Archived source document. Publication links are updated; scientific text and reported values are retained. Raw and complete data remain in the unpublished Zenodo preparation package.

Back to Model

Complete supporting record

Receptor analysis record

- Receptor data and prediction methods
- Structure preparation and trajectory checks
- Mutation and molecular-pathway methods
- Evidence record and comparison design

Receptor data and prediction methods

Full evidence and evaluation record

What was measured, and what did we compute?

The main benchmark contains 24 Drosophila receptors and 110 odorants: 2,640 receptor–odorant responses in spikes/s (neuronal action potentials recorded per second), with no missing response or duplicate pair ID in the curated table. A response threshold of 50 spikes/s was used for some classification comparisons; it is a project convention, not a universal activation threshold. ([1])

All chemical names were mapped to PubChem (public chemical database used here to identify odorant structures) ([17]). Eleven original names leave stereochemistry unspecified. The full docking (computationally placing and scoring ligand poses in a receptor) ([10]) screen excluded multicomponent ammonium hydroxide and therefore contains 24 × 109 = 2,616 pairs; the main predictive benchmark retains all 2,640 responses. These are different analysis populations.

The sequence resource includes 60 OR-family genes, including Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) ([6]), and 65 protein isoforms (different protein versions produced by a gene). One representative per gene and 60 monomer models from AlphaFold DB (public database of precomputed AlphaFold structure predictions) support feature extraction. A predicted monomer is not the complete functional OR/Orco complex. ([7]; [18])

Generalization is tested at more than one level

Five fixed split definitions were prepared. The main integrated-model comparisons emphasize random pairs, unseen chemical clusters and unseen receptors; defining a split does not mean every model was tested on it.

Table S1. Five prepared holdout definitions and the generalization question each asks. A defined split is not evidence that every model was evaluated on that split.

Split	Question
Random-pair	Predict a held-out combination within known receptor and odorant sets.
Odorant-cold	Predict an odorant excluded from training.
Chemical-cluster-cold	Predict an excluded cluster of chemically similar odorants.
Receptor-cold	Predict a receptor excluded from training.
Receptor-cluster-cold	Predict an excluded sequence-related receptor group.

DoOR checks data consistency

The DoOR (Database of Odorant Responses; integrates published fly odor-response profiles) comparison covers 107 shared odorants and 2,568 pairs. The sequence + ESM (protein language model used here to turn sequences into numerical features) ([9]) model has a median within-receptor Spearman (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage) correlation of 0.938. DoOR includes Hallem measurements, and overlapping Hallem pairs were used for training. This is a consistency check with shared sources, not independent external validation or 93.8% accuracy. ([19])

Full candidate selection and prediction comparisons

Which receptors and sequence changes are worth testing for a selected mold-related odor panel?

Evidence · What happened

Adding receptor sequence features improved prediction for held-out receptors. Or22a (one of the named Drosophila melanogaster odorant receptors in our design panel) ([1]) and Or35a showed useful target-versus-control response differences in the published data.

Learning · What it means

Sequence is informative, but additional structural features did not establish a reliable prediction gain beyond it.

Change · What we do next

Use predicted pockets (a local cavity in the receptor that can accommodate the ligand) to generate mutation hypotheses, then challenge their docking (computationally placing and scoring ligand poses in a receptor) ([10]) and contact explanations.

Explore Design → Build → Test → Learn → Redesign

Design question

Curate a benchmark of 24 Drosophila receptors and 110 odorants from published measurements. Compare a project-defined panel of mold-related volatiles with control odorants; these compounds are not assumed to be unique markers of mold. ([1])

Build the model

Combine molecular descriptors, receptor sequences and AlphaFold (AI system that predicts protein 3D structures from sequence)-derived structural features in matched prediction models. Generate Or22a (one of the named Drosophila melanogaster odorant receptors in our design panel) ([1]) / Or35a point-mutation models with ColabFold (protein-structure prediction workflow built around AlphaFold and fast sequence searches) and dock target and control odorants. The build consists of computational sequences and structures. ([7]; [8])

Test & results

For receptors excluded from training, R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) increased from 0.1125 with ligand (a molecule that can bind to a receptor; here, the tested odorant) features to 0.2325 with sequence features. The full representation reached 0.2356; its additional gain had an uncertainty interval spanning zero. An initial docking (computationally placing and scoring ligand poses in a receptor) ([10]) screen suggested small relative preference shifts for F34Y (residue 34: phenylalanine → tyrosine) and E196Q (residue 196: glutamate → glutamine).

Learn

A useful predictor can help prioritize experiments, but a small docking-score difference does not demonstrate improved selectivity. R^2 describes predictive fit, not percent accuracy.

Redesign

Move from ranking candidates to checking whether the proposed mutations actually contact the odorant and whether their ranking survives changes in docking setup.

Download plotted values (CSV) ↓ — Zenodo (publication pending)

Research evidence

Complete model comparisons

Matched feature ablation

The six feature sets use the same histogram gradient boosting (a tree-ensemble prediction model that successively corrects errors using binned features) settings, the same response pairs and fixed folds. L denotes ligand (a molecule that can bind to a receptor; here, the tested odorant) descriptors; S combines sequence descriptors and ESM2 (protein language model used here to turn sequences into numerical features) ([9]); P adds pocket (a local cavity in the receptor that can accommodate the ligand) features; D adds Vina (AutoDock Vina

docking score; lower scores are more favorable in this model, not a measured response) ([10]; [20]) and contact summaries; M adds multiple-pose features. PubChem (public chemical database used here to identify odorant structures) ([17]) identifiers, response-derived labels and split columns were excluded from the feature list; missing docking (computationally placing and scoring ligand poses in a receptor) descriptors follow the existing zero-fill convention.

Sequence information supplies the clearest gain in both evaluation settings. On unseen receptors, the full model adds only $\Delta R^2 = 0.0031$ over LS, with a receptor-grouped 95% bootstrap interval of −0.0316 to 0.0344. The interval does not establish an additional structural-feature gain.

Table S2. Matched feature ablation using the same response pairs, fixed folds and histogram-gradient-boosting settings. R^2 is reported for unseen chemical clusters and unseen receptors; the small full-model gain over ligand-plus-sequence features does not establish an additional structural-feature benefit.

Features	Chemical-cluster-cold R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse)	Receptor-cold R^2
L	0.1307	0.1125
LS	0.3855	0.2325
LSP	0.3840	0.2221
LSD	0.3781	0.2131
LSPD	0.3837	0.2154
LSPDM	0.3779	0.2356

How the earlier models informed redesign

These historical results use different representative models and feature versions. They are retained as method-development evidence and must not be combined into one model's performance. For the matched historical HGB comparison, $\Delta R^2 \approx 0.0175$ had an interval of −0.0015 to 0.0419. Bootstrap intervals use fixed out-of-fold predictions; they omit retraining and model-selection uncertainty.

Table S3. Historical model-development scores under different holdouts and representative model versions. R^2 and Spearman rank correlation are separate metrics; these rows must not be combined into the performance of one model.

Evaluation / model	R^2	Spearman ρ (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage)

Evaluation / model	R^2	Spearman ρ (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage)
Random-pair / stacked tree	0.569	0.739
Chemical-cluster-cold / global blend	0.420	0.659
Receptor-cold / original HGB	0.229	0.463
Receptor-cold / HGB with multiple poses	0.246	0.472

Docking features did not transfer uniformly

In an earlier ExtraTrees (an ensemble of decision trees with randomized splitting rules) analysis on 2,616 pairs, adding Vina/contact features raised chemical-cluster-cold R^2 from 0.066 to 0.220, while receptor-cold R^2 fell from 0.117 to 0.006. An eight-receptor pocket-redocking study and a small restrained implicit-solvent OpenMM (the software engine that integrates the atomic equations of motion) ([14]) study also failed to establish a stable unseen-receptor improvement. These results motivated better pocket placement, pose checks and membrane-based conformational sampling.

Feature hypotheses, comparison cohorts and interpretation

Purpose · connect both sides of a receptor–odorant pair

A receptor sequence describes the protein, while molecular weight, spectra and learned chemical features describe the odorant. They can enter the same response predictor before docking (computationally placing and scoring ligand poses in a receptor) ([10]) or MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]). Our baseline already combines 540 sequence descriptors, 320 ESM2 (protein language model used here to turn sequences into numerical features) features, 12 physicochemical descriptors including molecular weight, and a 2,048-bit Morgan fingerprint: 2,920 inputs in total. A Morgan fingerprint encodes local molecular connectivity; it is not a 3D structure. ([1]; [9])

We asked three successive questions: does measured IR (infrared absorption: how a molecule absorbs light at different vibrational frequencies) add information; do complete harmonic modes and EVA (a fixed-length descriptor formed by smoothing equally weighted normal-mode frequencies) ([11]) recover information outside that IR window; and can a pretrained odor-perception model supply a useful molecular representation? These are completed prediction experiments, separate from the OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6]) entry/exit simulations.

Method · compare the same eligible pairs within each experiment

HGB (a tree-ensemble prediction model that successively corrects errors using binned features) and ExtraTrees (an ensemble of decision trees with randomized splitting rules) are two tree-based regression methods. We kept their settings fixed and used the existing five-fold chemical-group and receptor holdouts. Chemical-group holdout tests responses to excluded chemical groups; receptor holdout tests excluded receptors with known odors. Neither is a simultaneous unseen-receptor-and-odor test. Random features of matched dimensionality provide a control.

Each eligible subset gets its own retrained baseline. Scores from the 82-, 85-, 67- and 109-odorant subsets are not directly interchangeable. R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) measures prediction error relative to a constant-mean reference; it is not a correlation coefficient or a percentage accuracy. Spearman ρ (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage) measures ranking, and MAE measures absolute response error in spikes/s (neuronal action potentials recorded per second).

Table S4. Eligible odorants, receptor–odorant pairs and chemical groups for each molecular-feature experiment. Each cohort receives its own retrained baseline, so scores from different rows are not directly interchangeable.

Experiment	Odorants	Response pairs	Chemical groups
Measured IR / size ablation	82	1,968	53
Quantum modes	85	2,040	56
Measured–computed overlap	67	1,608	45
OpenPOM (Principal Odor Map; OpenPOM is its open-source reproduction for learned molecular odor features) ([4]; [5])	109	2,616	69

Interpretation · exploratory comparisons with grouped uncertainty

Paired 95% intervals resample chemical groups or receptors using fixed out-of-fold predictions. They account for those grouped observations, but omit retraining, model-selection and assay (a defined experimental test of biological activity) uncertainty. A zero-crossing interval does not establish a reliable gain. These data and splits have been used in earlier work; this is exploratory evaluation, with no multiplicity adjustment for secondary comparisons.

Measured spectra, matched comparisons and controls

Purpose · test added information, not a vibration mechanism

Prior work relates receptor responses to molecular volume and uses physicochemical properties to predict odor perception. Molecular weight and volume are related but distinct quantities. We therefore tested removing weight alone and removing a size-related group, alongside adding measured infrared absorption spectra. ([21]; [22])

Method · 82 matched gas-phase spectra

Of 110 odorants, 82 passed full chemical-identifier, gas-phase source and spectral-coverage checks in the NIST Chemistry WebBook (US National Institute of Standards and Technology; its Chemistry WebBook supplies spectral records). Missing spectra were excluded from all compared models, not zero-filled. We subtracted the first-percentile background, clipped negative absorbance, normalized total area over 650–3,800 cm^{-1} and integrated 63 bands of 50 cm^{-1}. The 4 cm^{-1} source sampling interval is not an instrument-resolution claim. ([2])

All six variants retain the same receptor sequence features. “Sequence + IR (infrared absorption: how a molecule absorbs light at different vibrational frequencies) only” replaces the entire chemical block with IR. The size group is molecular weight, heavy-atom count and molar refractivity (a property related to polarizability and size). These removals test incremental predictive value in this model, not a causal biological effect.

Published measurements

Measured odorant spectra

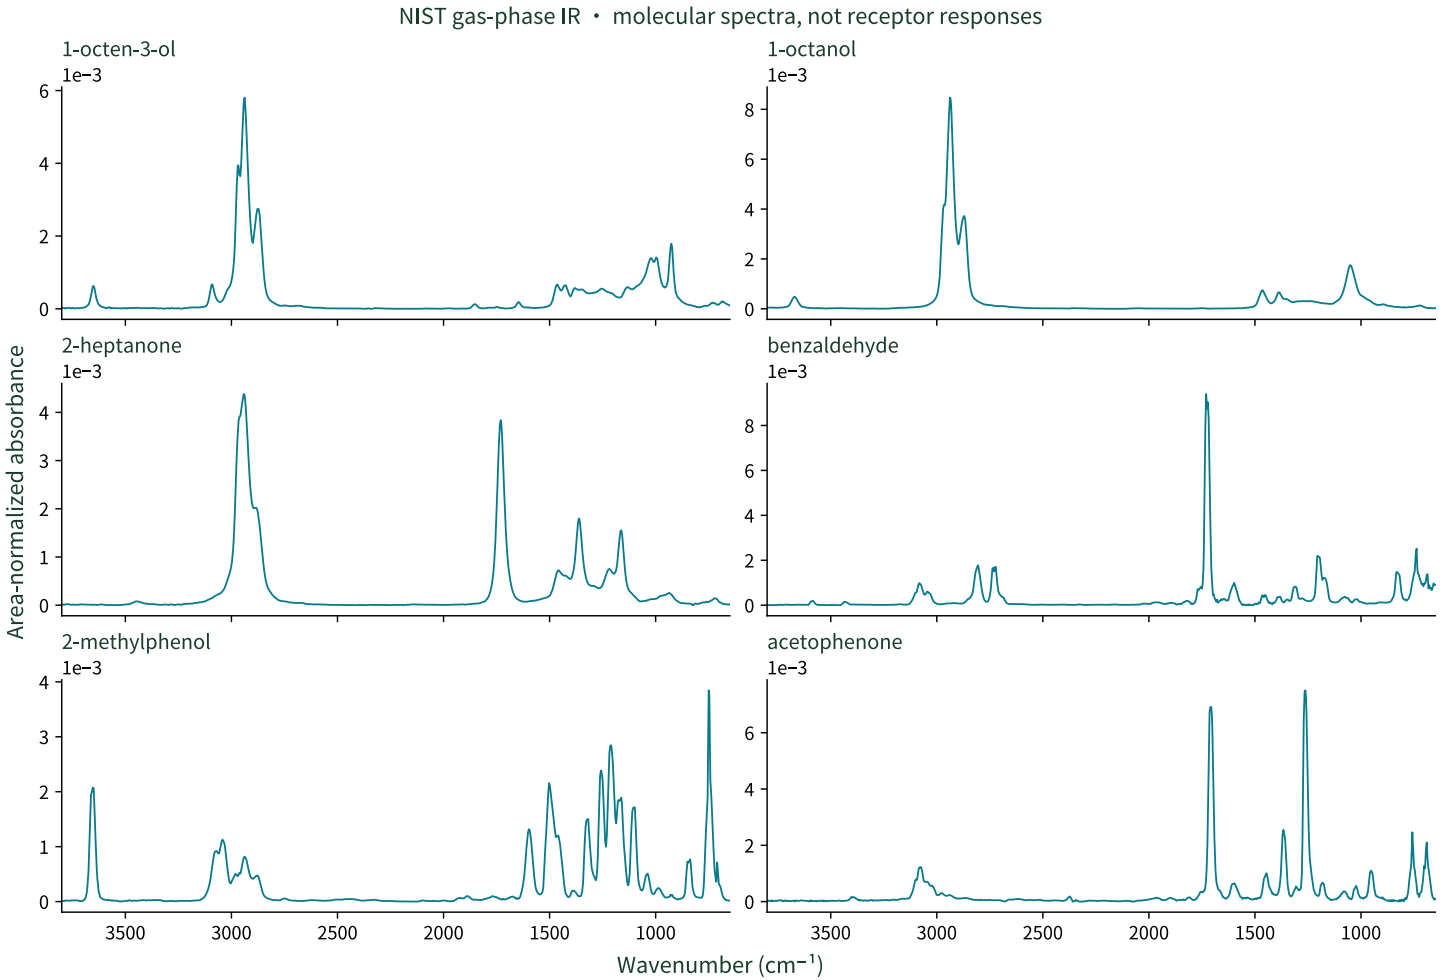

Enlarge ↗

Figure S1. Six accepted NIST (US National Institute of Standards and Technology; its Chemistry WebBook supplies spectral records) gas-phase spectra after background correction and area normalization. Peak patterns describe relative molecular absorption, not receptor activity or absolute absorption strength. ([2])

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

Results · measured IR did not improve new-chemical prediction

On unseen chemical groups, adding IR changed HGB (a tree-ensemble prediction model that successively corrects errors using binned features) R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) from 0.3930 to 0.3278 (Δ −0.0653; 95% interval −0.1127 to −0.0078) and ExtraTrees (an ensemble of decision trees with randomized splitting rules) from 0.3746 to 0.3458 (Δ −0.0288; −0.0564 to −0.0003). Receptor-holdout point gains were small: +0.0028 and +0.0102, with both intervals crossing zero.

