UNIVERSITY CEU - SAN PABLO

POLYTECHNIC SCHOOL

BIOMEDICAL ENGINEERING DEGREE

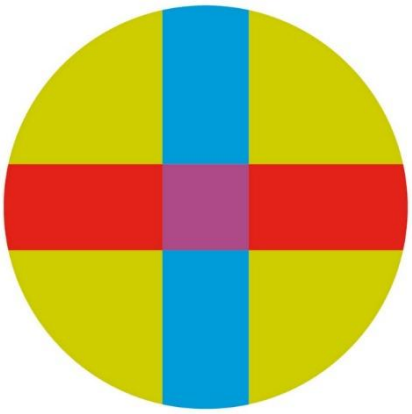

BACHELOR THESIS

Computational Optimization of Antibodies

Author: Javier Gallego Maestre
Supervisors: Carlos Oscar Sorzano Sánchez and Javier Tejedor Noguerales

January 2026

UNIVERSIDAD SAN PABLO-CEU
ESCUELA POLITÉCNICA SUPERIOR
División de Ingeniería

Datos del alumno

NOMBRE: Javier Gallego Maestre

Datos del Trabajo

TÍTULO DEL PROYECTO: Computational Optimization of Antibodies

Tribunal calificador

PRESIDENTE:	FDO.:
SECRETARIO:	FDO.:
VOCAL:	FDO.:

Reunido este tribunal el _____/_____/_____, acuerda otorgar al Trabajo Fin de Grado presentado por Don_____la calificación de _____.

ACKNOWLEDGMENTS

En primer lugar, me gustaría dar las gracias a Carlos Óscar Sorzano, por dedicarme su tiempo, por su disponibilidad constante y por su implicación a lo largo de todo el desarrollo de este Trabajo de Fin de Grado. También por darme la oportunidad de realizar un trabajo sobre un área que tanto me interesaba. Su orientación, sus consejos y su manera de guiar el trabajo han sido fundamentales para poder avanzar en cada una de las fases del proyecto y para afrontar con criterio los distintos retos que fueron surgiendo durante estos últimos meses de trabajo.

Quiero agradecer también a Javier Tejedor, segundo tutor del TFG, por sus aportaciones y por el seguimiento realizado a lo largo del trabajo. Sus comentarios y sugerencias han contribuido a mejorar el enfoque y a completar este estudio, aportando una visión complementaria muy valiosa.

Asimismo, deseo dar las gracias a todo el profesorado del grado que me ha acompañado durante estos cuatro años de formación universitaria. Los conocimientos adquiridos a lo largo de la carrera han sido esenciales para llegar hasta este momento y han sentado las bases necesarias para el desarrollo de este trabajo. También quiero dar a mis compañeros y compañeras de carrera, con quienes he compartido esta etapa, el esfuerzo compartido y por los momentos que hemos pasado juntos que han hecho este camino más enriquecedor tanto a nivel académico como personal. Sin ellos todos estos años no hubiesen sido lo mismo.

Por último, quiero expresar un agradecimiento muy especial a mi familia, por su apoyo incondicional, su paciencia y su confianza constante a lo largo de todo este tiempo. Han confiado en mi desde el primer momento y me han ayudado a conseguir todos los objetivos que me he propuesto a lo largo de estos años. Gracias por vuestra dedicación, y por apostar y confiar en mí durante toda la carrera.

ABSTRACT

Therapeutic antibodies are a central class of biopharmaceuticals used for the treatment of infectious diseases, cancer, and immune-related disorders. They usually require optimization to improve their affinity and stability. In this project, we developed a computational workflow in Scipion-Chem to analyze point mutations in antibodies. We tested our workflow with an antibody directed against the receptor-binding domain (RBD) of SARS-CoV-2.

The workflow included the import of the structure, the identification of paratope residues through a distance-based protocol, and the prediction of the energetic impact of mutations using two complementary tools: *SAAMBE-3D* (machine-learning model) and *FoldX* (empirical force field). Both tools generated independent ΔΔG predictions that were later integrated through the rank fusion protocol, using ZMUV normalization and the CombMED method to obtain a consensus ranking. The top-ranked mutations were A40F, A40D, A40K, A40Q and A59W, all of them located in key interfacial regions and predicted as stabilizing by both methods.

Overall, the study demonstrates that Scipion-Chem enables the systematic identification of potentially beneficial mutations and provides a reproducible framework for the computational optimization of antibodies.

RESUMEN

Los anticuerpos terapéuticos constituyen una clase central de biofármacos utilizados para el tratamiento de enfermedades infecciosas, cáncer y trastornos relacionados con el sistema inmunitario. Normalmente suelen requerir una optimización precisa de sus interfaces de unión para mejorar su afinidad y estabilidad. En este proyecto se desarrolló un flujo de trabajo computacional en Scipion-Chem para analizar mutaciones puntuales en anticuerpos. En particular, evaluamos un flujo de trabajo con un anticuerpo dirigido frente al dominio de unión al receptor (RBD) del SARS-CoV-2.

El flujo de trabajo incluyó la importación de la estructura, la identificación de los residuos del paratopo mediante un protocolo basado en distancias y la predicción del impacto energético de las mutaciones mediante dos herramientas complementarias: SAAMBE-3D (modelo de aprendizaje automático) y FoldX (campo de fuerza empírico). Ambas herramientas generaron predicciones independientes de ΔΔG que posteriormente se integraron mediante el protocolo de rank fusion, utilizando normalización ZMUV y el método CombMED para obtener una clasificación consensuada. Las mutaciones mejor posicionadas fueron A40F, A40D, A40K, A40Q y A59W, todas ellas localizadas en regiones clave del interfaz y predichas como estabilizadoras por ambos métodos.

En conjunto, el estudio demuestra que Scipion-Chem permite identificar de forma sistemática mutaciones potencialmente beneficiosas y proporciona un marco reproducible para la optimización computacional de anticuerpos.

.

INDEX

FIGURE INDEX

1 INTRODUCTION

1.1 Therapeutic antibodies: biological relevance, challenges, and the need for optimization

Nowadays in the biological drug development landscape, therapeutic antibodies have become a major player and a true growth engine for the pharmaceutical industry. Their biological effects are based on their ability to target specific molecular sites with impressive affinity and precision, allowing for incredibly accurate treatments across a range of diseases. Over the past three decades, antibodies have revolutionized how we approach inflammatory and infectious diseases, cancer, neurodegenerative disorders, and even metabolic syndromes. Their journey into clinical use has solidified their status as one of the most successful classes of biotherapeutics [1], [2], [3], [4].

Figure 1 illustrates the structural basis of antibody–antigen recognition, showing how the antigen binds specifically to the antibody paratope through a network of non-covalent interactions that define both binding affinity and specificity. Figure 2 complements this structural view by placing the interaction in a broader biological context, illustrating the role of antigens within the adaptive immune response.

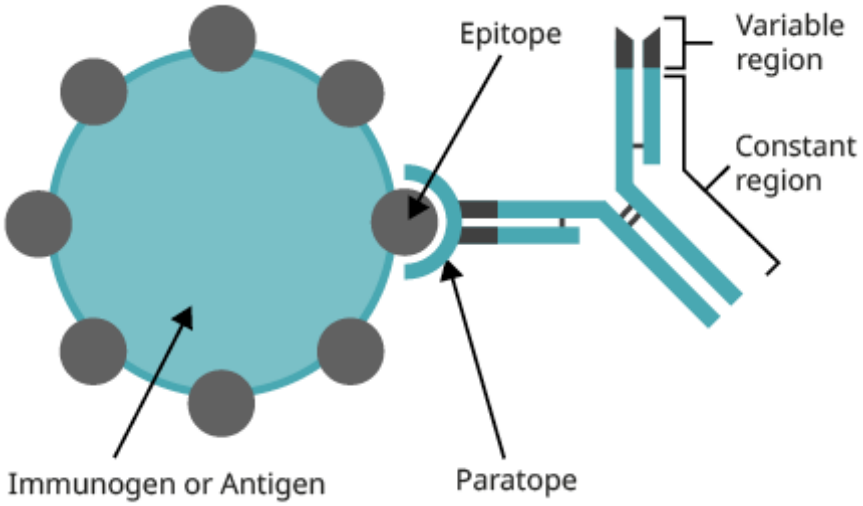

Figure 1 Structural representation of an antibody–antigen interaction.

