LOW RANK ADAPTATION OF CHEMICAL FOUNDATION MODELS GENERATES EFFECTIVE ODORANT REPRESENTATIONS

Anonymous authorsPaper under double-blind review

ABSTRACT

Featurizing odorants to enable robust prediction of their properties is difficult due to the complex activation patterns that odorants evoke in the olfactory system. Structurally similar odorants can elicit distinct activation patterns in both the sensory periphery (i.e., at the receptor level) and downstream brain circuits (i.e., at a perceptual level). Despite efforts to design or discover features for odorants to better predict how they activate the olfactory system, we lack a universally accepted way to featurize odorants. In this work, we demonstrate that featurebased approaches that rely on pre-trained foundation models do not significantly outperform classical hand-designed features, but that targeted foundation model fine-turning can increase model performance beyond these limits. To show this, we introduce a new model that creates olfaction-specific representations: LoRAbased Odorant-Receptor Affinity prediction with CROSS-attention (LORAX). We compare existing chemical foundation model representations to hand-designed physicochemical descriptors using feature-based methods and identify large information overlap between these representations, highlighting the necessity of finetuning to generate novel and superior odorant representations. We show that LO-RAX produces a feature space more closely aligned with olfactory neural representation, enabling it to outperform existing models on predictive tasks.

1 Introduction

A pervasive question in olfaction is: 'What is the best way to represent odorants?'. In other words, what are the best features to represent an odorant that are predictive of how that odorant will interact with the olfactory system? This question has remained elusive due to complex structure-odor relationships, where similar odorants can elicit divergent neural responses (Sell, 2006). Early efforts, inspired by chemoinformatics, focused on sets of hand-selected physicochemical descriptors (molecular weight, ring count, etc.) to best capture trends in neural data (Schmuker et al., 2007; Haddad et al., 2008a; Haddad et al., 2008b; Boyle et al., 2013; Gabler et al., 2013), and these descriptors remain a central tool in olfactory neuroscience (Pashkovski et al., 2020; Yang et al., 2023). More recently, deep learning has enabled the extraction of data-driven chemical features that facilitate a more accurate and comprehensive mapping of odor space (Seshadri et al., 2022; Lee et al., 2023; Shuvaev et al., 2024).

Within data-driven models, two broad approaches have emerged (Figure 1): supervised training of tailored architectures, such as graph neural networks (Figure 1B, left), and foundation models (Figure 1A and B, center), which are self-supervised approaches that leverage unlabeled data to generate useful chemical representations for post-hoc application (Shin et al., 2023; Taleb et al., 2024). The latter approach offers a promising solution to the challenges of constructing representations, particularly for olfaction where datasets are typically small. Moreover, given the estimated billions of potential odorants (Mayhew et al., 2022), self-supervised approaches may generalize more effectively by leveraging large unlabeled molecular datasets to capture trends not otherwise discernible.

Despite the potential of foundation models, their application to olfaction has been limited and not quantitatively and exhaustively compared to earlier approaches. Here, we evaluate several chemical foundation models and find they *do not* improve predictive performance over classical physicochem-

ical descriptors on odorant-receptor binding tasks (Section 3). Using two statistical shape analysis methods, canonical correlation analysis (CCA) and orthogonal Procrustes analysis, we find many of these representations have strong information overlap with physicochemical features, suggesting that self-supervised learning alone does not yield features well suited for odorant-receptor binding prediction.

To address this limitation, we introduce a model based on fine-tuning to produce novel odorant representations tailored to olfacation (Figure 1C). Our model, LoRA-based Odorant-Recptor Affinity prediction with CROSS-attention (LORAX; Figure 3), employs low-rank adaptations (Hu et al., 2021) of chemical foundation models to fine-tune their representations, making them more amenable for olfaction (Section 4). We demonstrate LORAX's odorant representation allows for better performance on odorant-receptor datasets and better generalization to unseen odorants. Furthermore, we show that the LORAX representation is more aligned with neural representation, providing insight into the cause of increased model performance.

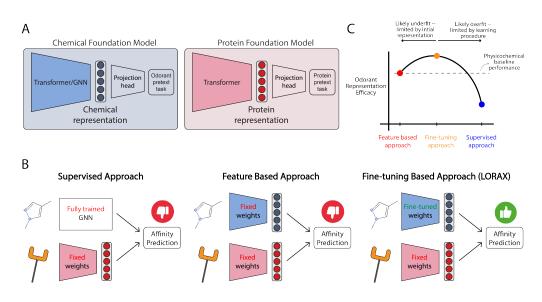


Figure 1: LORAX presents a new approach to incorporate chemical foundation models to predict odorant-receptor affinity. (A) Architecture of chemical and protein foundation models. (B) Current approaches to predicting affinity of odorant-receptor binding. (C) Putative thought model given our approach and results: We hypothesize that fine-tuning allows for creation of optimal odorant representations that outperform existing approaches.

2 RELATED WORKS

Physicochemical features to predict affinity. Guo & Kim (2010) predicted odorant-receptor relationships by representing odorants with GRIND descriptors (Pastor et al., 2000) and encoding receptors with multidimensional scaling vectors, using partial least squares. Boyle et al. (2013) and Gabler et al. (2013) use optimized physicochemical descriptors to predict receptor activation, employing hierarchical clustering, and support vector machines and random forests, respectively. Cong et al. (2022) uses physicochemical descriptors of odorants and multiple sequence alignment similarities of receptors to train a random forest that generates predictions. Berwal et al. (2025) uses molecular docking simulations with K-means clustering of odorants based on physicochemical descriptors to generate affinity predictions.

Supervised learning to predict affinity. Achebouche et al. (2022) develop graph and convolutional neural networks to predict receptor activation. Hladiš et al. (2023) and Chithrananda et al. (2024) use graph neural networks (GNNs) to represent odorants combined with a protein foundation model to predict odorant-receptor binding. None of these approaches investigate chemical foundation models

as odorant featurizers: while Chithrananda et al. (2024) and Hladiš et al. (2023) incorporate *protein* foundation models into odorant-receptor affinity prediction, they employ randomly initialized GNN architectures with supervised training rather than pre-trained models.

Pre-trained and fine-tuned chemical foundation models in olfaction. Shin et al. (2023) use the chemical foundation models MolCLR (Wang et al., 2022) and SMILES Transformer (Honda et al., 2019) to represent odorants to predict odor percepts. Similarly, Taleb et al. (2024) use MoLformer (Ross et al., 2022) to featurize odorants and predict their percept. Both groups found chemical foundation models offer valuable features that can aid odor prediction. However, these groups do not test chemical foundation models to predict odorant-receptor affinity and do not utilize as broad a selection of chemical foundation models as we do. To our knowledge, chemical foundation models have not been applied to predict odorant-receptor interactions, and different chemical foundation models have not been systematically benchmarked to evaluate their effectiveness. Furthermore, the LoRA-based fine tuning we introduce and the performance improvements it produces have not been documented for olfactory tasks.

3 DO CHEMICAL FOUNDATION MODELS CREATE EFFECTIVE FEATURES?

We benchmarked a variety of chemical foundation models to examine which, if any, enhance odorant-receptor affinity prediction. To do this, we compiled three datasets highlighted in the section below.

3.1 Datasets

Each dataset includes the amino acid sequence of the receptor, a chemical identifier of the odorant (e.g., SMILES string (Weininger, 1988)), and an experimental measurement of odorant-receptor affinity. These datasets span mammals and non-mammals and employ diverse methodologies for assessing odorant-receptor interactions.

Hallem Dataset. Hallem & Carlson (2006) reports the responses of 24 *D. melanogaster* receptors to 110 odorants using electrophysiology, yielding a total of 2,640 odorant-receptor pairs. Responses are average spikes per second across a range of concentrations. We normalize the data by *z*-scoring all responses, using the mean and standard deviation of the entire dataset (i.e., we don't *z*-score responses for each receptor individually).

Carey Dataset. Carey et al. (2010) reports the responses of 50 A. gambiae receptors to 110 odorants using electrophysiology, yielding a total of 5,500 odorant-receptor pairs. Responses are quantified as in the Hallem dataset. We normalize the dataset by z-scoring all responses, as in the Hallem dataset.

M2OR. Mammalian odorant-receptor interactions collected from 45 different sources (Lalis et al., 2024). The dataset contains 771 odorants and 1402 receptors across 16 species totaling 53,444 paired responses. We used a subset of 46,563 odorant-receptor pairs to compare LORAX to previous models. Each paired response is categorized as either one for responsive or zero for non-responsive.

