



MODEL PERFORMANCE SERIES

Astra AI on Ion Channels


How Orbion's five AI models perform across
2,700 ion channels.

AUTHORS

Çağlar Bozkurt^{1,*}, Aniruddh Goteti¹

¹Orbion GmbH, Berlin, Germany

*Correspondence: caglar.bozkurt@orbion.life



Executive Summary

Ion channels are among the most consequential — and most treacherous — protein classes in drug discovery. They are validated targets across pain, epilepsy, arrhythmia, and inflammation; and one of them, the cardiac potassium channel hERG, is the most prominent *anti*-target in small-molecule development, because off-target hERG block prolongs the QT interval and has withdrawn drugs from the market. Yet ion channels are structurally diverse and hard to characterize computationally: pore-forming bundles with selectivity filters and voltage-sensor domains, transmembrane topologies ranging from single-pass auxiliary subunits to 24-transmembrane voltage-gated giants, and drug-binding sites buried in the conduction pathway rather than in a tidy orthosteric pocket.

Orbion’s Astra AI Suite helps ion-channel teams **reduce the experimental search space** — prioritizing where to spend reagents and assay time across construct design, PTM interpretation, ligand-pocket and hERG/pore-block triage, and channelopathy-variant interpretation. This document is a transparent report on how the suite — five AI models — performs across these areas on ion channels. We ran every Astra AI model on **2,700 reviewed ion-channel proteins** from UniProt Swiss-Prot and compared each output against established public experimental references: Swiss-Prot literature annotation [4], PDB co-crystal contacts at 4 Å resolution [5], and Gene Ontology with ancestor-closure matching [6, 7]. Each section reports the headline performance for one prediction area, its known weaknesses, and where it is recommended for production use.

Headline Numbers Across Five Prediction Areas on Ion Channels

- **Protein Function and Family Classification.** 95.9 % correctly called non-enzyme; GO molecular-function term recovered in the top-5 predictions for **95.6 %** of channels; broad-category (“Transporters”) recall 80.9 %.
- **Topology and Membrane Class.** 84.6 % correctly classified as multi-pass membrane proteins — the remainder are genuinely single-pass auxiliary subunits or soluble Ca²⁺-sensors, reflecting true topological diversity.
- **Residue-Level Transmembrane Topology.** Per-residue prediction agrees with UniProt-annotated transmembrane segments at **AUROC 0.97, AUPRC 0.91, F1 0.87** (n = 307). The disorder model reaches AUROC 0.89.
- **Post-Translational Modification (PTM) Site Prediction.** F1 up to **0.90** for disulfide bonds and **0.88** for N-linked glycosylation; phosphorylation 0.45. All 39 modification classes covered; two operating points reported (high-precision and high-recall).
- **Ligand Binding Pocket Prediction.** The aggregate (**56 % ligand-identity recall**, 56 % pocket-success on the 181 channels with co-crystal data) masks a clean biological split, and the split is the result that matters. On **ligand-gated channels**, which have discrete pockets, the model localizes the pocket well. On **voltage-gated channels** — **including the hERG safety case** — it recovers the ligand identity but the pore-block site is structurally diffuse, so the panel should be read as a *ligand hypothesis set, not a residue-level contact map*. This is pocket-level triage, not atomic-contact prediction.

The story below the headlines is uneven — and we say where. Classification and topology are strong and broadly uniform, but somewhat lower than for receptor classes, because the ion-channel set genuinely contains enzymatic proteins (the chloride channel CFTR has ATPase activity) and ligand-gated channels that are equally validly called receptors.

Binding-pocket prediction tracks a clean biological split: it works on ligand-gated channels with discrete pockets and degrades on voltage-gated channels whose drug sites line the pore. Thermostability (ΔT_m) prediction is not included in this volume: experimentally measured thermal shifts for intact ion channels were not available — the one channel with such data (CFTR) is measured on an isolated soluble domain (§4).

The single most important section of this paper is §3, where we walk every Astra AI model through one channel — the cardiac potassium channel hERG (KCNH2, UniProt Q12809) — and show what a program team would do with the integrated output. The rest of the paper is the per-area evidence, and a decision framework (§5) mapping common ion-channel program questions to the prediction-area combinations that help.