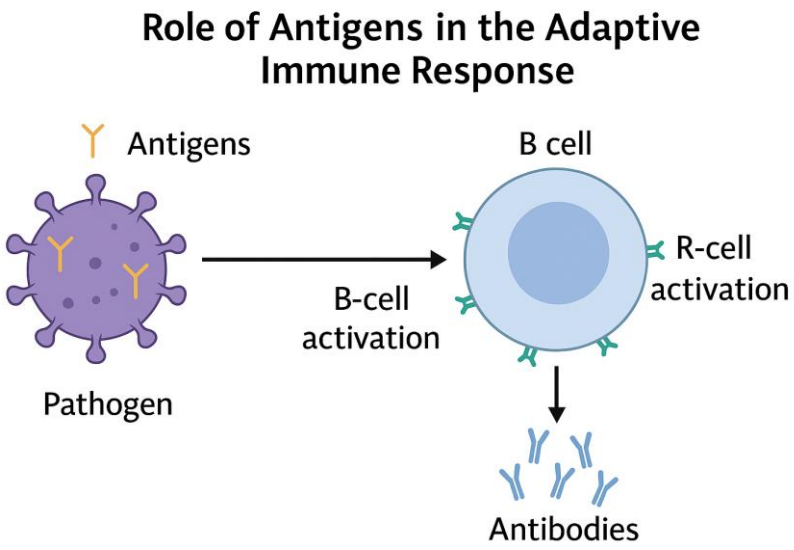

Figure 2 Role of antigens in the adaptive immune response.

The complementarity-determining regions of the variable domains are the site of a complex network of precisely calibrated molecular reactions that give rise to antibody function. These regions come together to form the so-called paratope, which is the part of the antibody that interacts with and binds to a specific area of an antigen known as the epitope. The underlying energetic and structural aspects of this interaction are quite intricate. Factors like hydrogen bonding, van der Waals forces, electrostatic interactions, hydrophobic packing, and the flexibility of the bound ligand all contribute uniquely to the overall affinity and specificity. Even a single amino acid can tip the scales significantly, meaning that even small changes can greatly affect stability, conformational dynamics, or how well an antigen is recognized [5], [6].

The immune system naturally produces antibodies, but these are not always perfect for therapeutic use. Most candidates need to demonstrate better affinity, reduced off-target binding, greater stability, or other improvements that make it easier to manufacture or formulate. The methods used to enhance the natural properties of these binding molecules are known as antibody optimization, and this process is crucial for transforming a natural binder into an effective therapeutic agent.

For many years, optimization primarily relied on experimental techniques like directed evolution, phage display, and error-prone PCR, along with high-throughput

mutagenesis screens. These approaches generated numerous mutants and enabled the selection of superior molecules through multiple rounds of testing. While they have been effective, they come with significant downsides. These methods can be slow and tedious, often requiring researchers to sort through thousands or even millions of mutations each time to pinpoint candidates with the desired biochemical or biophysical traits [7]. There is a practical limit to how much of the mutational space we can explore. Typically, the variable region of an antibody is made up of around two hundred amino acids, but the number of possible point mutations already surpasses three thousand. Trying to fully search for this space is just not feasible in most experimental scenarios [8].

A key challenge in antibody engineering is the vastness of the search space. Relying solely on brute-force experimentation to navigate these high-dimensional mutational landscapes just does not cut it. It also complicates our ability to predict how small changes can ripple through the protein structure or to pinpoint the role of individual residues at the interface. These limitations have led to a gradual shift towards methods that can guide experimental efforts with a solid understanding of structure, energy, and mechanics. As a result, computational antibody optimization has become an essential component of modern antibody design processes [9], [10].

1.2 Computational antibody engineering: principles, tools, and advantages

In the world of designing therapeutic antibodies, computational methods are playing an increasingly vital role. They research into the antibody–antigen interface, examining how each residue contributes to binding and predicting how specific mutations could impact affinity, stability, and overall developability. By providing a structured and mechanistic approach, these techniques not only save time and reduce costs but also ease the experimental burden by pinpointing the most promising mutations before any lab work begins [11], [12], [13].

When it comes to designing antibodies using computational methods, researchers typically follow a well-defined process. It all starts with techniques like crystallography or computational modeling to build a three-dimensional model of the antibody-antigen complex. Once they have that structure, the next step is to figure out

which residues are most important for interaction. This means identifying specific areas that are key for binding, making them ideal candidates for mutagenesis, all based on distance thresholds or energetic criteria.

The next step involves performing silico mutagenesis to determine the expected changes in binding free energy for each variant. These values, referred to as ΔΔG, give us insight into how a mutation might affect things. If the values come out positive, it indicates a negative impact, while negative values suggest that the mutation could stabilize or boost affinity. Once we crunch these numbers, we rank the mutations according to their anticipated energetic effects, which guides us in selecting the variants that are most likely to enhance the therapeutic properties of the antibody.

Machine learning has been gaining a lot of traction lately, especially since it can reveal intricate relationships between sequences and structures by utilizing large datasets. Deep learning models are also stepping up, playing a crucial role in optimizing potential therapies and predicting how antigens will behave based solely on their sequences [14]. In this context, we must make sure that model predictions are generated under consistent and well-defined conditions, allowing reliable comparison across different computational approaches [15], [16]. When using purely physics-based methods, these models often catch tiny interaction patterns that might slip under the radar otherwise.

When it comes to predicting ΔΔG, which is the change in binding free energy due to mutations at protein–protein interfaces, there is a whole range of tools designed for this crucial task. These tools are super helpful because they provide a clear idea of whether a mutation will boost or reduce binding. Some predictors rely on machine learning to estimate ΔΔG from structural or sequence data, while others take a more physics-based route, calculating energetic contributions directly from the structure itself. Often, using a mix of different predictors in a consensus approach yields more reliable estimates, since each method has its own unique strengths and weaknesses [6], [15], [16], [17].

The field of antibody engineering has really taken off thanks to the rise of computational techniques. Now, researchers can whip up a targeted list of candidates that are likely to boost performance, rather than shifting through thousands of variants in the

lab. This approach helps researchers focus on the designs that are most likely to lead to meaningful improvements. This is especially crucial when dealing with rapidly evolving pathogens like SARS-CoV-2, where quickly optimizing neutralizing antibodies can make all the difference in developing effective therapies. In these situations, computational pre-screening becomes an invaluable tool.

1.3 Workflow-based platforms for structural bioinformatics and the emergence of Scipion-Chem

Virtual drug screening (VDS) enables a rapid evaluation of large numbers of molecular variants before experimental testing. By combining structural data with computational scoring functions, VDS can estimate how mutations affect stability, affinity, or specificity, reducing the need for extensive laboratory assays. In antibody engineering, this approach is especially valuable because it allows efficient exploration of the mutational landscape of paratope residues and the prioritization of candidates with the greatest potential to improve antigen binding.

The landscape of tools for computational antibody design is still fragmented, despite its rapid evolution. Each software package has its own distinct data formats, dependencies, and methods of operation, whether that involves structural analysis, machine learning, or energy calculations. The end results depend not only on the algorithm itself but also on how it is set up, how files are managed, and the decisions made by the user. This diversity can make reproducing results quite challenging. Consequently, many computational pipelines feel like a black box, making it difficult to determine which software version was used, how intermediate files were handled, or how each result was derived [18].