Weight alone had model-dependent effects: retaining it changed chemical-holdout R^2 by −0.0121 in HGB but +0.0188 in ExtraTrees. Retaining the full size group improved point scores in both models (+0.0292 / +0.0751), although the HGB interval crossed zero. Weight was already available to the early response model; adding it is not an untested missing step.

The table reports R^2. The following harmonic-mode comparison extends beyond the measured spectral window. These predictor tests do not evaluate the biological mechanism of vibrational sensing.

Table S5. Matched R^2 results for measured infrared features, descriptor removals and random-feature controls on 82 odorants. Chemical and receptor columns denote separate holdouts; the comparison tests predictive value, not a biological vibration mechanism.

Inputs	Chemical · HGB	Chemical · ExtraTrees	Receptor · HGB	Receptor · ExtraTrees
Baseline	0.3930	0.3746	0.2080	0.1000
+ IR 63	0.3278	0.3458	0.2108	0.1103
Sequence + IR only	0.2135	0.3209	0.1679	0.1103
Remove weight	0.4052	0.3557	0.1792	0.1096
Remove size group	0.3638	0.2995	0.2053	0.0949
+ Random 63	0.3194	0.3241	0.2132	0.1078

molecular-features/ir/metrics.csv ↓ — Zenodo (publication pending)molecular-features/ir/paired_uncertainty.csv ↓ — Zenodo (publication pending)molecular-features/ir/coverage.csv ↓ — Zenodo (publication pending)

Full harmonic-mode methods, feature tests and matched cohorts

Purpose and acquisition · recover modes missed by the IR window

We reused published Gaussian quantum-chemistry outputs from Harada et al., a 2025 preprint, rather than claiming new quantum calculations. The method was gas-phase B3LYP/6-31G(d), a density-functional theory approximation and basis set. After checking calculation completion, mode counts, molecular connectivity and specified stereochemistry, 85 odorants were accepted: 4,767 harmonic modes, including 1,082 below 650 cm^{-1}. ([3])

Each molecule provides one source conformer's complete harmonic normal modes: coordinated atom-displacement patterns near an optimized geometry, with frequencies, IR (infrared absorption: how a molecule absorbs light at different vibrational frequencies) intensities, reduced masses and force constants. This does not cover all conformers, anharmonic motion, receptor-bound vibrations or full quantum dynamics. Twenty-two odorants lacked exact-match source files; two had imaginary frequencies; trans-caryophyllene required stereochemical review and was excluded.

Method · encode the same modes in two ways

EVA (a fixed-length descriptor formed by smoothing equally weighted normal-mode frequencies) gives every mode equal weight, smooths its frequency with a Gaussian of $\sigma = 10\ cm^{-1}$, and samples 1–3,996 cm^{-1} every 5 cm^{-1} to form 800 features. Computed IR uses the same modes and grid but weights each mode by its calculated IR intensity. Neither applies area normalization or frequency scaling. This classical EVA formulation is not a reproduction of the 27-dimensional EVA used by Gabler et al. ([11]; [23])

Alongside full 800-feature inputs, we tested PCA (principal component analysis, which compresses features along directions of greatest variation) compression to 16 components and random 800/16-feature controls. PCA was fitted only on distinct training odors within each fold. Its variance retention is not evidence that response information was retained.

Our computational results

Equal-mode EVA versus intensity-weighted IR

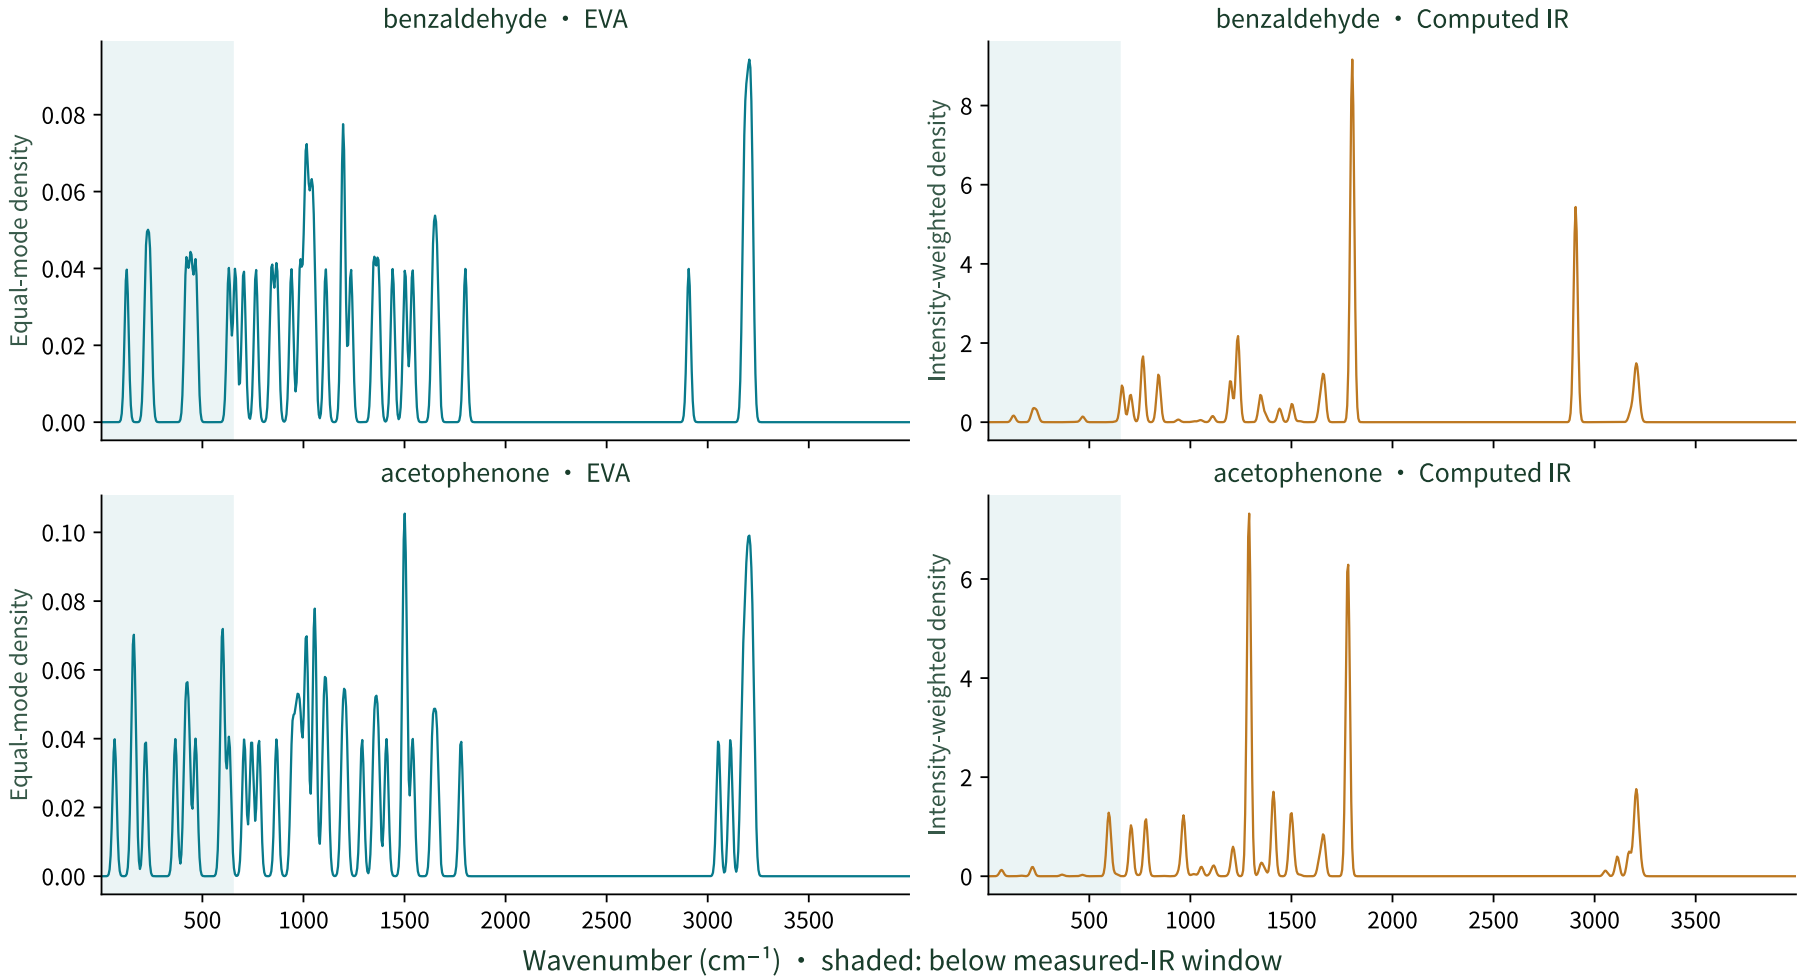

Enlarge ↗

Figure S2. Descriptors computed from two published harmonic mode sets, with the same Gaussian width and grid. The ordinate scales differ because EVA (a fixed-length descriptor formed by smoothing equally weighted normal-mode frequencies) weights modes equally and computed IR (infrared absorption: how a molecule absorbs light at different vibrational frequencies) uses intensities. Shading marks frequencies below the measured-IR window. This is not a measured spectrum or a receptor-bound calculation. ([11]; [3])

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

Results · full and compressed features, 85 odorants

The prespecified primary comparison was HGB (a tree-ensemble prediction model that successively corrects errors using binned features) with added EVA under chemical-group holdout: R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) 0.3810 → 0.3421, Δ −0.0389 (95% interval −0.0800 to +0.0022). The point estimate fell, but this interval crosses zero; it supports neither a reliable gain nor a confirmed decline. ExtraTrees (an ensemble of decision trees with randomized splitting rules) also fell, from 0.4225 to 0.3772. Compressed features came closer to the baseline without exceeding its chemical-holdout point score.

For unseen receptors, HGB rose slightly with full EVA or computed IR, but both intervals crossed zero and random features showed similar point gains. Thus no stable improvement across models and holdouts was established. All table values are R^2.

Table S6. R^2 comparison on 85 odorants for full 800-dimensional EVA or computed-IR features, 16-component training-fold PCA and random-feature controls. Point changes do not establish a stable improvement across models and holdouts.

Inputs	Chemical · HGB	Chemical · ExtraTrees	Receptor · HGB	Receptor · ExtraTrees
Baseline	0.3810	0.4225	0.2217	0.1435
＋EVA 800	0.3421	0.3772	0.2315	0.1175
+ Computed IR 800	0.3177	0.3755	0.2372	0.1014
＋EVA PCA 16	0.3683	0.4186	0.2186	0.1320
+ Computed IR PCA 16	0.3635	0.4164	0.2298	0.1272
+ Random 800	0.2265	0.2812	0.2338	0.1037
+ Random 16	0.3241	0.3987	0.2323	0.1287

Matched comparison · 67 odorants with both data sources

The measured-IR and quantum datasets overlap for 67 odorants. We retrained all baselines on these same 1,608 pairs. None of the spectral additions exceeded the chemical-holdout baseline in either model. Measured IR uses a different range, dimension and normalization, so this is not an isolated test of experimental versus quantum spectral accuracy.

Table S7. R^2 comparison after retraining on the same 67-odorant, 1,608-pair measured/computed-spectrum overlap. Differences in spectral range, dimension and normalization remain, so the comparison does not isolate source accuracy.

Inputs	Chemical · HGB	Chemical · ExtraTrees	Receptor · HGB	Receptor · ExtraTrees
Baseline	0.3744	0.4003	0.2075	0.1094
＋EVA 800	0.2847	0.3488	0.2076	0.0910
+ Computed IR 800	0.2391	0.3643	0.2047	0.0968
+ Measured IR 63	0.3360	0.3809	0.2127	0.1121

molecular-features/quantum/metrics.csv ↓ — Zenodo (publication pending)molecular-features/quantum/paired_uncertainty.csv ↓ — Zenodo (publication pending)molecular-features/quantum/audited_coverage.csv ↓ — Zenodo (publication pending)molecular-features/quantum/normal_modes.csv ↓ — Zenodo (publication pending)

What this says about vibrational theories

Vibrational explanations of smell remain disputed: isotope-related fly behavior has supported further investigation, while receptor experiments and an OR51E2/propionate computational study provide counterevidence in their tested systems. Predictive usefulness of a spectrum would not by itself demonstrate electron tunneling, and our negative feature results do not exclude every vibrational mechanism. Classical MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) here does not simulate electron tunneling. ([24]; [25]; [26])

Full OpenPOM methods, model comparisons and uncertainty

Purpose · from human odor perception to insect receptor responses

The Principal Odor Map learns molecular representations from human odor descriptions. OpenPOM (Principal Odor Map; OpenPOM is its open-source reproduction for learned molecular odor features) is an open-source reproduction, not the original POM model weights. We tested whether its learned chemical representation helps a sequence–odorant response predictor; a “floral” or “mushroom” label cannot replace a measured insect-receptor response. ([4]; [5])

Method · fixed 256-dimensional features, retained molecular weight

We extracted the official 256-dimensional penultimate linear representation from a prespecified public checkpoint in inference mode, without response-label fine-tuning. A second fixed checkpoint was tested separately as a sensitivity comparison. The variants either add POM to the baseline (3,176 total features) or replace only the Morgan fingerprint (1,128 total); all retain physicochemical descriptors including weight and both sequence blocks.

All variants use 109 odorants × 24 receptors. Ammonium hydroxide was excluded because its two isolated ions have no covalent bonds for this graph-model workflow; it was not replaced with ammonia. The upstream model does not encode stereochemistry: 109 odorants map to 106 distinct graphs, including three indistinguishable stereoisomer pairs. Identical graphs receive identical features.

A provenance audit found 108 of 109 non-stereochemical structures in the public 4,983-molecule candidate pretraining pool. Exact checkpoint training/test membership is unavailable. Chemical-group holdout therefore means unseen downstream Hallem responses, not a guarantee that the entire AI system has never encountered those molecular structures or human odor labels.

Results · the direction depends on the response model

The prespecified primary comparison, HGB (a tree-ensemble prediction model that successively corrects errors using binned features) plus POM on unseen response chemical groups, changed R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) from 0.3892 to 0.3232: Δ −0.0661 (95% interval −0.1327 to +0.0189). ExtraTrees (an ensemble of decision trees with randomized splitting rules) instead rose from 0.3572 to 0.3967: Δ +0.0395 (−0.0490 to +0.1334). Replacing the fingerprint and using the second checkpoint showed the same model-dependent direction; every POM R^2-difference interval crossed zero.

ExtraTrees' chemical-holdout R^2 gain did not improve ranking or absolute error: Spearman ρ (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage) changed 0.6465 → 0.6459 and MAE 25.2021 → 25.6071 spikes/s (neuronal action potentials recorded per second). On unseen receptors, adding POM gave small R^2 increases (HGB 0.2108 → 0.2332; ExtraTrees 0.1210 → 0.1321), but random features achieved 0.2312 / 0.1323. These results do not establish a POM-specific, stable gain. Table values are R^2.

Table S8. R^2 results for frozen 256-dimensional OpenPOM features on 109 odorants, including fingerprint replacement, feature addition, random controls and a second fixed checkpoint. Effects depend on the response model; downstream holdout does not guarantee unseen molecules in upstream pretraining.

