Identification of ligand-specific G-protein coupled receptor states and prediction of downstream efficacy via data-driven modeling

Oliver Fleetwood[1], Jens Carlsson[2], and Lucie Delemotte*[1]

[1] *Science for Life Laboratory, Department of Applied Physics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden*

[2] *Science for Life Laboratory, Department of Cell and Molecular Biology, Uppsala University, SE-751 05 Uppsala, Sweden*

** Corresponding author. E-mail: lucie.delemotte@scilifelab.se*

Abstract

G protein-coupled receptors (GPCRs) shift between inactive non-signalling states and active signalling states, to which intracellular binding partners can bind. Extracellular binding of ligands stabilizes different receptor states and modulates the intracellular response via a complex and not well understood allosteric process. Despite the recent advances in structure determination and spectroscopy techniques, a comprehensive view of the ligand-protein interplay remains a challenge. We derived the activation free energy of the β2 adrenergic receptor (β2AR) bound to ligands with different efficacy profiles using enhanced sampling molecular dynamics (MD) simulations. The resulting free energy landscapes reveal clear shifts towards active-like states at the G protein binding site for receptors bound to partial and full agonists compared to antagonists and inverse agonists. Not only do the ligands control the population of states, they also modulate the receptor's conformational ensemble by flipping allosteric protein microswitches. We find an excellent correlation between the conformation of the microswitches close to the ligand binding site and in the transmembrane region and experimentally reported cAMP signaling responses, highlighting the predictive power of our apprch. Using dimensionality reduction techniques, we could further assess the similarity between the unique states induced by different ligands. Two hotspots governing agonism on transmembrane helix 5 (TM5) and TM7, including the conserved NPxxY motif, formed the endpoints of an allosteric pathway between the binding sites. Our results demonstrate that computational methods can be used to assess GPCR drug candidates with specific efficacy profiles.

Introduction

G protein-coupled receptors (GPCRs) are membrane proteins which activate cellular responses by binding to and activating intracellular binding partners. This process is controlled by extracellular binding partners, such as ligands or peptides, and allosteric stimuli which increase the probability of intracellular activation, via receptor-mediated signal transduction across the cell membrane. GPCRs are vital in many physiological processes and constitute the most common class of drug targets[1].

Much of the current understanding of GPCR signalling at the molecular level can be attributed to the progress in GPCR structure determination during the last decade[2–7]. GPCRs interconvert between inactive (R) and active (R*) states, which can accommodate G protein binding to a conserved intracellular domain via conformational rearrangements amongst the seven transmembrane (TM) helices (Fig. 1a)[8]. This process is called basal activity. Ligands can bind to the orthosteric site in the receptor's extracellular domain and thereby control conformational rearrangements. Extracellular ligands are traditionally classified as either agonists, which promote activation, antagonists, which do not alter basal activity, or inverse agonists that reduce basal activity. However, this classical view of ligand efficacy is complicated by the fact that GPCRs can signal via several intracellular partners, e.g. G-proteins or β-arrestins. Most agonists will activate several signalling pathways, but agonists with the ability to activate one specific intracellular partner have also been identified, a phenomenon commonly referred to as *biased signalling*. Based on spectroscopy and structure determination studies, conformational changes in the receptor govern the activation of signalling pathways [6,9,10], although the underlying molecular mechanisms remain elusive. Characterization of the allosteric process guiding interactions with intracellular partners is a major challenge and can only be fully understood by using a combination of different methodologies.

The term *microswitch*, or *molecular switch,* describes local structural changes in the receptor believed to be important for activation and can, for example, involve sidechain rotamers, movement of two domains relative to each other, or a helix twist. Two microswitches implicated in activation of class A GPCRs are an outward displacement of the transmembrane helix 6 (TM6) and twist of the highly conserved N(7.49)P(7.50)xxY(7.53) motif (superscripts notation according to Ballesteros-Weinstein numbering [11]), which take part in the formation of the intracellular binding site. In the orthosteric site, microswitches are typically less conserved between receptors and depend on the type of ligand recognized by the receptor [8].

In this study, we focus on the prototypical β2 adrenergic receptor (β2AR), which interacts with Gs proteins to trigger a cyclic adenosine monophosphate (cAMP) response, and arrestins, which control endocytosis and kinase activation [12]. Both pathways are physiologically important and are modulated by therapeutic drugs. The β2AR is a common drug target for bronchoconstriction medication and was the first receptor crystallized in complex with a G protein [2,13]. Experimental studies, including crystallography [4,6,13], spectroscopy methods [10,14–17] and computational methods [18–21], have investigated the activation mechanism of the β2AR. Agonists resolved in experimental structures show a key interaction with S207(5.46) [22] which leads to an inward bulge of TM5 in the active state. In the TM domain between the orthosteric

site and G protein binding site*, a connector region* [23], or the P(5.50)I(3.40)F(6.44) motif, undergoes a rotameric change and interacts with the hydrated cavity surrounding the conserved Asp79(D2.50) (Imai et al. 2020), which in turn interacts with the conserved NPxxY motif in TM7. The combination of several microswitches lead to conformational changes that promote an outward movement of TM6 and G protein binding.

Considering the high dimensionality of a protein with over 300 interacting residues, it is difficult to identify relevant microswitches from the sequence or static experimental structures. Historically, sequence analysis and mutagenesis experiments [14,15,24–26] have been used to characterize motifs important for signalling, but this approach may overlook the roles of less conserved residues, solvent, and ions in ligand recognition and the activation mechanism [21]. Molecular dynamics (MD) simulations can generate trajectories from experimental starting structures and capture the dynamics of all microswitches, and allow us to derive the free energy landscapes governing the equilibrium between protein states. However, due to the computational cost of brute-force MD simulations, it is nearly impossible to obtain converged results unless enhanced sampling methods are used. Such methods introduce a bias in the simulations to explore the conformational landscape of the protein faster [27]. In a post-processing step, the bias can be adjusted for and it is possible to derive theoretically exact results at a fraction of the cost compared to unbiased simulations. However, the ever-increasing size of simulation data and diverse conformations sampled by MD makes it difficult to identify and judge the importance of microswitches. Data-driven and machine learning approaches can help to condense the data and reduce human bias in the interpretation of the result [28,29].

