# AlphaBridge: tools for the analysis of predicted macromolecular complexes

*Daniel Álvarez-Salmoral, Razvan Borza, Ren Xie, Robbie P. Joosten, Maarten L. Hekkelman,*
*Anastassis Perrakis*

## Abstract

Artificial intelligence (AI)-powered protein structure prediction methods have revolutionised how life scientists explore macromolecular function. Using AI to predict the structure of macromolecular complexes, is gaining attention for modelling known interactions and evaluating the likelihood of proteins forming multimers or interacting with other proteins or nucleic acids. There is a growing need for tools to efficiently evaluate these predicted models. We introduce new tools that use AlphaFold3's "predicted local-distance difference test" (pLDDT), "predicted aligned error" (PAE), and "predicted distance error" (PDE) matrices in a graph-based community clustering approach to label sequence motifs involved in binary interactions. The resulting "interaction image" is processed through a multidimensional image algorithm to cluster interacting sequence motifs in three-dimensional binary interfaces. This method allows us to present the interaction information between multiple proteins and nucleic acids back to two-dimensional space using chord diagrams. These "AlphaBridge" diagrams summarise predicted interfaces and intermolecular interactions, including prediction confidence and sequence conservation scores. They are valuable for efficient screening of predictions, ranking, and scoring the confidence of predicted interactions, prior to more detailed (and resource intensive) analysis.

## Introduction

Protein structure prediction algorithms have advanced rapidly extending also into predicting protein complexes. AlphaFold spearheaded this revolution: AlphaFold2 [1] initially showcased unprecedented accuracy in predicting individual protein structures. The AlphaFold Protein Structure Database established by Google DeepMind and EMBL-EBI [2] made these developments accessible to all. Concurrently, ColabFold offered a fast and accessible solution for predicting protein structures to all scientists through an intuitive user-friendly interface and by leveraging powerful computing resources using efficient search algorithms, matching AlphaFold's accuracy on critical benchmarks like CASP14 [3] and ClusPro4 [4]. Independent approaches most notably included RoseTTAFold [5] and the ESM Metagenomic Atlas [6].

Extending prediction from individual protein structures to the structures of protein complexes, rapidly followed. AlphaFold-Multimer was trained specifically for inputs of known stoichiometry, allowing modelling of protein-protein interactions [7]. Derivative approaches included the AlphaPulldown [8] idea to use prediction for enabling "virtual pull-down" experiments for validating plausible protein-protein interactions, as well as Predictomes [9] which demonstrated the power for this "virtual pull-down" approach for discovery of new biology, in a closed set of proteins involved in genome maintenance (Schmid & Walter, 2023). Importantly, the issue of homo-oligomerization of key proteomes has also been addressed [11]. In addition, the AlphaFill approach [10] provided a resource of potential complexes for predicted protein models with about 2,700 ligands that are present in experimental structures. Most recently, AlphaFold3 [12] not only provided algorithmic innovations to the protein structure prediction problem, but also allowed efficient and reliable prediction of protein complexes with not only other proteins but also with nucleic acids, and other common biological ligands. These latest AlphaFold developments, notwithstanding current limitations in code availability and licensing issues that preclude "pull-down" screening approaches, were incorporated in a user-friendly web server that popularised the use of protein structure predictions even further, making them available to an ever-growing user base. Similar developments, for example RoseTTAFoldNA [13] for protein-nucleic acid interactions, or NeuralPLexer [14] for

directly predicting protein–ligand complex structures using ligand molecular graph inputs, are further pushing the boundaries of deep learning applications for studying macromolecular complexes.

Predicted models are always closely coupled to metrics allowing to validate the quality of the prediction, locally and globally. While available reliability metrics are an excellent aid for experts and can be understood well by new keen users of the technology, new tools that validate and visualise specifically the reliability and nature of interactions and predicted protein complexes are of essence. The advent of AlphaFold3, with its straightforward availability, speed, and ability to predict protein complexes with proteins, DNA, RNA, as well as select ions, sugars, lipids and biological ligands, draws researchers from different fields into structural biology. We present AlphaBridge, a collection of tools to post-process and analyse information on interaction interfaces between predicted macromolecular complex components and visualise the most relevant information in an accessible and intuitive manner to scientists interested in macromolecular complexes.

