Human Cytokine and Coronavirus Nucleocapsid Protein Interactivity Using Large-Scale Virtual Screens

Phillip J. Tomezsko[1,*], Colby T. Ford[2,3,4,*], Avery E. Meyer[1], Adam M. Michaleas[1], and Rafael Jaimes III[1,✉]

[1]MIT Lincoln Laboratory, Lexington, MA, USA
[2]Tuple LLC, Charlotte, NC, USA
[3]University of North Carolina at Charlotte, Department of Bioinformatics and Genomics, Charlotte, NC, USA
[4]University of North Carolina at Charlotte, Center for Computational Intelligence to Predict Health and Environmental Risks (CIPHER), Charlotte, NC, USA

In the battle against the ever-changing SARS-CoV-2 landscape, understanding the interactions between viral proteins and the human immune system is paramount as it helps to explain potential factors contributing to diverse immunological responses in infected individuals. In this study, we employed state-of-the-art molecular docking tools to conduct large-scale virtual screens, predicting the binding affinities between 64 human cytokines against 17 coronavirus nucleocapsid proteins. Our comprehensive *in silico* analyses reveal specific changes in cytokine-nucleocapsid protein interactions, shedding light on potential modulators of the host immune response during infection. These findings offer valuable insights into the molecular mechanisms underlying viral pathogenesis and may guide the future development of targeted interventions. This manuscript serves as insight into the comparison of deep learning based AlphaFold2-Multimer and the semi-physicochemical based HADDOCK for protein-protein docking. We show the two methods are complementary in their predictive capabilities. We also introduce a novel algorithm for rapidly assessing the binding interface of protein-protein docks using graph edit distance: graph-based residue assessment function (G-RAF). The high-performance computational framework presented here will not only aid in accelerating the discovery of effective interventions against emerging viral threats, but extend to other applications of high throughput protein-protein screens.

SARS-CoV-2 | coronavirus | cytokines | nucleocapsid | AlphaFold2 | HADDOCK | protein modeling | immunology

Correspondence: ***rafael.jaimes@ll.mit.edu***

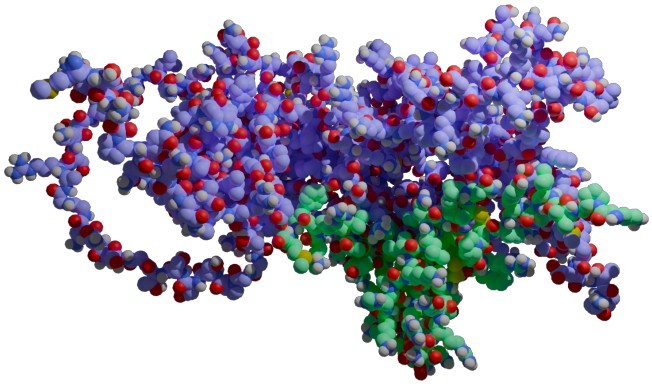

Fig. 1. Human cytokine CXCL12β (in green) docked against the SARS-CoV-2 WA-1 nucleocapsid protein (in purple). Rendered with MolecularNodes (1).

Introduction

SARS-CoV-2 evolution over the course of the pandemic has led to sustained continued waves of The variants of concern in have shown a high degree of mutation relative to the prevailing strain at the time of their emergence. Most research has focused on the impact of mutations on the spike proteins' ability to enter cells and evade antibodies, whether the antibodies are therapeutics or induced from vaccination or prior infection (2, 3). The main reason for the focus on spike is driven by the fact that the most selective pressure on the virus is applied to the spike and the mutations have consequences for the therapeutics that are available. Outside of spike, the nucleocapsid (N) gene of SARS-CoV-2 has had lineage defining mutations in each of the variants of concern (4). A robust platform for studying the effects of nucleocapsid mutations on viral RNA delivery and expression has identified critical mutations in the Alpha and Delta variants that may contribute to higher transmissibility(5).

Another critical function of SARS-CoV-2 N is to bind and sequester a subset of 11 cytokines in order to disrupt immune signaling (6). The set of cytokines that SARS-CoV-2 binds is shared with another distantly related human beta-coronavirus, HCoV-OC43, though HCoV-OC43 also binds to an additional 6 cytokines with high affinity (6, 7). Binding to at least one cytokine, CXCL12β, was determined experimentally for each human betacoronavirus. Binding of CXCL12β was shown to inhibit chemotaxis of macrophages in a transwell culture study. Other viruses also employ extracellular proteins that bind to specific subsets of cytokines with high affinity, most notably the SECRET-domain containing proteins of the poxviruses (8, 9). The presence or absence of certain SECRET-domain containing proteins can have important consequences on the pathogenicity of the poxvirus.

An *in silico* docking screen of the panel of 11 cytokines tested experimentally against SARS-CoV-2 original reference strain and key variants of concern was developed in the current study. The *in silico* screen was set up on a high-performance computing system. AlphaFold2-Multimer (10) and HADDOCK (11) were utilized in parallel to predict cytokine binding sites on N, as GDockScore (12) and PRODIGY (13) were used to further assess the properties of the predicted cytokine binding. The goals were to identify the cytokine binding sites

of the experimental cytokine hits and determine if their binding has been impacted by continuing evolution in the human host over the course of the pandemic. N proteins from human betacoronaviruses and close relatives of SARS-CoV-2 were included in the *in silico* screen as well to test the conservation of cytokine binding in the broader betacoronavirus family. Finally, we wanted to test the ability of the *in silico* tools to identify interactions with cytokines using the full 64 cytokine panel utilized experimentally. The full cytokine panel docking screen could be utilized in order to track how evolution of N impacts this function continuing forward in the pandemic as well test other viral proteins for the ability to sequester cytokines.

Results

The AlphaFold2 structural predictions matched the expected topology of the SARS-CoV-2 N protein (14). The predicted structure of the SARS-CoV-2 N WA1 strain had two structured domains that recapitulate two crystallized domains (PDB 6M3M for the N-terminal RNA binding domain, total RMSD 0.851Å and 6WZO for the dimerization domain, total RMSD 1.178Å) as well as three flexible domains with low AlphaFold2 confidence (pTM) that correspond with the three intrinsically disordered domains.

Computing Benchmarks. Running on the MIT SuperCloud high-performance computing (HPC) environment (15), significant throughput was achieved for the numerous computational docking tasks.

AlphaFold2 Performance. For AlphaFold2, performance benchmarks were collected on both standalone and high performance computing platforms.

For the standalone system, we used an NVIDIA DGX A100 80GB server. The system is equipped with dual AMD Rome 7742 CPUs, 2TB RAM, and 8 NVIDIA A100-SXM 80GB GPUs.

The MIT Lincoln Laboratory Tx-Green and MIT SuperCloud (15) systems were the environments used for the HPC-enabled AlphaFold pipeline prototyping and performance benchmarking. The compute nodes each consisted of a Intel Xeon Gold 6248 2.5 GHz CPU with 40 cores, 377GB RAM, and Intel Omni-Path with 2 NVIDIA Tesla V100 GPUs.

AlphaFold2-Multimer jobs took 493 $\pm$ 21 minutes of walltime per multimer run on nodes containing an Nvidia Volta V100 GPU, which contain 5,120 CUDA cores and 640 tensor cores. To complete the 1,088 AlphaFold2-Multimer experiments, 36 nodes were used, completing the entire set of experiments in approximately 12 days. A later subset was run on Ampere A100 GPUs to investigate if the newer generation GPU would increase speedups; the mean runtime decreased to 378 minutes $\pm$ 80 minutes (noting that the variability of duration increased). The A100 GPUs have more CUDA cores but fewer tensor cores, 6,912 and 432 respectively.

HADDOCK Performance. HADDOCK jobs took 707 $\pm$ 40 minutes of walltime on Intel Xeon Phi 7210 cluster nodes, which contain 64 physical cores. On Intel Xeon Platinum 8260 nodes, which contain 48 physical cores, HADDOCK took 133 $\pm$ 92 minutes of walltime. The speedup on the Xeon Platinum nodes was likely due to the much newer CPU architecture and higher maximum clockspeed of the Xeon Platinum (3.9 GHz) compared to the Xeon Phi (1.5 GHz). While HADDOCK performed as quickly as 74 min on the Xeon Platinum nodes, it had a wide range of runtimes, depending on the cytokine and number of surface residues selected against which to dock. Cytokines with multiple chains (IL-12p70, IL-23, IL-27, and IL-35) took approximately three times as long. The 1,088 HADDOCK docking experiments were run on 64 of the Xeon Platinum 8260 nodes, completing the entire set of experiments in approximately 36 hours.

Nucleocapsid Cytokine Binding Sites. HADDOCK and AlphaFold2-Multimer were used to identify potential binding sites of the cytokines with the nucleocapsid proteins. CXCL12β was chosen for initial representative structure modeling because it has been experimentally shown to bind to HCoV-OC43 and SARS-CoV-2 WA1 nucleocapsid in biolayer interferometry assays (6, 7). SARS-CoV, MERS-CoV, HCoV-OC43 and SARS-CoV-2 nucleocapsid were additionally shown to inhibit CXCL12β dependent migration of macrophages in a transwell assay (6, 7). For the six betacoronaviruses (using SARS-CoV-2 WA1 as the representative strain of SARS-CoV-2), HADDOCK docked CXCL12β to locations around the dimerization interface, with contributing contacts from other domains of the protein. AlphaFold2-Multimer complexed the cytokine consistently at the dimerization interface, including interacting residues in the beta sheets of both the nucleocapsid dimerization interface and CXCL12β. See Figure 2.

