Knowledge-based Search and Multi-objective Filters: Proposed Structural Models of GPCR Dimerization

Irina Hashmi
Dept. of Computer Science,
George Mason University,
Fairfax, VA 22030
ihashmi@gmu.edu

Daniel Veltri
School of Systems Biology,
George Mason University,
Manassas, VA 20110
dveltri@gmu.edu

Nadine Kabbani
Dept. of Molecular
Neuroscience,
Krasnow Institute,
George Mason University,
Fairfax, VA 22030
nkabbani@gmu.edu

Amarda Shehu*
Dept. of Computer Science,
Dept. of Bioengineering
School of Systems Biology
George Mason University,
Fairfax, VA 22030
amarda@gmu.edu

ABSTRACT

Many experimental studies point to the ubiquitous role of protein complexation in the cell while lamenting the lack of structural models to permit structure-function studies. This scarcity is due to persisting challenges in protein-protein docking. Methods based on energetic optimization have to handle vast and high-dimensional configuration spaces and inaccurate energy functions only to arrive at the wrong interface. Methods that employ learned models to replace or precede energetic evaluations are limited by the generality of these models. Computational approaches designed to be general often fail to provide realistic models on protein classes of interest in the wet laboratory. One such class are G protein-coupled receptors, which wet-lab studies suggest undergo complexation, possibly affecting drug efficacy. In this paper, we propose a computational protocol to address the unique challenges posed by these receptors. To deal with challenges, such as receptor size and inaccuracy of energy functions, the protocol takes a geometry-driven approach and integrates in the search geometric constraints posed by the environment where the receptors operate. Various filters are designed to handle the computational cost of energetic evaluation, and analysis techniques based on new scoring strategies, including multi-objective analysis, are employed to reduce the sampled ensemble to a few credible structural models. We demonstrate that dimeric models of the Dopamine D2 receptor targeted to treat psychotic disorders reproduce macroscopic knowledge extracted in the wet-laboratory and can be employed to further spur detailed structure-function studies.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Probabilistic algorithms; J.3 [**Life and Medical Sciences**]: Biology and genetics

General Terms

Algorithms

Keywords

G protein-coupled receptor assembly; Dopamine D2 receptor dimerization; protein-protein docking; geometry-driven search; multi-objective analysis; Pareto-based comparison

1. INTRODUCTION

Virtually all biological mechanisms involve molecular assemblies. Protein assemblies, in particular, are central components of cellular organization and function, involved in ion transport and regulation, signal transduction, protein degradation, transcriptional regulation, and more [8]. A growing body of wet-lab studies points to the important and varied role of protein assembly in the cell, all the while lamenting the lack of structural models [21]. Structural models of protein assemblies are often a crucial first step to permit detailed structure-function studies not only to better understand the role of an assembly in the healthy cell [16] but also to understand changes that assembly may introduce to the efficacy of a drug designed to interact with a single protein chain [17]. Yet, obtaining such models in silico, even when limited to protein-protein dimers, remains challenging. Size and cellular environment are all limiting factors for wet-lab techniques. Size is also a challenge for many computational methods that have to essentially explore a large bound configuration space in search of credible bound models.

Current protein-protein docking methods can be categorized as energy-driven, geometry-driven, and hybrid. Energy-driven methods approach protein-protein docking through

*Corresponding Author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
BCB'14, September 20–23, 2014, Newport Beach, CA, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2894-4/14/09 ...$15.00.
http://dx.doi.org/10.1145/2649387.2649391.

stochastic optimization algorithms based on Monte Carlo or evolutionary algorithms. Many such methods now exist [14], driven by the community-wide CAPRI experiment [22]. While great progress has been made, several CAPRI summaries make the case that high-accuracy pairwise docking remains challenging [20]. There is great difficulty, for instance, in locating the native interaction interface or even part of it, with top methods shown to predict only 30-58% of the correct interface in any given target [22]. An energy-based treatment is not guaranteed to drive the optimization process towards the right interface [28].

In response, a second group of methods delay consideration of a detailed energy function, choosing instead to sample rigid-body transformations (representing spatial arrangements) that superimpose geometrically-complementary regions on the molecular surfaces of the units participating in a dimer [14]. These methods are referred to as geometry-driven, but their accuracy is typically lower than that of energy-based methods if no energetic refinement is carried out afterwards on obtained models [26].

Hybrid methods have been proposed to improve accuracy and reduce computational cost. Such methods make use of additional information to determine whether a contact interface in a computed dimer is native-like or not, often providing a score to replace or precede detailed energetic evaluations and possible model refinements. This information can be as simple as an evolutionary conservation score [6] or as complex as a machine learning model learned a priori on known native dimers reviewed in [35]. The information can be used as a filter [13] or combined with a physics-based energy function to create a pseudo-energy function [18]. While research on finding what characterizes interaction interfaces remains active [24], methods that employ learned models are to a great extent limited by the generality of these models.

Computational approaches designed to be general often fail to provide realistic models on protein classes of interest in the wet laboratory. One such class are G protein-coupled receptors (GPCRs), which wet-lab studies indicate undergo assembly [30]. This evidence has prompted revision of traditional models of GPCR structure and function, raising questions on the nature of their interactions with ligands, drug compounds, and downstream signaling molecules [33]. Yet, structural models are not only difficult to obtain in silico, but often what is obtained is clearly inaccurate.

The point on the shortcomings of energy-driven methods is illustrated on the suspected dimerization of a specific GPCR, the Dopamine D2 receptor (D2R), whose dimerization has been suggested in the wet laboratory and possibly plays a role in altering the efficacy of antipsychotic drugs [7, 10]. Figure 1 shows the best predicted model for the D2R-D2R dimer by a thorough list of methods. Juxtaposition of these models shows that there is no consensus among them on even the overall topology of the dimer, which is known by wet-lab studies to have fewer options due to the presence of the membrane and to involve a specific trans-membrane region in the interface [10].

Prompted by the need for realistic structural models of GPCR dimerization and the inability of current methods to obtain such models, we propose here a computational protocol capable of addressing the unique challenges posed by these receptors. To deal with challenges, such as receptor size and inaccuracy of energy functions, the protocol takes a geometry-driven approach and integrates in the search geometric constraints posed by the membrane environment where the receptors operate. Various filters are designed to handle the computational cost of energetic evaluation, and analysis techniques based on new scoring strategies, including multi-objective analysis, are employed to reduce the sampled ensemble to a few credible structural models. We demonstrate that dimeric models of D2R-D2R targeted to treat psychotic disorders reproduce macroscopic knowledge extracted in the wet-laboratory and can be employed to further spur detailed structure-function studies. Due to the shared environmental constraints and structural topology among all five members of the family of dopamine receptors and other GPCRs, the proposed protocol is promising to obtain dimeric models of other GPCRs. Our focus on D2R-D2R in this paper is driven by the availability of experimental data that can be used to validate predicted models.