# Results

## Construction of the AlphaBridge diagrams

For a given biomolecular complex structure model (e.g. replication factor A, a complex between three protein chains, Figure 1A) AlphaFold3 provides three confidence metrics (Figure 1B): the predicted local distance difference test (pLDDT), which measures confidence in the placement of each atom in its local environment; the pairwise aligned error (PAE), quantifying the error in the relative position of two "tokens" (a residue for a protein, a nucleotide for DNA and RNA, and an atom for a ligand); and the predicted distance error (PDE), defined as the confidence that two tokens are in contact (i.e. closer than 8Å).

We combine the PAE and pLDDT information, in a new matrix we call the Predicted Merged Confidence (PMC) which is used to find "multi-component modules", non-sequential regions of complex components that are confidently predicted to form a distinct structural unit (Figure 1C). To detect such modules, we implemented a graph-based community clustering approach [15] based on the work of T. Croll [16] to detect protein domains implementing the Leiden community clustering algorithm. Here, rather than extracting protein domains from a single structure, we combine regions from multiple biomolecules in the predicted biomolecular complex. These multi-component modules are the basic units inside which we proceed to detect the interacting binary interfaces.

To define the interfaces, we first mask-out from the PDE matrix the regions outside the multi-component modules (Figure 1D). Next, we detect 'contact links', defined as sequential regions between two biomolecules that have a score higher than a set value in the unmasked regions of the PDE matrix (Figure 1D). For detecting contact links we follow a multidimensional image processing approach, using two-pass connected-component analysis. Sequential contact links that are less than two residues (or nucleotides) apart, are combined to a single link. The collection of all contact links between two components of each multi-component module, forms an interacting binary interface.

This process is iterated for all multi-component modules, until all binary interfaces are characterised (Figure 1E). This way, for each biomolecular complex, we define several interacting interfaces, and for each interacting interface we define contact links. The number of contact links and interfaces can vary depending on the cut-off value used for masking the PDE matrix. Varying this value can increase or decrease the confidence with which contacts and interfaces are predicted.

This evaluation procedure is useful two-fold: first, to visualise confidently predicted interfaces to help expert and non-expert users to better evaluate AlphaFold predictions; and second, to predict biomolecular assemblies based on the extent of local contacts providing a complementary criterion for scoring and ranking the quality of predicted interactions.

The assigned interfaces and contact links are then shown in a circular layout [17] as a chord diagram (Figure 1F). Each chain of the predicted complex is shown as a thick arc along the inner circumference, with lengths proportional to its sequence length. Interfaces are depicted as curves, or "bridges", between the contact links, with different colours for each interface. The outer circumference displays the pLDDT for each chain as a thick arc. This type visualisation integrates various complementary information sources into a single, easily interpretable visualisation allowing users to quickly assess the number and confidence of all interactions and the number of interfaces in each complex.

**Figure 1. Description of the AlphaBridge algorithm. (A)** The structure of the human replication protein A, a complex between three proteins (PDB: 8RK2) coloured by chain. **(B)** The AlphaFold3 model for the same complex coloured by the pLDDT score (left), the corresponding PAE (middle) and PDE (right) matrices. **(C)** the Predicted Merged Confidence (PMC) matrix with the regions assigned to the single detected module annotated by three blue bars on top (left) and the predicted structure with the detected module annotated (right). **(D)** A combination of the PMC matrix (upper right half) and the PDE matrix (lower left half) with the three interfaces defined between interacting sequential regions between binary combinations of the biomolecules within the module annotated (left) and the corresponding contact links for each interface (right). **(E)** The three interfaces mapped in the structure of the complex. **(F)** The AlphaBridge plot for this predicted complex.

## The AlphaBridge webserver

Online analysis of AlphaFold3 results can be performed by uploading the AlphaFold3 server "zip" file to https://alpha-bridge.eu/. All information is processed automatically and an interactive AlphaBridge plot is constructed to a viewport to the left, while an interactive 3D structure presentation is shown in the right viewport (Figure 2A). The two viewports are connected, so clicking to a contact link displays that interface. Mouse-over on contact links displays the residue numbers involved, while mouse-over on the outer arc shows information for the residue. The contact link details for each interface are also shown in a table with collapsible menus; the table is interactively linked to the 3D viewport, thus clicking to a link focuses on the relevant residues in the 3D viewer. The user can toggle between structure colouring by chain, by pLDDT, or by interface (Figure 2C-E). The AlphaBridge plot can be output as an SVG file for use in presentations or publications.