What This Document Is. A performance report on Orbion’s Astra AI Suite on the ion-channel class. Every number traces to a versioned source-data artifact, checked against public experimental references — Swiss-Prot, the PDB, and Gene Ontology — that any reader can independently cross-check; per-protein supplementary data is available on request (§6). The classification, topology and PTM metrics characterize how the deployed models behave on the ion-channel class in production — including proteins drawn from their training corpora; held-out generalization is reported in the per-model preprints (§6).




Figure 1: Headline classification and topology performance on the 2,700-ion-channel cohort. The amber bar is the residue-level transmembrane AUROC; the remaining bars are the function, family, and membrane-class classification outputs. (Ion channels fall under the model’s “Transporters” functional category — the taxonomy has no dedicated ion-channel class.) Full per-area numbers are reported in §2.1 and §2.2.

1 Why Ion Channels

Ion channels conduct the electrical signalling of every excitable cell, and they are targeted by drugs across cardiology, neurology, pain, and immunology — local anaesthetics and anti-arrhythmics on Na^+ channels, anticonvulsants and analgesics on Ca^{2+} channels, and a broad pharmacology on K^+ channels, TRP channels, and the ligand-gated receptor-channels. They are also the source of the most important safety liability in small-molecule discovery: the cardiac potassium channel hERG, whose promiscuous pore binds a chemically diverse set of drugs and whose blockade prolongs the QT interval. hERG screening is a standard, regulator-expected step in small-molecule development (ICH S7B). Characterizing ion channels — their topology, their modification state, and their drug-binding sites — is therefore both a target-discovery and a safety problem.

Structural Diversity. Unlike the uniform seven-transmembrane architecture of G protein-coupled receptors, ion channels span a wide range of folds. Voltage-gated K^+ channels assemble as tetramers of 6-transmembrane subunits around a central pore; voltage-gated Na^+ and Ca^{2+} channels fold a single chain into four 6-transmembrane repeats (24 transmembrane helices); TRP channels, ligand-gated cys-loop receptors, ionotropic glutamate receptors, and chloride channels each have distinct architectures; and the set includes single-pass auxiliary subunits and soluble Ca^{2+} -sensors that regulate the conducting subunits. A characterization model must handle this diversity without a single canonical template.

The Pore Is the Pharmacology. A channel’s drug-binding site is typically the conduction pathway itself — the central cavity below the selectivity filter, lined by the pore helices of all subunits — rather than a discrete surface pocket. This makes binding-site prediction structurally different from the receptor case, and it is the reason promiscuous pore-blockers (the hERG problem) are hard to anticipate from sequence alone.

Regulatory Complexity. Channel trafficking and gating are tuned post-translationally: phosphorylation of long cytoplasmic C-termini and loops, N-linked glycosylation of extracellular turrets that governs surface expression, and disulfide bonds that staple the pore turret. Predicting these sites from sequence is necessary to interpret channelopathy variants and to design stable constructs for structural work.

Why a Unified AI Approach. The Astra AI Suite is built on a shared protein-language-model foundation — a common sequence representation feeding task-appropriate model architectures for each prediction problem. The suite was validated on this premise: one consistent input pipeline, evaluated transparently per protein family. The remainder of this document reports how the suite performs on the 2,700 reviewed ion channels in UniProt Swiss-Prot [4], one prediction area at a time, with the experimental reference data and metrics for each.




Figure 2: Canonical tetrameric voltage-gated channel architecture and the prediction tasks the Astra suite addresses. Four subunits surround a central pore; the selectivity filter and pore-lining helices form the drug-binding conduction pathway, the voltage-sensor domains flank the pore module, extracellular turrets carry glycosylation and pore-stabilizing disulfides, and the cytoplasmic termini carry the regulatory phosphorylation sites.