2. METHODS

GPCRs have high structural similarity with one another and share key motifs for binding other molecules [29]. The polypeptide chain of a GPCR is over four hundred amino acids long and contains segments that operate outside and inside cellular membranes. All GPCRs contain an extracellular region, followed by seven transmembrane helices (TM1-7) which are connected by 3 intracellular and 3 extracellular loops (see Figure 2(a)). The TMs are arranged in a characteristic hollow cylindrical shape that fits inside a cellular membrane. A significant number of wet-lab studies show the involvement of these TMs in binding to form higher-order assemblies [9, 11, 23, 32]. In this paper, we propose a general docking protocol that is capable of taking into account experimental knowledge of putative interaction interfaces, additional geometric constraints due to the fact that GPCRs have to operate in membranes, and energetic constraints to compute thermodynamically-stable dimers of GPCRs.

2.1 Decoy Sampling

The method employed to sample dimeric decoys is probabilistic (as opposed to systematic). It is geometry-driven, as geometrically-complementary regions are first computed on the molecular surfaces of the two units to be docked. The novelty here is that experimental knowledge of putative interaction interfaces is integrated to narrow the focus to surface regions that comply with experimental knowledge. These regions are then sampled, and rigid-body transformations that superimpose them are computed to obtain docked units. Since a GPCR unit is large, a simple axis vector representation is employed to rapidly determine whether the resulting dimeric configuration complies with the additional geometric constraints placed by the membrane. Once this geometric filter passes, the entire all-atom representation of each GPCR unit is employed, and the all-atom dimeric configuration is evaluated according to its energy. Once the configuration passes the energetic filter, as well, it is added to a growing decoy ensemble. This protocol is summarized in pseudocode in algorithm 1. We now relate each of the main steps and then conclude with the various analyses of the decoy ensemble to select credible dimeric models.

2.1.1 Sampling-based Rigid-body Docking

Let us consider that two units, A and B, are to be docked, and their structures (cartesian coordinates of atoms) are available in C_A and C_B. One can consider unit A as the