**Figure 2.** The AlphaBridge web server showing a prediction for the structure of the IgG-BCR complex (PDB: 7WSO) between IGHG1, CD79A, CD79B proteins **(A)** The AlphaBridge diagram to the left, witch each interface in a different colour, and the 3D-view in the right **(B)** The list of all interactions by interface, with details on "interface 5" shown. **(C-F)** The 3D-viewport zoomed into that interface, coloured by chain, (C) by pLDDT value (D) highlighting the contact links within the interface, for the main chain (E) and all involved side chains (F).

## Variations in PDE filtering as a tool to evaluate interactions

A key parameter, is the cut-off value used for masking the PDE matrix to select the number of contact links and interfaces. To evaluate how different PDE cutoffs affect the ability to distinguish between reliable and less reliable predictions, we examined the prediction of human mismatch repair dimers binding to DNA. In prokaryotes, the MutS DNA mismatch repair complex is a homodimer. In humans, MutSα[18] is a dimer between MSH2 and MSH6 that preferentially recognises DNA mismatches and MutSβ[19] is a dimer between MSH2 and MSH3 which preferentially recognises looped out nucleotides. A dimer of MSH3 and MSH6 is not physiologically relevant. We analysed all three dimers using the same DNA sequence containing a single nucleotide mismatch.

The AlphaFold3 models look reasonable, and all three PAE matrices indicate that all three protein pairs and the mismatched DNA form a reliable complex (Figure 2A). We then run the AlphaBridge algorithm with three PDE cutoffs: 0.5, 0.7, and 0.9: using a cutoff of 0.5, all three dimers seems to interact with each other and with the mismatch; at a cutoff of 0.7, MutSβ does no longer recognise the mismatch, and the non-physiological dimer MSH2:MS3 does not recognise the mismatch and "loses" the interface between the N-terminal domains; finally, at a cutoff of 0.9 the contact links between the two "wrong" complexes diminish further, while the one for the "correct" MutSα complex is still clear. The persistence in predicting reliable dimerization between the C-terminal ATPase domains may be due to the region's high conservation across all members.

These examples and other empirical data suggest a cut-off value of 0.7 as default.

**Figure 3. Variations in PDE filtering for selecting physiologically relevant mismatch repair complexes:** from left to right: the models coloured by chain (MSH2:blue; MSH6:green; MSH3: brown; DNA: grey), the PAE matrix from AlphaFold3, and the AlphaBridge plot at three PDE cut-off values. The AlphaBridge diagrams at different PDE values suggest which complexes are more reliable and therefore more physiologically relevant.

**Protein-nucleic acids interactions**

To further explore AlphaBridge's ability to help distinguish between accurate and inaccurate predictions in the context of the interaction with different nucleotides, we tested it between two pumilio family members (PUM2 and PUM3) and RNA. PUM2 is known to bind single-stranded RNA (PDB:3Q0Q; [20]), while PUM3 has been reported to bind double-stranded DNA (PDB:4WZW; [21]). We then used AlphaFold3 to predict a model of each protein with single-stranded RNA. Both predictions look correct, with high pLDDT scores and good PAE values (Figure 4). However, the AlphaBridge algorithm, upon analysing the PDE matrix, delivers an AlphaBridge diagram that clearly shows the (correct) PUM2-RNA interaction, but does not positively evaluate the (incorrect) PUM3-RNA interaction. A prediction of the structure of PUM2 and PUM3 with or double-stranded DNA, shows the opposite result (Supplemental figure X).

**Figure 4. Differentiating between complexes with different nucleic acids.** From left to right: the prediction of the PUM2 and PUM3 complexes with RNA or DNA; the same coloured by pLDDT; the corresponding PAE matrix; the PDE matrix; and the AlphaBridge plot suggesting that only the complex of PUM2 with RNA is correct. The AlphaBridge diagrams correctly indicate that PUM2 binds RNA and PUM3 binds DNA.

