Generative artificial intelligence performs rudimentary structural biology modeling

Alexander M. Ille[1,2,3], Christopher Markosian[1,2,3], Stephen K. Burley[4,5,6,7], Michael B. Mathews[1,8], Renata Pasqualini[2,3 ‡ *], and Wadih Arap[2,9 ‡ *]

[1]School of Graduate Studies, Rutgers, The State University of New Jersey, Newark, New Jersey, USA
[2]Rutgers Cancer Institute of New Jersey, Newark, New Jersey, USA
[3]Division of Cancer Biology, Department of Radiation Oncology, Rutgers New Jersey Medical School, Newark, New Jersey, USA
[4]Research Collaboratory for Structural Bioinformatics Protein Data Bank, Institute for Quantitative Biomedicine, Rutgers, The State University of New Jersey, Piscataway, New Jersey, USA
[5]Department of Chemistry and Chemical Biology, Rutgers, The State University of New Jersey, Piscataway, New Jersey, USA
[6]Rutgers Cancer Institute of New Jersey, Robert Wood Johnson Medical School, New Brunswick, New Jersey, USA
[7]Research Collaboratory for Structural Bioinformatics Protein Data Bank, San Diego Supercomputer Center, University of California, San Diego, La Jolla, California, USA
[8]Division of Infectious Disease, Department of Medicine, Rutgers New Jersey Medical School, Newark, New Jersey, USA
[9]Division of Hematology/Oncology, Department of Medicine, Rutgers New Jersey Medical School, Newark, New Jersey, USA
‡ Equal contribution
** For correspondence: Renata Pasqualini (renata.pasqualini@rutgers.edu) or Wadih Arap (wadih.arap@rutgers.edu)*

Abstract

Natural language-based generative artificial intelligence (AI) has become increasingly prevalent in scientific research. Intriguingly, capabilities of generative pre-trained transformer (GPT) language models beyond the scope of natural language tasks have recently been identified. Here we explored how GPT-4 might be able to perform rudimentary structural biology modeling. We prompted GPT-4 to model 3D structures for the 20 standard amino acids and an α-helical polypeptide chain, with the latter incorporating Wolfram mathematical computation. We also used GPT-4 to perform structural interaction analysis between nirmatrelvir and its target, the SARS-

CoV-2 main protease. Geometric parameters of the generated structures typically approximated close to experimental references. However, modeling was sporadically error-prone and molecular complexity was not well tolerated. Interaction analysis further revealed the ability of GPT-4 to identify specific amino acid residues involved in ligand binding along with corresponding bond distances. Despite current limitations, we show the capacity of natural language generative AI to perform basic structural biology modeling and interaction analysis with atomic-scale accuracy.

Introduction

Artificial intelligence (AI)-based capabilities and applications in scientific research have made remarkable progress over the past few years [1, 2]. Advances in the field of protein structure prediction have been particularly impactful: AI-based dedicated structural biology tools such as AlphaFold2 and RoseTTAFold are capable of processing amino acid residue sequences only as input and modeling their corresponding protein structures with accuracy comparable to lower-resolution experimentally determined structures [3-5]. AlphaFold2 and RoseTTAFold are trained on datasets of protein sequence and structure datasets [6], and rely on neural network architectures specialized for modeling protein structures. Another category of AI-based tools for protein structure prediction are protein language models, which differ from AlphaFold2 and RoseTTAFold in that they are not trained on structures but rather on protein sequences [7-9]. Collectively, such protein structure prediction tools have been extensively used by researchers across various disciplines in the biological sciences and are expected to continue to add value alongside experimental structure determination [10-18].

On the other hand, generative AI language models, particularly the various generative pre-trained transformer (GPT) models powering ChatGPT [19-21], have garnered substantial interest with applications including for performance of autonomous and predictive chemical research [22, 23], drug discovery [24-26], bioinformatic analysis [27-30], synthetic biology [31], and the potential to aid in clinical diagnosis [32, 33], among others [34-36]. Unlike AlphaFold2, RoseTTAFold, and protein language models, the GPT family of language models [19, 20] are trained on natural language datasets and its neural network architecture is tailored towards understanding and generating natural language text rather than structural modeling. Our group has recently reported how GPT-4 interprets the central dogma of molecular biology and the genetic code [37]. While analogies can be made, the genetic code is not a natural language *per se*, yet GPT-4 seems to have an inherent capability of processing it. In a similar line of investigation, here we explored whether GPT-4 could perform rudimentary structural biology modeling and evaluate its

capabilities and limitations in this domain. Surprisingly, we found that GPT-4 is indeed capable of modeling the 20 standard amino acids and, with incorporation of the Wolfram plugin, a typical α-helical secondary structure element at the atomic level—albeit with some errors. Moreover, we used GPT-4 to perform structural analysis of interaction between the drug nirmatrelvir and its molecular target, the main protease of SARS-CoV-2 (the coronaviral etiology of COVID-19). To our knowledge, this is the first report to explore the specific capabilities of GPT-4 in rudimentary structural biology modeling and protein-ligand interaction analysis.

Results

Modeling of individual amino acid structures

Amino acid residues are the components of proteins, and their atomic composition and geometric parameters have been well characterized [38-40], making them suitable candidates for rudimentary structure modeling. We therefore prompted GPT-4 to model the 20 standard amino acids with minimal contextual information as input, including instructions for output in legacy Protein Data Bank (PDB) file format **(Table 1a, 2a, S1, and Fig. 1a)**. GPT-3.5 was included as a performance benchmark. Multiple iterations (n=5 for each amino acid) were run by using the same input prompt to monitor consistency (see Methods). For each individual amino acid, GPT-4 generated 3D structures with coordinate values for both backbone and sidechain atoms **(Fig. 1b)**. Generated structures contained all atoms specific to the amino acid prompted, except for a single iteration of cysteine which lacked the backbone O atom and a single iteration of methionine which lacked sidechain Cγ atom. Most amino acid structures (excluding achiral glycine) were modeled in L rather than D stereochemical configuration, while some were also modeled in planar configuration **(Fig. 1c)**. While the modeling favored the L-configuration, a more accurate distribution would be near exclusive L-configuration, given that D-amino acid residues are only rarely found in naturally occurring proteins [41].