In response to this fragmentation, workflow engines emerged to create unified frameworks that help organize diverse tools into coherent computational pipelines. One of the most recognized engines in structural biology is Scipion, which was originally designed to manage image-processing workflows for cryo-electron microscopy. With the addition of a graphical user interface and a modular plugin system, Scipion enables users with different levels of computational skills to run numerous interconnected protocols in a way that is both repeatable and traceable. It guarantees complete reproducibility and

auditability by automatically logging every action taken within Scipion, including the software versions, parameters used, and all intermediate outputs.

To expand the platform's reach into chemical and structural bioinformatics, Scipion-Chem was developed on this existing infrastructure. It utilizes the same workflow system as cryo-EM, seamlessly integrating computational tools that are essential for virtual screening, protein engineering, and drug design. This integration brings a host of advantages. Researchers can handle structural imports, mutagenesis, ΔΔG predictions, ranking, and visualization all from a single interface, eliminating the need to switch between different programs. Every time a protocol is executed (whether it involves *FoldX*, *SAAMBE-3D*, docking modules, or scoring techniques) it is automatically logged. The platform is also designed to be easily expanded with new plugins, allowing for the incorporation of the latest algorithms as they emerge. By standardizing data formats and parameter structures, the graphical environment significantly reduces the chances of user errors.

Scipion-Chem is the ideal environment for optimizing antibodies. It ensures complete traceability, streamlines the assessment of extensive mutational libraries, and allows for direct comparisons of different ΔΔG predictors (all crucial components of high-quality computational engineering workflows) [11].

Scipion-Chem is much more than just a collection of computational tools. It represents a workflow-driven approach that ensures the design of computational antibodies meets the same rigorous scientific standards as traditional experimental methods.

In addition to Scipion-Chem, several other platforms have been developed to support computational studies in structural bioinformatics. General workflow managers like Galaxy and Nextflow allow users to connect different tools in a reproducible way. However, they are not specifically designed for structure-based analyses. Other platforms focused on molecular modeling, such as Rosetta and HADDOCK, and offer strong methods for protein design, docking, and energy evaluation. However, they often focus on specific tasks and may need custom scripting or manual coordination between tools. In this setting, Scipion-Chem stands out by combining workflow automation with a

design focused on structure. This makes it easier to systematically compare different prediction methods within one reproducible framework.

1.4 Antibody optimization in the context of SARS-CoV-2 and the receptor-binding domain (RBD)

The swift advancement and fine-tuning of therapeutic antibodies have been essential, especially highlighted by the worldwide spread of SARS-CoV-2. The receptor-binding domain of the spike protein, which helps the virus enter our cells, has a direct interaction with the human ACE2 receptor. Neutralizing antibodies, whether produced during a natural infection or after vaccination, mainly focus on this specific area. To block the virus from attaching, a variety of therapeutic monoclonal antibodies have been developed to target the RBD.

On the other hand, SARS-CoV-2 changes remarkably quickly. Concerning variations like Alpha, Delta, Omicron, and their numerous sub lineages have accumulated many RBD mutations. A number of these modifications result in immune evasion with decreased neutralization potency and significantly altering antibody binding. Individual substitutions can alter the RBD surface, break hydrogen bond networks, or produce steric barriers that impair previously successful interactions with antibodies, according to structural and thermodynamic studies [19], [20].

SARS-CoV-2 serves as an excellent model for designing computational antibodies due to its rapid and ongoing evolution. We can pinpoint variants that are likely to maintain strong binding even as the virus diversifies by using predictive methods that assess the energetic effects of mutations. These tools can also help in the smart re-engineering of existing antibodies by identifying stabilizing mutations that lessen the impact of antigenic drift.

In this context, Scipion-Chem plays a vital role by offering an effective way to import the antibody (RBD complex, pinpoint energetic hotspots at the interface, and evaluate the thermodynamic effects of mutations through various ΔΔG predictors). In addition, users can easily select the variants that have the highest potential for improved stability or affinity by directly comparing the results on the platform.

The RBD serves as a robust benchmark for evaluating computational pipelines, thanks to its well-defined structure and the extensive research on its interactions with neutralizing antibodies. Insights gained from studying SARS-CoV-2 can be leveraged for any therapeutic antibody that requires affinity maturation, as well as for other rapidly changing viruses.

1.5 Objectives of the project

The main objective of this project is to computationally optimize an antibody targeting the receptor-binding domain (RBD) of SARS-CoV-2 by using Scipion-Chem as a workflow-based platform for structure-driven antibody engineering.

To achieve this general goal, we defined the following subobjectives:

- **Import and prepare the antibody–antigen atomic structure** within Scipion-Chem, ensuring correct organization of the structural model and accurate identification of protein chains.

- **Identify the structural regions of interest (ROIs)** corresponding to antibody paratope residues involved in the interaction with the antigen, using a distance-based interface detection protocol.

- **Evaluate the energetic impact of single-point mutations** at the antibody–antigen interface by applying ΔΔG prediction protocols available in Scipion-Chem.

- **Compare complementary prediction strategies**, combining empirical force-field calculations (*FoldX*) and machine-learning–based models (*SAAMBE-3D*) to assess mutation-induced stability changes.

- **Integrate and rank mutation predictions** through a consensus-based rank fusion strategy, enabling the identification of mutations consistently predicted as stabilizing across different algorithms.

- **Assess the suitability of Scipion-Chem as a reproducible and traceable platform** for computational antibody optimization, highlighting the advantages of workflow-based approaches for systematic mutational analysis.

Although the case study focuses on an antibody against SARS-CoV-2, the proposed methodology is designed to be transferable to other antibody–antigen systems.

2 MATERIAL AND METHODS

2.1 Material

2.1.1 Software

The Scipion-Chem platform (version 3.8.3) was utilized for all analyses. This platform allows the integration of external programs into traceable workflows, enhancing the Scipion workflow engine for tasks like virtual drug screening and protein engineering. Scipion-Chem simplifies the entire process, making it easy to repeat and reuse by meticulously tracking all parameters and intermediate files within a project database.

In this work, we employed several Scipion-Chem plugins and protocols. The analysis made use of the *Import Atomic Structure* protocol for structural data handling, the *Define Structural ROIs* protocol for interface residue identification, the *SAAMBE-3D* protocol from the *scipion-chem-alexov* plugin for ΔΔG prediction, the *FoldX* protocol from the *scipion-chem-foldxsuite* plugin for empirical energy evaluation, and the rank fusion protocol included in Scipion-Chem for consensus-based ranking of mutations.

2.1.2 Structural data

The structural data used in this study correspond to an antibody–antigen complex obtained from the Protein Data Bank (PDB). The structure represents the interaction between an antibody (or nanobody) and its target antigen, the receptor-binding domain (RBD) of the SARS-CoV-2 spike protein. The atomic model includes full three-dimensional coordinates for all residues involved in the interaction, providing a detailed description of the binding interface.

The PDB structure contains experimentally determined atomic coordinates, chain identifiers, residue numbering, and spatial arrangements that define the geometry of the antibody–antigen interface. These data allow the identification of interface residues and the evaluation of local physicochemical environments relevant for mutational analysis. The availability of a high-resolution structural model is essential for structure-based approaches, as ΔΔG prediction methods rely on accurate representations of residue contacts, distances, and interaction networks.

All analyses performed in this work were based on this single experimentally derived structure, which served as the structural reference for the study. Using a consistent atomic model ensures that energetic predictions and comparative analyses reflect differences arising from mutations rather than from variations in structural input, which is a standard practice in computational studies of protein–protein interactions.

2.2 Methodology

This project's computational workflow relies on the step-by-step execution of various Scipion-Chem protocols. It involves tasks like importing structures, defining interface residues, calculating ΔΔG values, and merging the resulting scores. The workflow transitioned fluently from structural preparation to the final energetic analysis, thanks to the fact that each step was carried out in a controlled and traceable way. All the phases of this workflow can be seen in Figure 3.