# Discussion

AlphaBridge has been designed to aid all scientists interpret predictions of the structure of macromolecular complexes. On the one hand, it provides an automated, reproducible, and reliable algorithm to detect linear sequence motifs (contact links) involved in binary interfaces between predicted structures of macromolecules, taking into account the metrics that describe the reliability of the predictions. On the second hand, it offers an interactive web interface connecting an intuitive way to present interactions in "AlphaBridge diagrams" (a useful way to present information in presentations and publications) together with live 3D visualisation (suitable for interactive analysis).

Currently, AlphaBridge works with results from the AlphaFold3 server. It is able to process predictions that include complexes between proteins and nucleic acids (DNA and RNA). An obvious extension includes the handling of post-translational modifications (PTMs) in proteins, and the processing of the various biological ligands available. The algorithm can be both integrated in additional tools (e.g. OpenFold, ColabFold or Predictomes) or can be straightforwardly enhanced to process output from such servers.

The probability contact matrix, allows to differentiate reliable interactions in an accessible and reproducible manner that goes beyond straight-forward inspection of the pLDDT and PAE matrix information. The delineation of the linear motifs that make the contact links, allows the straightforward inspection of regions of interaction for non-experts. In addition, the detection of these contacts using a single cut-off value as threshold holds promise for implementing new metrics that can rank and score the robustness of complex predictions, an issue that is already being addressed [9] but needs further development.

The circular layout visualisation we employed, is on one hand familiar to many users (as it has been used to show interactions derived from cross-link proteomics applications), has the possibility to represent increasingly complex data in two dimensions, and can also be used to integrate additional information. For example, the display of AlphaMissense predictions (for human proteins) and sequence conservation data (for all proteins) is readily implemented. Including these additional layers of information, can help analyse and understand the functional implications of interactions, identify functionally important residues, or indicate potential disease-associated variants, enriching the biological relevance of the interactions.

AlphaBridge significantly advances both the visualisation and interpretation of AlphaFold predictions of macromolecular complexes, aspiring to specifically engage non-experts.

# Online Methods

## Confidence-contact integration plot

For a given biomolecular structure, AlphaFold3 provides three confidence metrics: the per-atom local confidence (pLDDT), which measures confidence in the local structure, estimating how well the prediction would align with an experimental structure; the pairwise token-token aligned error (PAE), quantifying the error in the relative position of two tokens within the predicted structure; and the pairwise token-token contact confidence (predicted distance error, PDE), estimating the confidence that two relative tokens are in contact (with a distance of 8 Å or less between the representative atoms for each token).

The PAE serves as an effective metric for assessing the confidence in domain packing and the accuracy of the relative domain placement within the predicted structure. However, it is also important to consider the confidence in the local structure, estimating how well the prediction would agree with an experimental structure provided by the pLDDT value. Variations in the pLDDT score across the macromolecular structure highlight which parts of the predicted model are more reliable and which are less so. Such variations might signify regions of high flexibility or intrinsic disorder, where the structure is either ill-defined or lacks sufficient information for a confident prediction.

Another limitation arises while assessing the PAE value of two related tokens without considering whether these residues engage in meaningful interactions or belong to the same functional site, which poses a notable constraint. It is possible to have high confidence that two residues are far apart within the structure. Lower PAE scores suggest higher confidence in the relative position of two tokens within the predicted structure. However, it does not always infer whether those residues are in close contact or interacting with each other.

These limitations become especially significant when trying to understand the functional implications of residue interactions within a protein. To address this, we developed a reinterpretation and visualisation pipeline that compensates for the limitations of relying solely on PAE values.

The Predicted Merged Confidence (PMC) is a two-dimensional graph where each token is represented along both the x and y axes, which merges the information into one single matrix, the overall confidence in the local structure and the relative positions of two tokens. This is achieved by combining the Predicted Alignment Error (PAE) and Predicted Local Distance Difference Test (pLDDT) scores into a pairwise token-token confidence score. For a given pair of tokens the reversed pairwise pLDDT average is obtained by considering the pLDDT score of all atoms (N) involved in both pairs of residues.