A closer look at the HADDOCK and AlphaFold2-Multimer docking of CXCL12β shows that HADDOCK did not include any beta sheet residues from the dimerization interface, but did include polar contacts with residues on either side of the interface loop (S310, G335 and A336). Other interacting residues were located in the structured N-terminal RNA binding domain and the flexible C-terminal region. The AlphaFold2-Multimer complex included two polar contacts of residues directly involved in dimerization (R319 and I320), as well as resides in the flexible C-terminal region. Two residues on CXCL12β involved in glycosaminoglycan (GAG) binding (H25 and R47) were involved in the HADDOCK predicted polar contacts, whereas the GAG-binding residues showed hydrogen bonding but no polar contacts in the AlphaFold2-Multimer predicted complex (16, 17). The SARS-CoV-2 WA1 nucleocapsid CXCL12β interaction was previously shown to be competitive with heparin sulfate and chondroitin sulfate (6). See Figure 3.

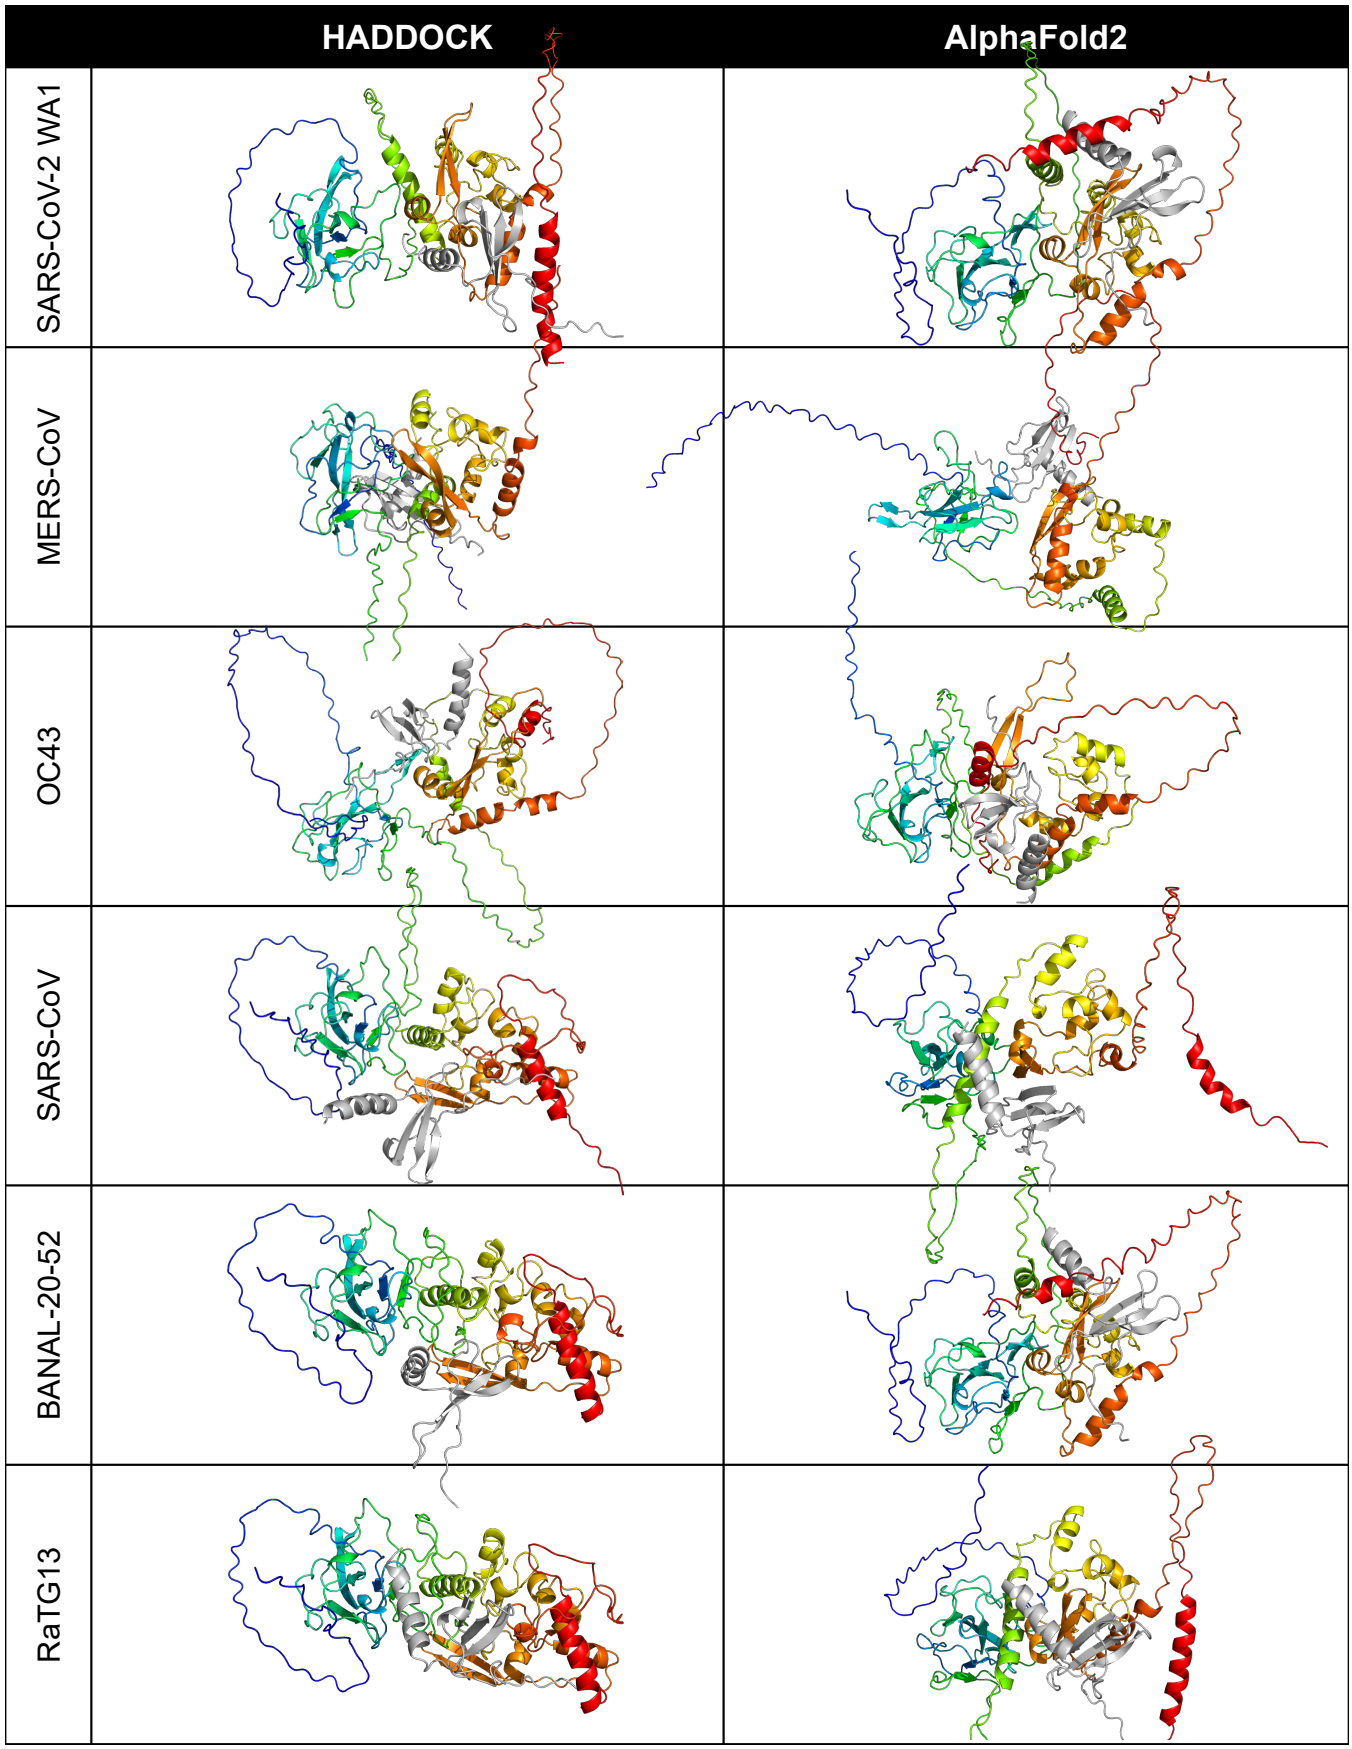

Fig. 2. Predicted binding site of CXCL12β across various coronaviruses. The proteins are shown in cartoon style with the N protein colored as rainbow (blue is the amino/N-terminus, red is the carboxyl/C-terminus) and the CXCL12β cytokine is colored in gray.