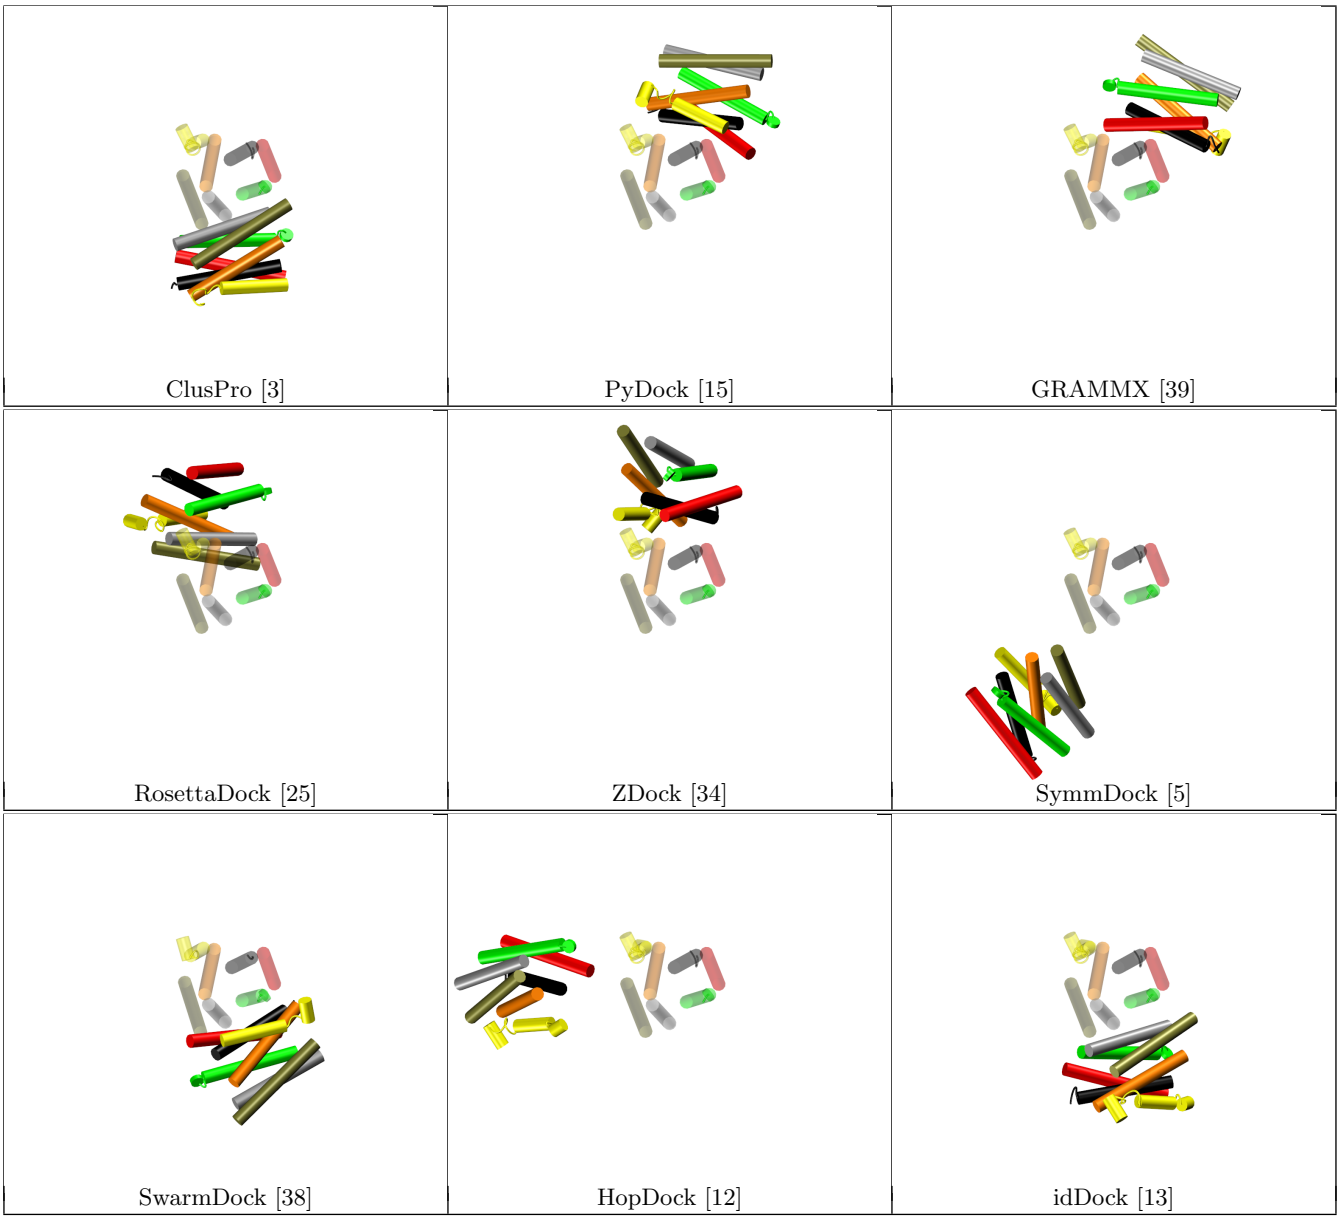

Figure 1: Several top methods, available as web servers, in pairwise docking (including our recent HopdDock [12] and idDock [13] methods) have been applied to obtain best models for a D2R-D2R assembly indicated to form in the living cell by wet-lab studies [7, 10]. The transmembrane regions (TMs) are color coded in each of the two chains (TM1 red, TM2 dark grey, TM3 orange, TM4 yellow, TM5 tan, TM6 light grey, TM7 green.) To show that there is no consensus among these models, all models are superimposed via least Root-Mean-Squared-Deviation (lRMSD) on unit/chain A, which is used as the base/reference unit. This unit is drawn in transparent and appears at the same position and orientation on all the subfigures. Due to the differences among the models, unit B, drawn in opaque, occupies various placements in space, clearly showing the lack of consensus among the models. Not only does the interaction interface not seem to involve TM 4, in contrast to wet-lab evidence [10], but in some models the actual contact interface is between the extra-cellular regions not drawn here in the interest of visibility.

Algo. 1 GPCR-GPCR Docking

1: **Input** :
3D Cartesian coordinates of base unit C_A and moving unit C_B
List of residues in desired interaction interface
Maximum relative orientation angle Θ
Maximum LJ potential value $E_{\max}$
Target number of models N
2: **Output**: Ensemble of dimeric configurations $\Omega = \{C_{AB}\}$
3: **Preprocessing**:
Generate Connolly Surface representation MS_A for C_A and MS_B for C_B
Sample Critical Point representation CP_A and CP_B from regions of MS_A and MS_B containing desired interface
Generate Triangular representation Δ_A from CP_A and Δ_B from CP_B
Define the axis vector v_A for C_A and v_B for C_B
4: **while** $|\Omega| \leq N$ **do**
5: Sample $\delta_A \in \Delta_A$ and $\delta_B \in \Delta_B$ ▷ δ_A and δ_B are geometrically-complementary triangles
6: Compute $T \in \mathrm{SE}(3)$ that superimposes δ_B on δ_A ▷ T: rigid-body transformation
7: $v_B^{'} \leftarrow T(v_B)$ ▷ Apply T on v_B to move only axis vector
8: $\theta \leftarrow \langle (v_A, v_B^{'})$ ▷ Compute angle between axis vectors
9: **if** $\theta \in [0, \Theta]$ **then**
10: $C_{AB} \leftarrow T(C_B)$ ▷ Apply T on C_B to obtain a dimeric configuration C_{AB}
11: Compute atoms in contact in C_{AB} and measure LJ potential e_{AB} over them
12: **if** $e_{AB} \leq E_{\max}$ **then** ▷ C_{AB} meets both geometric and energetic constraints
13: $\Omega \leftarrow \Omega \cup C_{AB}$ ▷ Add new configuration to Ω
14: **end if**
15: **end if**
16: **end while**

base, immobile unit, and B as the moving one to be docked onto A. While details of the geometry-driven approach to rigid-body docking can be found elsewhere [12], we summarize it here, focusing on the novel integration of experimental knowledge of the (putative) interaction interface.

First, molecular surfaces, MS_A and MS_B, are built from C_A and C_B. Geometrically-complementary regions, represented as triangles of (critical) points selected to represent a molecular surface, are computed on these surfaces. The critical points allow direct incorporation of additional constraints. For instance, given a desired (hypothesized or putative) list of residues, perhaps demonstrated to contain the interaction interface, only the subset of critical points CP_A and CP_B nearby the specified residues can be considered, thus narrowing the triangles of interest to a smaller desired "active" subset. Two such geometrically-complementary triangles, δ_A and δ_B are sampled, and a rigid-body transformation T in $SE(3)$ is computed to superimpose δ_B over δ_A (lines 5-6 in algorithm 1). Applying T to all the atoms in C_B results in a new configuration for unit B, now docked onto unit A, and resulting in a dimeric model or decoy C_{AB}. The rest of this section details the two filters employed to *efficiently* obtain *credible* dimeric models.

2.1.2 Geometric Filter

The geometric filter serves to ensure that all dimeric models allowed in the decoy ensemble Ω satisfy the environmental constraints posed on GPCRs. All TM regions of a GPCR need to be inside the membrane. This constraint needs to be satisfied in GPCRs assemblies, as well. For this reason, an axis vector representation is employed to record the placement of a GPCR unit with respect to the membrane. As illustrated in Figure 2(a), given a global coordinate frame where the y axis runs normal to the lipid bilayer (and z is the lateral axis), a GPCR in an ideal placement with respect to the membrane has the main axis of the hollow cylinder formed by its TM regions aligned with the y axis. We compute the main axis of this cylinder and use it to represent the placement of a GPCR with respect to the membrane. Let us refer to this axis vector for a GPCR unit A as v_A, computed as follows: Given C_A, the top residues of a TM region are those just before the chain becomes extracellular, and the bottom residues are the last residues just before the chain becomes intracellular. The center of mass of the bottom residues, $cm_{A,\text{bottom}}$, and that of the top residues, $cm_{A,\text{top}}$, are calculated. The axis vector $v_A = cm_{A,\text{top}} - cm_{A,\text{bottom}}$. Similarly, an axis vector v_B is defined for a second unit B to be docked onto a unit A.

This axis vector representation allows rapidly rejecting dimeric models where unit B would be in an invalid placement relative to the membrane. One does not need to apply a rigid-body transformation T to the entire C_B. Instead, applying it to just the vector v_B (see line 7 in algorithm 1) results in a new vector $V_B^{'}$, which summarizes the placement of the moving unit B without having to move all its atoms in C_B in space. A simple geometric constraint needs to be satisfied for $v_B^{'}$ to be considered valid. The angle θ between v_A and $v_B^{'}$ needs to be in a specific segment of $[0, 2\pi]$. This is illustrated in Figure 2(b). In an ideal placement, where two units in a dimer are perfectly aligned with the global y axis, the angle $\theta = 0^\circ$. If one considers unit A in this ideal alignment, what is the maximum value that θ can take before the less than ideally-placed unit B has TM regions exiting the membrane? Doing this calculation in a preprocessing stage, rotating a given GPCR unit in 5° increments before its TM regions exit the membrane gives the maximum misalignment allowed for a GPCR unit relative to the membrane. The maximum allowed value for θ is twice this value, as each of the units A and B can be in these

