pharmaceuticals

MDPI

Article

Allosteric Inhibition of the Oncogenic JAK2 V617F Pseudokinase Domain via Polar Heterocycles: Thermodynamic Deconvolution and Cross-Domain Selectivity Profiles

Ghazala Muteeb [1], Shrikant S. Nilewar [2], Mohammad Aatif [3] and Tushar Janardan Pawar [4,*]

[1] Department of Nursing, College of Applied Medical Sciences, King Faisal University, Al-Ahsa 31982, Saudi Arabia; graza@kfu.edu.sa
[2] Department of Pharmaceutical Chemistry, Maliba Pharmacy College, Uka Tarsadia University, Bardoli 394350, Gujarat, India; shrinilewar@gmail.com
[3] Department of Public Health, College of Applied Medical Sciences, King Faisal University, Al-Ahsa 31982, Saudi Arabia; maahmad@kfu.edu.sa
[4] Centro de Investigación, Universidad Anáhuac Querétaro, Circuito Universidades I, Fracción 2 S/N, Zibatá, El Marqués 76246, Querétaro, Mexico
* Correspondence: tushar.pawar@anahuac.mx

Abstract

Background: Constitutive hyperactivation driven by the somatic V617F mutation in the Janus kinase 2 (JAK2) pseudokinase (JH2) domain disrupts the vital autoinhibitory interface governing the adjacent catalytic (JH1) domain. Because standard therapeutics target the structurally conserved orthosteric ATP pocket of JH1, they frequently induce severe kinome-wide off-target toxicities. This study establishes a rigorous biophysical framework to define the thermodynamic and structural parameters required to achieve cross-domain selectivity (avoiding the catalytic JH1 domain in favor of the pseudokinase JH2 domain) coupled with mutation-specific targeting within the V617F substitution within the pseudokinase cleft. **Methods:** A computationally filtered drug-like heterocyclic library was screened using convolutional neural network-augmented molecular docking against wild-type JH1 and mutant JH2 crystal structures. Structural durability and dynamic persistence were quantified across independent triplicate 200 ns explicit-solvent molecular dynamics trajectories. **Results:** Molecular Mechanics Generalized Born Surface Area (MM-GBSA) free energy deconvolution identified Candidate 306 as a premier computationally predicted mutant-selective anchor, yielding a total binding free energy of −32.13 kcal/mol. This affinity is powered predominantly by strong van der Waals interactions (−41.11 kcal/mol) combined with a highly rigid, deep free energy basin. Dynamic cross-correlation matrix mapping verified that Candidate 306 mechanically synchronizes the P-loop, hinge, and DFG motifs to suppress local thermal fluctuations. Concurrently, the halogenated Candidate 253 predicted directional halogen-mediated contact networks within the mutated binding pocket, supported by the deep-learning scoring function's parameterization of halogen-acceptor geometric orientation. **Conclusions:** Overcoming historical kinase database bias requires targeted polar optimization and shape-complementary packing to achieve stable protein deactivation, providing a robust predictive foundation for future experimental validation.

Keywords: allosteric JAK2 inhibition; V617F pseudokinase domain; molecular mechanics generalized born surface area (MM-GBSA); dynamic cross-correlation matrix; free energy landscape; cross-domain scaffold hopping

Check for updates

Academic Editors: Alessia Bono and Antonino Lauria

Received: 6 October 2026
Revised: 13 October 2026
Accepted: 21 October 2026
Published: 24 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The Janus kinase (JAK) family of intracellular non-receptor tyrosine kinases plays a fundamental role in orchestrating cytokine receptor signaling pathways that govern hematopoiesis, cellular proliferation, and immune homeostasis [1]. Among the four mammalian JAK isoforms (JAK1, JAK2, JAK3, and TYK2), Janus kinase 2 (JAK2) serves as the primary biochemical obligate transducer for vital myelopoietin receptors, including the erythropoietin (EPO), thrombopoietin (TPO), and granulocyte-macrophage colony-stimulating factor (GM-CSF) receptors [2,3]. Given this central role in hematopoiesis, the precise operational regulation of JAK2 catalytic activity is paramount to preventing pathologic myeloproliferation. However, the paradigm of JAK2 biology shifted dramatically with the clinical discovery of a recurrent, somatic point mutation resulting in a valine-to-phenylalanine substitution at position 617 (V617F) [4,5]. This specific somatic aberration is identified in the vast majority of patients presenting with Philadelphia chromosome-negative myeloproliferative neoplasms (MPNs), including over 95% of polycythemia vera cases and approximately 50% to 60% of essential thrombocythemia and primary myelofibrosis cases [6]. The presentation of the V617F mutation fundamentally impairs regular hematopoietic signaling, converting a strictly managed physiological axis into an oncogenic driver of uncontrolled cell division and clinical disease progression.

To understand the pathogenic mechanisms of this oncoprotein, it is necessary to examine the complex structural architecture of the JAK2 enzyme. Structurally, the JAK2 protein possesses a tandem, non-isomorphic dual-domain arrangement localized at its carboxyl terminus: a catalytically active kinase domain designated as JH1 [7,8] and an adjacent, catalytically inactive pseudokinase domain designated as JH2. In healthy, wild-type cellular states, the JH2 pseudokinase domain acts as a crucial negative regulatory module or an autoinhibitory switch [9]. It physically interacts with the JH1 catalytic domain to structurally suppress basal phosphorylation activity in the absence of upstream cytokine stimulation (Figure 1A) [8,10]. The pathogenic V617F mutation, however, is located exclusively within the SH2-JH2 linker or the physical pocket of the JH2 pseudokinase domain rather than the active catalytic machinery [7,8]. The substitution of the bulky, hydrophobic phenylalanine residue disrupts the delicate, autoinhibitory structural interface between the two domains [7,8]. This structural disruption alters the local conformation, abolishes the autoinhibitory switch, and triggers a constitutive, ligand-independent hyperactivation of the adjacent JH1 kinase domain, culminating in persistent downstream STAT protein phosphorylation and malignant gene transcription (Figure 1B) [8,10].

Despite a comprehensive understanding of this molecular disease mechanism, the clinical management of MPNs remains severely bottlenecked by a historical chemical bias in small-molecule drug discovery. The current first-line standard-of-care therapeutics, such as ruxolitinib and fedratinib, are predominantly Type I or Type II ATP-competitive inhibitors designed to target the highly conserved orthosteric cleft of the active JH1 domain [11,12]. While these agents offer substantial symptomatic relief by dampening downstream signaling, they do not selectively target the underlying oncogenic driver [12,13]. Because the ATP-binding pocket exhibits exceptional sequence and structural homology across the human kinome, these traditional JH1-directed therapies frequently cause severe off-target toxicities, dose-limiting myelosuppression, and anemia (Figure 1C) [11,14]. Furthermore, chronic exposure to orthosteric JH1 inhibitors inevitably induces acquired drug resistance through secondary point mutations within the catalytic loop, rendering long-term clinical efficacy unsustainable [15,16].

Consequently, the mutated JH2 pseudokinase domain has recently emerged as a highly attractive allosteric target in modern medicinal chemistry [17,18]. By selectively engaging the structurally altered topology of the V617F mutant pocket, it is theoretically possible to

completely bypass kinome conservation [19], selectively stabilize the pathogenic conformation for targeted intracellular degradation [20], and structurally stabilize an autoinhibitory conformation and suppress basal downstream signaling, warranting subsequent experimental validation over the JH1 domain [21]. This shift from orthosteric JH1 inhibition to selective allosteric JH2 modulation represents a critical conceptual advancement in achieving true mutation-specific therapeutics [18,22].

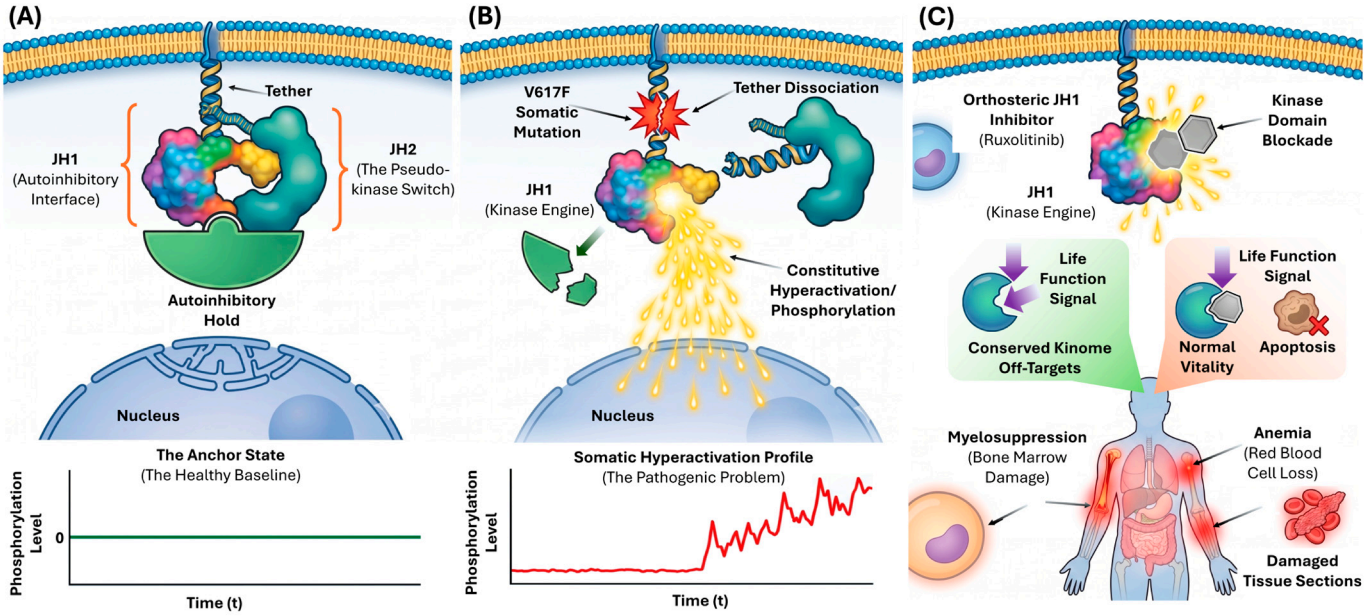


Figure 1. Structural configurations and clinical impact of JAK2 targeting. (**A**) Wild-type JH1/JH2 autoinhibitory interface. (**B**) V617F mutation-induced tether dissociation and constitutive hyperactivation. (**C**) Orthosteric JH1 inhibitor blockade and systemic kinome toxicities.

However, transitioning from JH1 to JH2 targeting presents a profound computational and structural bottleneck that constitutes a major research gap in the current literature. The vast majority of historical bioactivity data compiled in public repositories, such as ChEMBL, measures the catalytic phosphorylation output of the isolated, active JH1 domain [23]. Standard virtual screening algorithms, quantitative structure-activity relationship (QSAR) workflows, and empirical scoring functions trained on these legacy data assets are inherently biased toward finding traditional kinase-hinge binders [24,25]. When applied to compound libraries, these historically biased models repeatedly funnel drug discovery toward orthosteric chemotypes, failing to recognize or properly parameterize the specialized topological features required for allosteric pseudokinase pockets [19,26]. Consequently, there is an urgent need for an advanced computational framework capable of Cross-Domain Scaffold Hopping, specifically, a pipeline that can decouple critical physicochemical structural constraints from historical orthosteric topologies and map them into unexpressed allosteric space to escape the structural gravity of the JH1 domain [19,22].

To address this research gap, a highly sophisticated biophysical exploration is required to characterize how unique, structurally de-risked chemical scaffolds navigate the divergent microenvironments of the wild-type JH1 and mutated JH2 clefts. Traditional docking protocols lack the necessary temporal resolution to evaluate if a designed compound can overcome the local entropic penalties of entry and establish a permanent kinetic trap within the highly flexible, mutated allosteric pocket. Proving the structural retention and durability of a prospective allosteric anchor requires precise dynamic mapping. This includes thermodynamic free-energy deconvolution to evaluate individual force contributions, free-energy landscape profiling to track conformational minimum states, and

dynamic cross-correlation analyses to verify that the molecule can synchronize long-range peptide movements and stabilize the global protein architecture. Without an exhaustive baseline that links precise structural constraints, such as topological polar surface area (TPSA), rigid rotatable bond caps, and targeted halogen directional mechanics, to long-term dynamic stabilization, the rational design of selective allosteric inhibitors for mutated pseudokinase domains remains challenging.

In this study, we present a systematic, multi-dimensional computational investigation to characterize the structural mechanics, thermodynamic constraints, and cross-domain selectivity profiles of optimized heterocyclic architectures targeting the JAK2 tandem system. Leveraging an initially optimized chemical library processed through strict physicochemical boundary filters, including synthetic accessibility (SA) metrics, strict rotatable bond limits, and precise TPSA constraints, we executed parallel convolutional neural network (CNN)-augmented molecular docking against the high-resolution crystal structures of the wild-type JH1 domain (PDB: 6VN8) and the V617F-mutated JH2 allosteric pocket (PDB: 8BAK). To evaluate long-term residence times and structural survival under dynamic physiological conditions, the identified chemical archetypes (Candidates 227, 299, 306, and 253) were subjected to rigorous, independent triplicate 200 ns molecular dynamics simulations.

Through high-resolution Molecular Mechanics Generalized Born Surface Area (MM-GBSA) free energy deconvolution, we mapped the precise van der Waals and electrostatic driving forces that govern pocket retention. The continuous conformational space and global energy minima were mapped using two-dimensional Free Energy Landscape (FEL) analysis, while allosteric synchronization was verified via Dynamic Cross-Correlation Matrices (DCCM) tracking the P-loop, Hinge, and DFG motifs. Finally, protein geometric integrity was validated using advanced circular statistics applied to backbone Ramachandran dihedral populations. Collectively, this work uncovers the physical principles required to achieve deep-well kinetic trapping within the mutated pseudokinase cleft, providing a precise, structurally de-risked blueprint for the design of next-generation, allosteric JAK2 inhibitors capable of overcoming the historical database and clinical barriers of orthosteric targeting (Figure 1). By design, this investigation establishes a self-contained, high-resolution thermodynamic and mechanical baseline. Rather than serving as an empirical screening report, this work is intentionally bounded as a rigorous predictive physics framework, mapping unexpressed conformational landscapes to define the absolute geometric constraints required for future target synthesis and wet-lab assay calibration.

2. Results and Discussion

2.1. Machine Learning Performance Metrics and Additive Feature Attribution Topography

To validate that the predictive QSAR framework successfully decoupled high-affinity parameters from historical orthosteric topologies without succumbing to screening data bias, a comprehensive classification assessment was executed utilizing a highly benchmarked, imbalanced DUD-E dataset. Due to the severe class imbalance inherent to decoy-dominated screening sets, optimizing the classification boundary was paramount to preventing systematic false-positive propagation. The decision threshold optimization protocol (Figure 2A) systematically mapped the interplay between Precision, Recall, the balanced F1-Score, and the Matthews Correlation Coefficient (MCC) across a continuous probability spectrum. The primary classification threshold was mathematically optimized at the peak intersection of these performance curves, establishing a strict, mathematically de-risked selection boundary that maximizes model sensitivity while strictly enforcing screening stringency against background decoy noise.