Inputs	Chemical · HGB	Chemical · ExtraTrees	Receptor · HGB	Receptor · ExtraTrees
Baseline	0.3892	0.3572	0.2108	0.1210
POM replaces fingerprint	0.3251	0.4004	0.2267	0.1320
+ POM	0.3232	0.3967	0.2332	0.1321
+ Random 256	0.2821	0.2727	0.2312	0.1323
Replace: second weights	0.3515	0.4075	0.2057	0.1256
Add: second weights	0.3565	0.4103	0.2185	0.1351

Predicted versus measured response

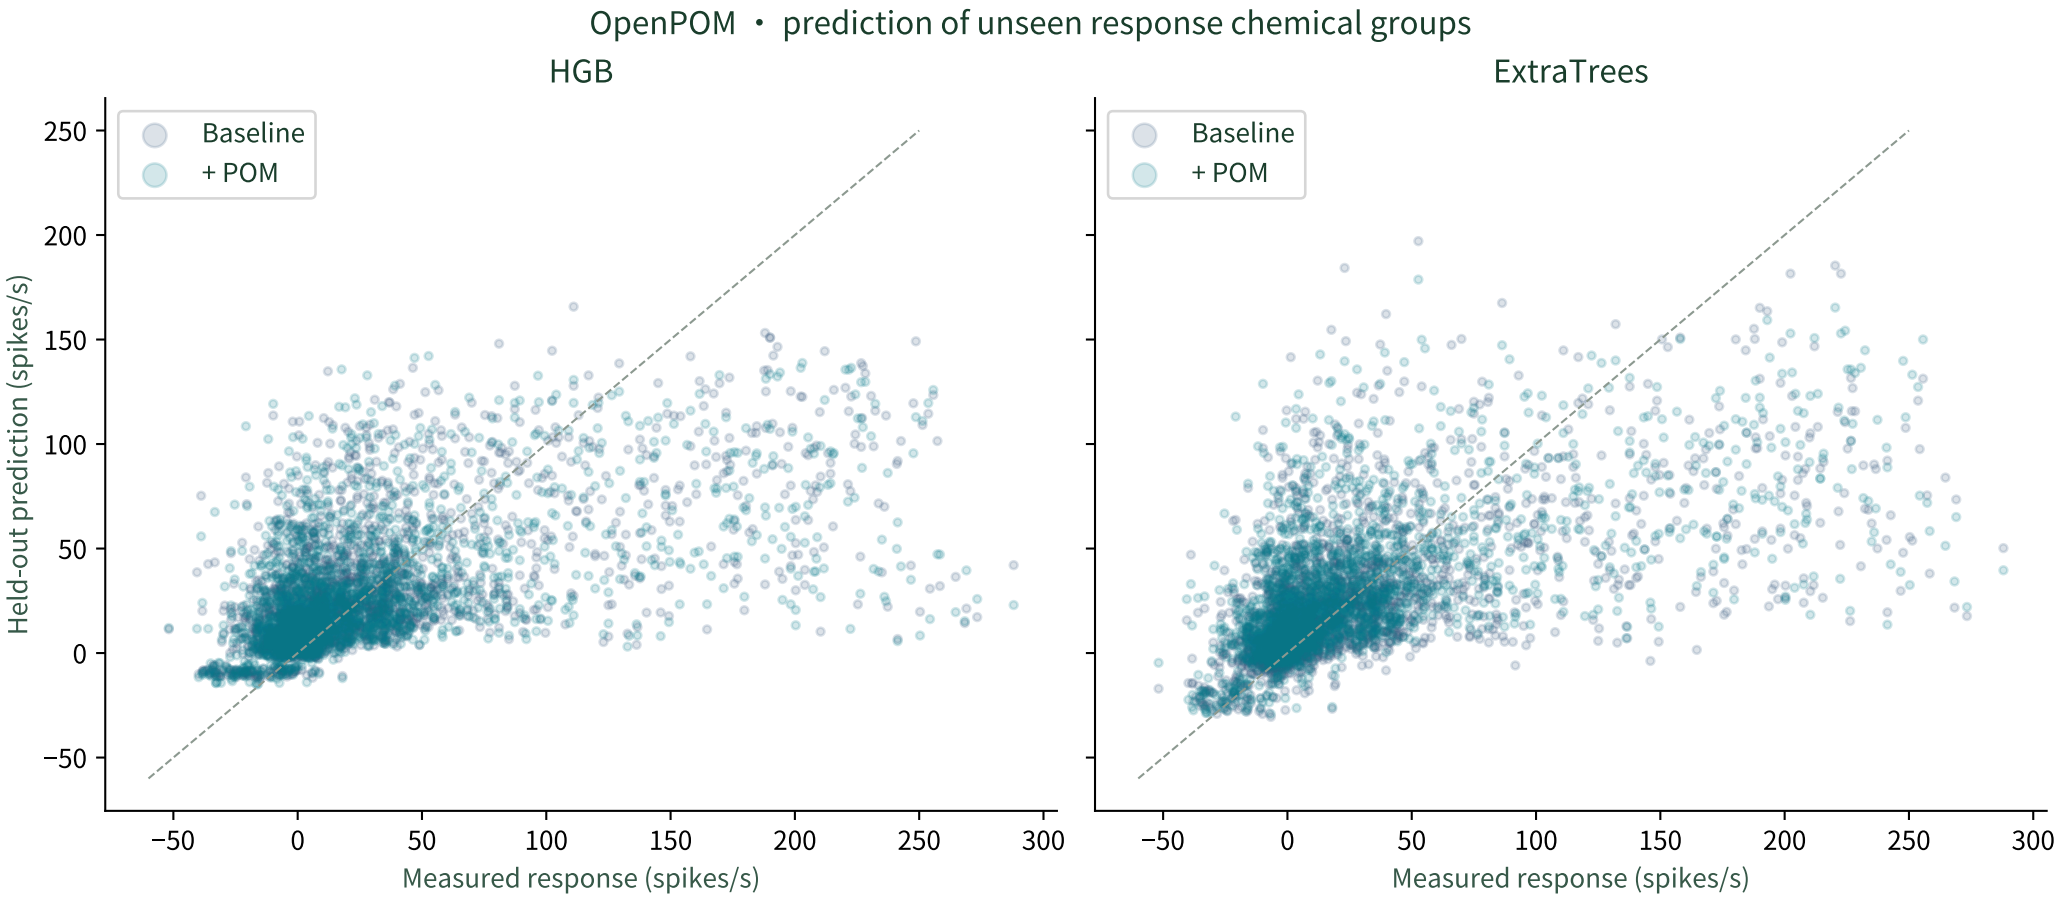

Enlarge ↗

Figure S3. Out-of-fold predictions for the chemical-group holdout, baseline versus added POM (Principal Odor Map; OpenPOM is its open-source reproduction for learned molecular odor features) ([4]; [5]). Each point is a receptor–odorant pair; shared receptors and chemicals mean points are not independent experimental replicates. The dashed line denotes perfect prediction. ([1])

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

molecular-features/pom/metrics.csv ↓ — Zenodo (publication pending)molecular-features/pom/paired_uncertainty.csv ↓ — Zenodo (publication pending)molecular-features/pom/coverage.csv ↓ — Zenodo (publication pending)

Learning and redesign · retain the baseline; validate adaptation next

Measured spectra, harmonic-mode descriptors and frozen AI features were tested. None established a stable general gain over the matched baseline; POM with ExtraTrees was not adopted as a replacement.

Related literature describes low-rank task adaptation of chemical foundation models. LORAX was not part of the completed comparisons reported here. [27]

molecular-features/benchmark-data.zip ↓ — Zenodo (publication pending)molecular-features/provenance.json ↓ — Zenodo (publication pending)

Full cross-species prediction and sequence-fragment tests

Design · distinguish response prediction from analogue matching

We tested whether fly-derived information can help prioritize mosquito receptors. Carey et al. provide 50 Anopheles gambiae receptors × 113 odorants, or 5,650 responses with SEM (standard error of the mean, describing uncertainty in an estimated mean); 40 odorants overlap the fly panel. This Anopheles transfer study is separate from the Aedes OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) molecular-dynamics reference system. ([12])

All 50 historical receptor names have candidate sequences: 46 grade A, two grade B and two grade C. Naming or duplicate-gene ambiguity affects AgOr12 (an Anopheles gambiae odorant receptor; Ag identifies the mosquito species) ([12]) and AgOr67. The main fragment analysis therefore uses 48 receptors with A/B evidence against 24 fly receptors, giving 1,152 receptor–receptor pairs. The initial direct-transfer baseline used all 50 mosquito receptors.

Test 1 · transfer the fly-trained response model

On the 2,000 mosquito receptor–odorant pairs sharing chemical identities with the fly benchmark, the ExtraTrees (an ensemble of decision trees with randomized splitting rules) transfer baseline produced R^2 (coefficient of determination; 1 is perfect, 0 matches the mean baseline, and negative values are worse) = −0.288 and Spearman ρ (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage) = −0.060. Negative R^2 means its squared prediction error exceeded the constant-mean reference on this evaluation set. This baseline does not support reliable species transfer.

Test 2 · match local sequence fragments

The alternative question is whether two receptors respond similarly across the shared 40-odorant panel. We compared full-length identity (fraction of aligned amino-acid positions that are identical) and ESM (protein language model used here to turn sequences into numerical features) ([9]) embeddings (a numerical vector representing a sequence or molecule) with local 21-, 31- and 41-amino-acid windows. Variants down-weight shared sequence framework and species signatures, subtract within-species embedding backgrounds, or require compatible relative position and hydropathy. These are sequence-based topology proxies, not a 3D pocket (a local cavity in the receptor that can accommodate the ligand) mapping.

Reciprocal matching asks whether each fragment selects its partner as the best match in both directions. Activity labels were not used to learn fragment weights. However, methods were compared and selected on the same response matrix, so the selected result still requires receptor-held-out confirmation.

Full-length identity gave ρ = 0.017. The 31-aa (amino-acid residues) mutual-best fraction reached ρ = 0.141 with a receptor-bootstrap (resampling observations or groups to assess statistical uncertainty) interval of 0.035–0.250. The unadjusted block-permutation p was 0.033, but BH (Benjamini–Hochberg multiple-testing correction; q is the adjusted p-value reported here) ([13]) correction across the tested methods gave q = 0.259. The weak exploratory signal did not meet the corrected 0.05 threshold; these correlations concern profile

similarity, not individual ligand (a molecule that can bind to a receptor; here, the tested odorant) activation.

Learn and redesign · validate when a local match is useful

The effect was receptor-dependent: the broader sensitivity analysis improved for AgOr4/AgOr56 and deteriorated for AgOr63/AgOr41. Removing a species mean alone did not solve transfer. The completed transfer comparison covers Drosophila and Anopheles.

For the third-species inventory, the functional literature is de Fouchier et al.'s moth-receptor study, and candidate amino-acid sequences were recovered from the supplementary material of Jacquin-Joly et al. These sources support the inventory, not a completed transfer-validation result. ([28]; [29])

Structure preparation and trajectory checks

Read the structure names

What are 8V00 and 8V02?

They are accession IDs in the Protein Data Bank (PDB), a public archive of molecular 3D structures. Each ID points to a deposited coordinate model and its experimental records. ([30])

Both structures were determined by cryo-electron microscopy (cryo-EM), which reconstructs molecular structures from electron images of frozen samples. Both contain one odor-selective OR10 subunit from Aedes aegypti (the yellow-fever mosquito) and three Orco co-receptor subunits from the fig wasp Apocrypta bakeri. Orco is the shared partner that helps form the odorant-gated ion channel. This is an experimentally assembled complex combining proteins from two species. ([6])

Table S9. Experimental cryo-EM references used for the OR10–Orco studies: odorant-free 8V00 and o-cresol-bound 8V02. Both combine Aedes aegypti OR10 with Apocrypta bakeri Orco; they are structural references rather than target Drosophila receptor measurements.

PDB entry	What was observed	Use in our study
8V00	Apo: no bound odorant. Experimental resolution: 2.88 Å (1 Å = 0.1 nm).	Odorant-free structural reference; a starting template for apo MD and static docking comparisons.
8V02	Odorant-bound (holo): o-cresol (2-methylphenol) is resolved in the OR10 pocket. Experimental resolution: 2.90 Å.	Reference for the bound ligand pose, bound-state MD, and candidate outward-path studies.

MD means molecular dynamics: calculating atomic motion over time. Docking searches and scores candidate ligand positions and orientations. The PDB entries are experimental coordinate models, not simulation trajectories or AlphaFold predictions (AI predictions of protein structures from sequence). Before MD we rebuilt unresolved regions and added a membrane, water and ions; those additions are computational model assumptions. A bound experimental structure establishes an observed pose, not the route the odorant took to reach it. ([6]; [7]; [14]; [10])

Research evidence

Build and check the membrane models

Correcting the membrane-system geometry resolved the observed short-range periodic-image issue, while rebuilt Orco regions remained uncertain structural hypotheses.

Uncertainty in rebuilt Orco regions

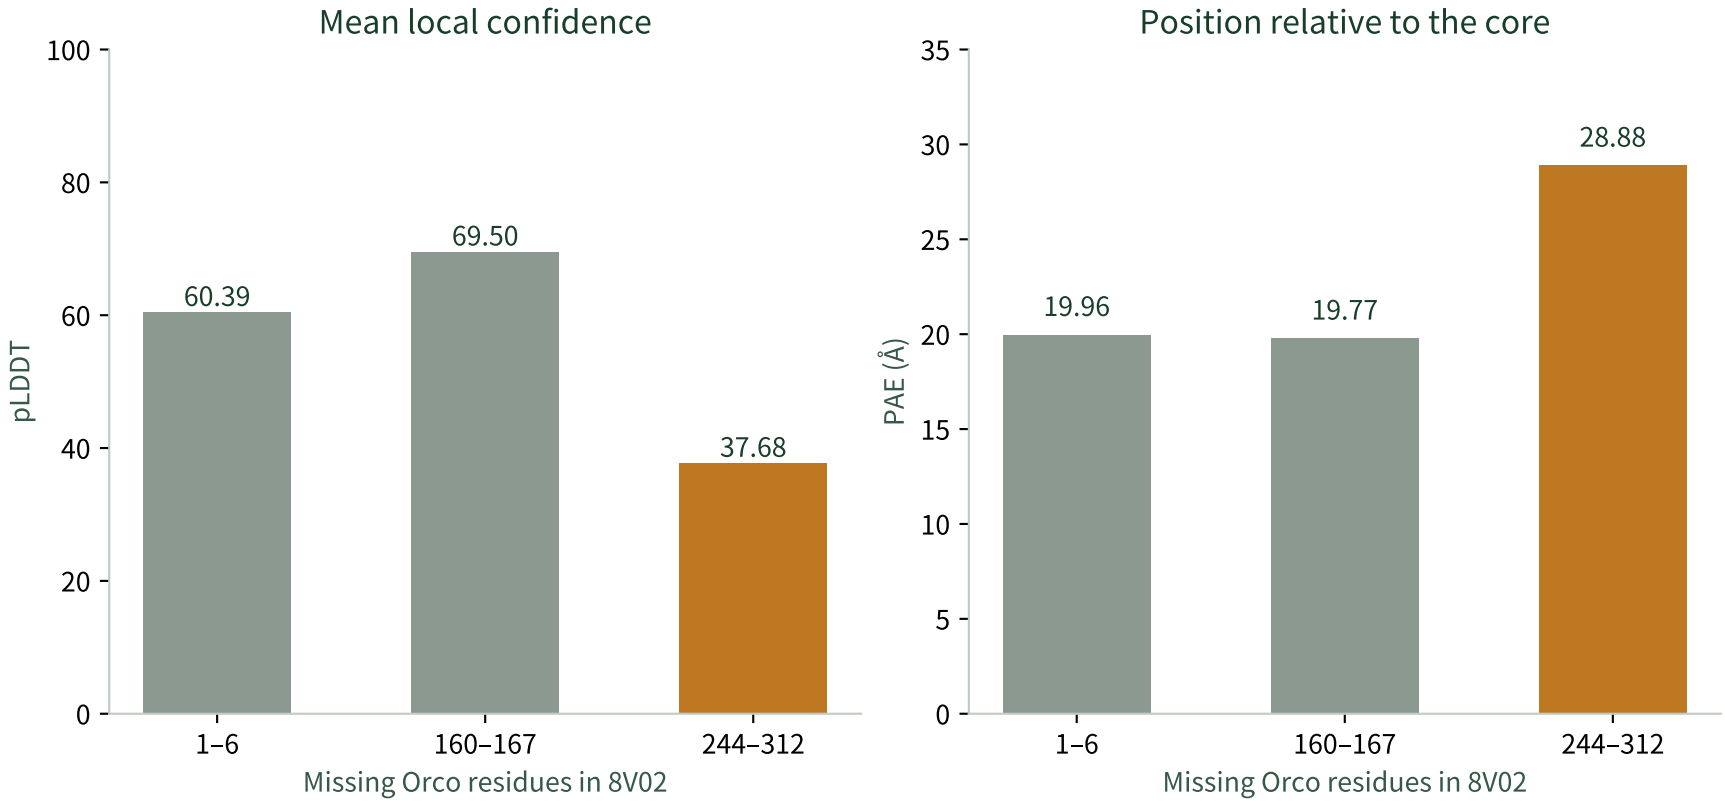

Enlarge ↗

Figure S4. AlphaFold (AI system that predicts protein 3D structures from sequence) confidence summaries for the three missing Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) segments in 8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex). Higher pLDDT (predicted local structure confidence on a 0–100 scale; higher is more confident) indicates greater local confidence; larger loop-to-core PAE (predicted aligned error, an estimate of uncertainty in relative residue placement; lower is more confident) indicates less certain relative placement. The long 244–312 loop has particularly uncertain placement. Confidence is inherited from the template prediction and does not validate the subsequently rebuilt, relaxed coordinates. Structure / method references: ([6]; [7]).