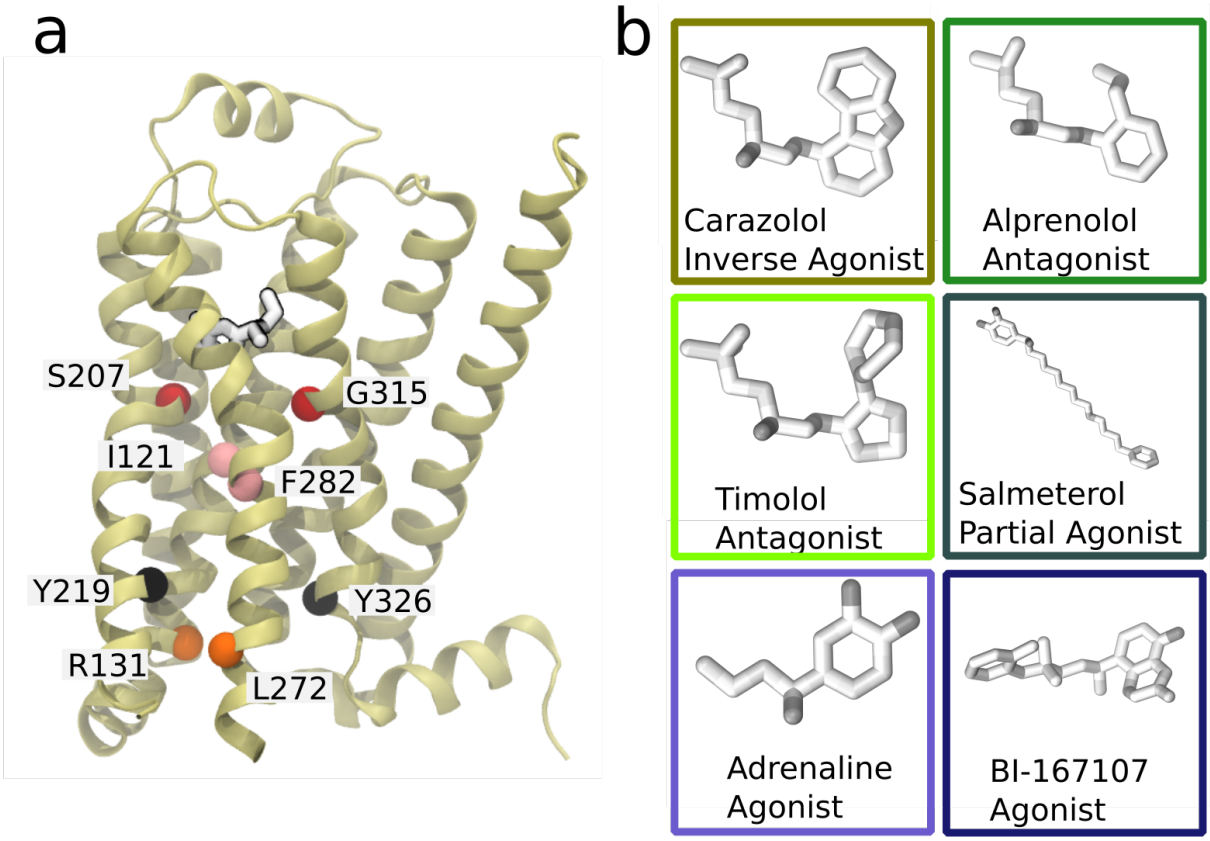

Figure 1. (a) An MD snapshot of the β2adrenergic receptor in complex with adrenaline in an active-like state (simulation starting from PDB 3P0G). Residue pairs reflecting important microswitches are highlighted as spheres: the TM5 bulge (red), measured as the closest heavy atom distance between S207(5.46) and G315(7.41); the connector region's conformational change (pink), measured as the difference in root-mean-square deviation (RMSD) between the active and inactive state structure of the residues I121(3.40) and F282(6.44); the Y-Y motif (black), measured as the C-ζ distance between Y219(5.58) and Y326(7.53) of the NPxxY motif; and the TM6-TM3 displacement (orange), measured as the closest heavy atom distance between L272(6.34) and R131(3.50). (b) Ligands considered in this study: agonists BI-167107 and adrenaline; biased partial agonists salmeterol; antagonists timolol and alprenolol; and the inverse agonist carazolol. Atoms are colored according to their partial charge.

In this study, we derived the activation free energy and characterized the details of the active-like state of the β2AR (Fig. 1a) in its ligand-free state and bound to ligands with different efficacy profiles (Fig. 1b) using swarm-of-trajectories enhanced sampling MD [30,31]. The set of ligands studied consisted of the full agonists BI-167107 and adrenaline, the G protein biased agonist salmeterol, the antagonists alprenolol and timolol (sometimes classified as a partial inverse agonist [7]), and the inverse agonist carazolol. Crystal structures of these ligands bound to the β2AR have been solved (Fig. 1b) and several are clinically approved drugs [32]. The free energy landscapes revealed a clear stabilization of active-like states for the receptor bound to agonists and a shift towards inactive-like states for the receptor bound to antagonists or inverse agonists. Remarkably, we obtained a strong quantitative correlation between experimentally measured intracellular cAMP response and the expectation values of the upper and transmembrane microswitches, highlighting the predictive power of our approach. In a second step, we developed an adaptive sampling protocol to quantitatively sample the most stabilized states kinetically accessible from the activated starting structure (which we will refer to as the active-like state). Using dimensionality reduction techniques we found that all ligands stabilize distinct receptor states and that ligands with similar pharmacological properties cluster together. Remarkably, many residues often considered to be of significance for GPCR activation, such the NPxxY motif and the extracellular end of TM5, were automatically identified as important with our protocol. Combined with the activation free energies, our results show how ligands control the population of states. They modulate the conformational equilibrium by flipping important allosteric microswitches, in particular near the G protein binding site. By inspecting the inter-residue contacts formed for different ligands, we identified an allosteric pathway between the two binding sites and a large heterogeneity of TM7 states. These results demonstrate how computational methods can be used to characterize GPCR drug candidates and provide insights into how ligands with specific efficacy profiles can be designed.

Methods

Simulation system configuration

We initiated simulation systems from a nanobody-bound active-state structure (PDB ID: 3P0G) [2] in CHARMM-GUI [33] with the CHARMM36m force field [34]. Following the protocol of a previous study [31], we reversed the N187E in the crystallized structure, protonated E122(3.41), and protonated the two histidines H172(4.64) and H178(4.70) at their epsilon positions. The receptor was embedded in a POPC membrane bilayer [35] of 180 molecules, then solvated in a 0.15M concentration of neutralizing sodium and chloride ions with 79 TIP3P water molecules [36] per lipid molecule. We performed the MD simulations with GROMACS 2018.6 [37] patched with PLUMED [38]. Ligands present in PDB structures 2RH1 [3], 3NYA [5], 3D4S [7], 6MXT [6] and 4LDO [4] were inserted into the 3P0G structure after alignment on residues interacting with the ligand. Input files required to run the simulations in this study are available online [39].

String method with swarms of trajectories

We used an optimized version of the string method with swarms of trajectories to enhance sampling and to estimate the free energy along various microswitches [30,31]. This method finds the most probable transition pathway between two end states in a high dimensional space spanned by a set of Collective Variables (CVs; in this context synonym with *reaction coordinates*). Given the initial guess of *points* distributed along a *string* in CV space, in this case the transition path from the previously published apo simulation, the system was equilibrated by running a 200 nanosecond steered MD simulation along the string with force constant per CV of 3366 kJ scaled by its estimated importance (more details in the following section), followed by a 7 nanosecond initial restrained equilibration at every point. Next, a *swarm* of 10 picosecond long trajectories were launched from the output coordinates of every restrained simulation. The average drift of the swarm was computed as the mean displacement of the short swarm trajectories in CV space, which is proportional to the gradient of the free energy landscape. Every swarm consisted of 16-32 trajectories and the exact number was determined adaptively to converge the drift vector. The points were displaced according to their drift and realigned along the string to maximize the number of transitions between neighbouring points. The string was updated iteratively with 30 ps of restrained equilibration per point, followed by a batch of swarms and string reparametrization. Gradually, the string relaxed into the most probable transition path connecting energetically stable intermediates between the two endpoints. We ran the simulations for 305 iterations, requiring an aggregated simulation time between 3-5.5 microseconds per ligand.

Kinetically trapped active-like state sampling

In order to quantitatively sample the most stabilized state kinetically accessible from the starting structure without applying an artificial force on the system, we developed an adaptive sampling protocol. A single swarm with 7.5 nanosecond long trajectories was launched from the same initial configuration as described in the previous paragraph. The swarm's center point, ***c***, in CV

space was taken as the mean of the trajectories' endpoint coordinates, $\boldsymbol{x}_i$. Next, we computed the average distance, d, from the endpoints to the center and assigned every replica, i, a weight, $w_i(\boldsymbol{x})=exp(-(|\boldsymbol{x}_i-\boldsymbol{c}| / d)$ ** 2. New trajectories were iteratively seeded by extending $n_i \propto w_i/\sum_j w_j$ copies of each replica, keeping the total number of trajectories fixed to 24. With this approach, only replicas close to the center were extended and the ensemble of trajectories eventually diffused around a single equilibrated state (Fig. S1).

Collective Variable selection

We derived the CVs for the string method with swarms of trajectories in a data-driven manner from the swarm coordinates of the apo simulation's final iteration using the software demystifying [28]. As features, we chose inverse inter-residue C-α distances and filtered them to only include those which sampled values in the interval 6-8 Å, then used the features to train a Restricted Boltzmann Machine[40] with two hidden components. The results were averaged over 50 iterations. Only CVs with an estimated importance above 0.33 were included in the final set (Table S1). Every CV was first scaled unitless in order to keep all values between 0 and 1, then rescaled according to its importance, so that the restraining force and the drift in the swarms of trajectories method would better emphasize the conformational changes along important degrees of freedom. Finally, we derived a new string in the resulting CV space by interpolating between the restrained points of the apo simulation. All string simulations used these CVs except for the previously published apo and BI-167107 systems.

Free energy estimation

Free energy landscapes were estimated during post processing by discretizing a grid along a chosen set of variables and constructing a regularized transition matrix from the swarm trajectories' transitions between bins. We then derived the resulting free energy landscape from the stationary probability distribution of the transition matrix using Boltzmann inversion [30,31]. To estimate the convergence of the free energy landscapes, we applied a Bayesian Markov chain Monte Carlo approach to sample a distribution of transition matrices [41] and thereby the corresponding probability distributions.

Microswitch expectation values

To quantify a ligand's effect on different microswitches, we computed the expectation value of individual microswitches from the probability distribution. We evaluated the correlation between the expectation values of previously identified key microswitches and experimental measurements of cellular response to ligand binding [42] using linear regression in SciPy [43]. The p-value was estimated using a Wald Test with t-distribution based on a null hypothesis that there was no correlation between the computational and experimental values (i.e. that the correlation coefficient R was zero). We applied the derived linear relationship to predict the efficacy of the two ligands not part of the experimental dataset, BI-167107 and carazolol, based on their microswitch expectation values.