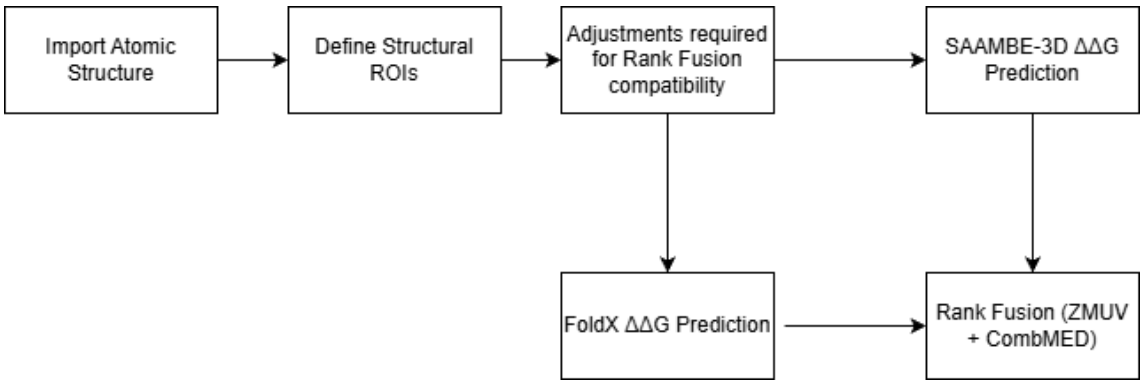

Figure 3 Computational workflow used for antibody optimization in Scipion-Chem.

2.2.1. Installation and preparation of the Scipion-Chem environment

The Scipion platform (version 3.8.3) was installed on a Linux system and configured to support all the protocols needed for this investigation before starting the computational workflow. After completing the base installation, the necessary plugins were added through the Scipion plugin manager. To enable chemical and structural bioinformatics workflows, we installed the scipion-chem plugin, which included

integrating the *SAAMBE-3D* ΔΔG predictor with scipion-chem-alexov, as well as incorporating the *FoldX* modules into Scipion using scipion-chem-foldxsuite.

Every plugin was carefully checked in the graphical user interface to make sure it was registered correctly and functioning properly. Once this verification was complete, the Scipion-Chem environment was fully equipped with all the necessary protocols for the workflow, covering everything from ΔΔG prediction and rank fusion to structure import and Region of Interest (ROI) definition. Thanks to Scipion's ability to manage internal dependencies, ensure protocol compatibility, and organize data automatically, there was no need for any extra scripting or handling of external files.

2.2.2 Import of the atomic structure

The first step in the workflow was to use the Import Atomic Structure protocol to load the antibody–antigen complex into the Scipion-Chem environment. This protocol is essential because it allows the software to manage the molecule internally in the later stages of the analysis. It is done by reading the coordinates from a PDB file and converting them into a Scipion atomic structure object.

Scipion took care of identifying the protein chains in the file and neatly organized them in the workspace as it went along. The import protocol automatically managed file validation and formatting, which meant there was no need for any manual editing or preprocessing of the PDB file. After that, the workflow moved on to defining structural regions of interest and running the ΔΔG prediction procedures, all using the imported structure.

2.2.3 Definition of structural ROIs

The Define Structural ROIs procedure in Scipion-Chem was used to identify the regions directly involved in the antibody–antigen interaction after the structure had been imported. By examining the spatial proximity between the two chains in the complex, this protocol enables the automatic detection of interface residues.

In this step, the antigen chain was identified as the key player, while the antibody or nanobody chain served as the reference point. Scipion took on the task of measuring

the distances between every atom in both chains and selected the residues whose atoms fell within a specific range. By employing this distance-based approach, the software successfully pinpointed the set of residues that were most likely to play a role in binding.

In this step, we identified the antigen chain as the interacting partner, while the antibody or nanobody chain served as the reference. Scipion analyzed the Euclidean distance between every atom in both chains and selected the residues whose atoms fell within a specific range. By employing this distance-based approach, the software successfully pinpointed the set of residues that were most likely to play a role in binding.

A list of interface residues that indicate the structural area where mutations are most likely to impact the interaction energy was the outcome of this analysis. All subsequent ΔΔG predictions made using the *SAAMBE-3D* and *FoldX* protocols were based on this list.

2.2.4 Prediction of mutation-induced ΔΔG using SAAMBE-3D

We used the *SAAMBE-3D* protocol, which is part of the scipion-chem-alexov plugin to carry out the initial evaluation of how mutations affect energy levels after identifying the interface residues. This protocol predicts the change in binding free energy (ΔΔG) associated with specific point mutations at the antibody-antigen interface.

Scipion started automatically generating candidate mutations for evaluation, using the imported structure along with a list of interface residues as its starting point. To figure out if a substitution would either stabilize or destabilize the interaction, *SAAMBE-3D* relied on its machine-learning model based on gradient boosting decision trees, trained on experimentally validated mutation datasets. The protocol then created a table with the expected ΔΔG values for all the mutations assessed. These results were saved in Scipion and later combined with the *FoldX* predictions during the consensus ranking phase of the rank fusion protocol. Additionally, we made a modification in the default output table of the *SAAMBE-3D* protocol to include an extra column containing the mutation identifier. This ensured that each ΔΔG value could be directly associated with its corresponding mutation during the downstream analysis.

Internally, *SAAMBE-3D* generates its predictions by extracting a wide set of structural, energetic and physicochemical descriptors from the three-dimensional geometry of the complex. For each mutation, the algorithm characterizes the local environment surrounding the substituted residue and quantifies variables such as solvent-accessible surface area, residue depth, hydrogen-bond networks, electrostatic contributions and pairwise interaction potentials. These features are incorporated into a machine-learning regression model that has been trained on experimentally measured ΔΔG values from mutational datasets. Learning the nonlinear relationships between local structural changes and their energetic consequences allows *SAAMBE-3D* to estimate whether a mutation weakens, preserves or reinforces the antibody–antigen interface. This integrative, feature-based approach allows the method to capture effects that are often difficult to model with purely physics-driven techniques, which makes *SAAMBE-3D* a useful and complementary tool for assessing mutational stability within protein–protein complexes.

2.2.5 Prediction of mutation-induced ΔΔG using FoldX

We also used the scipion-chem-foldxsuite plugin with the *FoldX* ΔΔG protocol to perform a second energetic evaluation alongside the *SAAMBE-3D* analysis. *FoldX* employs an empirical force-field model to assess how single-point mutations influence the stability of protein–protein complexes. It was provided with the same set of interface residues and the same imported atomic structure as those used in the *SAAMBE-3D* protocol. The protocol generated the relevant mutations and calculated the ΔΔG value for each residue within the area of interest. Positive values indicate a destabilizing effect, while negative values suggest that the mutation enhances the stability of the interaction.

Internally, *FoldX* predicts the ΔΔG values using a physics-based empirical force method that estimates the folding free energy (ΔG) as a weighted sum of different energy terms, as Equation 1 shows.

$$\Delta G = Wvdw \cdot \Delta GvdW + WsolvH \cdot \Delta GsolvH + WsolvP \cdot \Delta GsolvP + \Delta Gwb + \Delta Ghbond + \Delta Gel + \Delta GKon + Wmc \cdot T \cdot \Delta Smc + Wsc \cdot T \cdot \Delta Ssc \quad (1)$$

These terms refer to the main contributions to protein stability: ΔGvdw covers van der Waals interactions (with a weight Wvdw usually set to 0.33 because of differences in the reference states), ΔGsolvH and ΔGsolvP handle solvation penalties or gains for hydrophobic and polar groups when the protein folds. WsolvH an WsolvP are their weights. ΔGwb adds the extra stabilization from water bridges, ΔGhbond captures the balance between intra-protein hydrogen bonds and those with solvent, ΔGel includes electrostatics like charged groups and helix dipoles, ΔGKon deals with electrostatic effects on binding in complexes, and the last two terms (ΔSmc and ΔSsc) represent the entropic costs of restricting backbone and side-chain flexibility (multiplied by temperature T and their respective weights Wmc and Wsc) [22].

