

FAST AND INTERPRETABLE PROTEIN SUBSTRUCTURE ALIGNMENT VIA OPTIMAL TRANSPORT

Zhiyu Wang^{1,2} * Bingxin Zhou¹ * † Jing Wang¹ Yang Tan¹ Weishu Zhao¹
 Pietro Liò² Liang Hong¹ †

¹ Shanghai Jiao Tong University. ² University of Cambridge.

ABSTRACT

Proteins are essential biological macromolecules that execute life functions. Local structural motifs, such as active sites, are the most critical components for linking structure to function and are key to understanding protein evolution and enabling protein engineering. Existing computational methods struggle to identify and compare these local structures, which leaves a significant gap in understanding protein structures and harnessing their functions. This study presents PLASMA, a deep-learning-based framework for efficient and interpretable residue-level local structural alignment. We reformulate the problem as a regularized optimal transport task and leverage differentiable Sinkhorn iterations. For a pair of input protein structures, PLASMA outputs a clear alignment matrix with an interpretable overall similarity score. Through extensive quantitative evaluations and three biological case studies, we demonstrate that PLASMA achieves accurate, lightweight, and interpretable residue-level alignment. Additionally, we introduce PLASMA-PF, a training-free variant that provides a practical alternative when training data are unavailable. Our method addresses a critical gap in protein structure analysis tools and offers new opportunities for functional annotation, evolutionary studies, and structure-based drug design. Reproducibility is ensured via our official implementation at <https://github.com/ZW471/PLASMA-Protein-Local-Alignment.git>.

1 INTRODUCTION

Proteins are essential macromolecules responsible for life functions, from catalysis and signal transduction to structural support and transport. Local structural motifs (*e.g.*, catalytic residues, binding pockets, metal-binding sites) are critical for understanding mechanisms, designing therapeutics, and guiding protein engineering (Mills et al., 2018). Structural conservation is three to ten times stronger than sequence conservation across evolution, suggesting that local structural comparison can reveal functional relationships invisible to sequence-based methods (Hvidsten et al., 2009).

Despite their importance, existing computational methods primarily emphasize global structure comparison or sequence alignment. The inability to detect local structural motifs, *i.e.*, compact three-dimensional residue arrangements that often concentrate around catalytic pockets or interaction sites, prevents researchers from understanding protein evolution, predicting functions of uncharacterized proteins, and rationally designing proteins with desired properties. While large-scale resources like AFDB (Jumper et al., 2021; Varadi et al., 2022) open a unique opportunity to uncover conserved motifs across the protein universe, active sites often comprise spatially proximate residues that may be widely separated in sequence or embedded within different overall fold architectures (Liu et al., 2018). Addressing this gap is key to advancing our understanding of protein function and evolution.

The development of robust local structure alignment methods specifically targeting local structural motifs is not merely a technical challenge but a fundamental requirement for advancing multiple areas of biological research and application. Existing methods for protein substructure alignment

*Equal contribution first authors.

†Corresponding authors (bingxin.zhou@sjtu.edu.cn; hongl3liang@sjtu.edu.cn).

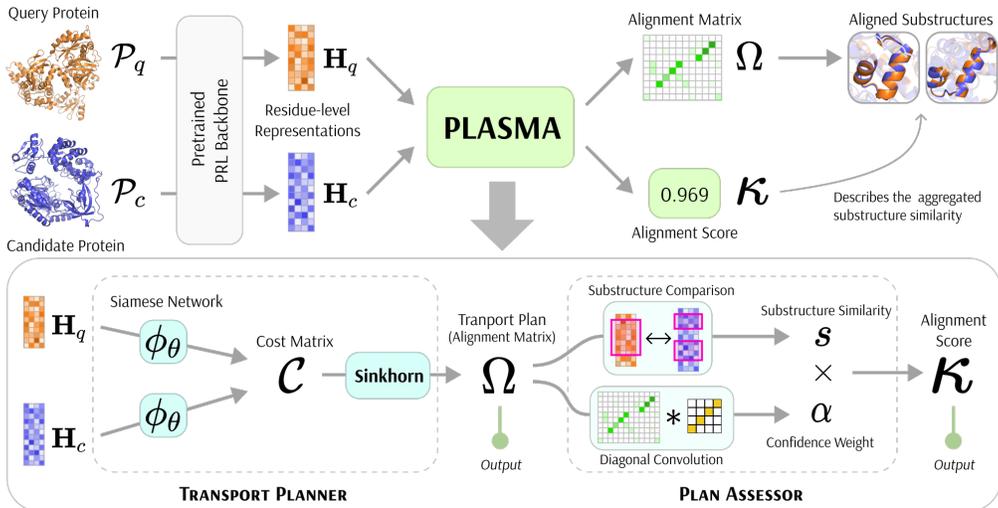


Figure 1: PLASMA Overview. PLASMA converts residue-level protein embeddings into substructure alignments using optimal transport. A *Transport Planner* learns cost matrices with Sinkhorn iterations, and a *Plan Assessor* produces similarity scores. The framework provides alignment matrices and quantitative scores without requiring model-specific designs.

can be broadly divided into three categories. The first relies on template-based searches, where predefined motifs are used to identify similar substructures (Bittrich et al., 2020; Kim et al., 2025). These approaches are effective for detecting well-characterized patterns but cannot uncover novel similarities, making them **unsuitable for pairing novel structural motifs**. The second category estimates substructure similarity based on the global similarity of entire protein structures. Several studies leverage structural superposition (Zhang, 2005) or structural tokenization (Holm, 2020) to produce residue-level matches with sequence alignment, but they are **computationally demanding and difficult to scale to large datasets**. More recent embedding-based methods (Hamamsy et al., 2024) are enabled by advances in protein representation learning, which make alignment faster and competitive for whole-protein comparison. However, they compress residue-level information into coarse embeddings, which causes **problems in producing interpretable local alignments**. The third category directly addresses substructure alignment by constructing pairwise similarity matrices and using dynamic programming to find matching regions. This approach captures local similarities more accurately than global methods and produces scores that reflect substructure correspondence (Kaminski et al., 2023; Liu et al., 2024; Pantolini et al., 2024). However, the results can be influenced by overall structural patterns, and **alignment matrices have limited interpretability** since they are optimized for algorithmic performance rather than clarity. Additionally, these methods are typically untrainable and cannot adapt to specific alignment tasks or incorporate domain knowledge, limiting their ability to improve through experience or be customized for particular biological contexts.

The challenges above point to the need for a novel protein substructure alignment method that combines accuracy, efficiency, and clarity. To this end, we explore optimal transport (OT), a mathematical framework proven effective in alignment problems (Mena et al., 2018). In particular, the differentiable Sinkhorn algorithm (Sinkhorn & Knopp, 1967; Cuturi, 2013) has shown strong ability to uncover meaningful correspondences in 3D shape analysis (Eisenberger et al., 2020) and subgraph matching (Ramachandran et al., 2024). Notably, these OT-based alignment methods assume strict one-to-one correspondences between all residues or that one set of residues is fully contained within the other. These constraints do not hold for protein substructure alignment, as functionally similar regions may only partially overlap and vary in length across proteins.

To address the aforementioned limitations, we reframe protein substructure alignment as an OT problem and introduce PLASMA (**Pluggable Local Alignment via Sinkhorn MAtrix**). As illustrated in Figure 1, PLASMA operates on residue-level embeddings from a pre-trained protein representation model and identifies the residue-level alignment between protein pairs. The *Transport Planner* computes the pairwise matching using a learnable cost matrix and differentiable Sinkhorn iterations (Section 3), and the *Plan Assessor* then summarizes the resulting alignment matrix into a single similarity score reflecting the overall similarity of the matched substructures (Section 4). PLASMA

functions as a lightweight, plug-and-play module for protein representation models. It is capable of efficiently aligning partial and variable-length matches between local structural regions.

Our work addresses these limitations through three contributions. First, we introduce a formulation of residue-level local structural alignment based on regularized optimal transport with a learnable geometric cost, which provides a principled and flexible way to define correspondence and enables efficient, fully parallel implementation. Second, this formulation enables clear and interpretable residue-residue correspondences and naturally supports partial, variable-length, and non-sequential motif alignments, resolving the difficulty of obtaining reliable local alignments. Third, PLASMA produces a normalized and interpretable similarity score through its OT-based objective, overcoming the limitations of existing approaches whose alignment matrices or similarity measures lack a consistent probabilistic meaning. Our experiments show strong generalization to low-homology structures, and the case studies demonstrate the biological interpretability and practical utility of the resulting alignments.

2 PROTEIN SUBSTRUCTURE ALIGNMENT VIA OPTIMAL TRANSPORT

Problem Formulation Consider a query protein $\mathcal{P}_q = \{r_{q,1}, \dots, r_{q,N}\}$ of N residues and a candidate protein $\mathcal{P}_c = \{r_{c,1}, \dots, r_{c,M}\}$ of M residues. Suppose the two proteins contain local structural motifs $\mathcal{F}_q = \{f_{q,1}, \dots, f_{q,n}\} \subseteq \mathcal{P}_q$ and $\mathcal{F}_c = \{f_{c,1}, \dots, f_{c,m}\} \subseteq \mathcal{P}_c$, where $n \leq N$ and $m \leq M$. The objective of protein substructure alignment is: (1) to identify the corresponding fragments \mathcal{F}_q and \mathcal{F}_c within \mathcal{P}_q and \mathcal{P}_c , and (2) to score their level of similarity.

The task is challenging for several reasons: the overall structures of \mathcal{P}_q and \mathcal{P}_c may differ substantially, the fragments \mathcal{F}_q and \mathcal{F}_c may vary in sequence length or composition, and alignments require remaining meaningful in a biological context. In particular, biologically relevant alignments should capture functional similarities, such as common enzymatic activities or conserved structural roles.

Optimal Transport Reformulation To address the protein substructure alignment problem, we reformulate it as an entropy-regularized OT problem between the residues of two proteins \mathcal{P}_q and \mathcal{P}_c . Each protein is represented as a set of residue embeddings that capture local biochemical and structural context. The OT solver then computes a soft alignment matrix $\Omega \in \mathbb{R}^{N \times M}$ by assigning weights between residues so as to minimize the overall transport cost \mathcal{C} . This formulation bypasses explicit fragment enumeration, naturally accommodates partial and variable-length matches, and produces interpretable alignment matrices that highlight the underlying substructures (Appendix A).

Overview of PLASMA We implement entropy-regularized OT and propose **PLASMA**, a module that transforms $\mathbf{H}_q \in \mathbb{R}^{N \times d}$ and $\mathbf{H}_c \in \mathbb{R}^{M \times d}$, residue-level d -dimensional hidden representations of \mathcal{P}_q and \mathcal{P}_c (e.g., from pre-trained protein language models), into a soft alignment matrix $\Omega \in \mathbb{R}^{N \times M}$ and a similarity score $\kappa \in [0, 1]$. In our experiments, we instantiate H_q and H_c with seven diverse protein representation backbones (Section 6), and observe consistent alignment behavior across them, indicating that PLASMA is not tied to a particular choice of encoder. Formally,

$$(\Omega, \kappa) = \text{PLASMA}(\mathbf{H}_q, \mathbf{H}_c). \quad (1)$$

PLASMA consists of two complementary components (visualized in Figure 1, with details introduced in the next two sections). The first component, the *Transport Planner*, produces Ω to highlight local correspondences between \mathcal{P}_q and \mathcal{P}_c . The second component, the *Plan Assessor*, summarizes this alignment matrix into a similarity score $\kappa \in [0, 1]$, providing a quantitative measure of alignment quality. The framework achieves a computational complexity of $O(N^2)$ (Appendix B).

3 TRANSPORT PLANNER

The Transport Planner module handles the core OT computation. It defines cost functions between residue pairs and solves the regularized OT problem to produce an Ω that captures residue-level matching between query and candidate proteins ($\mathcal{P}_q, \mathcal{P}_c$).

Cost Matrix We formulate a learnable cost matrix with a siamese network architecture to capture complex residue-level similarities. This approach enables PLASMA to learn task-specific represen-

tations that optimize alignment quality through end-to-end training. The cost from $r_{q,i}$ to $r_{c,j}$ is denoted by \mathcal{C}_{ij} in the learnable cost matrix, defined as

$$\mathcal{C}_{ij} = \left\| \left[\phi_{\theta}(\text{LN}(\mathbf{h}_{q,i})) - \phi_{\theta}(\text{LN}(\mathbf{h}_{c,j})) \right]_+ \right\|_1. \quad (2)$$

Here $\mathbf{h}_{q,i}$ and $\mathbf{h}_{c,j}$ denote the hidden representations of residues $r_{q,i}$ and $r_{c,j}$, respectively. The operator $[\cdot]_+$ applies a hinge non-linearity, shown to outperform dot-product similarity in subgraph matching tasks (Raj et al., 2025). The layer normalization $\text{LN}(\cdot)$ facilitates robust optimization dynamics with numerical stability and scale-invariant representations. The siamese network $\phi_{\theta}(\cdot)$ processes query and candidate residues using a twin architecture with shared parameters θ .

Learnable and Parameter-Free Implementations The siamese network architecture can be chosen flexibly, ranging from Transformer-based (Hamamsy et al., 2024) models to graph neural networks (Jamasp et al., 2024), depending on the inductive bias of the input data and the computational budget. Here we also provide a simple implementation using fully connected layers:

$$\phi_{\theta}(\mathbf{h}) = \text{ReLU}(\mathbf{h} \cdot \mathbf{W}_1) \cdot \mathbf{W}_2, \quad (3)$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times d'}$ and $\mathbf{W}_2 \in \mathbb{R}^{d' \times d'}$ are learnable transformation matrices with d' hidden dimension. For simplicity, we omit the subscript of \mathbf{H} as the siamese network applies the same set of parameters to both the query and candidate proteins. This lightweight design serves as an effective default while allowing more sophisticated architectures to be substituted without modifying the overall PLASMA architecture. In addition, for scenarios with a lack of labeled data, we introduce a parameter-free variant, **PLASMA-PF**, which bypasses the siamese network and operates directly on residue embeddings. The cost used in the OT objective follows (2) with no architectural components removed other than the encoder. PLASMA-PF preserves the fundamental alignment functionality and offers a fast baseline for substructure similarity evaluation. Notably, the learnable version remains preferable for improved stability and extrapolation (See Section 6.3 and Figure 4).

Sinkhorn Alignment Matrix Based on the cost matrix \mathcal{C} defined in (2), we formulate the corresponding OT problem (Appendix A) and solve it using the Sinkhorn algorithm (Cuturi, 2013). The algorithm approximates the OT plan by iteratively scaling the matrix to satisfy the marginal constraints with row and column normalizations, ensuring that the total alignment weights of each residue are properly distributed across residues of the other protein:

$$\Omega_{ij}^{(t+1)} = \frac{\mathbf{Z}_{ij}^{(t)}}{\sum_{v=1}^M \mathbf{Z}_{iv}^{(t)}}, \quad \text{where } \mathbf{Z}_{ij}^{(t)} = \frac{\Omega_{ij}^{(t)}}{\sum_{u=1}^N \Omega_{uj}^{(t)}}. \quad (4)$$

The iteration is initialized as $\Omega^{(0)} = \exp(-\mathcal{C}/\tau)$, where τ is a temperature parameter controlling the alignment sharpness (Appendix J). The optimal $\Omega^* = \Omega^{(T)}$ after T iterations serves as the Sinkhorn alignment matrix. For simplicity, we denote it as Ω in the subsequent discussions.

