova, P. Liò, G. Cucurull, A. Romero, and Y. Bengio, “Graph attention networks,” in *6th International Conference on Learning Representations, ICLR 2018—Conference Track Proceedings* (Vancouver, BC, 2017).

- ³⁵D. I. Spivak, T. Giesa, E. Wood, and M. J. Buehler, “Category theoretic analysis of hierarchical protein materials and social networks,” *PLoS One* **6**(9), e23911 (2011).
- ³⁶T. Giesa, D. I. Spivak, and M. J. Buehler, “Category theory based solution for the building block replacement problem in materials design,” *Adv. Eng. Mater.* **14**(9), 810 (2012).
- ³⁷M. Madani, K. Lin, and A. Tarakanova, “DSResSol: A sequence-based solubility predictor created with dilated squeeze excitation residual networks,” *Int. J. Mol. Sci.* **22**(24), 13555 (2021).
- ³⁸S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, “Basic local alignment search tool,” *J. Mol. Biol.* **215**(3), 403–410 (1990).
- ³⁹A. Datta, A. Ghosh, C. Airoidi, P. Sperandeo, K. H. Mroue, J. Jimenez-Barbero, P. Kundu, A. Ramamoorthy, and A. Bhunia, “Antimicrobial peptides: Insights into membrane permeabilization, lipopolysaccharide fragmentation and application in plant disease control,” *Sci. Rep.* **5**(1), 1–5 (2015).
- ⁴⁰J. Chen, S. Zheng, H. Zhao, and Y. Yang, “Structure-aware protein solubility prediction from sequence through graph convolutional network and predicted contact map,” *J. Cheminform* **13**(1), 1–10 (2021).
- ⁴¹T. Zhang, F. Wu, A. Katiyar, K. Q. Weinberger, and Y. Artzi, “Revisiting few-sample BERT fine-tuning,” in *ICLR 2021—9th International Conference on Learning Representations* (Virtual, 2020).
- ⁴²M. J. Buehler, “A computational building block approach towards multiscale architected materials analysis and design with application to hierarchical metal metamaterials,” *Model Simul. Mat. Sci. Eng.* **31**, 054001 (2023).
- ⁴³T. Ackbarow and M. J. Buehler, “Hierarchical coexistence of universality and diversity controls robustness and multi-functionality in protein materials,” *J. Comput. Theor. Nanosci.* **5**(7), 1193 (2008).
- ⁴⁴M. J. Buehler and T. Ackbarow, “Fracture mechanics of protein materials,” *Mater. Today* **10**(9), 46 (2007).
- ⁴⁵S. W. Cranford and M. J. Buehler, *Biomateriomics* (Springer, 2012).
- ⁴⁶C. H. Yu, W. Chen, Y. H. Chiang, K. Guo, Z. Martin Moldes, D. L. Kaplan, and M. J. Buehler, “End-to-end deep learning model to predict and design secondary structure content of structural proteins,” *ACS Biomater. Sci. Eng.* **8**(3), 1156–1165 (2022).
- ⁴⁷D. Hendrycks and K. Gimpel, “Gaussian Error Linear Units (GELUs),” [arXiv: abs/1606.08415](https://arxiv.org/abs/1606.08415) (2016).
- ⁴⁸C. J. Maddison, A. Mnih, and Y. W. Teh, “The concrete distribution: A continuous relaxation of discrete random variables,” in *5th International Conference on Learning Representations, ICLR 2017—Conference Track Proceedings* (Toulon, France, 2016).
- ⁴⁹E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with Gumbel-Softmax,” in *5th International Conference on Learning Representations, ICLR 2017—Conference Track Proceedings* (Toulon, France, 2016).
- ⁵⁰R. Wu, F. Ding, R. Wang, R. Shen, X. Zhang, S. Luo, C. Su, Z. Wu, Q. Xie, B. Berger, J. Ma, and J. Peng, “High-resolution de novo structure prediction from primary sequence,” [BioRxiv:2022.07.21.500999](https://arxiv.org/abs/2022.07.21.500999) (2022).
- ⁵¹R. P. Joosten, T. A. H. te Beek, E. Krieger, M. L. Hekkelman, R. W. W. Hoof, R. Schneider, C. Sander, and G. Vriend, “A series of PDB related databases for everyday needs,” *Nucleic Acids Res.* **39**, D411 (2011).
- ⁵²W. Kabsch and C. Sander, “Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features,” *Biopolymers* **22**(12), 2577–2637 (1983).
- ⁵³A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimsheine, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” *33rd Conference on Neural Information Processing Systems (NeurIPS2019)* (Vancouver, BC, 2019).
- ⁵⁴D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” [arXiv: abs/1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- ⁵⁵M. J. Buehler (2023). “Generative Pretrained Autoregressive Transformer Graph Neural Network applied to the Analysis and Discovery of Materials,” <https://github.com/lamm-mit/MateriomcTransformer>.