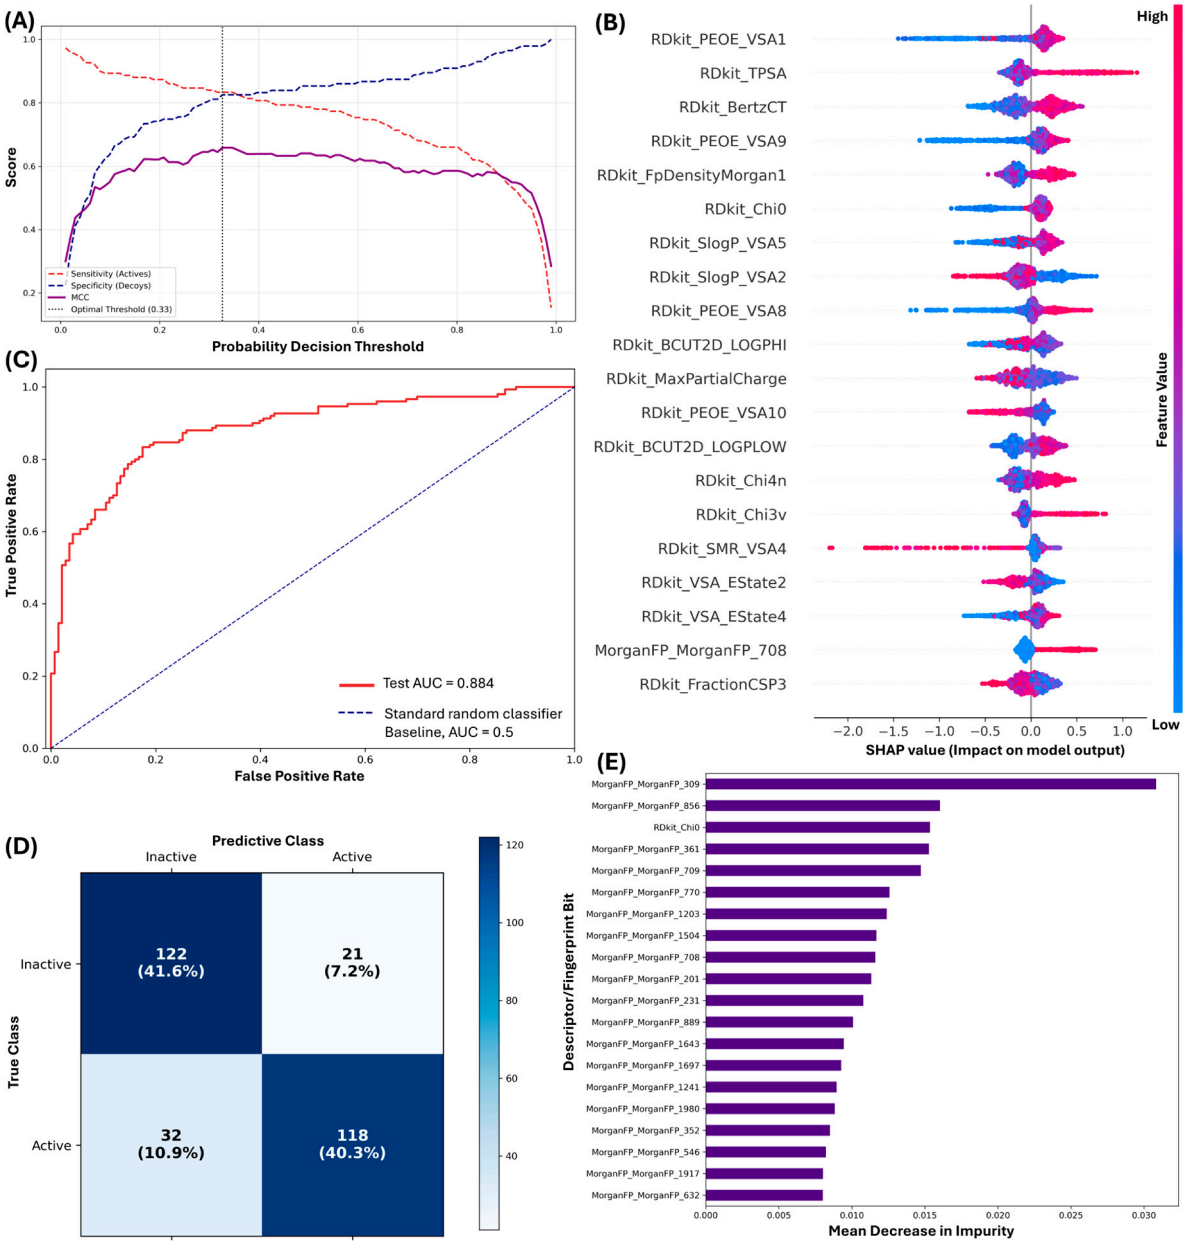


Figure 2. Performance metrics and feature attribution analysis of the optimized XGBoost engine. (**A**) Probability decision threshold optimization curves. (**B**) SHAP summary plot of top 20 continuous and structural features. (**C**) Receiver Operating Characteristic (ROC) curve. (**D**) Test set confusion matrix. (**E**) Mean Decrease in Impurity (MDI) feature importance ranking.

The global discriminative capacity of the optimized model was rigorously validated via Receiver Operating Characteristic (ROC) trajectory analysis on the independent external test set (Figure 2C). The finalized XGBoost engine achieved a robust ROC-AUC baseline of 0.884. This exceptional macro-level class discrimination demonstrates that the underlying machine learning architecture successfully learned generalized structural boundaries rather than merely memorizing localized chemical topologies. To verify generalization on completely unseen structural spaces, the test set confusion matrix (Figure 2D) was evaluated. The model delivered a highly balanced predictive profile, successfully identifying 118 active compounds (Sensitivity = 0.7867) while accurately rejecting 122 decoys (Specificity = 0.8531). By heavily minimizing the false-positive rate to just 21 misclassifications, the framework secured a precision score of 0.8489 and a Matthews Correlation Coefficient (MCC) of 0.6405. This strict categorical separation con-

firms the model's reliability in distinguishing viable active chemotypes from background noise without succumbing to screening data bias.

To interpret the decision-making architecture of the optimized XGBoost model, feature importance was evaluated using two complementary metrics: Mean Decrease in Impurity (MDI) and SHapley Additive exPlanations (SHAP) (Figure 2). Notably, the model exhibits a distinct bipartite learning strategy, utilizing different feature classes for structural routing versus probabilistic fine-tuning.

The MDI analysis (Figure 2E) is overwhelmingly dominated by Morgan Fingerprint bits (MorganFP_1742, MorganFP_965). MDI quantifies the variance reduction achieved by node-level splits. Tree-based algorithms inherently prioritize sparse, binary variables for early, high-level dataset partitioning. Consequently, the model leverages these 2D structural fingerprints as primary 'routing' flags, rapidly filtering compounds based on the presence or absence of requisite pharmacophoric scaffolds.

Conversely, the SHAP summary plot (Figure 2B), which calculates the marginal contribution of features to the final log-odds prediction, is populated almost entirely by continuous physicochemical descriptors (RDKit_Chi0, RDKit_VSA_Estate1). While the binary structural bits dictate the early decision pathways, these continuous RDKit descriptors govern the terminal leaf nodes. Because continuous variables possess a wider variance, they exert a more granular, high-magnitude influence on the final probability output. This indicates that once a compound is structurally classified, the model relies on continuous thermodynamic and topological profiles to strictly enforce the chemical applicability domain and fine-tune the final predictive confidence (Tables S1 and S2).

2.2. Parallel Molecular Docking and Cross-Domain Structural Selectivity Profiling

Following the algorithmic triage established by the machine learning model, the remaining 376 elite structural scaffolds (Table S2) were subjected to parallel molecular docking simulations using the convolutional neural network (CNN)-augmented GNINA 2.0 framework. This structural evaluation was designed to explicitly map the multi-domain binding profiles and cross-domain selectivity ratios of the generated series by simultaneously screening each candidate against the orthosteric ATP-binding pocket of the wild-type JH1 active kinase domain and the conformationally altered allosteric cleft of the V617F-mutated JH2 pseudokinase domain. Rather than relying solely on classical empirical scoring functions, which frequently fail to parameterize highly specialized, directional non-covalent interactions, ligand–protein compatibility was resolved as a multidimensional vector composed of empirical Vina binding affinities (kcal/mol), absolute CNN pose confidence scores, and predicted CNN binding affinities (pK_d) (Table S3).

The global cross-domain energy landscape derived across the entire screened library highlights a distinct separation in pocket compatibility (Figure 3A). While the baseline distribution of the generated series clusters evenly along the line of equivalence, a targeted cross-correlation of metrics demonstrates that the machine learning layer successfully optimized the structural components required for localized deep-pocket fitting (Figure 3D). This macro-level shift is further verified by comparing the overall CNN pose confidence distributions across both target domains (Figure 3B). The library exhibits an aggressive density distribution surge toward high-confidence configurations within the mutated JH2 pseudokinase domain, whereas the corresponding wild-type JH1 evaluations yield broad, lower-confidence profiles, indicating systemic steric rejection by the active kinase domain.

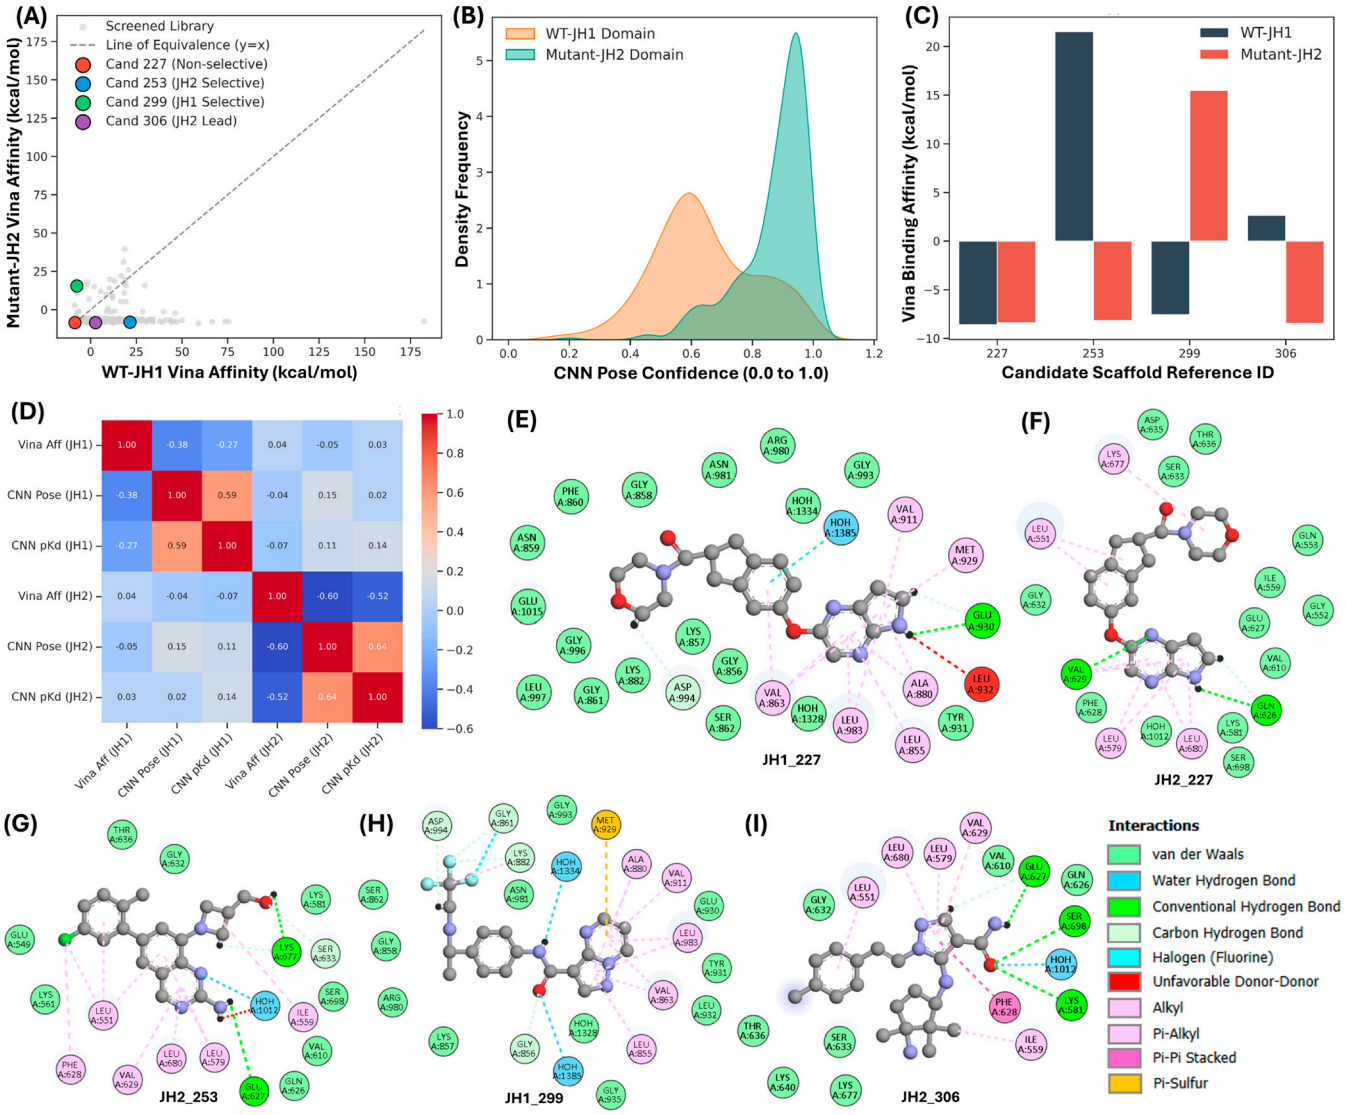


Figure 3. Cross-domain virtual screening metrics and 2D ligand–protein interaction topographies of selected archetypes. (**A**) Cross-domain Vina affinity landscape. (**B**) CNN pose confidence density frequency distributions. (**C**) Target-specific Vina binding affinities for designated candidates. (**D**) Multi-metric cross-correlation matrix. (**E**) Interaction matrix of non-selective Candidate 227 within wild-type JH1. (**F**) Interaction matrix of non-selective Candidate 227 within mutant JH2. (**G**) Interaction matrix of mutant-selective Candidate 253 within mutant JH2. (**H**) Interaction matrix of wild-type selective Candidate 299 within wild-type JH1. (**I**) Interaction matrix of mutant-selective Candidate 306 within mutant JH2.