The original Sinkhorn algorithm converges to a fully doubly stochastic matrix, forcing each query residue to distribute across all candidate residues (and vice versa). This strict matching is often biologically meaningless, as most residues lack relevant counterparts. PLASMA achieves implicit partial alignments via two mechanisms. First, *early termination* preserves sparsity by limiting Sinkhorn iterations, letting poorly matching residues retain low weights. Second, the *temperature parameter* τ controls alignment mass, with lower values producing sparser, focused alignments. Together, these mechanisms emphasize biologically relevant correspondences while avoiding forced matches, without hard constraints on the transport budget (Caffarelli & McCann, 2010; Figalli, 2010). Representative alignment matrices demonstrating these patterns are shown in Appendix I.

4 PLAN ASSESSOR

The Plan Assessor receives the alignment matrix Ω from the Transport Planner and transforms it into an interpretable single similarity score $\kappa \in [0, 1]$ that quantifies the existence and degree of similarity of the aligned substructures. This is computed by first calculating a substructure similarity score for the aligned regions, then adjusting it with a confidence weight to correct potential bias.

Substructure Similarity We calculate the alignment score on *matched substructure*. With a threshold ρ , a residue pair $r_{q,i} \in \mathcal{P}_q$ and $r_{c,j} \in \mathcal{P}_c$ is treated as matched if $\Omega_{ij} > \rho$. The matched residues then form two sets, $\mathcal{R}_q = \{r_{q,i} \mid \forall j, \Omega_{ij} > \rho\}$ and $\mathcal{R}_c = \{r_{c,j} \mid \forall i, \Omega_{ij} > \rho\}$. A matched substructure is a subset of these residues. The representation of the matched substructure can be approximated by summing the embeddings of residues from \mathcal{R}_q and \mathcal{R}_c . Therefore, the *substructure similarity score* $s \in [-1, 1]$ is defined as the cosine similarity between the summed representations:

$$s = \frac{\sum_{i \in \mathcal{R}_q} \mathbf{h}_{q,i} \cdot \sum_{j \in \mathcal{R}_c} \mathbf{h}_{c,j}}{\|\sum_{i \in \mathcal{R}_q} \mathbf{h}_{q,i}\| \cdot \|\sum_{j \in \mathcal{R}_c} \mathbf{h}_{c,j}\|}. \quad (5)$$

This substructure similarity score is effective when a sufficient number of residues are matched between the two proteins. However, it becomes less reliable when only a few residues are aligned or when the matched residues are dispersed along the sequence rather than forming a continuous region. In such cases, the score reduces to a residue-level similarity measure, which may appear deceptively high even though the aligned residues do not cluster into a structurally interpretable substructure. We thus introduce a *confidence weight* to adjust the initial similarity score.

Alignment Score with Confidence Weight Correction The *confidence weight* $\alpha \in [0, 1]$ is derived from Ω using a 2D convolution with an identity kernel $K = \mathbb{I}_k \in \mathbb{R}^{k \times k}$ of size k :

$$\alpha_{ij} = \sum_{u=0}^{k-1} \sum_{v=0}^{k-1} \Omega_{i+u, j+v} \cdot K_{uv} = \sum_{u=0}^{k-1} \Omega_{i+u, j+u}. \quad (6)$$

This convolution operation highlights continuous diagonal segments in Ω and emphasizes core regions where consecutive residues in the query align with consecutive residues in the candidate. A max-pooling layer then produces a scalar confidence weight $\alpha = \max_{i,j} \alpha_{ij}$, summarizing the strongest local alignment signal used to weight the similarity score and obtain the final *alignment score* $\kappa = \alpha \cdot s_+ \in [0, 1]$. Here s_+ is the non-negative substructure similarity score. This formulation provides an intuitive and interpretable measure: $\kappa = 0$ indicates no residue matches and $\kappa = 1$ represents perfect substructure alignment. We follow the convention of established alignment methods (e.g., TM-align (Zhang, 2005)) and exclude negative similarity values, since matched substructures with opposite orientations in the representation space lack meaningful biological interpretation. See Appendix I for visual examples of alignment matrices with different similarity scores.

5 MODEL OPTIMIZATION

PLASMA is trained with two complementary objectives: predicting the presence of aligned substructures via the alignment score κ and recovering precise residue-level matches via the alignment matrix Ω . Training data consists of protein pairs $(\mathcal{P}_q, \mathcal{P}_c)$, where a subset of pairs contains matched substructures with shared functions. For each input protein pair, two mask vectors $\mathcal{M}_q \in \{0, 1\}^N$ and $\mathcal{M}_c \in \{0, 1\}^M$ are respectively defined to indicate the position of target substructures \mathcal{F}_q and \mathcal{F}_c , where 1 marks the residues that belong to the substructure of interest.

Alignment Score Optimization The alignment score κ serves as the model’s prediction on whether the input protein pair contains aligned substructures. We define the ground truth $y = 1$ if the pair contains matched substructures and $y = 0$ otherwise. The prediction is optimized by $\mathcal{L}_{\text{BCE}} = -y \log(\sigma(\kappa)) - (1 - y) \log(1 - \sigma(\kappa))$, where $\sigma(\cdot)$ is the sigmoid function.

Alignment Matrix Optimization Unlike the alignment score, optimizing the alignment matrix is challenging because unlabeled residues may correspond to valid but unannotated matches. Treating these residues as negative examples would impose inappropriate penalties on the model. To address this, we propose the *Label Match Loss* (LML) to focus exclusively on the labeled substructures. Specifically, when $\|\mathcal{M}_c\|_1 > 0$ and $\|\mathcal{M}_q\|_1 > 0$, the LML for protein pairs is defined as

$$\mathcal{L}_{\text{LML}} = \|[\mathcal{M}_c - \Omega^\top \mathcal{M}_q]_+\|_1 / \|\mathcal{M}_c\|_1, \quad (7)$$

where $[\cdot]_+$ retains only non-negative elements, and $\|\cdot\|_1$ denotes the ℓ_1 norm. This loss evaluates how well the constructed alignment matrix Ω aligns the labeled substructures $(\mathcal{F}_q, \mathcal{F}_c)$ in $(\mathcal{P}_q, \mathcal{P}_c)$. For

each residue $r_j \in \mathcal{P}_c$, $(\Omega^\top \mathcal{M}_q)_j$ gives the alignment weight with respect to labeled residues in \mathcal{P}_q . The non-negative contributions by $[\mathcal{M}_c - \Omega^\top \mathcal{M}_q]_+$ are normalized by $\|\mathcal{M}_c\|_1$ across all labeled residues. When no labeled substructures exist, $\mathcal{L}_{\text{LML}} = 0$, which allows the model to focus on known substructures without penalizing unlabeled but potentially valid matches. This loss provides an optional bias toward annotated local structural motifs when such labels exist. These regions are typically small and structurally meaningful (*e.g.*, catalytic or binding motifs), and emphasizing them helps the model avoid being dominated by background alignments.

The final $\mathcal{L} = \mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{LML}}$ jointly detects substructure existence by κ and localizes known substructures by Ω , while staying robust to missing or incomplete labels in the training data.

6 EMPIRICAL ANALYSIS

We conduct extensive quantitative and qualitative evaluations to comprehensively assess the validity and advancement of PLASMA in local structural motif alignment tasks. All experiments are programmed with PyTorch v2.5.1 and run on NVIDIA RTX 4090 32 GB GPU.

6.1 EXPERIMENTAL SETUP

Prediction Tasks and Benchmark Datasets Our experiments are based on a residue-level functional alignment benchmark, **VenusX** (Tan et al., 2025a). We consider three common classes of functional substructures: activation sites, binding sites, and motifs. Across all test sets, the sequence identity between training and test proteins is kept below 50%. For quantitative evaluation, we design two levels of difficulty: (i) interpolation (`test_inter`), where the test set contains proteins from InterPro families already present in training; and (ii) extrapolation (`test_extra`), where the test set only includes novel substructures from unseen families. Further details are in Appendix C.1.

Baseline Methods We compare PLASMA with popular baselines in protein structure alignment, including structure-based methods (FOLDSEEK (Van Kempen et al., 2024), TM-ALIGN (Zhang, 2005), and TM-VEC (Hamamsy et al., 2024)) and embedding-based methods (EBA (Pantolini et al., 2024) and COSINESIM, a cosine similarity over protein embeddings). For all embedding-based methods, we implement seven popular pre-trained models to extract residue-level sequence and structure representations, including PROTT5 (Elnaggar et al., 2021), PROSTT5 (Heinzinger et al., 2024), ANKH (Elnaggar et al., 2023), ESM2 (Lin et al., 2023), PROTBERT (Brandes et al., 2022), TM-VEC (Hamamsy et al., 2024), and PROTSSN (Tan et al., 2025b). All baselines use the authors’ official code and checkpoints (see Appendices D for details).

Evaluation Metrics To assess the ability to detect the existence of local structural motifs, we use standard binary classification metrics, including ROC-AUC, PR-AUC, and F1-Max. Additionally, to evaluate alignment quality, we introduce the Label Match Score (LMS) by (7) with $\text{LMS} = 1 - \text{LML}$ to measure correspondence between predicted alignments and annotated functional regions.

6.2 QUANTITATIVE PERFORMANCE EVALUATION

Table 1 reports performance on `test_extra`, which contains functional substructures from protein families not seen during training. This setting evaluates the generalizability of the alignment framework, which is essential in practice because new functional substructures are continuously discovered. Full results on seven backbone models are provided in Appendix F, and all hyperparameter and dataset details are summarized in Appendix C.2. Corresponding interpolation results on `test_inter` are reported in Appendix E.

Across all three substructure detection tasks and all evaluation metrics, PLASMA achieves consistent top performance, highlighting its robustness in capturing fundamental local structural similarities for novel substructures beyond the training distribution. PLASMA-PF also performs strongly and remains competitive without task-specific training. However, unlike in the interpolation setting, PLASMA-PF does not surpass the learnable PLASMA variant on `test_extra`; this emphasizes the value of supervised examples in improving alignment accuracy for entirely new functional substructures. In contrast, baseline methods show large performance variation across backbone models.

Table 1: Model performance on `test_extra` (mean \pm std over three independent seeds). Colors indicate relative performance versus TM-ALIGN.

Metrics	Methods	Motif			Binding Site			Active Site		
		ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN
ROC-AUC	PLASMA	.98 \pm .008	.97 \pm .013	.96 \pm .016	.99 \pm .008	.98 \pm .013	.98 \pm .014	.98 \pm .012	.98 \pm .010	.97 \pm .011
	PLASMA-PF	.98 \pm .009	.93 \pm .004	.90 \pm .005	.99 \pm .006	.92 \pm .052	.96 \pm .012	.97 \pm .015	.96 \pm .006	.97 \pm .008
	EBA	.90 \pm .033	.92 \pm .021	.32 \pm .043	.99 \pm .007	.97 \pm .021	.30 \pm .060	.97 \pm .013	.97 \pm .012	.43 \pm .066
	Backbone	.85 \pm .019	.74 \pm .033	.79 \pm .018	.98 \pm .010	.72 \pm .060	.70 \pm .070	.96 \pm .012	.79 \pm .068	.76 \pm .033
	Foldseek		.89 \pm .033			.90 \pm .013			.87 \pm .022	
	TM-Align		.81 \pm .014			.91 \pm .040			.93 \pm .009	
PR-AUC	PLASMA	.98 \pm .011	.97 \pm .014	.96 \pm .017	.98 \pm .011	.97 \pm .019	.97 \pm .019	.97 \pm .014	.98 \pm .011	.97 \pm .012
	PLASMA-PF	.98 \pm .010	.95 \pm .005	.92 \pm .007	.98 \pm .012	.90 \pm .079	.95 \pm .026	.97 \pm .015	.96 \pm .006	.97 \pm .009
	EBA	.91 \pm .035	.93 \pm .019	.38 \pm .014	.98 \pm .012	.96 \pm .035	.28 \pm .063	.97 \pm .012	.97 \pm .012	.43 \pm .032
	Backbone	.86 \pm .023	.77 \pm .041	.82 \pm .027	.96 \pm .023	.67 \pm .093	.65 \pm .118	.96 \pm .016	.84 \pm .059	.80 \pm .038
	Foldseek		.84 \pm .031			.76 \pm .065			.81 \pm .026	
	TM-Align		.86 \pm .020			.89 \pm .064			.94 \pm .012	
F1-MAX	PLASMA	.97 \pm .009	.95 \pm .018	.92 \pm .022	.96 \pm .022	.95 \pm .030	.93 \pm .026	.98 \pm .013	.97 \pm .011	.97 \pm .011
	PLASMA-PF	.96 \pm .013	.90 \pm .006	.84 \pm .008	.96 \pm .027	.85 \pm .082	.90 \pm .031	.97 \pm .018	.94 \pm .016	.95 \pm .012
	EBA	.86 \pm .035	.87 \pm .024	.00 \pm .000	.97 \pm .021	.93 \pm .049	.00 \pm .000	.97 \pm .013	.97 \pm .008	.00 \pm .000
	Backbone	.79 \pm .008	.70 \pm .014	.73 \pm .013	.91 \pm .034	.62 \pm .087	.60 \pm .107	.92 \pm .020	.79 \pm .044	.71 \pm .018
	Foldseek		.91 \pm .046			.97 \pm .014			.96 \pm .015	
	TM-Align		.76 \pm .015			.87 \pm .063			.90 \pm .014	
LMS	PLASMA	.75 \pm .045	.69 \pm .019	.52 \pm .046	.82 \pm .062	.77 \pm .105	.65 \pm .088	.90 \pm .034	.87 \pm .038	.67 \pm .044
	PLASMA-PF	.78 \pm .055	.48 \pm .074	.23 \pm .021	.85 \pm .058	.49 \pm .082	.36 \pm .055	.94 \pm .029	.68 \pm .067	.43 \pm .032
Best		Baseline (TM-Align)								Worst

EBA performs reasonably well with sequence-based ANKH and ESM2 yet drops substantially with structure-based PROTSSN, especially under the extrapolation split. FOLDSEEK and TM-ALIGN remain consistently below PLASMA across nearly all conditions, reflecting the limited usefulness of global structural similarity for residue-level motif detection.

Beyond accuracy, PLASMA demonstrates exceptional computational efficiency. As shown in Figure 2, PLASMA achieves the best performance while requiring minimal time per protein pair—approximately 10ms for PLASMA and 7ms for PLASMA-PF. This represents a roughly 50 times speedup over global structure alignment methods like TM-Align and Foldseek, which require costly structural superposition, and about 3 times faster than EBA due to PLASMA’s fully differentiable OT formulation that is efficiently accelerated on GPUs, compared to EBA’s inherently sequential dynamic programming approach.