Backbone bond lengths and angles of the modeled structures varied in accuracy, yet clustered in approximation to experimentally determined reference values [38, 39] **(Fig. 1, d and e)**. Moreover, all reference values fell within the standard deviations of backbone bond lengths and angles of the modeled structures. Finally, sidechain bond lengths and bond angles also varied in accuracy, yet >89% of calculated bond lengths were within 0.1 Å and >79% of calculated bond angles were within 10° of experimentally determined reference values [40, 42], again indicating remarkable precision **(Fig. 1f-h, Fig. S1-S4)**. Sidechain bond lengths and bond angles outside of

these ranges generally occurred at random, but were notably more prevalent in the aromatic rings of histidine and tryptophan, along with the pyrrolidine component of proline. Although not entirely error-free, the ring structures of phenylalanine and tyrosine were more accurate, which may be due to the reduced complexity of their all-carbon ring composition. Across all parameters assessed, GPT-4 substantially outperformed GPT-3.5. Collectively, these findings demonstrate that GPT-4 is capable of structurally modeling single amino acid residues in a manner that resembles their experimentally-determined structures, though not without sporadic errors including incorrect stereochemistry and geometric distortion, which would require—at least presently—human operator curation or supervision to ensure fidelity.

Table 1. Prompts used for structural modeling. (A) Prompt used for modeling the structures of each of the 20 amino acids with GPT-4 and GPT-3.5. The same prompt was used for each amino acid by replacing "[amino acid]" with the full individual amino acid name. **(B)** Prompt used for modeling the α-helical polypeptide structure with GPT-4 running with the Wolfram plugin for enhanced mathematical computation. **(C)** Prompts used for structural drug interaction analysis of nirmatrelvir bound to the SARS-CoV-2 main protease (PDB ID: 7VH8).

(A) Amino acid structure modeling with GPT-4 and GPT-3.5	
Prompt: What are the typical distances and angles between the atoms of one [amino acid] residue in a protein? Based on these values, generate a structure in PDB file format for one [amino acid] residue. Ensure coordinate values have three decimal places and omit hydrogen atoms.	
(B) α-helix structure modeling with GPT-4 running the Wolfram plugin	
Initial Prompt: What are the typical geometric attributes (including distances and angles) between the backbone atoms of a typical alpha-helical polypeptide chain? Based on this information, generate a structure in PDB file format for an alpha-helical polypeptide chain 10 residues in length including only alpha carbon atoms. Ensure coordinate values have 3 decimal places and use the Wolfram plug-in for coordinate calculations but not for PDB file formatting.	
Refinement Prompts	The generated structure does not resemble an alpha-helix. Try again.
	The diameter of the helix is too large.
	The diameter of the helix is too small.
	The pitch (residues per turn) of the helix is too large.
	The pitch (residues per turn) of the helix is too small.
(C) Structural drug interaction analysis with GPT-4	
Ligand detection prompt:	Based on structural information, what is the ligand that is present in the attached PDB file?
Interaction detection prompt:	Analyze the structure to detect up to 5 amino acid residues in the protein chain which have important bond interactions with the "4WI" ligand without importing

	any external libraries. List the residues and the bond distances. Based on this information, predict potential mutations in the protein chain which would interfere with binding of the "4WI" ligand.

Table 2. Exemplary coordinate output from GPT-4. Within responses to modeling prompts **(Table 2, S1, and S2)** GPT-4 provided coordinates for generated structures in PDB file format, as shown below with output examples for **(A)** the arginine amino acid structure and **(B)** the α-helix structure.

(A) Output for arginine structure PDB file format	**(B) Output for α-helix structure PDB file format**
ATOM 1 N ARG A 1 0.000 0.000 0.000 ATOM 2 CA ARG A 1 1.460 0.000 0.000 ATOM 3 C ARG A 1 2.160 1.370 0.000 ATOM 4 O ARG A 1 2.160 2.170 -1.080 ATOM 5 CB ARG A 1 2.090 -1.180 -0.670 ATOM 6 CG ARG A 1 2.790 -2.440 0.110 ATOM 7 CD ARG A 1 3.420 -3.620 -0.560 ATOM 8 NE ARG A 1 4.120 -4.880 0.220 ATOM 9 CZ ARG A 1 4.750 -5.170 1.380 ATOM 10 NH1 ARG A 1 4.750 -4.450 2.490 ATOM 11 NH2 ARG A 1 5.350 -6.340 1.550	ATOM 1 CA ALA A 1 2.300 0.000 0.000 ATOM 2 CA ALA A 2 -0.399 2.265 1.500 ATOM 3 CA ALA A 3 -2.161 -0.787 3.000 ATOM 4 CA ALA A 4 1.150 -1.992 4.500 ATOM 5 CA ALA A 5 1.762 1.478 6.000 ATOM 6 CA ALA A 6 -1.762 1.478 7.500 ATOM 7 CA ALA A 7 -1.150 -1.992 9.000 ATOM 8 CA ALA A 8 2.161 -0.787 10.500 ATOM 9 CA ALA A 9 0.399 2.265 12.000 ATOM 10 CA ALA A 10 -2.300 0.000 13.500