$$\overline{pLDDT}_{ij} = \frac{1}{N}( pLDDT_{ij} + pLDDT_{ji} ) - 100$$

A single value of the PMC score is estimated by the difference between the $PAE_{ij}$ and the $\overline{pLDDT}_{ij}$ :

$$\text{PCS} = PAE_{ij} - \frac{1}{3} \overline{pLDDT}_{ij}$$

## Defining intra-protein structural modules

The graph-based community clustering approach [15] we use is based on the work of T. Croll [16]. This had been originally implemented to detect protein domains in a single sequence and the PAE matrix; here it is modified to work with multiple sequences and the PMC matrix. The algorithm is based upon the Leiden community clustering algorithm, to identify groups of nodes that are more densely connected to each other than to the rest of the network, greedily optimising modularity in the content of the division of a network into clusters.

By considering also the pLDDT scores, we aim to eliminate the interfaces with a low confidence of the local structure from the pool of potential candidates of interacting interfaces. Rather than extracting protein domains

from a single monomer structure, the adapted methodology offers a compilation of structural modules from multiple biomolecules (multi-component modules) in the predicted biomolecular complex that share the same biological environment within the predicted structure. These multi-component modules are the basic units inside which we proceed to detect the interacting binary interfaces.

## Assigning interfaces

Given an assigned multi-component module, each binary biomolecular combination is considered within the predicted structure. By evaluating this subset of interactions, the search space is significantly narrowed down from considering all potential binary interactions across all entities in the predicted structure, to focusing on each specific binary interaction inside the co-evolutionary domain. This approach aims to improve both the accuracy of identifying interacting interfaces and the efficiency of computational processing.

For a given binary biomolecular interaction, all the PDE scores were extracted within the range of tokens (residues for proteins, nucleotides for DNA and RNA) involved in the interaction. By default, everything above a 0.7 PDE value is considered a contact. A multidimensional image processing approach is applied to detect sequential regions of connected pairwise tokens or 'contact links' using a two-pass connected-component analysis. For this approach a graph is created using as an input the PDE matrix. Each vertex holds the necessary information for the comparison heuristic, while the edges represent connections to neighbouring vertices. An algorithm navigates through the graph, assigning labels to the vertices based on their connectivity and the relative values of their neighbours. This connected-component analysis uses the scipy.ndimage module from the SciPy package [22].

Within a binary interaction, an interacting interface is defined by all shared contact links between two distinct biomolecular entities. Sequential contact links that are less than two residues (or nucleotides) apart, are combined to a single link. If a binary interaction contains $n$ distinct groups of non-shared contact links, $n$ interacting interfaces will be defined for that specific binary interaction. The collection of all shared contact links between two components of each multi-component module, forms an interacting binary interface

## Circular layout diagrams

The circos-like layout figures were created using the pyCirclize package [23], designed for plotting circular figures like Circos Plots and Chord Diagrams in Python. pyCirclize employs a circular layout with sectors and tracks, allowing different data types to be assigned to each sector. Multiple tracks for data plotting can be freely arranged within each sector.

## Implementation and Code availability

The code was developed in Python 3.9 with Anaconda as the package manager to ensure a consistent and reproducible environment. Key modules used include SciPy for performing the two-pass connected-component analysis and PyCirclize for creating circular layout diagrams. The complete code, along with setup instructions and environment configuration, is available in the public repository https://github.com/PDB-REDO/AlphaBridge, which includes detailed guidance on installation and usage.

## Web server implementation

The server was built using the *libzeep* library, which provides essential tools for HTTP server management, HTML templating, and various components for web server construction in C++. To handle mmCIF files for handling macromolecular model coordinates we integrated *libcif++*. For model visualization, we used *Mol\** as an interactive web component embedded directly into the webpage. To manage JavaScript dependencies, we employed *Yarn*, an open-source package manager. For interactive data visualisation, specifically for generating circular layout diagrams, we utilised the free and open-source library *D3.js*.