2 The Astra AI Suite, Capability by Capability

The Astra AI Suite shares a common protein-language-model foundation (a sequence representation augmented with structural and physicochemical features), with each model built on the machine-learning technique best suited to its task — spanning transformer networks, graph neural networks, and regression models. The five models evaluated here were each run on the 2,700-ion-channel cohort against established public experimental references: Swiss-Prot literature annotation [4], PDB co-crystal heavy-atom contacts at 4 Å resolution [5], and Gene Ontology with ancestor-closure semantic matching [6, 7]. Methodological details, metric definitions, evaluation scope, and per-area cohorts are in §6. Thermostability (ΔT_m) prediction is omitted from this volume for lack of intact-channel thermal-shift reference data (§4).

Across the suite, the models are calibrated to be conservative: where signal is weak they tend to omit a prediction rather than fabricate one, so the dominant failure mode is a missed call rather than a false positive — the safer error for wet-lab triage. For binding-site prediction this property is reinforced by a seven-gate anti-hallucination audit (§2.5).

2.1 Protein Function and Family Classification

What We Predict. From sequence alone, the protein’s broad functional category (for ion channels, the expected label is “Transporters”), its Enzyme Commission family, its molecular-function Gene Ontology terms, and its pathway memberships. Evaluated on all 2,700 reviewed ion channels in the cohort.

Headline Numbers.

- **Enzyme-vs-non-enzyme accuracy:** **95.9 %**; Brier score 0.041.
- **GO molecular-function top-5 (ancestor closure):** **95.6 %**; top-1: 82.3 %.
- **“Transporters” category recall:** **80.9 %** (2,184 / 2,700); mean predicted probability when correct: 0.99.

Strong and Weak. Classification is solid but measurably harder than for the receptor classes, for biologically honest reasons. The 6.0 % of channels called “Receptors” rather than “Transporters” are dominated by ligand-gated channels — which *are* ionotropic receptors, so the call is defensible. The enzyme false-positive rate is also higher than for receptors because the cohort genuinely contains catalytic proteins: the chloride channel CFTR carries ABC-ATPase activity and is confidently (and not unreasonably) flagged as an enzyme. GO calibration is well-behaved and monotonic — accuracy rises smoothly from 35 % in the lowest-confidence bucket to 98 % in the highest.

Use as: production triage — confirms a sequence is a channel/transporter and assigns GO molecular function; treat the enzyme call with awareness that channels with catalytic domains exist.

2.2 Topology and Membrane Class

What We Predict. Membrane topology (soluble / single-pass / multi-pass), transmembrane-helix count, subcellular localization, and related structural classifications.

Headline Numbers.

- Membrane class “Multi-pass”: **84.6 %** (2,266 / 2,680).

Strong and Weak. Unlike GPCRs — where 100 % of the cohort is correctly multi-pass — the ion-channel cohort is genuinely topologically diverse: 11 % are single-pass (auxiliary β/γ subunits such as the SCN α B sodium-channel subunits and KCNE potassium-channel subunits), and 3 % are soluble (calmodulin and other Ca²⁺-sensors that regulate channels without spanning the membrane). The 84.6 % multi-pass call therefore reflects the true composition of the channel proteome rather than a model error; the pore-forming α -subunits are called multi-pass at high confidence.

Use as: production triage — structural-class confirmation, with the caveat that auxiliary subunits and Ca²⁺-sensors are correctly *not* multi-pass.

2.3 Residue-Level Transmembrane Topology and Disorder

What We Predict. For each residue, the probability of lipid-bilayer embedding, intracellular / extracellular orientation, intrinsic disorder, and amyloidogenicity. Evaluated on the 307 ion channels with UniProt transmembrane-feature annotations and the 186 with annotated disordered regions.

Headline Numbers.

- Transmembrane prediction: per-protein **mean AUROC 0.969**, **AUPRC 0.913**, F1 at threshold-optimal cutoff 0.872, recall at $k = n_{\text{TM residues}}$ 0.851, Brier 0.039.
- Disorder prediction ($n = 186$): mean AUROC 0.891, AUPRC 0.505.