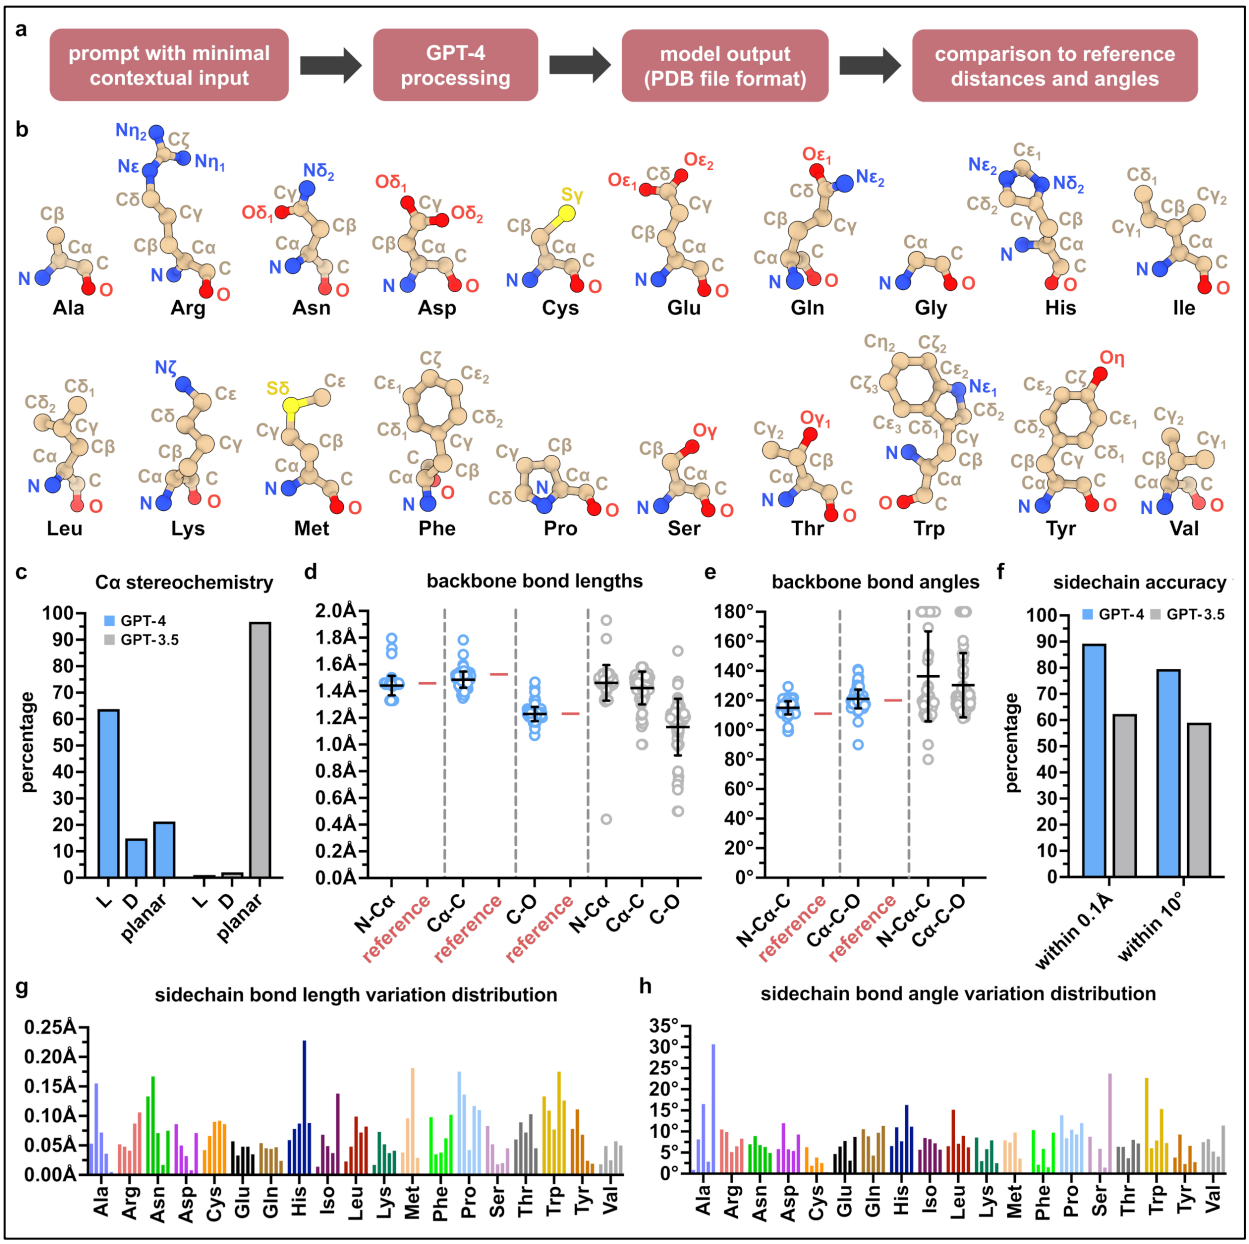

Fig. 1 Modeling the 3D structures of the 20 standard amino acids with GPT-4. (a) Procedure for structure modeling and analysis. **(b)** Exemplary 3D structures of each of the 20 amino acids modeled by GPT-4. **(c)** Cα stereochemistry of modeled amino acids including L and D configurations as well as nonconforming planar, n=5 per amino acid excluding achiral glycine and one GPT-4 iteration of cysteine (see methods). **(d-e)** Backbone bond lengths and angles of amino acids modeled by GPT-4 (blue) relative to experimentally determined reference values (red), n=5 per amino acid, excluding one iteration of cysteine (see methods). Corresponding values of amino acids modeled by GPT-3.5 shown adjacent (grey), n=5 per amino acid. Data shown as means ± SD. **(f)** Sidechain accuracy of modeled amino acid structures in terms of bond lengths (within 0.1 Å) and bond angles (within 10°) relative to experimentally determined reference values, n=5 per amino acid. See Methods for experimentally determined references. **(g-h)** Distributions of

sidechain bond length and angle variation relative to experimentally determined reference values for each amino acid generated by GPT-4, excluding glycine. Bars represent the mean bond length or angle variation for each of the five iterations per amino acid. One of the methionine iterations was excluded (see methods).

Modeling of an α-helix structure

The α-helix is the most commonly occurring and extensively studied secondary structure element found in proteins [43-46]. Thus, we next prompted GPT-4 and GPT-3.5 to model an α-helical polypeptide chain, but were unable to obtain accurate structures with either version, despite multiple attempts with various prompts. We then incorporated the Wolfram plugin, a mathematical computation extension developed by Wolfram-Alpha for use with GPT-4 [47]. GPT-4 used together with the Wolfram plugin was able to model a 10-residue α-helical structure and output the result in legacy PDB file format with minimal contextual information as input **(Table 1b, 2b, S2, and Fig. 2a)**. Multiple iterations were run by using the same input prompt to monitor consistency, and up to two prompt-based refinements after the first attempt were permitted per iteration for improved accuracy (see Methods). To reduce complexity, only Cα atoms were modeled.