Supervised and unsupervised feature extraction and learning

We analyzed the trajectories from the last iteration of the adaptive state sampling protocol and the swarms of trajectories with various dimensionality reduction methods. To identify conformational differences induced by the ligands, we performed supervised and unsupervised feature extraction with the software demystifying [28] and projected important features onto snakeplot templates downloaded from GPCRDB [44]. As features, we used scaled inverse closest-heavy atom distances. Furthermore, we performed unsupervised dimensionality reduction in Scikit-learn [45] with Principal Component Analysis (PCA) [46], Multi-dimensional scaling (MDS) [47] with a Euclidean distance metric and T-distributed Stochastic Neighbor Embedding (t-SNE) [48], and projected the simulation snapshots onto the reduced feature spaces. We constructed a similarity by taking the average Eucledian distance between all snapshots in two classes, and normalized the class similarities between 0 and 1, with higher values representing similar classes (i.e. close frames).

Software code to reproduce the results in this manuscript is available online [39,49,50].

Results

Ligands control efficacy by reshaping microswitches' probability distributions

We derived the free energy landscape along the most probable activation pathway of the β2AR bound to different ligands (Fig. 1b) with 3-5.5 microsecond long swarms of trajectories simulations. Based on the short swarm trajectories we calculated free energies landscapes along different microswitches identified previously [31] (Fig. 2a, b and S2): the TM5 bulge, which is an indicator of contraction in the ligand binding site; the Connector ΔRMSD, a rotameric switch between residues I121(3.40) and F282(6.44) in the TM region; the TM6-TM3 distance showing the outwards movement of TM6, measured as the closest heavy atom distance between L272(6.34) and R131(3.50); and the Y-Y motif as the C-ζ distance between Y219(5.58) and Y326(7.53), which is an indicator of the twist of the NPxxY motif and a slight reorientation of TM5.

Fig. 2a-b and S2 show different projections of the free energy landscapes for the ligand-receptor complexes. The free energy landscapes along the Connector ΔRMSD (Fig. 2a) reveal two states. In agreement with what could be expected, agonists lower the relative free energy of the active state (R*) of this microswitch. For the TM5 bulge in the orthosteric site and the TM6-TM3 distance in the G protein binding site (Fig. 2b), the free energy landscapes show how the two domains are coupled. Agonists generally shift the energy balance to favor active-intermediate states.

It is not straightforward to predict a ligand's efficacy by mere visual inspection of it's free energy landscape, since the width and depth of the basins together determine the most probable state projected onto that space. To circumvent this issue, we computed expectation values of the metrics representing each of these microswitches and compared them to functional experiments

measuring the maximal G protein-mediated cAMP production (Emax) (Fig. 2c-f)[42]. The expectation values associated with the upper microswitches were strongly correlated to the experimental values, in particular the TM5 bulge (Fig. 2c; R=-0.95; p=0.013) and the Connector ΔRMSD (Fig. 2d; R=0.93; p=0.022). Based on the inferred linear relationships between these microswitches and the experimental values, we predicted BI-167107 to have a cAMP Emax value slightly higher than adrenaline and salmeterol, and carazolol to have an Emax similar to the values of the ligand free receptor and inverse agonist timolol (Fig. 2c, d). Although no experimental data was available for these two compounds, the results are inline with expectations; BI-167107 is indeed a known full agonist [2,51]. Encouragingly, these results indicate that our simulations accurately captured the relative stability of states, and should therefore be able to provide insights into how ligands with different efficacy profiles control the conformational ensemble of the receptor.

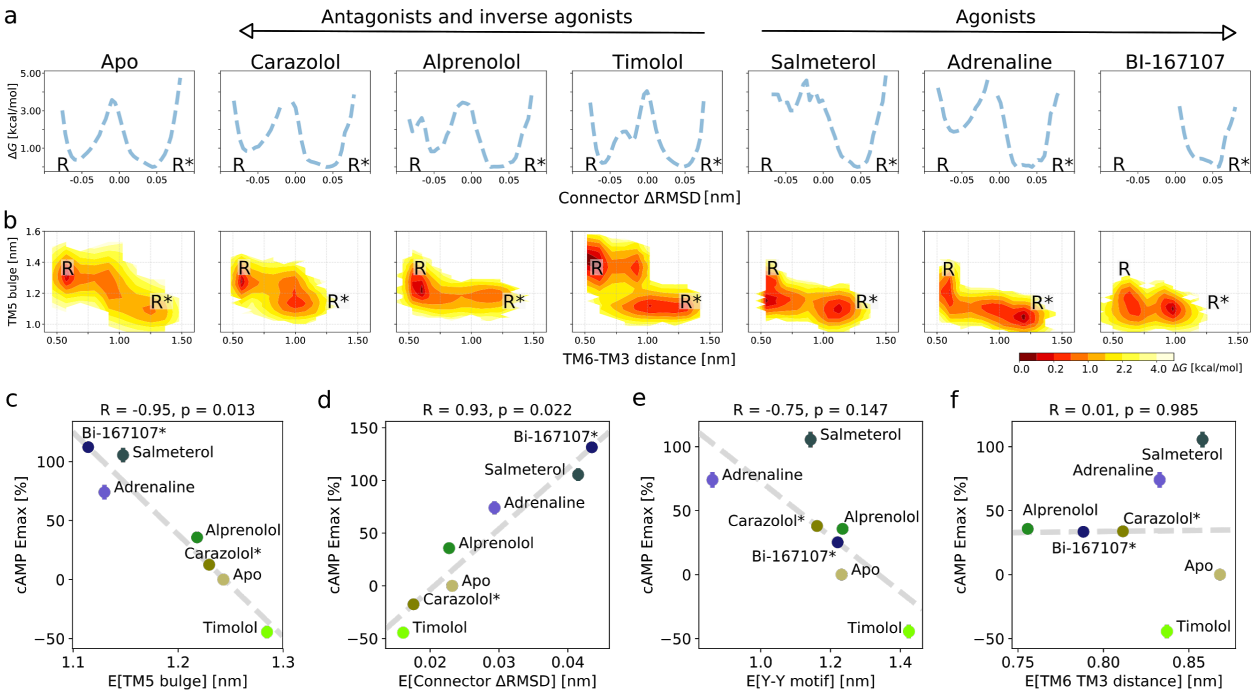

Figure 2. (a) Free energy landscapes for the different ligands projected along the connector ΔRMSD, located in the transmembrane region. The inactive and active states are marked by R and R* respectively. (b) Free energy projected onto the TM5 Bulge CV in the orthosteric site and the TM6-TM3 distance in the G protein binding site. (c)-(e) Correlation between experimental values of downstream cAMP signalling and the expectation value of different microswitches for the receptor bound to different ligands, (c) for the TM5 bulge microswitch (d) for the Connector ΔRMSD, (e) for the Y-Y motif microswitch, and (f) for the TM6-TM3 distance. The cAMP values for Carazolol and BI-167107 (marked with an asterisk) which were not available in the literature are inferred from the linear regression.

For the microswitches that were closer to the intracellular G protein binding site, the expectation values were less correlated to the experimental data, with a p-value indicating only loose coupling between the Y-Y motif and the cAMP response (Fig. 2e; R=-0.75; p=0.147) and no coupling between the TM6-TM3 distance and the cAMP response (Fig. 2f; R=-0.01; p=0.985). Note that the ligands stabilized active-like states of different TM6-TM3 displacements, which may weaken the correlation for this microswitch. For example, the apo receptor stabilized an active state with a larger TM6-TM3 closest heavy atom distance than that of the receptor bound to the agonist BI-167107 (Fig. 3b). Although the agonist's active state was of lower free energy, this conformational difference leads to a poor prediction of its efficacy (Fig. 3f). In general, since GPCRs only assume their fully active state in the presence of an intracellular binding partner—a condition not met in the simulations carried out in this work—we expect a loose coupling between intracellular microswitches and the cellular response.

Data-driven analysis reveals that ligands stabilize unique states.

To pinpoint the molecular basis of signalling, we reduced the high dimensional datasets to a more compressed representation using methods from machine learning (Fig. 3). We used three different methods considering as input all inter-residue contacts: PCA, MDS, and t-SNE. All these approaches were designed to find a low dimensional embedding of the high dimensional data, but differ in their underpinnings. PCA seeks a linear transformation of the input data into an orthogonal basis of Principal Components (PCs) and is designed to cover as much of the variance in the data as possible (Fig. 3a). MDS projects the high dimensional space into a low dimensional representation using a non-linear transformation which preserves the distance between points (Fig. 3b). T-SNE is a visualization technique which seeks to disentangle a high-dimensional data set by transforming it into a low-dimensional embedding where similar points are near each other and dissimilar objects have a high probability of being distant (Fig. 3c).

We evaluated two datasets: the single equilibrated active-like state ensemble and the swarm trajectories from the final iteration of the converged string. For the latter, which represent the converged activation pathways, the dimensionality reduction techniques created embeddings which separated snapshots by their progression along the activation path (Fig. S3). This is expected since unsupervised dimensionality reduction methods tend to emphasize large scale amplitude motions of activation, such as the displacement of TM6 [28]. We thus confirm that this feature is shared amongst all activation pathways, regardless of which ligand the receptor was bound to.