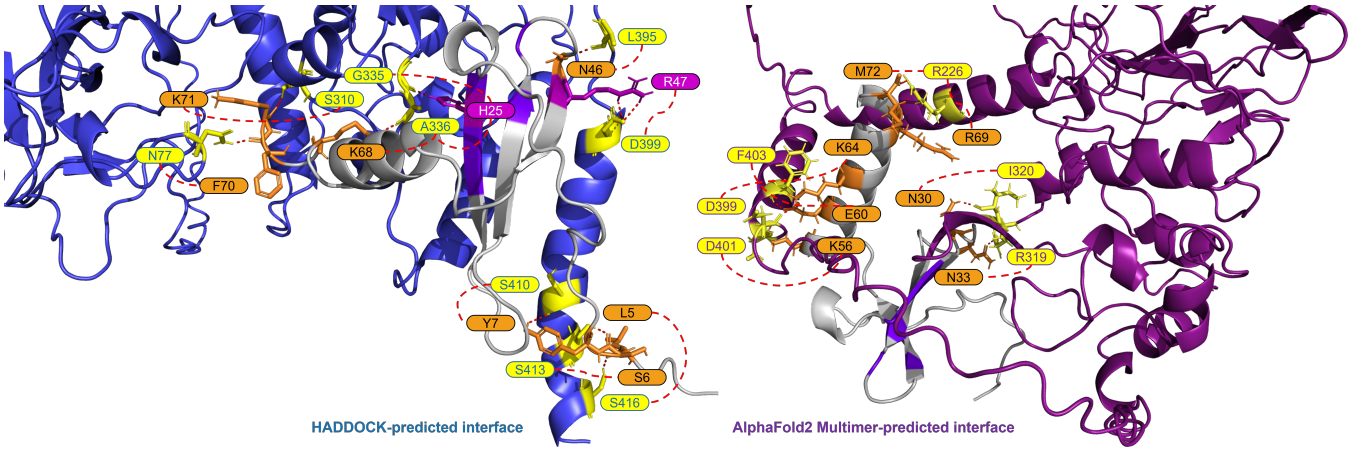

Fig. 3. Predicted interface between the SARS-CoV-2 WA-1 N protein and CXCL12β. The proteins are shown in cartoon style with the N protein colored as periwinkle in the HADDOCK-predicted complex and dark magenta in the AlphaFold2 Multimer-predicted complex (with polar interacting residues shown as yellow sticks). The CXCL12β protein is shown in cartoon style in gray (with polar interacting residues shown as orange sticks). Residues in the GAG binding site on CXCL12β are shown in purple (with polar interacting residues shown as magenta sticks).

Cytokine Strength. Across the 64 cytokines that were tested, binding affinities were highly correlated between the AlphaFold2-Multimer predictions and HADDOCK. However, the ranges of the scores may vary significantly. For example, the range of the HADDOCK van der Waals energies is [-109.82, -38.45] compared to [-56.76, -10.46] for the AlphaFold2-Multimer predictions (measured by FoldX). See Table 1.

p-values / R^2		HADDOCK			AlphaFold2 Multimer		
		van der Waals Energy	PRODIGY ΔG	FoldX ΔG	van der Waals Energy	PRODIGY ΔG	FoldX ΔG
HADDOCK	van der Waals Energy		**0.9867**	0.0830	0.0037	0.0122	0.0384
	PRODIGY ΔG	**<0.0001**		**0.6095**	**0.8560**	**0.8935**	0.0366
	FoldX ΔG	**0.0210**	**<0.0001**		0.0515	0.0419	0.0582
AlphaFold2 Multimer	van der Waals Energy	0.6349	**<0.0001**	0.0713		**0.5741**	0.1605
	PRODIGY ΔG	0.3852	**<0.0001**	0.1045	**<0.0001**		0.0161
	FoldX ΔG	0.1206	0.1267	0.0548	**0.0010**	0.3178	

Table 1. Correlations between AlphaFold2-Multimer and HADDOCK docking predictions across various binding metrics. Significant correlations are shown in **bold**.

Genetic Distance. Correlating the cytokine binding affinity metrics with Fitch distance of each SARS-CoV-2 variants' N protein from the SARS-CoV-2 WA1 N protein, there are significant positive correlations for some cytokines. Namely, CCL28, CXCL10, CXCL11, and CXCL14 from the HADDOCK predictions show that the binding affinity worsens as the genetic distance of the N protein increases from SARS-CoV-2 WA1. For some cytokines, while not statistically significant, there are some negative trends that would imply an improvement in N-cytokine binding due to mutations in newer SARS-CoV-2 variants over time. See Figures 4 and 5.

Breadth of CXCL12β Binding. Graph Edit Distance (GED) provides a view into how the N-protein receptor (i.e., the binding pocket) to CXCL12β changed between different N-proteins. Both AlphaFold2 and HADDOCK structures concurred with relatively low GED of SARS-CoV-2-B.1.1 and SARS-CoV-2-B.1.617.2-DeltaA (Figures 7A, 7B, respectively). AlphaFold2 had high GED for MERS-CoV and SARS-CoV-2-BQ.1, suggesting far evolutionary distance in the binding pocket compared to the SARS-COV2-WA1 baseline (Figure 7A).

There was marked disagreement between AlphaFold2 and HADDOCK for SARS-CoV-2-BQ.1 (Figure 7C). This was partly due to the discrepancy in the number of interfacing residues on the AlphaFold2 structure compared to the HADDOCK structure (Figure 7D). AlphaFold2 and HADDOCK reported a similar number of interfacing residues on other N proteins; both agreed the number of interfacing residues on SARS-CoV-2-XBB and SARS-CoV-2-P.1 as relatively low. AlphaFold2 reported 6 residues on SARS-CoV-2-WA1 compared to HADDOCK's 8 residues.

Discussion

In silico structural prediction and molecular docking tools were utilized in order to interrogate the potential evolutionary change of the interaction between SARS-CoV-2 nucleocapsid proteins and 64 human cytokines. Other betacoronaviruses were also included in the analysis, as all human betacoronaviruses had been shown to have a similar cytokine inhibitory phenotype. Additionally, HCoV-OC43 was shown to bind 16 human cytokines, including all 11 cytokines that SARS-CoV-2 bound despite being distantly related (6, 7). We identified the nucleocapsid dimerization domain as an important site of multiple cytokine interactions, with AlphaFold2 consistently identifying the dimerization loop specifically as the interaction site. We also identified four cytokines (CCL28, CXCL10, CXCL11, CXCL14) that had significantly reduced binding to N as SARS-CoV-2 evolved by at least one *in silico* metric. Finally, we identified 11 cytokines that did not bind SARS-CoV-2 WA1 N but had increased interactions with N variants. These cytokines could be potential new targets of inhibition by N.

We expected that interactions with N and cytokines would generally get stronger as SARS-CoV-2 co-evolved with humans. Contrary to our hypothesis, we found that the many