M2OR (EC50). The M2OR dataset has a variety of data quality. Much of the data comes from 'primary' or 'secondary' screening, which are not as precise as data collected in the Hallem and Carey datasets. For benchmarking, we use the M2OR 'EC50' data, as it is the highest quality and foundation model comparison does not require the entire dataset. This subset of the data contains 474 odorants and 503 receptors, yielding a total of 5,834 odorant-receptor pairs.

3.2 BENCHMARKING MODELS

To effectively benchmark chemical foundation models against standard methods, we select six physicochemical representations, five transformer-based foundation models, three GNN-based foundation models, and one randomly generated representation as a lower bound (see Appendix A for details). We use three feature-based models (i.e., models that are not fine-tuned) to benchmark these representations. For models that incorporate a protein foundation model, we use ESM (Rives et al., 2021). As we are predominantly interested in odorant representations, we used only one protein foundation model.

Molecule only model (MO). A ridge regression model using only odorant representations. This model is trained on every receptor individually and results are averaged over all receptors. This model assesses the effectiveness of the odorant representations alone, testing whether only chemical information can predict receptor activation.

Molecule + protein model (MP). A ridge regression model² that uses an odorant representation concatenated with a protein representation to predict affinity. This model tests the ability of the odorant representation to predict binding with additional receptor target information.

ProSmith. We adapt a state-of-the-art multi-modal transformer (Kroll et al., 2024) to test odorant representations for affinity prediction.

Further details of these benchmarking models are provided in Appendix B. Performance was evaluated using 5-fold cross-validation with random splits³. Averaged prediction metrics are reported for the Carey, Hallem, and M2OR (EC50) datasets.

3.3 CHEMICAL FOUNDATION MODELS DO NOT IMPROVE ODORANT-RECEPTOR AFFINITY PREDICTION

Benchmarking results are shown in Table 1 and Appendix E. We focus on the Carey dataset but the observations and conclusions are consistent across the other datasets as well.

Table 1: Affinity prediction performance across odorant representations for the Carey dataset. Reported as mean $R^2 \pm {\rm standard}$ deviation across 5-fold cross validation. Best performing representations are bolded.

Odorant representation	MO	MP	ProSmith
CATS	0.110 ± 0.227	0.271 ± 0.013	0.637 ± 0.019
MordredDescriptors	0.178 ± 0.261	0.295 ± 0.022	0.719 ± 0.021
Pharmacophore2D	0.175 ± 0.261	0.282 ± 0.018	0.548 ± 0.014
RDKitDescriptors2D	0.034 ± 0.404	0.264 ± 0.021	$\textbf{0.720} \pm \textbf{0.026}$
ScaffoldKeyCalculator	0.005 ± 0.280	0.246 ± 0.016	0.594 ± 0.022
ECFP	0.207 ± 0.410	0.297 ± 0.020	0.664 ± 0.034
Roberta-Zinc480M-102M	0.181 ± 0.237	0.295 ± 0.020	0.672 ± 0.030
GPT2-Zinc480M-87M	0.121 ± 0.266	0.296 ± 0.023	0.668 ± 0.027
ChemGPT-19M	-0.029 ± 1.046	0.285 ± 0.028	0.685 ± 0.024
MolT5	0.215 ± 0.279	0.302 ± 0.018	0.676 ± 0.022
ChemBERTa-77M-MTR	0.218 ± 0.298	0.290 ± 0.019	0.704 ± 0.017
gin_supervised_infomax	0.243 ± 0.273	$\textbf{0.303} \pm \textbf{0.021}$	0.671 ± 0.016
gin_supervised_edgepred	0.189 ± 0.308	0.295 ± 0.019	0.687 ± 0.027
gin_supervised_contextpred	$\textbf{0.257} \pm \textbf{0.255}$	$\textbf{0.303} \pm \textbf{0.021}$	0.699 ± 0.033
random	-0.071 ± 0.137	0.294 ± 0.023	0.586 ± 0.037

From these results, we draw three conclusions. First, molecular information alone is insufficient to do prediction. This conclusion is supported by the MO column in Table 1. We see low R^2 scores with high variance across folds, indicating that solely using the odorant representations is insufficient for accurate prediction. Both physicochemical and chemical foundation model representations lack sufficient contextual information to predict binding.

Second, incorporating protein information is essential for achieving reliable prediction. This can be seen by the large performance increase from the MO to MP models, showing that incorporating receptor information increases R^2 scores and decreases variance across folds. Additionally, we obtain very high predictive power using the ProSmith model when both receptor and odorant representations are combined through a transformer. This highlights the importance of multi-modal approaches in this domain.

¹For the M2OR (EC50) dataset, we exclude receptors with fewer than 50 entries when training the MO model, as this amount of data is too limited to ensure reliable training.

²The MP and MO models are a logistic regression with an L2 penalty for the M2OR (EC50) dataset.

³We use a random stratified split for the MO model on the M2OR (EC50) dataset.

Third, the choice of odorant representation has little effect on predictive performance. We performed a one-way ANOVA followed by a Bonferroni-corrected post-hoc Tukey test on our ProSmith results. Our analysis showed, with the exception of 'Pharmacophore2D' and 'ScaffoldKeyCalculator', all representations performed equivalently. This is surprising, especially given the success of foundation models in other domains. One would expect that chemical foundation model representations would provide richer molecular context for predicting odorant-receptor affinity, but given results highlighted here, it seems these representations do not offer important features that enhance predictability.

3.4 CHEMICAL FOUNDATION MODEL REPRESENTATIONS ARE ALIGNED WITH PHYSICOCHEMICAL FEATURES

To understand the failure of foundation models to improve prediction, we analyzed foundation model odorant representations by calculating dissimilarity metrics between each representation (Williams et al., 2022). We embedded all odorants in the Carey dataset into their respective feature spaces and calculated both orthogonal Procrustes distance metrics (Figure 2a) and CCA distance metrics (Figure 2b) between each pair of representations (see Appendix C for details).

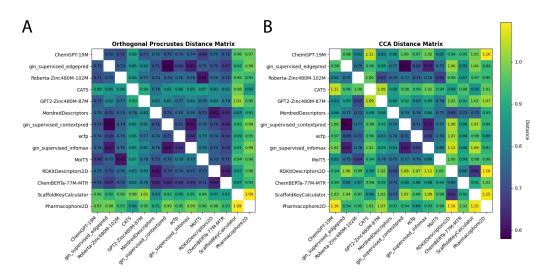


Figure 2: Distance metrics between odorant representations in the Carey dataset. Shown are (A) orthogonal Procrustes distance and (B) CCA distance metrics. See Williams et al. (2022) and Appendix C for details on these metrics.

From these distance matrices, particularly the Procrustes matrix, we see that many of the odor spaces are aligned. The only two representations that are visibly dissimilar ('ScaffoldKeyCalculator' and 'Pharmacophore2D') are the worst performing representations (Table 1), while every other representation forms a block of low similarity scores in the upper left. The odorant representations that are the most similar also share similar performance in Table 1. In addition to performance trends, this analysis uncovers other patterns present in the representations. For example, the most similar representations, 'gin_supervised_contextpred' and 'gin_supervised_edgepred', are generated from two instantiations of the same GNN (Hu et al., 2020). We also see that the 'ChemBERTa-77M-MTR' representation, a foundation model trained on 77 million SMILES strings to predict RDKit descriptors, is similar to the 'RDKitDescriptors2D' representation. Analyzing these spaces, and their dissimilarities, explains why many of these descriptors perform the same: there is heavy overlap of information content, making the representations fairly homogeneous and redundant.

4 LORAX: Fine tuning chemical foundation models with LorA

Given the observed redundancy among chemical representations, we developed a model that is able to fine-tune, and adapt, these chemical foundation models to create an *olfaction-tailored represen*-

tation. All models tested in Section 3 are feature-based, relying on fixed chemical (and protein) representations rather than adapting them through fine-tuning. A key advantage of fine-tuning is that it produces a new, task-tailored representation rather than depending solely on pre-trained embeddings. Prior work has demonstrated that fine-tuning is generally more effective than feature-based methods (Devlin et al., 2019), and adapter techniques, such as LoRA (Hu et al., 2021), enable efficient fine-tuning of large models. Given that many of the foundation model representations have large information overlap, we predicted that refining these representations would yield more informative and task relevant features for olfactory tasks. Motivated by these considerations, we introduce LORAX (Figure 3).