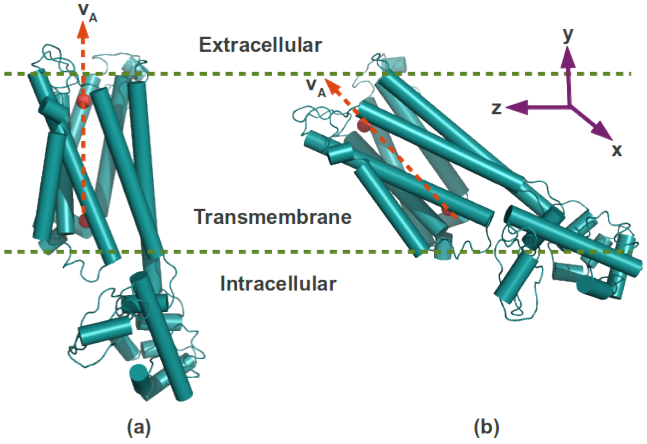

Figure 2: (a) An axis vector v_A is defined on a unit A to track the unit placement relative to the membrane (y is chosen to be normal to the lipid bilayer). (b) The axis vector allows determining the maximum allowed unit misplacement relative to the membrane, prior to TM regions exiting it. An angular constraint can be defined to force specific relative placements of two units in a dimer.

maximally-misaligned but still valid placements relative to the membrane. Lines $9-10$ show that if $0 \leq \theta \leq \Theta$, only then is the rigid-body transformation T applied to the entire C_B. The resulting dimeric model C_{AB} is next subjected to an energetic filter.

2.1.3 *Energetic Filter*

Amino acids of each unit in a geometrically-valid decoy C_{AB}, where CA atoms are within a d_{contact}Å of each-other, are considered part of the contact interface in C_{AB}. Rather than evaluating a computationally-intensive energy function, only the Lennard-Jones (LJ) term is computed first on atoms in contact. The goal is to filter out decoys C_{AB} with significant steric clashes that cannot be removed by energetic refinement or slight structural fluctuations. We implement the LJ potential based on the CHARMM22 force field [2] but downweight the contribution from the repulsion. We conduct a simple analysis: 1000 dimeric decoys are collected after they pass the geometric filter, and their distribution of energies is recorded. Three decoys with low, medium, and high energy values, according to the obtained distribution, are fed to Firedock [1], an energetic refinement protocol. Decoys where steric collisions could not be removed upon slight backbone and side-chain fluctuations by Firedock were noted in order to determine both a maximum energetic value $E_{\max}$ and a credible value by which to downplay the repulsion term. As shown in lines $11-12$ in algorithm 1, a decoy is considered valid and finally added to the ensemble Ω only if its LJ value is below $E_{\max}$.

2.2 Decoy Selection

We investigate different selection mechanisms, as it is generally challenging to determine a singular criterion by which to further discriminate sampled decoys and offer a subset for prediction. While most efforts either cluster decoys (based on structure comparison) or devise new scoring functions to further discriminate among them (or both), here we proceed as follows. First, we take an energy landscape approach, under which we assume that the different ways in which two GPCRs can be docked correspond to local optima or basins in a landscape. Rather than cluster decoys by structure comparison, we organize them according to the TMs involved in the interface. Essentially, we seek basins in the landscape that correspond to the distinct TMs involved in an interface (given x TMs, this means that the decoys are organized in $\binom{x}{2}$ states). We conduct density of state comparisons, where states correspond to these TM-based groupings of decoys, to determine whether there is a specific TM pairing that is most favored. Our analysis focuses only on decoys that meet specific criteria. We consider three such criteria here, one based on total energy, another based on a new scoring function, and a last one based on multi-objective analysis.

2.2.1 *Selection Based on Total Interaction Energy*

A detailed energetic evaluations can be performed on the decoy ensemble Ω to select decoys based on total interaction energy. We employ a popular physics-based force field, FoldX [37], which is defined as follows:

$$\Delta G = W_{LJ} \cdot \Delta G_{LJ} + W_{solvH} \cdot \Delta G_{solvH} + W_{solvP} \cdot \Delta G_{solvP} + \Delta G_{hbond} + \Delta G_{el} + all_{else} \quad (1)$$

where, ΔG_{LJ} is the sum of the LJ contributions from all atoms, ΔG_{solvH} and ΔG_{solvP} are the differences in solvation energy for polar and non polar groups, correspondingly, when transferring them from the unfolded to folded stage. ΔG_{hbond} is the free energy difference between the formation of an intra-molecular hydrogen bond and that of an inter-molecular hydrogen-bond (with solvent), and ΔG_{el} is the electrostatic contribution from charged groups. The all_{else} groups the rest of the terms in FoldX. W_{LJ}, W_{solvH} and W_{solvP} are the weights applied to the terms. Using FoldX, one can focus only on the lowest-energy decoys. Distributing these decoys into the TM-based states described above allows then comparing basins in the energy landscape. Various statistics can be measured, the most basic of which is number of decoys in a basin, or density of state. Such comparison allows determining which ones are the most populated basins to determine whether certain TM pairings in the interaction interface are more thermodynamically-stable (energetically-favored) over others.

2.2.2 *Selection Based on Combined Score (CS)*

Alternatively, the density of state comparison can be conducted on decoys that may not be lowest in terms of interaction energy but some other scoring function. As it is known that even the most sophisticated energy functions are inaccurate, it is important to supplement the analysis in decoy selection with other scoring functions. We design one such here. When comparing two basins, depth (lowest energy) is not the most important consideration. In fact, focusing on lowest energies may promote outliers, which may arise either due to artifacts in the energy function or lack of sampling. We design a "Combined Score" (CS) to penalize outliers. For each state/basin, we record not only the lowest energy D_{state}, but also the Z-score, $Z_{\text{state}} = (D_{\text{state}} - \mu_{\text{state}})/\delta_{\text{state}}$, where μ_{state} and δ_{state} are the mean and standard deviation over energy values of decoys in a given state. The penalty term, C_{state}, in equation 2, estimates high structural diversity via the maximum least root mean square deviation (lRMSD) [27] between any two decoys in a state. Diversity in a state is penalized for the following reason: if

a specific TM pairing is observed among low-energy dimers, that pairing is more likely to be native if there is consistency among the obtained decoys; that is, the sampling algorithm repeatedly reproduces it.

$$\mathrm{CS}_{\mathrm{state}} = \frac{\mathrm{D}_{\mathrm{state}}}{\mathrm{Z}_{\mathrm{state}} \cdot \mathrm{C}_{\mathrm{state}}} \quad (2)$$

2.2.3 Selection Based on Multi-objective Analysis

The terms summed in an energy function, even as sophisticated as FoldX, are conflicting optimization criteria. Analysis of correlations between the energy terms in FoldX reveals which ones are conflicting and need to be in separate groups/criteria, and which can be grouped together. There are four important terms in physics-based interaction energy function, such as FoldX: solvation, hydrogen-bonding, LJ, and electrostatic. We group the solvation and hydrogen-bonding terms together, obtaining thus 3 groups. Our analysis (data not shown) reveals that these groups have either negative or no correlation with each-other, whereas the solvation and hydrogen-bonding terms have a positive correlation and thus can be grouped together. The 3 groups can be treated as separate optimization objectives, and decoys can be evaluated and compared across all these groups, based on the concept of Pareto dominance.