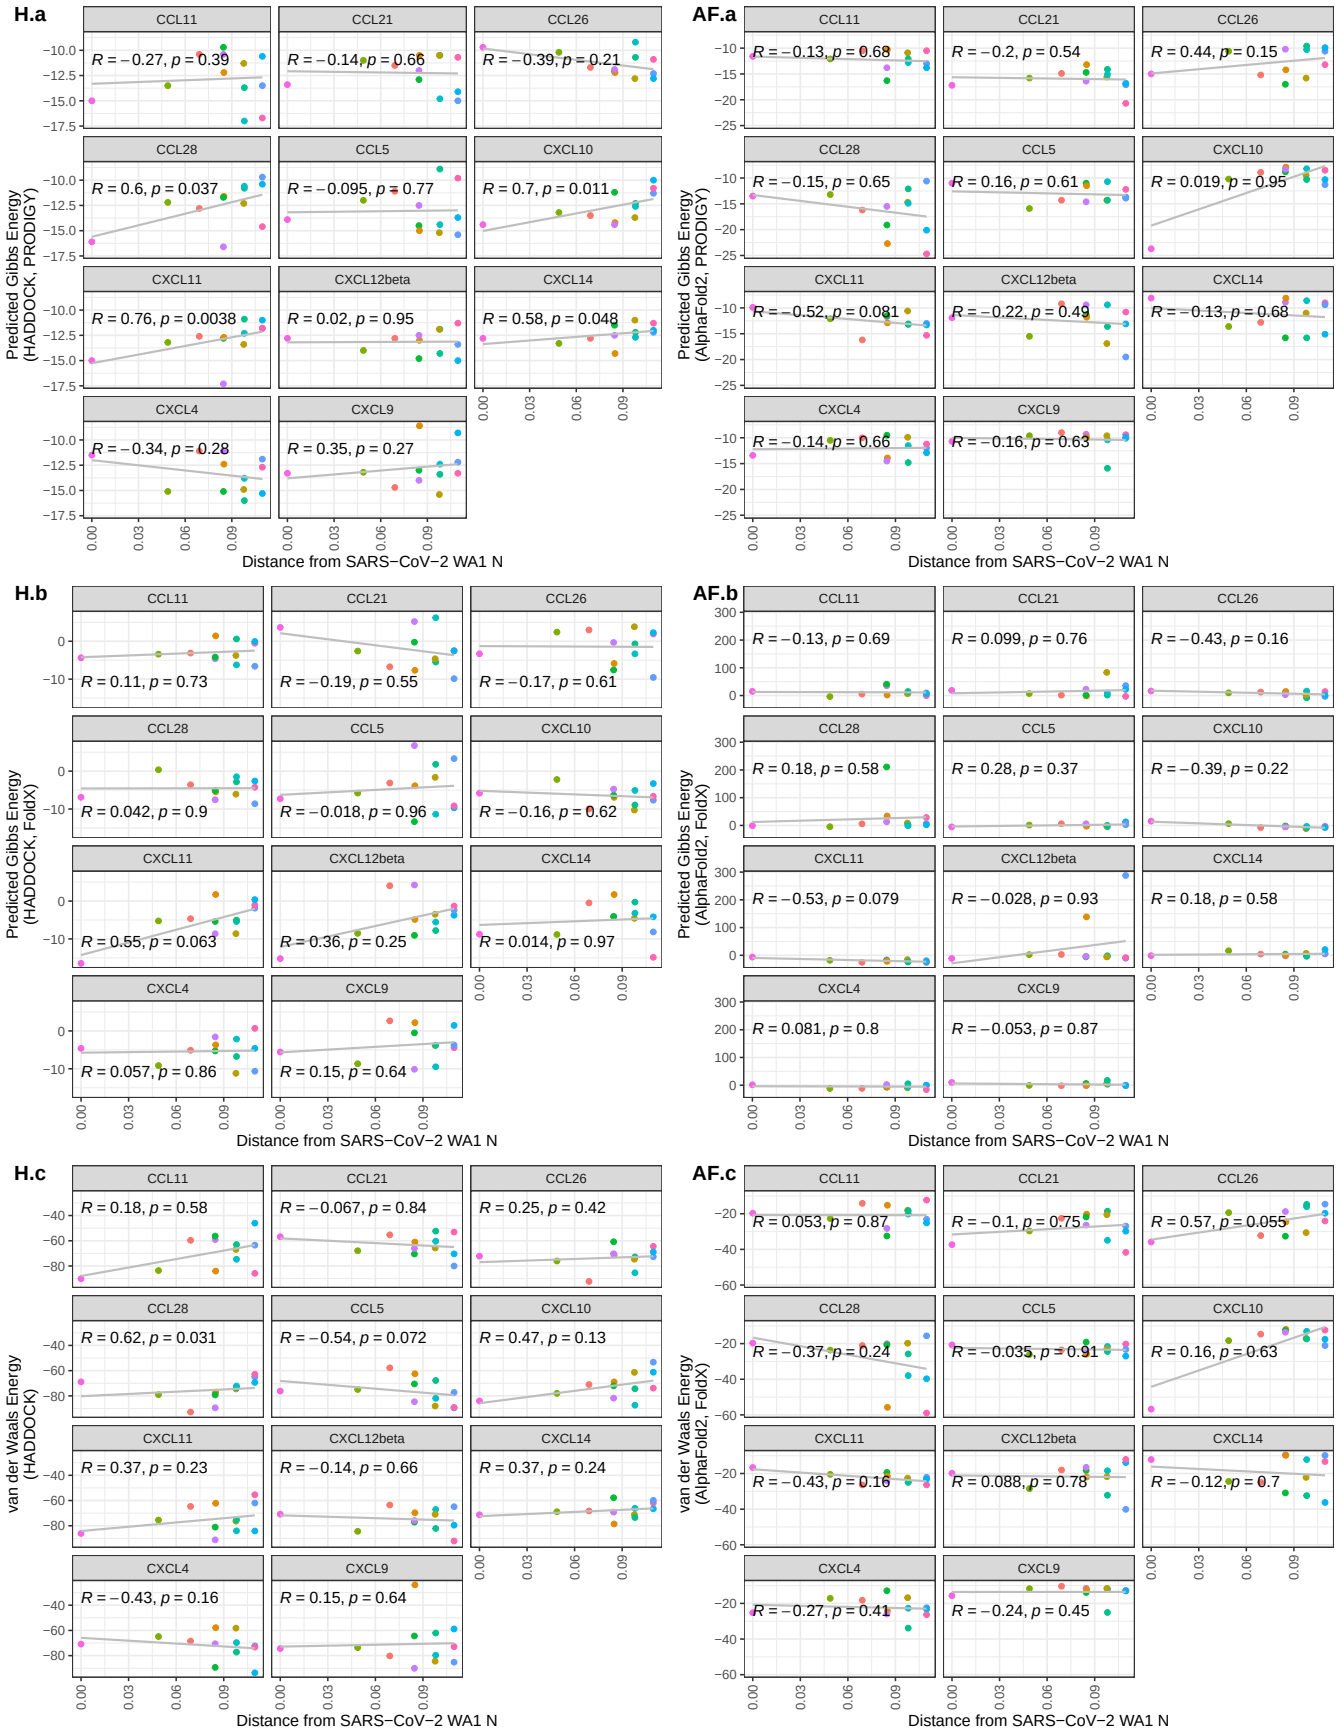

Fig. 4. Scatterplots of cytokine binding affinity for cytokines referenced in López-Muñoz et al. (2022). Shown is the relationship between the Fitch distance from the SARS-CoV-2 WA1 N variant vs. HADDOCK and AlphaFold2 predicted binding affinity metrics. Each point represents a SARS-CoV-2 variant.