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

Membrane construction, rebuilt regions and geometry checks

Experimental coordinates, AlphaFold and missing regions

8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex) is the o-cresol (2-methylphenol, the odorant resolved in 8V02)-bound Aedes aegypti OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) complex with three Apocrypta bakeri Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) subunits; 8V00 (PDB ID for the experimentally determined odorant-free OR10–Orco complex) is its unbound structural reference. AlphaFold (AI system that predicts protein 3D structures from sequence) supports the missing-region hypotheses and fly-receptor features, while ColabFold (protein-structure prediction workflow built around AlphaFold and fast sequence searches) supplies the mutation models. The resolved OR10/Orco coordinates used for MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) originate from cryo-EM (cryo-electron microscopy: reconstructing molecular structures from electron images of frozen samples), not a de novo AlphaFold complex prediction. ([6]; [7]; [8])

Table S10. Unresolved Orco sequence segments in the experimental 8V02 and 8V00 references, with the total residues rebuilt across each complex. Rebuilt coordinates are starting hypotheses rather than experimentally resolved loop conformations.

Reference	Unresolved Orco segments per subunit	Rebuilt residues across the complex
8V02	1–6, 160–167, 244–312	249
8V00	1–3, 160–167, 244–312	240

A complete sequence is not a validated loop conformation

The complete complex has 1,797 protein residues. The Orco 244–312 loop has mean pLDDT (predicted local structure confidence on a 0–100 scale; higher is more confident) ([7]) 37.68 and mean loop-to-core PAE (predicted aligned error, an estimate of uncertainty in relative residue placement; lower is more confident) 28.88 Å (ångström; 1 Å = 0.1 nm) in the AlphaFold reference. Core alignment and local rebuilding produce testable hypotheses, not observed loop positions. Geometrically rejected candidate structures were excluded from MD, and accepted apo (the odorant-free state) hypotheses retain common modeling ancestry.

Physical model and engineering corrections

OpenMM (the software engine that integrates the atomic equations of motion) integrates classical atomic motion in a POPC (a phosphatidylcholine lipid used to build the model membrane) membrane with TIP3P (a three-site water model used in molecular simulations) ([31]) water and ions, using ff14SB (the Amber parameter set for protein atomic interactions) ([32])/Lipid17 (the Amber parameter set for lipid atomic interactions) ([33]) and GAFF 2.2.20 (General Amber Force Field; here version 2.2.20 supplies small-molecule parameters) ([34]) with AM1-BCC (a method for assigning approximate partial charges to ligand atoms) ([35]) ligand (a molecule that can bind to a receptor; here, the tested odorant) charges. Preparation uses pH 7 protonation, not constant-pH dynamics. The baseline is 300 K and 1 bar with PME (particle-mesh Ewald, a method for long-range electrostatics in periodic boxes) and a 1 nm (nanometers; 1 nm = 10^{-9} meters) short-range cutoff; bonds to hydrogen are constrained, with no hydrogen-mass repartitioning. Steering studies use the documented 300 K NVT (constant particle number, volume and temperature) setup. ([14])

A restrained 200 ps (picoseconds, 10^{-12} seconds) pilot checked system preparation. Full-sequence 5 ns (nanoseconds, 10^{-9} seconds) diagnostics exposed a minimum protein periodic-image separation of 0.877 nm, below the short-range cutoff. Increasing minimum water padding to 2.5 nm produced a 658,984-atom system with 1,066 POPC molecules and 162,148 waters. In the corrected 5 ns runs, checked minimum image separations were 3.227–4.113 nm; this fixes the observed short-range issue without excluding long-range finite-size effects.

The first 1 ns includes 20 ps heating and 980 ps staged restraint release; the remaining 4 ns are unrestrained. Integration uses 1 fs (femtoseconds, 10^{-15} seconds) during heating and 2 fs thereafter, distinct from saved-frame spacing. Frame-level checks cover 1,944 protein stereocentres, heavy-atom bonds and periodic images. Geometry passing is necessary for interpretation but is not evidence of equilibrium. ([36])

Research evidence

Bound-state MD: from 5 ns to 20 ns

Extending the bound-state simulations to 20 ns preserved checked geometry but revealed continued drift in rebuilt regions, so equilibrium was not established.

Full bound-state trajectories and stability diagnostics

What remained stable, and what kept changing?

Matched bound-state trajectories were extended from 5 to 20 ns (nanoseconds, 10^{-9} seconds) without positional restraints during the extension. The prespecified summary windows are 3–5 ns and 15–20 ns. All completed trajectories passed the saved-frame geometry checks without detected chiral inversion; one rebuilt-region RMSD (root-mean-square deviation: positional difference after the stated structural alignment) changed by 0.2591 nm (nanometers; 1 nm = 10^{-9} meters) between the 10–15 and 15–20 ns windows, exceeding the provisional 0.1 nm drift threshold.

Each summary first averages within a trajectory. Variation across those averages describes sampling conditional on the same starting model; correlated frames are not independent experiments. Rebuilt-region RMSD measures relaxation from a modeled starting hypothesis, not error against a known native loop.

Table S11. Bound-state MD readouts in the prespecified 3–5 ns and 15–20 ns windows, summarized as mean ± SD across trajectory-level averages. Frames are correlated; continued rebuilt-region drift prevents these diagnostics from establishing equilibrium.

Readout	3–5 ns · mean ± SD (standard deviation, describing spread among observations)	15–20 ns · mean ± SD
Temperature (K)	300.2847 ± 0.0416	300.2873 ± 0.0363
Box volume (nm^3)	6464.2306 ± 0.5542	6463.5634 ± 0.5781
Lateral area (nm^2)	384.2783 ± 0.5193	378.9845 ± 1.8196

Readout	3–5 ns · mean ± SD (standard deviation, describing spread among observations)	15–20 ns · mean ± SD
Membrane phosphate thickness (nm)	3.8239 ± 0.0094	3.8575 ± 0.0128
Resolved-core Cα (alpha-carbon atoms, used to trace the protein backbone) RMSD (nm)	0.2189 ± 0.0051	0.2500 ± 0.0212
OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6]) Cα RMSD (nm)	0.1269 ± 0.0149	0.1432 ± 0.0180
Rebuilt-region Cα RMSD (nm)	1.0383 ± 0.0445	1.5917 ± 0.1706
Ligand (a molecule that can bind to a receptor; here, the tested odorant) RMSD after OR10 fit (nm)	0.0758 ± 0.0039	0.1142 ± 0.0434
Ser133 (serine at residue 133) H-bond geometry fraction	0.9567 ± 0.0252	0.7613 ± 0.1792

holo-window-statistics.csv ↓ — Zenodo (publication pending)

Ligand contacts are structural readouts

The Ser133 hydrogen-bond geometry requires O···O <0.35 nm and donor–hydrogen–acceptor angle >135°, checking both donor directions. Its mean fraction changes from 0.9567 in the early comparison window to 0.7613 in the later window. These geometric fractions do not measure affinity, residence time or activation. Together with continued loop drift, the results support local structural diagnostics rather than a claim of equilibrium.

Mutation and molecular-pathway methods

Full mutation design, docking checks and structural interpretation

Do the leading mutation hypotheses have a stable score advantage and a plausible structural explanation?

Evidence · What happened

F34Y (residue 34: phenylalanine → tyrosine) and E196Q (residue 196: glutamate → glutamine) lacked direct mutated-side-chain contacts in the inspected poses. Their relative preference changed with docking (computationally placing and scoring ligand poses in a receptor) ([10]) setup; no robust positive winner emerged.

Learning · What it means

A favorable score can reflect pose placement or search choices without supporting the proposed mutation mechanism.

Change · What we do next

Improve pocket (a local cavity in the receptor that can accommodate the ligand) and conformation assessment before expanding mutation claims; study access in an experimentally resolved reference complex.

Explore Design → Build → Test → Learn → Redesign

Design question

Require both an interpretable ligand (a molecule that can bind to a receptor; here, the tested odorant)–residue relationship and a preference that is stable across reasonable docking (computationally placing and scoring ligand poses in a receptor) ([10]) choices.

Build the model

Align mutant and wild-type structures, compare a common search region with an alternative template-informed center, and inspect side-chain contacts. Define preference as the docking change for controls minus the change for mold-related odorants.

Test & results

Using a 4.5 Å (ångström; 1 Å = 0.1 nm) contact criterion, neither F34Y (residue 34: phenylalanine → tyrosine) nor E196Q (residue 196: glutamate → glutamine) supported a direct-contact explanation in the inspected poses. Candidate preference directions were sensitive to docking setup and odor-panel composition. No mutation maintained a robust positive preference across these checks.

Learn

The first screen generated hypotheses, not validated engineered receptors. Lack of direct contact does not rule out indirect effects, but those effects were not demonstrated by these models.

Redesign

Separate the functional-prediction question from the physical-access question. Use the OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6])/Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) complex to test what docking alone leaves unresolved.

Research evidence

Mutation design and the odor panel

Priority receptors and improved mutants are separate decisions

Our comparison of published responses found higher mean responses to the selected mold-related panel than to the control panel for unmodified Or22a (one of the named Drosophila melanogaster odorant receptors in our design panel) and Or35a. This supports prioritizing these receptors for functional testing; their selection does not depend on a proposed mutation succeeding. The selected-panel comparison does not establish mold specificity or a detection limit. ([1])

Mutation screening asks a further question: does changing an amino acid improve preference relative to the original receptor? F34Y (residue 34: phenylalanine → tyrosine) and E196Q (residue 196: glutamate → glutamine) showed small favorable shifts in the initial screen, but subsequent checks found that the direction and ranking depended on the calculation setup. Thus receptor candidates have been identified, while the proposed mutation benefits remain uncertain even before wet-lab validation.

Computational constructs

We proposed 35 single-residue substitutions across Or35a, Or22a, Or85b, Or98a and Or67a and checked the original amino acid at every position. The first structural screen used 12 substitutions plus Or22a/Or35a wild types, giving 14 ColabFold (protein-structure prediction workflow built around AlphaFold and fast sequence searches) ([8]) models. These are sequence and structure designs; they do not represent completed cloning or functional expression.

Mean pLDDT (predicted local structure confidence on a 0–100 scale; higher is more confident) ([7]) was about 91.6–92.2 for Or22a and 84.0–86.0 for Or35a. Local Cα (alpha-carbon atoms, used to trace the protein backbone) RMSD (root-mean-square deviation: positional difference after the stated structural alignment) within five residues on each side of the substitution was about 0.06–0.25 Å (ångström; 1 Å = 0.1 nm). Similar predicted backbones do not establish protein stability or receptor function.

Table S12. Or22a and Or35a substitutions included in the first computational structural screen. These are modeled constructs; similar predicted structures do not establish expression, stability or receptor function.

Receptor	Substitutions in the first structural screen
Or22a	E196Q, E196A (residue 196: glutamate → alanine), E196K (residue 196: glutamate → lysine), S33A (residue 33: serine → alanine), S33T (residue 33: serine → threonine), F34A (residue 34: phenylalanine → alanine), F34Y
Or35a	I87A (residue 87: isoleucine → alanine), I87V (residue 87: isoleucine → valine), I87F (residue 87: isoleucine → phenylalanine), W305F (residue 305: tryptophan → phenylalanine), W305Y (residue 305: tryptophan → tyrosine)

What “mold-related preference” means here

The 14 structures were docked against 16 odorants. ΔVina (mutant docking score minus the corresponding wild-type score) = score(mutant) − score(WT (wild type: the unmutated reference sequence)); panel-preference shift = mean(ΔVina_control) − mean(ΔVina_mold-related). A positive shift favors this selected panel in the score comparison. It is not a receptor activation, EC50 (concentration producing half the maximal response in a specified functional assay) or detection-limit estimate.

Table S13. The selected eight mold-related and eight control odorants used in the docking panel-preference calculation. Preference refers to this score comparison and panel, not demonstrated mold specificity or receptor activation.

Group	Odorants
Mold-related	1-octen-3-ol; 2-heptanone; 1-octanol; 1-butanol; 3-methylbutanol; 4-ethyl guaiacol; methyl octanoate; ethyl octanoate
Controls	acetophenone; benzaldehyde; ethyl acetate; ethyl butyrate; geranyl acetate; linalool; linalool oxide; methyl acetate

Why the first ranking was revised

Initial shifts were +0.0581 for F34Y and +0.0543 for E196Q. Later common-box means were −0.0234 and +0.0142, with search ranges crossing zero. Both mutations lacked direct side-chain contacts in the inspected original and recalculated poses. For 1-octen-3-ol, the historical closest side-chain distances were about 10.6 and 10.4 Å.

The audit aligned mutant and WT Cα atoms, inspected all saved docking (computationally placing and scoring ligand poses in a receptor) ([10]) modes, and compared a common 26 Å cubic search box with a template-informed alternative center. The 8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex) ([6]) footprint was transferred using US-align (software that aligns 3D molecular structures) ([37]) only as a heterologous spatial hypothesis: some original poses lay about 24 Å from that center. It is not an experimentally mapped Drosophila pocket (a local cavity in the receptor that can accommodate the ligand).

Leave-one-odorant-out and within-panel resampling also left F34Y/E196Q uncertainty spanning zero. Input preparation retained aromatic-bond perception warnings, and historical and revised docking settings were analyzed separately. F34Y loses priority as a selectivity candidate; E196Q remains a low-confidence comparison; I87V is useful as a contact-positive case with a negative relative preference. No first-round mutation is a validated positive winner.

Full ensemble docking, functional panels and cycle record

Does an ensemble of receptor conformations explain odor responses better than a single static pocket (a local cavity in the receptor that can accommodate the ligand)?

Evidence · What happened

MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) supplied alternative pockets, but their docking (computationally placing and scoring ligand poses in a receptor) ([10])–response correlations did not exceed the static apo (the odorant-free state) reference in the expanded odor panel.

Learning · What it means

More conformations do not automatically yield a better functional predictor. Docking also assumes the ligand (a molecule that can bind to a receptor; here, the tested odorant) has already reached the pocket.

Change · What we do next

Investigate how the odorant reaches and leaves the pocket, alongside how it fits once placed there.

Explore Design → Build → Test → Learn → Redesign

Design question

Use experimental unbound 8V00 (PDB ID for the experimentally determined odorant-free OR10–Orco complex) and o-cresol (2-methylphenol, the odorant resolved in 8V02)-bound 8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex) structures as OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito)/Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) references. Reconstructed unresolved loops are model assumptions, not experimentally observed coordinates. ([6])

Build the model

Prepare the receptor complex in a hydrated POPC (a phosphatidylcholine lipid used to build the model membrane) membrane with ions. Correct geometric and periodic-box problems before extending bound-state diagnostics to 20 ns (nanoseconds, 10^{-9} seconds) and sampling unbound pockets (a local cavity in the receptor that can accommodate the ligand) with 5 ns MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]). Cluster the pockets and dock odorants into representative conformations.

Test & results

In an expanded panel of 29 dockable odorants, Spearman (rank correlation, from −1 to +1; it measures ordering, not prediction accuracy as a percentage) correlations were 0.407 and 0.338 for two MD-derived model ensembles, compared with 0.450 for static 8V00 and 0.504 for static 8V02. Responses were digitized from a published calcium assay (a defined experimental test of biological activity). An exploratory aromatic-ring-count baseline reached 0.790, highlighting chemical-panel effects. ([15])

Learn

These retrospective correlations do not establish sensitivity, EC50 (concentration producing half the maximal response in a specified functional assay) or affinity. Different functional assays were kept separate, and short MD does not establish equilibrium or accuracy of the rebuilt regions.