Pareto-based multi-objective analysis is popular in evolutionary computation, and we employ it here as our final strategy for analysis. A configuration C_i is said to dominate another configuration C_j, if every energy term of C_i is lower than that of C_j and C_j is said to be dominated by C_i. This is also noted as strong dominance, and is what we employ here. The Pareto count PC_i of a configuration C_i denotes the total number of configurations that C_i dominates. The Pareto rank PR_i of a configuration C_i is the total number of configurations that dominates C_i. If C_i is not dominated by any configuration in the ensemble Ω, then C_i has Pareto rank 0. The set of all such configurations are also referred to as the Pareto front. Pareto rank and Pareto count can be additionally employed to compare decoys in the ensemble. Given two decoys, the one with lower Pareto rank may be preferred in a selection procedure. If two decoys have the same Pareto rank, the one with lower Pareto count may be preferred, as lower Pareto count ensures energetic diversity [40]. Using these metrics, Pareto rank, Pareto count, in addition to a total energy value, one can define a sorted order of the decoys in the ensemble Ω and use this order to select a subset of credible decoys for prediction.

3. RESULTS

We demonstrate the proposed protocol and its usefulness on offering credible models of dimerization of D2R, a central GPCR. A consensus is emerging from wet-lab studies that interaction interfaces for higher-order assembly of D2Rs involve a subset of TMs, namely, $TM1$, $TM4$ and $TM5$ [9, 11, 23]. We emphasize that the application setup of the proposed protocol is a blind prediction setting, as no structural models exist for GPCR dimerization. This paper constitutes some of the early efforts in this direction. In addition, no X-ray structure of a D2R unit exists. However, we exploit the high structural similarity among GPCRs to obtain a a credible model for a D2R unit. We employ I-TASSER [36], a top performer in the community-wide **C**ritical **A**ssessment of Protein **S**tructure **P**rediction (CASP) [31] and the human X-ray structure of D3R, the closest GPCR to D2R available to date, as a template model for I-TASSER to build a model for a D2R unit.

3.1 Implementation Details

Table 1: Parameter Settings

Parameter	Values
Size (Nr. of Atoms)	7198
Active TMs	$1, 4, 5$
d_{contact} (Å)	6.5
Θ(degree)	$[0, 80]$
E_{max} (kcal/mol)	$100,000$
w_{attr}	1.0
w_{repul}	0.1
Ω	10,000

Parameter values used in application of our protocol to D2R are shown in Table 1. Row 1 shows the size of a D2R unit (number of atoms). Row 2 shows that, in keeping with experimental knowledge [9–11, 23], the active regions used are those on TM1, TM4, and TM5. Row 3 shows that d_{contact} to determine amino acids in contact is set to 6.5Å, as in [19]. Row 4 shows the value for the angular threshold Θ used by the geometric filter. This value was obtained by observing the angle, in 5°, at which TMs in a D2R exit the membrane. Beyond 40°, a significant portion of regions in TMs exit the membrane. So, Θ is set to twice this value, 80°. Row 5 shows the $E_{\max}$ value used by the energetic filter. Rows 6−7 show the weights used for the attraction and repulsion terms in the LJ potential. We recall that the reason for downplaying the contribution from the repulsion is to allow slight inter-unit penetrations that can be resolved with short energetic refinement protocols and anticipated structural changes, particularly in the possibly highly-flexible loop regions. The results presented here for D2R dimerization are those with an ensemble Ω of $10,000$ sampled decoys (last row in Table 1). A second ensemble of $50,000$ decoys has been generated. The same results have been obtained, which allows concluding that an ensemble of $10,000$ decoys is sufficient to reach conclusions on D2R dimerization.

3.2 Distribution of Lowest-energy Decoys

Decoys are evaluated based on their total FoldX energy value, and only those with negative energies are retained in the analysis here. These decoys are grouped into the different TM-based states, and the population density of these states is shown in Figure 3. Comparison of these densities shows that the three most populous states are those where TM4 is involved in the interaction interface. These three states (TM1-TM4, TM4-TM4, and TM4-TM5) also contain some of the lowest-energy decoys. Figure 4 summarizes the density of state comparison in a heatmap color-coded by density. The heatmap shows that two most populous states are those corresponding to the TM1-TM4 and TM4-TM4 interface, with TM4-TM5 following third.

These results agree strongly with cross-linking wet-lab studies in [9–11, 23], which propose a central role for TM4 in D2R dimerization. Work in [9] proposes that TM4-TM4 is the specific interface, while others [11, 23] suggest that TM4 may interface with TM1 and TM5, as well, and form the base for higher-order assemblies. Our results in Figures 3 and 4 indicate that all interfaces where TM4 is involved are

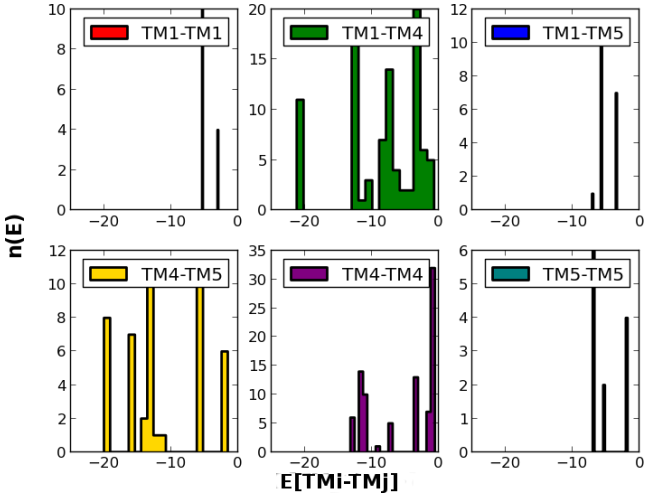

Figure 3: Histogram of lowest-energy decoys grouped in the different TMi-TMj states.

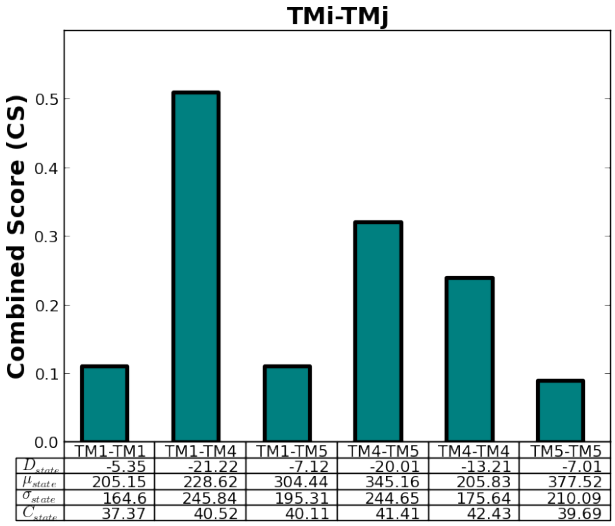

	TM1-TM1	TM1-TM4	TM1-TM5	TM4-TM5	TM4-TM4	TM5-TM5
D_{state}	-5.35	-21.22	-7.12	-20.01	-13.21	-7.01
μ_{state}	205.15	228.62	304.44	345.16	205.83	377.52
σ_{state}	164.6	245.84	195.31	244.65	175.64	210.09
C_{state}	37.37	40.52	40.11	41.41	42.43	39.69

Figure 5: CS is shown for each state. Values of various metrics used in CS are shown in the table.

energetically-favorable. This suggests that TM4 may be involved in both stable and transient interfaces, possibly providing flexibility for formation of higher-order assemblies.