Strong and Weak. Transmembrane prediction is excellent and uniform across most channel families (per-family mean AUROC 0.96–0.997 for voltage-gated K⁺ and Na⁺, Ca-activated K⁺, ligand-gated cation and anion channels, ionotropic glutamate, CNG/HCN, and TRP). The notable exception is the **voltage-gated Ca²⁺ channels** (CACNA family), where mean AUROC drops to 0.82 — these are the 24-transmembrane single-chain giants whose four repeated domains and very long sequences are the hardest topology in the class. The evaluable set is smaller than the full cohort (307 of 2,700) because only that subset carries human-quality UniProt transmembrane annotation; auxiliary subunits and intracellular channels are not part of it.

Use as: production triage for transmembrane topology on pore-forming subunits; hypothesis generation for the disorder prediction.




Figure 3: Distribution of per-protein transmembrane-residue AUROC across the 307 ion channels with UniProt transmembrane annotations.

2.4 Post-Translational Modification (PTM) Site Prediction

What We Predict. For each residue, the probability that it carries any of 39 post-translational modifications. Two operating points are emitted in parallel: a high-precision call for confident wet-lab handoff and a high-recall call for hypothesis generation [1].

	<i>n</i>	Transp. recall	TM-residue AUROC	Phospho AUROC
Voltage-gated K ⁺ (Kv)	150	1.000	0.988	0.873
Ca-activated / other K ⁺	158	0.962	0.993	0.779
Voltage-gated Na ⁺	25	1.000	0.989	0.804
Voltage-gated Ca ²⁺	77	0.767	0.819	0.876
Ligand-gated cation	107	0.679	0.992	0.705
Ligand-gated anion	85	0.795	0.993	0.826
Ionotropic glutamate	61	0.667	0.997	0.718
TRP	70	0.797	0.989	0.709
Chloride (incl. CFTR)	68	1.000	0.976	0.895
CNG / HCN	39	0.897	0.974	0.972
ENaC / ASIC	50	0.920	0.962	0.913

Cell shade ∝ value (white = low, amber = high). Voltage-gated Ca²⁺ (24-TM) is the weakest topology; ligand-gated families pull lower Transporters recall (they are also ionotropic receptors).

Figure 4: Performance across the ion-channel super-family, broken down by family. Cells shaded by value (white = low, amber = high). Transmembrane prediction is strong and uniform except on the 24-transmembrane voltage-gated Ca²⁺ channels. Ligand-gated families are classified as “Receptors” more often (they are ionotropic receptors), and the chloride family carries elevated enzyme probability (CFTR).

Evaluated on the subset of the cohort with ≥1 Swiss-Prot-curated positive per modification class.

Headline Numbers. Figure 5 reports F1 at the threshold-optimal operating point per modification class. The high-recall operating point dominates the high-precision operating point on the modification classes with sufficient ground truth.

	High-precision F1	High-recall F1	
Disulfide bond	0.64	0.90	<i>n</i> =140
N-linked glycosylation	0.45	0.88	<i>n</i> =199
Phosphorylation	0.41	0.45	<i>n</i> =150
Ubiquitination	0.23	0.36	<i>n</i> =11
S-palmitoylation	0.13	0.20	<i>n</i> =21

Cell shade ∝ F1 score; bold values are the higher of the two operating points.

Figure 5: PTM site prediction F1 at the threshold-optimal operating point on ion-channel Swiss-Prot experimental reference, per modification class.

Strong and Weak. Strongest on the structural modifications: **disulfide bonds** (*n* = 140; F1 0.64 high-precision, **0.90** high-recall) and **N-linked glycosylation** (*n* = 199; 0.45 / **0.88**) — the extracellular pore-turret modifications that govern channel surface expression. Phosphorylation (*n* = 150) is weaker (0.41 / 0.45): ion channels carry long

cytoplasmic termini with many literature-curated phospho-sites, and the conservative output budget recovers only part of them. Most other modification classes have too few annotated ion-channel sites to evaluate.

Use as: production triage at the high-precision operating point for disulfide and glycosylation sites (construct and expression engineering); hypothesis generation at the high-recall operating point.