6.3 QUALITY OF PREDICTED ALIGNMENTS

Beyond quantitative metrics, we assess PLASMA’s robustness in identifying biologically meaningful substructures by examining both alignment scores and alignment matrices.

PLASMA effectively distinguishes proteins with shared local functional substructures even when overall structural similarity is low. Figure 3 provides evidence from two perspectives, with all embedding-based methods obtaining protein representations from ANKH. Figure 3A compares similarity score distributions for protein pairs from `test_inter`, where PLASMA and PLASMA-PF clearly separate positive and negative pairs. This advantage comes from the OT framework, which emphasizes local correspondences independent of overall similarity. In contrast, EBA and COSINESIM show substantial overlap between positive and negative distributions. EBA in particular lacks an upper bound on its scores, making them difficult to interpret and subject to calibration problems (*i.e.*, scores cannot be directly used as probabilities and lead to unstable thresholds). Figure 3B further groups test-set alignment scores by TM-score to assess performance under different levels of global similarity for protein pairs. Although all methods degrade as TM-score decreases, PLASMA and PLASMA-PF consistently maintain high ROC-AUC values above 0.9, whereas baseline EBA, COSINESIM, Foldseek, and TM-align deteriorate sharply on low-similarity samples when TM-score is sufficiently small (*e.g.*, lower than 0.5).

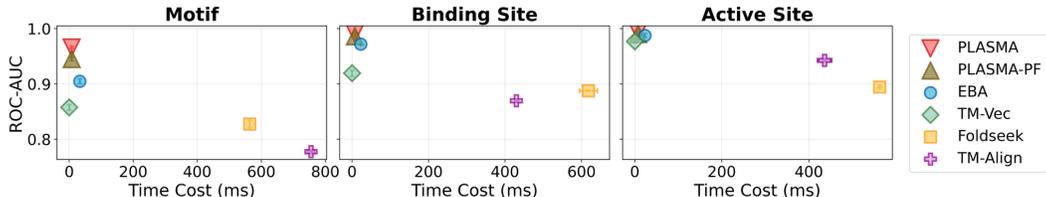


Figure 2: Performance versus computational efficiency comparison. ROC-AUC scores plotted against inference time (milliseconds) for motif and binding/active site detection using PROSTT5 embeddings. Points represent averages across three splits with standard error bars on both axes.

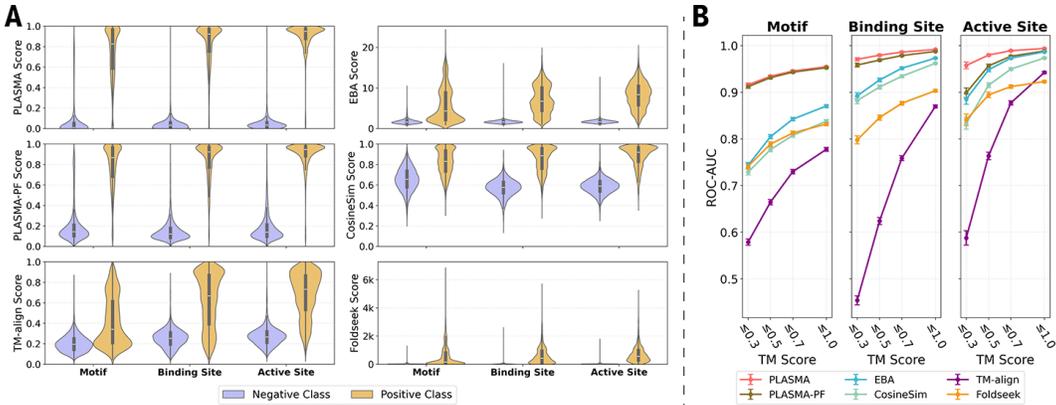


Figure 3: Alignment quality analysis across different approaches. **A.** Distribution of alignment scores for positive and negative protein pairs. **B.** ROC-AUC score trend at different global structural similarity levels.

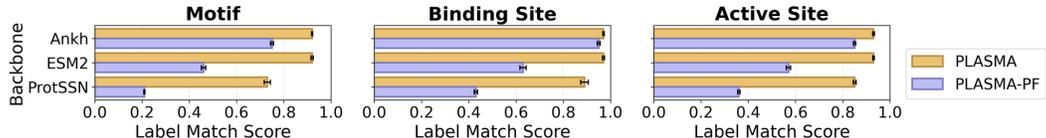


Figure 4: Label Match Score comparison between PLASMA and PLASMA-PF across different substructure types, demonstrating the improved alignment quality achieved through training.

While both PLASMA variants demonstrate strong performance in score-based discrimination, their alignment quality differs. This is evident in Figure 4, which compares their performance using the LMS score to evaluate correspondence between predicted alignments and annotated regions. PLASMA consistently outperforms PLASMA-PF across motifs, binding sites, and active sites, demonstrating that learning improves the prediction of local structural motifs. By contrast, while EBA also produces alignment matrices, it cannot be meaningfully assessed with LMS: its unconstrained formulation yields a maximal LMS of 1.0 regardless of true alignment accuracy.

6.4 REPRESENTATIVE ALIGNMENT EXAMPLES

The next experiment evaluates PLASMA’s utility in real biological applications. We examine three protein pairs of different substructure sizes (independent of the training set), including simple local motifs, complex cofactor-binding domains, and extended multi-element substructures. In each case, we provide UniProt identifiers, functional descriptions, alignment results, and visualizations from PLASMA and EBA, and corresponding analyses. Appendix N provides additional visualizations that further illustrate the generality of these conclusions. Collectively, these cases show PLASMA detects biologically meaningful local similarities across diverse sequences, structures, and functions.

Conserved Small Helical Motifs Across Functionally Diverse Protein Structures The first case matches local structures between P40343 (Vps27, a yeast ESCRT-0 complex component) and

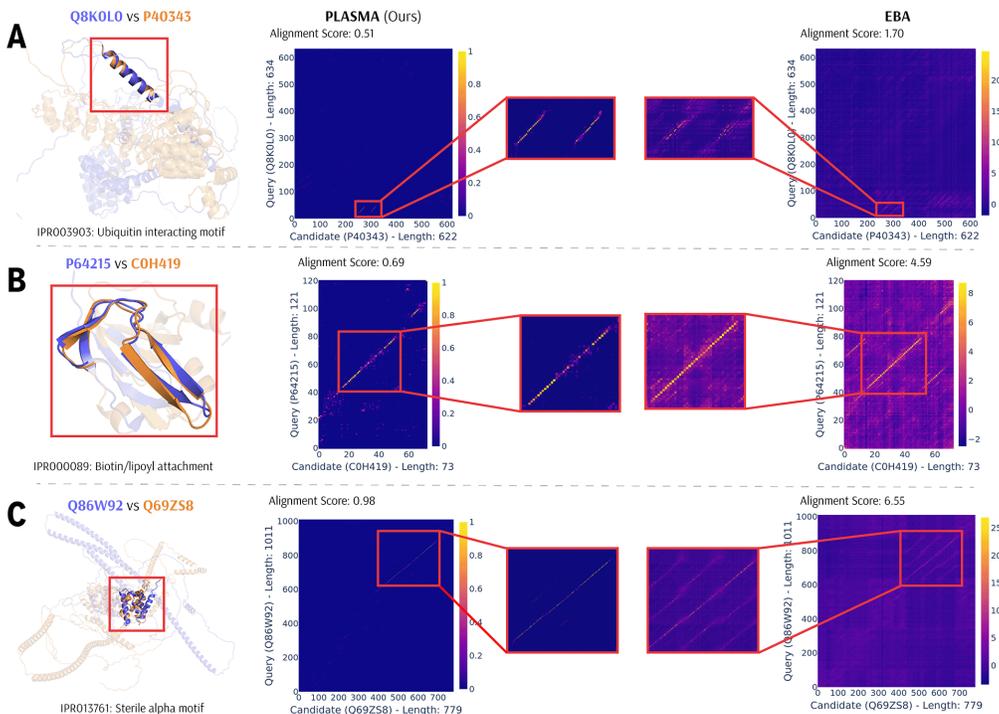


Figure 5: Representative alignment examples across three protein pairs. **A**, P40343 vs Q8K0L0. **B**, P64215 vs C0H419. **C**, Q69ZS8 vs Q86W92. Left: 3D structures with highlighted aligned regions. Center and right: alignment matrices from PLASMA and EBA with zoomed insets. A higher resolution version of this figure can be found at Appendix H.

Q8K0L0 (ASB2, a mouse E3 ubiquitin ligase substrate-recognition component). The two proteins share no apparent sequence homology (21.0% identity) and participate in distinct cellular processes (endosomal sorting versus proteasomal degradation), yet both use analogous helical arrangements for protein-protein interactions: Vps27’s GAT domain forms coiled-coils for ESCRT-I recruitment (Curtiss et al., 2007), whereas ASB2 employs ankyrin repeat helices for substrate recognition in the E3 ligase complex. PLASMA assigns high-confidence scores to residues mediating these interactions (Figure 5A). The 3D structure visualization also confirms the alignment of the conserved Leu-X-X-Leu-Leu motif for both proteins (Ren et al., 2008), with an aligned RMSD of 0.18 Å. This finding suggests potential convergent evolution of helical protein-binding interfaces across distinct cellular machineries. By contrast, EBA identifies multiple helices, but most correspond to nonfunctional scaffold regions rather than the relevant interaction motifs.

Structurally and Functionally Relevant motifs of Different Sizes and Metabolic Contexts The second case examines P64215 (GcvH, glycine cleavage system H protein from *Mycobacterium tuberculosis*) and C0H419 (YngHB, biotin/lipoyl attachment protein from *Bacillus subtilis*) (Cui et al., 2006). These proteins have different overall sequences (25.2% sequence identity) and metabolic functions: GcvH shuttles methylamine groups in glycine catabolism, while YngHB accommodates both biotin and lipoic acid in a single-domain architecture. Despite these differences, both bind similar cofactors and exhibit conserved β -sheet arrangements necessary for post-translational modification. As shown in Figure 5B, PLASMA successfully aligns the four-stranded β -barrel architectures, highlighting the critical lysine-containing β -turns with an overall alignment score of 0.69 and RMSD of 0.83, whereas the baseline EBA misaligns nonfunctional regions. The alignment of complex conserved structural motifs across protein families demonstrates the potential of PLASMA in revealing modular evolution and conserved cofactor-binding architectures.

Extended Multi-Element Substructures in Cell Adhesion Regulators The third case investigates Q69ZS8 (Kazrin, a scaffold protein in *Mus musculus*) and Q86W92 (Liprin- β 1/PPFIBP1, a human focal adhesion regulator). Despite their different cellular localizations and interaction part-

ners, they regulate distinct but mechanistically related aspects of cell-cell adhesion: Kazrin organizes desmosomal components in keratinocytes, and Liprin- β 1 modulates focal adhesion disassembly and cell migration. Yet both proteins rely on extended α -helical regions for protein-protein interactions (Groot et al., 2004). As in Figure 5C, PLASMA successfully aligns complex multi-coil substructures spanning multiple helical segments interspersed with flexible linkers, with an overall alignment score of 0.98 and RMSD 0.82 Å. The alignment highlights conserved leucine-rich motifs and hinge regions that stabilize oligomerization interfaces, revealing analogous scaffolding strategies. In contrast, EBA identifies plausible structures but often misaligns helices or matches nonfunctional scaffold regions, failing to capture more than just biologically meaningful substructures.

7 RELATED WORKS

Protein Global Structure Alignment Global structure alignment methods evaluate overall protein similarity. Classic approaches like TM-ALIGN (Zhang, 2005) are foundational, while modern methods increase efficiency by abstracting structures into 1D sequences (FOLDSEEK (Van Kempen et al., 2024)), representing them as fixed vectors for rapid search (TM-VEC (Hamamsy et al., 2024)), or using advanced spatial indexing (GTALIGN (Margelevičius, 2024)). The field has also expanded to align multiple structures (MTM-ALIGN (Dong et al., 2018)), multi-chain complexes (MM-ALIGN (Mukherjee & Zhang, 2009)), and diverse macromolecules universally (US-ALIGN (Zhang et al., 2022)). However, their global nature limits the detection of conserved motifs in dissimilar proteins.

Substructure and Sequence-based Alignment To find local similarities, substructure-based methods use graph-based residue embeddings (Tan et al., 2024), focus on active-site environments (Castillo & Ollila, 2025), or apply linear-assignment formulations (Zhang et al., 2025). PLM-based residue representations are also widely used from raw embedding similarity scoring (Kaminski et al., 2023; Liu et al., 2024) to learned alignment models and embedding-aware dynamic programming (Llinares-López et al., 2023; Iovino & Ye, 2024). OT-based differentiable graph matching has been used to learn structure/function-aware substitution matrices (Pellizzoni et al., 2024), with a primary focus on learning matching costs. PLASMA instead targets residue-level local substructure alignment, producing explicit mappings with practical speed and interpretability. Meanwhile, embedding-score-based alignment methods remain hard to interpret quantitatively, as their scores are essentially unbounded (Pantolini et al., 2024).

8 CONCLUSION AND DISCUSSION

This work presents PLASMA, a local structural motif alignment framework leveraging regularized optimal transport to detect biologically meaningful local similarities across proteins with diverse sequences, structures, and functions. PLASMA consistently outperforms baseline methods in accuracy, efficiency, and interpretability, capturing subtle structural correspondences often invisible to global alignments. Its trainable variant benefits from supervision to improve alignment precision, while the training-free variant achieves robust performance without task-specific labels.

Beyond quantitative performance, PLASMA provides clear, residue-level alignment matrices that support mechanistic insights into protein function, evolutionary relationships, and structure-guided protein engineering. Its ability to handle varying substructure sizes and complexities (*e.g.*, from short helices to extended multi-element domains) demonstrates versatility and practical relevance. Overall, PLASMA establishes a new standard for accurate, efficient, interpretable, and practically applicable protein local structural motif alignment.

ACKNOWLEDGMENTS

This work was supported by the grants from National Science Foundation of China (Grant Number 92451301; 62302291), the AI for Science Program by Shanghai Municipal Commission of Economy and Informatization (2025-GZL-RGZN-BTBX-02009), the National Key Research and Development Program of China (2024YFA0917603), and Computational Biology Key Program of Shanghai Science and Technology Commission (23JS1400600).

Reproducibility Statement To promote reproducibility, we release all source code and trained models under an open-source license, which is available at <https://github.com/ZW471/PLASMA-Protein-Local-Alignment.git>. Details of data sources are provided in Appendix C.1. Task definitions, evaluation protocols, and hyperparameter settings are described in Sections 6.1 and Appendices C.2. Implementation details and instructions for reproducing experiments are included in the project repository to facilitate independent verification.