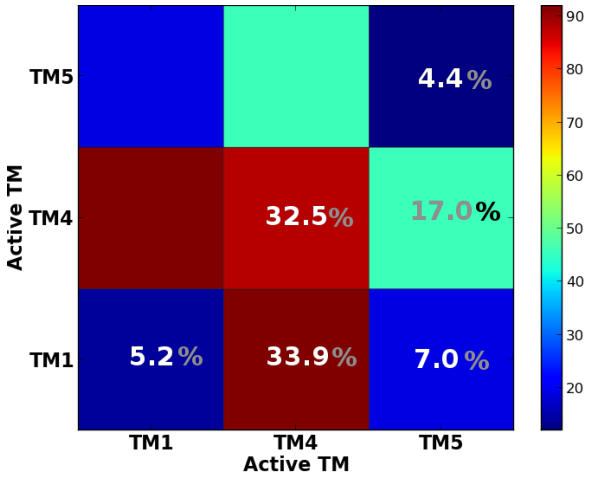

Figure 4: States are color-coded based on their density (focusing only on negative-energy decoys in Ω). Actual distribution of population is indicated over each cell/state in the heatmap.

3.3 Distribution of Highest-CS Decoys

The CS of each state is computed and shown in the bar diagram in Figure 5. According to this score, the same three interfaces where TM4 is involved rank highest. The score favors TM1-TM4. While the structural diversity in each of these three states is similar, TM1-TM4 has more configurations with low energies (as evidenced by the mean energy value shown in the table below the bar diagram). The score also slightly favors TM4-TM5 over TM4-TM4. These results support those shown above, that it is perhaps TM1-TM4 that promotes a stable dimer, but other TM4-based interfaces may be transient dimers possibly giving D2R the flexibility to form a rich set of higher-order assemblies.

3.4 Distribution of Decoys Based on Multi-objective Analysis

Decoys are now compared according to the Pareto rank and Pareto count measures, as well. The following two analyses are conducted. First, decoys are sorted based on their Pareto rank (low to high), and those with the same Pareto rank are further sorted based on their energies (low to high). The configurations in the lowest $p\%$ in this sorted ordering are retained in what we refer to as the $\Omega_{PR:E_{p\%}}$ subensemble. The decoys in this ensemble are grouped into the possible 6 TM-based states, and the density of state analysis is conducted. This is repeated with a different sorted ordering, where decoys with the same Pareto rank are then sorted by their Pareto count (low to high), and those with the same Pareto count are further sorted by their total energy (low to high). The retained subensemble is now referred to as $\Omega_{PR:PC:E_{p\%}}$ ensemble. The bar diagram in Figure 6 shows the density of state analysis for each sorted ordering for $p \in \{5, 10\}$.

Figure 6 shows that two specific states, TM1-TM4 and TM4-TM5, are consistently well-populated with changing p and the Pareto metrics. The additional consideration of Pareto count in the decoy selection favors TM5, with TM4-TM5 ranking higher than TM1-TM4. These results suggest that TM5 adds specific energetic contributions to the interaction interface, albeit they may not result in the lowest interaction energy.

3.5 Proposed Models for D2R Dimerization

Taken together, the results in this work conclusively demonstrate that TM4 is part of the interaction interface in D2R dimerization. Moreover, a specific interface, TM1-TM4, achieves some of the lowest energy values. Since low energy states the fact that it has high affinity, we propose TM1-TM4 as the core interface for a long-lived, stable dimer. The emergence of other states, such as TM4-TM4 and TM4-TM5, is proposed here to be seminal for transient, shorter-lived dimers, which may be less stable than TM1-TM4, but may be employed as alternative ways to link D2R units in higher-order assemblies.

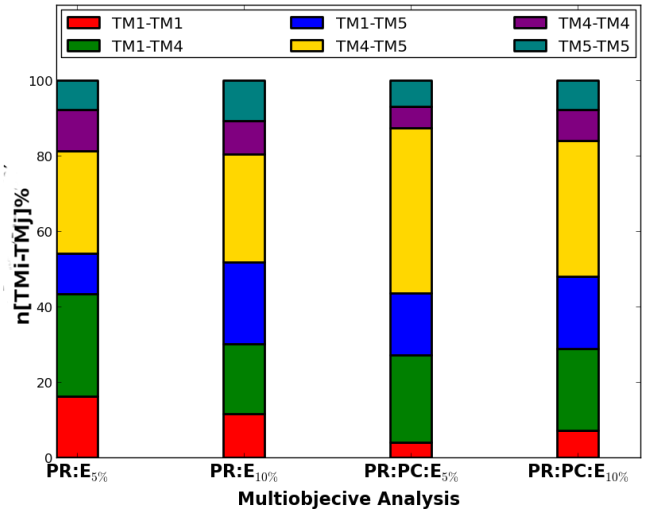

Figure 6: Decoys in the $\Omega_{PR:E_{p\%}}$ ensemble are organized into the 6 possible states, and the density analysis is repeated in the first two bar diagrams for $p = 5\%$ and $p = 10\%$. The density analysis is repeated on decoys in the $\Omega_{PR:PC:E_{5\%}}$ and $\Omega_{PR:PC:E_{10\%}}$ ensemble in the next two bars.

We show here in structural detail some of the actual D2R-D2R models we propose to occur in the cell. Figures 7(a)-(b) show the representative model/decoy in the TM1-TM4 state proposed here for a stable D2R-D2R dimer. It is interesting to note that a TM1-TM4 interface promotes a secondary symmetric TM4-TM1 interface (see top view in Figure 7(a)), which additionally lowers the potential energy and further explains why a TM1-TM4 interface may be the most stable in the cell. Furthermore, a side view in Figure 7(b) shows that this interface allows enough room for the intracellular loops to move around without energetic penalty. High entropy, due to the the mobility of these loop regions, while not directly considered in our calculations, may add to the free energy of the TM1-TM4 state, further lending credibility to this state as the most stable one for a D2R-D2R dimer. In contrast, Figure 7(c)-(d) shows a representative model for the TM4-TM4, and Figure 7(e)-(f) shows a representative model for the TM4-TM5 interface, which we predict as transient. It is clearly evident that these interfaces do not readily allow for a secondary symmetric interface. Moreover, the intracellular loop regions are closer to each-other than in the TM1-TM4 state. While this may promote secondary interactions between the loop regions, the entropic cost may be too high in these alternative dimeric states.