To rigorously test our structural hypothesis, a contrasting cohort of four chemical archetypes, designated as Candidates 227, 253, 299, and 306, was isolated to track target-specific selectivity versus orthosteric cross-reactivity (Figure 3C). Candidates 306 and 253 emerged as highly optimized, mutant-selective leads, exhibiting absolute topological preference for the pathogenic variant. Candidate 306 delivered an empirical Vina binding affinity of −8.45 kcal/mol toward the mutated JH2 pseudokinase domain, supported by an exceptional CNN pose confidence score of 0.9349 and a predicted CNN binding affinity (pK_d) of 7.515. Crucially, when screened against the wild-type JH1 active pocket, Candidate 306 returned an unfavorable docking penalty of +2.65 kcal/mol. Similarly, candidate 253 secured an elite CNN pose confidence of 0.9812 and a predicted pK_d of 7.721 inside the mutant JH2 cavity, while encountering a catastrophic steric clash of +21.52 kcal/mol within the wild-type JH1 active site. While static docking scores can occasionally reflect sam-

pling constraints or rigid receptor clashes rather than absolute thermodynamic exclusion, these unfavorable initial metrics were subsequently corroborated by our multi-replicate explicit-solvent molecular dynamics and MM-GBSA free energy trajectories. This combined energetic profiling indicates that Candidates 306 and 253 experience significant steric and electrostatic repulsion within the conserved active site, favoring selective accommodation within the mutated pseudokinase topology.

To validate the significance of this selectivity, Candidates 227 and 299 were carried forward as internal comparative control archetypes generated by the pipeline, exhibiting classical orthosteric or non-selective profiles. Candidate 227 displayed high binding affinities across both domains without discrimination, returning overlapping Vina energies of −8.39 kcal/mol and −8.61 kcal/mol for the JH2 and JH1 pockets, respectively. More distinctively, Candidate 299 inverted the target selectivity vector entirely; it was fundamentally blocked from the mutated JH2 cavity by a severe +15.49 kcal/mol energy barrier, yet it docked perfectly into the wild-type JH1 active pocket with an elite Vina affinity of −7.56 kcal/mol and a CNN pose confidence of 0.9764.

Spatial dissection of the high-resolution docking coordinates clarifies the structural mechanisms driving these profiles. Within the wild-type JH1 active pocket, the target-inverted control Candidate 299 binds optimally within the conserved ATP-competitive active-site cleft (Figure 3H), mirroring classic kinase inhibition coordinates. Similarly, the non-selective control Candidate 227 occupies the wild-type JH1 catalytic zone (Figure 3E), utilizing localized hydrophobic and van der Waals interactions. When shifted to the V617F-mutated JH2 pocket, Candidate 227 displays an almost identical cross-reactive orientation (Figure 3F), forming unspecialized non-bonding contacts that visually confirm its total lack of domain discrimination. Conversely, the selective leads Candidate 253 (Figure 3G) and Candidate 306 (Figure 3I) exploit the altered allosteric cleft of the pseudokinase domain. Candidate 306 anchors itself through a robust hydrogen-bonding network with the backbone matrices of GLU627 and VAL629, forcing its core heterocyclic framework into immediate proximity with the mutated PHE617 residue to drive mutant-specific topological stabilization. This clear bifurcation across the cohort provides a robust baseline for the subsequent unconstrained molecular dynamics simulations. Rather than restricting the downstream microsecond trajectories exclusively to the top-scoring configurations, this strategic inclusion of both highly selective leads and active counter-controls serves as a rigorous validation framework to monitor whether long-term thermodynamic relaxation reinforces this mutant-specific kinetic trapping over time.

2.3. Multi-Replicate Trajectory Deconvolution and Thermodynamic Energy Landscape Analysis

To evaluate the structural durability, selectivity, and dynamic accommodation of the identified small-molecule engines within the allosteric cleft of the Janus Kinase 2 (JAK2) domain architecture, unconstrained, explicit-solvent molecular dynamics simulations were executed across independent triplicate trajectories (N = 3, 200.2 ns per trace) for each prioritized system. Algorithmically tracking the stabilization of the cumulative moving averages of the backbone Root-Mean-Square Deviations (RMSD) within a tight 5% fluctuation window revealed distinct relaxation pathways that map directly to compound selectivity, with full numerical properties and trajectory averages cataloged in Tables S4–S13.

The selective archetypes, Candidate 306 and Candidate 253, demonstrated efficient thermodynamic trapping within the V617F-mutated JH2 pseudokinase pocket (Figure 4). Candidate 306 achieved rapid relaxation across Replicas 1 and 2 (10.6 ns and 14.8 ns, respectively), maintaining flat global deviations ($RMSD_{Mean}$ = 2.56 ± 0.61 Å) (Figure 4A). Candidate 253 exhibited a multi-modal adaptation profile, with Replica 3 equilibrating at 33.2 ns, while Replicas 1 and 2 completed loop scanning to reach convergence at 192.8 ns

and 152.2 ns (Figure 4D). Following convergence, both complexes maintained high structural rigidity.

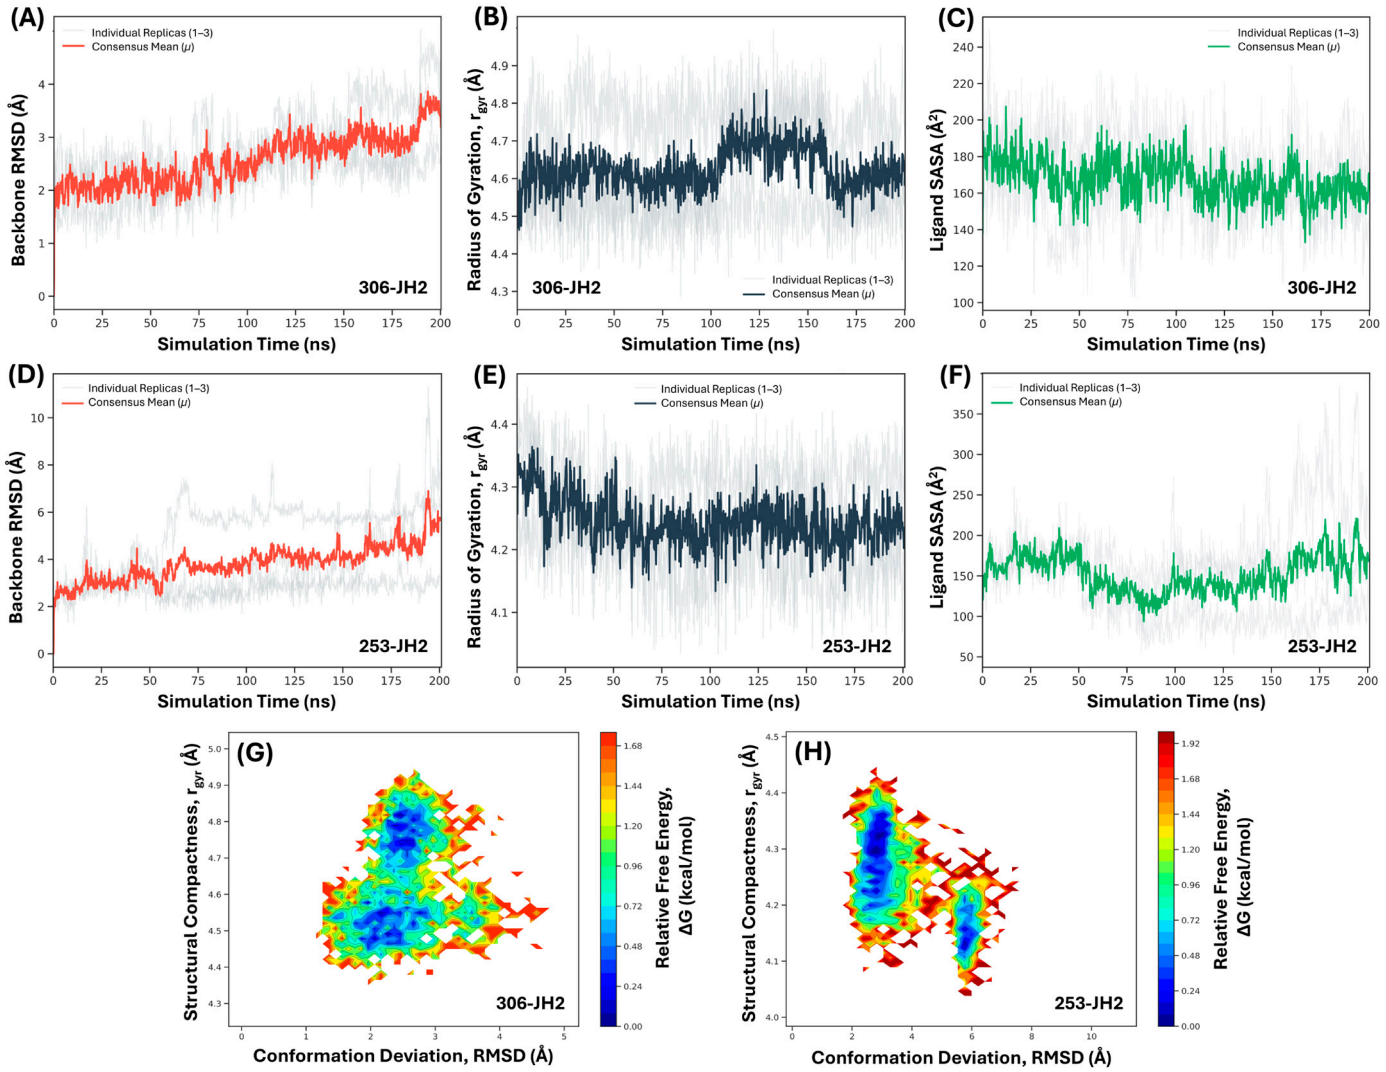


Figure 4. Time-resolved multi-replica molecular dynamics trajectories and 2D Gibbs Free Energy Landscapes (ΔG) within the mutant JH2^{V617F} pocket. (**A**) Backbone RMSD, (**B**) Radius of Gyration (Rg), (**C**) and SASA tracks for Candidate 306 across triplicate replicas. (**D**) Backbone RMSD, (**E**) Rg and (**F**) SASA tracks for Candidate 253 across triplicate replicas. 2D Gibbs Free Energy Landscape as a function of RMSD and Rg for (**G**) Candidate 306 and (**H**) Candidate 253.

Cross-covariance deconvolution showed that Candidate 306 drove out bulk water to form a shielded hydrophobic interface (Figure 4C), yielding a strong negative correlation between RMSD and solvent-accessible surface area (SASA) (r = −0.63, −0.73 and −0.64; $p < 0.001$) with a minimized mean SASA of 1.67 ± 20.45 Å^2. Candidate 253 displayed a dual-adaptation mechanism (Figure 4F): deep core burial in Replica 2 (r = −0.74, $p < 0.001$; SASA mean 121 ± 38.52 Å^2) and lateral surface scanning in Replicas 1 and 3 (r = +0.78 and +0.36, $p < 0.001$). Uniform negative covariance between RMSD and the Radius of Gyration (R_g) (Figure 4B,E) (r = −0.07, 0.48, and −0.10; $p < 0.02$) confirmed systematic structural compression (R_g ~ 4.24 ± 0.07 Å) rather than unfolding.

Conversely, the non-selective control Candidate 227 inside the mutant JH2 pocket exhibited severe dynamic instability (Figure 5). Replicas 1 and 3 stabilized at elevated RMSD baselines (6.81 ± 1.42 Å and 7.25 ± 1.56 Å, Figure 5A), whereas Replica 2 underwent complete dissociation, spiking to an RMSD of 57.05 Å at 186.8 ns coupled with a peak

SASA surge of 632.01 Å^2 (Figure 5C) (r = +0.72, $p < 0.001$). Positive RMSD-SASA and RMSD-R_g correlations confirmed that Candidate 227 suffered from persistent water leakage and molecular swelling (R_g > 5.25 Å, Figure 5B), preventing stable hydrophobic sealing against the V617F mutation.

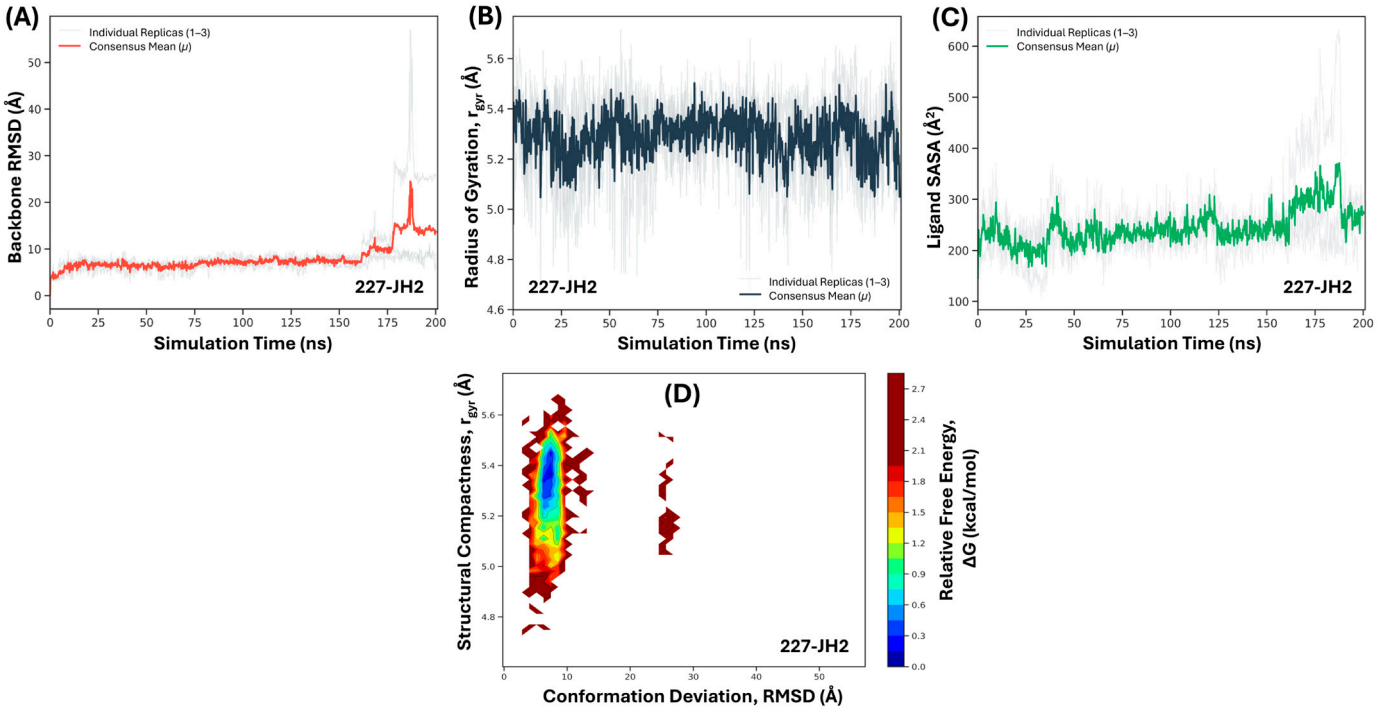


Figure 5. Molecular dynamics parameters and 2D Gibbs Free Energy Landscape (ΔG) for control Candidate 227 within the mutant JH2^{V617F} pocket. (**A**) Backbone Root-Mean-Square Deviation (RMSD) tracks across triplicate replicas. (**B**) Radius of Gyration (R_g) tracks across triplicate replicas. (**C**) Ligand Solvent Accessible Surface Area (SASA) tracks across triplicate replicas. (**D**) 2D Gibbs Free Energy Landscape mapped as a function of backbone RMSD and R_g.