When FoldX analyses a mutation, it builds a model of the mutant by introducing the new residue and then optimizes the local structure around it. It calculates ΔG for both the wild-type and the mutant, and the difference between them gives ΔΔG. Breaking down the energy into these individual terms it can be seen what is driving the change, which gives a nice insight into how the mutation affects the protein-protein interface.

Since FoldX is based on this fixed, physics-inspired model rather than being trained on large datasets, it provides a different angle compared to machine learning tools like *SAAMBE-3D*, and combining both makes the predictions more robust.

The *FoldX* protocol produces a structured table that outlines the ΔΔG predictions for each mutation we are evaluating. As in the previous case, we adapted the *FoldX* ΔΔG protocol so that the resulting table included a column with the mutation name. This modification allowed a clear correspondence between each mutation and its predicted ΔΔG value.

2.2.6 Rank Fusion: consensus scoring of ΔΔG predictions

The Rank Fusion protocol was used to integrate the ΔΔG predictions generated independently by *SAAMBE-3D* and *FoldX*. This approach allows Scipion-Chem to merge the heterogeneous scoring outputs produced by both methods into a single ranking that offers a consensus view of which mutations are most likely to stabilize or destabilize the antibody–antigen complex. The protocol begins by applying Z-score normalization

(ZMUV) to the ΔΔG values from each predictor so that the scores become comparable despite the different scales and statistical behaviours of the underlying algorithms, and its mathematical calculation is shown in Equation 2. This normalization scales the scores so that their mean becomes zero and their variance 1.

$$ZMUVNorm(s) = \frac{s - s_{mean}}{s_{std}} \quad (2)$$

where s represents the original score assigned to a mutation by a given predictor, s_{mean}corresponds to the mean value of the score distribution, and s_{std}denotes its standard deviation [23].

After normalization, the rankings are combined using the CombMED rule, a fusion strategy that prioritizes mutations consistently ranked among the best candidates by multiple predictors while reducing the impact of outliers. In this approach, each predictor assigns a relative ranking position $r_k(m)$ to every mutation m according to its predicted energetic impact. It can be seen in Equation 3. The CombMED strategy then computes the median of these ranking positions to obtain a consensus score, defined as:

$$\text{CombMED}(m) = \text{median}\big(r_1(m), r_2(m), \dots, r_K(m)\big) \quad (3)$$

where Krepresents the number of predictors considered. This rank-based fusion method emphasizes agreement between independent predictors and avoids direct comparison of energetic scales that may differ across models by relying on the median of rankings rather than on absolute ΔΔG values. As a result, mutations that are consistently ranked among the top candidates are prioritized, while the influence of method-specific outliers is reduced.

The final ranking, therefore, highlights mutations for which both *SAAMBE-3D* and *FoldX* show strong agreement, increasing confidence in the stability-enhancing candidates identified through the workflow. The table produced by the Rank Fusion protocol represents the final output of the energetic analysis and serves as the basis for selecting the top-scoring mutations for downstream interpretation.

2.2.7 Adjustments required for Rank Fusion compatibility

During workflow development, a problem occurred when merging predictions from *SAAMBE-3D* and *FoldX* using the Rank Fusion protocol. Although both tools successfully computed ΔΔG values for all target mutations, their default output tables did not include the mutation identifier next to each prediction. Instead of listing the mutations by name, such as A40F or A59W, the tables indexed the results numerically. This prevented Rank Fusion from matching the ΔΔG values from the two predictors to the correct mutation during the fusion step.

To solve this issue, both protocols were adjusted so that their output tables included an additional column containing the mutation name. This modification created a direct link between each ΔΔG value and its corresponding mutation, allowing Rank Fusion to correctly align, normalize and combine the datasets. With this change in place, the consensus ranking was generated without errors, and the final list of stabilizing mutations became more accurate, reproducible and easier to interpret.

3 RESULTS

3.1 Structural data import (Import Atomic Structure)

Bringing the antibody-antigen complex into the Scipion-Chem environment kicked off our workflow. We started by loading the PDB file that holds the antibody and the SARS-CoV-2 RBD, transforming it into an atomic structure object that Scipion could work with using the Import Atomic Structure protocol. Scipion automatically identified the protein chains in the file and organized them within the project to make things easier for the following steps.

The import protocol handles all the file validation and formatting, so there was no need for any manual adjustments to the structure. Thanks to the successful importation of the complex, we were able to visually examine the chains, confirming that the structure could indeed be used as a basis for defining ROI and predicting ΔΔG. Figure 4 illustrates the imported structure as seen in the Scipion viewer, showing the recognized chains of the complex.

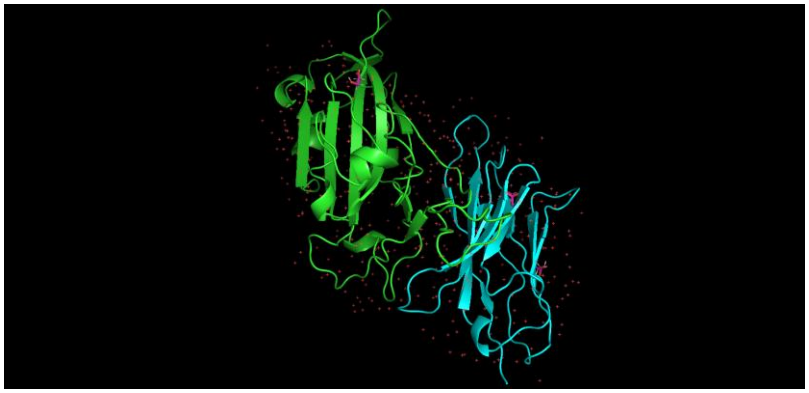

Figure 4 Imported antibody–antigen complex.

3.2 Definition of ROIs

After importing the atomic structure, the next step was to identify the residues involved in the antibody-antigen interaction. For this, we utilized the Define Structural ROIs protocol available in Scipion-Chem. This handy tool automatically detects interface residues by examining the spatial relationships between the selected chains, all based on a set distance threshold.

In this case, the antigen chain was identified as the interaction partner, while the antibody chain was selected as the reference chain. To compile the list of interface residues that form the structural region of interest, Scipion calculated the minimum distances between all atoms from both chains and picked the residues that fell within the specified cutoff. During the subsequent ΔΔG prediction protocols, these residues became the targets for mutation. The results from the Define Structural ROIs protocol, which includes the list of interface residues identified by Scipion-Chem, are shown in Figure 5.

Region	Residues ([aa][pos])
1–2	Q1, V2
30–33	R30, T31, C32, N33
39–46	Q39, A40, P41, G42, K43, E44, L45, E46
50–52	S50, I51, S52
54–65	G54, S55, T56, N57, Y58, A59, G60, S61, V62, K63, G64, R65
87	E87
91	V91
93	Y93
98–100	I98, S99, R100
103–114	S103, L104, W105, C106, E107, E108, Y109, W110, G111, Q112, G113, T114

Figure 5 Residues identified at the antibody–antigen interface using the *Define Structural ROIs* protocol in Scipion-Chem.

3.3 Generation of mutant models

Scipion-Chem created a complete set of single-point mutations that will be evaluated in the upcoming ΔΔG prediction steps, following the identification of interface residues through the ROI definition protocol. The workflow produced every possible amino acid substitution for each residue of interest.

The output of this step was a well-organized list detailing each mutation created from the selected residues. Each entry in the list clearly indicates the location of the residue, the original amino acid, and the proposed substitution. This mutational library was then used as input for both *SAAMBE-3D* and *FoldX*, ensuring that all prediction tools evaluated the same set of variants consistently. Figure 6 displays an excerpt of the list of mutations generated in Scipion, which includes all the single-point variants derived from the residues located at the antibody–antigen interface.