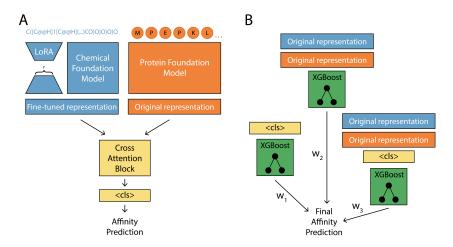


Figure 3: LORAX model. (A) A multi-modal LoRA adapted transformer model with a cross attention block for odorant-receptor prediction. r is the rank of the low rank adapter matrices. (B) An ensemble of gradient boosted decision tree (XGBoost) models are used to enhance performance while also providing model interpretability. A weighted sum of these XGBoost models is used to get the final prediction. 'Original representation' means the non-LoRA fine-tuned foundation model representation. Orange represents the protein foundation model and blue represents the chemical foundation model.

LORAX is a multi-modal transformer that incorporates a protein and a chemical foundation model to predict affinity (Figure 3a). LORAX uses LoRA to update the chemical representation over training. The protein and chemical foundation model can be any model on Huggingface, but for all analysis in this work, we use ESM (Rives et al., 2021) to represent receptors and ChemBERT-77M-MTR (Ahmad et al., 2022) to represent odorants. We take inspiration from Kroll et al. (2024) and incorporate an ensemble of gradient boosted decision tree models (XGBoost, (Chen & Guestrin, 2016)) to make final predictions (Figure 3b). Each XGBoost model in the ensemble uses different inputs: (1) the transformer's penultimate representation, which we call the <cls> token, (2) the raw, concatenated chemical and protein foundation model representations. Architecture and training details are available in Appendix D.

4.1 LORAX CONSTRUCTS A BETTER ODORANT REPRESENTATION

We trained LORAX on the Carey dataset and found it performs almost identically to ProSmith (LORAX R^2 : 0.712 ± 0.032 , ProSmith R^2 : 0.703 ± 0.017 , p=0.264 paired t-test, data not shown). However, when we interrogate the trained models, we uncovered nuanced differences between them. Because ProSmith and LORAX have similar architectures and both employ identical XGBoost ensembles, a direct comparison between them is possible. When examining the ensemble weights (Figure 4B), we found that ProSmith places high confidence in the XGBoost model that uses only the original chemical and protein foundation model representations. In other words, the additional representation produced by the transformer (the <cls>token) is largely ignored. This is supported by the relatively weak standalone validation accuracy (Figure 4A, gray) showing ProSmith's trans-

former lacks predictive performance. When we compute these metrics for LORAX (Figure 4, blue), we see the validation accuracy of its' multi-modal transformer is much higher and the weights placed on XGBoost models without the <cls> token are lower. We conclude that LORAX extracts more relevant information for the task from the protein and chemical foundation models.

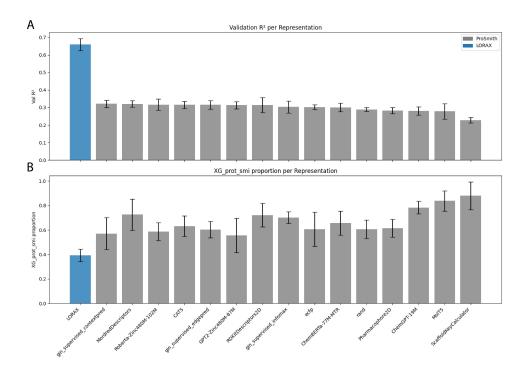


Figure 4: LORAX and ProSmith representation comparison after training on the Carey dataset. (A) Validation \mathbb{R}^2 scores for both models. (B) The weight assigned to the XGBoost model that takes only the original chemical foundation model and protein foundation model representations.

With this context, we next examined whether the differences in representation weighting between ProSmith and LORAX translate into improved odorant representations across different scenarios. We first wanted to assess the generalizability of LORAX. Using the Carey dataset, we trained both LORAX and ProSmith⁴ on two scenarios: 1) generalize to unseen receptors (i.e., receptors in the test set are never seen in training or validation), and 2) generalize to unseen odorants (i.e., odorants in the test set never appear in training or validation). Table 2 shows that LORAX generalizes better to unseen odorants but does not generalize better to unseen receptors. LORAX's improved performance on odorant generalization is noteworthy; although not quite statistically significant (p=0.06, paired t-test), it suggests that the LORAX representation captures more salient chemical information for odorant-receptor binding prediction, facilitating better generalization. Nevertheless, overall generalization performance was limited for both models.

Table 2: LORAX vs. ProSmith performance under two generalization settings. Shown are \mathbb{R}^2 scores for the test set for each random split. Bolded are the best performing models.

Scenario	Method	rand_split_1	rand_split_2	rand_split_3	rand_split_4	rand_split_5	avg	std
Unseen odorants	ProSmith	-2.098	-0.812	0.364	0.246	0.288	-0.402	1.064
	LORAX	-1.039	0.201	0.472	0.292	0.298	0.045	0.614
Unseen receptors	ProSmith	0.072	0.246	0.263	-0.137	-0.003	0.088	0.169
	LORAX	-0.032	0.051	0.031	0.133	0.117	0.060	0.067

⁴We train ProSmith with the ChemBERT-77M-MTR model for appropriate comparison.

To further interrogate LORAX, we examine how the model performs relative to other models on the full M2OR dataset, a larger dataset that presents greater diversity and complexity, offering a more rigorous test of model performance. We utilize the weighting scheme outlined in Hladiš et al. (2023) to account for differing data quality. From Table 3, we can see that LORAX outperforms Hladiš et al. (2023) and ProSmith on this larger dataset. This confirms that fine-tuning chemical foundation models offers an important avenue to improve odorant-receptor prediction. We note that LORAX does not outperform MolOR (Chithrananda et al., 2024) but their results are reported as AUROC over *random seeds* rather than across folds of M2OR, making direct comparison inappropriate.

Table 3: Comparison of MolOR, Hladiš et al., ProSmith, and LORAX across multiple metrics. LORAX (C) fine-tunes only the chemical foundation model; LORAX (P + C) fine-tunes both chemical and protein foundation models. Values for MolOR and Hladiš et al. are taken from respective publications. Uncertainties for Hladiš et al., ProSmith, and LORAX are standard deviation across 5 fold cross validation. Uncertainties for MolOR are standard deviation across 5 random seeds. Bold indicates highest value in each column. n/a entries indicate that the authors did not report that metric.

	AUROC	AveP	Precision	Recall	F-score	MCC
MolOR (Weighted)	86.78 ± 0.80	n/a	n/a	n/a	n/a	n/a
Hladiš et. al.	n/a	$\textbf{0.780} \pm \textbf{0.01}$	0.689 ± 0.02	0.698 ± 0.04	0.693 ± 0.02	0.605 ± 0.02
ProSmith	$\textbf{90.47} \pm \textbf{0.01}$	0.584 ± 0.03	$\textbf{0.764} \pm \textbf{0.06}$	0.671 ± 0.06	0.712 ± 0.02	0.641 ± 0.03
LORAX(P+C)	82.43 ± 0.01	0.776 ± 0.03	0.729 ± 0.06	0.727 ± 0.03	0.727 ± 0.02	0.650 ± 0.03
LORAX (C)	83.24 ± 0.01	0.778 ± 0.03	0.710 ± 0.05	$\textbf{0.754} \pm \textbf{0.02}$	$\textbf{0.730} \pm \textbf{0.03}$	$\textbf{0.651} \pm \textbf{0.04}$

We also tested a version of LORAX that uses both a LoRA adapted protein foundation model and a LoRA adapted chemical foundation model (LORAX (P + C)). Interestingly, this model performs slightly worse than the model with only chemical foundation model fine-tuning, even though LORAX (P + C) has \sim 2 million more parameters than LORAX (C). We conclude that the task benefits mainly from LORAX's improved chemical representations, while the original protein representations are already sufficient. In summary, we found that LORAX fine-tuning produced models with more refined <cls> tokens, translating to improved prediction on M2OR and improved generalizability to unseen odorants.

4.2 LORAX ODOR SPACE IS SIMILAR TO UNDERLYING NEURAL REPRESENTATION

Given LORAX's promising capabilities, we analyzed the odorant representations it produces. We plot CCA distances of LORAX's odor space for the Carey dataset in Figure 5.