2.5 Ligand Binding Pocket Prediction

What We Predict. Given a target sequence, a ranked panel of plausible ligand candidates with each ligand’s predicted binding-residue set [3]. A seven-gate anti-hallucination audit verifies position validity, confidence bounds, ligand-ID realism, reproducibility, and biological plausibility on every production batch. Evaluated on the 181 ion channels with both predictions and PDB co-crystal contacts at 4 Å resolution.

How to Read These Numbers. This module is built for *ligand and pocket triage* — which ligands are plausible and roughly where they bind — not atomic-level contact prediction. The recall and pocket-success rates below measure that triage; the residue-level F1 is deliberately conservative because the model reports broader pockets than the strict 4 Å contact set.

Headline Numbers.

- **Ligand-identity recall: 56 %** — of PDB-observed ligand codes per channel, the model predicts 56 % of them.
- **Pocket-success rate: 56 %** — 56 % of model predictions overlap at least one PDB-observed ligand-contact residue set.
- Residue F1 on ligand-ID-matched predictions: 0.17 (median 7 candidate ligands per channel vs. 2 PDB-observed) — as for receptors, the model returns broader pockets than the strict 4 Å ground truth, the expected signature of pocket-level rather than contact-level prediction.

Strong and Weak. Binding-pocket prediction tracks a clean biological split. It is strongest on **ligand-gated channels with discrete binding pockets** — P2X3 purinergic (F1=0.40 on 3 of 4 matched ligands), HCN2 (0.40 on 8 of 9), the GluD2 glutamate-receptor channel (0.36), the nicotinic CHRNA4 (0.33), and the IP₃ receptor ITPR3 (0.31). It is weakest on **voltage-gated channels** (KCNQ1, KCNK potassium channels), where the model correctly identifies the ligand but the drug site lines the diffuse conduction pore rather than a compact pocket, so residue overlap is low. This split mirrors the underlying biology: discrete pockets are predictable, pore-block sites are not yet.

Use as: ligand-panel triage and pocket localization on ligand-gated and small-molecule-targeted channels; for voltage-gated pore-block liabilities (including the hERG safety case), treat the ligand panel as a hypothesis set rather than a residue-level map.




Figure 6: Top 10 ion channels by binding-residue F1 on ligand-identity-matched predictions — a pocket-level triage metric, not atomic-contact prediction. Ligand-gated channels (purinergic, nicotinic, glutamate, IP₃-receptor) dominate the top of the ranking.

3 Worked Example: hERG Across the Full Suite

The headline numbers in §2 aggregate across thousands of channels. In practice, a discovery or safety team works one channel at a time. This section walks every Astra AI model through a single channel — the cardiac potassium channel hERG (gene KCNH2, UniProt Q12809, 1,159 residues) — and shows what the integrated output looks like on the protein that matters most in small-molecule safety. Drug-induced hERG block is the leading single cause of cardiac small-molecule attrition; its promiscuous pore binds antiarrhythmics, antihistamines, antipsychotics, antibiotics, and kinase inhibitors, and predicting which scaffolds will hit it remains an open problem. A model that can characterize hERG from sequence — topology, modification state, and the pore ligand panel — is directly useful to every program that must clear a hERG screen.

Figure 7 summarizes the per-area output on hERG in a single structural view. The function-and-family classification returns **Transporters at P = 0.999** and **Non-enzyme at P = 1.000**. The topology classification returns **multi-pass membrane at P = 0.995** and a transmembrane-helix-count call of **4–6 helices at P = 0.97**, correct for the 6-transmembrane Kv-type subunit fold. The residue-level transmembrane prediction is excellent: per-residue **AUROC 0.994** against the 126 UniProt-annotated transmembrane residues, cleanly resolving the six membrane spans of the pore module and voltage sensor.

The PTM site prediction returns a large phosphorylation set concentrated in hERG’s long (~400-residue) cytoplasmic C-terminus and N-terminal PAS domain — consistent with the dense regulatory phosphorylation literature on hERG trafficking and gating — together with the extracellular pore-turret modifications that govern surface expression.