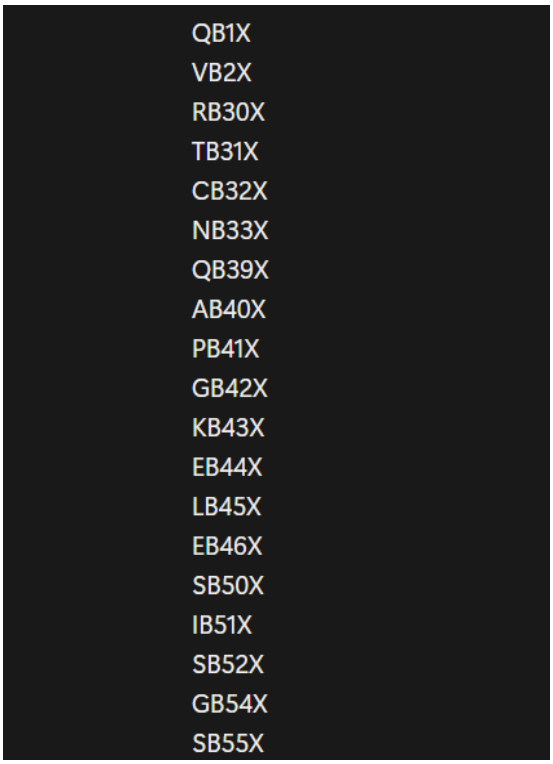

Figure 6 Mutational library generated from the ROI residues.

3.4 ΔΔG predictions obtained with SAAMBE-3D

We used the *SAAMBE-3D* protocol integrated into Scipion-Chem to evaluate how single-point mutations affect energy levels after specifying the structural regions of interest. This method allows for an initial evaluation of how each residue change might influence the stability of the antibody-antigen interaction by generating predictions specific to each mutation regarding changes in binding free energy (ΔΔG).

Additionally, we created a table displaying the predicted ΔΔG values for each mutation we evaluated using the protocol. Each row highlights a unique mutation at one of the interface residues, along with the numerical value assigned by the *SAAMBE-3D* model. If the ΔΔG value is negative, it suggests a stabilizing effect on the interaction, while positive values point to destabilizing mutations. We successfully generated results for every mutation listed in the ROI.

The predicted ΔΔG values for the examined mutations are shown in an excerpt of the SAAMBE-3D output table in Figure 7. The first column shows the mutation name, the second one the ΔΔG values calculated with the protocol, and the last one the values normalized.

	mutation	ddg	zscore
1	AB40F	-1.85	-3.49
2	AB40D	-1.7	-3.26
3	AB40K	-1.61	-3.13
4	AB40Q	-1.58	-3.08
5	AB40L	-1.36	-2.76
6	AB40N	-1.35	-2.74
7	AB40E	-1.34	-2.73
8	AB40C	-1.32	-2.7
9	AB40I	-1.3	-2.67
10	AB40R	-1.28	-2.64
11	AB59F	-1.26	-2.61
12	AB40M	-1.21	-2.53
13	AB40H	-1.11	-2.38
14	AB40Y	-1.02	-2.25

Figure 7 SAAMBE-3D ΔΔG predictions.

3.5 ΔΔG predictions obtained with FoldX

Using the *FoldX* ΔΔG protocol, which is integrated through the scipion-chem-foldxsuite plugin, we conducted a second independent evaluation of how mutations affect the antibody–antigen interaction. This protocol utilized the empirical force-field model from *FoldX* to estimate the changes in binding free energy for each mutation within the designated area of interest.

The protocol generated a result table listing all evaluated mutations and their corresponding ΔΔG values. Just like the output from *SAAMBE-3D*, positive values point to destabilizing mutations, while negative values indicate stabilizing effects on the interaction. This table provides the second independent ranking required for the upcoming fusion analysis and enables a direct comparison of *FoldX* predictions across all candidate mutations.

An excerpt portion of the *FoldX* results table produced by Scipion-Chem, which lists the ΔΔG values connected to each mutation, is shown in Figure 8.

	mutation	ddg	zscore
1	YB109E	-2.35	-3.06
2	SB61R	-2.06	-2.73
3	YB58T	-1.68	-2.28
4	SB55W	-1.54	-2.13
5	EB44H	-1.29	-1.84
6	SB55Y	-1.28	-1.82
7	YB58I	-1.24	-1.78
8	SB52E	-1.17	-1.69
9	NB33L	-1.14	-1.66
10	AB59W	-1.12	-1.64
11	NB33M	-1.06	-1.56
12	GB42W	-1.06	-1.56
13	EB46Y	-1.01	-1.51
14	EB44R	-0.98	-1.47

Figure 8 ΔΔG values predicted by the FoldX protocol.

3.6 Consensus ranking with Rank Fusion

Using the Rank Fusion protocol that is part of Scipion-Chem, we conducted a consensus analysis to merge the predictions from *SAAMBE-3D* and *FoldX*. This protocol employs normalization and aggregation techniques to bring together multiple ranked lists, ultimately producing a single ranking that reflects the consensus among the different ΔΔG predictors.

The protocol started by applying a zero-mean, unit-variance transformation (ZMUV) to standardize the scores after it received the output tables from *SAAMBE-3D* and *FoldX*. This step ensured that both predictors were now aligned on a similar numerical scale. Once the normalization was complete, the individual rankings for each mutation were merged into a single consensus score using the CombMED aggregation rule. Both methods suggest that mutations with lower consensus values are likely to have a more significant stabilizing effect.

The protocol produced a ranked list that featured every mutation evaluated, along with its final position in the combined ranking and consensus score. Figure 9 shows the table generated by the rank fusion protocol, which contains the total scores and the resulting order of mutations.

	mutation	RanxScore
1	AB40F	-1.9
2	AB40D	-1.79
3	AB40K	-1.72
4	AB40Q	-1.7
5	AB59W	-1.67
6	AB40N	-1.57
7	AB40E	-1.54
8	AB40L	-1.53
9	AB40C	-1.5
10	AB40R	-1.48
11	AB40I	-1.47
12	AB40M	-1.42
13	AB59F	-1.41
14	AB40H	-1.31

Figure 9 Consensus ranking obtained with the Rank Fusion protocol.

3.7 Global analysis of mutational effects

We evaluated a total of 940 single-point mutations for each predictor, corresponding to the complete set of variants generated from the residues selected in the region-of-interest analysis. For each mutation, both *SAAMBE-3D* and *FoldX* provided an estimated change in binding free energy (ΔΔG), together with a normalized score expressed as a z-score. This comprehensive dataset allowed a global analysis of the mutational landscape at the antibody–antigen interface beyond the inspection of individual top-ranked variants.

The overall distribution of predicted ΔΔG values indicates that most mutations produce either neutral or destabilizing effects on the interaction, while only a limited subset is predicted to be stabilizing. This behavior is consistent with the general expectation for protein–protein interfaces, where only a small fraction of possible substitutions improves binding affinity, and the majority either disrupt favorable interactions or have a negative energetic impact. As a result, prioritization strategies are required to distinguish meaningful candidates from background noise.

Comparison between *SAAMBE-3D* and *FoldX* predictions reveals partial agreement across the full set of mutations. While individual ΔΔG values may differ between methods due to their distinct modelling approaches, both predictors tend to converge on a small group of mutations with consistently favorable energetic profiles. This observation supports the use of a consensus-based ranking strategy, as agreement between independent predictors increases confidence in the relevance of the selected candidates.

The addition of z-score normalization further facilitated comparison between predictors by placing their outputs on a common scale. Normalization reduces the influence of differences in score distributions and ensures that rankings are driven by relative performance rather than absolute numerical values. Consequently, the subsequent rank fusion step emphasizes mutations that perform well across both predictors instead of those supported by a single method.