GPT-4 arbitrarily assigned all residues as alanine, which was likely done for the sake of simplicity, but nevertheless aligns well with the fact that alanine has the greatest α-helix propensity of all 20 standard amino acids [46]. Remarkably, accuracy of the modeled α-helix was comparable to an experimentally determined α-helical structure consisting of 10 consecutive alanine residues (PDB ID: 1L64) [48] **(Fig. 2, b and c)**. Forty-four percent of modeled structures had a root-mean-square deviation (RMSD) of < 0.5 Å relative to the reference experimental structure on the first attempt, and 88% had a RMSD of < 0.5 Å after two prompt-based refinements **(Fig. 2d)**. The structures generated by GPT-4 were also compared to poly-alanine α-helix structures modeled by AlphaFold2, ChimeraX, and PyMOL, and the lowest RMSDs (*i.e.*, greatest structural similarity) were found between the AlphaFold2 and ChimeraX structures **(Fig. 2b)**. Taken together, these results demonstrate the capability of GPT-4, with a seamless incorporation of the Wolfram plugin, to predict the atomic level structure of an α-helix.

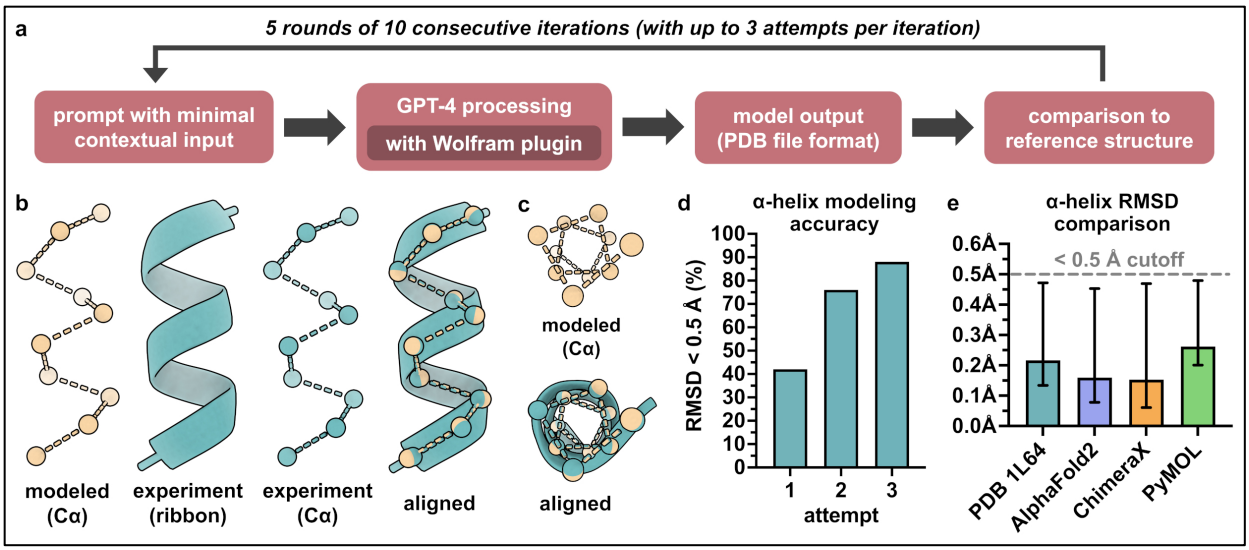

Fig. 2 Modeling the 3D structure of an α-helical polypeptide structure with GPT-4. (a) Procedure for structure modeling and analysis. **(b)** Exemplary 3D structure of a modeled α-helix (beige), an experimentally determined α-helix reference structure (PDB ID 1L64) (teal), and their alignment (RMSD = 0.147 Å). **(c)** Top-down view of modeled and experimental α-helices from panel B. **(d)** Accuracy of α-helix modeling as measured by number of attempts (including up to two refinements following the first attempt) required to generate a structure with RMSD < 0.5 Å relative to the experimentally determined reference structure, n = 5 rounds of 10 consecutive iterations (total n = 50 models). **(e)** Comparison of RMSDs between GPT-4 α-helix structures and the experimentally determined α-helix structure, the AlphaFold2 α-helix structure, the ChimeraX α-helix structure, and the PyMOL α-helix structure. Only structures with RMSD < 0.5 Å (dashed grey line) relative to each reference structure are included (88% included in reference to PDB ID 1L64; 90% to AlphaFold; 90% to ChimeraX; 88% to PyMOL). Data shown as means ± range.

Structural interaction analysis

Structural interaction between drugs and proteins is a key aspect of molecular biology with basic, translational and clinical implications. For instance, binding of the Paxlovid (ritonavir-boosted nirmatrelvir) protease inhibitor compound, nirmatrelvir, to the SARS-CoV-2 main protease is of particular clinical relevance [49, 50], especially given the concern that the mutation-prone SARS-CoV-2 leads to treatment resistance [51]. Thus, we used GPT-4 to perform qualitative structural analysis of drug binding within the nirmatrelvir-SARS-CoV-2 drug-protein paradigm. We first provided the PDB file input of a crystal structure of nirmatrelvir bound to the SARS-CoV-2 main protease (PDB ID: 7VH8) [49] and prompted GPT-4 to detect the nirmatrevir

ligand, followed by a subsequent prompt for interaction detection and interaction-interfering mutation prediction **(Table 1c, S3, and Fig. 3a)**.