The binding-pocket prediction is the most striking result on this target. AstraBIND returns a 12-ligand panel for hERG that recovers **all three of its PDB-cocrystallized ligands** (chemical-component codes XB7, A1L2J, and 1II) and, in the same panel, names two mechanistically distinct classes of hERG ligand. First, **small-molecule central-cavity blockers** — **E4031**, the methanesulfonanilide reference blocker that binds the

S6-lined inner cavity below the selectivity filter (the residues that define hERG drug pharmacology). Second, **outer-pore peptide toxins** — the scorpion toxins **BeKm-1** and **ErgTx-1**, which bind the extracellular pore turret rather than the inner cavity. The two sites are distinct, and the predicted contact residues span both — the S6 central cavity and the selectivity-filter turret — which together make up hERG’s promiscuous drug-binding surface.




Figure 7: Integrated view of all model predictions on the cardiac potassium channel hERG (KCNH2, UniProt Q12809). Four subunits surround the central pore; the pore-lining S6 helices (amber) form the central drug-binding cavity and the selectivity-filter turret forms the outer-pore site, where AstraBIND localizes its ligand panel (recovering the three PDB ligands plus the central-cavity blocker E4031 and the outer-pore scorpion toxins BeKm-1/ErgTx-1). The residue-level topology model resolves the six transmembrane spans at AUROC 0.994; phosphorylation sites cluster in the cytoplasmic N- and C-termini. All predictions are produced from sequence alone, with no structural input and no per-target fine-tuning.

What a Program Team Would Do With This Output. A team facing a hERG liability question, or engineering hERG for a structural or electrophysiology campaign, can move directly from this output:

- Use the AstraBIND pore-ligand panel — which recovers known hERG ligands across central-cavity blockers (E4031) and outer-pore toxins (BeKm-1), together with the crystallographic ligands — as a structural hypothesis for which chemotypes engage the pore, and as a starting set for a counter-screen against off-target hERG block.
- Focus pore-cavity mutagenesis on the predicted S6 / selectivity-filter binding residues when probing block mechanism or engineering a block-resistant construct.
- Preserve the extracellular pore-turret disulfide and N-glycosylation sites during construct engineering — both govern hERG surface trafficking, a frequent failure mode in heterologous expression.
- Treat the C-terminal phosphorylation cluster as a hypothesis map for trafficking and gating regulation rather than a confirmed site list.

A team working hERG from sequence alone arrives at a defensible structural and pharmacological picture — pore topology, the promiscuous drug cavity with named blockers, and the trafficking-relevant modification sites — before any wet-lab commitment.

4 Limits and Operating Envelope

Three operating boundaries are worth making explicit. Each is a scope statement, not a recovery roadmap.

Binding-Pocket Prediction Scope — Discrete Pockets, Not Pore-Block Sites.

The 56 % ligand-identity recall and 56 % pocket-success rate in §2.5 are aggregate metrics that mask a clean biological split. On ligand-gated channels with discrete binding pockets (purinergic P2X, nicotinic, ionotropic glutamate, IP₃ receptors) the model localizes the pocket well. On voltage-gated channels — including the hERG safety case — the drug-binding site lines the conduction pore and is structurally diffuse; the model recovers the ligand identity but its residue-level pocket overlaps the strict 4 Å contact set poorly. For voltage-gated pore-block questions, the ligand panel should be used as a hypothesis set, not an atomic-level contact map.

Thermostability (ΔT_m) Prediction — Not Evaluated on This Class. Orbion’s thermostability model is validated on receptor and soluble-protein mutation panels. The only ion-channel thermal-shift data we hold is for CFTR, and it measures the isolated NBD1 cytoplasmic domain (a soluble fragment) rather than the intact membrane channel — a single protein, on a domain readout, and not run through the thermostability model. That is not a representative basis for a channel-thermostability benchmark, so we make no thermostability claim for ion channels in this volume. Closing this gap — with thermal-shift data on intact channels of the kind used to enable channel cryo-EM — is the priority for a future revision.