Whereas the activation path ensembles contained inactive, intermediate and active-like states for every ligand-receptor complex, the individual state simulations revealed that the active-like states sampled in presence of different ligands differed substantially. Indeed, after eight iterations with an accumulated simulation time of 1.4 microseconds per ligand (see Methods), the method for finding single equilibrated states generated trajectories which diffused around the most stabilized state kinetically accessible from the starting structure. Unexpectedly, the details of those varied amongst ligands (Fig. 3a-c). Agonist-bound conformations tended to be

grouped together for all three dimensionality reduction methods, but each of them generally also led to a distinct conformational ensemble. The simulation snapshots with the agonist adrenaline bound were generally close to the full agonist BI-167107 and the partial agonist salmeterol. For the other ligands, the receptor explored a different conformational space than with agonists, but the ensembles were more diverse. In agreement with the projections, the similarity matrix based on the average distance between snapshots in the full feature space (Fig. 3d) showed that agonists and non agonists stabilize significantly different states.

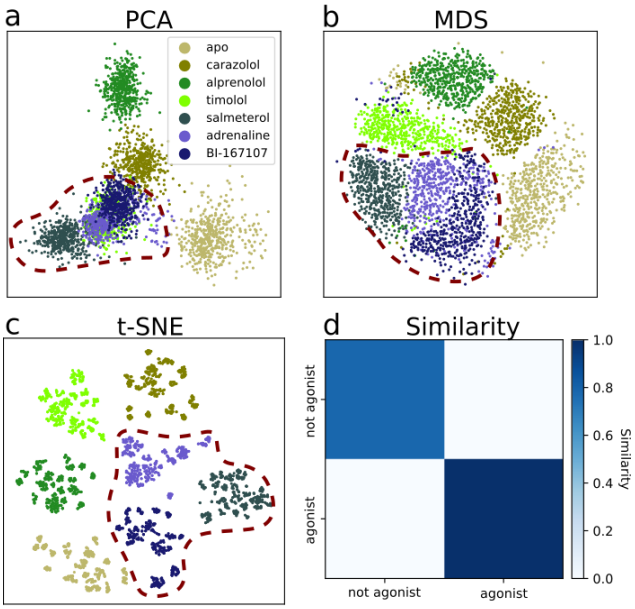

Figure 3. Dimensionality reduction techniques applied to the active-like simulation ensembles. Each point represents a simulation snapshot, colored according to the ligand bound to the receptor. Red dashed lines highlight regions where agonists cluster. The features are computed as the inverse closest heavy atom distances between residues. (a) Principal Component Analysis (PCA) projection onto the first two Principal Components (PCs), (b) Multi-dimensional scaling (MDS) and (c) T-distributed Stochastic Neighbor Embedding (t-SNE). (d) The similarity between conformations sampled when agonist (BI-167107, adrenaline and salmeterol) and non-agonist ligands are bound, measured as the average distance between configurations in the full feature space.

Using PCA, we note that the first two PCs are not sufficient to completely separate the dataset according to the ligands present (Fig. 3a), but including more PCs in the projection led to a satisfactory separation (Fig. S4). The non-linear methods separated the classes well in two dimensions. As expected, a few points deviated from the other snapshots in the same class due to sampling slightly outside defined free energy basins. We also note that although t-SNE generates an embedding with perfect separation between classes, the micro-clusters depend on parametrization of the method and their exact placement varies with random instantiations [52].

To summarize, analysis of the simulations by machine learning shows that ligands share many overall features of activation, but stabilize unique local states. Together with the results from the previous section, our findings support that ligands control the relative time a receptor spends in active-like states, and induces small conformational state-specific signatures throughout the protein.

Ligands control residues near the G protein binding site.

To characterize molecular differences between the active-like states controlled by the different ligands, we applied supervised learning on our dataset. With this approach, we derived features discriminating between the classes and thereby identified residues which may be important for activation. As in the previous section, we considered inter-residue distances as input features. Importance profiles were computed for discriminating between agonists and non-agonists (Fig. 4a) and identifying features to distinguish all ligands from each other (Fig. 4b) using a symmetrized version of the Kullback–Leibler (KL) divergence [28,53]. As a control, we also evaluated the important features learned by a Random Forest Classifier (SI Text). Remarkably, both data-driven methods identified allosteric sites as the most important regions for classification: the NPxxY motif and the intracellular part of TM6 in the intracellular binding site (Fig. 4a,b and S5). These results underpin experimental evidence that differences in this region are related to biased signalling [9,10,15,24,54]. The TM5 bulge was particularly important for discriminating between agonists and non-agonists (Fig. 4a). This region does not show up as important when differentiating between all ligands (Fig. 4b), which means that the TM5 conformations within the two groups of ligands were so similar that this region could not be used to, for example, discriminate agonists from each other. The NPxxY motif, on the other hand, assumed a unique conformation for every individual ligand (Fig. 4b).

Our finding that the ligands share features of activation and potentially have overlapping states (Fig. 3b) suggested that unsupervised learning could capture important characteristics of receptor activation [28]. We therefore applied PCA on the swarms of trajectories datasets representing the activation ensemble (Fig. 4c). This analysis also identified TM6 and TM7 near the G-protein binding site as particularly important (Fig. 4c and S7), adding further support for the importance of these microswitches for activation.

Taken together, we arrived at a conceptual model to describe how different ligands control the G protein binding site (Fig. 4d). Ligands exercise direct control on TM5, where agonists and non-agonists stabilize different states. In turn, residues approximately one helical turn below the orthosteric site, including the connector region which was identified as a good predictor of downstream response (Fig. 2d), couple to the conformations in the orthosteric site. This leads to distinct interaction patterns between TM6 and TM7 in the transmembrane domain. The importance of TM7 is further enhanced by direct ligand interactions, generating a variety of ligand-specific NPxxY motifs. By favoring distinct TM7 states and modulating the probability for TM6 to be in an active conformation, ligands hence control the G protein binding site.

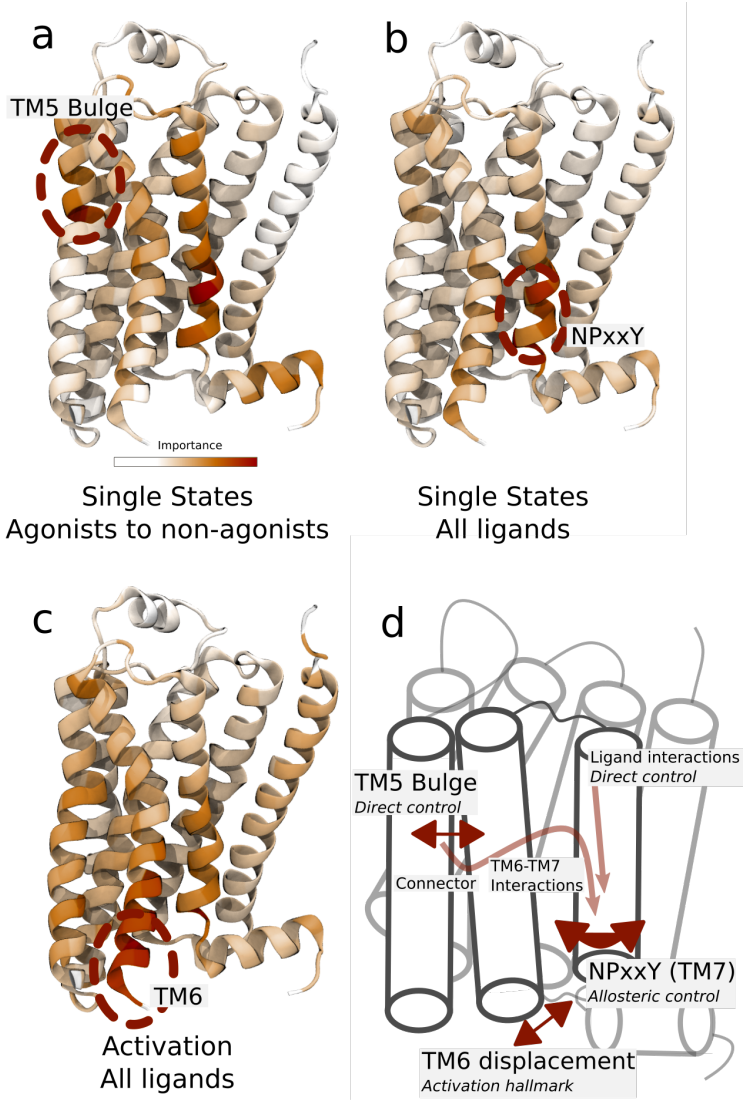

Figure 4. (a)-(b) Residues identified to be important for classification of the equilibrated active-like states, seen from the transmembrane side. Importance was derived by computing the Kullback–Leibler (KL) divergence along all features, followed by averaging per residue. (a) Comparison of agonists to non-agonists. One signaling hotspot is located at the TM5 Bulge domain and another on TM7 near the NPxxY motif. (b) Important residues to discriminate between all ligands. The main signaling hotspot is located near the NPxxY motif. (c) Important residues for the activation ensemble from the swarms of trajectories method, extracted with PCA. The importance per feature was computed as the product of the PC's weights and the PCs projection onto the input feature. The intracellular end of TM6, which undergoes significant conformational change upon activation, is marked as important. Inverse closest heavy atom distances were used as input features in all figures. (d) Conceptual model describing allosteric transmission between the signalling hotspots. Ligands exercise direct control of TM5, which is stabilized in different states by agonists and non-agonists. In turn, residues

approximately one helical turn below the orthosteric site, including the connector region, couple to the conformations in the orthosteric site. This leads to distinct interaction patterns between TM6 and TM7 in the transmembrane domain. The importance of the TM7 is further enhanced by direct ligand interactions. By favoring distinct conformational TM7 states and modulating the probability for TM6 to be in an active position, ligands control the G protein binding site. A snakeplot representation of the same data is shown in Fig. S6.