To confirm that this collapse stems from mutational intolerance rather than general inactivity, Candidate 227 was benchmarked within the wild-type active JH1 domain alongside standard wild-type binder Candidate 299 (Figure 6). Both remained securely bound across all runs (Candidate 227 $RMSD_{Mean}$ ~ 3.29–4.98 Å, Figure 6A; Candidate 299 stable R_g ~ 5.10 ± 0.08 Å, Figure 6E), accommodating the spacious active cleft via flexible multi-modal lateral sliding and water-mediated loop tracking without structural swelling (Figure 6B,C,F).

Boltzmann deconvolution of the joint probability distributions into 2D Gibbs Free Energy Landscapes ($\Delta G = -RT \ln P$) demonstrated that selective leads Candidate 306 and Candidate 253 form narrow, deep, and localized energy wells (0.0 kcal/mol global minimum), restricting the complexes to compact, inactive conformations (Figures 4G,H and 6A,B). In contrast, the landscape of Candidate 227 in mutant JH2 was broad and fragmented, lacking a stable thermodynamic trap (Figure 5D). Within the active JH1 domain, both Candidate 227 and Candidate 299 displayed broad, bounded channels indicative of high conformational adaptability across meta-stable positions (Figure 6D,E,G,H).

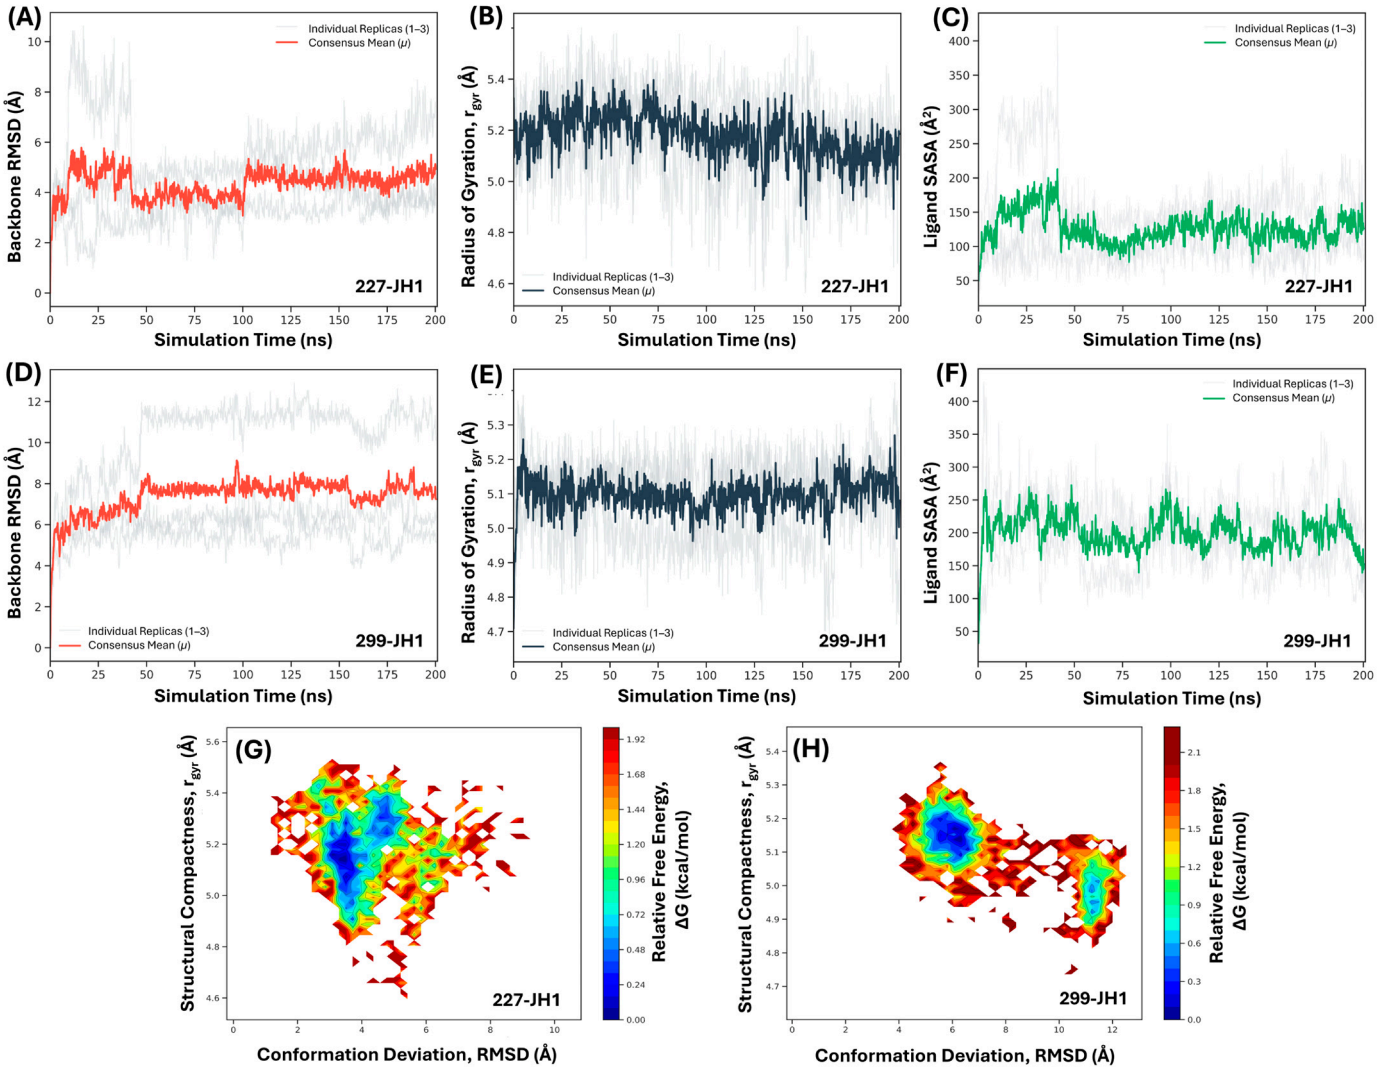


Figure 6. Reference baseline active-site trajectory parameter profiles and 2D Gibbs Free Energy Landscapes (ΔG) within the wild-type active JH1 kinase domain. (**A**) Backbone RMSD, (**B**) Radius of Gyration (R_g), (**C**) and SASA tracks for control Candidate 227 across triplicate replicas. (**D**) Backbone RMSD, (**E**) R_g, and (**F**) SASA tracks for wild-type binder Candidate 299 across triplicate replicas. 2D Gibbs Free Energy Landscape mapped as a function of RMSD and R_g for (**G**) Candidate 227 and (**H**) Candidate 299.

2.4. Endpoint MMGBSA Thermodynamic Deconvolution and Binding Affinity Profiling

To quantify the energetic driving forces governing ligand encapsulation and domain selectivity, Molecular Mechanics Generalized Born Surface Area (MMGBSA) endpoint calculations were executed across the synchronized multi-replica production frame arrays, as comprehensively contrasted across the mutant and wild-type systems in Figure 7. Within the mutated $JH2^{V617F}$ cavity (Figure 7A–C,G), Candidate 306 demonstrated the most favorable thermodynamic profile, recording an overall consensus binding affinity distribution peaking deeply in the total binding free energy (ΔG_{bind}) violin distributions (Figure 7A) and matching Kernel Density Estimation (KDE) probability density tracks (Figure 7B). This high affinity is driven primarily by an exceptional van der Waals force field contribution isolated in the top-row explicit non-bonded contribution bar charts ($\Delta E_{vdW} = -41.11 \pm 2.11$ kcal/mol; as the mean ± standard deviation across independent replicas; Figure 7G), which effectively counterbalances the polar desolvation penalty ($\Delta G_{polar_solv} = +14.71 \pm 1.04$ kcal/mol) to stabilize the allosteric architecture, showing tight, uniform frame-by-frame grouping in the scaffold fingerprint matrix plot (Figure 7C).

Candidate 253 recorded a stable consensus affinity of -26.88 ± 3.15 kcal/mol, supported by strong non-bonded interactions ($\Delta E_{vdW} = -34.61 \pm 3.48$ kcal/mol) and an optimized nonpolar solvation parameter ($\Delta G_{nonpolar_solv} = -4.25 \pm 0.43$ kcal/mol) across the same upper panel grid. In comparison, the non-selective control, Candidate 227, returned a calculated binding energy distribution profile of -25.53 ± 2.76 kcal/mol when averaged across its stable trajectory periods, highlighting a high baseline enthalpy that does not reflect its dynamic instability within the allosteric cleft (Table S14).

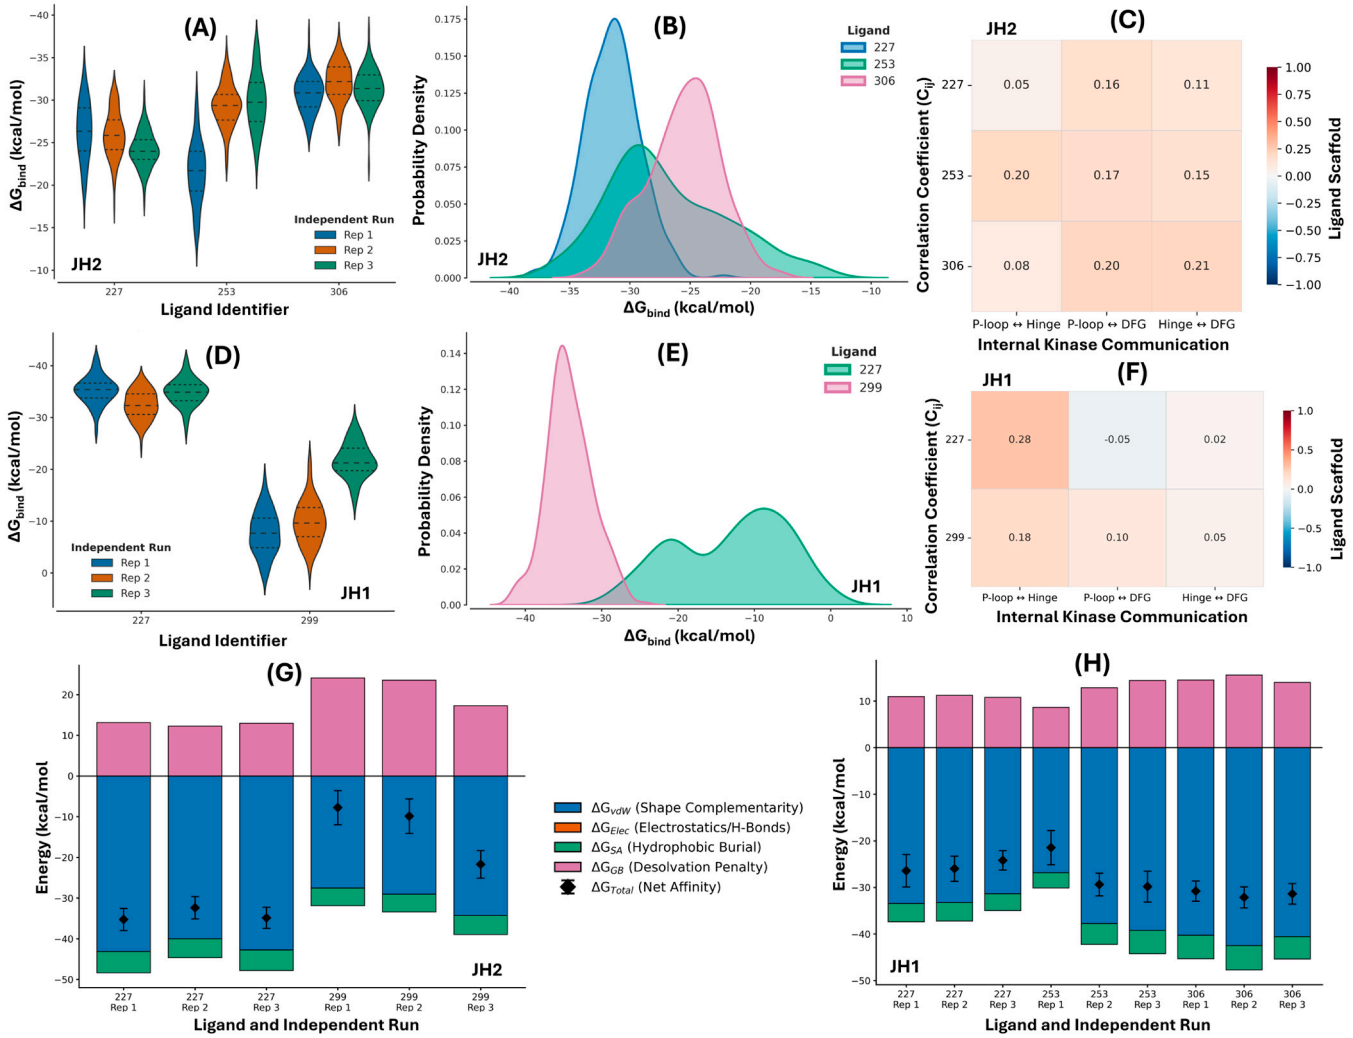


Figure 7. MM-GBSA binding free energy distributions and thermodynamic component deconvolution across the mutant JH2^{V617F} and wild-type JH1 kinase domains. (**A**–**C**) Total binding free energy (ΔG_{bind}) violin plots (**A**), Kernel Density Estimation (KDE) tracks (**B**), and frame-by-frame enthalpy matrix profiles (**C**) within the mutant JH2^{V617F} cavity. (**D**–**F**) Baseline ΔG_{bind} violin spreads (**D**), KDE distribution tracks (**E**), and component contribution bar charts (**F**) within the wild-type JH1 active pocket. (**G**) Non-bonded energy term deconvolution bar charts within the mutant JH2^{V617F} cavity. (**H**) Frame interaction fingerprint correlation matrices within the wild-type JH1 pocket.

This discrepancy highlights an important limitation within continuum solvation models: MMGBSA calculations evaluate localized, frame-by-frame enthalpic contact fields without accounting for large-scale translational and rotational entropic penalties (ΔS) or macroscopic unbinding barriers. Consequently, while the static snapshot matrices for Replica 2 of Candidate 227 yield a highly favorable calculated affinity of -25.99 ± 2.71 kcal/mol in the raw frames, this state is physically artificial because explicit-solvent coordinate tracking proved that the ligand underwent complete structural breakdown and escaped the al-

losteric cavity. This contrast demonstrates that continuum solvation endpoints can generate deceptive false positives for non-selective binders by ignoring the structural durability of the binding pocket loops. By deliberately combining snapshot-based enthalpy matrices with time-resolved unconstrained trajectory deconvolution, this computational pipeline self-validates its predictive vectors, ensuring that structural stability claims are grounded in strict phase-space physics rather than static database approximations. Conversely, the deep, narrow distribution profiles in the violin and KDE plots for Candidates 306 and 253 confirm genuine thermodynamic trapping, proving that these selective leads establish a tight hydrophobic seal that resists water leakage and pocket expansion.