Functional Classification — Catalytic and Ligand-Gated Channels. The enzyme-vs-non-enzyme call is 95.9 % accurate, with the residual error concentrated on channels that genuinely carry catalytic domains (the ABC-ATPase chloride channel CFTR is confidently flagged as an enzyme). Likewise, ligand-gated channels are sometimes assigned to the “Receptors” category rather than “Transporters”. Both behaviours are biologically defensible rather than failures, but a pipeline that gates strictly on these labels should account for them.

5 Decision Framework: Which Models for Which Question

The Astra AI Suite is most useful when used in combination. The table below maps four common ion-channel program questions to the prediction-area combinations that address them, the operating points to use, and the expected output for a discovery or safety team. Numbers in parentheses reference the per-area evidence in §2.

Table 1: Decision framework mapping common ion-channel program questions to combinations of Orbion’s Astra AI Suite. The integrated suite is most valuable when used as a workflow; single-prediction-area use is appropriate for triage but generally leaves information on the table.

Program Question	Recommended Workflow Across Orbion’s Astra AI Suite
1. Assessing a hERG / Pore-Block Liability	The ligand binding pocket prediction returns the pore-ligand panel for the channel; on hERG it recovers known blockers and the crystallographic ligands (§3). Cross-check the predicted pore-lining residues against the residue-level topology prediction to confirm they fall on the S6 / selectivity-filter segments. Treat the panel as a structural hypothesis set for which chemotypes engage the pore, not an atomic contact map (see §4).
2. Pocket Identification on a Ligand-Gated Channel	The ligand binding pocket prediction localizes the discrete orthosteric pocket on ligand-gated channels (purinergic, nicotinic, ionotropic glutamate; pocket-success strongest in this family). Cross-check predicted pocket residues against the residue-level topology prediction to separate extracellular-domain pockets from transmembrane sites. Output is a ranked candidate-ligand panel for screening or de-novo design.
3. Engineering a Channel Construct for Structure / Electrophysiology	The residue-level topology prediction maps the transmembrane spans and flags disordered termini for truncation; PTM site prediction (high-precision) flags the extracellular pore-turret disulfides and N-glycosylation sites that govern surface expression and must be preserved; the function and family classification confirms the construct’s identity. Output is a construct-boundary and modification-aware design.
4. Interpreting a Channelopathy Variant	The residue-level topology prediction places the variant in the pore, voltage sensor, or cytoplasmic domain; PTM site prediction flags variants that disrupt a conserved modification site (<i>e.g.</i> loss of a glycosylation or disulfide position); the ligand binding pocket prediction indicates whether the variant lies in the drug-binding pathway. ClinVar or gnomAD missense variants can be triaged by predicted structural context.

The recurring pattern across all five workflows is the same: a primary prediction area produces a candidate set, and one or two complementary areas filter or contextualize the candidates against other protein properties. This is the integration story — Orbion’s Astra AI Suite is built so that one sequence-input call returns enough complementary outputs to make multi-criteria decisions.

About Orbion

Orbion is an AI-powered protein engineering platform that compresses the workflow from sequence to experiment-ready protocol. The Astra AI Suite described in this document is the prediction layer of the platform — functional annotation, post-translational modification site prediction, ligand binding-site identification, residue-level topology and

disorder, and thermostability prediction across ΔT_m and $\Delta\Delta G$. Sitting alongside it on the same platform are AlphaFold2 / AlphaFold-Multimer structure prediction, a construct engineering workspace with composite scoring against an organization’s expression vector library, automated experimental protocol generation across expression / purification / crystallization / cryo-EM / stabilization, and a mutation engine for thermostabilization campaigns. Orbion was founded in 2024, is headquartered in Berlin, and works with contract research organizations, pharmaceutical and biotech teams, and academic groups. Free access is available to academic users with a partnership agreement.

Platform. app.orbion.life **Web.** orbion.life **Contact.** contact@orbion.life **Academic access.** orbion.life/academic

6 Methods

Cohort. 2,700 reviewed ion-channel proteins from UniProt Swiss-Prot [4] by ion-channel and channel-activity keywords (KW-1071 / KW-0407 and related), all organisms, de-duplicated against sequence-integrity checks. Per-area evaluation subsets are stated inline in each §2 subsection; they vary because not every channel carries the experimental reference needed for a given task (notably, only 307 carry human-quality transmembrane annotation and 181 have PDB co-crystal contacts).