4. CONCLUSION

We have proposed here a complete protocol to predict models of GPCR dimerization. The protocol has been demonstrated on D2R, a central GPCR involved in addiction and psychotic disorders. Due to its central role in various disorders, D2R participation in higher-order assemblies has been investigated in many wet laboratories, which provides us both with information that is integrated in the protocol to improve model credibility, as well as validation for the models proposed by the protocol. In particular, we propose a specific model that pairs TM1 with TM4 as representative of the stable and long-lived D2R-D2R dimer. Other transient dimeric states are proposed, all involving TM4. These alternative shorter-lived states may be instrumental to formation of higher-order assemblies of D2Rs. The proposed protocol is general, as many GPCRs share both structural similarity and experience similar placement constraints imposed by the cellular membrane. Future work will focus on investigating both higher-order assemblies of D2R, as well as oligomerization of other GPCRs. Taking into account the membrane environment is another direction of future work.

Acknowledgment

Many of the computations were performed on ARGO, a research computing cluster provided by the Office of Research Computing at George Mason University. Funding for this work is provided in part by NSF IIS CAREER Award No. 1144106 to AS and a Virginia Youth Tobacco Program Award to NK and AS.

5. REFERENCES

[1] N. Andrusier, R. Nussinov, and H. J. Wolfson. Firedock: fast interaction refinement in molecular docking. *Proteins: Struct. Funct. Bioinf.*, 69(1):139–159, 2007.

[2] B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, and M. Karplus. Charmm: A program for macromolecular energy, minimization, and dynamics calculations. *Journal of computational chemistry*, 4(2):187–217, 1983.

[3] S. R. Comeau, D. W. Gatchell, S. Vajda, and C. J. Camacho. ClusPro: a fully automated algorithm for protein-protein docking. *Nucl. Acids Res.*, 32(S1):W96–W99, 2004.

[4] W. L. DeLano. Pymol. *DeLano Scientific, San Carlos, CA*, 700, 2002.

[5] D. Duhovny-Schneidman, Y. Inbar, R. Nussinov, and H. J. Wolfson. PatchDock and SymmDock: servers for rigid and symmetric docking. *Nucl. Acids Res.*, 33(S2):W363–W367, 2005.

[6] S. Engelen, A. T. Ladislas, S. Sacquin-More, R. Lavery, and A. Carbone. A large-scale method to predict protein interfaces based on sequence sampling. *PLoS Comp Bio*, 5(1):e1000267, 2009.

[7] S. R. George and B. F. O'Dowd. A novel dopamine receptor signaling unit in brain: heterooligomers of D1 and D2 dopamine receptors. *Scientific World J*, 7:58–63, 2011.

[8] D. S. Goodsell and A. J. Olson. Structural symmetry and protein function. *Annu. Rev. Biophys. and Biomolec. Struct.*, 29:105–153, 2000.

[9] W. Guo, L. Shi, M. Filizola, H. Weinstein, and J. A. Javitch. Crosstalk in g protein-coupled receptors: changes at the transmembrane homodimer interface determine activation. *Proceedings of the National Academy of Sciences of the United States of America*, 102(48):17495–17500, 2005.

[10] W. Guo, L. Shi, and J. A. Javitch. The fourth transmembrane segment forms the interface of the dopamine D2 receptor homodimer. *J. Biol. Chem.*, 278:4385–4388, 2003.

[11] W. Guo, E. Urizar, M. Kralikova, J. C. Mobarec, L. Shi, M. Filizola, and J. A. Javitch. Dopamine d2

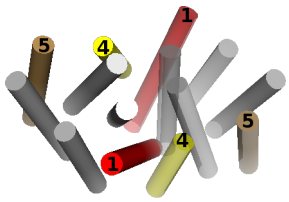

(a) TM1-TM4 top view

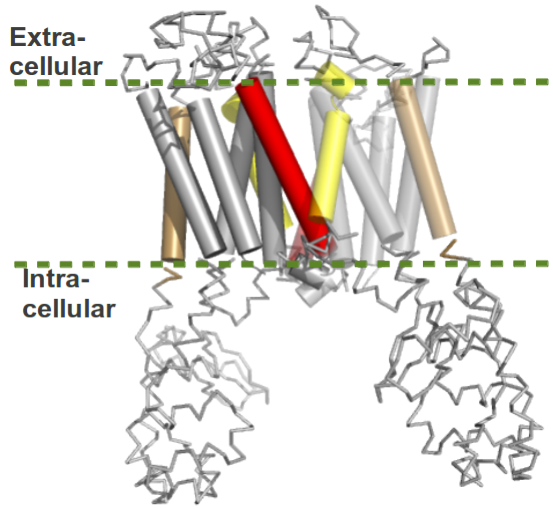

(b) TM1-TM4 side view

(c) TM4-TM4 top view

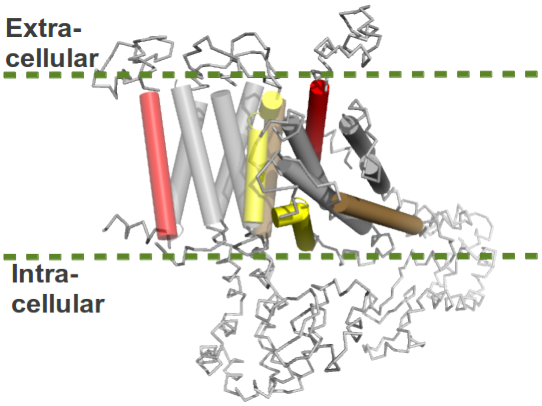

(d) TM4-TM5 side view

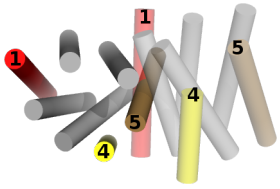

(e) TM4-TM5 top view

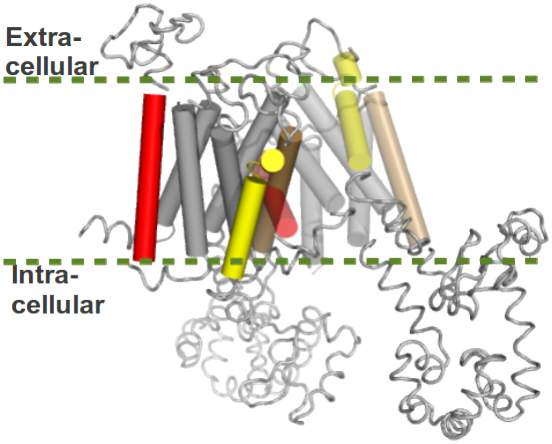

(f) TM4-TM5 side view

Figure 7: (a)-(f) Representative models are drawn with pyMol [4], with the TM regions shown as cylinders. TM1 is drawn in red, TM4 in yellow, and TM5 in tan. For ease of visualization, the base unit is drawn in transparent, while the other one in opaque. (a)-(b) show a representative model of the TM1-TM4 state, whereas (c)-(d) show one such model of the TM4-TM4 state, and (e)-(f) show one such model of the TM4-TM5 state. Top views are shown in (a), (c) and (e), and side views in (b), (d) and (f)

receptors form higher order oligomers at physiological expression levels. *The EMBO journal*, 27(17):2293–2304, 2008.