GPT-4 correctly identified the nirmatrelvir ligand, which in the input PDB file is designated as "4WI" **(Table S3)**. For interaction detection, GPT-4 listed five amino acid residues within the substrate-binding pocket of the protein, four of which directly bind the nirmatrelvir ligand (Cys145 forms a covalent bond, His163 and His164 each form hydrogen bonds, and Glu166 forms three separate hydrogen bonds) [49] **(Fig. 3b)**. The fifth residue (Thr190) does not form a bond with the ligand, but is located within the binding pocket [49]. Moreover, the distances provided by GPT-4 for the four binding residues correspond precisely to the distances between the interacting atoms, information which is not inherent in the input PDB file. GPT-4 also described several mutations which may interfere with binding **(Table S3)**, and while most were plausible, others would likely be inconsequential. Notably, however, the suggested mutation of Glu166 to a residue lacking negative charge has been documented to be critically detrimental to nirmatrelvir binding [52-54] and confers clinical therapeutic resistance [55, 56]. Altogether, this exercise reveals the ability of GPT-4 to perform basic structural analysis of ligand interaction in a manner which, in conjunction with molecular analysis software such as ChimeraX, highlights its practical utility.

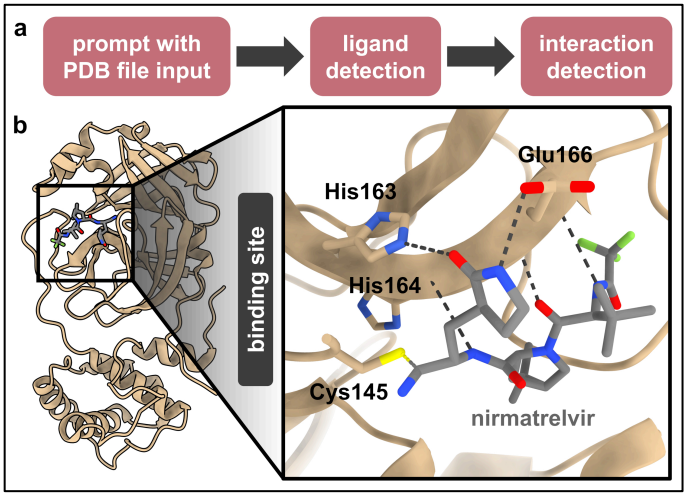

Fig. 3 Structural analysis of interaction between nirmatrelvir and the SARS-CoV-2 main protease. (a) Procedure for performing ligand interaction analysis. **(b)** Crystal structure of nirmatrelvir bound to the SARS-CoV-2 main protease (PDB ID: 7VH8) with bond-forming residues detected by GPT-4, and their bonds depicted with ChimeraX (inset). Distances between interacting atom pairs were 1.81Å (Cys145 Sγ–C3), 2.68 Å (His163 Nε2–O1), 2.77Å (Glu166 O–N4), 3.02Å (Glu164 O–N1), as determined by GPT-4 and 1.814Å (Cys145 Sγ–C3), 2.676 Å (His163 Nε2–O1), 2.767Å (Glu166 O–N4), 2.851Å (Glu166 N–O3), 3.019Å (Glu166 Oε1–N2), 3.017Å (Glu164 O–N1), as determined with ChimeraX. Note that distance values

corresponding to the Glu166 N–O3 and Glu166 Oε1–N2 atom pair interactions were not provided by GPT-4.

Discussion

The proof-of-concept findings reported here show the current capabilities and limitations of GPT-4, a natural language-based generative AI, for rudimentary structural biology modeling and drug interaction analysis. This presents a unique aspect of novelty, given the inherent distinction between natural language models and other dedicated AI tools commonly used for structural biology, including AlphaFold, RosettaFold, and protein language models. While such tools are unequivocally far more advanced in terms of the scale of molecular complexity they are able to process, GPT-4 sets the stage for a broadly accessible and computationally distinct avenue for use in structural biology and protein-ligand interaction analysis.

The performance of GPT-4 for modeling of the 20 standard amino acids was favorable in terms of atom composition, bond lengths, and bond angles. However, stereochemical configuration propensity and modeling of ring structures require improvement. Performance for α-helix modeling, with a seamless incorporation of advanced mathematical computation from the Wolfram plugin, was also favorable. While the requirement of prompt-based refinements may be viewed as a limitation, they may also serve as a means and opportunity by which the user can optimize and modify a structure. Nonetheless, improvements will be required in the capacity to model more complex all-atom structures, not only Cα backbone atoms.

Of note, these structural modeling capabilities also raise the question of modeling methodology, especially since GPT-4 was not explicitly developed for this specialized purpose. It would be challenging to provide a precise answer for this, and several (proprietary or non-proprietary) computational methods may be involved. For instance, GPT-4 may be utilizing pre-existing atomic coordinate information present in its broad training dataset, which includes "publicly available data (such as internet data) and data licensed from third-party providers" [20]. However, this reasoning does not adequately explain the geometric variability observed in the predicted structures, and why structural complexity appears to be a limiting factor. The modeling may also be performed *ab initio*, given that the generated responses often articulate geometric parameters (*e.g.*, specific bond lengths and angles, number of amino acid residues per α-helix turn, α-helix diameter, etc.) in addition to providing atomic coordinates **(Table S1 and S2)**. Alternatively, the modeling methodology may involve both the use of pre-existing coordinates plus *ab initio* computation.

Of note, the comparison of the α-helix model generated by GPT-4 with those generated by other computational tools was quite revealing. AlphaFold2, as mentioned above, predicts structures based on training data consisting of protein sequences and 3D structures, and was developed specifically for modeling protein structures. In addition to their dedicated molecular analysis capabilities, ChimeraX and PyMOL may be used to model basic, idealized secondary structure elements in a manner which narrowly considers precise predefined geometries, thus providing accurate α-helix structures. Despite not being explicitly developed to model atomic coordinates for α-helical segments of protein chains, GPT-4 was able to generate an α-helix with accuracy comparable to the structures modeled by the above tools. The requirement of adding the Wolfram plugin likely suggests that mathematical computation is heavily relied upon by GPT-4 for α-helix modeling. Yet, α-helix structural properties and self-instruction are generated by GPT-4 prior to engaging the Wolfram plugin **(Table S2)**, suggesting that some degree of intrinsic "reasoning" might perhaps be involved. So-called reasoning, in this regard, refers to the documented performance of benchmark reasoning exercises by GPTs [19, 20], and it should be noted that there is ongoing debate about what constitutes reasoning as it pertains to AI [57, 58].