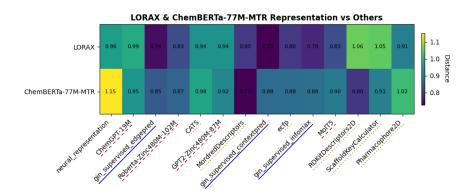


Figure 5: CCA distance matrix showing the LORAX and ChemBERTa-77M-MTR representations on the Carey dataset. 'neural_representation' encodes each odorant as a feature vector based on the responses it elicits from each neuron. Red dashed lines indicate transformer-based foundation model representations, blue solid lines indicate GNN-based foundation model representations, and orange dotted lines indicate physicochemical representations. Full distance matrices shown in Figure 7.

Additionally, for this distance matrix, we include a 'neural_representation', in which each odorant is represented as a vector of the neural responses it elicits. We compare the LORAX representation with 'ChemBERTa-77M-MTR', the foundation model that is adapted using LoRA during training, for the Carey dataset. The 'ChemBERTa-77M-MTR' representation can be considered the initial state before fine-tuning while 'LORAX' represents the final state after adaptation.

We found the LORAX representation is more dissimilar from many of the physicochemical descriptors than ChemBERTa-77M-MTR and more similar to all graph based methods (ECFP and all the GNN foundation models). Additionally, the LORAX representation is more aligned with the neural representation than the ChemBERTa-77M-MTR representation. These analyses reveal that LORAX is generating a unique representation, distinct from the original foundation model, that is better suited to describe odorant-receptor relationships and therefore generating a more informative odor space.

5 DISCUSSION

Identifying representations that best characterize odorants for olfactory prediction has been a major research focus. We systematically evaluated representations from several pre-trained chemical foundation models and found they did not provide an improvement over hand-tuned physicochemical descriptors. To move beyond this limitation, we introduced a new model, LORAX, for fine-tuning chemical foundation models to produce richer and more powerful representations. Collectively, we present two novel and interconnected findings: 1) To our knowledge, we performed the first and most comprehensive analysis of chemical foundation models applied to olfaction to demonstrate their effective equivalence with physicochemical descriptors; 2) We introduced the first application of LoRA fine-tuning of chemical foundation models for olfaction, demonstrating the improvement that fine-tuning offers. Our results demonstrate that while existing chemical foundation models may capture overlapping information for olfactory datasets, targeted fine-tuning creates specialized and potentially more powerful representations than those obtained from pre-trained models alone. We argue that these models have significant potential for olfactory neuroscience, particularly in addressing key challenges such as limited data and vast chemical search spaces.

We hypothesize the targeted fine-tuning LoRA provides identifies a 'sweet spot' that balances the strengths of alternative approaches (Figure 1C). Although speculative, support for this hypothesis is found in Table 3. Supervised approaches that lack the context of pre-trained models underperform likely due to overfitting (Hladiš et al., 2023). Pre-trained models alone likewise underperform because a limited initial representation leads to underfitting. How then can we *selectively* modify a pre-trained representation to better generalize to a task without running into problems that come with non-selective fine-tuning or supervised learning alone? The selective nature of LoRA's low-rank assumption enables robust generalization while not entirely erasing the benefits that come from pre-training. To the best of our knowledge, this work presents the first application of LoRA to olfaction, and given its strength and natural approach to balancing over- and underfitting when fine-tuning pre-trained models, we expect it to provide impactful and wide-ranging future uses in olfaction.

There remain many exciting avenues for future exploration with this modeling framework. A dominant question in olfaction is how do chemical features produce olfactory perception? (Zhou et al., 2018; Seshadri et al., 2022; Lee et al., 2023; Yang et al., 2023; Qian et al., 2023; Shuvaev et al., 2024). Shin et al. (2023) and Taleb et al. (2024) found that pre-trained models with indiscriminate or no fine-tuning improved performance on perception tasks; our results suggest further improvement can be achieved with LoRA-based fine-tuning. Additionally, LORAX is modular by design, allowing for the interchange of protein and chemical foundation models. While our analysis utilizes ChemBERTa-77M-MTR (Ahmad et al., 2022), the field of chemical foundation modeling is rapidly advancing with larger and more diverse models. Systematic benchmarking and further training with these new models may yield even more optimal odorant representations.

Despite our progress, several important questions remain. For example, how robust are the LORAX-learned representations when transferring to different olfactory datasets? Which components of LORAX contribute the most to prediction performance? Why does simultaneous fine-tuning of both protein and chemical foundation models hamper performance? We leave these to future work, with the anticipation that continued research will deepen our understanding of how chemical foundation models can be effectively leveraged in olfaction.

REPRODUCIBILITY STATEMENT

All datasets used in this work are freely available and from published work linked in the text. Code of the benchmarking work in Section 3 and LORAX model will be uploaded to github upon publication. Default configurations and training regimes are outlined in the main text and appendices for all models used in this work. Statistical shape analysis methods in this work are freely available to use here.

USE OF LLMS IN THIS WORK

We use LLMs to to aid and polish writing throughout the text.

REFERENCES

- Rachid Achebouche, Aurore Tromelin, Karine Audouze, and Olivier Taboureau. Application of artificial intelligence to decode the relationships between smell, olfactory receptors and small molecules. *Scientific Reports*, 12:18817, 2022.
- Wasi Ahmad, Ethan Simon, Sambhavi Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta-2: Towards chemical foundation models. *arXiv preprint arXiv:2209.01712*, 2022. doi: 10.48550/arXiv.2209.01712. URL https://doi.org/10.48550/arXiv.2209.01712.
- Bhavika Berwal, Pinaki Saha, and Ritesh Kumar. A fully in silico protocol to understand olfactory receptor-odorant interactions. *ACS Omega*, 10(23):24030–24049, 2025. doi: 10.1021/acsomega. 4c08181. URL https://doi.org/10.1021/acsomega.4c08181.
- S. M. Boyle, S. McInally, and A. Ray. Expanding the olfactory code by in silico decoding of odor-receptor chemical space. *eLife*, 2:e01120, 2013.
- Alexander F Carey, Guirong Wang, Chih-Ying Su, Laurence J Zwiebel, and John R Carlson. Odorant reception in the malaria mosquito anopheles gambiae. *Nature*, 464(7285):66–71, 2010.
- Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, New York, NY, USA, 2016. doi: 10.1145/2939672.2939785.
- Shashank Chithrananda, Javier Amores, and Kevin K Yang. Mapping the combinatorial coding between olfactory receptors and perception with deep learning. *bioRxiv preprint*, pp. 2024.09.16.613334, 2024. doi: 10.1101/2024.09.16.613334. URL https://doi.org/10.1101/2024.09.16.613334.
- Xiaowei Cong, Xiangyu Chen, Yanyan Jiang, Ruqiang Li, Liang Li, Zhen Xu, Yinan Zhang, Jian Yu, Shuo Tang, Jian Zhang, et al. Large-scale g protein-coupled olfactory receptor-ligand pairing. *ACS Central Science*, 8(3):379–387, 2022.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2019. doi: 10.48550/arXiv.1810.04805.
- Chase Edwards, Tina Lai, Kevin Ros, Galen Honke, Kyunghyun Cho, and Heng Ji. Translation between molecules and natural language. *arXiv preprint arXiv:2204.11817*, 2022. doi: 10.48550/arXiv.2204.11817.
- Peter Ertl. Identification of bioisosteric scaffolds using scaffold keys. *ChemRxiv*, 2021. doi: 10. 26434/chemrxiv.13525457.v2. Preprint.
- Nathaniel Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor Coley, and Vijay Gadepally. Neural scaling of deep chemical models. *ChemRxiv*, 2022. doi: 10.26434/chemrxiv-2022-3s512. Preprint.
 - S. Gabler, J. Soelter, T. Hussain, S. Sachse, and M. Schmuker. Physicochemical vs. vibrational descriptors for prediction of odor receptor responses. *Molecular Informatics*, 32:855–865, 2013.