[12] I. Hashmi and A. Shehu. Hopdock: A probabilistic search algorithm for decoy sampling in protein-protein docking. *Proteome Sci*, 11(Suppl 1):S6, 2013.

[13] I. Hashmi and A. Shehu. Informatics-driven protein-protein docking. In *ACM Conf on Bioinf and Comp Biol Workshops (BCBW)*, pages 772–779, Washington, D. C., September 2013.

[14] S.-Y. Huang. Search strategies and evaluation in protein–protein docking: principles, advances and challenges. *Drug discovery today*, 2014.

[15] B. Jiménez-García, C. Pons, and J. Fernández-Recio. pydockweb: a web server for rigid-body protein–protein docking using electrostatics and desolvation scoring. *Bioinformatics*, 29(13):1698–1699, 2013.

[16] N. Kabbani, J. C. Nordman, B. Corgiat, D. Veltri, A. Shehu, and D. J. Adams. Are nicotinic receptors coupled to g proteins? *BioEssays*, 35(12):1025–1034, 2013.

[17] N. Kabbani, M. P. Woll, J. C. Nordman, and R. Levenson. Dopamine receptor interacting proteins: targeting neuronal calcium sensor-1/D2 dopamine receptor interaction for antipsychotic drug development. *Curr. Drug Targets*, 13(1):72–79, 2012.

[18] E. Kanamori, Y. Murakami, Y. Tsuchiya, D. Standley, H. Nakamura, and K. Kinoshita. Docking of protein molecular surfaces with evolutionary trace analysis. *proteinssfb*, 69:832–838, 2007.

[19] O. Keskin, B. Ma, and R. Nussinov. Hot regions in protein-protein interactions: the organization and contribution of structurally conserved hot spot residues. *J. Mol. Biol.*, 345(5):1281–1294, 2005.

[20] D. Kozakov, D. Beglov, T. Bohnuud, S. E. Mottarella, B. Xia, D. R. Hall, and S. Vajda. How good is automated protein docking? *Proteins: Structure, Function, and Bioinformatics*, 81(12):2159–2166, 2013.

[21] F. Lennartz, T. Hoenen, M. Lehmann, A. Groseth, and W. Garten. The role of oligomerization for the biological functions of the arenavirus nucleoprotein. *Arch. Virol.*, 2013. pre-print.

[22] M. F. Lensink and S. J. Wodak. Blind predictions of protein interfaces by docking calculations in CAPRI. *Proteins: Struct. Funct. Bioinf.*, 78(15):3085–3095, 2010.

[23] Y. Liang, D. Fotiadis, S. Filipek, D. A. Saperstein, K. Palczewski, and A. Engel. Organization of the g protein-coupled receptors rhodopsin and opsin in native membranes. *Journal of Biological Chemistry*, 278(24):21655–21662, 2003.

[24] Q. Liu and J. Li. Protein binding hot spots and the residue-residue pairing preference: a water exclusion perspective. *BMC Bioinf*, 11:244, 2010.

[25] S. Lyskov and J. J. Gray. The RosettaDock server for local protein-protein docking. *Nucl. Acids Res.*, 36(S2):W233–W238, 2008.

[26] E. Mashiach, R. Nussinov, and H. J. Wolfson. Fiberdock: Flexible induced-fit backbone refinement in molecular docking. *Proteins: Struct. Funct. Bioinf.*, 78(6):1503–1519, 2010.

[27] A. D. McLachlan. A mathematical procedure for superimposing atomic coordinates of proteins. *Acta Crystallogr. A.*, 26(6):656–657, 1972.

[28] N. Moitessier, P. Englebienne, D. Lee, J. Lawandi, and C. R. Corbeil. Towards the development of universal, fast and highly accurate docking/scoring methods: a long way to go. *British J Pharmacology*, 153(S1):S7–S27, 2009.

[29] S. Mondal, J. M. Johnston, H. Wang, G. Khelashvili, M. Filizola, and H. Weinstein. Membrane driven spatial organization of gpcrs. *Scientific reports*, 3, 2013.

[30] E. Moreno, H. Hoffmann, M. Gonzalez-Sepúlveda, G. Navarro, C. V., A. Cortés, J. Mallol, M. Vignes, P. J. McCormick, E. I. Canela, C. Lluís, R. Moratalla, S. Ferré, J. Ortiz, , and R. Franco. Dopamine D1-histamine H3 receptor heteromers provide a selective link to MAPK signaling in gabaergic neurons of the direct striatal pathway. *J. Biol. Chem.*, 286:5846–6854, 2011.

[31] J. Moult, K. Fidelis, A. Kryshtafovych, and A. Tramontano. Critical assessment of methods of protein structure prediction (CASP) - round IX. *Proteins: Structure, Function, and Bioinformatics*, 79(S10):1–5, 2011.

[32] G. Y. Ng, B. F. O'Dowd, S. P. Lee, H. T. Chung, M. R. Brann, P. Seeman, and S. R. George. Dopamine d2 receptor dimers and receptor-blocking peptides. *Biochemical and biophysical research communications*, 227(1):200–204, 1996.

[33] M. L. Perrault, A. Hasbi, O. B. F., and S. R. George. Heteromeric dopamine receptor signaling complexes: emerging neurobiology and disease relevance. *Neuropsychopharmacology*, 39(1):156–168, 2014.

[34] B. G. Pierce, K. Wiehe, H. Hwang, B.-H. Kim, T. Vreven, and Z. Weng. Zdock server: interactive docking prediction of protein-protein complexes and symmetric multimers. *Bioinformatics*, page btu097, 2014.

[35] D. W. Ritchie. Recent progress and future directions in protein-protein docking. *Current protein & peptide science*, 9(1):1, 2008.

[36] A. Roy, A. Kucukural, and Y. Zhang. I-tasser: a unified platform for automated protein structure and function prediction. *Nature protocols*, 5(4):725–738, 2010.

[37] J. Schymkowitz, J. Borg, F. Stricher, R. Nys, F. Rousseau, and L. Serrano. The foldx web server: an online force field. *Nucl. Acids Res.*, 33(Web server issue):W382–W388, 2005.

[38] M. Torchala, I. H. Moal, R. A. Chaleil, J. Fernandez-Recio, and P. A. Bates. Swarmdock: a server for flexible protein–protein docking. *Bioinformatics*, 29(6):807–809, 2013.

[39] A. Tovchigrechko and I. A. Vakser. GRAMM-X public web server for protein-protein docking. *Nucl. Acids Res.*, 34(Web Server issue):W310–4, 2006.

[40] J. A. Vrugt, H. V. Gupta, L. A. Bastidas, W. Bouten, and S. Sorooshian. Effective and efficient algorithm for multiobjective optimization of hydrologic models. *Water Resources Research*, 39(8), 2003.