Experimental Reference. UniProt feature table for sequence-level features (transmembrane segments, modified residues, glycosylation sites, disulfide bonds, regions). Gene Ontology molecular-function annotations expanded to the `is_a` / `part_of` ancestor closure via QuickGO [6, 7]. PDB ligand-contact reference (§2.5): heavy-atom contacts at 4.0 Å between each ligand and protein residue, aggregated across all deposited co-crystal structures per UniProt accession [5].

Metrics. AUROC, AUPRC [8], F1 at threshold-optimal cutoff, F1 at 0.5, recall at $k = n_{\text{positives}}$, Brier score, and Pearson / Spearman correlations. Per-protein metrics aggregated by mean; distributions reported where shape is informative.

Evaluation Scope. The classification, topology, and PTM metrics in §2.1–§2.4 measure how the deployed Astra models perform on the ion-channel class as they are run in production. The cohort is drawn from the reviewed proteome, which overlaps the corpora these models were trained on; the figures therefore characterize within-distribution behavior on real targets rather than held-out generalization. Held-out test-split performance for each model is reported in its respective preprint [1–3]. The contribution of this report is the per-family ion-channel performance breakdown and the integrated single-channel workflow (§3). Thermostability (ΔT_m) prediction is not evaluated in this volume: experimentally measured thermal shifts were available only for an isolated soluble channel domain (CFTR’s NBD1), not intact channels (§4).

Reproducibility and Data Availability. Every aggregate number traces to a versioned source-data artifact in Orbion’s internal evaluation repository, alongside the published methods [1–3]. The underlying experimental references are public: the cohort is the reviewed ion-channel set in UniProt Swiss-Prot, binding ground truth is PDB co-crystal structures, and functional ground truth is Gene Ontology via QuickGO. As is normal in a commercial setting, we do not release the full internal evaluation data; but the cohort accession list, ground-truth sources, and per-protein supplementary data are available on request, so the headline numbers can be independently cross-checked.

References

- [1] Bozkurt, Ç., Vasilyeva, A., Goteti, A. AstraPTM2: A Context-Aware Transformer for Broad-Spectrum PTM Prediction. *bioRxiv* 2025.10.03.680341 (2025). doi:10.1101/2025.10.03.680341.
- [2] Bozkurt, Ç., Vasilyeva, A., Goteti, A. AstraROLE2 & AstraSUIT2: Multi-Task Annotation Models for Functional Profiling of Proteins. *bioRxiv* 2025.06.21.660734 (2025). doi:10.1101/2025.06.21.660734.
- [3] Goteti, A., Vasilyeva, A., Bozkurt, Ç. AstraBIND: Graph Attention Network for Predicting Ligand Binding Sites. *bioRxiv* 2025.11.10.687555 (2025). doi:10.1101/2025.11.10.687555.
- [4] The UniProt Consortium. UniProt: the Universal Protein Knowledgebase in 2023. *Nucleic Acids Research*, **51**(D1):D523–D531 (2023). doi:10.1093/nar/gkac1052.
- [5] Burley, S. K. *et al.* RCSB Protein Data Bank: biological macromolecular structures enabling research and education in fundamental biology, biomedicine, biotechnology and energy. *Nucleic Acids Research*, **47**(D1):D464–D474 (2019). doi:10.1093/nar/gky1004.
- [6] Binns, D. *et al.* QuickGO: a web-based tool for Gene Ontology searching. *Bioinformatics*, **25**(22):3045–3046 (2009). doi:10.1093/bioinformatics/btp536.
- [7] Ashburner, M. *et al.* Gene Ontology: tool for the unification of biology. *Nature Genetics*, **25**:25–29 (2000). doi:10.1038/75556.
- [8] Davis, J., Goadrich, M. The relationship between Precision-Recall and ROC curves. *Proceedings of the 23rd International Conference on Machine Learning*, 233–240 (2006). doi:10.1145/1143844.1143874.