- Shanhe Guo and Jae Kim. Dissecting the molecular mechanism of drosophila odorant receptors through activity modeling and comparative analysis. *Proteins*, 78(2):381–399, 2010.
- R. Haddad, A. Lanjuin, Y. Weissberg, D. Harel, E. Feinstein, E. Zucker, L. D. Giorgianni, and S. 'Others'. A metric for odorant comparison. *Nature Methods*, 5:425–429, 2008a.
 - R. Haddad, H. Lapid, D. Harel, and N. Sobel. Measuring smells. *Current Opinion in Neurobiology*, 18:438–444, 2008b.
- Elissa A Hallem and John R Carlson. Coding of odors by a receptor repertoire. *Cell*, 125(1): 143–160, 2006.
 - Karl Heyer. Gpt2 zinc 87m, 2023a. URL https://github.com/kheyer/gpt2_zinc_ 87m. Accessed: 2025-05-15.
 - Karl Heyer. Roberta zinc 480m, 2023b. URL https://github.com/kheyer/roberta_zinc_480m. Accessed: 2025-05-15.
 - Matej Hladiš, Maxence Lalis, Sébastien Fiorucci, and Jérémie Topin. Matching receptor to odorant with protein language and graph neural networks. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*, 2023. URL https://openreview.net/forum?id=q9VherQJd8_.
 - Satoshi Honda, Sheng Shi, and Hiroshi R. Ueda. SMILES transformer: Pre-trained molecular fingerprint for low data drug discovery. *arXiv preprint arXiv:1911.04738*, 2019. URL https://arxiv.org/abs/1911.04738.
 - Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. doi: 10.48550/arXiv.2106.09685. URL https://doi.org/10.48550/arXiv.2106.09685.
 - Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. *arXiv preprint arXiv:1905.12265*, 2020. doi: 10.48550/arXiv.1905.12265. URL https://doi.org/10.48550/arXiv.1905.12265.
 - Damjan Kalajdzievski. A rank stabilization scaling factor for fine-tuning with lora, 2023. URL https://arxiv.org/abs/2312.03732.
 - Alexander Kroll, Sahasra Ranjan, and Martin J. Lercher. A multimodal transformer network for protein-small molecule interactions enhances predictions of kinase inhibition and enzyme-substrate relationships. *PLOS Computational Biology*, 20(5):e1012100, 2024. doi: 10.1371/journal.pcbi.1012100. URL https://doi.org/10.1371/journal.pcbi.1012100.
 - Manon Lalis, Adrien Comte, Julien Allali, Paul Le Dú, Sebastiano Daberdaku, Quentin Gaudry, Jeremie Golebiowski, Alexandre Bassée, Laurent Charlier, Charlyne Grison, et al. M2or: a database of olfactory receptor–odorant pairs for understanding the molecular mechanisms of olfaction. *Nucleic Acids Research*, 52(D1):D1370–D1379, 2024.
 - Brian K. Lee, Emily J. Mayhew, Benjamin Sanchez-Lengeling, Jennifer N. Wei, Wesley W. Qian, Kelsie A. Little, Matthew Andres, Britney B. Nguyen, Theresa Moloy, Jacob Yasonik, Jane K. Parker, Richard C. Gerkin, Joel D. Mainland, and Alexander B. Wiltschko. A principal odor map unifies diverse tasks in olfactory perception. *Science*, 381:999–1006, 2023.
- E. J. Mayhew, C. J. Arayata, R. C. Gerkin, B. K. Lee, J. M. Magill, L. L. Snyder, K. A. Little, C. W. Yu, and J. D. Mainland. Transport features predict if a molecule is odorous. *Proceedings of the National Academy of Sciences*, 119(15):e2116576119, 2022. doi: 10.1073/pnas.2116576119. URL https://doi.org/10.1073/pnas.2116576119.
 - Hirotomo Moriwaki, Yusuke-Saito Tian, Norihito Kawashita, and Tatsuya Takagi. Mordred: A molecular descriptor calculator. *Journal of Cheminformatics*, 10(1):4, 2018. doi: 10.1186/s13321-018-0258-y.

Emmanuel Noutahi, Cas Wognum, Hadrien Mary, Honoré Hounwanou, Kyle M. Kovary, Desmond Gilmour, thibaultvarin r, Jackson Burns, Julien St-Laurent, t, DomInvivo, Saurav Maheshkar, and rbyrne momatx. datamol-io/molfeat: 0.9.4 (0.9.4), 2023. URL https://doi.org/10. 5281/zenodo.8373019.

- Stan L Pashkovski, Giovanni Iurilli, Dean Brann, Daniel Chicharro, Kathryn Drummey, Kevin M Franks, Stefano Panzeri, and Sandeep R Datta. Structure and flexibility in cortical representations of odour space. *Nature*, 583(7815):253–258, 2020.
- Manuel Pastor, Giorgio Cruciani, Ian McLay, Stuart Pickett, and Stefano Clementi. Grid-independent descriptors (GRIND): A novel class of alignment-independent three-dimensional molecular descriptors. *Journal of Medicinal Chemistry*, 43(17):3233–3243, 2000.
- Wenwen Qian, Emma J. Mayhew, Alexander B. Wiltschko, Yoav Zeevi, Joel D. Mainland, et al. Metabolic activity organizes olfactory representations. *eLife*, 12:e82502, 2023.
- Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo, Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15):e2016239118, 2021. doi: 10.1073/pnas.2016239118. URL https://doi.org/10.1073/pnas.2016239118.
- David Rogers and Mathew Hahn. Extended-connectivity fingerprints (ECFP). *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010. doi: 10.1021/ci100050t.
- Jonathan Ross, Katherina Atz, Yibo Gao, Stefan Jaeger, Boris Lipp, Daniel Rufa, Paul Schrader, Robert Berendt, and Alán Aspuru-Guzik. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(11):1256–1264, 2022.
- M. Schmuker, M. de Bruyne, M. Hähnel, and G. Schneider. Predicting olfactory receptor neuron responses from odorant structure. *Chemistry Central Journal*, 1:11, 2007.
- G. Schneider, W. Neidhart, T. Giller, and G. Schmid. Scaffold-Hopping by topological pharmacophore search: A contribution to virtual screening. *Angewandte Chemie International Edition*, 38(19):2894–2896, 1999. doi: 10.1002/(SICI)1521-3773(19991004)38:19\(\cappa 2894:: AID-ANIE2894\)\)3.0.CO;2-F.
- Charles S. Sell. On the unpredictability of odor. *Angewandte Chemie International Edition*, 45: 6254–6261, 2006.
- A. Seshadri, H. A. Gandhi, G. P. Wellawatte, and A. D. White. Why does that molecule smell? ChemRxiv, 2022. doi: 10.26434/chemrxiv-2022-vs838. URL https://doi.org/10. 26434/chemrxiv-2022-vs838.
- Daniel Shin, Gao Pei, Priyadarshini Kumari, and Tarek R. Besold. Optimizing learning across multimodal transfer features for modeling olfactory perception. In *Proceedings of the International Workshop on Multimodal Learning at KDD 2023*, 2023. URL https://priyadarshini-k.com/assets/pdf/Multimodal_KDD_2023.pdf.
- Sergei Shuvaev, Kevin Tran, Ksenia Samoilova, Charles Mascart, and Alexei Koulakov. Deepnose: An equivariant convolutional neural network predictive of human olfactory percepts. In 2024 58th Asilomar Conference on Signals, Systems, and Computers, pp. 536–543, 2024. doi: 10. 1109/IEEECONF60004.2024.10942624.
- Mateusz Sypetkowski et al. On the scalability of gnns for molecular graphs. arXiv preprint arXiv:2404.11568, 2024. doi: 10.48550/arXiv.2404.11568. URL https://doi.org/10.48550/arXiv.2404.11568.
- Fares Taleb, Miguel Vasco, Andre H. Ribeiro, Mats Björkman, and Danica Kragic. Can transformers smell like humans? *arXiv preprint arXiv:2411.03038*, 2024. doi: 10.48550/arXiv.2411.03038. URL https://doi.org/10.48550/arXiv.2411.03038.

- Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R. Devon Hjelm. Deep graph infomax. *arXiv preprint arXiv:1809.10341*, 2018. doi: 10.48550/arXiv.1809. 10341.
- Yao Wang, Jun Wang, Zhaocheng Cao, and Amir Barati Farimani. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3):279–287, 2022.
- David Weininger. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. doi: 10.1021/ci00057a005.
- Alex H. Williams, Erik Kunz, Simon Kornblith, and Scott W. Linderman. Generalized shape metrics on neural representations. *arXiv preprint arXiv:2110.14739*, 2022. doi: 10.48550/arXiv.2110.14739. URL https://doi.org/10.48550/arXiv.2110.14739.
- Jia-Ying Yang et al. Restructuring of olfactory representations in the fly brain around odor relationships in natural sources. *bioRxiv preprint*, 2023. doi: 10.1101/2023.02.15.528627. URL https://doi.org/10.1101/2023.02.15.528627.
- Gengmo Zhou, Chen Shen, Huaixiu Steven Tong, Qiangfeng Cliff Dai, Shuangjia Zhang, Jinyu Li, Shengjie Li, Yufei Du, Ziqi Xiong, Fankai Liang, Jie Zhang, Wei Liu, Feng Zhang, Qiang Zhang, Tao Xu, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. arXiv preprint arXiv:2210.01733, 2022. URL https://arxiv.org/abs/2210.01733.
- Yang Zhou, Brian H. Smith, and Tatyana O. Sharpee. Hyperbolic geometry of the olfactory space. *Science Advances*, 4:eaaq1458, 2018.