The exercise exploring the capability of GPT-4 to perform structural analysis of ligand-protein binding showed promise, especially given the clinical relevance of the protein binding interaction between nirmatrelvir and the SARS-CoV-2 main protease. Ligand detection was expected to be a straightforward task, as PDB files include unique designations for various molecular entities. Interaction detection was surprisingly well-handled, considering the complexity of locating amino acid residues with spatial proximity to the ligand and providing precise distances between interacting atoms. Based on the generated response **(Table S3)**, it is likely that proximity was the primary criteria used by GPT-4 for interaction detection. While proximity is important, the analysis would benefit from additional criteria such as hydrophobicity, electrostatic potential, solvent effects, etc. These criteria would likely improve the prediction of interaction-interfering mutations.

Considering both strengths and weaknesses, the structural modeling capabilities of GPT represent an intriguing aspect of the unprecedented advancement of natural language-based generative AI, a transformative technology still in its infancy. While this modeling remains rudimentary and is currently of limited practical utility, it establishes an immediate and direct precedent for applying this technology in structural biology as generative AI natural language models undergo continued development and specialization. Concurrently, this broadly-accessible technology presents opportunity for structural analysis of drug-protein interaction. In the interim, further research on the capabilities and limitations of generative AI is merited, not only in structural biology but also for other potential applications in the biological sciences.

Methods

Prompt-based modeling with GPT-4

Modeling of individual amino acid structures was performed by challenging GPT-4 through the ChatGPT interface [20, 21] with a single prompt **(Table 1a)**, one amino acid residue at a time. For each individual amino acid, the same prompt was used for five consecutive iterations with each iteration initiated in a new dialog. GPT-4 was run in classic mode without “browser” and “analysis” features enabled, formerly known as “web browser” and “code interpreter” plug-ins, respectively [59]. Classic mode limits processing to GPT-4 with no additional capabilities. Amino acid modeling was also performed with GPT-3.5 in the same manner. However, GPT-3.5 would frequently generate PDB file output with missing or extra atoms. In such cases, responses were regenerated within each GPT-3.5 dialog until PDB file output contained the correct number of atoms required for analysis. Modeling of α-helix structures was performed by challenging GPT-4 running the Wolfram plugin [60] through the ChatGPT interface with an initial prompt followed by up to two refinement prompts in the same dialog, for a total of up to three attempts **(Table 1b)**. The same prompt was used for five rounds of ten consecutive iterations with each iteration initiated in a new dialog.

Analysis of generated structures

Structures were analyzed by using the UCSF ChimeraX [42]. For amino acid structures, the “distance” (for bond lengths) and “angle” (for bond angles) commands were used with atom specification tailored to each amino acid type. Experimentally determined amino acid backbone bond lengths (N-Cα, 1.458Å; Cα-C, 1.525Å; C-O 1.231Å) [38], backbone bond angles (N-Cα-C, 111.0°; Cα-C-O, 120.1°) [39], and sidechain bond lengths and bond angles **(Table S4)** [40] served as references for evaluating predicted amino acid structures. Sidechain bond lengths and bond angles, obtained from a backbone-dependent rotamer library built into ChimeraX, represent those in which the backbone dihedral angles are ω=180° (typical of trans peptide bonds), φ=180°, and ψ=180° [40, 42]. For GPT-4, one iteration of cysteine lacked the backbone O atom and one iteration of methionine lacked the sidechain Cγ atom. Thus, these single iterations (n=1) from were excluded from analyses involving the missing atoms.

For α-helix structures, the matchmaker tool was used for alignment and RMSD determination across all atom pairs by using its default parameters (*i.e.*, best-aligning pair). An experimentally determined α-helical structure consisting of 10 consecutive alanine residues, detected within an engineered form of bacteriophage T4 lysozyme (PDB ID 1L64) [48], was used as the reference for evaluating the modeled α-helix structures. The AlphaFold2 α-helix structure was modeled using ColabFold [61] through ChimeraX run with default parameters (*i.e.*, minimization and without templates usage) using an elongated polyalanine sequence as input **(Table S5 and Fig. S5a)**. The ChimeraX and PyMOL (version 2.5.7) α-helix structures were modeled by using the build structure command and fab command, respectively, each by using a 10-residue alanine sequence as input and run with default α-helix parameters (*i.e.*, backbone dihedral angles set to φ=-57° and ψ=-47°) **(Fig. S5, b and c)**. All data were analyzed using GraphPad Prism 10.1.0 (GraphPad Software). Statistical details are reported in the figure legends and statistical measurements include mean, mean ± SD, and mean ± range.

Prompt-based interaction analysis with GPT-4

Structural analysis of binding interaction was performed by providing GPT-4 with an input PDB file and prompting as described (Table 1B) through the ChatGPT interface. The PDB file used as input was unmodified as obtained from the PDB entry for PDB ID: 7VH8 [49]. It should be noted that PDB ID: 7VH8 refers to nirmatrelvir as PF-07321332. ChimeraX was used to analyze amino acid residues detected by GPT-4 to interact with nirmatrelvir. The “contacts” and “distance” commands were used to verify interactions. For this exercise, GPT-4 was not limited to classic mode. Rather the “browser” and “analysis” features were enabled within the ChatGPT interface to enable file input, a feature available for GPT-4 but not GPT-3.5. Only the “analysis” feature was engaged for the responses generated by GPT-4 **(Table S3)**.

Acknowledgments

This work was supported by core services from the Cancer Center Support Grant of the Rutgers Cancer Institute of New Jersey (P30CA072720), by the National Institutes of Health (R01CA226537 to R.P. and W.A.), and by the Levy-Longenbaugh Donor-Advised Fund (to R.P. and W.A.). RCSB Protein Data Bank is jointly funded by the National Science Foundation (DBI-1832184, PI: S.K. Burley), the US Department of Energy (DE-SC0019749, PI: S.K. Burley), and the National Cancer Institute, the National Institute of Allergy and Infectious Diseases, and the

National Institute of General Medical Sciences of the National Institutes of Health (R01GM133198, PI: S.K. Burley). Molecular graphics and analyses performed with UCSF ChimeraX, developed by the Resource for Biocomputing, Visualization, and Informatics at the University of California, San Francisco, with support from National Institutes of Health R01-GM129325 and the Office of Cyber Infrastructure and Computational Biology, National Institute of Allergy and Infectious Diseases.