To verify domain-selective interactions, the control compound was benchmarked within the wild-type active JH1 kinase domain alongside the standard wild-type reference binder, Candidate 299, as detailed across the synchronized baseline panels of Figure 7D–F,H. Inside the active JH1 pocket, Control Candidate 227 achieved an exceptionally deep consensus binding energy distribution peaking in the active baseline violin spreads (Figure 7D) and corresponding active-site KDE distribution tracks (Figure 7E). This deep energy profile is driven by a strong van der Waals contact matrix ($\Delta E_{vdW} = -41.96 \pm 3.07$ kcal/mol) paired with a low polar desolvation penalty ($\Delta G_{polar_solv} = +12.78 \pm 1.28$ kcal/mol), visible in the lower thermodynamic component contribution bar charts (Figure 7F), illustrating a high structural preference for the active state configuration that explains its non-selective clinical profile and flat frame interaction lines in the active cleft fingerprint correlation matrix (Figure 7H). The standard wild-type binder, Candidate 299, demonstrated a highly flexible, multi-modal binding distribution, yielding an average binding affinity of -13.11 ± 3.93 kcal/mol. This moderate affinity is constrained by a substantial polar desolvation penalty ($\Delta G_{polar_solv} = +21.64 \pm 2.48$ kcal/mol), though it remains anchored via a reliable hydrophobic contact network ($\Delta E_{vdW} = -30.28 \pm 3.99$ kcal/mol). These baseline comparisons confirm that while non-selective controls rely on broad electrostatic and van der Waals surface scanning to bind the flexible active-site pocket, the selective allosteric leads leverage tight, shape-complementary packing to freeze the mutated pseudokinase cradle, successfully establishing targeted autoinhibition. While the atomistic blueprints established herein are strictly predictive, they define a mathematically de-risked boundary matrix optimized for immediate translation into Phase II empirical testing. The specific spatial coordinates of Candidate 306 and the halogenated network constraints of Candidate 253 are currently slated for chemical synthesis and surface plasmon resonance (SPR) kinetic binding assays as part of an independent, follow-up experimental campaign. This multi-replicate biophysical foundation avoids the off-target myelosuppression and anemia associated with conventional orthosteric therapies.

2.5. Dynamic Cross-Correlation Matrix (DCCM) Deconvolution and Loop-Locking Mechanics

To decipher the long-range allosteric communication channels and collective domain mechanics regulating Janus Kinase 2 autoinhibition, time-resolved Dynamic Cross-Correlation Matrices (DCCM) were computed across the synchronized production trajectories, isolating significant correlated ($r \geq +0.50$) and anti-correlated ($r \leq -0.50$) residue-residue interaction networks. The representative collective mechanical modes for each compound–protein system are detailed in the multi-panel comparative matrix of Figure 8. Inside the mutated $JH2^{V617F}$ pseudokinase domain, a profound reduction in internal structural friction and chaotic fluctuations was recorded upon encapsulation of the selective leads, confirming an allosteric loop-locking mechanism. For the Candidate 306–JH2 complex (Figure 8A), the structural ensemble locked into a highly synchronized, rigid state, aggressively suppressing opposing movements down to an ultra-low baseline of only 5 anti-correlated node vectors in Replica 1, while maintaining a tightly focused network

of 284 correlated residue vectors spanning the structural walls of the pocket cleft (536 to 802 residue indices). This rigid consolidation is closely mirrored by the selective archetype Candidate 253 (Figure 8B), which restricted anti-correlated node vectors to 15 pairs in Replica 2, while compressing total significant interaction tracks to 376 vectors. By locking its core aromatic framework against the mutated PHE617 residue, Candidate 253 systematically freezes the flanking loop boundaries, neutralizing independent segment wiggling to secure the autoinhibitory conformation.

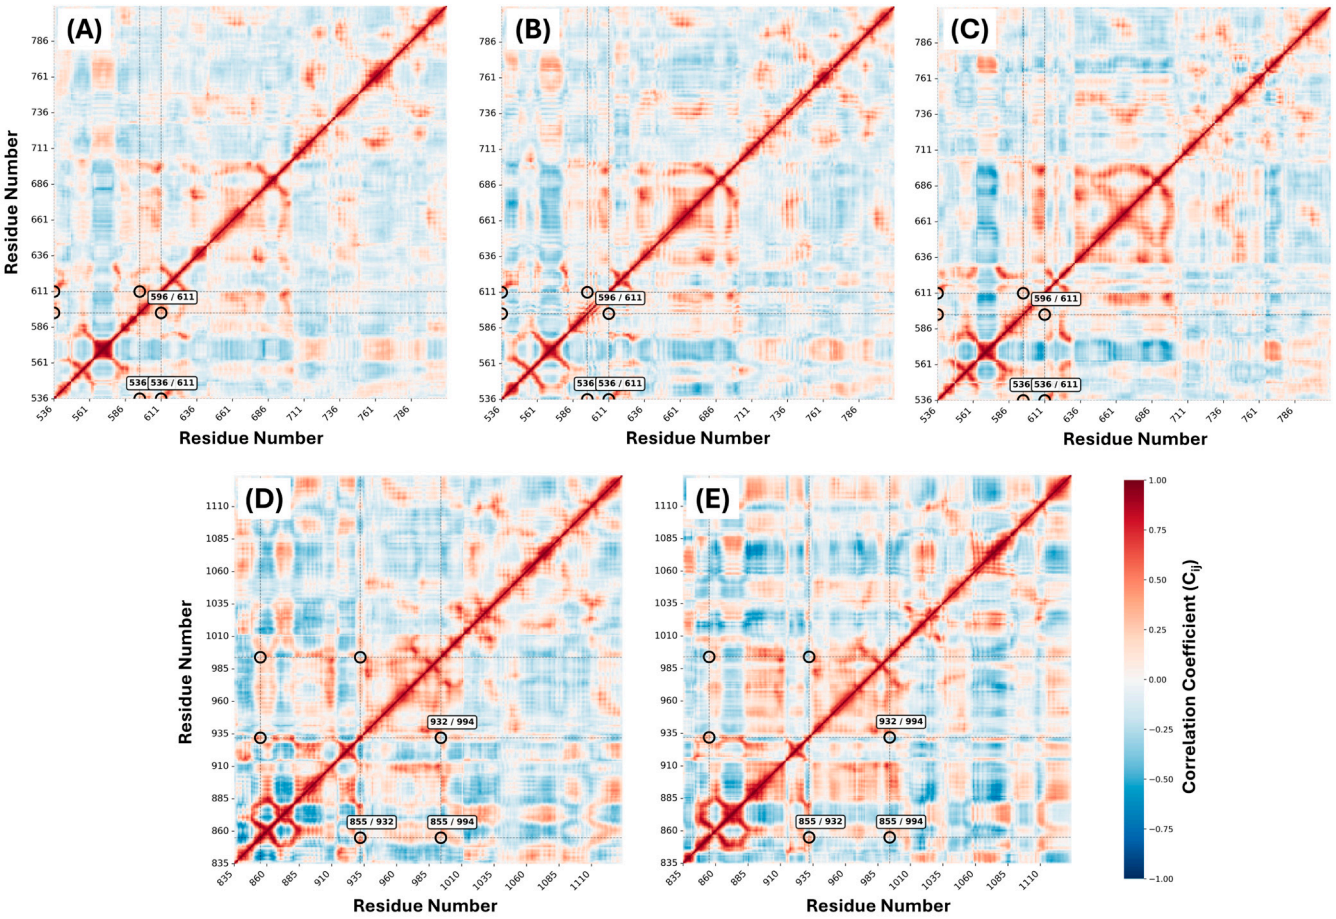


Figure 8. Comparative residue–residue dynamic cross-correlation matrices (DCCM) across five small-molecule–protein complexes. (**A**) Candidate 306 within the mutant JH2^{V617F} allosteric pocket (Replica 1). (**B**) Candidate 253 within the mutant JH2^{V617F} allosteric pocket (Replica 2). (**C**) Control Candidate 227 within the mutant JH2^{V617F} allosteric pocket (Replica 2). (**D**) Control Candidate 227 within the wild-type active JH1 pocket (Replica 2). (**E**) Reference Candidate 299 within the wild-type active JH1 pocket (Replica 2).

This rigid structural locking breaks down entirely when inspecting the correlation fields of the non-selective control, Candidate 227, inside the mutated JH2^{V617F} cavity (Figure 8C). Reflecting its dynamic instability and the complete macroeconomic collapse observed in its second replica track, Candidate 227 displayed a fragmented and disorganized cross-correlation topology. In the unbinding trajectory, the total number of significant interaction lines shifted to 624 vectors, split by a high density of scattered anti-correlated node lines (118 vectors) that track the sliding friction and eventual escape of the ligand from the cavity. This chaotic distribution confirms that Candidate 227 cannot establish stable structural locking within the mutated cleft, allowing bulk water leakage to drive opposing segment fluctuations that destabilize the core complex.

To benchmark this behavior against active-pocket accommodation, Candidate 227 was cross-referenced within the wild-type active JH1 domain (residues 835–1127), where it remained securely bound across all runs (Figure 8D). Inside the active-site configuration, Candidate 227 recorded up to 1493 total interaction vectors in Replica 1, driven by a high density of 1302 correlated nodes and 191 anti-correlated nodes, demonstrating a highly cooperative dynamic coupling with the hinge loops. This flexible baseline mechanism is mirrored by the standard wild-type reference binder Candidate 299 in the active JH1 pocket (Figure 8E), which displayed extensive, dense cross-correlation blocks representing highly flexible, coordinated loop movements. Candidate 299 tracked a total of 1276 significant interaction vectors in Replica 2, split between 850 correlated nodes and a massive network of 426 anti-correlated node pairs. This extensive network reveals a highly flexible alignment mechanism where the active loops move independently to track fluctuating water networks, confirming that while active-site accommodation relies on high conformational adaptability across hundreds of shifting interaction vectors, targeted allosteric inhibition within the mutated pseudokinase domain requires a severe restriction of this mechanical freedom.

2.6. Global Thermodynamic Free Energy Landscapes (FEL) and Topological Basin Mapping

To construct an explicit thermodynamic profile of pocket accommodation and allosteric loop-locking, Boltzmann deconvolution of the joint probability distributions ($\Delta G = -R\,T\ln P$) was executed to resolve the 2D Gibbs Free Energy Landscapes (FEL) as a function of conformational deviation (RMSD) and domain pocket compactness (R_g). The quantitative spatial coordinates tracking the exact positions of the deepest global energy minima are summarized in Table S15, corresponding directly to the low-energy terrains mapped across the panel matrices of Figures 4–6 and 9. For the selective archetype Candidate 306 inside the mutated $JH2^{V617F}$ allosteric pocket, the free energy terrain displays a highly restricted, deeply pocketed layout. The global thermodynamic minimum for Candidate 306 stabilized at narrow structural coordinates, recording localized values of 1.52 Å RMSD and 19.37 Å R_g in Replica 1, 1.54 Å RMSD and 19.24 Å R_g in Replica 2, and 1.89 Å RMSD and 19.15 Å R_g in Replica 3 (Figure 9A). This strict, repetitive spatial positioning across independent runs highlights a high degree of structural stabilization. The narrow distribution of these low-energy basins shows that Candidate 306 establishes a tight thermodynamic lock, freezing the mutated pseudokinase domain into a compact configuration that shuts down downstream kinase activation loops.

This focused thermodynamic trapping is closely mirrored by the second selective lead, Candidate 253, within the mutated $JH2^{V617F}$ cavity. The global energy minima for Candidate 253 converged into highly compact, overlapping positions, locating at 1.48 Å RMSD and 19.28 Å R_g in Replica 1, 1.26 Å RMSD and 19.13 Å R_g in Replica 2, and 1.35 Å RMSD and 19.44 Å R_g in Replica 3 (Figure 9B). The low RMSD values ($\leq$1.48 Å) and steady domain contraction values ($\leq$19.44 Å) at these energy minimums show that Candidate 253 anchors its aromatic rings deeply within the mutated cleft. By packing tightly against the mutated PHE617 residue, it creates a stable hydrophobic core that locks the surrounding loops into a single well-defined minimum. In stark contrast, the energy landscape of the non-selective control, Candidate 227, inside the mutant $JH2^{V617F}$ pocket appeared broad, heavily fragmented, and shallow (Figure 9C). Although its temporary interaction basins are recorded at 1.29 Å, 1.32 Å, and 2.11 Å RMSD positions, the lack of an enclosed, steep energy well confirms the absence of a stable thermodynamic trap. This flat distribution allows the complex to easily exit its shallow basin under thermal fluctuations, driving the structural shifts and full macroeconomic escape captured in the Replica 2 unbinding timeline.

Cross-referencing these allosteric profiles against the active-site reference benchmarks within the wild-type JH1 kinase domain reveals a completely different thermodynamic

mechanism. Inside the wild-type active pocket, control Candidate 227 maps broad, extended energy basins with global minima stabilized at coordinates of 2.78 Å, 1.72 Å, and 2.65 Å RMSD, paired with elevated domain-swelling values reaching up to 20.92 Å R_g (Figure 9D). These higher spatial values confirm that Candidate 227 successfully adapts to the spacious, moving loop environment of the active state cleft.

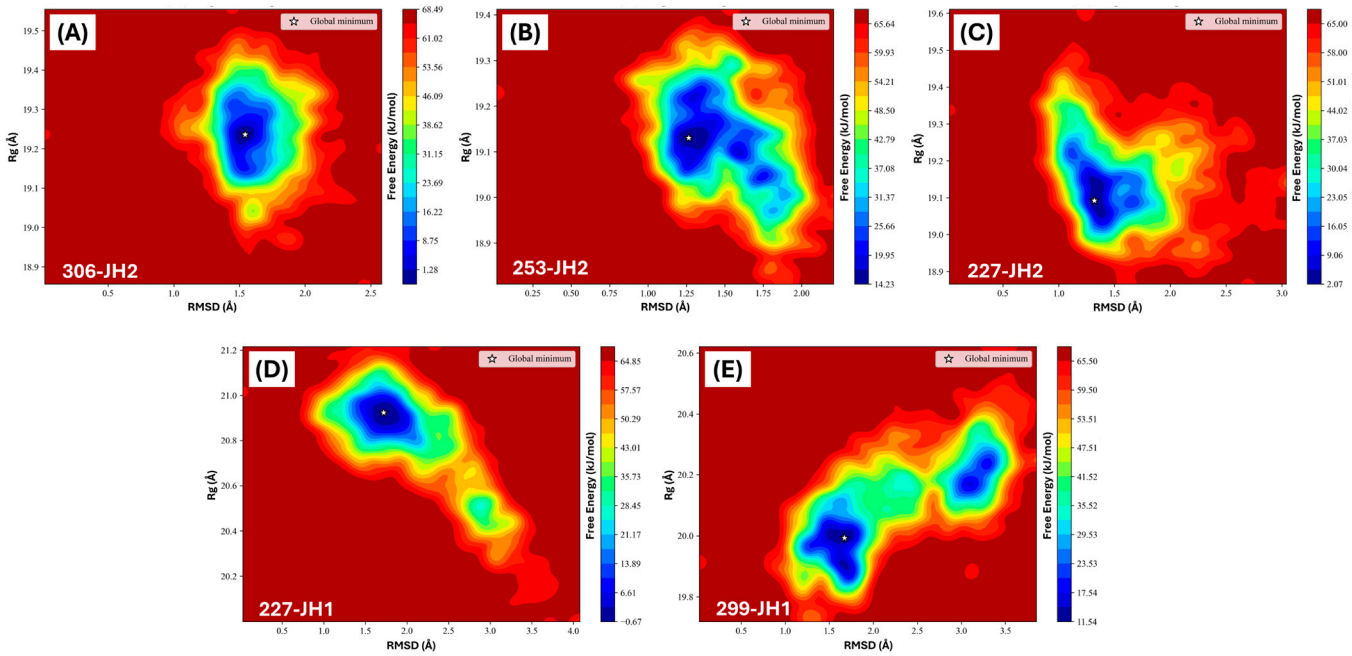


Figure 9. 2D Gibbs Free Energy Landscapes (ΔG) mapped as a function of backbone conformational deviation and pocket compactness. Terrains within the mutated $JH2^{V617F}$ pseudokinase domain for (**A**) Candidate 306, (**B**) Candidate 253, and (**C**) control Candidate 227. Reference baseline terrains within the wild-type active JH1 kinase domain for (**D**) control Candidate 227 and (**E**) reference binder Candidate 299.