Redesign

Test molecular access explicitly, beginning with unbiased trajectories of o-cresol placed outside the apo (the odorant-free state) receptor.

Experimental structure

Where is the odorant in the receptor complex?

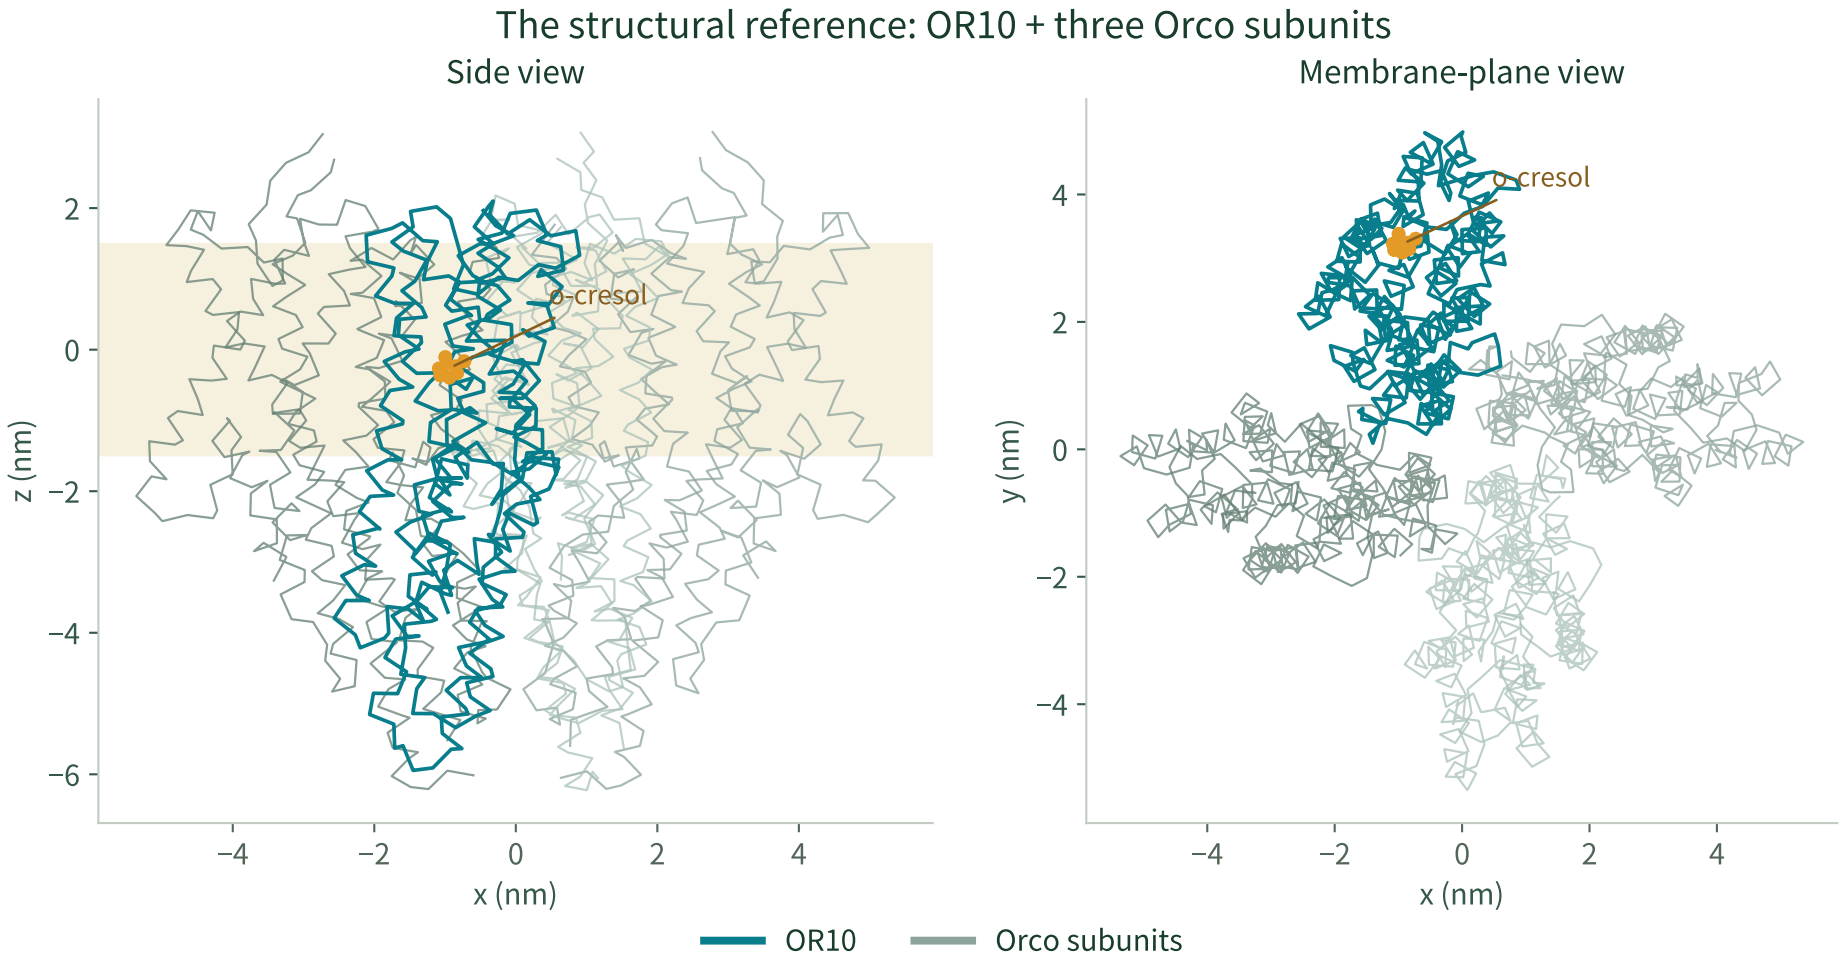

Enlarge ↗

Figure S5. Cα (alpha-carbon atoms, used to trace the protein backbone) traces and o-cresol (2-methylphenol, the odorant resolved in 8V02) heavy atoms from experimental PDB (Protein Data Bank, the archive of molecular 3D structures) ([30]) 8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex). Teal is OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito), gray-green the three Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) subunits, and orange the odorant. Missing resolved-chain segments remain disconnected. The tan band is a homology-transferred membrane-orientation reference, not observed lipids. This figure shows the structural starting reference, not an MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) trajectory. Experimental structure source: ([6]).

Vector SVG ↓

Our computational results

See the odorants behind the correlation

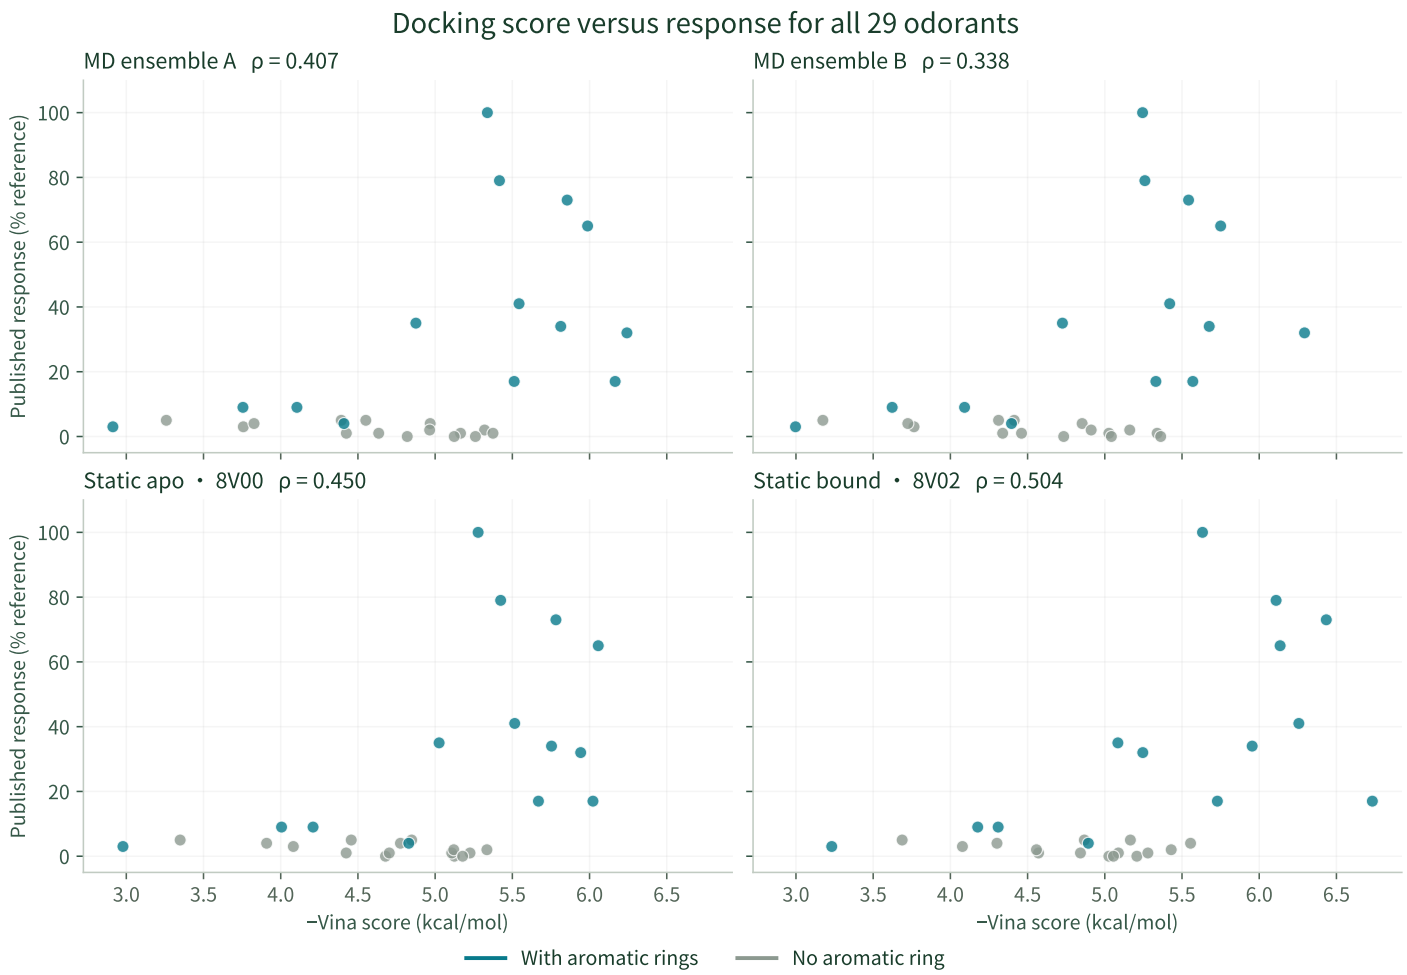

Enlarge ↗

Figure S6. Each panel contains the same 29 odorants. Our docking (computationally placing and scoring ligand poses in a receptor) ([10]) scores are compared with AaOR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6]) calcium responses digitized from Bohbot et al. (2011), Figure 2; these are literature measurements, not our wet-lab results. The horizontal axis reverses the Vina (AutoDock Vina docking score; lower scores are more favorable in this model, not a measured response) ([20]) score so stronger predicted docking lies to the right. Color identifies aromatic chemistry. The scatter exposes substantial response variation at similar scores; ensemble docking has not established a calibrated sensitivity or EC50 (concentration producing half the maximal response in a specified functional assay) predictor. Published data source: ([15]).

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

Download plotted values (CSV) ↓ — Zenodo (publication pending)

Research evidence

Keep the two functional assays separate

The initial seven-odorant reference

The initial functional dataset contains seven odorants and 63 published oocyte-current observations with AaegOR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6])/AaegOrco (the odorant-receptor co-receptor from Aedes aegypti). Screening concentration remains unresolved in the inspected records, and dose-response tables have normalization/labeling questions. The structures instead contain AbOrco (the odorant-receptor co-receptor from the fig wasp Apocrypta bakeri). These differences prevent treating the comparison as a calibrated sensitivity or EC50 (concentration producing half the maximal response in a specified functional assay) model. ([38])

Static 8V00 (PDB ID for the experimentally determined odorant-free OR10–Orco complex) ([6]) and ligand (a molecule that can bind to a receptor; here, the tested odorant)-removed 8V02 (PDB ID for the experimentally determined o-cresol-bound OR10–Orco complex) docking (computationally placing and scoring ligand poses in a receptor) ([10]) gave rank correlations of 0.072 and 0.595. The 8V02 o-cresol (2-methylphenol, the odorant resolved in 8V02) pose control reproduced the native pose with symmetry-corrected heavy-atom RMSD (root-mean-square deviation: positional difference after the stated structural alignment) of 1.17–1.22 Å (ångström; 1 Å = 0.1 nm); o-cresol had no functional label in this seven-odorant set and was excluded from the response correlation. After apo (the odorant-free state) sampling and pocket (a local cavity in the receptor that can accommodate the ligand) clustering, the two MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) ensembles gave 0.432 and 0.072. These small retrospective comparisons do not establish generalization.

The first functional panel and static docking controls

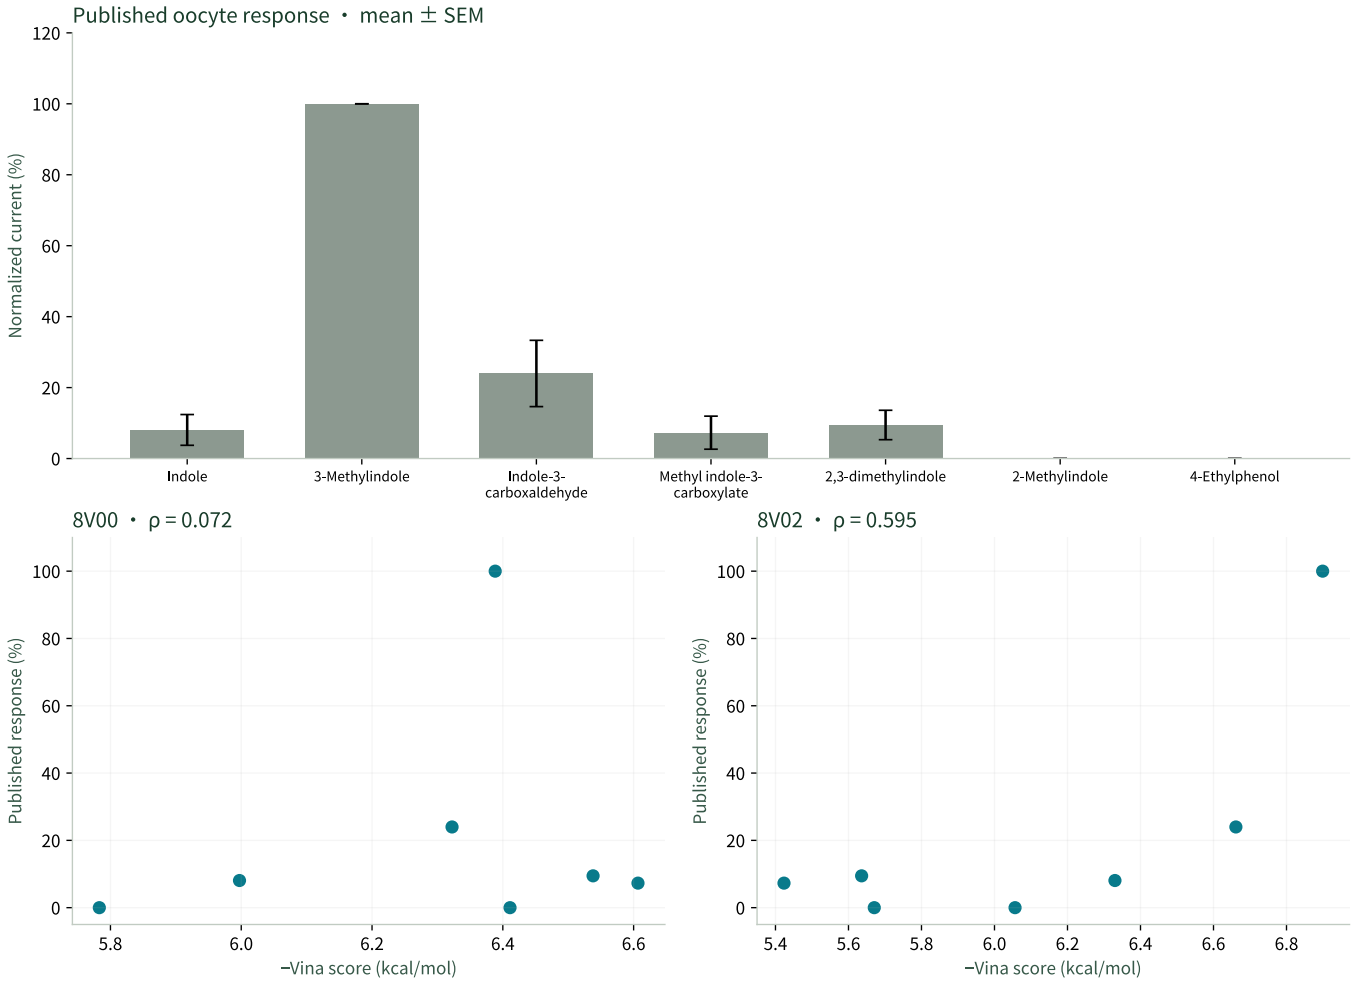

Enlarge ↗

Figure S7. Top: published AaegOR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6])/AaegOrco (the odorant-receptor co-receptor from Aedes aegypti) oocyte currents normalized to 3-methylindole, with mean ± SEM (standard error of the mean, describing uncertainty in an estimated mean) from nine oocytes per odorant. A zero SEM for the normalization reference is not evidence of noise-free measurement. Bottom: this project's static docking (computationally placing and scoring ligand poses in a receptor) ([10]) scores versus the same seven response means. The assay (a defined experimental test of biological activity) uses a different Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) from the structural complex, and screening concentration remains unresolved. This panel is analyzed separately from the 29-odorant calcium assay. Published data source: ([38]).