Molecular basis for allosteric transmission from the orthosteric to the intracellular binding site

To explore the atomistic basis of our conceptual model (Fig. 4), we systematically inspected the most important features connecting the two hotspots near the TM5 bulge and the NPxxY motif. With this approach we identified inter-residue distances of high KL divergence which best highlighted the differences between the equilibrated active-like states with ligands. Although the identified residue-pairs do not necessarily reflect the causality of molecular interactions driving the receptor's conformational change, key features of activation were captured by this automated approach.

We first identified features shared between agonists near the orthosteric site. We found that V206(5.46) could form van der Waals interactions with T118(3.37) only in the presence of agonists (Fig. 5a and b). This interaction is probably caused by hydrogen bonding between S207(5.46), which defines the TM5 bulge microswitch, and the ligand. In the transmembrane domain, agonists induced a contraction between L284(6.46) and F321(7.48) (Fig. 5a and b) compared to non-agonists. Both these residues face the lipid bilayer and are only weakly interacting in the simulations with agonists bound, but are located in a region that undergoes significant change upon activation. L284(6.46) is located just above the region where TM6 kinks upon activation. The identified feature connects the binding site and PIF motif to the NPxxY motif. F282(6.44) of the PIF motif neighbors L284(6.46). T118(3.37), which was identified as important near the TM5 bulge (Fig. 5b), is only one helix turn above I121(3.40) of the PIF motif in the Connector region. Thus, the Connector region is likely a driving factor behind the allosteric transmission between the two binding sites, which influences the region around F321(7.48). F321(7.48) is located right above the NPxxY motif, which undergoes a twist upon activation, and belongs to the important hotspot on TM7. In this region we found that the backbone carbonyl of S319(7.46) formed a hydrogen bond with the sidechain of N51(1.50) on TM1 for agonists, whereas this interaction was destabilized for the other ligands (Fig. 5d). N51(1.50) is one of the most conserved residues across class A GPCRs [44] and supports a stabilizing water network together with D79(2.50) and Y326(7.53) in the inactive receptor [3,55]. Thus, this agonist-specific interaction likely promotes dehydration of the water-filled cavity around conserved residue D79(2.50) and stabilizes a twist of the NPxxY motif. Overall, agonists favored contractions between local inter-residue distances compared to non-agonists. By inspecting the most substantially changing large scale distances, we also identified a contraction of the entire protein for agonist-bound receptors, as reflected by the decrease in distance between S203(5.43) and E338(8.56) on helix 8 (H8) and between S207(5.46) and V307(7.33) across the orthosteric binding site (Fig. 5a and c).

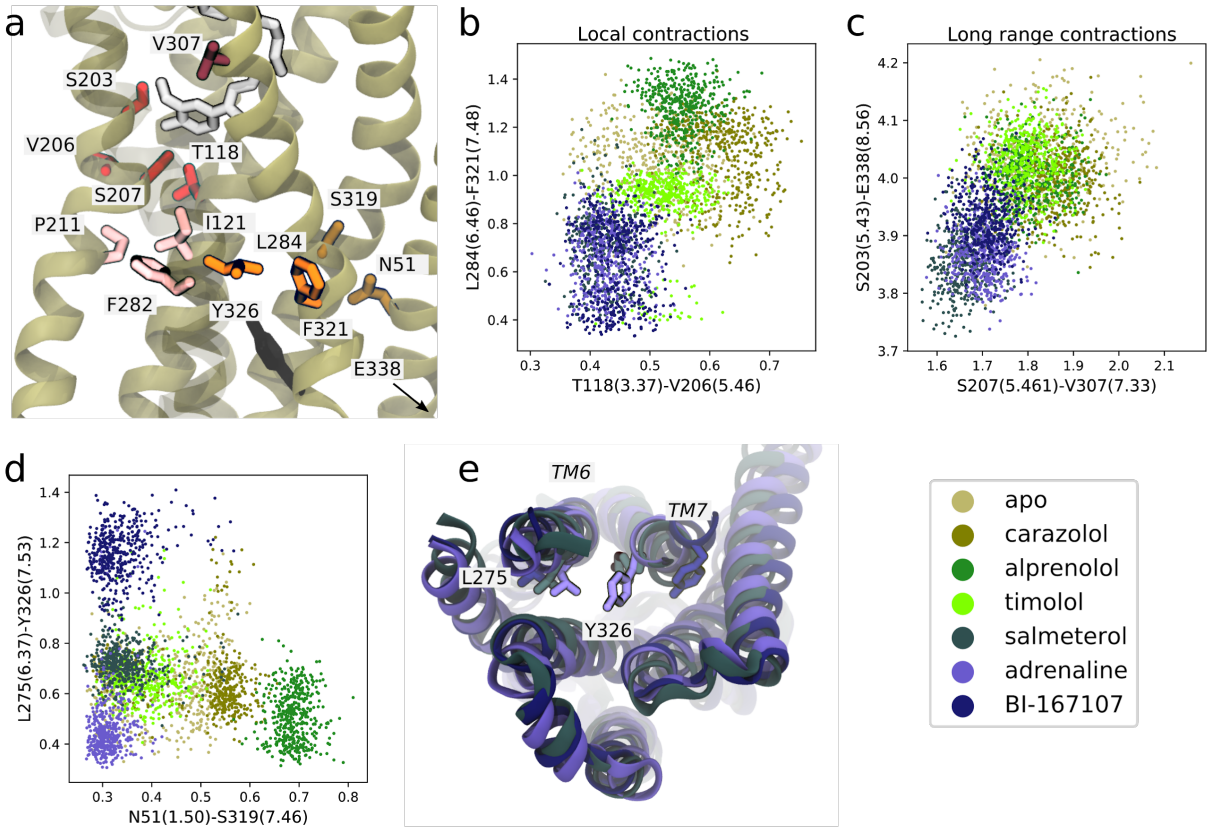

Figure 5. (a) Molecular basis for agonists' allosteric control of receptor activation. S203(5.43) and S207(5.46) (red sticks) are part of the TM5 bulge, which form direct contacts with the ligand. V206(5.46) forms van der Waals interactions with T118(3.37) (red), which is located above I121(3.40) of the PIF motif in the Connector region (pink sticks). TM6 and TM7, highlighted as L284(6.46) (orange sticks) close to F282(6.44) of the PIF motif and F321(7.48) (orange sticks) above the NPxxY motif, move closer together in the presence of agonists. Half a helix turn above F321(7.48), S319(7.46) forms a backbone interaction with the sidechain of N51(1.50) for agonist-bound receptors, whereas water molecules interact with these residues for non-agonists. Together with TM7-ligand contacts in the orthosteric site, these interaction pathways stabilize the second important hotspot on TM7 close to the NPxxY motif (Y326(7.53) shown in black). (b) The T118(3.37)-V206(5.46) distance near the orthosteric site against the L284(6.46)-F321(7.48) distance in the transmembrane region. Agonists contract both of these regions. (c) The distance across the orthosteric site between S207(5.461) and V307(7.33) (dark red in (a)) against the S203(5.43)-E338(8.56) distance across the transmembrane domain. Agonists stabilize more compact receptor conformations. (d) The N51(1.50)-S319(7.46) distance against the L275(6.37)-Y326(7.53) distance. Agonists share the common feature of stabilizing the N51(1.50)-S319(7.46) backbone interaction, but form different NPxxY orientations, shown as the distance from Y326(7.53) to L275(6.37). (e) The three agonists stabilize slightly different TM6 and TM7 orientations, here shown as the distance between L275(6.37) and Y326(7.53). Adrenaline (purple) induces an active-like NPxxY motif

whereas BI-167107 (dark blue) stabilizes an inactive-like motif. The salmeterol-bound receptor (slate gray) adopts a distinct Y326(7.53) orientation.