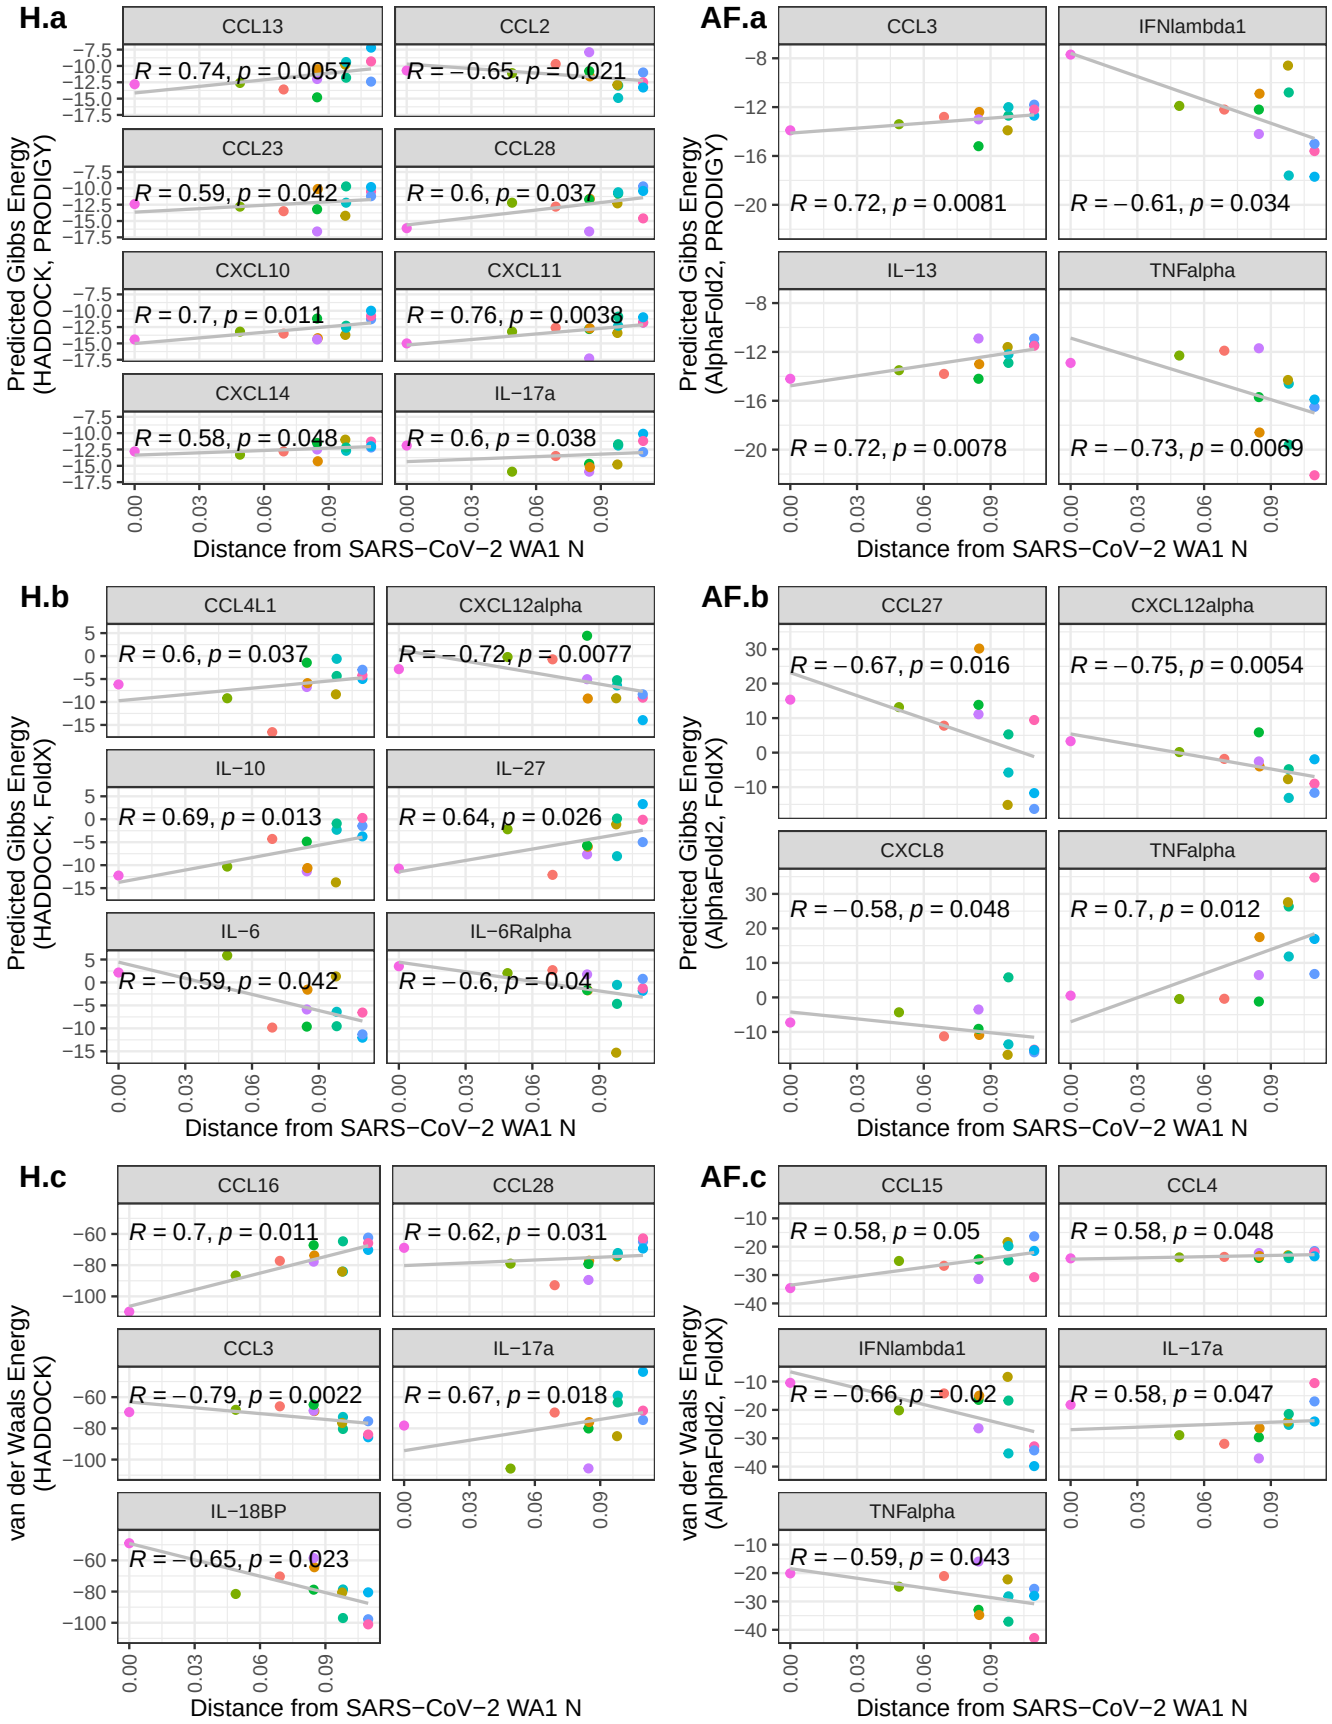

Fig. 5. Scatterplots of cytokine binding affinity for all cytokines with a statistically-significant correlation to genetic distance. Shown is the relationship between the Fitch distance from the SARS-CoV-2 WA1 N variant vs. HADDOCK and AlphaFold2 predicted binding affinity metrics. Each point represents a SARS-CoV-2 variant, with the most distant points representing Omicron subvariants.

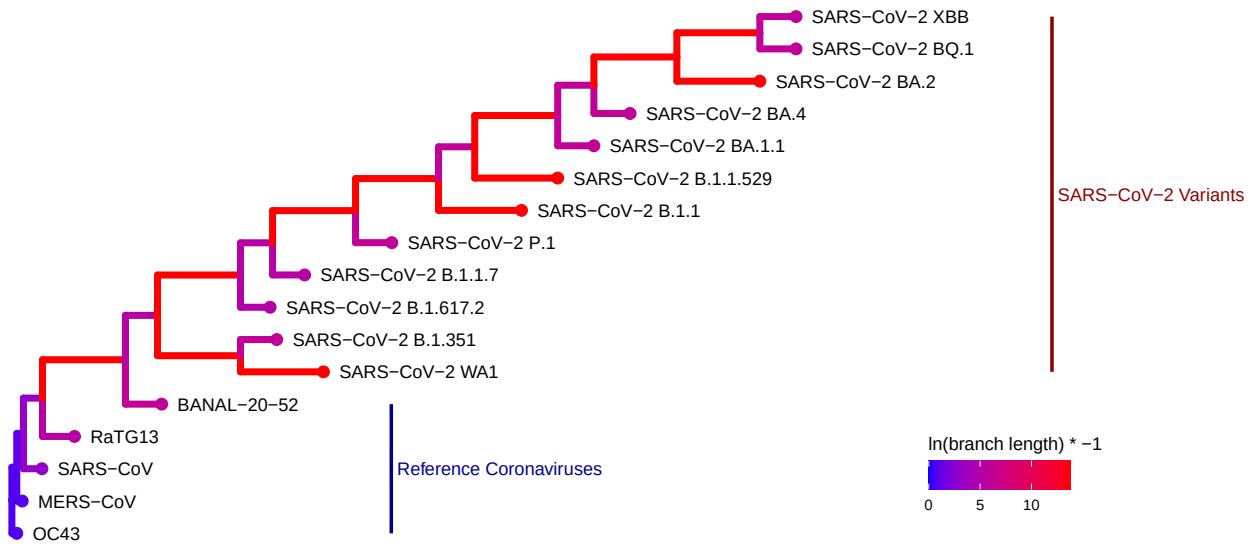

Fig. 6. Phylogenetic tree generated with RAxML depicting the genetic distance of the 17 coronavirus N proteins used in this study. (Colors represent the log-transformed branch length: $ln(\text{branch length}) \times -1$.)

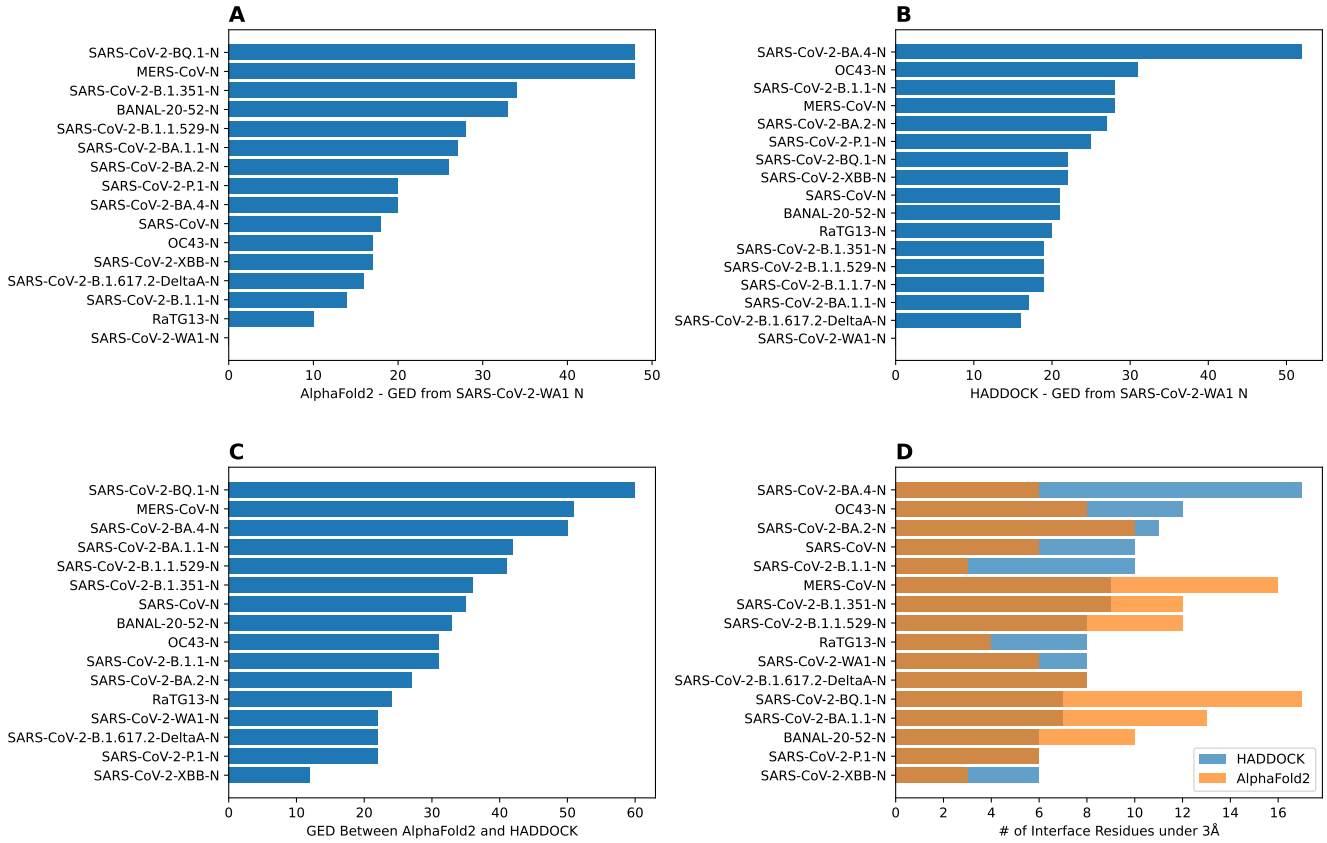

Fig. 7. Bar charts depicting the graph edit distance for CXCL12β. A) AlphaFold2 GED from SARS-CoV-2 WA1 N. B) HADDOCK GED from SARS-CoV-2 WA1 N. C) GED between AlphaFold2 and HADDOCK SARS-CoV-2 WA1 N. D) Number of interface residues under 3Å.

statistically significant changes were weakening of the cytokine interaction as mutations accumulated. Multiple variants of the Omicron group were included in the analysis and were the most distant from the reference WA1 strain. The decreasing capacity for inhibition of cytokine signaling could be at least partially responsible for the decreased severity of disease observed with Omicron. Of the four cytokine that HADDOCK predicted would bind with decreased affinity, CXCL10 stands out as it is a strong predictor of disease severity (18–20). Decreased disregulation of CXCL10 by Omicron N could be driving any related phenotypic change towards lower severity.

In addition to interrogating how the binding might change over the course of the pandemic for the cytokines that were experimentally shown to bind to SARS-CoV-2 WA1 N, we also identified a series of cytokines that may be inhibited by Omicron that were not inhibited by the original SARS-CoV-2. These include predicted increased binding of CXCL8, CXCL12α, CCL27, IFNλ1, and TNFα by AlphaFold2 and CCL2, CCL3, CXCL12α, IL-6, IL6Rα and IL-18BP by HADDOCK. Though the predicted affinity and/or van der Waals associated is lower over the course of viral evolution, it is not certain that these cytokines would be inhibited because it is unclear if the lower affinity is enough to translate to physiological binding and the predicted binding site may not be inhibitory.