A DATA REPRESENTATION

Here we describe the various methods we use to featurize odorants and receptors. For odorant featurization, we wanted a diverse array of chemical foundation models as well as physicochemical descriptors to get a good sense of what features need to go into an odorant representation to be able to effectivity predict how that odorant binds with a given receptor. The odorant representations we used are described in Table 4. For all odorant representations in Section 3 we use Molfeat (Noutahi et al., 2023) to generate representations. In Section 4, we use Huggingface to generate molecular representations as we focus on finetuning the original models when we train LORAX.

We would also like to note that the generation of chemical foundation models is a rapidly progressing field, and we realize that some of the most recent chemical foundation models such as Uni-Mol (Zhou et al., 2022) and MolGPS (Sypetkowski et al., 2024) are not present in this study. However, we would like to highlight that we cover a variety of architectures, pre-training tasks, and model sizes in our selection of chemical foundation models.

As we are mainly interested in exploring the efficacy of different odorant representations, we only use a single protein foundation model ESM (Rives et al., 2021). For benchmarking in Section 3 we use 'esm1b_t33_650M_UR50S', the esm model used by default in ProSmith (Kroll et al., 2024). In Section 4, we use a newer version of the ESM model 'esm2_t33_650M_UR50D'.

B PROBLEM STATEMENT AND MODEL DESCRIPTIONS FOR BENCHMARKING STUDY

We represent any featurization of an odorant as a vector $\mathbf{x} \in \mathbb{R}^n$, where n is the embedding dimension of the representation.⁵ Similarly, we represent a receptor as a matrix $\mathbf{P} \in \mathbb{R}^{l_p \times m}$, where l_p is the sequence length of the receptor (number of amino acids)⁶. Given an odorant–receptor pair

⁵Some transformer models produce representations that scale with the sequence length of the molecule. To compare these fairly with physicochemical descriptors and GNN-based embeddings, we mean-pool across the sequence dimension.

 $^{^6}$ The receptor embedding dimension m is determined by the protein foundation model.

Table 4: Summary of molecular descriptors used in this work.

Descriptor	Source	Description
CATS	Physicochemical	A representation designed for drug discovery applications that uses a histogram of distances between atom pairs to represent a molecule (Schneider et al., 1999). We use the 2D version.
MordredDescriptors	Physicochemical	A collection of 1,800 features representing the physicochemical, topological, geometrical, and constitutional properties of molecules (Moriwaki et al., 2018). We remove NaN values and z-score each feature.
Pharmacophore2D	Physicochemical	Computes topological distances between all pairs of atoms within a molecule. Atoms are featurized with a variety of physical properties. These feature pairs and distances are encoded into a fixed-length bit vector; we use a 2048-length vector.
RDKitDescriptors2D	Physicochemical	All 2D physicochemical descriptors from RDKit, an open-source cheminformatics toolkit for molecular representation, manipulation, and analysis.
ScaffoldKeyCalculator	Physicochemical	Encoding of simple substructure features of molecular scaffolds, such as hydrogen bond donors, acceptors, and aromatic rings, into a fixed-length binary vector (Ertl, 2021). We use a 42-length vector.
ECFP	Physicochemical	Circular molecular fingerprints that represent a molecule by iteratively encoding the local atomic environments around a defined radius (Rogers & Hahn, 2010). We use a 2000-length vector.
Roberta-Zinc480M-102M	Transformer	A RoBERTa style transformer with 102M parameters trained with a masked language model objective on 480M SMILES strings (Heyer, 2023b).
GPT2-Zinc480M-87M	Transformer	A GPT2 style transformer with 87M parameters trained on a next character prediction task on 480M SMILES strings (Heyer, 2023a).
ChemGPT-19M	Transformer	A GPT3 style transformer with over 1B parameters trained on a next character prediction task on 10M SELFIES strings (Frey et al., 2022).
MolT5	Transformer	A T5 style transformer that is trained on both SMILES strings and natural language text (Edwards et al., 2022). This model uses a 'replace corrupted spans' task, where the model samples natural language text and a SMILES string simultaneously, masks various parts of each, then trains the model to predict masked sections.
ChemBERTa-77M-MTR	Transformer	A RoBERTa style transformer trained to predict a set of 200 molecular properties on 77M SMILES strings (Ahmad et al., 2022).
gin_supervised_infomax	GNN	A graph isomorphism network (GIN) trained on 2M molecules to maximize the mutual information between local node representations and whole graph representations (Hu et al., 2020; Veličković et al., 2018).
gin_supervised_edgepred	GNN	A GIN trained on 2M molecules to predict masked edge attributes of molecular graphs (Hu et al., 2020).
gin_supervised_contextpred	GNN	A GIN trained on 2M molecules to use subgraphs of the molecules to predict their surrounding graph structures (Hu et al., 2020).
random	random	As a baseline, we generate a random vector representation for every odorant. This gives us a lower bound on performance.

 $\{x, P\}$, our goal is to predict the experimentally measured affinity y with various odorant embeddings:

$$f(\{\boldsymbol{x}, \boldsymbol{P}\}) = y. \tag{1}$$

We use three different models to assess the efficacy of chemical foundation models. A schematic of all three models can be seen in Figure 6.

B.1 MOLECULE ONLY MODEL (MO)

A ridge regression model that predicts affinity using only chemical foundation model representations. For each receptor, the model takes the form

$$f(\boldsymbol{x}|\boldsymbol{P}) = y \tag{2}$$

where only the odorant representation x is used. For each receptor, we run a nested 5-fold cross validation. We conduct a grid search over regularization parameters from 21 values spaced evenly on a logarithmic scale between 10^{-10} to 10^{10} . For the M2OR dataset, as it is a binary classification task instead of a regression, we use a L2 penalized logistic regression instead of a ridge regression. For M2OR we use a stratified split to ensure there is both positive and negative examples in each of

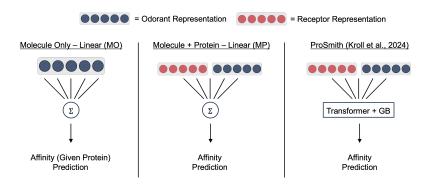


Figure 6: Schematic showing all the benchmarking models used in Section 3. Odorant representation can be any representation highlighted in Table 4. Protein representation is ESM.

the splits. Additionally, we ignore proteins with <50 odorants, as there is not enough data associated with those odorants to ensure reliable prediction.

B.2 MOLECULE + PROTEIN MODEL (MP)

A ridge regression model that uses both molecular and protein embeddings. Specifically, we concatenate the odorant embedding with a mean-pooled protein embedding

$$(\boldsymbol{x}||\boldsymbol{p}) \in \mathbb{R}^{n+m}, \quad \boldsymbol{p} = \text{meanpool}(\boldsymbol{P})$$
 (3)

where || denotes concatenation, and train

$$f(\boldsymbol{x}||\boldsymbol{p}) = y \tag{4}$$

This model, similarly to the MO model, uses a nested 5-fold cross validation scheme. We again use logistic regression with an L2 penalty to train this model on M2OR. We no longer take out any data from the M2OR dataset, as we did in the MO model.