Near the NPxxY motif, we found that the agonists stabilized different TM6 and TM7 orientations (Fig. 5d and e). Adrenaline favored the most active-like NPxxY motif, which was also maintained throughout its activation path (Fig. S8), with Y326(7.53) closer to L275(6.37) (Fig. 5d). Salmeterol stabilized a distinct NPxxY conformation, which was also observed in the activation path ensembles (Fig. S9), where Y326(7.53) underwent a sidechain rotation, bringing its head group further into the interface between TM6 and TM7. This pose is reminiscent of rotameric conformations suggested by ^{19}F NMR studies on the β1-adrenergic receptor[9]. BI-167107 stabilized an inactive-like NPxxY motif with Y326(7.53) twisting away from L275(6.37) (Fig. 5d). These results reveal that, against our expectations, protein residues far away from the orthosteric site can experience the largest ligand-induced conformational heterogeneity.

Discussion

The derivation of free energy landscapes and the corresponding microswitch expectation values provide a tool for estimating the stability of activation states, and thus also the relative efficacy of different ligands. Given the high correlations between microswitch expectation value and experimental data, we anticipate that microswitches located in the orthosteric and connector regions can be used for future predictions of ligand efficacy. Additionally, an advantage of the methods used in this study is that results were derived from several simulation replicas, which reduces the statistical error related to the stochastic nature of MD simulations on short time scales.

Following the progress in GPCR research, it has become evident that a simple two-state model of activation is an oversimplification with considerable limitations. To take biased and partial agonism into account, there is a need for a more comprehensive model. Many ligands have been characterized as full, partial or biased agonists [42]. However, the molecular mechanisms which transmit this allosteric communication across the cell membrane remain elusive. Researchers have successfully managed to discriminate between arrestin and G protein bias using spectroscopic probes [10,15], but the conformational changes and dynamics of many microswitches cannot be captured in a single measurement. MD has the potential to provide additional insights [56]. Whereas the free energy landscapes clearly reveal a shift in the most stabilized microswitch values upon ligand binding (Fig. 2a, b and S2), a direct comparison between free energy profiles may be misleading if the other orthogonal microswitches are ignored. For example, except for a small TM6-TM3 offset in the BI-167107 (agonist)- and alprenolol (antagonist) -bound landscapes (Fig. 2b and S2), the basins are located in similar positions, signifying similar outwards movement of TM6 as the receptor activates. However, such projections onto single variables can be misleading and obscure major differences in other microswitches. To address this limitation, we used data-driven analysis methods, which are better suited for handling high-dimensional data than mere visual inspection, and found that there were indeed significant conformational differences in the states stabilized by the different ligands. Remarkably, this protocol automatically identified both receptor-specific and conserved motifs considered to be of significance for GPCR activation, such the TM5 bulge and the NPxxY

motif, as important. The different dimensionality reduction techniques found similar partitioning of the data, which strongly indicates that the results are not due to fortunate parameter tuning or method choice. One of the remaining enigmas in GPCR research is to understand how the same overall activation mechanism can be conserved in spite of the fact that very different ligands are recognized by the family. Our machine learning-inspired data analysis protocol provides an unbiased approach to identify key features of activation for different receptor types.

In a reversible system, extracellular microswitches might be loosely coupled to intracellular binding similar to how the intracellular microswitches couple to ligand binding. For pharmacological purposes, the ligand control of receptor activation is of great interest. The conformational selection of intracellular binding partners enables construction of molecules or nanobodies that have a high binding affinity only in combination with a specific ligand [57], which has been utilized in β2AR structure determination [2,4,6,58]. Similarly, a well-designed nanobody or allosteric modulator could enhance or block binding of a specific compound [58]. However, in this study, the ligands were bound to the receptor throughout the simulations and no intracellular binding partner was considered. GPCRs only assume their fully active state in the presence of an intracellular binding partner [14,25,59]. The native state intracellular binding partner-free receptor undergoes basal activity, where it fluctuates between active-like and inactive states [14,24], which can be inferred from the relatively low free energy difference between the basins in the energy landscapes (Fig. 2a-b). While extracellular and transmembrane microswitches only require an agonist to maintain their active state conformations, intracellular microswitches interact with an intracellular binding partner directly and are more affected by its absence, which is likely why correlations between microswitch expectation value and Emax is less good for the intracellular microswitches.

The ligand being bound to the receptor throughout the simulation corresponds to having a very high concentration of compound present in the extracellular solvent, justifying comparing our results to the experimental maximum attainable response, Emax, of the dose response curves. Other important aspects of receptor-ligand interactions, such as identifying the binding pose of novel compounds or estimating a ligand's binding affinity, cannot be estimated with this approach. There are other, complementary, techniques such as Free Energy Perturbation methods [60], which can be used to rigorously estimate binding strength. Combined with enhanced sampling MD, we move towards having a complete toolkit for *in silico* drug design for development of high affinity GPCR ligands with a specific efficacy.

Despite the increasing number of β2AR structures, the receptor has not been solved in a form bound to arrestin. However, conformations observed in our simulations share properties with other receptors' arrest-signalling states. For example, the supervised learning methods identified the C-terminal and H8 below the NPxxY motif as relatively important (Fig. 3a, S6a, S7 and S10). This domain is known to be stabilized in ligand-dependent states for the angiotensin II type 1 receptor [54,61]. While the reorganization of H8 may be a secondary effect due to modulation of the NPxxY motif, this region could be important for dimerization or arrestin recruitment [62,63]. Salmeterol's distinct NPxxY state (Fig. S9) only formed in combination with a lost interaction between the salmeterol and S207(5.46) and S203(5.43) (Fig. S9d), which is remarkable since the two binding sites are believed to be only loosely coupled [19,31]. A similar

phenomenon has been reported in a recent structure of β1AR [64], where the corresponding serines in the orthosteric site experienced weakened interactions to the biased agonist formoterol. Seeing that Y(5.53) is known to form contacts with arrestin for β1AR [64], our derived β2AR conformation may have biased signaling properties.

Our results show that the activation pathways as well as the stabilized states are significantly altered upon ligand binding, and that ligands with shared efficacy profiles generate similar, albeit not identical, ensembles of states. It therefore cannot be taken for granted that two ligands which lead to a similar downstream response necessarily stabilize the same receptor conformations. As we increased our dataset with more compounds in this study, similarities and differences between different compound classes emerged. The results in this study provide a good starting point for further analysis and allowed us to catch a glimpse of the complexity underlying receptor signalling. A thorough quantification of biased and partial agonism will require us to further extend our dataset with ligands that stimulate various signalling pathways to different extents.

Conclusion

In this study, we derived the activation free energy of the β2AR bound to ligands with different efficacy profiles using enhanced sampling MD simulations. We found a strong correlation between cellular response and the computed expectation values of the upper and transmembrane microswitches, which demonstrates that our approach holds predictive power. Not only did the results show how ligands control the population of states, they also modulate the conformational ensemble of states by flipping important allosteric microswitches in the vicinity of the G protein binding site. By inspecting the contacts formed for agonists and non-agonist ligands, we identified an allosteric pathway between the two binding sites and a large heterogeneity of TM7 states. Our results show how MD simulations in combination with data-driven analysis provides the means for generating a comprehensive view of the complex signalling landscape of GPCRs. We anticipate that our protocol can be used together with other computational methods to understand GPCR signalling at the molecular level and provide insights that make it possible to design ligands with specific efficacy profiles.

Acknowledgements

This work was supported by grants from the Gustafsson Foundation and Science for Life Laboratory to JC and LD. The work was also supported by grants from the Swedish Research Council (2017-4676) and the Swedish strategic research program eSSENCE to JC. The simulations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC) at PDC Centre for High Performance Computing (PDC-HPC).

Abbreviations

GPCRs, G protein-coupled receptors; MD, Molecular dynamics; β2AR, β2 adrenergic receptor; β1AR, β1 adrenergic receptor; TM1-7, Transmembrane helix 1-7; H8, Helix 8; cAMP, cyclic adenosine monophosphate; CV, Collective variable; PCA, Principal component analysis; PC,

principal component; MDS, Multi-dimensional scaling; t-SNE, T-distributed stochastic neighbor embedding; RMSD, Root-mean-square deviation; KL divergence, Kullback–Leibler (KL) divergence

Supporting Information Available

The Supporting Information is available free of charge online.