This global analysis emphasizes the highly selective nature of stabilizing mutations at the antibody–antigen interface and justifies the need for ranking and fusion

strategies when exploring large mutational datasets. Rather than focusing only on absolute ΔΔG values, the workflow captures general trends across hundreds of variants, enabling a robust prioritization of mutations for further analysis.

4 DISCUSSION

The analysis combining *SAAMBE-3D* and *FoldX* identified a small subset of mutations within the antibody paratope that consistently show stabilizing effects on the interaction with the SARS-CoV-2 RBD. The top-ranked variants (AB40F, AB40D, AB40K, AB40Q and AB59W) correspond to substitutions at positions 40 and 59 of chain B, both of which belong to interface regions previously detected by the ROI selection protocol.

The four highest-scoring mutations were all located at residue 40, originally an alanine. The substitutions A40F, A40D, A40K and A40Q seem to decrease the binding free energy ($\Delta\Delta G < 0$) leading to their prioritization in the fused ranking. This consistent agreement between predictors suggests that position 40 plays a relevant role within the antibody–antigen interface, and that altering its physicochemical properties may enhance the stability of the complex [10], [17].

In addition to residue 40, the mutation A59W also ranked among the best candidates. Position 59 is part of a different paratope region, indicating that multiple interface segments contain residues with potential for stabilization through single-point substitution. Although *SAAMBE-3D* and *FoldX* rely on different modelling principles, both identified this mutation as energetically favorable, reinforcing the validity of the consensus-based prioritization.

From a structural perspective, the concentration of stabilizing mutations within paratope regions is consistent with the central role of these residues in antigen recognition. The paratope constitutes the antibody surface directly involved in binding and is therefore particularly sensitive to physicochemical changes introduced by point mutations. Even subtle substitutions at the interface can modulate local packing, electrostatic complementarity, or solvent exposure, leading to measurable energetic effects. The identification of multiple stabilizing mutations within distinct paratope segments suggests that the antibody–antigen interface contains several positions amenable to optimization rather than a single dominant hotspot.

These results collectively indicate that positions 40 and 59 concentrate the mutations with the strongest predicted stabilizing effects and therefore represent

promising targets for future experimental evaluation. The fact that two independent ΔΔG predictors converged on these residues increases the confidence in the computational predictions, although experimental validation would still be required to confirm their actual impact on binding affinity.

The use of the strategy based on ROI was important in focusing the mutational analysis on residues most likely to influence binding. By restricting the search space to interface residues identified through distance-based criteria, the workflow avoided the exploration of mutations with limited structural relevance. This targeted approach not only improves computational efficiency but also increases the interpretability of the results, as predicted energetic changes can be directly linked to spatial proximity to antigen. The consistency between ROI selection and the location of the top-ranked mutations further supports the validity of this strategy.

The convergence of *SAAMBE-3D* and *FoldX* predictions on the same set of top-ranked mutations highlights the utility of combining ΔΔG predictors within a consensus-based workflow. While each method relies on different modelling assumptions (machine-learning features in the case of *SAAMBE-3D* and an empirical force-field prediction in *FoldX*) the agreement observed for residues 40 and 59 suggests that the stabilizing trends detected are not artefacts of a single algorithm. This reinforces the robustness of the rank fusion strategy, as mutations prioritized by both approaches are more likely to represent genuine energetic improvements rather than method-specific biases.

Similar consensus-based strategies have been increasingly adopted in computational protein engineering to mitigate the limitations of individual predictors. Single-method approaches may be biased by their underlying assumptions or training data, whereas combining complementary models can improve robustness and confidence in the selected candidates. In this context, the agreement observed between a machine-learning–based predictor and a physics-based force-field model aligns with current best practices in computational antibody optimization, where consensus scoring is often preferred over reliance on a single algorithm.

Despite these results, there are several limitations in the present study. The predictions are based on a single static structure of the antibody–antigen complex and do

not account for conformational flexibility or long-timescale dynamics that may influence binding. Furthermore, ΔΔG estimates obtained from computational models provide relative trends rather than absolute affinity changes, and their accuracy depends on the quality of the underlying structural data and training datasets. The stabilizing effect predicted for the selected mutations should be interpreted as hypotheses that require experimental validation for this reason.

Despite these limitations, the results obtained provide valuable guidance for rational antibody design. Computational prioritization of stabilizing mutations can significantly reduce the experimental burden by narrowing down the number of variants that need to be tested in vitro. In particular, mutations consistently predicted as stabilizing across multiple models represent strong candidates for subsequent experimental characterization, such as binding affinity measurements or stability assays. As such, the present workflow can be seen as an effective filtering step within a broader antibody engineering pipeline.

Finally, although this study focuses on a single antibody–antigen system, the proposed workflow is not system-specific. The modular design of Scipion-Chem allows the same pipeline to be applied to other antibody targets or protein–protein interactions with minimal adaptation. Additional predictors or alternative ranking strategies could also be incorporated in future applications, further enhancing the flexibility of the approach.

Beyond the specific mutations identified, this work demonstrates the applicability of Scipion-Chem as a reproducible workflow-based platform for antibody optimization. The integration of structure preparation, interface definition, ΔΔG prediction, and rank fusion within a single environment simplifies the exploration of mutational landscapes and reduces manual intervention. This modular approach facilitates the comparison of multiple predictors and can be readily extended to other antibody–antigen systems or additional scoring methods.

5 CONCLUSIONS

This project demonstrates the applicability of Scipion-Chem as an integrated and reproducible platform for the computational optimization of antibody–antigen interactions. Combining structure-based analysis, interface residue selection, ΔΔG prediction, and consensus ranking within a single workflow, we established a systematic approach to explore the energetic impact of single-point mutations at the antibody paratope. This integration allowed the entire mutational analysis to be carried out in a transparent and traceable manner, minimizing manual intervention and ensuring consistency across all computational steps.

The use of two complementary ΔΔG predictors, *SAAMBE-3D* and *FoldX*, enabled the evaluation of mutational effects from distinct methodological perspectives. While *SAAMBE-3D* relies on machine-learning models trained on structural and physicochemical descriptors of protein–protein interfaces, FoldX applies an empirical force-field framework to estimate energetic contributions. The agreement observed between both methods for a subset of mutations underscores the value of combining heterogeneous predictors when assessing protein stability and binding affinity. Rather than relying on a single scoring function, the consensus-based strategy implemented through rank fusion increased confidence in the prioritization of candidate mutations by reducing method-specific biases.

Through this consensus approach, a small group of mutations affecting residues 40 and 59 of chain B was consistently identified as potentially stabilizing the interaction with the SARS-CoV-2 receptor-binding domain. These residues were previously detected as part of the antibody–antigen interface during the region-of-interest selection step, supporting the relevance of the structural filtering strategy employed. Although the present study does not provide experimental validation, the convergence of predictions across independent models suggests that these positions represent promising targets for further investigation in vitro or in vivo.

Beyond the identification of specific mutations, one of the main contributions of this work lies in the establishment of a generalizable computational workflow for antibody optimization. The modular design of Scipion-Chem allows each stage of the

analysis (structure import, interface definition, mutation generation, energy evaluation, and ranking) to be adapted or extended with minimal changes. As a result, the same pipeline could be applied to other antibody–antigen systems or expanded to include additional predictors, scoring strategies, or structural models, depending on the needs of future studies.

In addition to the specific mutations identified, the present study highlights several strengths of the proposed computational approach. One of the main advantages is the integration of all analysis steps within a single workflow-based platform, which ensures reproducibility and traceability throughout the entire process. By combining structure import, interface definition, mutational analysis, energetic prediction, and consensus ranking in a unified environment, the workflow minimizes manual intervention and reduces the risk of inconsistencies between intermediate results. Furthermore, the use of complementary predictors and rank-based aggregation increases the robustness of prioritization, as candidate mutations are selected based on consistent trends rather than isolated predictions from a single method.