Ethics Statement All experiments are conducted on publicly available protein sequence and structure databases. We follow established ethical guidelines in data usage and acknowledge that historical biases present in these resources may be reflected in our results, which is independent of model development.

The Use of Large Language Models (LLM) In the preparation of this manuscript, GPT-5 and GPT-4o were utilized as writing assistants. The usage was strictly limited to improving grammar, clarity, and overall readability. All scientific ideas, experimental results, and conclusions were conceived and formulated exclusively by the authors. All text polished or modified by the LLM was subsequently reviewed and edited by the authors to ensure that the original scientific meaning was accurately preserved.

REFERENCES

- Sebastian Bittrich, Stephen K. Burley, and Alexander S. Rose. Real-time structural motif searching in proteins using an inverted index strategy. *PLOS Computational Biology*, 16(12):e1008502, 2020. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1008502.
- Matthias Blum, Antonina Andreeva, Laise Cavalcanti Florentino, Sara Rocio Chuguransky, Tiago Grego, Emma Hobbs, Beatriz Lazaro Pinto, Ailsa Orr, Typhaine Paysan-Lafosse, Irina Ponomareva, Gustavo A Salazar, Nicola Bordin, Peer Bork, Alan Bridge, Lucy Colwell, Julian Gough, Daniel H Haft, Ivica Letunic, Felipe Llinares-López, Aron Marchler-Bauer, Laetitia Meng-Papaxanthos, Huaiyu Mi, Darren A Natale, Christine A Orengo, Arun P Pandurangan, Damiano Piovesan, Catherine Rivoire, Christian J A Sigrist, Narmada Thanki, Françoise Thibaud-Nissen, Paul D Thomas, Silvio C E Tosatto, Cathy H Wu, and Alex Bateman. InterPro: The protein sequence classification resource in 2025. *Nucleic Acids Research*, 53(D1):D444–D456, 2025. ISSN 0305-1048, 1362-4962. doi: 10.1093/nar/gkae1082.
- Nadav Brandes, Dan Ofer, Yam Peleg, Nadav Rappoport, and Michal Linial. ProteinBERT: A universal deep-learning model of protein sequence and function. *Bioinformatics*, 38(8):2102–2110, 2022. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btac020.
- Luis Caffarelli and Robert McCann. Free boundaries in optimal transport and monge-ampère obstacle problems. *Annals of Mathematics*, 171(2):673–730, 2010. ISSN 0003-486X. doi: 10.4007/annals.2010.171.673.
- Sandra Castillo and Osmo Henri Samuli Ollila. Actseek: fast and accurate search algorithm of active sites in alphafold database. *Bioinformatics*, 41(8):btaf424, 2025.
- Gaofeng Cui, Beiyan Nan, Jicheng Hu, Yiping Wang, Changwen Jin, and Bin Xia. Identification and solution structures of a single domain biotin/lipoyl attachment protein from bacillus subtilis. *Journal of Biological Chemistry*, 281(29):20598–20607, 2006.
- Matt Curtiss, Charles Jones, and Markus Babst. Efficient cargo sorting by escrt-i and the subsequent release of escrt-i from multivesicular bodies requires the subunit mvb12. *Molecular Biology of the Cell*, 18(2):636–645, 2007.
- Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *NIPS’13*, volume 2. Curran Associates Inc., 2013. doi: 10.5555/2999792.2999868.
- Runze Dong, Zhenling Peng, Yang Zhang, and Jianyi Yang. mTM-align: an algorithm for fast and accurate multiple protein structure alignment. *Bioinformatics*, 34(10):1719–1725, 2018.
- Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. *Advances in Neural Information Processing Systems*, 33:10491–10502, 2020.
- Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. ProfTrans: towards cracking the language of life’s code through self-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:7112–7127, 2021.
- Ahmed Elnaggar, Hazem Essam, Wafaa Salah-Eldin, Walid Moustafa, Mohamed Elkerdawy, Charlotte Rochereau, and Burkhard Rost. Ankh: Optimized protein language model unlocks general-purpose modelling. *arXiv:2301.06568*, 2023.
- Alessio Figalli. The optimal partial transport problem. *Archive for Rational Mechanics and Analysis*, 195(2):533–560, 2010. ISSN 0003-9527, 1432-0673. doi: 10.1007/s00205-008-0212-7.
- Karen R Groot, Lisa M Sevilla, Kazunori Nishi, Teresa DiColandrea, and Fiona M Watt. Kazrin, a novel periplakin-interacting protein associated with desmosomes and the keratinocyte plasma membrane. *The Journal of Cell Biology*, 166(5):653–659, 2004.

- Tymor Hamamsy, James T. Morton, Robert Blackwell, Daniel Berenberg, Nicholas Carriero, Vladimir Gligorijevic, Charlie E. M. Strauss, Julia Koehler Leman, Kyunghyun Cho, and Richard Bonneau. Protein remote homology detection and structural alignment using deep learning. *Nature Biotechnology*, 42(6):975–985, 2024. ISSN 1087-0156, 1546-1696. doi: 10.1038/s41587-023-01917-2.
- Michael Heinzinger, Konstantin Weissenow, Joaquin Gomez Sanchez, Adrian Henkel, Milot Mirdita, Martin Steinegger, and Burkhard Rost. Bilingual language model for protein sequence and structure. *NAR Genomics and Bioinformatics*, 6(4):lqae150, 2024. ISSN 2631-9268. doi: 10.1093/nargab/lqae150.
- Liisa Holm. Using Dali for protein structure comparison. In Zoltán Gáspári (ed.), *Structural Bioinformatics*, volume 2112, pp. 29–42. Springer US, 2020. ISBN 978-1-0716-0269-0 978-1-0716-0270-6. doi: 10.1007/978-1-0716-0270-6_3.
- Torgeir R Hvidsten, Astrid Læg Reid, Andriy Kryshchak, Gunnar Andersson, Krzysztof Fidelis, and Jan Komorowski. A comprehensive analysis of the structure-function relationship in proteins based on local structure similarity. *PLoS One*, 4(7):e6266, 2009.
- Benjamin Giovanni Iovino and Yuzhen Ye. Protein embedding based alignment. *BMC Bioinformatics*, 25(1):85, 2024.
- Arian Rokkum Jamasb, Alex Morehead, Chaitanya K. Joshi, Zuobai Zhang, Kieran Didi, Simon V Mathis, Charles Harris, Jian Tang, Jianlin Cheng, Pietro Lio, and Tom Leon Blundell. Evaluating representation learning on the protein structure universe. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=sTYuRVrdK3>.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislaw Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596:583–589, 2021. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-021-03819-2.
- Kamil Kaminski, Jan Ludwiczak, Kamil Pawlicki, Vikram Alva, and Stanislaw Dunin-Horkawicz. pLM-BLAST: Distant homology detection based on direct comparison of sequence representations from protein language models. *Bioinformatics*, 39(10):btad579, 2023. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btad579.
- Hyunbin Kim, Rachel Seongeun Kim, Milot Mirdita, and Martin Steinegger. Structural motif search across the protein-universe with FoldDisco. *bioRxiv*, pp. 2025–07, 2025.
- Michael Knudsen and Carsten Wiuf. The CATH database. *Human Genomics*, 4(3):207, 2010. ISSN 1479-7364. doi: 10.1186/1479-7364-4-3-207.
- Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan Dos Santos Costa, Maryam Fazel-Zarandi, Tom Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.ade2574.
- Wei Liu, Ziyi Wang, Ronghui You, Chenghan Xie, Hong Wei, Yi Xiong, Jianyi Yang, and Shan-feng Zhu. PLMSearch: Protein language model powers accurate and fast sequence search for remote homology. *Nature Communications*, 15(1):2775, 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-46808-5.
- Yang Liu, Qing Ye, Liwei Wang, and Jian Peng. Learning structural motif representations for efficient protein structure search. *Bioinformatics*, 34(17):i773–i780, 2018.

- Felipe Llinares-López, Quentin Berthet, Mathieu Blondel, Olivier Teboul, and Jean-Philippe Vert. Deep embedding and alignment of protein sequences. *Nature Methods*, 20(1):104–111, 2023.
- Mindaugas Margelevičius. GTalign: Spatial index-driven protein structure alignment, superposition, and search. *Nature Communications*, 15(1):7305, 2024.
- Gonzalo Mena, David Belanger, Scott Linderman, and Jasper Snoek. Learning latent permutations with gumbel-sinkhorn networks. In *International Conference on Learning Representations*, 2018.
- Caitlyn L Mills, Rohan Garg, Joslynn S Lee, Liang Tian, Alexandru Suciuc, Gene D Cooperman, Penny J Beuning, and Mary Jo Ondrechen. Functional classification of protein structures by local structure matching in graph representation. *Protein Science*, 27(6):1125–1135, 2018.
- Srayanta Mukherjee and Yang Zhang. MM-align: a quick algorithm for aligning multiple-chain protein complex structures using iterative dynamic programming. *Nucleic Acids Research*, 37(11):e83–e83, 2009.
- Lorenzo Pantolini, Gabriel Studer, Joana Pereira, Janani Durairaj, Gerardo Tauriello, and Torsten Schwede. Embedding-based alignment: Combining protein language models with dynamic programming alignment to detect structural similarities in the twilight-zone. *Bioinformatics*, 40(1):btad786, 2024. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btad786.
- Paolo Pellizzoni, Carlos Oliver, and Karsten Borgwardt. Structure-and function-aware substitution matrices via learnable graph matching. In *International Conference on Research in Computational Molecular Biology*, pp. 288–307. Springer, 2024.
- Vaibhav Raj, Indradyumna Roy, Ashwin Ramachandran, Soumen Chakrabarti, and Abir De. Charting the design space of neural graph representations for subgraph matching. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=5pd78GmXC6>.
- Ashwin Ramachandran, Vaibhav Raj, Indradyumna Roy, Soumen Chakrabarti, and Abir De. Iteratively refined early interaction alignment for subgraph matching based graph retrieval. In *Advances in Neural Information Processing Systems*, 2024.
- Jihui Ren, Natasha Pashkova, Stanley Winistorfer, and Robert C Piper. Doa1/ufd3 plays a role in sorting ubiquitinated membrane proteins into multivesicular bodies. *Journal of Biological Chemistry*, 283(31):21599–21611, 2008.
- Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967. ISSN 0030-8730, 0030-8730. doi: 10.2140/pjm.1967.21.343.
- Baris E. Suzek, Hongzhan Huang, Peter McGarvey, Raja Mazumder, and Cathy H. Wu. UniRef: Comprehensive and non-redundant UniProt reference clusters. *Bioinformatics*, 23(10):1282–1288, 2007. ISSN 1367-4811, 1367-4803. doi: 10.1093/bioinformatics/btm098.
- Yang Tan, Lirong Zheng, Bozita Zhong, Liang Hong, and Bingxin Zhou. Protein representation learning with sequence information embedding: Does it always lead to a better performance? In *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 233–239. IEEE, 2024.
- Yang Tan, Wenrui Gou, Bozita Zhong, Liang Hong, Huiqun Yu, and Bingxin Zhou. VenusX: Unlocking fine-grained functional understanding of proteins. *arXiv:2505.11812*, 2025a.
- Yang Tan, Bingxin Zhou, Lirong Zheng, Guisheng Fan, and Liang Hong. Semantical and geometrical protein encoding toward enhanced bioactivity and thermostability. *eLife*, 13:RP98033, 2025b.
- Michel Van Kempen, Stephanie S. Kim, Charlotte Tumescheit, Milot Mirdita, Jeongjae Lee, Cameron L. M. Gilchrist, Johannes Söding, and Martin Steinegger. Fast and accurate protein structure search with foldseek. *Nature Biotechnology*, 42(2):243–246, 2024. ISSN 1087-0156, 1546-1696. doi: 10.1038/s41587-023-01773-0.

Mihaly Varadi, Stephen Anyango, Mandar Deshpande, Sreenath Nair, Cindy Natassia, Galabina Yordanova, David Yuan, Oana Stroe, Gemma Wood, Agata Laydon, et al. Alphafold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Research*, 50(D1):D439–D444, 2022.

Chengxin Zhang, Morgan Shine, Anna Marie Pyle, and Yang Zhang. US-align: universal structure alignments of proteins, nucleic acids, and macromolecular complexes. *Nature Methods*, 19(9): 1109–1115, 2022.

Xuechen Zhang, Zhuoyang Chen, Junyu Li, Qiong Luo, Longjun Wu, and Weichuan Yu. eplsap-align: a non-sequential protein structural alignment solver with entropy-regularized partial linear sum assignment problem formulation. *Bioinformatics*, 41(6):btaf309, 2025.

Y. Zhang. TM-align: A protein structure alignment algorithm based on the TM-score. *Nucleic Acids Research*, 33(7):2302–2309, 2005. ISSN 1362-4962. doi: 10.1093/nar/gki524.

This appendix provides additional details, analyses, and results that complement the main paper.

- Appendix A gives the full derivation of our OT objective.
- Appendix B presents a more precise discussion of computational cost.
- Appendix C describes the benchmark datasets (**VenusX**) and the hyperparameter configuration.
- Appendix D summarizes all comparison methods, including global structure alignment, global embedding-based alignment, local embedding-based alignment, and the backbone models.
- Appendix E and Appendix F report complete quantitative results for all backbones on `test_inter` and `test_extra` split.
- Appendix G provides further insight into the contribution of individual components.
- Appendices H-I contains additional visualizations of alignment matrices and case studies.
- Appendices J-M offer more detailed quantitative analyses of model behaviour under different settings.

A OPTIMAL TRANSPORT FORMULATION FOR PROTEIN ALIGNMENT

To circumvent the computational bottleneck of explicit fragment enumeration, we reframe the alignment problem as finding optimal correspondences between individual residues rather than pre-defined fragments. This approach leverages optimal transport theory, which provides a principled framework for finding the most efficient assignment between two sets of points based on their similarity and a transportation cost function.

Specifically, we model protein substructure alignment as an entropy regularized optimal transport problem that determines how to optimally redistribute alignment weights from query residues to candidate residues. Instead of relying solely on explicit structural coordinates, this formulation operates on learned residue representations that encode local neighborhood properties, biochemical characteristics, and structural context. The optimal transport solver then identifies which residues should be matched by minimizing the total transportation cost—effectively the sum of dissimilarities between matched residue pairs—across the embedding space.

This approach naturally produces soft, many-to-many alignments where functionally and structurally similar residues are preferentially matched, while simultaneously identifying the corresponding aligned fragments without explicit enumeration. Mathematically, we formulate this as the following optimal transport problem with entropic constraints:

$$\min_{\Omega} \sum_{i=1}^N \sum_{j=1}^M \Omega_{ij} C_{ij} - \lambda \sum_{i=1}^N \sum_{j=1}^M \Omega_{ij} \log(\Omega_{ij}) \quad (8)$$

$$\text{subject to: } \sum_{j=1}^M \Omega_{ij} = 1, \quad \forall i \in \{1, \dots, N\} \quad (9)$$

$$\sum_{i=1}^N \Omega_{ij} = 1, \quad \forall j \in \{1, \dots, M\} \quad (10)$$

$$\Omega_{ij} \geq 0, \quad \forall i, j \quad (11)$$

Here, $\Omega \in \mathbb{R}^{N \times M}$ is the transport plan (alignment matrix), C_{ij} represents the cost of aligning query residue i to candidate residue j , and λ is the entropic regularization parameter that controls the smoothness of the alignment. This optimization seeks to find the optimal transport plan that minimizes the total alignment cost while the entropic regularization term ($-\lambda$ term) encourages smooth, distributed assignments rather than hard one-to-one mappings. The equality constraints ensure each query residue distributes 1 total weight and each candidate residue receives 1 total weight.