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)Supporting data CSV ↓ — Zenodo (publication pending)

Why the expanded panel answers a different question

Bohbot et al. provide a separate 30-odorant screen using HEK293 calcium signaling with AgOrco. We matched identities to Supplemental Table I and digitized AaOR10 mean responses in Figure 2; these are not original replicate measurements. Excluding ammonium hydroxide leaves 29 dockable odorants. They do not overlap the initial seven labeled odorants, yielding 36 distinct dockable odorants across two assays (a defined experimental test of biological activity), not one pooled response scale. ([15])

Four organic acids were evaluated in charged and neutral forms, and two stereochemically unspecified odorants were expanded into alternative isomers. The primary analysis uses rounded digitized means; low bars are not definitive inactive labels. MD pocket representatives are actual 3–5 ns (nanoseconds, 10^{-9} seconds) medoid (an actual sampled conformation selected to represent a cluster) frames, and equal weighting is descriptive rather than thermodynamic.

For the 29-odorant panel, MD A/B correlations of 0.407/0.338 were below the static apo value of 0.450. Ligand-only controls gave 0.081 for heavy-atom count and −0.012 for neutral cLogP (calculated octanol–water partition coefficient on a logarithmic scale, a lipophilicity descriptor). The later aromatic-ring-count comparison reached 0.790. Chemical composition and post-selection optimism therefore remain plausible explanations; the exploratory intervals omit experimental error and chemical dependence.

Full entry simulations, placement methods and trajectory criteria

Is failure to observe entry simply caused by waiting for the odorant to encounter the receptor?

Evidence · What happened

Odorants contacted the membrane or receptor, but the analyzed trajectories did not meet the sustained pocket (a local cavity in the receptor that can accommodate the ligand)-entry criterion, including near-entrance simulations lasting 10 ns (nanoseconds, 10^{-9} seconds).

Learning · What it means

Bringing a ligand (a molecule that can bind to a receptor; here, the tested odorant) closer or increasing its local availability did not resolve pocket entry within the sampled time.

Change · What we do next

Start at the known bound position and explore candidate exits to identify where access is obstructed.

Explore Design → Build → Test → Learn → Redesign

Design question

Distinguish reaching the receptor surface from actually occupying its binding pocket (a local cavity in the receptor that can accommodate the ligand).

Build the model

Place o-cresol (2-methylphenol, the odorant resolved in 8V02) ([6]) in external water, at the water–membrane interface, and near the candidate S3/S6 (labels for the third and sixth membrane-spanning helices of OR10) entrance. Compare low- and high-occupancy solvent preparations to explore encounter frequency, then reduce the initial diffusion distance directly.

Test & results

Track pocket distance, receptor contacts and sustained occupancy. The available trajectories showed surface and membrane encounters but no sustained entry. The concentration exploration is preliminary and does not establish a concentration–response relationship.

Learn

A nearby starting position does not guarantee a passable entrance. A missing event over nanoseconds cannot prove that natural entry is impossible.

Redesign

Use the experimentally observed bound pose to explore outward motion and test receptor-side opening motions separately.

Full steering methods, saved structures and release trajectories

Can outward steering or a more apo (the odorant-free state)-like entrance geometry produce a complete exit?

Evidence · What happened

Initial outward steering encountered resistance. Guiding S3/S6 (labels for the third and sixth membrane-spanning helices of OR10) ([6]) distances and extending the observation to 70 ns (nanoseconds, 10^{-9} seconds) changed entrance geometry but produced no sustained exit.

Learning · What it means

An entrance-distance change is not equivalent to a continuous passage for a molecule of finite size.

Change · What we do next

Search the saved structures for full-ligand (a molecule that can bind to a receptor; here, the tested odorant) routes and identify local clashes along those routes.

Explore Design → Build → Test → Learn → Redesign

Design question

Separate direct ligand (a molecule that can bind to a receptor; here, the tested odorant) pulling from receptor-side bias, and compare guided entrance motion with an unsteered control under the same observation schedule.

Build the model

Use steered molecular dynamics (SMD (steered molecular dynamics: MD with an applied guiding force) ([16])) for exploratory outward motion. In a separate test, bias S3/S6 (labels for the third and sixth membrane-spanning helices of OR10) ([6]) reference distances toward apo (the odorant-free state) values while leaving o-cresol (2-methylphenol, the odorant resolved in 8V02) free of direct pulling. Extend the receptor-guidance protocol to a total of 70 ns (nanoseconds, 10^{-9} seconds), with the final 1 ns after bias removal.

Test & results

Assess entrance distances and ligand departure separately. Neither the guided nor the unsteered 70 ns trajectories met the sustained-exit criterion. The initial outward SMD attempts did not yield an accepted exit route for reverse-entry testing.

Learn

Longer observation improved the opportunity to sample rearrangements but did not reveal an exit. The biased simulations do not measure a natural residence time, and 70 ns is insufficient to rule out slower events or other pathways.

Redesign

Replace a single opening-distance proxy with explicit position-and-orientation searches for the whole odorant, including nearby lipids.

Method schematic

Two different targets for an applied force

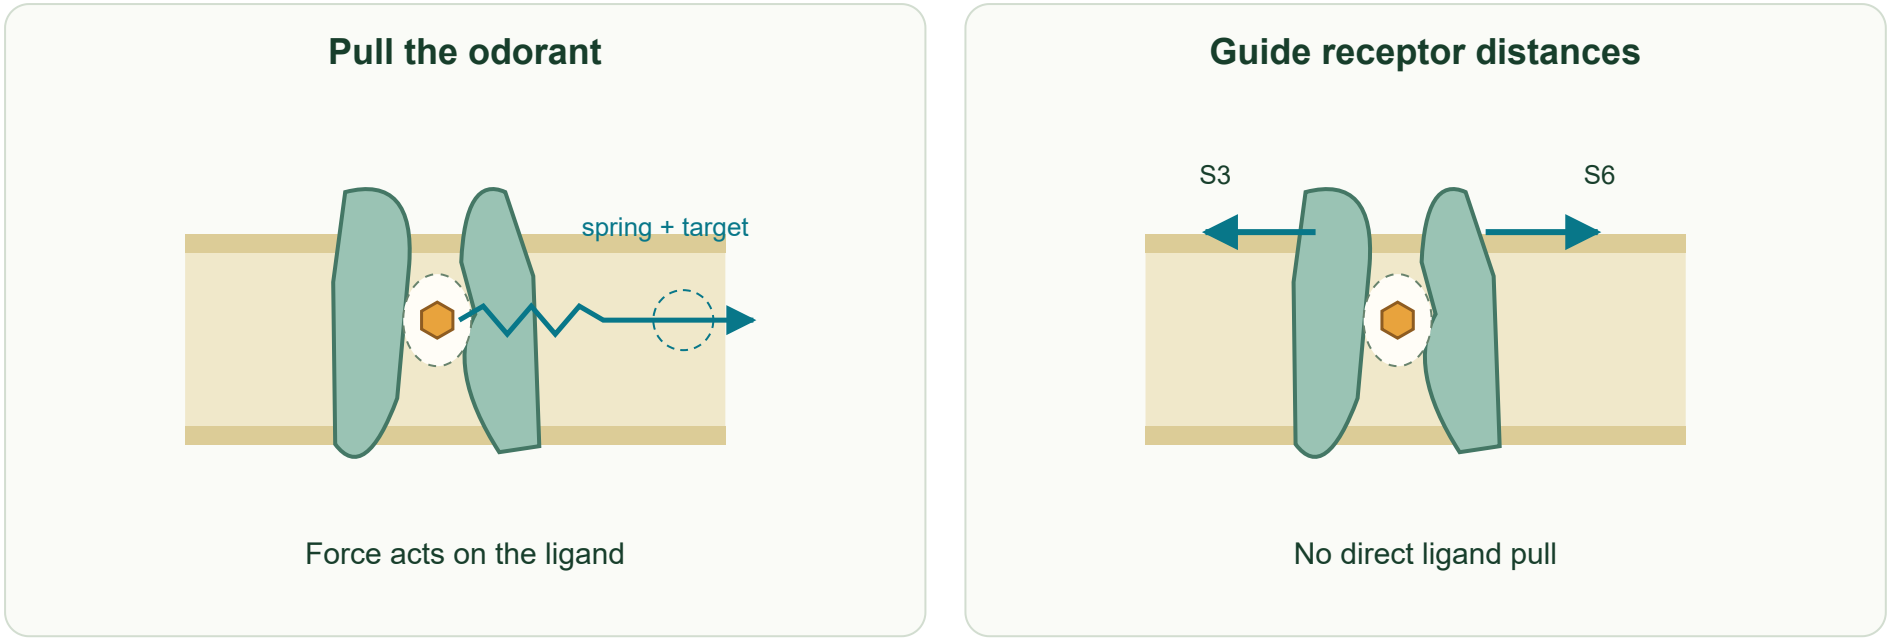

Enlarge ↗

Figure S8. Left: SMD (steered molecular dynamics: MD with an applied guiding force) ([16]) couples the odorant to a moving target using a spring. Right: a separate protocol guides S3/S6 (labels for the third and sixth membrane-spanning helices of OR10) ([6]) reference distances while applying no direct pulling force to the odorant. Blue arrows indicate imposed guidance, not observed trajectories. Both protocols test constrained conditions and do not measure natural entry or exit rates.

Vector SVG ↓

Research evidence

How the steering experiments evolved

Outward exploration before reverse entry

The first outward SMD attempts pulled o-cresol along candidate guidance lines. All encountered excessive target lag before satisfying exit criteria, leaving no accepted outward path for a reverse-entry analysis. [16] [6]

SMD uses a spring between ligand (a molecule that can bind to a receptor; here, the tested odorant) coordinates and a moving target, with transverse guidance around the route. Work associated with target motion and changing the bias is tracked separately. Neither non-equilibrium work nor transient retention measures PMF (potential of mean force: a free-energy profile along a chosen coordinate), binding free energy, sensitivity or natural entry kinetics.

Guide the receptor instead of pulling the ligand

Receptor guidance controls nine distances between three S3 (labels for the third and sixth membrane-spanning helices of OR10) ([6]) Cα (alpha-carbon atoms, used to trace the protein backbone) segments (123–127, 130–134, 136–140) and three S6 segments (289–293, 282–286, 277–281), moving them toward experimental apo (the odorant-free state) reference distances. Each distance uses a spring constant of 3000/9 kJ mol^{-1} nm (nanometers; 1 nm = 10^{-9} meters)$^{-2}$. These collective distances are not tunnel radii. Short guided-opening tests did not meet the prespecified distance and directional-progress requirements for subsequent ligand pulling.

Table S14. Durations of baseline, receptor-distance guidance, target holding and bias removal in the three guided protocols. These schedules control receptor distances rather than pulling the ligand directly.

Stage	Short protocol	Extended protocol	70 ns (nanoseconds, 10^{-9} seconds) protocol
Baseline	0.1 ns	0.1 ns	0.1 ns
Change the target	1 ns	5 ns	5 ns
Hold the target	0.5 ns	2.5 ns	63.9 ns
After bias removal	1 ns	1 ns	1 ns
Total trajectory time	2.6 ns	8.6 ns	70 ns

Compare opening geometry with matched controls

Slower target changes and longer holding reduced entrance-distance error more for starting conformation A than B in the short-to-extended comparison. The paired duration contrasts were −0.0663 and −0.0187 nm, respectively. Timing changes affected both ramp speed and hold duration, so those effects cannot be separated. Guided and unsteered 70 ns trajectories still showed no sustained exit.

An exit candidate must be at least 1.5 nm from the pocket (a local cavity in the receptor that can accommodate the ligand) reference, have zero designated pocket-residue contacts and persist for at least 100 ps (picoseconds, 10^{-12} seconds) within one stage. Bias status must accompany any event. Reaching an apo-like distance or briefly satisfying a force threshold is not equivalent to molecular passage.

Our computational results

Short versus extended entrance guidance

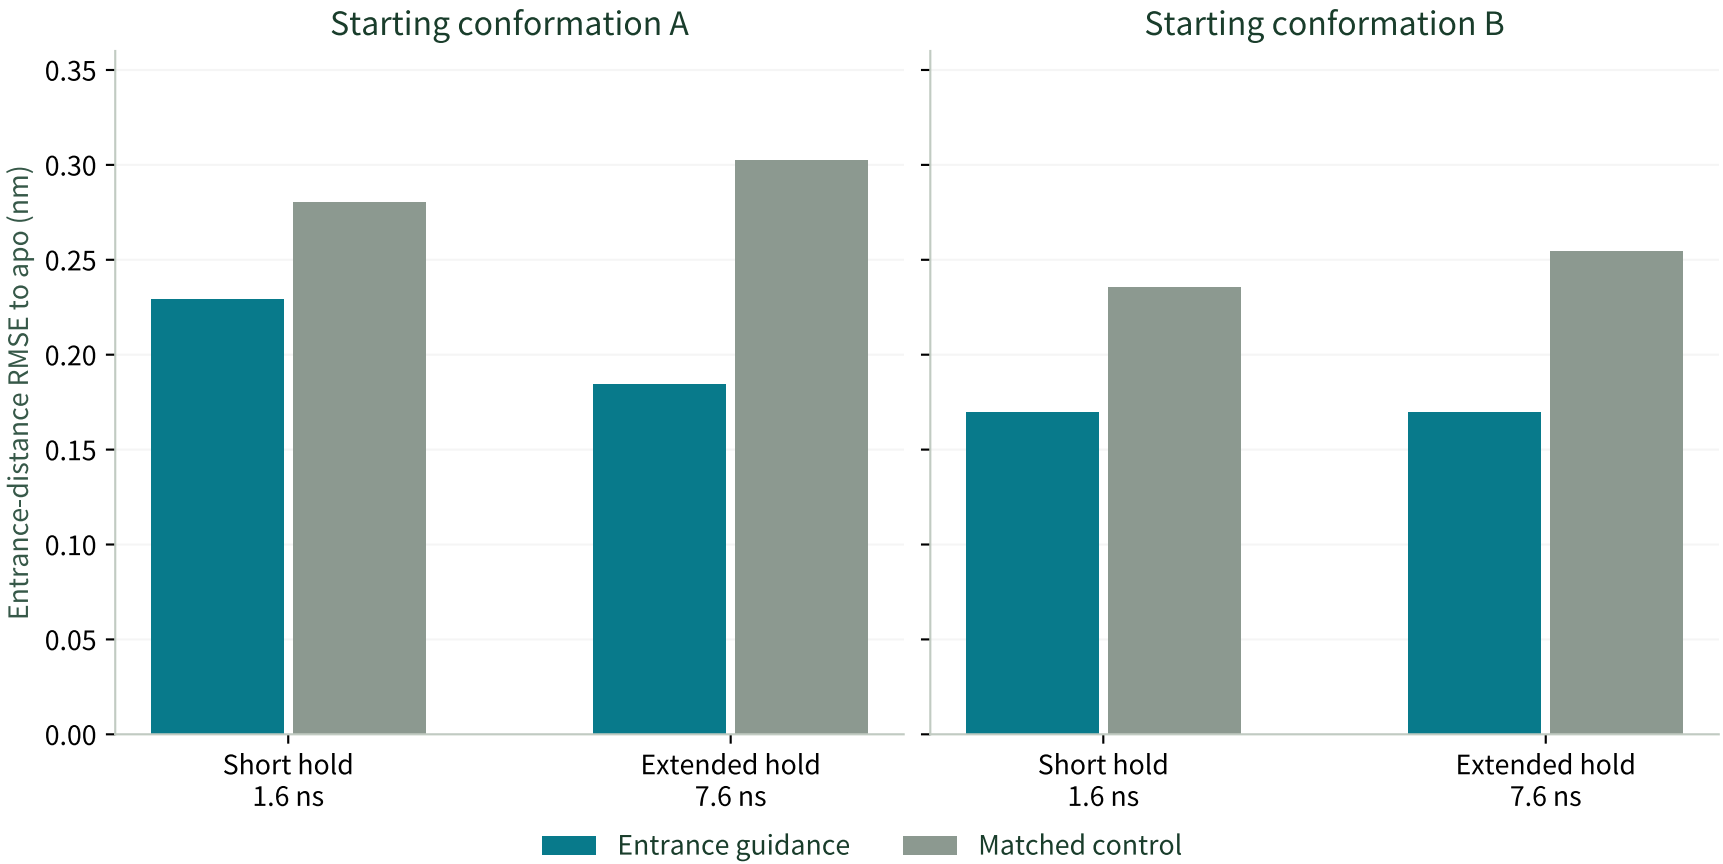

Enlarge ↗

Figure S9. Entrance-distance RMSE (root-mean-square error relative to the selected reference values) over the last 100 ps (picoseconds, 10^{-12} seconds) before bias removal, compared with time-matched unsteered controls. Each starting conformation is shown separately. The extended protocol changes both target-ramp duration and holding time. Lower RMSE means closer to the selected apo (the odorant-free state) distances, not a wider or passable tunnel. The comparison is descriptive and uses correlated trajectories.