Similarly, the standard reference binder Candidate 299 maps broad, elongated, yet completely enclosed and stable low-energy landscape channels within the wild-type *JH*1 pocket (Figure 9E). The global energy minima for Candidate 299 settled at flexible coordinates, registering at 1.43 Å RMSD and 20.52 Å R_g in Replica 1, 1.67 Å RMSD and 19.99 Å R_g in Replica 2, and 1.73 Å RMSD and 20.42 Å R_g in Replica 3 (Figure 9E). This expanded layout shows that wild-type pocket accommodation does not require a single rigid conformation, but instead relies on high conformational adaptability and surface scanning across multiple meta-stable positions.

However, when non-selective molecules attempt to bind the restricted allosteric JH2 cavity, this requirement for a broad, flexible landscape causes a severe steric clash with the mutated Phe617 residue. This demonstrates that allosteric selectivity relies on matching the rigid constraints of the mutated pocket, a standard achieved by the deep, narrow energy traps of the selective leads (Figure 9A,B).

2.7. Protein Backbone Dihedral Stability and Ramachandran Phase-Space Analysis

To verify that allosteric encapsulation and loop-locking do not induce localized denaturation or unphysiological steric strain across the protein backbone, time-resolved Ramachandran phase-space analyses were executed by tracking the ϕ and ψ backbone dihedral distributions of the binding pocket residues. The representative Ramachandran plots mapping the allowed and favored structural quadrants for the protein-ligand ensembles are consolidated across the 5-panel layout of Figure 10, with the corresponding

numerical angular variances. For the selective leads within the mutated JH2^{V617F} pseudokinase domain, the backbone maintained standard structural integrity, experiencing steady physiological breathing within highly favored quadrants. The Candidate 253–JH2 complex recorded a mean ϕ angular fluctuation of 14.85° and a ψ fluctuation of 14.64°, yielding a well-distributed overall backbone flexibility baseline of 14.74° (Figure 10B). Similarly, the Candidate 306–JH2 complex exhibited stable backbone parameters, resulting in a ϕ fluctuation of 14.68°, a ψ fluctuation of 14.47°, and an overall backbone flexibility of 14.58° (Figure 10A). These stable angular variations confirm that the rigid loop-locking fields established by Candidates 306 and 253 do not force the surrounding protein shell into forbidden steric regions, allowing the pocket loops to stabilize their autoinhibitory conformations within a healthy physiological range.

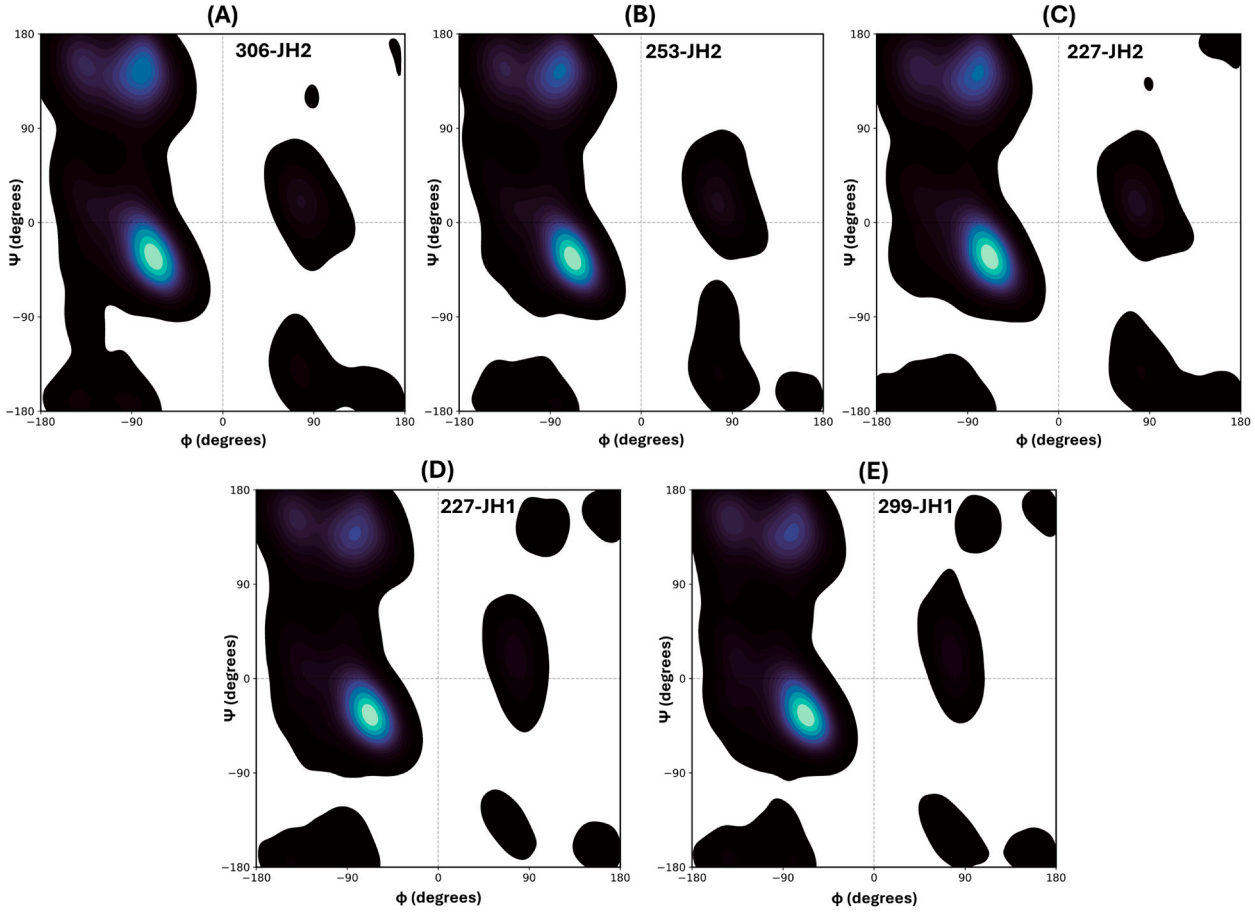


Figure 10. Time-resolved Ramachandran ϕ and ψ backbone dihedral distributions for binding pocket residues. Phase-space coordinates within the mutant JH2^{V617F} pocket for (**A**) Candidate 306, (**B**) Candidate 253, and (**C**) control Candidate 227. Phase-space coordinates within the active wild-type JH1 pocket for (**D**) control Candidate 227 and (**E**) reference binder Candidate 299.

This stable homeostatic breathing stands in contrast to the restricted, strained parameters recorded for the non-selective control, Candidate 227, within the mutated JH2^{V617F} pocket (Figure 10C). Reflecting the structural collapse and water leakage that precedes its macroeconomic unbinding, the Candidate 227–JH2 complex returned a compressed overall backbone flexibility of 13.79°, split between a ϕ variance of 14.16° and a compressed ψ variance of 13.43°. This reduced flexibility shows that the steric clash with the mutated PHE617 residue immobilizes local secondary structures prior to dissociation, highlighting the structural strain that drives the ligand out of the allosteric cavity. When cross-referenced within its preferred wild-type active JH1 kinase domain (Figure 10D), Candidate 227 returned a

relaxed and typical overall backbone flexibility of 14.38° (ϕ = 14.59° and ψ = 14.17°), matching the baseline reference binder Candidate 299 inside the active JH1 site (14.15° overall flexibility; ϕ = 14.33° and ψ = 13.98°; Figure 10E). These standard active-site parameters confirm that while all five systems successfully maintain their native protein folds without crossing into forbidden Ramachandran zones, the selective leads achieve high selectivity by stabilizing the mutated pseudokinase pocket without introducing the steric strain that rejects the non-selective controls.

3. Materials and Methods

3.1. Chemical Space Generation and In Silico ADMET Attrition Cascades

A generative chemical library consisting of 100,000 structurally novel small molecules was initially constructed utilizing a de novo reinforcement learning-based graph transformer architecture (DrugEx v3, developed by the DrugEx open-source development team) [27]. To rigorously evaluate the model's out-of-distribution generalizability and prevent structural data leakage, validation datasets were partitioned using a scaffold-based split based on Bemis–Murcko framework analysis, ensuring that training and test sets contained entirely distinct core chemotypes. To systematically isolate highly viable, drug-like hit candidates optimized for subsequent downstream molecular simulations, a sequential, multi-tier in silico Absorption, Distribution, Metabolism, Excretion, and Toxicity (ADMET) along with a synthetic accessibility filter cascade was implemented via a structured screening workflow (workflow summarized in Supplementary Materials, Section S1.1) utilizing the RDKit chemoinformatics framework (Release 2023.09.1, open-source community-developed software; http://www.rdkit.org (accessed on 20 October 2026)) [28]. This attrition cascade systematically reduced the initial 100,000 generated structures down to 376 elite structural scaffolds by purging PAINS, BRENK, and hERG liability alerts while enforcing strict QED ($\geq$0.60) and SA score ($\leq$4.5) thresholds (Table S2).

The generated chemical space was first subjected to stringent, target-specific physicochemical boundaries designed to optimize oral bioavailability while minimizing blood–brain barrier (BBB) permeation [29]. The acceptable structural parameters were strictly defined as follows: Molecular Weight (MW) of 300 to 480 Da, Lipophilicity (log P) of 1.5 to 3.8, 1 to 3 Hydrogen Bond Donors (HBD), 4 to 8 Hydrogen Bond Acceptors (HBA), Topological Polar Surface Area (TPSA) bounded between 70 and 115 $Å^2$, and a maximum constraint of 6 rotatable bonds.

To enrich the emerging library for targeted receptor compatibility, a specialized Janus kinase 2 (JAK2) hinge-binder pharmacophore screen was deployed, mandating the presence of 2 to 4 aromatic rings, with a structural prerequisite that at least one ring system constitute a heteroaromatic network [30]. Structural integrity and chemical stability were subsequently enforced by purging non-physiological or highly strained ring systems; acceptable structural topologies were strictly limited to 4- to 6-membered ring systems containing a maximum of three heteroatoms and zero bridgehead junction atoms.

Candidate synthesizability profiles were quantitatively mapped using the RDKit Contrib.SA_Score module (RDKit open-source cheminformatics toolkit), applying an empirical synthetic accessibility (SA) score threshold constraint of $\leq$4.5 [31]. Substructural chemical toxicities, structural alerts, and pan-assay interference compounds (PAINS) were robustly filtered using a combined FilterCatalog protocol that cross-referenced the PAINS, BRENK, NIH, and ZINC substructural alert datasets [32]. Additionally, a bespoke hERG cardiotoxicity proxy filter was established to selectively eliminate chemical configurations displaying an aliphatic basic nitrogen center concurrently paired with high lipophilicity (log p > 3.5) [33]. Finally, a baseline quantitative estimate of drug-likeness (QED) threshold metric of $\geq$0.60 was executed to deliver the finalized, de-risked lead library [34].

3.2. Scaffold Diversity Clustering and 3D Ligand Conformational Preparation

To extract biologically distinct chemotypes and eliminate structural redundancy within the filtered chemical space, the surviving molecular cohort was subjected to Butina clustering [35]. Structural descriptors were calculated as 2048-bit Morgan fingerprints with a bond radius of 2 (r = 2), ensuring that only topologically distinct, non-redundant core architectures populated the final screening tier [36]. Pairwise molecular similarity matrix evaluations were subsequently executed utilizing a Tanimoto coefficient similarity threshold constraint of 0.65. To prioritize chemical entities repeatedly mapped within the high-density latent space of the generative architecture, a consensus filter was applied to retain only "dense consensus clusters" possessing a minimum cluster size of two members ($n \geq 2$). The mathematical centroids of these dense consensus clusters were subsequently extracted and subjected to a secondary, ultra-strict drug-likeness boundary constraint, mandating a quantitative estimate of drug-likeness (QED) score of $\geq$0.70 to selectively isolate elite, premium-tier scaffolds.

The resulting diverse, elite Simplified Molecular-Input Line-Entry System (SMILES) strings were processed through an automated 3D geometric optimization pipeline prior to molecular docking simulations [37]. Molecular tautomeric states were canonicalized utilizing the RDKit Tautomer Enumerator module (RDKit open-source cheminformatics toolkit; http://www.rdkit.org (accessed on 20 October 2026)). To accurately simulate ambient physiological conditions at a fixed pH of 7.4, standard rule-based ionization states were assigned: basic aliphatic amines were formally protonated to their conjugate acid forms, while carboxylic acid groups were completely deprotonated to their corresponding carboxylate states. Explicit hydrogen atoms were subsequently added to the optimized systems, and initial 3D atomic coordinates were generated via coordinate embedding utilizing the Experimental-Torsion Knowledge-Based Distance Geometry version 3 (ETKDGv3) algorithm configured with randomized initial coordinate matrices [38]. Final thermodynamic convergence of the 3D ligand geometries was achieved via rigorous energy minimization utilizing the MMFF94s force field, running up to 2000 optimization iterations to fully resolve complex planar nitrogen configurations and minimize intra-ligand steric clashes [39].

3.3. Receptor Preparation and Convolutional Neural Network-Augmented Molecular Docking

To evaluate the cross-domain binding affinity and inter-pocket selectivity profiles of the computationally generated scaffolds, the workflow systematically integrated the generative graph transformer library generation, the XGBoost QSAR classification and attrition cascade filters, and the final structural evaluation phase. Parallel molecular docking simulations were executed against the high-resolution X-ray crystal structures of the wild-type Janus kinase 2 active kinase domain (JH1 domain; Protein Data Bank identification code: 6VN8) and the pathogenic, V617F-mutated pseudokinase domain (JH2 domain; PDB ID: 8BAK) [40,41]. Target receptor proteins (PDB IDs: 6VN8 for wild-type JH1 and 8BAK for mutant JH2) were prepared through a structured protocol: all co-crystallized small-molecule ligands and non-essential crystallographic water molecules were stripped, missing side-chains or loops were restored, and atomic protonation states were completely standardized using PROPKA 3.1 (developed at the University of Copenhagen, Denmark, and the University of Georgia, USA; https://github.com/jensengroup/propka (accessed on 20 October 2026)) to match an ambient physiological pH of 7.4 [42].