B COMPLEXITY ANALYSIS

PLASMA achieves optimal $O(N^2)$ complexity while maintaining full differentiability. The cost matrix computation dominates computational requirements, requiring $O(N \cdot M \cdot D) = O(N^2 \cdot D)$ operations for the hinge non-linearity between proteins of lengths N and M , where D represents the embedding dimension. The siamese network contributes $O(N \cdot D^2)$ operations per protein (if using a two-layer MLP), yielding total $O(N \cdot D^2)$ since $D \ll N$ in practice. The Sinkhorn algorithm requires $O(T \cdot N^2)$ operations where T represents the number of iterations (typically $T \ll N$). The Plan Assessor contributes $O(N^2)$ for substructure similarity computation and $O(K^2 \cdot N^2)$ for confidence weight calculation via diagonal convolution with kernel size $K \ll N$. The overall complexity remains $O(N^2)$, matching the best achievable complexity of the methods based on dynamic programming.

C DETAILED EXPERIMENTAL SETUP

C.1 BENCHMARK DATASETS: VENUSX

We construct our evaluation datasets from the **VenusX** (Tan et al., 2025a) benchmark (<https://github.com/ai4protein/VenusX>), which provides protein pairs with annotated biologically important substructures curated from the InterPro (Blum et al., 2025) database. We focus on three substructure types: activation sites, binding sites, and motifs, corresponding to the VenusX_Res_{Act/BindI/Motif}_MP50 datasets where protein pairs share less than 50% sequence similarity. These datasets present increasing difficulty due to their substructure sizes: active sites (18.7 ± 7.0 residues), binding sites (26.6 ± 21.7 residues), and motifs (80.23 ± 73.8 residues). From each VenusX dataset, we generate 20,000 protein pairs with balanced labels: half sharing the same InterPro family ID (positive pairs, $y = 1$) and half from different families (negative pairs, $y = 0$). Each sample is represented as $(\mathcal{P}_q, \mathcal{P}_c, \mathbf{l}_q, \mathbf{l}_c, y)$, where \mathcal{P}_q and \mathcal{P}_c are the protein pair, \mathbf{l}_q and \mathbf{l}_c are their respective substructure annotations, and y indicates family membership.

To evaluate all the embedding based methods’ generalization capability across different evolutionary contexts, we create two complementary test scenarios using three different random seeds for robust evaluation. This dual evaluation is crucial for protein analysis since biological systems constantly encounter both familiar protein families with slight variations and entirely novel protein architectures through evolution, horizontal gene transfer, and structural convergence. First, we randomly exclude 10% of InterPro family IDs and split the remaining data into training (75%), validation (5%), and `test_inter` (20%). `test_inter` evaluates *interpolation* performance—the model’s ability to recognize substructure similarities within the distribution of known protein families, mimicking scenarios where researchers analyze variants of well-characterized proteins. Second, we create `test_extra` by sampling an equivalent number of protein pairs exclusively from the excluded InterPro families (maintaining the same 50–50 balance between positive and negative pairs). `test_extra` evaluates *extrapolation* performance—the model’s ability to identify functional similarities in completely novel protein families, which is critical for annotating newly discovered proteins, understanding convergent evolution, and predicting function in understudied organisms. For each test scenario, the data exclusion and splitting procedure is repeated across three different seeds (1, 42, and 100) to ensure statistical reliability.

C.2 HYPERPARAMETER CONFIGURATION

For both PLASMA and PLASMA-PF variants, we employ the following hyperparameters: the siamese network uses a hidden dimension of 512 to balance expressiveness with computational efficiency. To ensure computational feasibility while maintaining statistical significance, our training sets only use 1500 protein pairs by sampling 10% of the full training set. The Sinkhorn temperature parameter τ is set to 0.1 to encourage sparse, focused alignments that highlight the most relevant correspondences. The diagonal convolution kernel size $K = 10$ captures sequential patterns in alignment matrices, while the residue matching threshold $\rho = 0.5$ defines when transport weights indicate meaningful correspondences between residue pairs. See Appendix M for detailed sensitivity analysis and justification of these choices.

D BASELINES

D.1 GLOBAL STRUCTURE ALIGNMENT METHODS

Traditional structural biology approaches rely on atomic coordinates to identify protein similarities:

- TM-ALIGN (Zhang, 2005) represents the gold standard for protein structure alignment based on Template Modeling scores. This method performs geometric alignment of protein backbones to identify structurally similar regions.
- FOLDSEEK (Van Kempen et al., 2024) performs structural alignment using 3Di tokenizations, converting 3D structural information into sequence-like representations for comparison.

D.2 GLOBAL EMBEDDING-BASED ALIGNMENT

COSINESIM methods employ direct cosine similarity between globally aggregated protein embeddings from the backbone models discussed in Appendix D.4, similar to the approach used in TM-Vec (Hamamsy et al., 2024). This approach provides a baseline for embedding-based similarity without explicit residue-level alignment, representing proteins as single vectors and measuring their similarity through cosine distance.

D.3 LOCAL EMBEDDING-BASED ALIGNMENT

EBA (Pantolini et al., 2024) represents the current state-of-the-art in local embedding-based alignment, combining statistical alignment with neural embeddings to identify similar substructures. This method performs local alignment at the residue level using learned representations.

D.4 BACKBONES

We evaluate PLASMA with seven popular protein sequence and structure representation models, using the following specific versions and configurations:

- ANKH (Elnaggar et al., 2023): We employ the `base` model variant, which is a compact encoder-decoder architecture optimized for protein sequences with 110 million parameters. This model was trained on protein sequences using a masked language modeling objective and represents one of the most parameter-efficient protein language models. *Available at:* <https://huggingface.co/ElnaggarLab/ankh-base>
- ESM2 (Lin et al., 2023): We utilize the `t33_650M_UR50D` variant, a 650-million parameter encoder-only transformer model with 33 layers. This model was trained on the UniRef50 database and represents one of the largest and most comprehensive protein language models available, providing rich contextual representations for protein analysis. *Available at:* https://huggingface.co/facebook/esm2_t33_650M_UR50D
- PROSTT5 (Heinzinger et al., 2024): We use the `AA2fold` checkpoint, which is specifically fine-tuned for protein folding applications. This bilingual language model can process both amino acid sequences and structural information, making it particularly well-suited for structure-aware protein analysis tasks. *Available at:* <https://huggingface.co/Rostlab/ProstT5>
- PROT5 (Elnaggar et al., 2021): We employ the `xl_half_uniref50-enc` model, which uses only the encoder component of the T5 architecture. This variant was trained on UniRef50 (Suzek et al., 2007) sequences and provides balanced performance between computational efficiency and representation quality with approximately 3 billion parameters. *Available at:* https://huggingface.co/Rostlab/prot_t5_xl_half_uniref50-enc
- PROTSSN (Tan et al., 2025b): We utilize the `k20_h512` configuration, which combines sequence and structural information through a hybrid architecture. The model uses $k = 20$ nearest neighbors for structural context and hidden dimensions of 512, enabling it to capture both sequential and geometric protein properties. *Available at:* <https://github.com/tyang816/ProtSSN>
- TM-VEC (Hamamsy et al., 2024): We employ the `cath_model_large` variant, which was specifically trained on the CATH structural classification database (Knudsen & Wiuf, 2010). This

model specializes in learning structure-aware representations and is particularly effective for detecting remote homology relationships based on structural similarity. *Available at:* https://figshare.com/articles/dataset/TMvec_DeepBLAST_models/25810099

- PROTBERT (Brandes et al., 2022): We use the bfd checkpoint, which was trained on the Big Fantastic Database (Jumper et al., 2021) containing over 2.1 billion protein sequences. This BERT-based model provides robust protein representations through bidirectional context modeling and large-scale pretraining. *Available at:* https://huggingface.co/Rostlab/prot_bert_bfd

E FULL INTERPOLATION PERFORMANCE COMPARISON

This section presents comprehensive experimental results using seven backbone protein representation learning models (PROSTT5, PROTT5, ANKH, ESM2, PROTSSN, TM-VEC, and PROTBERT) across three substructure alignment tasks (motifs, binding sites, and active sites) on the `test_inter` dataset. The key findings demonstrate that both PLASMA and PLASMA-PF consistently achieve superior performance across all backbone-task combinations, highlighting the robustness of our optimal transport framework regardless of the underlying protein representation model. Additionally, the Label Match Score (LMS) results show that the trainable PLASMA variant significantly outperforms the parameter-free PLASMA-PF in predicting precise locations of aligned substructures, validating the benefits of supervised learning for accurate residue-level alignment localization.

Table 2: Comprehensive motif detection results on `test_inter` dataset across seven protein representation models.

Metrics	Methods	Motif						
		PROSTT5	PROTT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.97\pm.002	.97\pm.002	.95\pm.002	.96\pm.002	.96\pm.001	.92\pm.004	.87\pm.004
	PLASMA-PF	.94 \pm .003	.96 \pm .002	.95 \pm .003	.93 \pm .004	.91 \pm .003	.87 \pm .001	.85 \pm .004
	EBA	.90 \pm .004	.91 \pm .004	.87 \pm .005	.88 \pm .003	.44\pm.002	.88 \pm .004	.73 \pm .006
	COSINESIM	.82 \pm .008	.87 \pm .003	.84 \pm .006	.73 \pm .009	.75 \pm .006	.86 \pm .005	.57\pm.014
	FOLDSEEK	.83 \pm .007						
	TM-ALIGN	.78 \pm .003						
	PR-AUC	PLASMA	.96\pm.002	.96\pm.003	.95\pm.002	.97\pm.001	.96\pm.001	.93\pm.004
PLASMA-PF	.95 \pm .003	.96 \pm .002	.95 \pm .003	.94 \pm .002	.92 \pm .001	.88 \pm .002	.87 \pm .003	
EBA	.92 \pm .004	.93 \pm .004	.90 \pm .004	.90 \pm .004	.45\pm.004	.91 \pm .004	.78 \pm .005	
COSINESIM	.85 \pm .005	.88 \pm .002	.86 \pm .005	.76 \pm .008	.78 \pm .006	.88 \pm .002	.63\pm.016	
FOLDSEEK	.78 \pm .008							
TM-ALIGN	.83 \pm .004							
F1-MAX	PLASMA	.92\pm.001	.93\pm.001	.93\pm.002	.93\pm.004	.91\pm.000	.88\pm.004	.80 \pm .001
	PLASMA-PF	.90 \pm .005	.93 \pm .002	.93 \pm .004	.89 \pm .004	.84 \pm .004	.84 \pm .002	.77 \pm .002
	EBA	.84 \pm .006	.86 \pm .003	.80 \pm .005	.81 \pm .003	.00\pm.000	.82 \pm .003	.69 \pm .006
	COSINESIM	.74 \pm .007	.79 \pm .004	.76 \pm .003	.69 \pm .001	.70 \pm .002	.78 \pm .005	.67 \pm .003
	FOLDSEEK	.84 \pm .007						
	TM-ALIGN	.70 \pm .002						
	LMS	PLASMA	.91\pm.007	.92\pm.001	.92\pm.002	.92\pm.005	.73\pm.013	.76\pm.005
PLASMA-PF		.57 \pm .003	.37 \pm .006	.75 \pm .006	.46 \pm .009	.21 \pm .002	.45 \pm .001	.39 \pm .009

Table 3: Comprehensive binding site detection results on `test_inter` dataset across seven protein representation models.

Metrics	Methods	Binding Site						
		PROSTT5	PROTT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.99\pm.001	.99\pm.000	.99\pm.000	.99\pm.001	.99\pm.001	.96\pm.003	.98\pm.001
	PLASMA-PF	.99 \pm .001	.99 \pm .001	.99 \pm .000	.96 \pm .003	.97 \pm .001	.92 \pm .004	.90 \pm .003
	EBA	.97 \pm .001	.97 \pm .001	.97 \pm .001	.97 \pm .002	.40 \pm .005	.95 \pm .000	.84 \pm .006
	COSINESIM	.87 \pm .005	.88 \pm .004	.96 \pm .002	.79 \pm .009	.75 \pm .008	.92 \pm .006	.66 \pm .008
	FOLDSEEK	.89 \pm .001						
	TM-ALIGN	.87 \pm .003						
PR-AUC	PLASMA	.99\pm.001	.99\pm.001	.99\pm.000	.99\pm.001	.99\pm.001	.97\pm.002	.98\pm.001
	PLASMA-PF	.99 \pm .001	.99 \pm .001	.99 \pm .000	.97 \pm .002	.98 \pm .001	.93 \pm .004	.93 \pm .001
	EBA	.98 \pm .000	.98 \pm .001	.98 \pm .001	.98 \pm .001	.42 \pm .004	.96 \pm .001	.87 \pm .003
	COSINESIM	.90 \pm .005	.90 \pm .003	.97 \pm .002	.83 \pm .007	.78 \pm .005	.94 \pm .004	.70 \pm .006
	FOLDSEEK	.83 \pm .002						
	TM-ALIGN	.91 \pm .002						
F1-MAX	PLASMA	.98\pm.002	.98\pm.001	.98\pm.001	.98\pm.002	.97\pm.002	.95\pm.002	.94\pm.002
	PLASMA-PF	.96 \pm .001	.97 \pm .001	.97 \pm .001	.92 \pm .003	.94 \pm .001	.91 \pm .005	.83 \pm .002
	EBA	.94 \pm .001	.94 \pm .001	.94 \pm .001	.93 \pm .002	.00 \pm .000	.93 \pm .001	.78 \pm .007
	COSINESIM	.80 \pm .008	.80 \pm .005	.91 \pm .005	.73 \pm .006	.69 \pm .006	.86 \pm .007	.67 \pm .001
	FOLDSEEK	.94 \pm .001						
	TM-ALIGN	.84 \pm .005						
LMS	PLASMA	.93\pm.002	.93\pm.003	.93\pm.004	.93\pm.003	.85\pm.006	.86\pm.002	.84\pm.003
	PLASMA-PF	.80 \pm .008	.59 \pm .008	.85 \pm .005	.57 \pm .009	.36 \pm .005	.60 \pm .008	.44 \pm .004

Table 4: Comprehensive active site detection results on `test_inter` dataset across seven protein representation models.