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)

Full tunnel-search method and trajectories

Can a complete odorant translate and rotate through the candidate corridors as its surroundings change?

Evidence · What happened

Two candidate directions emerged, but whole-ligand (a molecule that can bind to a receptor; here, the tested odorant) searches across 24 saved conformations found no collision-free complete route. Dynamic route-following and longer local waiting also produced no complete exit.

Learning · What it means

Local side-chain and lipid arrangements matter. Coordinated motion is a plausible explanation to test, but ligand-induced opening has not been demonstrated.

Change · What we do next

Design controlled, small-step movement and relaxation tests that track the ligand and its environment together.

Explore Design → Build → Test → Learn → Redesign

Design question

Distinguish a point-sized probe, a full rigid molecule and a flexible molecular environment. A geometrical centerline alone cannot establish passage.

Build the model

Identify a lateral S3/S6 (labels for the third and sixth membrane-spanning helices of OR10) ([6]) corridor and a lower-membrane-coordinate corridor. Search the translation and rotation of all 16 o-cresol (2-methylphenol, the odorant resolved in 8V02) atoms with the receptor and lipids fixed. Test route-following SMD (steered molecular dynamics: MD with an applied guiding force) ([16]) with mobile surroundings, extend local waiting from 0.5 to 2 ns (nanoseconds, 10^{-9} seconds), then repeat the rigid searches across 24 conformations from the 70 ns trajectories.

Test & results

No tested static conformation provided a complete collision-free route. Flexible pulling sometimes showed transient catch-up to the target, but no sustained release. In the least-repulsive selected protein-plus-lipid route, L128 (leucine at residue 128), Q292 (glutamine at residue 292) and S291 (serine at residue 291) and nearby lipids remained collision sources. The repulsion score diagnoses clashes; it is not a free-energy barrier.

Learn

Two finite corridors were sampled in related conformations. No continuous molecular passage was observed, including in SMD with receptor motion.

Redesign

The resulting comparison design pairs small outward ligand steps and local relaxation with a no-pull control, separating ligand motion from side-chain, lipid and water rearrangement.

Research evidence

Trace a path for the whole molecule

Search candidate geometry in saved structures

A widest-path graph search on saved conformations identified candidate space near S3/S6 (labels for the third and sixth membrane-spanning helices of OR10). The 47.6 ns (nanoseconds, 10^{-9} seconds) snapshot suggested a lateral route near I132 (isoleucine at residue 132)/Y285 (tyrosine at residue 285)/M288 (methionine at residue 288); the 59 ns snapshot suggested a lower-membrane-coordinate route near L128 (leucine at residue 128)/I184 (isoleucine at residue 184)/Q292 (glutamine at residue 292). “Lower” refers to the chosen coordinate system, not a demonstrated intracellular or extracellular exit. The lateral structural channel was already described in the experimental work.

At 0.05 nm (nanometers; 1 nm = 10^{-9} meters) grid spacing, candidate spherical-probe radius lower bounds were 0.089, 0.102 and 0.119 nm at 7.6, 47.6 and 59 ns using protein obstacles. Adding lipids and reselecting from the common candidate pool gave 0.089, 0.077 and 0.086 nm. The selected routes can differ between obstacle models. These are radii of particular candidate centerlines, not ligand (a molecule that can bind to a receptor; here, the tested odorant) dimensions or upper bounds on every possible route.

Grid spacing changes the result: the 59 ns protein candidate changes from 0.091 nm on the coarse grid to 0.119 nm on the fine grid. Alternative direction-cone searches and a larger search extent tested sensitivity. Multiple endpoint directions may share a bottleneck and do not establish separate channels.

Why a point-probe route was insufficient

Placing full o-cresol (2-methylphenol, the odorant resolved in 8V02) ([6]) at candidate positions with multiple orientations exposed atom overlaps beyond the pocket (a local cavity in the receptor that can accommodate the ligand). In one 47.6 ns candidate, the most congested position beyond 0.35 nm from the start retained about 0.16 nm hard-sphere overlap even at its best tested orientation, versus about 0.05 nm at the starting pose. Native poses also overlap under a hard-sphere model, so zero overlap is not by itself a physical pass/fail criterion.

The next static search used the full ligand’s translation and orientation in a 63.2 ns structure, explicitly checking consecutive poses and both protein-only and protein-plus-lipid repulsion. Flexible route-following SMD (steered molecular dynamics: MD with an applied guiding force) ([16]) then tested the candidates at 0.2 and 0.5 nm/ns. Target progress stalled relative to actual ligand movement, and extending fixed-target waiting to 2 ns did not establish sustained exit.

The ensemble analysis then evaluated 24 saved conformations across both starting structures and guided/control conditions, two corridors and both obstacle objectives. No evaluated selected route established a low-clash passage. The least severe combined peak was still about 3.34×10^4 kJ/mol versus about 23 kJ/mol for its native pose. Positive Lennard–Jones (an atomic interaction potential; only its positive repulsive part is summed in our clash scans) repulsion is only a clash diagnostic. Paths from different conformations cannot be spliced into an observed continuous trajectory.

The initial 63.2 ns whole-ligand search

This fixed-structure search expanded the corridor radius from 0.30 to 0.50 nm and extended candidates to 2.0 nm from the initial ligand center. Intermediate poses limited each atom’s inter-sample movement arc to 0.02 nm. Internal ligand torsions were fixed. The following values describe this particular snapshot, separately from the later ensemble analysis.

The native pose scored about 15.8 kJ/mol under the same positive-repulsion diagnostic. Protein bottlenecks involved L128/S291 (serine at residue 291)/Q292 or L128/M288/I132; POPC (a phosphatidylcholine lipid used to build the model membrane) dominated the combined maximum. These scores omit attraction, electrostatics, water and entropy. They are neither free-energy barriers nor the work needed for passage, and the search covers only two finite corridors.

Table S15. Centerline lengths and peak positive-repulsion scores for two whole-ligand corridors in the fixed 63.2 ns snapshot. Protein-only and protein-plus-lipid scores diagnose clashes; they are not free-energy barriers or observed ligand passage.

Corridor	Centerline arc (nm)	Protein peak (kJ/mol)	Protein + lipid peak (kJ/mol)
Lower membrane coordinate	3.27	12,005	555,394
S3/S6 lateral	4.03	29,838	103,961

Flexible pulling: target motion versus ligand motion

Route-following SMD moved the ligand center while allowing rotation and receptor, lipid and water motion; it did not impose the rigid search's full orientation sequence. Target advancement was held when lag reached 0.40 nm or force reached 180 kJ mol^{-1} nm^{-1}. Advancement could resume only below 0.20 nm lag and 100 kJ mol^{-1} nm^{-1} force. Holding the target leaves MD (molecular dynamics: computing atomic motion step by step from a force field) ([14]) and the spring force active.

Across the tested pulling conditions, the ligand's maximum pocket distance was only 0.139–0.190 nm, with at least eight pocket residues still in contact. Targets traversed 0.469–0.947 nm along the curved routes, and peak forces were about 274–304 pN. Target arc length is not ligand displacement or an exit percentage. A 2 ns unsteered control reached 0.126 nm; unequal observation durations make these descriptive extrema, not a matched rate comparison.

Locate the obstruction along a candidate path

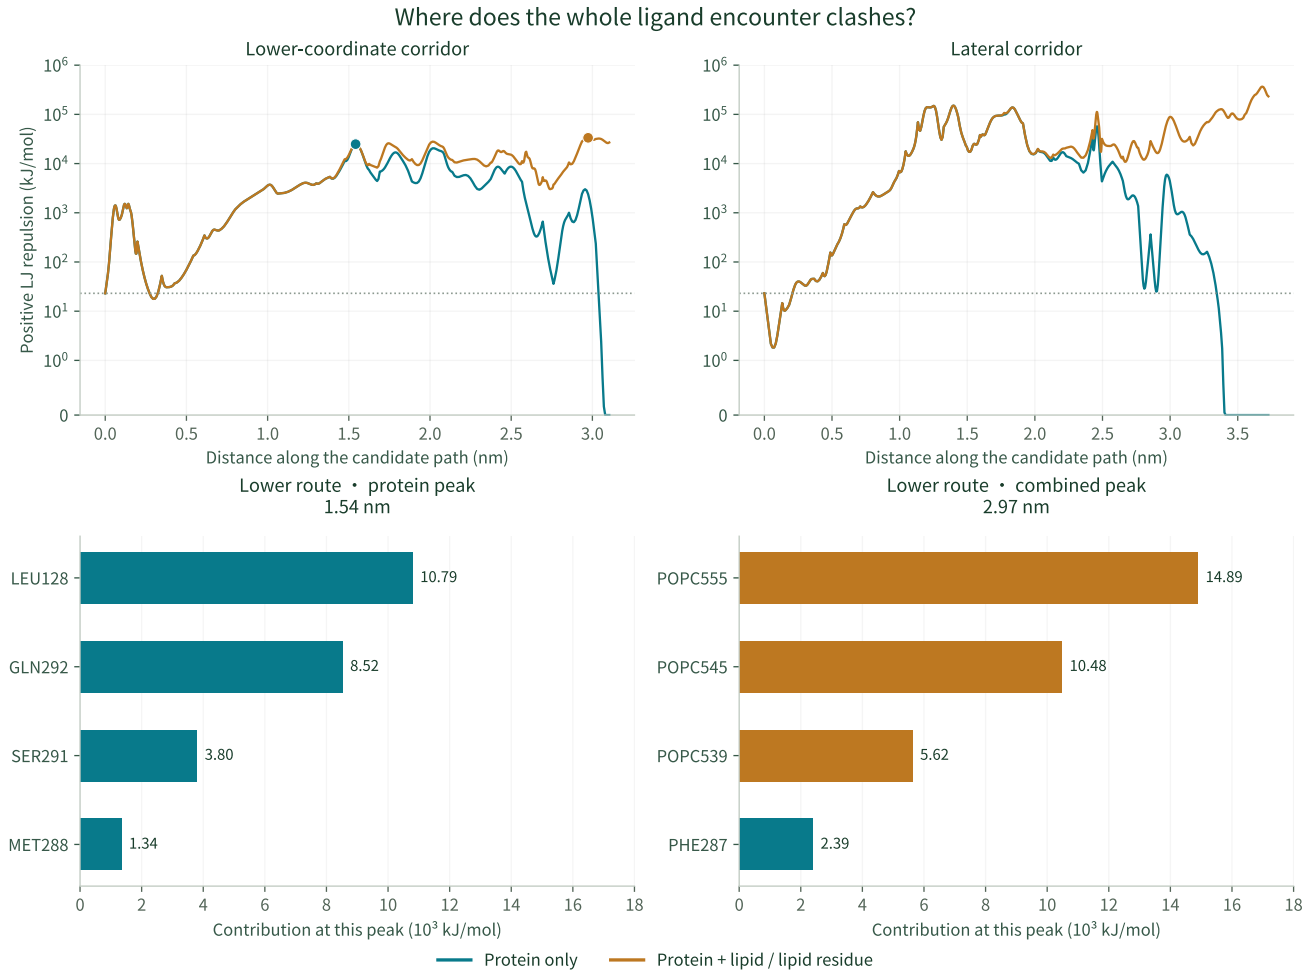

Enlarge ↗

Figure S10. Detailed static placement scan for the 69 ns (nanoseconds, 10^{-9} seconds) entrance-guided A snapshot: the lower corridor has the lowest combined repulsion peak among the evaluated selected paths. Top panels score the same ligand (a molecule that can bind to a receptor; here, the tested odorant) poses with protein alone and with protein plus lipids; the dotted line is the native-pose reference (~23 kJ/mol). Bottom panels identify the four largest contributors at two different peaks of the lower route: protein at 1.54 nm (nanometers; 1 nm = 10^{-9} meters) and combined at 2.97 nm. Protein residues and later lipids both obstruct this candidate. Distances are path arc lengths, not simulation time; positive LJ (an atomic interaction potential; only its positive repulsive part is summed in our clash scans) repulsion is a clash diagnostic, not a free-energy profile.

Vector SVG ↓Data CSV ↓ — Zenodo (publication pending)Contributors CSV ↓ — Zenodo (publication pending)

Molecular availability: one versus twenty odorants

The incomplete one-versus-twenty-odorant comparison found some protein-surface contacts but no sustained pocket entry, without establishing a concentration–response relationship.

Molecular counts, completed runs and concentration assumptions

Controlled placement and concentration meaning

The concentration exploration used the same initial receptor snapshot and box for paired 1- and 20-o-cresol (2-methylphenol, the odorant resolved in 8V02) ([6]) conditions. The first ligand (a molecule that can bind to a receptor; here, the tested odorant) position was shared and the additional molecules were placed in the external water layer with collision exclusions. Initial boxes were about 19 × 20 × 17 nm (nanometers; 1 nm = 10^{-9} meters), corresponding to nominal whole-box concentrations of about 0.26 and 5.2 mM. These are not effective water-phase concentrations or airborne ppm; the periodic box does not maintain separate concentration reservoirs.

The available completed 2 ns (nanoseconds, 10^{-9} seconds) observations form an incomplete, unbalanced exploratory comparison. Greater molecular availability produced some surface contacts but no sustained pocket (a local cavity in the receptor that can accommodate the ligand) entry. The table reports mean per-molecule protein-contact time fractions from completed, checked trajectories only; unfinished observations are excluded. Surface contact includes the protein beyond the OR10 (odorant receptor 10 from Aedes aegypti, the yellow-fever mosquito) ([6]) pocket.

Table S16. Mean per-molecule protein-contact time fractions for completed, checked 2 ns trajectories with one or twenty o-cresol molecules in starting models A and B. This incomplete exploratory comparison measures surface contact, not sustained pocket entry, a concentration–response curve or an association rate.

Starting model	1 odorant	20 odorants
A	0.000	0.036
B	0.000	0.021

What the comparison can establish

Molecules sharing one receptor are not independent experiments. The nearest of twenty molecules will tend to be closer than a single molecule even without stronger binding. Entry requires the same molecule to satisfy distance, residue-contact and duration criteria; contacts from different molecules cannot be joined into one event. These observations do not establish a concentration–response curve or a physical association rate.

The near-entrance study then shortened the initial distance directly and extended unrestrained observation to 10 ns. It remained distinct from the external-water concentration comparison. No sustained entry was seen, motivating outward path exploration from the known bound pose.