Author Contributions

A.M.I. conceptualization; A.M.I., C.M., and S.K.B. methodology; A.M.I. and C.M. investigation; A.M.I. data curation; A.M.I., C.M., S.K.B., M.B.M., R.P., and W.A. formal analysis; A.M.I. and C.M. writing - original draft; S.K.B., M.B.M., R.P., and W.A. writing - review & editing; S.K.B., R.P., and W.A. funding acquisition; R.P. and W.A. overall project supervision.

Conflict of Interest

A.M.I. is a co-founder and partner of North Horizon. R.P. and W.A. are founders and equity shareholders of PhageNova Bio. R.P. is Chief Scientific Officer and a paid consultant of PhageNova Bio. R.P. and W.A. are founders and equity shareholders of MBrace Therapeutics; R.P. and W.A. serve as paid consultants for MBrace Therapeutics. R.P. and W.A. have Sponsored Research Agreements (SRAs) in place with PhageNova Bio and with MBrace Therapeutics. These arrangements are managed in accordance with the established institutional conflict-of-interest policies of Rutgers, The State University of New Jersey. This study falls outside of the scope of these SRAs.

References

1. Jordan, M.I. and T.M. Mitchell, *Machine learning: Trends, perspectives, and prospects.* Science, 2015. **349**(6245): p. 255-60.
2. Wang, H., et al., *Scientific discovery in the age of artificial intelligence.* Nature, 2023. **620**(7972): p. 47-60.
3. Jumper, J., et al., *Highly accurate protein structure prediction with AlphaFold.* Nature, 2021. **596**(7873): p. 583-589.
4. Baek, M., et al., *Accurate prediction of protein structures and interactions using a three-track neural network.* Science, 2021. **373**(6557): p. 871-876.
5. Kryshtafovych, A., et al., *Critical assessment of methods of protein structure prediction (CASP)-Round XIV.* Proteins, 2021. **89**(12): p. 1607-1617.
6. Burley, S.K. and H.M. Berman, *Open-access data: A cornerstone for artificial intelligence approaches to protein structure prediction.* Structure, 2021. **29**(6): p. 515-520.
7. Elnaggar, A., et al., *ProtTrans: Toward Understanding the Language of Life Through Self-Supervised Learning.* IEEE Trans Pattern Anal Mach Intell, 2022. **44**(10): p. 7112-7127.
8. Chowdhury, R., et al., *Single-sequence protein structure prediction using a language model and deep learning.* Nat Biotechnol, 2022. **40**(11): p. 1617-1623.
9. Lin, Z., et al., *Evolutionary-scale prediction of atomic-level protein structure with a language model.* Science, 2023. **379**(6637): p. 1123-1130.
10. Bordin, N., et al., *AlphaFold2 reveals commonalities and novelties in protein structure space for 21 model organisms.* Commun Biol, 2023. **6**(1): p. 160.
11. Bordin, N., et al., *Novel machine learning approaches revolutionize protein knowledge.* Trends Biochem Sci, 2023. **48**(4): p. 345-359.
12. Mosalaganti, S., et al., *AI-based structure prediction empowers integrative structural analysis of human nuclear pores.* Science, 2022. **376**(6598): p. eabm9506.
13. Fontana, P., et al., *Structure of cytoplasmic ring of nuclear pore complex by integrative cryo-EM and AlphaFold.* Science, 2022. **376**(6598): p. eabm9326.
14. Read, R.J., et al., *AlphaFold and the future of structural biology.* Acta Crystallogr F Struct Biol Commun, 2023. **79**(Pt 7): p. 166-168.
15. Edich, M., et al., *The impact of AlphaFold2 on experimental structure solution.* Faraday Discuss, 2022. **240**(0): p. 184-195.
16. Varadi, M., et al., *AlphaFold Protein Structure Database in 2024: providing structure coverage for over 214 million protein sequences.* Nucleic Acids Res, 2023.
17. Varadi, M. and S. Velankar, *The impact of AlphaFold Protein Structure Database on the fields of life sciences.* Proteomics, 2023. **23**(17): p. e2200128.
18. Burley, S.K., W. Arap, and R. Pasqualini, *Predicting Proteome-Scale Protein Structure with Artificial Intelligence.* N Engl J Med, 2021. **385**(23): p. 2191-2194.
19. Brown, T.B., et al., *Language Models are Few-Shot Learners.* arXiv, 2020.
20. OpenAI, *GPT-4 Technical Report.* arXiv, 2023.
21. OpenAI. *Introducing ChatGPT*. 2022 [cited 2023 June 15, 2023]; Available from: https://openai.com/blog/chatgpt.
22. Boiko, D.A., et al., *Autonomous chemical research with large language models.* Nature, 2023. **624**(7992): p. 570-578.
23. Jablonka, K.M., et al., *Leveraging large language models for predictive chemistry.* Nature Machine Intelligence, 2024.
24. Gurwitz, D. and N. Shomron, *Artificial intelligence utility for drug development: ChatGPT and beyond.* Drug Dev Res, 2023.
25. Wang, R., H. Feng, and G.W. Wei, *ChatGPT in Drug Discovery: A Case Study on Anticocaine Addiction Drug Development with Chatbots.* J Chem Inf Model, 2023. **63**(22): p. 7189-7209.
26. Chakraborty, C., M. Bhattacharya, and S.S. Lee, *Artificial intelligence enabled ChatGPT and large language models in drug target discovery, drug discovery, and development.* Mol Ther Nucleic Acids, 2023. **33**: p. 866-868.
27. Lubiana, T., et al., *Ten quick tips for harnessing the power of ChatGPT in computational biology.* PLoS Comput Biol, 2023. **19**(8): p. e1011319.
28. Shue, E., et al., *Empowering beginners in bioinformatics with ChatGPT.* Quant Biol, 2023. **11**(2): p. 105-108.
29. Karkera, N., S. Acharya, and S.K. Palaniappan, *Leveraging pre-trained language models for mining microbiome-disease relationships.* BMC Bioinformatics, 2023. **24**(1): p. 290.
30. Toufiq, M., et al., *Harnessing large language models (LLMs) for candidate gene prioritization and selection.* J Transl Med, 2023. **21**(1): p. 728.