Metrics	Methods	Active Site						
		PROSTT5	PROTT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.99\pm.001	.99\pm.001	.99\pm.001	.99\pm.001	.99\pm.002	.99\pm.003	.99\pm.004
	PLASMA-PF	.99 \pm .002	.99 \pm .003	.99 \pm .003	.96 \pm .002	.98 \pm .002	.98 \pm .003	.94 \pm .006
	EBA	.99 \pm .003	.99 \pm .003	.99 \pm .003	.99 \pm .003	.43 \pm .005	.99 \pm .003	.90 \pm .005
	COSINESIM	.91 \pm .004	.91 \pm .003	.97 \pm .002	.78 \pm .009	.74 \pm .006	.98 \pm .002	.66 \pm .003
	FOLDSEEK	.89 \pm .001						
	TM-ALIGN	.94 \pm .003						
PR-AUC	PLASMA	.99\pm.000	.99\pm.001	.99\pm.001	.99\pm.000	.99\pm.001	.99\pm.002	.99\pm.003
	PLASMA-PF	.99 \pm .001	.99 \pm .002	.99 \pm .002	.97 \pm .001	.99 \pm .001	.98 \pm .003	.95 \pm .004
	EBA	.99 \pm .003	.99 \pm .002	.99 \pm .002	.99 \pm .002	.43 \pm .006	.99 \pm .003	.92 \pm .003
	COSINESIM	.93 \pm .002	.92 \pm .001	.98 \pm .001	.83 \pm .004	.79 \pm .002	.98 \pm .001	.70 \pm .007
	FOLDSEEK	.83 \pm .006						
	TM-ALIGN	.96 \pm .001						
F1-MAX	PLASMA	.98\pm.003	.98\pm.003	.99\pm.003	.98\pm.001	.99\pm.002	.98\pm.003	.96 \pm .004
	PLASMA-PF	.98 \pm .003	.98 \pm .004	.98 \pm .003	.93 \pm .004	.96 \pm .003	.97 \pm .004	.89 \pm .005
	EBA	.97 \pm .005	.98 \pm .004	.97 \pm .003	.97 \pm .003	.00 \pm .000	.97 \pm .005	.84 \pm .004
	COSINESIM	.85 \pm .004	.83 \pm .002	.94 \pm .003	.71 \pm .006	.68 \pm .001	.93 \pm .002	.67 \pm .006
	FOLDSEEK	.97 \pm .005						
	TM-ALIGN	.90 \pm .003						
LMS	PLASMA	.97\pm.004	.97\pm.004	.97\pm.003	.97\pm.004	.89\pm.016	.93\pm.006	.89\pm.008
	PLASMA-PF	.91 \pm .010	.68 \pm .003	.95 \pm .006	.63 \pm .013	.43 \pm .007	.77 \pm .011	.52 \pm .004

Table 5: Model performance on `test_inter` (mean \pm std over three independent seeds). Colors indicate relative performance versus TM-Align, percentage values report the associated specific relative performance difference.

Metrics	Methods	Motif			Binding Site			Active Site		
		ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN
ROC-AUC	PLASMA	.95 ^{↑21.8%} _{±.002}	.96 ^{↑23.1%} _{±.002}	.96 ^{↑23.1%} _{±.001}	.99 ^{↑13.8%} _{±.000}	.99 ^{↑13.8%} _{±.001}	.99 ^{↑13.8%} _{±.001}	.99 ^{↑5.3%} _{±.001}	.99 ^{↑5.3%} _{±.001}	.99 ^{↑5.3%} _{±.002}
	PLASMA-PF	.95 ^{↑21.8%} _{±.003}	.93 ^{↑19.2%} _{±.004}	.91 ^{↑16.7%} _{±.003}	.99 ^{↑13.8%} _{±.000}	.96 ^{↑10.3%} _{±.003}	.97 ^{↑11.5%} _{±.001}	.99 ^{↑5.3%} _{±.003}	.96 ^{↑2.1%} _{±.002}	.98 ^{↑4.3%} _{±.002}
	EBA	.87 ^{↑11.5%} _{±.005}	.88 ^{↑12.8%} _{±.003}	.44 ^{↓43.6%} _{±.002}	.97 ^{↑11.5%} _{±.001}	.97 ^{↑11.5%} _{±.002}	.40 ^{↓54.0%} _{±.005}	.99 ^{↑5.3%} _{±.003}	.99 ^{↑5.3%} _{±.003}	.43 ^{↓54.3%} _{±.005}
	Backbone	.84 ^{↑7.7%} _{±.006}	.73 ^{↓6.4%} _{±.009}	.75 ^{↓3.8%} _{±.006}	.96 ^{↑10.3%} _{±.002}	.79 ^{↓9.2%} _{±.009}	.75 ^{↓13.8%} _{±.008}	.97 ^{↑3.2%} _{±.002}	.78 ^{↓17.0%} _{±.009}	.74 ^{↓21.3%} _{±.006}
	Foldseek		.83 ^{↑6.4%} _{±.007}			.89 ^{↑2.3%} _{±.001}			.89 ^{↑5.3%} _{±.001}	
	TM-Align		.78 _{±.003}			.87 _{±.003}			.94 _{±.003}	
PR-AUC	PLASMA	.95 ^{↑14.5%} _{±.002}	.97 ^{↑16.9%} _{±.001}	.96 ^{↑15.7%} _{±.001}	.99 ^{↑8.8%} _{±.000}	.99 ^{↑8.8%} _{±.001}	.99 ^{↑8.8%} _{±.001}	.99 ^{↑3.1%} _{±.001}	.99 ^{↑3.1%} _{±.000}	.99 ^{↑3.1%} _{±.001}
	PLASMA-PF	.95 ^{↑14.5%} _{±.003}	.94 ^{↑13.3%} _{±.002}	.92 ^{↑10.8%} _{±.001}	.99 ^{↑8.8%} _{±.000}	.97 ^{↑6.6%} _{±.002}	.98 ^{↑7.7%} _{±.001}	.99 ^{↑3.1%} _{±.002}	.97 ^{↑1.0%} _{±.001}	.99 ^{↑3.1%} _{±.001}
	EBA	.90 ^{↑8.4%} _{±.004}	.90 ^{↑8.4%} _{±.004}	.45 ^{↓45.8%} _{±.004}	.98 ^{↑7.7%} _{±.001}	.98 ^{↑7.7%} _{±.001}	.42 ^{↓53.8%} _{±.004}	.99 ^{↑3.1%} _{±.002}	.99 ^{↑3.1%} _{±.002}	.43 ^{↓55.2%} _{±.006}
	Backbone	.86 ^{↑3.6%} _{±.005}	.76 ^{↓8.4%} _{±.008}	.78 ^{↓6.0%} _{±.006}	.97 ^{↑6.6%} _{±.002}	.83 ^{↓8.8%} _{±.007}	.78 ^{↓14.3%} _{±.005}	.98 ^{↑2.1%} _{±.001}	.83 ^{↓13.5%} _{±.004}	.79 ^{↓17.7%} _{±.002}
	Foldseek		.78 ^{↓6.0%} _{±.008}			.83 ^{↓8.8%} _{±.002}			.83 ^{↓13.5%} _{±.006}	
	TM-Align		.83 _{±.004}			.91 _{±.002}			.96 _{±.001}	
F1-MAX	PLASMA	.93 ^{↑32.9%} _{±.002}	.93 ^{↑32.9%} _{±.004}	.91 ^{↑30.0%} _{±.000}	.98 ^{↑16.7%} _{±.001}	.98 ^{↑16.7%} _{±.002}	.97 ^{↑15.5%} _{±.002}	.99 ^{↑10.0%} _{±.003}	.98 ^{↑8.9%} _{±.001}	.99 ^{↑10.0%} _{±.002}
	PLASMA-PF	.93 ^{↑32.9%} _{±.004}	.89 ^{↑27.1%} _{±.004}	.84 ^{↑20.0%} _{±.004}	.97 ^{↑15.5%} _{±.001}	.92 ^{↑9.5%} _{±.003}	.94 ^{↑11.9%} _{±.001}	.98 ^{↑8.9%} _{±.003}	.93 ^{↑3.3%} _{±.004}	.96 ^{↑6.7%} _{±.003}
	EBA	.80 ^{↑14.3%} _{±.005}	.81 ^{↑15.7%} _{±.003}	.00 ^{↓100.0%} _{±.000}	.94 ^{↑11.9%} _{±.001}	.93 ^{↑10.7%} _{±.002}	.00 ^{↓100.0%} _{±.000}	.97 ^{↑7.8%} _{±.003}	.97 ^{↑7.8%} _{±.003}	.00 ^{↓100.0%} _{±.000}
	Backbone	.76 ^{↑8.6%} _{±.003}	.69 ^{↓1.4%} _{±.001}	.70 ^{↓0.0%} _{±.002}	.91 ^{↑8.3%} _{±.005}	.73 ^{↓13.1%} _{±.006}	.69 ^{↓17.9%} _{±.006}	.94 ^{↑4.4%} _{±.003}	.71 ^{↓21.1%} _{±.006}	.68 ^{↓24.4%} _{±.001}
	Foldseek		.84 ^{↑20.0%} _{±.007}			.94 ^{↑11.9%} _{±.001}			.97 ^{↑7.8%} _{±.005}	
	TM-Align		.70 _{±.002}			.84 _{±.005}			.90 _{±.003}	

F FULL EXTRAPOLATION PERFORMANCE COMPARISON

This section evaluates PLASMA’s generalization capability on the `test_extra` dataset, which contains substructures never encountered during training. These experiments are crucial for assessing applicability in detecting unknown substructures. The results demonstrate that PLASMA maintains superior performance even when confronted with completely unseen substructures, achieving the highest scores for both detecting the existence of similar substructures and accurately localizing their positions for most of the cases. This robust extrapolation performance further validates that our optimal transport framework captures fundamental protein substructure similarity patterns that transcend specific training examples, making it highly valuable for analyzing newly discovered proteins and understudied organisms.

Table 6: Comprehensive motif detection results on `test_extra` dataset across seven protein representation models.

Metrics	Methods	Motif						
		PROST5	PROT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.97\pm.015	.98\pm.012	.98\pm.008	.97\pm.013	.96\pm.016	.95\pm.023	.79 \pm .022
	PLASMA-PF	.97 \pm .014	.98 \pm .010	.98 \pm .009	.93 \pm .004	.90 \pm .005	.88 \pm .039	.82 \pm .016
	EBA	.94 \pm .017	.95 \pm .009	.90 \pm .033	.92 \pm .021	.32 \pm .043	.94 \pm .016	.76 \pm .025
	COSINESIM	.84 \pm .029	.89 \pm .024	.85 \pm .019	.74 \pm .033	.79 \pm .018	.83 \pm .050	.62 \pm .080
	FOLDSEEK	.89 \pm .033						
	TM-ALIGN	.81 \pm .014						
	PLASMA	.97\pm.017	.97\pm.018	.98\pm.011	.97\pm.014	.96\pm.017	.95\pm.025	.84 \pm .014
PR-AUC	PLASMA-PF	.97 \pm .015	.97 \pm .016	.98 \pm .010	.95 \pm .005	.92 \pm .007	.88 \pm .040	.86\pm.012
	EBA	.94 \pm .018	.96 \pm .010	.91 \pm .035	.93 \pm .019	.38 \pm .014	.95 \pm .014	.80 \pm .029
	COSINESIM	.85 \pm .028	.90 \pm .017	.86 \pm .023	.77 \pm .041	.82 \pm .027	.86 \pm .036	.66 \pm .090
	FOLDSEEK	.84 \pm .031						
	TM-ALIGN	.86 \pm .020						
F1-MAX	PLASMA	.95\pm.011	.96\pm.010	.97\pm.009	.95\pm.018	.92\pm.022	.92\pm.022	.72 \pm .017
	PLASMA-PF	.93 \pm .019	.96 \pm .006	.96 \pm .013	.90 \pm .006	.84 \pm .008	.85 \pm .041	.75 \pm .017
	EBA	.88 \pm .027	.90 \pm .014	.86 \pm .035	.87 \pm .024	.00 \pm .000	.87 \pm .019	.73 \pm .008
	COSINESIM	.77 \pm .020	.82 \pm .025	.79 \pm .008	.70 \pm .014	.73 \pm .013	.77 \pm .040	.68 \pm .015
	FOLDSEEK	.91 \pm .046						
	TM-ALIGN	.76 \pm .015						
	LMS	PLASMA	.72\pm.022	.70\pm.022	.75 \pm .045	.69\pm.019	.52\pm.046	.60\pm.021
	PLASMA-PF	.62 \pm .042	.38 \pm .057	.78\pm.055	.48 \pm .074	.23 \pm .021	.44 \pm .026	.41 \pm .066

Table 7: Comprehensive binding site detection results on `test_extra` dataset across seven protein representation models.

Metrics	Methods	Binding Site						
		PROSTT5	PROT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.98\pm.009	.98 \pm .009	.99\pm.008	.98\pm.013	.98\pm.014	.98\pm.008	.92\pm.019
	PLASMA-PF	.98 \pm .008	.98 \pm .010	.99 \pm .006	.92 \pm .052	.96 \pm .012	.95 \pm .019	.87 \pm .032
	EBA	.98 \pm .013	.99\pm.009	.99 \pm .007	.97 \pm .021	.30\pm.060	.98 \pm .014	.83 \pm .072
	COSINESIM	.89 \pm .038	.86 \pm .059	.98 \pm .010	.72 \pm .060	.70 \pm .070	.94 \pm .021	.56\pm.029
	FOLDSEEK	.90 \pm .013						
	TM-ALIGN	.91 \pm .040						
PR-AUC	PLASMA	.98\pm.011	.98\pm.010	.98\pm.011	.97\pm.019	.97\pm.019	.97\pm.012	.90\pm.043
	PLASMA-PF	.98 \pm .013	.98 \pm .014	.98 \pm .012	.90 \pm .079	.95 \pm .026	.93 \pm .022	.84 \pm .078
	EBA	.98 \pm .014	.98 \pm .014	.98 \pm .012	.96 \pm .035	.28\pm.063	.97 \pm .020	.79 \pm .115
	COSINESIM	.86 \pm .076	.82 \pm .099	.96 \pm .023	.67 \pm .093	.65 \pm .118	.93 \pm .029	.49\pm.076
	FOLDSEEK	.76 \pm .065						
	TM-ALIGN	.89 \pm .064						
F1-MAX	PLASMA	.97\pm.016	.97\pm.011	.96 \pm .022	.95 \pm .030	.93 \pm .026	.96 \pm .014	.83 \pm .046
	PLASMA-PF	.96 \pm .023	.97 \pm .017	.96 \pm .027	.85 \pm .082	.90 \pm .031	.93 \pm .018	.76 \pm .073
	EBA	.96 \pm .021	.96 \pm .026	.97\pm.021	.93 \pm .049	.00\pm.000	.94 \pm .034	.73 \pm .108
	COSINESIM	.78 \pm .081	.76 \pm .089	.91 \pm .034	.62 \pm .087	.60 \pm .107	.86 \pm .046	.55\pm.092
	FOLDSEEK	.97 \pm .014						
	TM-ALIGN	.87 \pm .063						
LMS	PLASMA	.84\pm.050	.83\pm.051	.82 \pm .062	.77\pm.105	.65\pm.088	.75\pm.071	.56\pm.075
	PLASMA-PF	.79 \pm .098	.55 \pm .079	.85\pm.058	.49 \pm .082	.36 \pm .055	.65 \pm .070	.43 \pm .038

Table 8: Comprehensive active site detection results on `test_extra` dataset across seven protein representation models.