1.    Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583-589. doi:10.1038/s41586-021-03819-2

2.    Tunyasuvunakool K, Adler J, Wu Z, et al. Highly accurate protein structure prediction for the human proteome. *Nature*. 2021;596(7873):590-596. doi:10.1038/s41586-021-03828-1

3.    Kryshtafovych A, Schwede T, Topf M, Fidelis K, Moult J. Critical assessment of methods of protein structure prediction (CASP)—Round XIV. *Proteins: Structure, Function and Bioinformatics*. 2021;89(12):1607-1617. doi:10.1002/prot.26237

4.    Vajda S, Yueh C, Beglov D, et al. New additions to the ClusPro server motivated by CAPRI. *Proteins: Structure, Function and Bioinformatics*. 2017;85(3):435-444. doi:10.1002/prot.25219

5.    Baek M, Dimaio F, Anishchenko I, et al. *Accurate Prediction of Protein Structures and Interactions Using a Three-Track Neural Network*. Vol 373.; 2021. https://predictioncenter.org/casp14/

6.    Lin Z, Akin H, Rao R, et al. *Evolutionary-Scale Prediction of Atomic-Level Protein Structure with a Language Model*. Vol 379.; 2023. https://www.science.org

7.    Evans R, O'neill M, Pritzel A, et al. *Protein Complex Prediction with AlphaFold-Multimer*.; 2022.

8.    Yu D, Chojnowski G, Rosenthal M, Kosinski J. AlphaPulldown—a python package for protein–protein interaction screens using AlphaFold-Multimer. *Bioinformatics*. 2023;39(1). doi:10.1093/bioinformatics/btac749

9.    Schmid EW, Walter JC. *Predictomes: A Classifier-Curated Database of AlphaFold-Modeled Protein-Protein Interactions*.; 2023.

10.   Hekkelman ML, de Vries I, Joosten RP, Perrakis A. AlphaFill: enriching AlphaFold models with ligands and cofactors. *Nat Methods*. 2023;20(2):205-213. doi:10.1038/s41592-022-01685-y

11.   Schweke H, Pacesa M, Levin T, et al. An atlas of protein homo-oligomerization across domains of life. *Cell*. 2024;187(4):999-1010.e15. doi:10.1016/j.cell.2024.01.022

12.   Abramson J, Adler J, Dunger J, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*. 2024;630(8016):493-500. doi:10.1038/s41586-024-07487-w

13.   Baek M, McHugh R, Anishchenko I, Jiang H, Baker D, DiMaio F. Accurate prediction of protein–nucleic acid complexes using RoseTTAFoldNA. *Nat Methods*. 2024;21(1):117-121. doi:10.1038/s41592-023-02086-5

14.   Qiao Z, Nie W, Vahdat A, Miller TF, Anandkumar A. State-specific protein-ligand complex structure prediction with a multi-scale deep generative model. Published online September 29, 2022. http://arxiv.org/abs/2209.15171

15.   Traag VA, Waltman L, van Eck NJ. From Louvain to Leiden: guaranteeing well-connected communities. *Sci Rep*. 2019;9(1). doi:10.1038/s41598-019-41695-z

16.   Tristan Croll, Augustin Zidek. Graph-based community clustering approach to extract protein domains from a predicted aligned error matrix. https://github.com/tristanic/pae_to_domains.

17.   Krzywinski M, Schein J, Birol I, et al. *Circos: An Information Aesthetic for Comparative Genomics*.; 2009. http://mkweb.bcgsc.ca/circos.

18. Warren JJ, Pohlhaus TJ, Changela A, Iyer RR, Modrich PL, Beese LSS. Structure of the Human MutSα DNA Lesion Recognition Complex. *Mol Cell*. 2007;26(4):579-592. doi:10.1016/j.molcel.2007.04.018

19. Gupta S, Gellert M, Yang W. Mechanism of mismatch recognition revealed by human MutSβ bound to unpaired DNA loops. *Nat Struct Mol Biol*. 2012;19(1):72-79. doi:10.1038/nsmb.2175

20. Lu G, Hall TMT. Alternate modes of cognate RNA recognition by human PUMILIO proteins. *Structure*. 2011;19(3):361-367. doi:10.1016/j.str.2010.12.019

21. Qiu C, McCann KL, Wine RN, Baserga SJ, Hall TMT. A divergent Pumilio repeat protein family for pre-rRNA processing and mRNA localization. *Proc Natl Acad Sci U S A*. 2014;111(52):18554-18559. doi:10.1073/pnas.1407634112

22. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17(3):261-272. doi:10.1038/s41592-019-0686-2

23. Yuki Shimoyama. pyCirclize: Circular visualization in Python.