Evidence record and comparison design

Full evidence and evaluation record

Results and evidence scope

The work establishes curated data, controlled feature comparisons, mutation audits, a cross-species transfer test and a membrane-MD (molecular dynamics: computing atomic motion step by step from a force field) ([14])/pathway workflow. Sequence features improve within-species prediction. Completed molecular-weight ablation, measured IR (infrared absorption: how a molecule absorbs light at different vibrational frequencies), quantum-mode and OpenPOM (Principal Odor Map; OpenPOM is its open-source reproduction for learned molecular odor features) ([4]; [5]) comparisons show model-dependent effects without a stable general gain from spectra or frozen odor AI. Added structural features, proposed selectivity mutations and direct cross-species transfer remain unvalidated. Reference MD constrains candidate access routes without establishing ligand (a molecule that can bind to a receptor; here, the tested odorant)-induced opening or complete passage.

Table S17. Inference boundaries for prediction, mutation, starting structures and molecular pathways, paired with the comparisons defined to address them.

Defined comparison	Question it can answer
Fold-contained AI adaptation, 3D odorant descriptors and independent data	Can richer features improve responses beyond the retained baseline?
Alternative loop hypotheses and local drift analysis	How dependent are the conclusions on uncertain starting regions?
Target Drosophila WT (wild type: the unmutated reference sequence)/mutant OR–Orco (odorant-receptor co-receptor, a shared partner that forms the channel with the odor-selective OR) ([6]) complexes	Do candidate-specific dynamic differences reproduce?
Independent functional data and nested species-transfer evaluation	Does performance persist on genuinely unseen receptors and assays (a defined experimental test of biological activity)?
Matched expression and functional dose-response experiments	Do designed receptors change activity, selectivity or sensitivity?
Continuous flexible pathways and unbiased follow-up	Is a guided route physically sustainable after removing the force?

Data and reproducibility

Evidence tables retain units and analysis windows; the provenance manifest records source hashes. Model comparisons use stored out-of-fold predictions, while docking and pathway analyses retain pose-level outputs. Saved trajectory selections vary by study; water and lipid-tail coordinates were not continuously recorded in every study. [10]

provenance.json ↓ — Zenodo (publication pending)research-coverage.json ↓ — Zenodo (publication pending)

Comparison design

Coordinated-motion comparison design

The pathway results informed the comparison schematic below: ligand guidance and local relaxation are paired with a matched no-pull control.

- The feature record preserves the fixed baseline and cites related task-adaptation methods separately. [27]
- Matched conditions: small-step ligand guidance and no-pull control, using the same local structural measurements.
- Specified route criteria: continuous molecular positions, clash diagnostics and persistence after force removal.
- Functional endpoint: matched receptor dose–response measurements.

The deliverable is the comparison design; no experiment is represented by this schematic.

Comparison design

Coordinated motion with a matched control

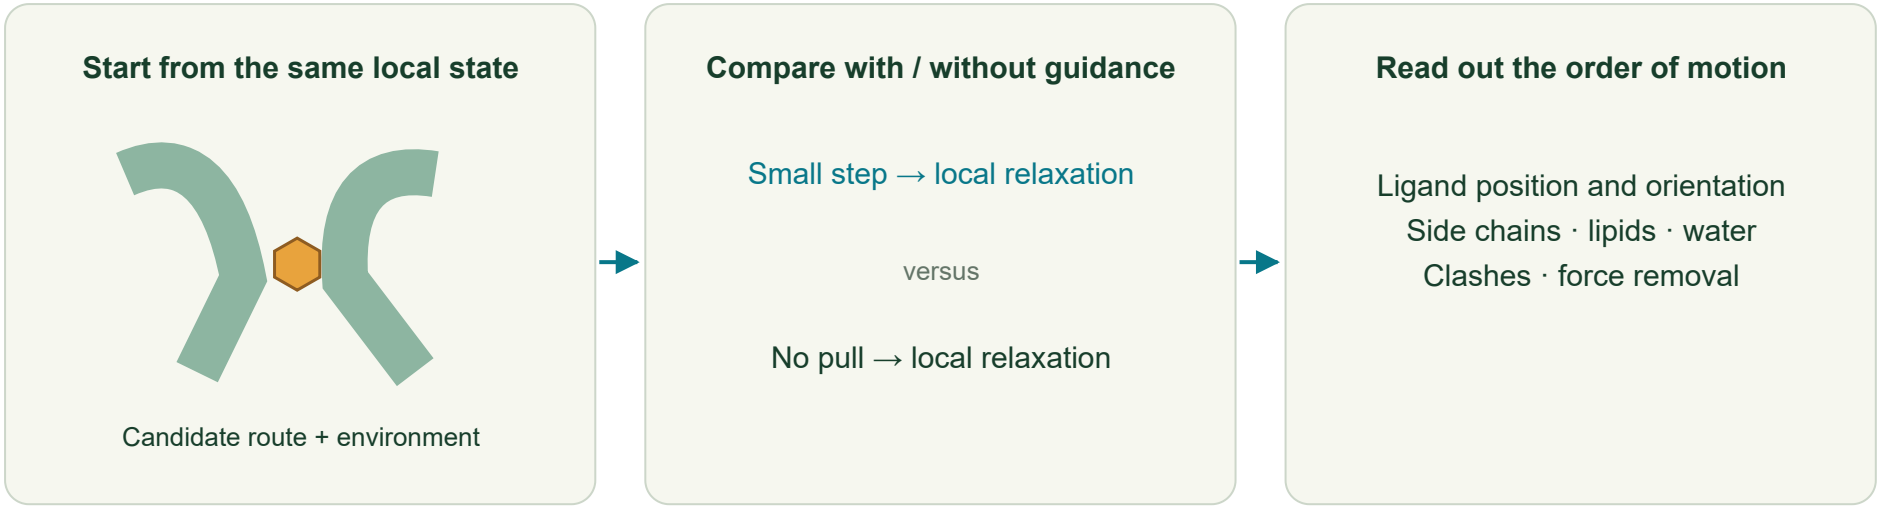

Enlarge ↗

Figure S11. The comparison design pairs small guided ligand movements and local relaxation with a matched no-pull condition. It records side-chain, lipid and water rearrangements relative to ligand motion. This schematic specifies the comparison; it contains no experimental result.

Vector SVG ↓

Receptor engineering sources

Publications, datasets and software cited in the receptor-engineering analysis.

1. Hallem, E. A., & Carlson, J. R. (2006). Coding of Odors by a Receptor Repertoire. *Cell*, 125(1), 143–160. Source ↗
   Back to citation

2. National Institute of Standards and Technology. (n.d.). NIST Chemistry WebBook, SRD 69: Infrared spectra and data guide [Database]. Retrieved October 11, 2026.
   Source ↗
   Back to citation

3. Harada, Y., Maeda, S., Shen, J., Misonou, T., Hori, H., & Nakamura, S. (2025). Motional representation; the ability to predict odor characters using molecular vibrations [Preprint]. arXiv, 2509.16245. Source ↗
Back to citation

4. Lee, B. K., Mayhew, E. J., Sanchez-Lengeling, B., Wei, J. N., Qian, W. W., Little, K. A., Andres, M., Nguyen, B. B., Moloy, T., Yasonik, J., Parker, J. K., Gerkin, R. C., Mainland, J. D., & Wiltschko, A. B. (2023). A principal odor map unifies diverse tasks in olfactory perception. Science, 381(6661), 999–1006. Source ↗
Back to citation

5. Barsainyan, A. A., Kumar, R., Saha, P., & Schmuker, M. (2023). OpenPOM — Open Principal Odor Map [Software; revision 5434c4b7badebca8e7d180aa34d2f787402121ee]. Retrieved October 11, 2026. Source ↗
Back to citation

6. Zhao, J., Chen, A. Q., Ryu, J., & del Mármol, J. (2024). Structural basis of odor sensing by insect heteromeric odorant receptors. *Science*, 384(6703), 1460–1467. Source ↗
Back to citation

7. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. Source ↗
Back to citation

8. Mirdita, M., Schütze, K., Moriwaki, Y., Heo, L., Ovchinnikov, S., & Steinegger, M. (2022). ColabFold: making protein folding accessible to all. *Nature Methods*, 19(6), 679–682. Source ↗
Back to citation

9. Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli, O., Shmueli, Y., dos Santos Costa, A., Fazel-Zarandi, M., Sercu, T., Candido, S., & Rives, A. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637), 1123–1130. Source ↗
Back to citation

10. Trott, O., & Olson, A. J. (2010). AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2), 455–461. Source ↗
Back to citation

11. Turner, D. B., & Willett, P. (2000). The EVA spectral descriptor. European Journal of Medicinal Chemistry, 35(4), 367–375. Source ↗
Back to citation

12. Carey, A. F., Wang, G., Su, C. Y., Zwiebel, L. J., & Carlson, J. R. (2010). Odorant reception in the malaria mosquito Anopheles gambiae. *Nature*, 464(7285), 66–71. Source ↗

Back to citation

13. Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 57(1), 289–300. Source ↗
Back to citation

14. Eastman, P., Galvelis, R., Peláez, R. P., Abreu, C. R. A., Farr, S. E., Gallicchio, E., Gorenko, A., Henry, M. M., Hu, F., Huang, J., Krämer, A., Michel, J., Mitchell, J. A., Pande, V. S., Rodrigues, J. P., Rodriguez-Guerra, J., Simmonett, A. C., Singh, S., Swails, J., Turner, P., Wang, Y., Zhang, I., Chodera, J. D., De Fabritiis, G., & Markland, T. E. (2024). OpenMM 8: Molecular Dynamics Simulation with Machine Learning Potentials. *The Journal of Physical Chemistry B*, 128(1), 109–116. Source ↗
Back to citation

15. Bohbot, J. D., Jones, P. L., Wang, G., Pitts, R. J., Pask, G. M., & Zwiebel, L. J. (2011). Conservation of Indole Responsive Odorant Receptors in Mosquitoes Reveals an Ancient Olfactory Trait. *Chemical Senses*, 36(2), 149–160. Source ↗
Back to citation

16. Israelewitz, B., Baudry, J., Gullingsrud, J., Kosztin, D., & Schulten, K. (2001). Steered molecular dynamics investigations of protein function. *Journal of Molecular Graphics and Modelling*, 19(1), 13–25. Source ↗
Back to citation

17. Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B. A., Thiessen, P. A., Yu, B., Zaslavsky, L., Zhang, J., & Bolton, E. E. (2025). PubChem 2025 update. *Nucleic Acids Research*, 53(D1), D1516–D1525. Source ↗
Back to citation

18. Varadi, M., Anyango, S., Deshpande, M., Nair, S., Natassia, C., Yordanova, G., Yuan, D., Stroe, O., Wood, G., Laydon, A., Žídek, A., Green, T., Tunyasuvunakool, K., Petersen, S., Jumper, J., Clancy, E., Green, R., Vora, A., Lutfi, M., Figurnov, M., Cowie, A., Hobbs, N., Kohli, P., Kleywegt, G., Birney, E., Hassabis, D., & Velankar, S. (2022). AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research*, 50(D1), D439–D444. Source ↗
Back to citation

19. Münch, D., & Galizia, C. G. (2016). DoOR 2.0 - Comprehensive Mapping of Drosophila melanogaster Odorant Responses. *Scientific Reports*, 6, 21841. Source ↗
Back to citation

20. Eberhardt, J., Santos-Martins, D., Tillack, A. F., & Forli, S. (2021). AutoDock Vina 1.2.0: New Docking Methods, Expanded Force Field, and Python Bindings. *Journal of Chemical Information and Modeling*, 61(8), 3891–3898. Source ↗
Back to citation

21. Saberi, M., & Seyed-allaei, H. (2016). Odorant receptors of Drosophila are sensitive to the molecular volume of odorants. Scientific Reports, 6, 25103. Source ↗
Back to citation

22. Chacko, R., Jain, D., Patwardhan, M., Puri, A., Karande, S., & Rai, B. (2020). Data based predictive models for odor perception. Scientific Reports, 10, 17136. Source ↗
Back to citation

23. Gabler, S., Soelter, J., Hussain, T., Sachse, S., & Schmuker, M. (2013). Physicochemical vs. Vibrational Descriptors for Prediction of Odor Receptor Responses. Molecular Informatics, 32(9–10), 855–865. Source ↗
Back to citation

24. Franco, M. I., Turin, L., Mershin, A., & Skoulakis, E. M. C. (2011). Molecular vibration-sensing component in Drosophila melanogaster olfaction. Proceedings of the National Academy of Sciences, 108(9), 3797–3802. Source ↗
Back to citation

25. Block, E., et al. (2015). Implausibility of the vibrational theory of olfaction. Proceedings of the National Academy of Sciences, 112(21), E2766–E2774. Source ↗
Back to citation

26. Williams, J. Z., Marszalek, P. E., & Yang, W. (2026). Evidence from computational infrared spectroscopy against vibrational detection of propionate by human olfactory receptor OR51E2. European Biophysics Journal, 55(4), 631–638. Source ↗
Back to citation

27. McConachie, G., Duniec, E., Guerina, F., Younger, M., & DePasquale, B. (2026). Low rank adaptation of chemical foundation models generate effective odorant representations. International Conference on Learning Representations (ICLR 2026). Source ↗
Back to citation

28. de Fouchier, A., Walker, W. B., Montagné, N., Steiner, C., Binyameen, M., Schlyter, F., Chertemps, T., Maria, A., François, M. C., Monsempes, C., Anderson, P., Hansson, B. S., Larsson, M. C., & Jacquin-Joly, E. (2017). Functional evolution of Lepidoptera olfactory receptors revealed by deorphanization of a moth repertoire. *Nature Communications*, 8, 15709. Source ↗
Back to citation

29. Jacquin-Joly, E., Legeai, F., Montagné, N., Monsempes, C., François, M. C., Poulain, J., Gavory, F., Walker III, W. B., Hansson, B. S., & Larsson, M. C. (2012). Candidate chemosensory Genes in Female Antennae of the Noctuid Moth Spodoptera littoralis. *International Journal of Biological Sciences*, 8(7), 1036–1050. Source ↗
Back to citation

30. Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N., & Bourne, P. E. (2000). The Protein Data Bank. *Nucleic Acids Research*, 28(1), 235–242. Source ↗
Back to citation

31. Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W., & Klein, M. L. (1983). Comparison of simple potential functions for simulating liquid water. *The Journal of Chemical Physics*, 79(2), 926–935. Source ↗
Back to citation

32. Maier, J. A., Martinez, C., Kasavajhala, K., Wickstrom, L., Hauser, K. E., & Simmerling, C. (2015). ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *Journal of Chemical Theory and Computation*, 11(8), 3696–3713. Source ↗
Back to citation

33. OpenMM contributors. (n.d.). OpenMM User Guide 7.7: Running simulations, AMBER14 force fields [Software documentation; includes ff14SB and Lipid17 parameter-file descriptions]. Retrieved October 11, 2026. Source ↗
Back to citation

34. Wang, J., Wolf, R. M., Caldwell, J. W., Kollman, P. A., & Case, D. A. (2004). Development and testing of a general amber force field. *Journal of Computational Chemistry*, 25(9), 1157–1174. Source ↗
Back to citation

35. Jakalian, A., Jack, D. B., & Bayly, C. I. (2002). Fast, efficient generation of high-quality atomic charges. AM1-BCC model: II. Parameterization and validation. *Journal of Computational Chemistry*, 23(16), 1623–1641. Source ↗
Back to citation

36. Grossfield, A., Patrone, P. N., Roe, D. R., Schultz, A. J., Siderius, D., & Zuckerman, D. M. (2018). Best Practices for Quantification of Uncertainty and Sampling Quality in Molecular Simulations [Article v1.0]. *Living Journal of Computational Molecular Science*, 1(1), 5067. Source ↗
Back to citation

37. Zhang, C., Shine, M., Pyle, A. M., & Zhang, Y. (2022). US-align: universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nature Methods*, 19(9), 1109–1115. Source ↗
Back to citation

38. Pullmann-Lindsley, H., Huff, R. M., Boyi, J., & Pitts, R. J. (2024). Odorant receptors for floral- and plant-derived volatiles in the yellow fever mosquito, Aedes aegypti (Diptera: Culicidae). *PLOS ONE*, 19(5), e0302496. Source ↗
Back to citation

Prediction comparison — Zenodo (publication pending) · Odor ranking — Zenodo (publication pending) · Cycle content (English / Traditional Chinese) · Full research details (English / Traditional Chinese) · Terminology & references