Metrics	Methods	Active Site						
		PROSTT5	PROT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC	PLASMA	.98\pm.011	.98\pm.010	.98\pm.012	.98\pm.010	.97\pm.011	.97\pm.013	.95\pm.026
	PLASMA-PF	.98 \pm .010	.98 \pm .011	.97 \pm .015	.96 \pm .006	.97 \pm .008	.97 \pm .014	.93 \pm .024
	EBA	.98 \pm .012	.98 \pm .012	.97 \pm .013	.97 \pm .012	.43\pm.066	.97 \pm .013	.91 \pm .027
	COSINESIM	.87 \pm .032	.91 \pm .011	.96 \pm .012	.79 \pm .068	.76 \pm .033	.96 \pm .013	.71\pm.012
	FOLDSEEK	.87 \pm .022						
	TM-ALIGN	.93 \pm .009						
PR-AUC	PLASMA	.97 \pm .014	.98\pm.010	.97\pm.014	.98\pm.011	.97\pm.012	.97\pm.016	.96\pm.019
	PLASMA-PF	.98\pm.013	.98 \pm .011	.97 \pm .015	.96 \pm .006	.97 \pm .009	.96 \pm .017	.95 \pm .017
	EBA	.97 \pm .013	.97 \pm .014	.97 \pm .012	.97 \pm .012	.43\pm.032	.97 \pm .014	.93 \pm .019
	COSINESIM	.90 \pm .031	.92 \pm .017	.96 \pm .016	.84 \pm .059	.80 \pm .038	.96 \pm .015	.75\pm.010
	FOLDSEEK	.81 \pm .026						
	TM-ALIGN	.94 \pm .012						
F1-MAX	PLASMA	.97\pm.012	.98\pm.013	.98\pm.013	.97\pm.011	.97\pm.011	.97\pm.015	.92 \pm .036
	PLASMA-PF	.97 \pm .015	.97 \pm .020	.97 \pm .018	.94 \pm .016	.95 \pm .012	.96 \pm .011	.89 \pm .032
	EBA	.97 \pm .014	.97 \pm .013	.97 \pm .013	.97 \pm .008	.00\pm.000	.97 \pm .020	.87 \pm .026
	COSINESIM	.83 \pm .033	.84 \pm .013	.92 \pm .020	.75 \pm .044	.71 \pm .018	.92 \pm .010	.68\pm.008
	FOLDSEEK	.96 \pm .015						
	TM-ALIGN	.90 \pm .014						
LMS	PLASMA	.89 \pm .044	.83\pm.030	.90 \pm .034	.87\pm.038	.67\pm.044	.84\pm.053	.60\pm.024
	PLASMA-PF	.90\pm.043	.70 \pm .014	.94\pm.029	.68 \pm .067	.43 \pm .032	.78 \pm .048	.50 \pm .021

Table 9: Model performance on `test_extra` (mean \pm std over three independent seeds). Colors indicate relative performance versus TM-Align, percentage values report the associated specific relative performance difference.

Metrics	Methods	Motif			Binding Site			Active Site		
		ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN	ANKH	ESM2	PROTSSN
ROC-AUC	PLASMA	.98 ^{↑21.0%} _{±.008}	.97 ^{↑19.8%} _{±.013}	.96 ^{↑18.5%} _{±.016}	.99 ^{↑8.8%} _{±.008}	.98 ^{↑7.7%} _{±.013}	.98 ^{↑7.7%} _{±.014}	.98 ^{↑5.4%} _{±.012}	.98 ^{↑5.4%} _{±.010}	.97 ^{↑4.3%} _{±.011}
	PLASMA-PF	.98 ^{↑21.0%} _{±.009}	.93 ^{↑14.8%} _{±.004}	.90 ^{↑11.1%} _{±.005}	.99 ^{↑8.8%} _{±.006}	.92 ^{↑1.1%} _{±.052}	.96 ^{↑5.5%} _{±.012}	.97 ^{↑4.3%} _{±.015}	.96 ^{↑3.2%} _{±.006}	.97 ^{↑4.3%} _{±.008}
	EBA	.90 ^{↑11.1%} _{±.033}	.92 ^{↑13.6%} _{±.021}	.32 ^{↓60.5%} _{±.043}	.99 ^{↑8.8%} _{±.007}	.97 ^{↑6.6%} _{±.021}	.30 ^{↓67.0%} _{±.060}	.97 ^{↑4.3%} _{±.013}	.97 ^{↑4.3%} _{±.012}	.43 ^{↓53.8%} _{±.066}
	Backbone	.85 ^{↑4.9%} _{±.019}	.74 ^{↑8.6%} _{±.033}	.79 ^{↑2.5%} _{±.018}	.98 ^{↑7.7%} _{±.010}	.72 ^{↓20.9%} _{±.060}	.70 ^{↓23.1%} _{±.070}	.96 ^{↑3.2%} _{±.012}	.79 ^{↑15.1%} _{±.068}	.76 ^{↑18.3%} _{±.033}
	Foldseek		.89 ^{↑9.9%} _{±.033}			.90 ^{↑1.1%} _{±.013}			.87 ^{↑6.5%} _{±.022}	
	TM-Align		.81 _{±.014}			.91 _{±.040}			.93 _{±.009}	
PR-AUC	PLASMA	.98 ^{↑14.0%} _{±.011}	.97 ^{↑12.8%} _{±.014}	.96 ^{↑11.6%} _{±.017}	.98 ^{↑10.1%} _{±.011}	.97 ^{↑9.0%} _{±.019}	.97 ^{↑9.0%} _{±.019}	.97 ^{↑3.2%} _{±.014}	.98 ^{↑4.3%} _{±.011}	.97 ^{↑3.2%} _{±.012}
	PLASMA-PF	.98 ^{↑14.0%} _{±.010}	.95 ^{↑10.5%} _{±.005}	.92 ^{↑7.0%} _{±.007}	.98 ^{↑10.1%} _{±.012}	.90 ^{↑1.1%} _{±.079}	.95 ^{↑6.7%} _{±.026}	.97 ^{↑3.2%} _{±.015}	.96 ^{↑2.1%} _{±.006}	.97 ^{↑3.2%} _{±.009}
	EBA	.91 ^{↑5.8%} _{±.035}	.93 ^{↑8.1%} _{±.019}	.38 ^{↓55.8%} _{±.014}	.98 ^{↑10.1%} _{±.012}	.96 ^{↑7.9%} _{±.035}	.28 ^{↓68.5%} _{±.063}	.97 ^{↑3.2%} _{±.012}	.97 ^{↑3.2%} _{±.012}	.43 ^{↓54.3%} _{±.032}
	Backbone	.80 ^{↓0.0%} _{±.023}	.77 ^{↑10.5%} _{±.041}	.82 ^{↑4.7%} _{±.027}	.96 ^{↑7.9%} _{±.023}	.67 ^{↓24.7%} _{±.093}	.65 ^{↓27.0%} _{±.118}	.96 ^{↑2.1%} _{±.016}	.84 ^{↑10.6%} _{±.059}	.80 ^{↑14.9%} _{±.038}
	Foldseek		.84 ^{↓2.3%} _{±.031}			.76 ^{↓14.6%} _{±.065}			.81 ^{↑13.8%} _{±.026}	
	TM-Align		.86 _{±.020}			.89 _{±.064}			.94 _{±.012}	
F1-MAX	PLASMA	.97 ^{↑27.6%} _{±.009}	.95 ^{↑25.0%} _{±.018}	.92 ^{↑21.1%} _{±.022}	.96 ^{↑10.3%} _{±.022}	.95 ^{↑9.2%} _{±.030}	.93 ^{↑6.9%} _{±.026}	.98 ^{↑8.9%} _{±.013}	.97 ^{↑7.8%} _{±.011}	.97 ^{↑7.8%} _{±.011}
	PLASMA-PF	.96 ^{↑26.3%} _{±.013}	.90 ^{↑18.4%} _{±.006}	.84 ^{↑10.5%} _{±.008}	.96 ^{↑10.3%} _{±.027}	.85 ^{↓2.3%} _{±.082}	.90 ^{↑3.4%} _{±.031}	.97 ^{↑7.8%} _{±.018}	.94 ^{↑4.4%} _{±.016}	.95 ^{↑5.6%} _{±.012}
	EBA	.86 ^{↑13.2%} _{±.035}	.87 ^{↑14.5%} _{±.024}	.00 ^{↓100.0%} _{±.000}	.97 ^{↑11.5%} _{±.021}	.93 ^{↑6.9%} _{±.049}	.00 ^{↓100.0%} _{±.000}	.97 ^{↑7.8%} _{±.013}	.97 ^{↑7.8%} _{±.008}	.00 ^{↓100.0%} _{±.000}
	Backbone	.79 ^{↑3.9%} _{±.008}	.70 ^{↑7.9%} _{±.014}	.73 ^{↑3.9%} _{±.013}	.91 ^{↑4.6%} _{±.034}	.62 ^{↓28.7%} _{±.087}	.60 ^{↓31.0%} _{±.107}	.92 ^{↑2.2%} _{±.020}	.75 ^{↓16.7%} _{±.044}	.71 ^{↓21.1%} _{±.018}
	Foldseek		.91 ^{↑19.7%} _{±.046}			.97 ^{↑11.5%} _{±.014}			.96 ^{↑6.7%} _{±.015}	
	TM-Align		.76 _{±.015}			.87 _{±.063}			.90 _{±.014}	

G ABLATION STUDY

This section analyzes the contribution of the two plan-assessor components: the local-motif loss (LML) and the weight-correction term (WC) derived from the diagonal kernel. The combined ROC-AUC and LMS results across seven protein backbones and three tasks show two clear trends.

First, both LML and WC improve PLASMA’s alignment quality. Adding LML yields consistently higher ROC-AUC, confirming that it helps the model concentrate alignment mass on the task-relevant functional substructures it is trained to detect. We also observe that LML can slightly reduce performance on `test_extra`, indicating a mild trade-off between specialization and generalization.

Second, WC is essential for ensuring stable alignment behavior, especially for the parameter-free PLASMA-PF variant. Removing WC causes a substantial performance drop on several backbones (notably ESM2 and ProtBERT), demonstrating that continuity weighting is crucial for suppressing fragmented correspondences and producing coherent alignment plans.

Overall, these results show that LML shapes the model toward identifying the desired functional motifs, while WC is indispensable for robust and stable alignment across architectures, particularly in the parameter-free setting.

Table 10: Ablation study results. Here we ablate two cases: not using the Label Matching Loss (w/o LML) and not using weight correction (w/o WC).

Task	Method	PROST5	PROT5	ANKH	ESM2	PROTSSN	TM-VEC	PROTBERT
ROC-AUC								
Motif	PLASMA	.97\pm.002	.97\pm.002	.95\pm.002	.96\pm.002	.96\pm.001	.92\pm.004	.87\pm.004
	PLASMA-PF	.94 \pm .003	.96 \pm .002	.95\pm.003	.93 \pm .004	.91 \pm .003	.87 \pm .001	.85 \pm .004
	PLASMA (w/o LML)	.95 \pm .008	.95 \pm .006	.93 \pm .004	.91 \pm .022	.89 \pm .018	.89 \pm .033	.85 \pm .004
	PLASMA (w/o WC)	.91 \pm .005	.95 \pm .004	.91 \pm .004	.87 \pm .019	.84 \pm .003	.86 \pm .006	.73 \pm .009
	PLASMA-PF (w/o WC)	.74 \pm .004	.87 \pm .002	.85 \pm .006	.44 \pm .009	.75 \pm .009	.84 \pm .007	.60 \pm .012
Binding Site	PLASMA	.99\pm.001	.99\pm.000	.99\pm.000	.99\pm.001	.99\pm.001	.96 \pm .003	.98\pm.001
	PLASMA-PF	.99\pm.001	.99\pm.001	.99\pm.000	.96 \pm .003	.97 \pm .001	.92 \pm .004	.90 \pm .003
	PLASMA (w/o LML)	.99\pm.002	.99\pm.001	.99\pm.002	.97 \pm .001	.99\pm.000	.98\pm.001	.98\pm.001
	PLASMA (w/o WC)	.99\pm.002	.99\pm.001	.99\pm.001	.98 \pm .001	.92 \pm .008	.97 \pm .002	.77 \pm .004
	PLASMA-PF (w/o WC)	.91 \pm .002	.97 \pm .001	.97 \pm .002	.49 \pm .003	.85 \pm .006	.96 \pm .003	.67 \pm .006
Active Site	PLASMA	.99\pm.001	.99\pm.001	.99\pm.001	.99\pm.001	.99\pm.002	.99\pm.003	.99\pm.004
	PLASMA-PF	.99\pm.002	.99\pm.003	.99\pm.003	.96 \pm .002	.98 \pm .002	.98 \pm .003	.94 \pm .006
	PLASMA (w/o LML)	.99\pm.001	.99\pm.000	.99\pm.001	.99\pm.005	.99\pm.000	.98 \pm .009	.99\pm.000
	PLASMA (w/o WC)	.99\pm.001	.99\pm.001	.99\pm.001	.98 \pm .000	.93 \pm .008	.99\pm.001	.81 \pm .033
	PLASMA-PF (w/o WC)	.95 \pm .002	.97 \pm .001	.98 \pm .001	.55 \pm .008	.87 \pm .005	.99\pm.001	.67 \pm .009
LMS								
Motif	PLASMA	.91\pm.007	.92\pm.001	.92\pm.002	.92\pm.005	.73\pm.013	.76\pm.005	.71 \pm .007
	PLASMA (w/o LML)	.66 \pm .135	.65 \pm .142	.92\pm.012	.77 \pm .170	.48 \pm .136	.68 \pm .167	.74\pm.012
Binding Site	PLASMA	.93\pm.002	.93\pm.003	.93\pm.004	.93\pm.003	.85 \pm .006	.86 \pm .002	.84 \pm .003
	PLASMA (w/o LML)	.87 \pm .080	.84 \pm .110	.79 \pm .081	.49 \pm .004	.88\pm.000	.90\pm.011	.89\pm.012
Active Site	PLASMA	.97\pm.004	.97\pm.004	.97\pm.003	.97\pm.004	.89 \pm .016	.93\pm.006	.89\pm.008
	PLASMA (w/o LML)	.89 \pm .080	.84 \pm .131	.91 \pm .065	.79 \pm .187	.90\pm.000	.79 \pm .143	.89\pm.007

H CASE STUDY

To provide a clearer view of the residue-level alignment patterns, we include enlarged versions of the alignment matrices corresponding to Figure 5 in the main text. These zoomed-in visualizations highlight how PLASMA identifies coherent local structural motifs across proteins with different folds, lengths, and sequence identities.