- Supplementary Text, Tables S1–S2, and Figures S1–S10 (PDF)

References

Bibliography

(1) Hauser, A. S., Attwood, M. M., Rask-Andersen, M., Schiöth, H. B., and Gloriam, D. E. (2017) Trends in GPCR drug discovery: new agents, targets and indications. *Nat. Rev. Drug Discov. 16*, 829–842.
(2) Rasmussen, S. G. F., Choi, H.-J., Fung, J. J., Pardon, E., Casarosa, P., Chae, P. S., Devree, B. T., Rosenbaum, D. M., Thian, F. S., Kobilka, T. S., Schnapp, A., Konetzki, I., Sunahara, R. K., Gellman, S. H., Pautsch, A., Steyaert, J., Weis, W. I., and Kobilka, B. K. (2011) Structure of a nanobody-stabilized active state of the β(2) adrenoceptor. *Nature 469*, 175–180.
(3) Cherezov, V., Rosenbaum, D. M., Hanson, M. A., Rasmussen, S. G. F., Thian, F. S., Kobilka, T. S., Choi, H.-J., Kuhn, P., Weis, W. I., Kobilka, B. K., and Stevens, R. C. (2007) High-resolution crystal structure of an engineered human beta2-adrenergic G protein-coupled receptor. *Science 318*, 1258–1265.
(4) Ring, A. M., Manglik, A., Kruse, A. C., Enos, M. D., Weis, W. I., Garcia, K. C., and Kobilka, B. K. (2013) Adrenaline-activated structure of β2-adrenoceptor stabilized by an engineered nanobody. *Nature 502*, 575–579.
(5) Wacker, D., Fenalti, G., Brown, M. A., Katritch, V., Abagyan, R., Cherezov, V., and Stevens, R. C. (2010) Conserved binding mode of human beta2 adrenergic receptor inverse agonists and antagonist revealed by X-ray crystallography. *J. Am. Chem. Soc. 132*, 11443–11445.
(6) Masureel, M., Zou, Y., Picard, L.-P., van der Westhuizen, E., Mahoney, J. P., Rodrigues, J. P. G. L. M., Mildorf, T. J., Dror, R. O., Shaw, D. E., Bouvier, M., Pardon, E., Steyaert, J., Sunahara, R. K., Weis, W. I., Zhang, C., and Kobilka, B. K. (2018) Structural insights into binding specificity, efficacy and bias of a β2AR partial agonist. *Nat. Chem. Biol. 14*, 1059–1066.
(7) Hanson, M. A., Cherezov, V., Griffith, M. T., Roth, C. B., Jaakola, V.-P., Chien, E. Y. T., Velasquez, J., Kuhn, P., and Stevens, R. C. (2008) A specific cholesterol binding site is established by the 2.8 A structure of the human beta2-adrenergic receptor. *Structure 16*, 897–905.
(8) Manglik, A., and Kruse, A. C. (2017) Structural Basis for G Protein-Coupled Receptor Activation. *Biochemistry 56*, 5628–5634.

(9) Frei, J. N., Broadhurst, R. W., Bostock, M. J., Solt, A., Jones, A. J. Y., Gabriel, F., Tandale, A., Shrestha, B., and Nietlispach, D. (2020) Conformational plasticity of ligand-bound and ternary GPCR complexes studied by 19F NMR of the β1-adrenergic receptor. *Nat. Commun. 11*, 669.
(10) Liu, J. J., Horst, R., Katritch, V., Stevens, R. C., and Wüthrich, K. (2012) Biased signaling pathways in β2-adrenergic receptor characterized by 19F-NMR. *Science 335*, 1106–1110.
(11) Ballesteros, J. A., and Weinstein, H. (1995) [19] Integrated methods for the construction of three-dimensional models and computational probing of structure-function relations in G protein-coupled receptors, in *Receptor Molecular Biology*, pp 366–428. Elsevier.
(12) Jean-Charles, P.-Y., Kaur, S., and Shenoy, S. K. (2017) G Protein-Coupled Receptor Signaling Through β-Arrestin-Dependent Mechanisms. *J. Cardiovasc. Pharmacol. 70*, 142–158.
(13) Rasmussen, S. G. F., DeVree, B. T., Zou, Y., Kruse, A. C., Chung, K. Y., Kobilka, T. S., Thian, F. S., Chae, P. S., Pardon, E., Calinski, D., Mathiesen, J. M., Shah, S. T. A., Lyons, J. A., Caffrey, M., Gellman, S. H., Steyaert, J., Skiniotis, G., Weis, W. I., Sunahara, R. K., and Kobilka, B. K. (2011) Crystal structure of the β2 adrenergic receptor-Gs protein complex. *Nature 477*, 549–555.
(14) Gregorio, G. G., Masureel, M., Hilger, D., Terry, D. S., Juette, M., Zhao, H., Zhou, Z., Perez-Aguilar, J. M., Hauge, M., Mathiasen, S., Javitch, J. A., Weinstein, H., Kobilka, B. K., and Blanchard, S. C. (2017) Single-molecule analysis of ligand efficacy in β2AR-G-protein activation. *Nature 547*, 68–73.
(15) Lamichhane, R., Liu, J. J., White, K. L., Katritch, V., Stevens, R. C., Wüthrich, K., and Millar, D. P. (2020) Biased Signaling of the G-Protein-Coupled Receptor β2AR Is Governed by Conformational Exchange Kinetics. *Structure*.
(16) Kofuku, Y., Ueda, T., Okude, J., Shiraishi, Y., Kondo, K., Maeda, M., Tsujishita, H., and Shimada, I. (2012) Efficacy of the β_2-adrenergic receptor is determined by conformational equilibrium in the transmembrane region. *Nat. Commun. 3*, 1045.
(17) Imai, S., Yokomizo, T., Kofuku, Y., Shiraishi, Y., Ueda, T., and Shimada, I. (2020) Structural equilibrium underlying ligand-dependent activation of β2-adrenoreceptor. *Nat. Chem. Biol. 16*, 430–439.
(18) Tikhonova, I. G., Selvam, B., Ivetac, A., Wereszczynski, J., and McCammon, J. A. (2013) Simulations of biased agonists in the β(2) adrenergic receptor with accelerated molecular dynamics. *Biochemistry 52*, 5593–5603.
(19) Dror, R. O., Arlow, D. H., Maragakis, P., Mildorf, T. J., Pan, A. C., Xu, H., Borhani, D. W., and Shaw, D. E. (2011) Activation mechanism of the β2-adrenergic receptor. *Proc Natl Acad Sci USA 108*, 18684–18689.
(20) Kohlhoff, K. J., Shukla, D., Lawrenz, M., Bowman, G. R., Konerding, D. E., Belov, D., Altman, R. B., and Pande, V. S. (2014) Cloud-based simulations on Google Exacycle reveal ligand modulation of GPCR activation pathways. *Nat. Chem. 6*, 15–21.
(21) Chen, K.-Y. M., Keri, D., and Barth, P. (2020) Computational design of G Protein-Coupled Receptor allosteric signal transductions. *Nat. Chem. Biol. 16*, 77–86.
(22) Chan, H. C. S., Filipek, S., and Yuan, S. (2016) The Principles of Ligand Specificity on beta-2-adrenergic receptor. *Sci. Rep. 6*, 34736.
(23) Weis, W. I., and Kobilka, B. K. (2018) The Molecular Basis of G Protein-Coupled Receptor Activation. *Annu. Rev. Biochem. 87*, 897–919.