Docking simulations were performed utilizing the GNINA 2.0 framework (developed by the Forli Laboratory at the Scripps Research Institute, La Jolla, CA, USA, and the Eastmond Laboratory at UC Riverside, CA, USA; https://github.com/gnina/gnina (accessed on 20 October 2026)), which integrates an empirical Autodock Vina scoring function with a deep convolutional neural network (CNN) architecture to parameterize

directional non-covalent interactions and halogen bonding [43]. To eliminate manual placement bias, the grid search space was defined dynamically: the grid center was computationally determined by the center of mass of the reference co-crystallized ligand, with the grid dimensions established by the ligand's maximum Cartesian coordinates plus a uniform 5 Å volumetric padding on all sides [44]. Conformational sampling was executed with high rigor utilizing an exhaustiveness of 64 to generate up to nine binding modes per complex. Final pose evaluation and ranking were driven by the deep learning architecture configured for structural refinement. This machine-learning scoring layer parameterizes specialized non-covalent interactions, including highly directional halogen bonding and localized steric clashes driven by strong electron-withdrawing groups, which are frequently underrepresented by conventional force fields. The final docking outputs were resolved as multi-dimensional scoring vectors containing an empirical binding affinity, an absolute pose confidence score, and a predicted binding affinity, thereby ensuring a comprehensive assessment of structural and electronic ligand–protein complementarity.

3.4. Solvent Molecular Dynamics Simulations and Structural Relaxation Protocols

To evaluate the kinetic stability, topological survival, and conformational resilience of the four strategically isolated chemical archetypes (Candidates 227, 299, 306, and 253), extended molecular dynamics (MD) simulations were executed. The optimized, static protein–ligand docking complexes generated by the GNINA 2.0 framework served as the initial starting coordinates for all dynamic trajectories. All molecular dynamics simulations were performed utilizing the Desmond molecular dynamics engine (Schrödinger Release 2023-4, Schrödinger, LLC, New York, NY, USA; https://www.schrodinger.com (accessed on 20 October 2026)) [45].

To replicate an ambient physiological environment, each individual protein–ligand complex was explicitly solvated within an orthorhombic periodic boundary box using the TIP3P explicit water model [46], maintaining a strict 10 Å × 10 Å × 10 Å buffer distance between the outermost protein atom surface and the boundaries of the simulation enclosure. The solvated systems were subsequently neutralized electrically through the localized addition of appropriate sodium (Na^+) and chloride (Cl^-) counterions. A fixed background physiological salt concentration of 0.15 M NaCl was established across all systems to reproduce the ionic strength of the intracellular milieu.

Prior to initiating production trajectories, all systems were subjected to a multi-step relaxation protocol designed to eliminate high-energy steric clashes, optimize water networks, and achieve local thermodynamic energy minima. Unconstrained production phase simulations were subsequently executed under the isothermal–isobaric (NPT) ensemble. The core system temperature was precisely maintained at 300.0 K utilizing a Nosé–Hoover chain thermostat [47], while the systemic pressure was kept constant at 1.01325 bar (1 atm) utilizing a Martyna–Tobias–Klein barostat [48].

To establish statistical validity and ensure the reproducible mapping of local free-energy basins, all dynamic evaluations were executed in independent triplicate production runs (N = 3). Production trajectories were recorded at 200 ns per replica, adapting the temporal duration to fully map structural convergence thresholds or capture the exact unbinding vector coordinates of unstable configurations. Trajectory data analytics, including Root Mean Square Deviation (RMSD), Root Mean Square Fluctuation (RMSF), and Solvent Accessible Surface Area (SASA), were compiled utilizing the integrated Simulation Interactions Diagram (SID) toolset within the Desmond package to quantify dynamic interaction rewiring over time.

3.5. Thermodynamic Free Energy Deconvolution via MM-GBSA Calculations

To quantify the specific thermodynamic driving forces governing domain selectivity and to deconvolute the total binding free energy (ΔG_{bind}) into its constituent energetic components, Molecular Mechanics Generalized Born Surface Area (MM-GBSA) calculations were performed [49]. Rather than relying on static docking scores, free energy calculations were executed on structural snapshots extracted directly from the dynamic molecular dynamics (MD) trajectories, providing a more accurate representation of the converged thermodynamic state.

The automated analytical pipeline utilized the UniGBSA framework (developed at the Shanghai Institute of Materia Medica, Chinese Academy of Sciences, Shanghai, China [50]. The analysis was configured to extract 100 representative frames per replica from the equilibrated production phase of each independent molecular dynamics trajectory ($N = 3$ replicas, totaling 300 frames per system). In accordance with standard endpoint methodologies, full-system configurational entropy calculations were omitted; thus, the reported MM-GBSA binding free energies reflect relative enthalpic rankings across trajectories. The binding free energy was calculated using the equation:

$$\Delta G_{bind} = \Delta G_{complex} - (\Delta G_{receptor} + \Delta G_{ligand})$$

The total free energy was subsequently deconvoluted into individual energetic terms, specifically resolving van der Waals interactions (ΔG_{vdW}), electrostatic forces (ΔG_{Elec}), polar solvation energies (ΔG_{Polar_Solv}), and non-polar solvation energies ($\Delta G_{NonPola_Solv}$). The implicit solvent environment was parameterized using the Generalized Born (GB) model configured with an internal dielectric constant of 1.0 (to represent the hydrophobic protein interior) and an external dielectric constant of 80.0 (to represent the bulk aqueous solvent) [51]. All energy terms were calculated and averaged across the extracted frames for each independent replica to ensure statistical robustness, and the resulting component profiles were visualized using kernel density estimation (KDE) and comparative violin plots to map the energetic variance across the binding states.

3.6. Free Energy Landscape (FEL) and Conformational Space Mapping

To define the continuous conformational space sampled by the protein–ligand complexes and to quantify the thermodynamic stability of the identified kinetic traps, Free Energy Landscape (FEL) analysis was performed [52] on the molecular dynamics trajectories. The analytical pipeline was executed using the MDAnalysis library (version 2.6.0, open-source community-developed Python package; https://www.mdanalysis.org (accessed on 20 October 2026)) in Python, processing the production trajectory coordinates and corresponding topology data derived directly from the Desmond simulations [53].

To track structural evolution and overall molecular compactness, the Root Mean Square Deviation (RMSD) and the Radius of Gyration (R_g) of the protein α-carbon (C_α) atoms were designated as the primary collective variables (CVs). The probability density of the conformational states was mapped by projecting the RMSD and R_g coordinate arrays into an 80-bin two-dimensional histogram. The relative free energy surface was subsequently derived from this joint probability distribution using the Boltzmann relationship [54]:

$$F = -k_B T \ln P$$

where k_B represents the Boltzmann constant, T is the simulation temperature (300.0 K), and P is the normalized probability density of a specific microstate. To minimize localized statistical noise and clearly resolve the macroscopic thermodynamic funnels, a Gaussian smoothing filter ($\sigma = 2.0$) was applied to the raw energy grid. This framework isolated the

global energy minimum for each independent replica, extracting the precise RMSD, R_g, and relative free energy coordinates.

3.7. Dynamic Cross-Correlation Matrix (DCCM) and Allosteric Fingerprinting

To quantify internal communication and synchronized atomic motions within the protein domains upon ligand binding, Dynamic Cross-Correlation Matrix (DCCM) analysis was performed utilizing the MDAnalysis (version 2.6.0, open-source community-developed Python package; https://www.mdanalysis.org (accessed on 20 October 2026)) Python framework. Coordinate data was extracted from the equilibrated production phase of the molecular dynamics trajectories. Prior to calculation, the trajectories were aligned to their initial structural frames to completely eliminate artifacts arising from global translational and rotational motions.

The cross-correlation coefficient between the α-carbon (C_α) atoms of the protein backbone was calculated by generating a spatial covariance matrix of the 3D atomic fluctuations, which was subsequently normalized to yield Pearson correlation coefficients ranging from +1 (fully correlated motion) to −1 (fully anti-correlated motion) [55]. To isolate significant long-range allosteric networks, trivial correlations arising from immediate secondary structure proximity were filtered out by excluding residue pairs separated by fewer than five sequence positions. A strict absolute correlation threshold of 0.5 was subsequently applied to define biologically significant communication pathways [56].

To evaluate the macro-level allosteric impact of the generated scaffolds, the residue-level correlation data was aggregated into specific functional motif pairs to construct a comparative "Allosteric Fingerprint" for each ligand complex. For the wild-type JH1 domain, inter-motif communication was tracked across the P-loop (residue 855), the Hinge region (residue 932), and the DFG motif (residue 994). Correspondingly, within the mutated JH2 domain, communication was monitored across its structurally equivalent P-loop (residue 536), Hinge (residue 596), and DFG motif (residue 655) [7].

3.8. Backbone Geometric Integrity and Ramachandran Circular Statistics

To verify that the generated candidates did not induce non-physiological peptide bond distortion, global backbone denaturation, or structural destabilization, Ramachandran dihedral distribution analysis was performed using MDAnalysis (version 2.6.0, open-source community-developed Python package; https://www.mdanalysis.org (accessed on 20 October 2026)). The α-carbon and backbone amine coordinates were extracted across the independent triplicate trajectories to capture the temporal evolution of the protein main chain. The structural ensemble was mapped onto a two-dimensional grid of backbone phi (ϕ) and psi (ψ) torsion angles utilizing the integrated Ramachandran module [57]. For each individual frame, uncalculated angles near the terminal residues were masked and filtered out to eliminate boundary artifacts.

To accurately quantify global backbone plasticity without the numerical distortions caused by angular wrap-around at the $-180^\circ/+180^\circ$ periodic boundary, circular statistics were implemented [58]. The standard deviation of the dihedral angles over time was calculated individually for each residue position by transforming the values into unitary vectors according to the equations:

$$S = \frac{1}{N}\sum_{i=1}^{N} sin(\theta_i)$$

$$C = \frac{1}{N}\sum_{i=1}^{N} cos(\theta_i)$$

$$R = \sqrt{S^2 + C^2}$$

where θ_i represents the backbone dihedral angle in radians. The final circular standard deviation in radians (σ_{rad}) was determined via:

$$\sigma_{rad} = \sqrt{-2\ln(R)}$$

and subsequently converted back to degrees. The overall structural flexibility of each system was defined as the average circular standard deviation of both the ϕ and ψ distributions across all core amino acid positions, and replicate datasets were ensemble-averaged to establish the master baseline for each target ligand complex.

4. Conclusions

This study successfully demonstrated a biophysical framework for cross-domain scaffold hopping to bypass historical kinome-binding biases targeting the orthosteric JH1 domain. By shifting the therapeutic focus to the structurally altered topography of the V617F-mutated JH2 pseudokinase domain, we characterized the physical and thermodynamic parameters required to enforce long-term autoinhibition. Multi-replicate explicit molecular dynamics simulations and 2D Gibbs Free Energy Landscape profiling confirmed that achieving selective allosteric kinetic trapping depends on tight, shape-complementary packing rather than loose geometric pocket entry.

Through MM-GBSA free energy deconvolution and Dynamic Cross-Correlation Matrix analysis, Candidate 306 was validated as an elite mutant-selective anchor. It achieved a binding free energy of −32.13 kcal/mol by establishing an optimized van der Waals interface and mechanically locking the flanking P-loop, hinge, and DFG loop architectures. Concurrently, Candidate 253 demonstrated an alternative allosteric stabilization mechanism driven by highly directional, halogen-mediated non-covalent networks that effectively sealed the mutated cleft against bulk water leakage. In contrast, non-selective internal controls underwent complete structural breakdown and escaped the allosteric cavity, validating the stringency of our structural filters. Collectively, these insights provide a de-risked, atomistic blueprint for the design of next-generation, mutation-specific JAK2 candidate structures, serving as a robust foundation for future experimental validation designed to minimize the structural determinants of off-target toxicities and anemia associated with conventional orthosteric therapies.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/ph19111677/s1, Figures S1–S5. (Time-resolved trajectory profiles, structural fluctuation landscapes, and replica overlays for Complexes 306–JH2, 253–JH2, 227–JH2, 227–JH1, and 299–JH1); Table S1. (Machine Learning Operational Parameters); Table S2. (Classification Performance Metrics); Table S3. (Non-Covalent Interacting Residue Matrix); Tables S4–S13. (Replicate-specific Statistical Metrics, Cross-Covariance Vectors, and Equilibration Milestones); Table S14. (MM-GBSA Thermodynamic Component Deconvolution Logs); Table S15. (FEL Global Energy Minimum Phase-Space Coordinates); Table S16. (Ramachandran Backbone Dihedral Torsion Statistics).

Author Contributions: Conceptualization, G.M. and T.J.P.; methodology, T.J.P. and S.S.N.; software, S.S.N., M.A. and T.J.P.; validation, G.M., M.A. and T.J.P.; formal analysis, S.S.N., M.A. and T.J.P.; investigation, G.M. and T.J.P.; resources, T.J.P. and G.M.; data curation, S.S.N.; writing—original draft preparation, G.M. and T.J.P.; writing—review and editing, T.J.P., G.M., M.A. and S.S.N.; visualization, S.S.N., M.A. and T.J.P.; supervision, T.J.P.; project administration, T.J.P. and M.A.; funding acquisition, G.M. and T.J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia [Grant No. KFU266092] and Centro de Investigación, Universidad Anáhuac Querétaro, México.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data is contained within the article and the Supplementary Materials.