Despite its strengths, this work also presents several limitations inherent to structure-based computational approaches. The analysis was based on a single static structural model of the antibody–antigen complex and therefore does not account for conformational flexibility, induced fit effects, or long-timescale dynamics that may influence binding. In addition, ΔΔG predictions obtained from computational models provide relative energetic trends rather than absolute affinity measurements. Consequently, the stabilizing effects predicted for the selected mutations should be interpreted as hypotheses that require experimental validation through biochemical or biophysical assays.

Nevertheless, within these limitations, the proposed workflow represents an effective strategy for analysing large mutational spaces and guiding experimental efforts toward the most promising candidates. The approach can significantly reduce the experimental burden associated with antibody engineering and accelerate early-stage optimization by prioritizing mutations with consistent stabilizing predictions across multiple models.

From a broader perspective, the methodology developed in this project can serve as a flexible framework for future computational antibody optimization studies. The same workflow could be readily applied to other antibody–antigen systems or adapted to include additional ΔΔG predictors and alternative ranking strategies. Future work could also extend the analysis to account for conformational flexibility, for example through molecular dynamics simulations, or explore combinations of mutations to assess potential cooperative effects. In this context, the present study provides a solid starting point for iterative optimization pipelines in which computational predictions guide experimental validation, contributing to more efficient and rational antibody engineering strategies.

In conclusion, this project highlights the potential of workflow-driven computational platforms such as Scipion-Chem to support rational antibody design. The combination of reproducibility, scalability, and consensus-based evaluation provides a robust framework for exploring antibody–antigen interactions and identifying candidate mutations for further development. As computational methods continue to evolve and integrate more accurate predictive models, these workflows are expected to play an increasingly important role in the design and optimization of next-generation therapeutic antibodies.

6 REFERENCES

[1] A. M. Hummer, B. Abanades, and C. M. Deane, «Advances in computational structure-based antibody design», *Curr. Opin. Struct. Biol.*, vol. 74, p. 102379, Jun. 2022, doi: 10.1016/j.sbi.2022.102379.

[2] T. Jain *et al.*, «Biophysical properties of the clinical-stage antibody landscape», *Proc. Natl. Acad. Sci. U. S. A.*, vol. 114, no. 5, pp. 944-949, Jan. 2017, doi: 10.1073/pnas.1616408114.

[3] K. M. McCoy, M. E. Ackerman, and G. Grigoryan, «A comparison of antibody–antigen complex sequence-to-structure prediction methods and their systematic biases», *Protein Sci. Publ. Protein Soc.*, vol. 33, no. 9, p. e5127, Aug. 2024, doi: 10.1002/pro.5127.

[4] M. I. J. Raybould *et al.*, «Five computational developability guidelines for therapeutic antibody profiling», *Proc. Natl. Acad. Sci. U. S. A.*, vol. 116, no. 10, pp. 4025-4030, Mar. 2019, doi: 10.1073/pnas.1810576116.

[5] G. Liu *et al.*, «Antibody complementarity determining region design using high-capacity machine learning», *Bioinformatics*, vol. 36, no. 7, pp. 2126-2133, Apr. 2020, doi: 10.1093/bioinformatics/btz895.

[6] R. Guerois, J. E. Nielsen, and L. Serrano, «Predicting changes in the stability of proteins and protein complexes: a study of more than 1000 mutations», *J. Mol. Biol.*, vol. 320, no. 2, pp. 369-387, Jul. 2002, doi: 10.1016/S0022-2836(02)00442-4.

[7] J. Adolf-Bryfogle *et al.*, «RosettaAntibodyDesign (RAbD): A general framework for computational antibody design», *PLoS Comput. Biol.*, vol. 14, no. 4, p. e1006112, Apr. 2018, doi: 10.1371/journal.pcbi.1006112.

[8] R. Akbar *et al.*, «In silico proof of principle of machine learning-based antibody design at unconstrained scale», *mAbs*, vol. 14, no. 1, p. 2031482, Jan-Dec. 2022 doi: 10.1080/19420862.2022.2031482.

[9] J. A. Ruffolo, L.-S. Chu, S. P. Mahajan, and J. J. Gray, «Fast, accurate antibody structure prediction from deep learning on massive set of natural antibodies», *Nat. Commun.*, vol. 14, p. 2389, Apr. 2023, doi: 10.1038/s41467-023-38063-x.

[10] J.-E. Shin *et al.*, «Protein design and variant prediction using autoregressive generative models», *Nat. Commun.*, vol. 12, p. 2403, Apr. 2021, doi: 10.1038/s41467-021-22732-w.

[11] D. Del Hoyo, M. Salinas, A. Lomas, E. Ulzurrun, N. E. Campillo, and C. O. Sorzano, «Scipion-Chem: An Open Platform for Virtual Drug Screening», *J. Chem. Inf. Model.*, vol. 63, no. 24, pp. 7873-7885, Dec. 2023, doi: 10.1021/acs.jcim.3c01085.

[12] F. Gaudreault, T. Sulea, and C. R. Corbeil, «AI-augmented physics-based docking for antibody-antigen complex prediction», *Bioinformatics*, vol. 41, no. 4, p. btaf129, Mar. 2025, doi: 10.1093/bioinformatics/btaf129.

[13] Y. W. Lim, A. S. Adler, and D. S. Johnson, «Predicting antibody binders and generating synthetic antibodies using deep learning», *mAbs*, vol. 14, n.º 1, p. 2069075, Dec. 2022, doi: 10.1080/19420862.2022.2069075.

[14] A. M. Wollacott *et al.*, «Quantifying the nativeness of antibody sequences using long short-term memory networks», *Protein Eng. Des. Sel.*, vol. 32, no. 7, pp. 347-354, Dec. 2019, doi: 10.1093/protein/gzz031.

[15] E. K. Makowski *et al.*, «Co-optimization of therapeutic antibody affinity and specificity using machine learning models that generalize to novel mutational space», *Nat. Commun.*, vol. 13, Jul. 2022, doi: 10.1038/s41467-022-31457-3.

[16] D. M. Mason *et al.*, «Optimization of therapeutic antibodies by predicting antigen specificity from antibody sequence via deep learning», *Nat. Biomed. Eng.*, vol. 5, pp. 600-612, 2021, Jun. 2021 doi: 10.1038/s41551-021-00699-9.

[17] J. C. F. Ng *et al.*, «Tuning antibody stability and function by rational designs of framework mutations», *mAbs*, vol. 17, no. 1, p. 2532117, Dec. 2025, doi: 10.1080/19420862.2025.2532117.

[18] A. Madushanka, E. Laird, C. Clark, and E. Kraka, «SmartCADD: AI-QM Empowered Drug Discovery Platform with Explainability», *J. Chem. Inf. Model.*, vol. 64, no. 17, pp. 6799-6813, Sep. 2024, doi: 10.1021/acs.jcim.4c00720.

[19] J. Lee, C. Seok, S. Ham, and S. Chong, «Atomic-level thermodynamics analysis of the binding free energy of SARS-CoV-2 neutralizing antibodies», *Proteins*, p. 10.1002/prot.26458, Jan. 2023, doi: 10.1002/prot.26458.

[20] L. Li *et al.*, «Machine learning optimization of candidate antibody yields highly diverse sub-nanomolar affinity antibody libraries», *Nat. Commun.*, vol. 14, Jun. 2023, doi: 10.1101/2022.10.07.502662.

[21] «Welcome to Scipion Chem's documentation! — Scipion Chem 1.0.0 documentation». Accessed: Dec. 28, 2025. [Online]. Available: https://scipion-chem.github.io/docs/index.html

[22] «Products | FoldX». Accessed: Dec. 22, 2025. [Online]. Available: https://foldxsuite.crg.eu/products#foldx

[23] M. E. Renda and U. Straccia, Web metasearch: rank vs. score based rank aggregation methods», in *Proceedings of the 2003 ACM symposium on Applied computing*, Melbourne Florida: ACM, Mar. 2003, pp. 841-846. doi: 10.1145/952532.952698.