The predicted SARS-CoV-2 N binding to cytokine shares several features with the better characterized SECRET-cytokine binding. The binding to N typically involves a beta sheet to beta sheet interaction. One of the two main beta sheets involved in cytokine binding is the dimerization interface of SARS-CoV-2 N. The beta sheet interaction frequently involves the oligomerization interface of the cytokine. The binding both the GAG-binding domain oligomerization interface stabilized by a flexible arm describe the interaction motif of the SECRET proteins, at least one interaction of which has been crystallized (PDB 3ONA) (9).

Computational approaches to modeling binding are an attractive solution for variant tracking and biosurveillance in order to deal with the huge influx of sequences. Utilization of structural and docking predictions could be used to test the functional interactions as new variant sequences are identified. Previous computational studies have provided accurate early insights into viral mutations and their affects on human health (21–25).

In this work, we have not only showcased the utility of *in silico*-based protein modeling, but also the importance of scalability through high-performance computing. HPC-enabled frameworks, as used here, allow for tremendous throughput improvement, reducing costly and slow lab-based workloads.

This approach will be important for proactively determining when a pathogen has acquired new phenotypes, such as increased transmissibility, pathogenicity, or zoonotic potential. In the long term, robust computational assays for several functions would need to be developed in order to track those higher-order viral features such as transmissibility.

Methods

Overall Approach. The docking experiment set in this study contains 64 human cytokines and the nucleocapsid (N) protein from 17 different coronaviruses, 12 of which are from SARS-CoV-2 variants. The experiment generation consists of all possible combinations of the cytokines and N proteins for a total of 1,088 cytokine-N protein complexes.
Then, each experiment was submitted through both AlphaFold2-Multimer and HADDOCK docking systems. These tools generate numerous PDB files of predicted protein complex conformations. From these outputs, the best representative complex PDB structure was selected through various filtering techniques (described below) and then compared across the full experiment set. See Figure 8.

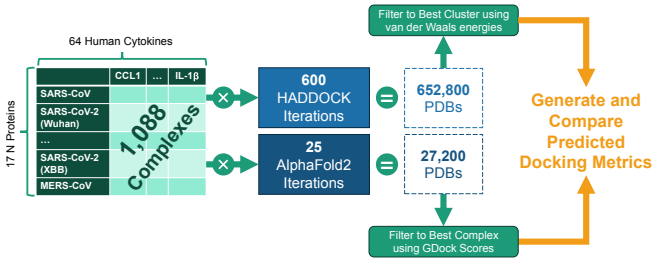

Fig. 8. Experiment generation workflow for the 1,088 cytokine-N complexes.

AlphaFold2 Protein Complex Prediction. AlphaFold2 (version 2.3.2) was compiled into a Singularity container and queried against a reference database constructed on 2023-04-12 according to the author's instructions (26). The database was compressed into a SquashFS file and bound directly to the Singularity container. The specification for the AlphaFold2 Singularity container can be found at: https://github.com/mit-ll/AlphaFold.
AlphaFold2 was run in multimer mode and used default settings to construct 5 randomly seeded models with 5 predictions each, for a total of 25 models (PDBs) per complex. Amber relaxation was used on all resulting 25 models. The inference was performed on an Nvidia Volta V100 GPU.
Signal peptides were identified and removed from cytokine sequences using SignalP6.0 web portal[1] before structure prediction with AlphaFold2.

AlphaFold2 Confidence. AlphaFold2 confidence was given as the interface predicted template model score plus predicted template model score (ipTM+pTM). The ipTM was weighted by 80% and the pTM was weighted by 20%, as described by the authors (10). The total AlphaFold2 confidence (ipTM+pTM) ranged from [0, 1], where 1 was the highest confidence.

GDock Score. *GDockScore* is a graph-based deep learning model to assess the docking of two proteins (12). The model was pre-trained by the original authors on docking outputs generated from Protein Data Bank, RosettaDock, HADDOCK decoys, and the ZDOCK Protein Docking

[1] SingalIP6.0 web portal: https://services.healthtech.dtu.dk/services/SignalP-6.0/

Benchmark- to include a wide variety of protein complexes and ensure generalization. *GDockScore* achieves state-of-the-art on the CAPRI Score Set, a challenging dataset for developing docking scoring functions (27). *GDockScore* ranges from 0 to 1, where a score of 0 coincides to an unfeasible complex and a score of 1 coincides with high probability that the protein complex is comparable to a high quality CAPRI complex. The authors show a *GDockScore* of under 0.2 approximates to a complex that is unlikely to exist (i.e., bad protein-protein dock) and a *GDockScore* range from 0.2 to 1 coincides with an increasing predicted quality.

Best AlphaFold2 Structure Selection. AlphaFold2 ranks the 25 structures by confidence (AF2 conf.). We added the raw AF2 conf. to the GDock score (i.e., total score) for all 25 structures and selected the highest score for each multimer experiment. The highest total score was not necessarily the best ranked from AF2 conf. alone. The best 1,088 structures were used for all downstream analysis. See Figure 9

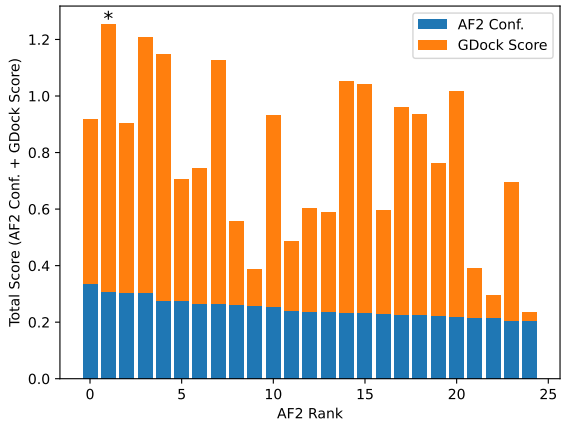

Fig. 9. Example total score calculation for SARS-CoV-2-WA1 bound to CXCL12β. AlphaFold2 confidence was calculated by ipTM+pTM (see text for details). The AlphaFold2 rank 0 had a slightly higher confidence than the rank 1, however the rank 1 structure had a significantly higher GDock score which led to a higher total score. In this case, the rank 1 structure was selected as the best structure.

HADDOCK Docking. For each of the 1,088 experiments generated in this study, HADDOCK v2.4 (11) was used to dock the N protein and cytokine in each experiment.

HADDOCK, High Ambiguity Driven protein-protein DOCKing, is a biomolecular modeling software that provides docking predictions for provided structures using an information-driven flexible docking approach.

For this study, we utilized a Docker containerized version of HADDOCK, which contains all of the software dependencies to allow HADDOCK to run more readily in an HPC environment. HADDOCK was run on both 64 physical cores (Intel Xeon Phi 7210) and 48 physical cores (Intel Xeon Platinum 8260). HADDOCK needs to be compiled against the number of physical cores; more information can be found at: https://github.com/colbyford/HADDOCKer, or on DockerHub at: https://hub.docker.com/r/cford38/haddock.

HADDOCK Experiment Setup. Run parameters for each experiment were generated programmatically, defining the N protein as the static object around which the cytokine protein is positioned and measured. Other experiment files were also copied or created programmatically that include the scripts to run the docking process, define restraints, and specify input PDB files.

HADDOCK requires the definition of active/inactive residue restraints (AIRs) to help guide the protein docking process. To avoid biasing the docking placement of the cytokine, all surface residues were selected as "active" and were then included in the AIR file on which to dock.

The logic for this programmatic generation of HADDOCK experiment files is available in the supplementary GitHub repository.

HADDOCK Outputs. The HADDOCK docking process consists of three steps:

it0: Randomization of orientations and rigid-body minimization

it1: Semi-flexible simulated annealing through molecular dynamics simulations in torsion angle space

itw: Refinement by energy minimization in Cartesian space with explicit solvent (i.e., in water)

Each of the above steps generates 200 PDB files of the proteins in a complex for a total of 600 PDB files. From the final, water-refined (*itw*) set of outputs, the "best" PDB file was retrieved from the cluster of predictions with the lowest van der Waals energy. This representative PDB of each N-cytokine complex was then used in subsequent analyses and comparisons.

HADDOCK Metrics. The HADDOCK system outputs multiple metrics for the predicted binding affinities and an output set of PDB files containing the N protein docked against the cytokine protein. Some main metrics include:

- van der Waals intermolecular energy (*Evdw*)
- Electrostatic intermolecular energy (*Eelec*)
- Desolvation energy (*Edesol*)
- Restraints violation energy (*Eair*)
- Buried surface area (*BSA*)
- HADDOCK score: $1.0Evdw + 0.2Eelec + 1.0Edesol + 0.1Eair$

PRODIGY ΔG Prediction. PRODIGY, a tool to predict the binding affinity of protein-protein complexes from structural data, was used on each complex in this study (13). The predicted binding affinities are reported as Gibbs energy, shown as ΔG (in Kcal/mol units).
The predicted ΔG values are calculated by counting the number of various interfacial contacts (ICs) between the chains of the input complex (N protein and cytokine) along with some properties of the non-interacting surfaces (NIS) using the following equation:

$$\begin{aligned}\Delta G_{predicted} = & -0.09459 \cdot ICs_{charged/charged} \\ & -0.10007 \cdot ICs_{charged/apolar} \\ & +0.19577 \cdot ICs_{polar/polar} \\ & -0.22671 \cdot ICs_{polar/apolar} \\ & +0.18681 \cdot NIS\%_{apolar} \\ & +0.38100 \cdot NIS\%_{charged} \\ & -15.9433\end{aligned}$$

FoldX ΔG Prediction. FoldX (28), a tool that evaluates protein-protein complex interactions and their stability, was also used to predict ΔG for all of the complexes. In the command line, the AF2 and HADDOCK structures were first repaired using the PDB repairing function, RepairPDB, which minimizes the complex energy by rotating identified side chain residues.
FoldX calculates ΔG for the repaired structures using a linear combination of empirically derived energy terms including van der Waals , solvation energy for apolar and polar groups, and electrostatic contribution of charged groups.