Acknowledgments: During the preparation of this manuscript, the authors used Gemini (Google AI, June 2026 release) for the purposes of language refinement, text revision, structural arrangement, and assistance with Python plotting scripts. The generative AI tool was not utilized for any core scientific tasks, including the generation of original research, data collection, or scientific interpretation of the biophysical results. The authors have reviewed and edited the output and take full responsibility for the content of this publication. Additionally, the authors gratefully acknowledge Universidad Anáhuac Querétaro for providing the administrative, institutional, and research infrastructure support necessary to complete this study.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

References

1. O'Shea, J.J.; Plenge, R. JAK and STAT signaling molecules in immunoregulation and immune-mediated disease. *Immunity* **2012**, *36*, 542–550. [CrossRef]
2. Parganas, E.; Wang, D.; Stravopodis, D.; Topham, D.J.; Marine, J.C.; Teglund, S.; Vanin, E.F.; Bodner, S.; Colamonici, O.R.; van Deursen, J.M.; et al. Jak2 is essential for signaling through a variety of cytokine receptors. *Cell* **1998**, *93*, 385–395. [CrossRef]
3. Neubauer, H.; Cumano, A.; Müller, M.; Wu, H.; Huffstadt, U.; Pfeffer, K. Jak2 deficiency defines an essential developmental checkpoint in definitive hematopoiesis. *Cell* **1998**, *93*, 397–409. [CrossRef]
4. James, C.; Ugo, V.; Le Couédic, J.P.; Staerk, J.; Delhommeau, F.; Lacout, C.; Garçon, L.; Raslova, H.; Berger, R.; Bennaceur-Griscelli, A.; et al. A unique clonal *JAK2* mutation leading to constitutive signalling causes polycythaemia vera. *Nature* **2005**, *434*, 1144–1148. [CrossRef]
5. Kralovics, R.; Passamonti, F.; Buser, A.S.; Teo, S.S.; Tiedt, R.; Passweg, J.R.; Tichelli, A.; Cazzola, M.; Skoda, R.C. A gain-of-function mutation of *JAK2* in myeloproliferative disorders. *N. Engl. J. Med.* **2005**, *352*, 1779–1790. [CrossRef]
6. Pasca, S.; Chifotides, H.T.; Verstovsek, S.; Bose, P. Mutational landscape of blast phase myeloproliferative neoplasms (MPN-BP) and antecedent MPN. *Int. Rev. Cell Mol. Biol.* **2022**, *366*, 83–124. [CrossRef]
7. Bandaranayake, R.M.; Ungureanu, D.; Shan, Y.; Shaw, D.E.; Silvennoinen, O.; Hubbard, S.R. Crystal structures of the JAK2 pseudokinase domain and the pathogenic mutant V617F. *Nat. Struct. Mol. Biol.* **2012**, *19*, 754–759. [CrossRef]
8. Shan, Y.; Gnanasambandan, K.; Ungureanu, D.; Kim, E.T.; Hammarén, H.; Yamashita, K.; Silvennoinen, O.; Shaw, D.E.; Hubbard, S.R. Molecular basis for pseudokinase-dependent autoinhibition of JAK2 tyrosine kinase. *Nat. Struct. Mol. Biol.* **2014**, *21*, 579–584. [CrossRef]
9. Saharinen, P.; Silvennoinen, O. The pseudokinase domain is required for suppression of basal activity of Jak2 and Jak3 tyrosine kinases and for cytokine-inducible activation of signal transduction. *J. Biol. Chem.* **2002**, *277*, 47954–47963. [CrossRef]
10. Hammarén, H.M.; Virtanen, A.T.; Raivola, J.; Silvennoinen, O. The regulation of JAKs in cytokine signaling and its breakdown in disease. *Cytokine* **2019**, *118*, 48–63. [CrossRef]
11. Coltoff, A.; Mascarenhas, J. Fedratinib in 2025 and beyond: Indications and future applications. *Blood Adv.* **2025**, *9*, 1907–1915. [CrossRef]
12. Ross, D.M.; Babon, J.J.; Tvorogov, D.; Thomas, D. Persistence of myelofibrosis treated with ruxolitinib: Biology and clinical implications. *Haematologica* **2021**, *106*, 1244–1253. [CrossRef]
13. Pardanani, A.; Tefferi, A. How I treat myelofibrosis after failure of JAK inhibitors. *Blood* **2018**, *132*, 492–500. [CrossRef]
14. Daver, N.; Cortes, J.; Newberry, K.; Jabbour, E.; Zhou, L.; Wang, X.; Pierce, S.; Kadia, T.; Sasaki, K.; Borthakur, G.; et al. Ruxolitinib in combination with lenalidomide as therapy for patients with myelofibrosis. *Haematologica* **2015**, *100*, 1058–1066. [CrossRef]
15. Deshpande, A.; Reddy, M.M.; Schade, G.O.M.; Ray, A.; Chowdary, T.K.; Griffin, J.D.; Sattler, M. Kinase domain mutations confer resistance to novel inhibitors targeting JAK2V617F in myeloproliferative neoplasms. *Leukemia* **2012**, *26*, 708–715. [CrossRef]

16. Gorantla, S.P.; Prince, G.; Osius, J.; Dinesh, D.C.; Boddu, V.; Duyster, J.; von Bubnoff, N. Type II mode of JAK2 inhibition and destabilization are potential therapeutic approaches against the ruxolitinib resistance driven myeloproliferative neoplasms. *Front. Oncol.* **2024**, *14*, 1430833. [CrossRef]
17. Preuss, F.; Chatterjee, D.; Dederer, V.; Knapp, S.; Mathea, S. Enabling pseudokinases as potential drug targets. *Methods Enzymol.* **2022**, *667*, 663–683. [CrossRef]
18. Wrobleski, S.T.; Moslin, R.; Lin, S.; Zhang, Y.; Spergel, S.; Kempson, J.; Tokarski, J.S.; Strnad, J.; Zupa-Fernandez, A.; Cheng, L.; et al. Highly Selective Inhibition of Tyrosine Kinase 2 (TYK2) for the Treatment of Autoimmune Diseases: Discovery of the Allosteric Inhibitor BMS-986165. *J. Med. Chem.* **2019**, *62*, 8973–8995. [CrossRef]
19. Ouassaf, M.; Zekri, A.; Khan, S.U.; Rengasamy, K.R.R.; Alhatlani, B.Y. Deep learning-guided discovery of selective JAK2-JH2 allosteric inhibitors: Integration of MLP predictive modeling, BREED-based library design, and computational validation. *Front. Chem.* **2025**, *13*, 1646784. [CrossRef]
20. Békés, M.; Langley, D.R.; Crews, C.M. PROTAC targeted protein degraders: The past is prologue. *Nat. Rev. Drug Discov.* **2022**, *21*, 181–200. [CrossRef]
21. Primi, M.C.; Tavares, M.T.; Hatcher, J.M.; Kang, H.; Kresina, T.A.C.; Chakraborty, S.; Leroy, E.; Schmoker, A.M.; Jeon, H.; Weisberg, E.L.; et al. Allosteric inhibition of JAK2 with lysine-reactive compounds that bind the pseudokinase domain. *Eur. J. Med. Chem.* **2026**, *302*, 118274. [CrossRef]
22. Liosi, M.E.; Krimmer, S.G.; Newton, A.S.; Dawson, T.K.; Puleo, D.E.; Cutrona, K.J.; Suzuki, Y.; Schlessinger, J.; Jorgensen, W.L. Selective Janus Kinase 2 (JAK2) Pseudokinase Ligands with a Diaminotriazole Core. *J. Med. Chem.* **2020**, *63*, 5324–5340. [CrossRef]
23. Gaulton, A.; Hersey, A.; Nowotka, M.; Bento, A.P.; Chambers, J.; Mendez, D.; Mutowo, P.; Atkinson, F.; Bellis, L.J.; Cibrián-Uhalte, E.; et al. The ChEMBL database in 2017. *Nucleic Acids Res.* **2017**, *45*, D945–D954. [CrossRef]
24. Huang, N.; Shoichet, B.K.; Irwin, J.J. Benchmarking sets for molecular docking. *J. Med. Chem.* **2006**, *49*, 6789–6801. [CrossRef]
25. van Westen, G.J.; Gaulton, A.; Overington, J.P. Chemical, target, and bioactive properties of allosteric modulation. *PLoS Comput. Biol.* **2014**, *10*, e1003559. [CrossRef]
26. Noonan, T.; Denzinger, K.; Talagayev, V.; Chen, Y.; Puls, K.; Wolf, C.A.; Liu, S.; Nguyen, T.N.; Wolber, G. Mind the Gap—Deciphering GPCR Pharmacology Using 3D Pharmacophores and Artificial Intelligence. *Pharmaceuticals* **2022**, *15*, 1304. [CrossRef]
27. Liu, X.; Ye, K.; van Vlijmen, H.W.T.; IJzerman, A.P.; van Westen, G.J.P. DrugEx v3: Scaffold-constrained drug design with graph transformer-based reinforcement learning. *J. Cheminform.* **2023**, *15*, 24. [CrossRef]
28. Landrum, G. RDKit: Open-Source Cheminformatics. Available online: http://www.rdkit.org (accessed on 2 March 2026).
29. Pajouhesh, H.; Lenz, G.R. Medicinal chemical properties of successful central nervous system drugs. *NeuroRx* **2005**, *2*, 541–553. [CrossRef]
30. Fabbro, D.; Cowan-Jacob, S.W.; Moebitz, H. Ten things you should know about protein kinases: IUPHAR Review 14. *Br. J. Pharmacol.* **2015**, *172*, 2675–2700. [CrossRef]
31. Ertl, P.; Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *J. Cheminform.* **2009**, *1*, 8. [CrossRef]
32. Brenk, R.; Schipani, A.; James, D.; Krasowski, A.; Gilbert, I.H.; Frearson, J.; Wyatt, P.G. Lessons learnt from assembling screening libraries for drug discovery for neglected diseases. *ChemMedChem* **2008**, *3*, 435–444. [CrossRef]
33. Jamieson, C.; Moir, E.M.; Rankovic, Z.; Wishart, G. Medicinal chemistry of hERG optimizations: Highlights and hang-ups. *J. Med. Chem.* **2006**, *49*, 5029–5046. [CrossRef]
34. Bickerton, G.R.; Paolini, G.V.; Besnard, J.; Muresan, S.; Hopkins, A.L. Quantifying the chemical beauty of drugs. *Nat. Chem.* **2012**, *4*, 90–98. [CrossRef]
35. Butina, D. Unsupervised Data Base Clustering Based on Daylight's Fingerprint and Tanimoto Similarity: A Fast and Automated Way To Cluster Small and Large Data Sets. *J. Chem. Inf. Comput. Sci.* **1999**, *39*, 747–750. [CrossRef]
36. Rogers, D.; Hahn, M. Extended-connectivity fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754. [CrossRef]
37. Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J. Chem. Inf. Comput. Sci.* **1988**, *28*, 31–36. [CrossRef]
38. Wang, S.; Witek, J.; Landrum, G.A.; Riniker, S. Improving Conformer Generation for Small Rings and Macrocycles Based on Distance Geometry and Experimental Torsional-Angle Preferences. *J. Chem. Inf. Model.* **2020**, *60*, 2044–2058. [CrossRef]
39. Halgren, T.A. MMFF VI. MMFF94s option for energy minimization studies. *J. Comput. Chem.* **1999**, *20*, 720–729. [CrossRef]
40. Virtanen, A.T.; Haikarainen, T.; Sampathkumar, P.; Palmroth, M.; Liukkonen, S.; Liu, J.; Nekhotiaeva, N.; Hubbard, S.R.; Silvennoinen, O. Identification of Novel Small Molecule Ligands for JAK2 Pseudokinase Domain. *Pharmaceuticals* **2023**, *16*, 75. [CrossRef]
41. Davis, R.R.; Li, B.; Yun, S.Y.; Chan, A.; Nareddy, P.; Gunawan, S.; Ayaz, M.; Lawrence, H.R.; Reuther, G.W.; Lawrence, N.J.; et al. Structural Insights into JAK2 Inhibition by Ruxolitinib, Fedratinib, and Derivatives Thereof. *J. Med. Chem.* **2021**, *64*, 2228–2241. [CrossRef]

42. Olsson, M.H.M.; Søndergaard, C.R.; Rostkowski, M.; Jensen, J.H. PROPKA3: Consistent Treatment of Internal and Surface Residues in Empirical pKa Predictions. *J. Chem. Theory Comput.* **2011**, *7*, 525–537. [CrossRef]
43. McNutt, A.T.; Li, Y.; Meli, R.; Aggarwal, R.; Koes, D.R. GNINA 1.3: The next increment in molecular docking with deep learning. *J. Cheminform.* **2025**, *17*, 28. [CrossRef]
44. Trott, O.; Olson, A.J. AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *J. Comput. Chem.* **2010**, *31*, 455–461. [CrossRef]
45. Bowers, K.J.; Chow, D.E.; Xu, H.; Dror, O.; Eastwood, M.P.; Gregersen, B.A.; Klepeis, J.L.; Kolossvary, I.; Moraes, M.A.; Sacerdoti, F.D.; et al. Scalable Algorithms for Molecular Dynamics Simulations on Commodity Clusters. In Proceedings of the 2006 ACM/IEEE Conference on Supercomputing (SC '06), Tampa, FL, USA, 11–17 November 2006; p. 43. [CrossRef]
46. Jorgensen, W.L.; Chandrasekhar, J.; Madura, J.D.; Impey, R.W.; Klein, M.L. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.* **1983**, *79*, 926–935. [CrossRef]
47. Martyna, G.J.; Klein, M.L.; Tuckerman, M. Nosé–Hoover chains: The canonical ensemble via continuous dynamics. *J. Chem. Phys.* **1992**, *97*, 2635–2643. [CrossRef]
48. Martyna, G.J.; Tobias, D.J.; Klein, M.L. Constant pressure molecular dynamics algorithms. *J. Chem. Phys.* **1994**, *101*, 4177–4189. [CrossRef]
49. Kollman, P.A.; Massova, I.; Reyes, C.; Kuhn, B.; Huo, S.; Chong, L.; Lee, M.; Lee, T.; Duan, Y.; Wang, W.; et al. Calculating structures and free energies of complex molecules: Combining molecular mechanics and continuum models. *Acc. Chem. Res.* **2000**, *33*, 889–897. [CrossRef]
50. Yang, M.; Bo, Z.; Xu, T.; Bo, X.; Wang, D.; Zheng, H. Uni-GBSA: An open-source and web-based automatic workflow to perform MM/GB(PB)SA calculations for virtual screening. *Brief. Bioinform.* **2023**, *24*, bbad218. [CrossRef]
51. Onufriev, A.; Bashford, D.; Case, D.A. Exploring protein native states and large-scale conformational changes with a modified generalized born model. *Proteins Struct. Funct. Bioinform.* **2004**, *55*, 383–394. [CrossRef]
52. Frauenfelder, H.; Sligar, S.G.; Wolynes, P.G. The energy landscapes and motions of proteins. *Science* **1991**, *254*, 1598–1603. [CrossRef]
53. Michaud-Agrawal, N.; Denning, E.J.; Woolf, T.B.; Beckstein, O. MDAnalysis: A toolkit for the analysis of molecular dynamics simulations. *J. Comput. Chem.* **2011**, *32*, 2319–2327. [CrossRef]
54. Alavi, S. Statistical Mechanics: Theory and Molecular Simulation. By Mark E. Tuckerman. *Angew. Chem. Int. Ed.* **2011**, *50*, 12138–12139. [CrossRef]
55. Ichiye, T.; Karplus, M. Collective motions in proteins: A covariance analysis of atomic fluctuations in molecular dynamics and normal mode simulations. *Proteins Struct. Funct. Genet.* **1991**, *11*, 205–217. [CrossRef]
56. Sethi, A.; Eargle, J.; Black, A.A.; Luthey-Schulten, Z. Dynamical networks in tRNA:Protein complexes. *Proc. Natl. Acad. Sci. USA* **2009**, *106*, 6620–6625. [CrossRef]
57. Ramachandran, G.N.; Ramakrishnan, C.; Sasisekharan, V. Stereochemistry of polypeptide chain configurations. *J. Mol. Biol.* **1963**, *7*, 95–99. [CrossRef]
58. Mardia, K.V.; Jupp, P.E. *Directional Statistics*; John Wiley & Sons, Inc.: New York, NY, USA, 1999. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.