31. Xiao, Z., et al., *Generative Artificial Intelligence GPT-4 Accelerates Knowledge Mining and Machine Learning for Synthetic Biology.* ACS Synth Biol, 2023. **12**(10): p. 2973-2982.
32. Eriksen, A.V., S. Möller, and J. Ryg, *Use of GPT-4 to Diagnose Complex Clinical Cases.* NEJM AI, 2023. **1**(1).
33. Hirosawa, T., et al., *ChatGPT-Generated Differential Diagnosis Lists for Complex Case-Derived Clinical Vignettes: Diagnostic Accuracy Evaluation.* JMIR Med Inform, 2023. **11**: p. e48808.
34. van Dis, E.A.M., et al., *ChatGPT: five priorities for research.* Nature, 2023. **614**(7947): p. 224-226.
35. Graf, A. and R.E. Bernardi, *ChatGPT in Research: Balancing Ethics, Transparency and Advancement.* Neuroscience, 2023. **515**: p. 71-73.
36. Thapa, S. and S. Adhikari, *ChatGPT, Bard, and Large Language Models for Biomedical Research: Opportunities and Pitfalls.* Ann Biomed Eng, 2023. **51**(12): p. 2647-2651.
37. Ille, A.M. and M.B. Mathews, *AI interprets the Central Dogma and Genetic Code.* Trends Biochem Sci, 2023. **48**(12): p. 1014-1018.
38. Engh, R.A. and R. Huber, *Accurate bond and angle parameters for X-ray protein structure refinement.* Acta Crystallographica Section A, 1991. **47**(4): p. 392-400.
39. Berkholz, D.S., et al., *Conformation dependence of backbone geometry in proteins.* Structure, 2009. **17**(10): p. 1316-25.
40. Shapovalov, M.V. and R.L. Dunbrack, Jr., *A smoothed backbone-dependent rotamer library for proteins derived from adaptive kernel density estimates and regressions.* Structure, 2011. **19**(6): p. 844-58.
41. Fujii, N. and T. Saito, *Homochirality and life.* Chem Rec, 2004. **4**(5): p. 267-78.
42. Meng, E.C., et al., *UCSF ChimeraX: Tools for structure building and analysis.* Protein Sci, 2023. **32**(11): p. e4792.
43. Doig, A.J., et al., *Structure, stability and folding of the alpha-helix.* Biochem Soc Symp, 2001(68): p. 95-110.
44. Eisenberg, D., *The discovery of the alpha-helix and beta-sheet, the principal structural features of proteins.* Proc Natl Acad Sci U S A, 2003. **100**(20): p. 11207-10.
45. Pauling, L., R.B. Corey, and H.R. Branson, *The structure of proteins; two hydrogen-bonded helical configurations of the polypeptide chain.* Proc Natl Acad Sci U S A, 1951. **37**(4): p. 205-11.
46. Pace, C.N. and J.M. Scholtz, *A helix propensity scale based on experimental studies of peptides and proteins.* Biophys J, 1998. **75**(1): p. 422-7.
47. Wolfram, S., *ChatGPT Gets Its "Wolfram Superpowers"!* 2023.
48. Heinz, D.W., W.A. Baase, and B.W. Matthews, *Folding and function of a T4 lysozyme containing 10 consecutive alanines illustrate the redundancy of information in an amino acid sequence.* Proc Natl Acad Sci U S A, 1992. **89**(9): p. 3751-5.
49. Zhao, Y., et al., *Crystal structure of SARS-CoV-2 main protease in complex with protease inhibitor PF-07321332.* Protein Cell, 2022. **13**(9): p. 689-693.
50. Hammond, J., et al., *Oral Nirmatrelvir for High-Risk, Nonhospitalized Adults with Covid-19.* N Engl J Med, 2022. **386**(15): p. 1397-1408.
51. Chatterjee, S., et al., *Resistance to nirmatrelvir due to mutations in the Mpro in the subvariants of SARS-CoV-2 Omicron: Another concern?* Mol Ther Nucleic Acids, 2023. **32**: p. 263-266.
52. Hu, Y., et al., *Naturally Occurring Mutations of SARS-CoV-2 Main Protease Confer Drug Resistance to Nirmatrelvir.* ACS Cent Sci, 2023. **9**(8): p. 1658-1669.
53. Iketani, S., et al., *Multiple pathways for SARS-CoV-2 resistance to nirmatrelvir.* Nature, 2023. **613**(7944): p. 558-564.
54. Zhou, Y., et al., *Nirmatrelvir-resistant SARS-CoV-2 variants with high fitness in an infectious cell culture system.* Sci Adv, 2022. **8**(51): p. eadd7197.
55. Zuckerman, N.S., et al., *Nirmatrelvir resistance - de novo E166V/L50V mutations in an immunocompromised patient treated with prolonged nirmatrelvir/ritonavir monotherapy leading to clinical and virological treatment failure - a case report.* Clin Infect Dis, 2023.
56. Hirotsu, Y., et al., *Multidrug-resistant mutations to antiviral and antibody therapy in an immunocompromised patient infected with SARS-CoV-2.* Med, 2023. **4**(11): p. 813-824 e4.
57. Eisenstein, M., *A test of artificial intelligence.* Nature, 2023.
58. Biever, C., *ChatGPT broke the Turing test - the race is on for new ways to assess AI.* Nature, 2023. **619**(7971): p. 686-689.
59. OpenAI, *ChatGPT plugins*. 2023.
60. Wolfram, S., *ChatGPT Gets Its "Wolfram Superpowers"!* 2023.
61. Mirdita, M., et al., *ColabFold: making protein folding accessible to all.* Nat Methods, 2022. **19**(6): p. 679-682.