B.3 PROSMITH

A state of the art small molecule protein multi-modal transformer model (Kroll et al., 2024). The odorant and receptor representations are projected into a shared q-dimensional space using respective MLPs, $MLP_c(\cdot)$ and $MLP_p(\cdot)$, and concatenated as

$$MLP_p(\mathbf{P})||MLP_c(\mathbf{x}) \in \mathbb{R}^{(l_p+1)\times q}$$
 (5)

Similarly to the LORAX model, the transformer f processes this sequence and produces a joint representation <cls>

$$f(MLP_n(\mathbf{P})||MLP_c(\mathbf{x})) = MLP(\langle cls \rangle) = y_{int}$$
(6)

where y_{int} is an intermediate odorant-receptor affinity prediction. The <cls> token is then combined with molecular and protein representations via an ensemble of XGBoost models:

$$y_{fin} = c_1 * h_1(\langle cls \rangle) + c_2 * h_2(\mathbf{p}||\mathbf{x}) + c_3 * h_3(\langle cls \rangle||\mathbf{p}||\mathbf{x})$$
 (7)

where y_{fin} is the final prediction of the model, $h_{1,2,3}$ are the XGBoost models, and the ensemble weights c_i satisfy $\sum_i c_i = 1$. This model is also trained using nested 5-fold cross validation. It is trained in a two step fashion, where the best performing transformer model on the validation

data is saved and used to generate the <cls> token. Next, each XGBoost is tuned independently via hyperopt, using random search over 500 iterations. The hyperparameter space and ranges are detailed in Table 5. After hyperparameter selection, the XGBoost models are trained on the training set and evaluated on the validation set to generate ensemble weights c_1 , c_2 , and c_3 . Finally, the weighted ensemble is used to predict affinities.

C METRIC SPACE DISTANCE CALCULATION DETAILS

Here we provide a brief overview of how the distance metrics in the paper were calculated. Consider two odorant representations (e.g., CATS and ECFP) that map an odorant to feature vectors $\boldsymbol{x}_i \in \mathbb{R}^n$ and $\boldsymbol{x}_j \in \mathbb{R}^m$. Embedding an entire dataset of N odorants with each representation yields matrices $\boldsymbol{X}_i \in \mathbb{R}^{N \times n}$ and $\boldsymbol{X}_j \in \mathbb{R}^{N \times m}$. To project each representation into the same space, we use PCA to reduce the dimensionality of both matrices to k=42, resulting in $\boldsymbol{X}_i, \boldsymbol{X}_j \in \mathbb{R}^{N \times k}$. We then optimize the following

$$Q^* = \underset{Q \in \mathcal{O}}{\operatorname{argmin}} ||X_i - X_j Q|| \tag{8}$$

where \mathcal{O} is the set of $k \times k$ orthogonal matrices. This optimization problem can be solved by doing a singular value decomposition on $X_i^{\top} X_j$ where $Q^* = UV^{\top}$. We calculate distance between the two spaces with

$$d(\mathbf{X}_i, \mathbf{X}_j) = \arccos\left(\frac{1}{n} \sum_{i=1}^n \sigma_i\right)$$
 (9)

where σ are the singular values of $X_i^{\top} X_j$. Optimizing equation 8 and calculating distance from equation 9 gives rise to the orthogonal Procrustes distance metric mentioned in the main text. If we instead whiten each representation as

$$\boldsymbol{X}_{i}^{\theta} = \boldsymbol{X}_{i} (\boldsymbol{X}_{i}^{\top} \boldsymbol{X}_{i})^{1/2}, \quad \boldsymbol{X}_{j}^{\theta} = \boldsymbol{X}_{j} (\boldsymbol{X}_{j}^{\top} \boldsymbol{X}_{j})^{1/2}$$

$$(10)$$

and solve

$$Q^* = \underset{Q \in \mathcal{O}}{\operatorname{argmin}} ||X_i^{\theta} - X_j^{\theta} Q|| \tag{11}$$

the resulting distances (again calculated via equation 9) correspond to the CCA distance metrics reported in the main text. These metrics are calculated using the package 'netrep' from Williams et al. (2022).

D LORAX MODEL DETAILS

In this section we describe, in depth, model details of our LORAX model. LORAX uses the Huggingface and PEFT libraries to adapt both the chemical and protein foundation models. All results shown in this work (unless otherwise specified) only adapt the chemical foundation model. The chemical foundation model adapter was run with r=8 rank adapter matrices on the query, key and value matrices with $\alpha=8$. We do not train the bias parameters. We have a dropout rate of 0.1 in the adapter matrices during training. We do not use rslora (Kalajdzievski, 2023) as we are interested in starting from the original chemical foundation model representation to see how the representation changes over training. The cross attention block of the model is composed of two multi-headed attention layers. Both of these layers use 8 heads by default. The SMILES multiheaded attention layer takes the last hidden state of the LoRA adapted chemical foundation model (smi_{lora}) as the query, and the last hidden state of the protein foundation model (prot_{ori}) as the key and value matrices. The protein multi-headed attention layer does the opposite taking prot_{ori}

⁷We use 42 here because the largest representation vector, 'ScaffoldKeyCalculator', is 42 length.

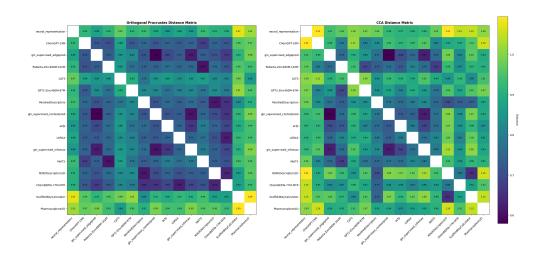


Figure 7: Full distance matrices using Procrustes (right) and CCA (left) distance metrics corresponding to the distance matrices shown in Figure 5.

as the query matrix and smi_{lora} as the key and value matrices. In both of these multi-headed attention layers there is a dropout rate of 0.1 during training. After the cross attention block, there is a residual connection where the output of the SMILES (protein) multi-headed attention layer is added to smi_{lora} (prot_{ori}). These representations are then passed through a layer norm. Both the protein and odorant representations are then mean pooled across their sequence dimensions and concatenated together. This concatenation is the <cls> token referenced in the main text. The <cls> token is then put through an MLP with default dimension 512 to predict affinity between the odorant and receptor. During training, the model is run with a batch size of 12 for 50 epochs for the Carey, Hallem, and M2OR (EC50) datasets and batch size 21 for 15 epochs for the M2OR dataset. By default the training is conducted at a learning rate of 1e-5 unless otherwise specified. We use Adam as our optimizer.

Throughout training, the loss on the validation set is monitored, and the weights of multi-modal transformer model with the lowest validation loss are saved. Then, <cls> tokens are generated for all data using that multi-modal transformer model. We then train 3 separate XGBoost models. One model takes, as input, the <cls> token; one takes, as input, the original, non LoRA-optimized, chemical foundation model concatenated to the protein foundation model representation; and the last model takes, as input, the <cls> token concatenated with the non LoRA-optimized chemical foundation model representation and the protein foundation model representation. Each XGBoost model has a set of hyperparameters that we optimize on the validation set using hyperopt. The hyperparameters, as well as the range of their values, are highlighted in Table 5. The parameters are optimized using a random search over a default of 500 iterations. After optimizing hyperparameters for all three XGBoost models, they weighted according to their accuracy on the validation dataset. Finally, all XGBoost models are used to predict the test set values using the confidence weights. Code for this model will be following publication.

E ALL DATASETS AND METRICS STUDIED

This section presents all metrics and datasets omitted from Section 3. Tables 6 and 7 provide affinity prediction results and additional ProSmith metrics, respectively, for the Hallem dataset. Table 8 reports further metrics for the ProSmith model trained on the Carey dataset. Finally, Tables 9 and 10 show affinity prediction performance and other relevant metrics for the ProSmith model on the M2OR (EC50) dataset.

Table 5: Hyperparameters and their search ranges for XGBoost optimization.

Hyperparameter	Search Range
learning_rate	0.01 - 0.5
max_depth	6, 7, 8, 9, 10, 11, 12, 13, 14
reg_lambda	0 - 5
reg_alpha	0 - 5
max_delta_step	0 - 5
min_child_weight	0.1 - 15
num_rounds	30 - 1000
weight	0.01 – 0.99

Table 6: Affinity prediction performance across odorant representations for the Hallem dataset. Reported as mean $R^2 \pm \text{standard}$ deviation across 5 fold cross validation. Best performing representations are highlighted in bold.

ions are mightighted in bold.			
Odorant representation	MO	MP	ProSmith
CATS	-0.129 ± 1.819	0.367 ± 0.033	0.561 ± 0.032
MordredDescriptors	0.137 ± 0.341	0.382 ± 0.042	0.579 ± 0.036
Pharmacophore2D	0.192 ± 0.247	0.389 ± 0.037	0.530 ± 0.059
RDKitDescriptors2D	0.044 ± 0.201	0.302 ± 0.043	$\textbf{0.601} \pm \textbf{0.057}$
ScaffoldKeyCalculator	-0.014 ± 0.219	0.272 ± 0.040	0.501 ± 0.044
ECFP	0.182 ± 0.282	0.389 ± 0.036	0.558 ± 0.036
Roberta-Zinc480M-102M	0.183 ± 0.252	$\textbf{0.398} \pm \textbf{0.039}$	0.564 ± 0.041
GPT2-Zinc480M-87M	0.095 ± 0.205	0.367 ± 0.041	0.521 ± 0.067
ChemGPT-19M	0.118 ± 0.240	0.359 ± 0.035	0.559 ± 0.051
MolT5	0.156 ± 0.249	0.389 ± 0.035	0.555 ± 0.046
ChemBERTa-77M-MTR	0.198 ± 0.349	0.373 ± 0.039	0.574 ± 0.040
gin_supervised_infomax	0.198 ± 0.287	0.396 ± 0.041	0.530 ± 0.044
gin_supervised_edgepred	0.177 ± 0.252	0.390 ± 0.040	0.550 ± 0.031
gin_supervised_contextpred	$\textbf{0.203} \pm \textbf{0.411}$	0.396 ± 0.043	0.547 ± 0.061
random	-0.069 ± 0.119	0.382 ± 0.039	0.439 ± 0.073

Table 7: Other metrics for the ProSmith model trained on the Hallem dataset. Error indicates standard deviation across 5 fold cross validation. Best performing representations are highlighted in bold.