Graph-based Residue Assessment Function (G-RAF). The interfaces of all complexes were first processed with *INTERCAAT* (29) to produce a full list of all interacting residue pairs between the N-protein and bound cytokine. A bigraph was created using Python NetworkX v3.1 (30) where the nodes were defined as each residue and the edges were defined as the predicted distance between the residues (in angstroms) by *INTERCAAT*. Distances longer than 3 angstroms were discarded. A lookup table was constructed for the 17 N-proteins using a consensus sequence from multiple sequence alignment (Geneious, Dotmatics, Inc) so that each residue number was aligned with the consensus. Baseline graphs were generated for SARS-CoV-2-WA1 with CXCL12β from both AlphaFold2 and HADDOCK predicted complexes. Graphs were then generated for the other N-proteins with CXCL12β. Graph Edit Distance (GED) was computed between each of the other N-proteins with SARS-CoV-2-WA1. GED was then computed between each N:CXCL12β complex from AlphaFold2 and HADDOCK.
The number of interfacing residues on N-protein from both AlphaFold2 and HADDOCK were enumerated. GED between a pair of graphs was defined as follows:

$$GED(g_1, g_2) = \min_{(e_1,\ldots,e_k)\in\mathcal{P}(g_1,g_2)} \Sigma_{i=1}^{k} c(e_i)$$

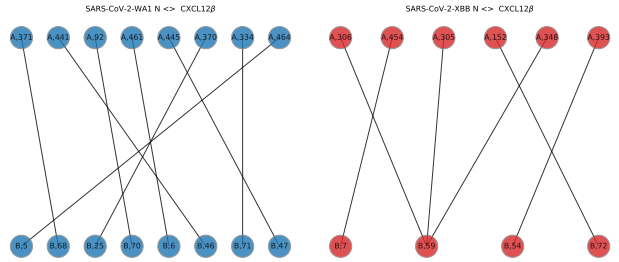

Fig. 10. Example bigraphs for SARS-CoV-2-WA1 (A chain, left) and SARS-CoV-2-XBB (A chain, right) bound to CXCL12β (B chains) on models predicted by HADDOCK.

GED minimization was capped at 10 seconds as we saw that further significant cost minimization over the set of edit paths did not occur beyond this point. GED costs were equally weighted a value of 1 for substitutions, insertions, or deletions.

Protein Structure and Results Visualization. Protein complexes were visualized using using PyMOL (31), except Figure 1. Also, PyMol's analysis capabilities were employed to detect and show interfacing residues (polar interactions within 3.0Å) between the N protein and cytokine in each complex.
Graphs were rendered using the ggplot2 v3.4.2 R package (32) or Matplotlib v3.8.1 (33). The phylogenetic tree figure was rendered using the ggtree v3.6.2 R package (34).

N Protein Alignment, Genetic Distance, and Tree Generation. The sequences for the N protein of each variant were aligned using Muscle v3.8.425 (35) and was written in FASTA format. From the aligned FASTA file, the pairwise distances were computed using the Fitch matrix (36) from the seqinr v4.2.30 R package (37). For Figures 4 and 5, the distance from the SARS-CoV-2 WA1 isolate was used as the reference point.
A maximum likelihood phylogenetic tree of the N protein alignment was generated with RAxML v8.2.12 (38). This tree was generated using the "rapid bootstrap" mode with 100 replicates and rooted on the OC43 taxon.

Author Contributions

PJT and RJ3 conceptualized and obtained funding for the project. PJT and CTF drafted the original manuscript draft. RJ3 and AMM developed the pipeline to run AlphaFold2 and provided assistance with debugging the HADDOCK software container. CTF developed the pipeline to run HADDOCK and PRODIGY, generated the PDB model figures, performed correlations for cytokine binding affinity and Fitch distance, and the phylogenetic tree. RJ3 developed the GDockScore pipeline to select the AlphaFold2 structures and Graph Edit Distance (G-RAF) algorithm. AEM performed FoldX and statistical analysis. PJT, CTF, RJ3, and AEM all analyzed data and edited the manuscript.

Acknowledgements

The authors would like to thank Alberto D. López-Muñoz and Jonathan Yewdell at the National Institutes of Health for guidance on understanding coronavirus nucleocapsid and cytokine interactions. The authors also want to acknowledge Darrell Ricke (MIT LL) for initial help with debugging AlphaFold2, Keegan Quiqley (MIT LL) and Menghua Rachel Wu (MIT CSAIL) for helpful suggestions, the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing resources and consultation, the MIT BeaverWorks facilities, and Jon Halter from the UNC Charlotte University Research Computing group.

Data availability statement

All code, data, results, and additional analyses are openly available on GitHub at:
https://github.com/tuplexyz/SARS-CoV-2_N-Cytokine_Docking.
This repository includes the PDB files for the top 1,088 complexes for both AlphaFold2 and HADDOCK experiments, all output metrics, experiment generation and data preparation logic, HPC submission scripts, and the code for and post-processing analyses and generating figures.

Conflict of interest disclosure

Author CTF is the owner of Tuple, LLC, a biotechnology consulting firm. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Legal and funding statement

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.
This material is based upon work supported by the Department of the Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Air Force.