(24) Lamichhane, R., Liu, J. J., Pljevaljcic, G., White, K. L., van der Schans, E., Katritch, V., Stevens, R. C., Wüthrich, K., and Millar, D. P. (2015) Single-molecule view of basal activity and activation mechanisms of the G protein-coupled receptor β2AR. *Proc Natl Acad Sci USA 112*, 14254–14259.
(25) Manglik, A., Kim, T. H., Masureel, M., Altenbach, C., Yang, Z., Hilger, D., Lerch, M. T., Kobilka, T. S., Thian, F. S., Hubbell, W. L., Prosser, R. S., and Kobilka, B. K. (2015) Structural Insights into the Dynamic Process of β2-Adrenergic Receptor Signaling. *Cell 161*, 1101–1111.
(26) Picard, L.-P., Schonegge, A.-M., and Bouvier, M. (2019) Structural Insight into G Protein-Coupled Receptor Signaling Efficacy and Bias between Gs and β-Arrestin. *ACS Pharmacol. Transl. Sci. 2*, 148–154.
(27) Harpole, T. J., and Delemotte, L. (2018) Conformational landscapes of membrane proteins delineated by enhanced sampling molecular dynamics simulations. *Biochim. Biophys. Acta Biomembr. 1860*, 909–926.
(28) Fleetwood, O., Kasimova, M. A., Westerlund, A. M., and Delemotte, L. (2020) Molecular Insights from Conformational Ensembles via Machine Learning. *Biophys. J. 118*, 765–780.
(29) Hu, X., Wang, Y., Hunkele, A., Provasi, D., Pasternak, G. W., and Filizola, M. (2019) Kinetic and thermodynamic insights into sodium ion translocation through the μ-opioid receptor from molecular dynamics and machine learning analysis. *PLoS Comput. Biol. 15*, e1006689.
(30) Pan, A. C., Sezer, D., and Roux, B. (2008) Finding transition pathways using the string method with swarms of trajectories. *J. Phys. Chem. B 112*, 3432–3440.
(31) Fleetwood, O., Matricon, P., Carlsson, J., and Delemotte, L. (2020) Energy Landscapes Reveal Agonist Control of G Protein-Coupled Receptor Activation via Microswitches. *Biochemistry 59*, 880–891.
(32) Woo, T. M., and Robinson, M. V. (2015) Pharmacotherapeutics for advanced practice nurse prescribers 4th ed. F. A. Davis Company.
(33) Lee, J., Cheng, X., Swails, J. M., Yeom, M. S., Eastman, P. K., Lemkul, J. A., Wei, S., Buckner, J., Jeong, J. C., Qi, Y., Jo, S., Pande, V. S., Case, D. A., Brooks, C. L., MacKerell, A. D., Klauda, J. B., and Im, W. (2016) CHARMM-GUI Input Generator for NAMD, GROMACS, AMBER, OpenMM, and CHARMM/OpenMM Simulations Using the CHARMM36 Additive Force Field. *J. Chem. Theory Comput. 12*, 405–413.
(34) Huang, J., Rauscher, S., Nawrocki, G., Ran, T., Feig, M., de Groot, B. L., Grubmüller, H., and MacKerell, A. D. (2017) CHARMM36m: an improved force field for folded and intrinsically disordered proteins. *Nat. Methods 14*, 71–73.
(35) Klauda, J. B., Venable, R. M., Freites, J. A., O'Connor, J. W., Tobias, D. J., Mondragon-Ramirez, C., Vorobyov, I., MacKerell, A. D., and Pastor, R. W. (2010) Update of the CHARMM all-atom additive force field for lipids: validation on six lipid types. *J. Phys. Chem. B 114*, 7830–7843.
(36) Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W., and Klein, M. L. (1983) Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys. 79*, 926.
(37) Abraham, M. J., Murtola, T., Schulz, R., Páll, S., Smith, J. C., Hess, B., and Lindahl, E. (2015) GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX 1–2*, 19–25.
(38) Tribello, G. A., Bonomi, M., Branduardi, D., Camilloni, C., and Bussi, G. (2014) PLUMED 2: New feathers for an old bird. *Comput. Phys. Commun. 185*, 604–613.

(39) Fleetwood, O. (2019) delemottelab/gpcr-string-method-2019 initial-submission. *Zenodo*.
(40) Smolensky, P. (1986) Information Processing in Dynamical Systems: Foundations of Harmony Theory.
(41) Harrigan, M. P., Sultan, M. M., Hernández, C. X., Husic, B. E., Eastman, P., Schwantes, C. R., Beauchamp, K. A., McGibbon, R. T., and Pande, V. S. (2017) Msmbuilder: statistical models for biomolecular dynamics. *Biophys. J. 112*, 10–15.
(42) van der Westhuizen, E. T., Breton, B., Christopoulos, A., and Bouvier, M. (2014) Quantification of ligand bias for clinically relevant β2-adrenergic receptor ligands: implications for drug taxonomy. *Mol. Pharmacol. 85*, 492–509.
(43) Millman, K. J., and Aivazis, M. (2011) Python for Scientists and Engineers. *Comput. Sci. Eng. 13*, 9–12.
(44) Isberg, V., Vroling, B., van der Kant, R., Li, K., Vriend, G., and Gloriam, D. (2014) GPCRDB: an information system for G protein-coupled receptors. *Nucleic Acids Res. 42*, D422-5.
(45) Pedregosa, F., Varoquaux, G., and Gramfort, A. (2011) Scikit-learn: Machine learning in Python. *… of machine learning …*.
(46) Tipping, M. E., and Bishop, C. M. (1999) Probabilistic Principal Component Analysis. *J. Royal Statistical Soc. B 61*, 611–622.
(47) (2005) Modern Multidimensional Scaling. Springer New York, New York, NY.
(48) van der Maaten, L., and Hinton, G. (2008) Visualizing Data using t-SNE. *J Mach Learn Res 9*.
(49) Fleetwood, O. (2019) delemottelab/demystifying: Updated benchmarks. *Zenodo*.
(50) Fleetwood, O. (2020) delemottelab/state-sampling initial-release. Computer software, zenodo.
(51) Ueda, T., Kofuku, Y., Okude, J., Imai, S., Shiraishi, Y., and Shimada, I. (2019) Function-related conformational dynamics of G protein-coupled receptors revealed by NMR. *Biophys. Rev. 11*, 409–418.
(52) Schubert, E. (2017, February 28) interpretation - Clustering on the output of t-SNE - Cross Validated. *interpretation - Clustering on the output of t-SNE - Cross Validated*.
(53) Kullback, S., and Leibler, R. A. (1951) On information and sufficiency. *Ann. Math. Statist. 22*, 79–86.
(54) Suomivuori, C.-M., Latorraca, N. R., Wingler, L. M., Eismann, S., King, M. C., Kleinhenz, A. L. W., Skiba, M. A., Staus, D. P., Kruse, A. C., Lefkowitz, R. J., and Dror, R. O. (2020) Molecular mechanism of biased signaling in a prototypical G protein-coupled receptor. *Science 367*, 881–887.
(55) Venkatakrishnan, A. J., Ma, A. K., Fonseca, R., Latorraca, N. R., Kelly, B., Betz, R. M., Asawa, C., Kobilka, B. K., and Dror, R. O. (2019) Diverse GPCRs exhibit conserved water networks for stabilization and activation. *Proc Natl Acad Sci USA 116*, 3288–3293.
(56) Lamim Ribeiro, J. M., and Filizola, M. (2019) Allostery in G protein-coupled receptors investigated by molecular dynamics simulations. *Curr. Opin. Struct. Biol. 55*, 121–128.
(57) Sencanski, M., Glisic, S., Šnajder, M., Veljkovic, N., Poklar Ulrih, N., Mavri, J., and Vrecl, M. (2019) Computational design and characterization of nanobody-derived peptides that stabilize the active conformation of the β2-adrenergic receptor (β2-AR). *Sci. Rep. 9*, 16555.
(58) Staus, D. P., Strachan, R. T., Manglik, A., Pani, B., Kahsai, A. W., Kim, T. H., Wingler, L.

M., Ahn, S., Chatterjee, A., Masoudi, A., Kruse, A. C., Pardon, E., Steyaert, J., Weis, W. I., Prosser, R. S., Kobilka, B. K., Costa, T., and Lefkowitz, R. J. (2016) Allosteric nanobodies reveal the dynamic range and diverse mechanisms of G-protein-coupled receptor activation. *Nature 535*, 448–452.
(59) Nygaard, R., Zou, Y., Dror, R. O., Mildorf, T. J., Arlow, D. H., Manglik, A., Pan, A. C., Liu, C. W., Fung, J. J., Bokoch, M. P., Thian, F. S., Kobilka, T. S., Shaw, D. E., Mueller, L., Prosser, R. S., and Kobilka, B. K. (2013) The dynamic process of β(2)-adrenergic receptor activation. *Cell 152*, 532–542.
(60) Cournia, Z., Allen, B., and Sherman, W. (2017) Relative binding free energy calculations in drug discovery: recent advances and practical considerations. *J. Chem. Inf. Model. 57*, 2911–2937.
(61) Wingler, L. M., Elgeti, M., Hilger, D., Latorraca, N. R., Lerch, M. T., Staus, D. P., Dror, R. O., Kobilka, B. K., Hubbell, W. L., and Lefkowitz, R. J. (2019) Angiotensin Analogs with Divergent Bias Stabilize Distinct Receptor Conformations. *Cell 176*, 468-478.e11.
(62) Staus, D. P., Wingler, L. M., Choi, M., Pani, B., Manglik, A., Kruse, A. C., and Lefkowitz, R. J. (2018) Sortase ligation enables homogeneous GPCR phosphorylation to reveal diversity in β-arrestin coupling. *Proc Natl Acad Sci USA 115*, 3834–3839.
(63) Lally, C. C. M., Bauer, B., Selent, J., and Sommer, M. E. (2017) C-edge loops of arrestin function as a membrane anchor. *Nat. Commun. 8*, 14258.
(64) Lee, Y., Warne, T., Nehmé, R., Pandey, S., Dwivedi-Agnihotri, H., Edwards, P. C., García-Nafría, J., Leslie, A. G. W., Shukla, A. K., and Tate, C. G. (2020) Molecular determinants of β-arrestin coupling to formoterol-bound β_1 -adrenoceptor. *BioRxiv*.
(65) Ho, T. K. (1995) Random decision forests. *Proceedings of 3rd international conference on document analysis and recognition 1*, 278–282.