Odorant representation	MSE	rm2	CI
CATS	0.462 ± 0.042	0.548 ± 0.034	0.818 ± 0.006
MordredDescriptors	0.444 ± 0.052	0.571 ± 0.043	0.824 ± 0.009
Pharmacophore2D	0.494 ± 0.068	0.526 ± 0.052	0.806 ± 0.007
RDKitDescriptors2D	$\textbf{0.421} \!\pm \textbf{0.075}$	$\textbf{0.587} \!\pm \textbf{0.061}$	$\textbf{0.830} \!\pm \textbf{0.007}$
ScaffoldKeyCalculator	0.524 ± 0.050	0.491 ± 0.044	0.813 ± 0.011
ECFP	0.466 ± 0.049	0.551 ± 0.034	$0.821 {\pm} 0.004$
Roberta-Zinc480M-102M	0.459 ± 0.053	0.554 ± 0.046	0.824 ± 0.005
GPT2-Zinc480M-87M	0.505 ± 0.081	0.511 ± 0.062	0.810 ± 0.011
ChemGPT-19M	0.465 ± 0.068	0.551 ± 0.058	0.821 ± 0.008
MolT5	0.470 ± 0.066	0.546 ± 0.052	0.818 ± 0.010
ChemBERTa-77M-MTR	0.449 ± 0.055	0.563 ± 0.045	$0.821 {\pm} 0.004$
gin_supervised_infomax	0.494 ± 0.052	$0.522 {\pm} 0.045$	0.820 ± 0.007
gin_supervised_edgepred	0.474 ± 0.045	0.541 ± 0.031	0.826 ± 0.003
gin_supervised_contextpred	0.478 ± 0.074	0.533 ± 0.060	0.819 ± 0.007
random	0.590 ± 0.086	0.434 ± 0.071	0.793 ± 0.012

Table 8: Other metrics for the ProSmith model trained on the Carey dataset. Error indicates standard deviation across 5 fold cross validation. Best performing representations are highlighted in bold.

ii across 5 roid cross varidation	. Dest performing	s representations	are mighingined m
Odorant representation	MSE	rm2	CI
CATS	0.378 ± 0.036	0.631 ± 0.016	0.815 ± 0.008
MordredDescriptors	0.293 ± 0.037	0.717 ± 0.027	$\textbf{0.822} \pm \textbf{0.005}$
Pharmacophore2D	0.471 ± 0.038	0.543 ± 0.011	0.801 ± 0.008
RDKitDescriptors2D	$\textbf{0.292} \pm \textbf{0.039}$	$\textbf{0.718} \pm \textbf{0.028}$	0.821 ± 0.009
ScaffoldKeyCalculator	0.422 ± 0.031	0.590 ± 0.023	0.796 ± 0.012
ECFP	0.351 ± 0.046	0.658 ± 0.030	0.805 ± 0.010
Roberta-Zinc480M-102M	0.341 ± 0.037	0.666 ± 0.031	0.817 ± 0.011
GPT2-Zinc480M-87M	0.347 ± 0.045	0.661 ± 0.030	0.813 ± 0.011
ChemGPT-19M	0.328 ± 0.034	0.680 ± 0.024	0.816 ± 0.014
MolT5	0.338 ± 0.038	0.669 ± 0.021	0.812 ± 0.005
ChemBERTa-77M-MTR	0.309 ± 0.032	0.700 ± 0.020	0.822 ± 0.010
gin_supervised_infomax	0.343 ± 0.034	0.667 ± 0.018	0.812 ± 0.013
gin_supervised_edgepred	0.326 ± 0.035	0.681 ± 0.028	0.820 ± 0.011
gin_supervised_contextpred	0.315 ± 0.047	0.696 ± 0.035	$\textbf{0.822} \pm \textbf{0.008}$
random	0.431 ± 0.047	0.582 ± 0.038	0.796 ± 0.011

Table 9: Affinity prediction performance across odorant representations for the M2OR (EC50) dataset. Reported as mean Matthew's correlation coefficient (MCC) \pm standard deviation. Best performing representations are highlighted in bold.

Odorant representation	MO	MP	ProSmith
CATS	0.090 ± 0.262	0.562 ± 0.030	0.674 ± 0.029
MordredDescriptors	0.113 ± 0.291	0.568 ± 0.027	$\textbf{0.709} \pm \textbf{0.018}$
Pharmacophore2D	0.097 ± 0.270	0.570 ± 0.030	0.656 ± 0.032
RDKitDescriptors2D	0.066 ± 0.241	0.515 ± 0.038	0.694 ± 0.024
ScaffoldKeyCalculator	0.053 ± 0.198	0.583 ± 0.025	0.646 ± 0.032
ECFP	0.096 ± 0.293	0.581 ± 0.030	0.663 ± 0.016
Roberta-Zinc480M-102M	0.131 ± 0.310	0.568 ± 0.029	0.697 ± 0.027
GPT2-Zinc480M-87M	0.109 ± 0.286	0.563 ± 0.026	0.678 ± 0.023
ChemGPT-19M	0.103 ± 0.295	0.574 ± 0.035	0.683 ± 0.017
MolT5	0.113 ± 0.294	0.589 ± 0.035	0.685 ± 0.030
ChemBERTa-77M-MTR	$\textbf{0.172} \pm \textbf{0.351}$	0.591 ± 0.030	0.696 ± 0.033
gin_supervised_infomax	0.109 ± 0.300	$\textbf{0.595} \pm \textbf{0.030}$	0.684 ± 0.030
gin_supervised_edgepred	0.118 ± 0.312	0.593 ± 0.021	0.692 ± 0.040
gin_supervised_contextpred	0.103 ± 0.266	0.589 ± 0.031	0.695 ± 0.026
random	0.004 ± 0.120	0.550 ± 0.029	0.640 ± 0.027

Table 10: Other metrics for the ProSmith model trained on the M2OR (EC50) dataset. Error indicates standard deviation across 5 fold cross validation. Best performing representations are highlighted in bold.

Odorant representation	Accuracy	AUROC
CATS	0.877 ± 0.013	0.917 ± 0.006
MordredDescriptors	$\textbf{0.892} \pm \textbf{0.006}$	$\textbf{0.933} \pm \textbf{0.007}$
Pharmacophore2D	0.872 ± 0.012	0.913 ± 0.007
RDKitDescriptors2D	0.886 ± 0.008	0.931 ± 0.007
ScaffoldKeyCalculator	0.870 ± 0.012	0.910 ± 0.005
ECFP	0.873 ± 0.005	0.923 ± 0.009
Roberta-Zinc480M-102M	0.889 ± 0.010	0.927 ± 0.005
GPT2-Zinc480M-87M	0.882 ± 0.008	0.923 ± 0.005
ChemGPT-19M	0.883 ± 0.006	0.923 ± 0.006
MolT5	0.884 ± 0.010	0.925 ± 0.007
ChemBERTa-77M-MTR	0.888 ± 0.012	0.927 ± 0.003
gin_supervised_infomax	0.884 ± 0.010	0.925 ± 0.006
gin_supervised_edgepred	0.887 ± 0.013	0.928 ± 0.008
gin_supervised_contextpred	0.887 ± 0.010	0.925 ± 0.004
random	0.870 ± 0.010	0.902 ± 0.007