Bibliography

1. Brady Johnston, Yuxuan Zhuang, Yinying Yao, William McCorkindale, Johannes Elferich, Patrick Kunzmann, Rich, Olivier Laprevote, Thibault Tubiana, Domenico Marson, James Hooker, Jessica A. Nash, and Joyce Kim. BradyAJohnston/MolecularNodes: v2.10.0 for Blender 3.5+, September 2023.
2. Willett BJ, Grove J, MacLean OA, Wilkie C, De Lorenzo G, Furnon W, Cantoni D, Scott S, Logan N, Ashraf S, Manali M, Szemiel A, Cowton V, Vink E, Harvey WT, Davis C, Asamaphan P, Smollett K, Tong L, Orton R, Hughes J, Holland P, Silva V, Pascall DJ, Puxty K, da Silva Filipe A, Yebra G, Shaaban S, Holden MTG, Pinto RM, Gunson R, Templeton K, Murcia PR, Patel AH, Klenerman P, Dunachie S, PITCH Consortium, COVID-19 Genomics UK (COG-UK) Consortium, Haughney J, Robertson DL, Palmarini M, Ray S, and Thomson EC. SARS-CoV-2 Omicron is an immune escape variant with an altered cell entry pathway. *Nat Microbiol*, 7(8):1161–1179, 08 2022. doi: 10.1038/s41564-022-01143-7.
3. Emma C Wall, Mary Wu, Ruth Harvey, Gavin Kelly, Scott Warchal, Chelsea Sawyer, Rodney Daniels, Philip Hobson, Emine Hatipoglu, Yenting Ngai, Saira Hussain, Jerome Nicod, Robert Goldstone, Karen Ambrose, Steve Hindmarsh, Rupert Beale, Andrew Riddell, Steve Gamblin, Michael Howell, George Kassiotis, Vincenzo Libri, Bryan Williams, Charles Swanton, Sonia Gandhi, and David LV Bauer. Neutralising antibody activity against SARS-CoV-2 VOCs B.1.617.2 and B.1.351 by BNT162b2 vaccination. *The Lancet*, 397(10292):2331–2333, 2021. ISSN 0140-6736. doi: https://doi.org/10.1016/S0140-6736(21)01290-3.
4. Emma B. Hodcroft. CoVariants: SARS-CoV-2 Mutations and Variants of Interest, 2021.
5. Abdullah M. Syed, Taha Y. Taha, Takako Tabata, Irene P. Chen, Alison Ciling, Mir M. Khalid, Bharath Sreekumar, Pei-Yi Chen, Jennifer M. Hayashi, Katarzyna M. Soczek, Melanie Ott, and Jennifer A. Doudna. Rapid assessment of SARS-CoV-2–evolved variants using virus-like particles. *Science*, 374(6575):1626–1632, 2021. doi: 10.1126/science.abl6184.
6. Alberto Domingo López-Muñoz, Ivan Kosik, Jaroslav Holly, and Jonathan W. Yewdell. Cell surface SARS-CoV-2 nucleocapsid protein modulates innate and adaptive immunity. *Science Advances*, 8(31):eabp9770, 2022. doi: 10.1126/sciadv.abp9770.
7. Alberto Domingo López-Muñoz, Jefferson J. S. Santos, and Jonathan W. Yewdell. Cell surface nucleocapsid protein expression: A betacoronavirus immunomodulatory strategy. *Proceedings of the National Academy of Sciences*, 120(28):e2304087120, 2023. doi: 10.1073/pnas.2304087120.
8. Alí Alejo, M. Begoña Ruiz-Argüello, Yin Ho, Vincent P. Smith, Margarida Saraiva, and Antonio Alcami. A chemokine-binding domain in the tumor necrosis factor receptor from variola (smallpox) virus. *Proceedings of the National Academy of Sciences*, 103(15):5995–6000, 2006. doi: 10.1073/pnas.0510462103.
9. Xiaoguang Xue, Qingyu Lu, Hui Wei, Dongli Wang, Dongwei Chen, Guangjun He, Li Huang, Hanzhong Wang, and Xinquan Wang. Structural Basis of Chemokine Sequestration by CrmD, a Poxvirus-Encoded Tumor Necrosis Factor Receptor. *PLOS Pathogens*, 7(7):1–13, 07 2011. doi: 10.1371/journal.ppat.1002162.
10. Richard Evans, Michael O'Neill, Alexander Pritzel, Natasha Antropova, Andrew Senior, Tim Green, Augustin Žídek, Russ Bates, Sam Blackwell, Jason Yim, Olaf Ronneberger, Sebastian Bodenstein, Michal Zielinski, Alex Bridgland, Anna Potapenko, Andrew Cowie, Kathryn Tunyasuvunakool, Rishub Jain, Ellen Clancy, Pushmeet Kohli, John Jumper, and Demis Hassabis. Protein complex prediction with AlphaFold-Multimer. *bioRxiv*, 2022. doi: 10.1101/2021.10.04.463034.
11. G.C.P. van Zundert, J.P.G.L.M. Rodrigues, M. Trellet, C. Schmitz, P.L. Kastritis, E. Karaca, A.S.J. Melquiond, M. van Dijk, S.J. de Vries, and A.M.J.J. Bonvin. The HADDOCK2.2 Web Server: User-Friendly Integrative Modeling of Biomolecular Complexes. *Journal of Molecular Biology*, 428(4):720–725, 2016. doi: 10.1016/j.jmb.2015.09.014.
12. Matthew McFee and Philip M Kim. GDockScore: a graph-based protein–protein docking scoring function. *Bioinformatics Advances*, 3(1):vbad072, 06 2023. ISSN 2635-0041. doi: 10.1093/bioadv/vbad072.
13. Anna Vangone and Alexandre MJJ Bonvin. Contacts-based prediction of binding affinity in protein–protein complexes. *eLife*, 4:e07454, jul 2015. ISSN 2050-084X. doi: 10.7554/eLife.07454.
14. Y Peng, N Du, Y Lei, S Dorje, J Qi, T Luo, GF Gao, and H Song. Structures of the SARS-CoV-2 nucleocapsid and their perspectives for drug design. *EMBO J.*, 39(20):e105938, 10 2020. doi: 10.15252/embj.2020105938.
15. Albert Reuther, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Lauren Milechin, Julia Mullen, Andrew Prout, Antonio Rosa, Charles Yee, and Peter Michaleas. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.
16. Nydia Panitz, Stephan Theisgen, Sergey A Samsonov, Jan-Philip Gehrcke, Lars Baumann, Kathrin Bellmann-Sickert, Sebastian Köhling, M Teresa Pisabarro, Jörg Rademann, Daniel Huster, and Annette G Beck-Sickinger. The structural investigation of glycosaminoglycan binding to CXCL12 displays distinct interaction sites. *Glycobiology*, 26(11):1209–1221, 11 2016. ISSN 0959-6658. doi: 10.1093/glycob/cww059.
17. Helena Crijns, Vincent Vanheule, and Paul Proost. Targeting Chemokine—Glycosaminoglycan Interactions to Inhibit Inflammation. *Frontiers in Immunology*, 11, 2020. ISSN 1664-3224. doi: 10.3389/fimmu.2020.00483.
18. Guo-Lin Wang, Hui-Xia Gao, Yu-Ling Wang, Xiao Wei, Yu-Zhen Liu, Jian-Hua Lu, Li Li, Hai-Bin Wang, Lei Zhao, Yan-Xiao Rong, Lin Yao, Li-Juan Duan, Benjamin D. Anderson, Mai-Juan Ma, Er-Hei Dai, and Xiang-Na Zhao. Serum IP-10 and IL-7 levels are associated with disease severity of coronavirus disease 2019. *Cytokine*, 142:155500, 2021. ISSN 1043-4666. doi: https://doi.org/10.1016/j.cyto.2021.155500.
19. Y Yang, C Shen, J Li, J Yuan, J Wei, F Huang, F Wang, G Li, Y Li, L Xing, L Peng, M Yang, M Cao, H Zheng, W Wu, R Zou, D Li, Z Xu, H Wang, M Zhang, Z Zhang, GF Gao, C Jiang, L Liu, and Y. Liu. Plasma IP-10 and MCP-3 levels are highly associated with disease severity and predict the progression of COVID-19. *J Allergy Clin Immunol.*, 146(1):119–127.e4, 07 2020. doi: 10.1016/j.jaci.2020.04.027.
20. Y Chen, J Wang, C Liu, L Su, D Zhang, J Fan, Y Yang, M Xiao, J Xie, Y Xu, Y Li, and S. Zhang. IP-10 and MCP-1 as biomarkers associated with disease severity of COVID-19. *Mol Med.*, 29(26):97, 10 2020. doi: 10.1186/s10020-020-00230-x.
21. Colby T. Ford, Denis Jacob Machado, and Daniel A. Janies. Predictions of the SARS-CoV-2 Omicron Variant (B.1.1.529) Spike Protein Receptor-Binding Domain Structure and Neutralizing Antibody Interactions. *Frontiers in Virology*, 2, 2022. ISSN 2673-818X. doi: 10.3389/fviro.2022.830202.
22. Colby T. Ford, Shirish Yasa, Denis Jacob Machado, Richard Allen White, and Daniel A. Janies. Predicting changes in neutralizing antibody activity for SARS-CoV-2 XBB.1.5 using *in silico* protein modeling. *Frontiers in Virology*, 3, 2023. ISSN 2673-818X. doi: 10.3389/fviro.2023.1172027.
23. Laura A. VanBlargan, John M. Errico, Peter J. Halfmann, Seth J. Zost, James E. Crowe, Lisa A. Purcell, Yoshihiro Kawaoka, Davide Corti, Daved H. Fremont, and Michael S. Diamond. An infectious SARS-CoV-2 B.1.1.529 Omicron virus escapes neutralization by therapeutic monoclonal antibodies. *Nature Medicine*, 28(3):490–495, Mar 2022. ISSN 1546-170X. doi: 10.1038/s41591-021-01678-y.
24. Aki Sugano, Junon Murakami, Haruyuki Kataguchi, Mika Ohta, Yoshiaki Someya, Shigemi Kimura, Akira Kanno, Yoshimasa Maniwa, Toshihide Tabata, Kazuyuki Tobe, and Yutaka Takaoka. *In silico* binding affinity of the spike protein with ACE2 and the relative evolutionary distance of S gene may be potential factors rapidly obtained for the initial risk of SARS-CoV-2. *Microbial Risk Analysis*, 25:100278, 2023. ISSN 2352-3522. doi: https://doi.org/10.1016/j.mran.2023.100278.
25. Sakshi Piplani, Puneet Kumar Singh, David A. Winkler, and Nikolai Petrovsky. *In silico* comparison of SARS-CoV-2 spike protein-ACE2 binding affinities across species and implications for virus origin. *Scientific Reports*, 11(1):13063, Jun 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-92388-5.
26. John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A A Kohl, Andrew J Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, August 2021.
27. Marc F. Lensink and Shoshana J. Wodak. Score_set: A CAPRI benchmark for scoring protein complexes. *Proteins: Structure, Function, and Bioinformatics*, 82(11):3163–3169, 2014. doi: https://doi.org/10.1002/prot.24678.
28. Joost Schymkowitz, Jesper Borg, Francois Stricher, Robby Nys, Frederic Rousseau, and Luis Serrano. The FoldX web server: an online force field. *Nucleic Acids Research*, 33 (suppl_2):W382–W388, 07 2005. ISSN 0305-1048. doi: 10.1093/nar/gki387.
29. Steven Grudman, J Eduardo Fajardo, and Andras Fiser. INTERCAAT: identifying interface residues between macromolecules. *Bioinformatics*, 38(2):554–555, 09 2021. ISSN 1367-4803. doi: 10.1093/bioinformatics/btab596.
30. Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
31. Schrödinger, LLC. The PyMOL molecular graphics system, version 1.8. November 2015.
32. Hadley Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer International Publishing, 2016. doi: 10.1007/978-3-319-24277-4.
33. J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.
34. Guangchuang Yu, David Smith, Huachen Zhu, Yi Guan, and Tommy Tsan-Yuk Lam. ggtree: an R package for visualization and annotation of phylogenetic trees with their covariates and other associated data. *Methods in Ecology and Evolution*, 8:28–36, 2017. doi: 10.1111/2041-210X.12628.
35. Robert C. Edgar. MUSCLE: a multiple sequence alignment method with reduced time and space complexity. *BMC Bioinformatics*, 5(1):113, Aug 2004. ISSN 1471-2105. doi: 10.1186/1471-2105-5-113.
36. Walter M. Fitch. An improved method of testing for evolutionary homology. *Journal of Molecular Biology*, 16(1):9–16, 1966. ISSN 0022-2836. doi: https://doi.org/10.1016/S0022-2836(66)80258-9.
37. D. Charif and J.R. Lobry. SeqinR 1.0-2: a contributed package to the R project for statistical computing devoted to biological sequences retrieval and analysis. In U. Bastolla, M. Porto, H.E. Roman, and M. Vendruscolo, editors, *Structural approaches to sequence evolution: Molecules, networks, populations*, Biological and Medical Physics, Biomedical Engineering, pages 207–232. Springer Verlag, New York, 2007. ISBN : 978-3-540-35305-8.
38. Alexandros Stamatakis. RAxML version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics*, 30(9):1312–1313, 01 2014. ISSN 1367-4803. doi: 10.1093/bioinformatics/btu033.