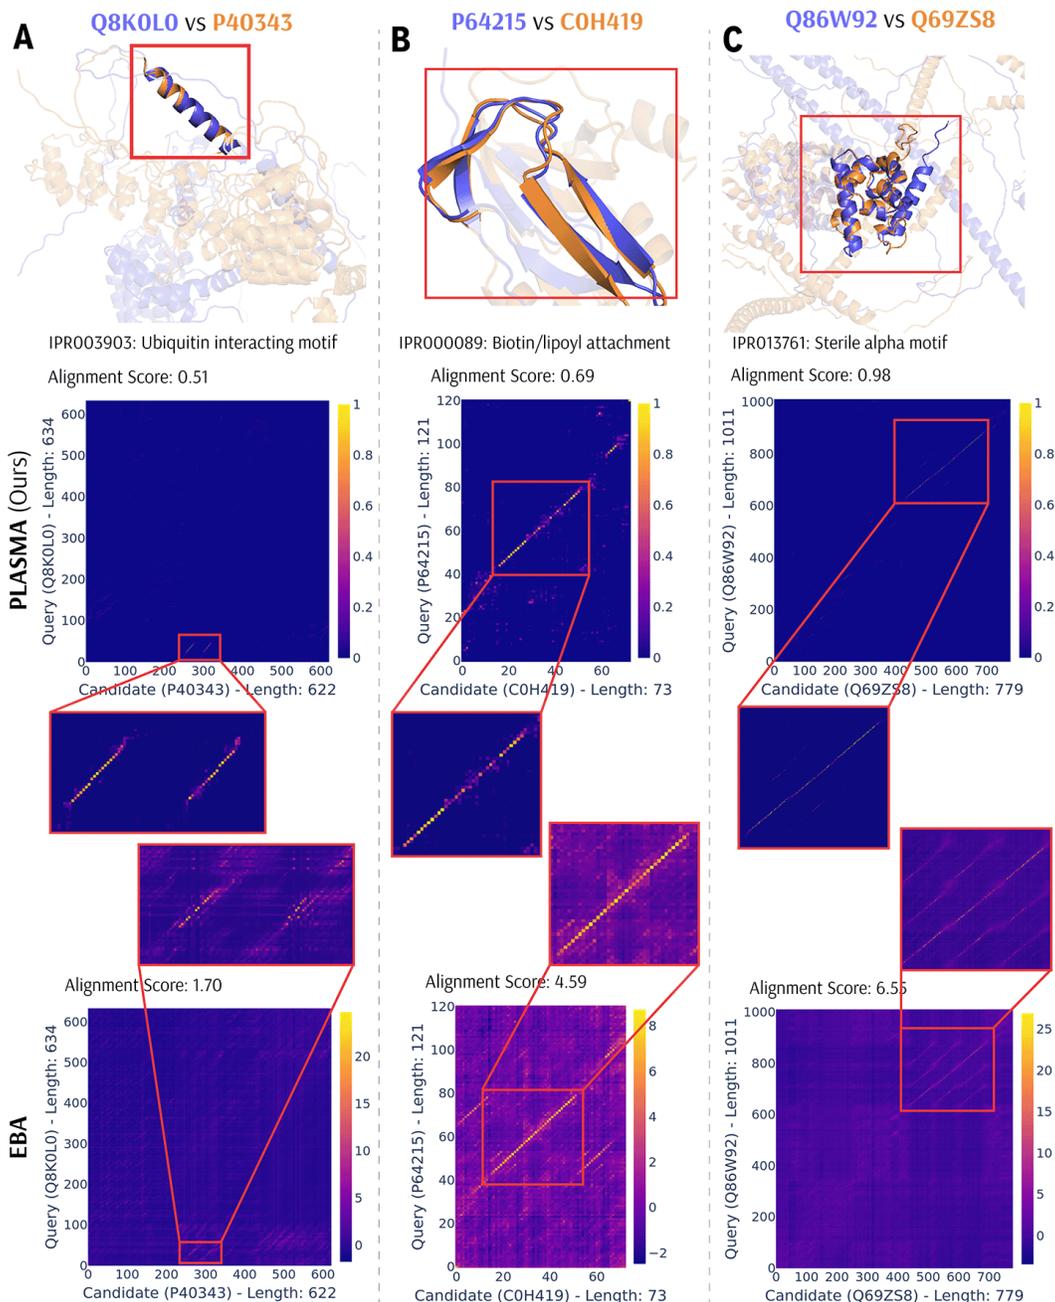


Figure 6: Representative alignment examples across three protein pairs. **A**, P40343 vs Q8K0L0. **B**, P64215 vs C0H419. **C**, Q69ZS8 vs Q86W92. Left: 3D structures with highlighted aligned regions. Center and right: alignment matrices from PLASMA and EBA with zoomed insets. This figure is the higher resolution version of Figure 5.

I ALIGNMENT MATRIX VISUALIZATIONS

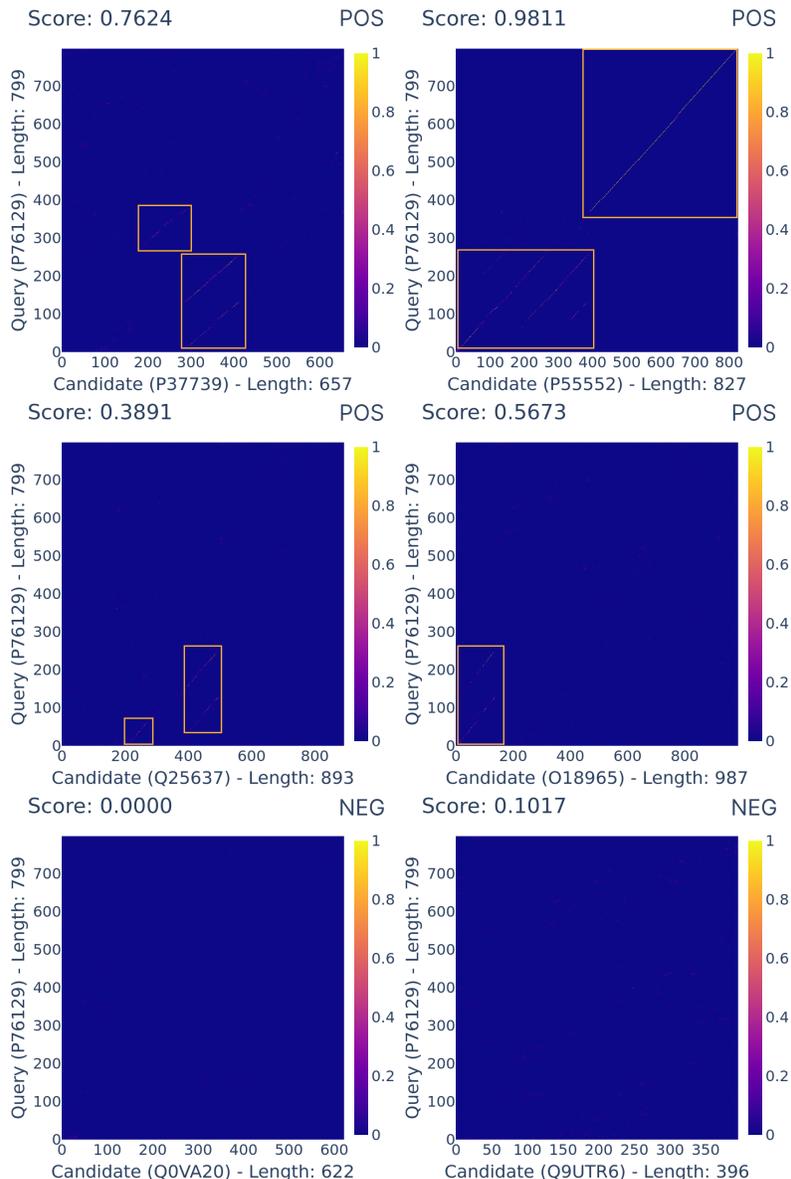


Figure 7: Representative alignment matrices comparing query protein P76129 against six candidate proteins. The visualization shows four positive pairs (POS) with shared substructures and two negative pairs (NEG) without substructure similarity. Orange regions highlight aligned substructures.

Figure 7 demonstrate PLASMA’s interpretability by showing clear patterns that correspond to different levels of substructure similarity. The matrices were generated by comparing a single query protein (InterPro ID: P76129) against six different candidate proteins, including four positive pairs sharing functional substructures and two negative pairs without similar functional substructures. The orange-highlighted regions indicate aligned substructures, where larger and more intensely colored blocks correspond to stronger and more extensive alignments. Notably, positive pairs exhibit prominent diagonal patterns reflecting substructure correspondences, while negative pairs show minimal coherent structures and low alignment scores. This visualization validates that PLASMA’s alignment scores accurately reflect the underlying biological relationships between protein substructures.

J TEMPERATURE PARAMETER ANALYSIS

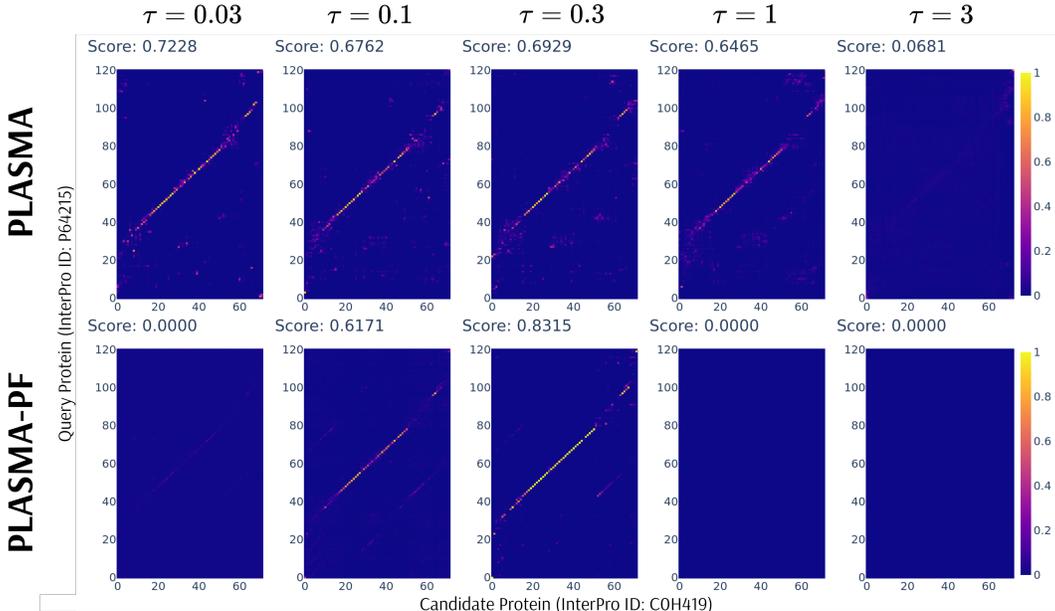


Figure 8: Effect of Sinkhorn temperature parameter τ on alignment matrix and score for both PLASMA and PLASMA-PF variants.

Figure 8 illustrates how the Sinkhorn temperature parameter τ impacts the alignment matrix in both PLASMA variants. The supervised PLASMA variant demonstrates greater stability and maintains meaningful alignment patterns across a wider range of temperature settings compared to PLASMA-PF, highlighting the robustness benefits of end-to-end training.

K PERFORMANCE EVALUATION AT DIFFERENT STRUCTURAL SIMILARITY THRESHOLD

We report the detailed values of the performance at different TM-score thresholds visualized in Figure 3. PLASMA consistently outperforms other baseline methods, especially in low similarity settings (*e.g.*, TM-score ≤ 0.5 and TM-score ≤ 0.3).

Table 11: Numerical results of the ROC-AUC Performance at different TM-Align thresholds.

Task	TM Score	PLASMA	PLASMA-PF	EBA	COSINESIM	TM-ALIGN	FOLDSEEK
Motif	≤ 1.0	.96 \pm .002	.95 \pm .002	.87 \pm .003	.84 \pm .004	.78 \pm .004	.83 \pm .004
	≤ 0.7	.95 \pm .002	.94 \pm .002	.84 \pm .004	.81 \pm .004	.73 \pm .005	.81 \pm .004
	≤ 0.5	.93 \pm .003	.93 \pm .003	.81 \pm .005	.78 \pm .005	.66 \pm .006	.79 \pm .005
	≤ 0.3	.92 \pm .004	.91 \pm .004	.74 \pm .006	.73 \pm .006	.58 \pm .007	.74 \pm .006
Binding Site	≤ 1.0	.99 \pm .001	.99 \pm .001	.97 \pm .002	.96 \pm .002	.87 \pm .003	.90 \pm .003
	≤ 0.7	.99 \pm .001	.98 \pm .002	.95 \pm .003	.93 \pm .003	.76 \pm .006	.88 \pm .004
	≤ 0.5	.98 \pm .002	.97 \pm .003	.93 \pm .004	.91 \pm .004	.62 \pm .007	.85 \pm .006
	≤ 0.3	.97 \pm .004	.96 \pm .004	.89 \pm .007	.88 \pm .007	.45 \pm .010	.80 \pm .009
Active Site	≤ 1.0	.99 \pm .001	.99 \pm .001	.99 \pm .001	.97 \pm .001	.94 \pm .002	.92 \pm .003
	≤ 0.7	.99 \pm .001	.98 \pm .002	.97 \pm .002	.95 \pm .003	.88 \pm .005	.91 \pm .004
	≤ 0.5	.98 \pm .003	.96 \pm .004	.95 \pm .004	.92 \pm .005	.76 \pm .008	.89 \pm .006
	≤ 0.3	.96 \pm .007	.90 \pm .010	.89 \pm .011	.83 \pm .013	.59 \pm .016	.84 \pm .013

L SEQUENCE SIMILARITY ANALYSIS

To further examine whether PLASMA’s alignment performance is influenced by unintended global similarity, we analyze how PLASMA’s alignment score relates to the sequence similarity of the aligned residues. Same as before, we define sequence similarity as the percentage of aligned residue pairs that share the same amino acid type.

Figure 9 presents the distribution of alignment scores and sequence-similarity values across all test pairs. The results show that high alignment scores typically coincide with high alignment coverage rather than high sequence similarity. Many correctly aligned substructures exhibit low sequence similarity despite high PLASMA scores, indicating that the method is driven by shared local 3D geometry rather than residue identity. For negative test pairs, the sequence-similarity values appear highly dispersed, which arises from their extremely low alignment coverage; with very few aligned residue pairs, the resulting sequence-similarity statistic becomes unstable and effectively uninformative. The upper-right region of the plot remains sparse, reflecting our dataset construction protocol that limits the global sequence identity of all test proteins to below 50%.

Overall, this analysis demonstrates that strong PLASMA alignment scores do not depend on high sequence similarity. The method therefore does not rely on global homology signals and is not affected by unintended data leakage.

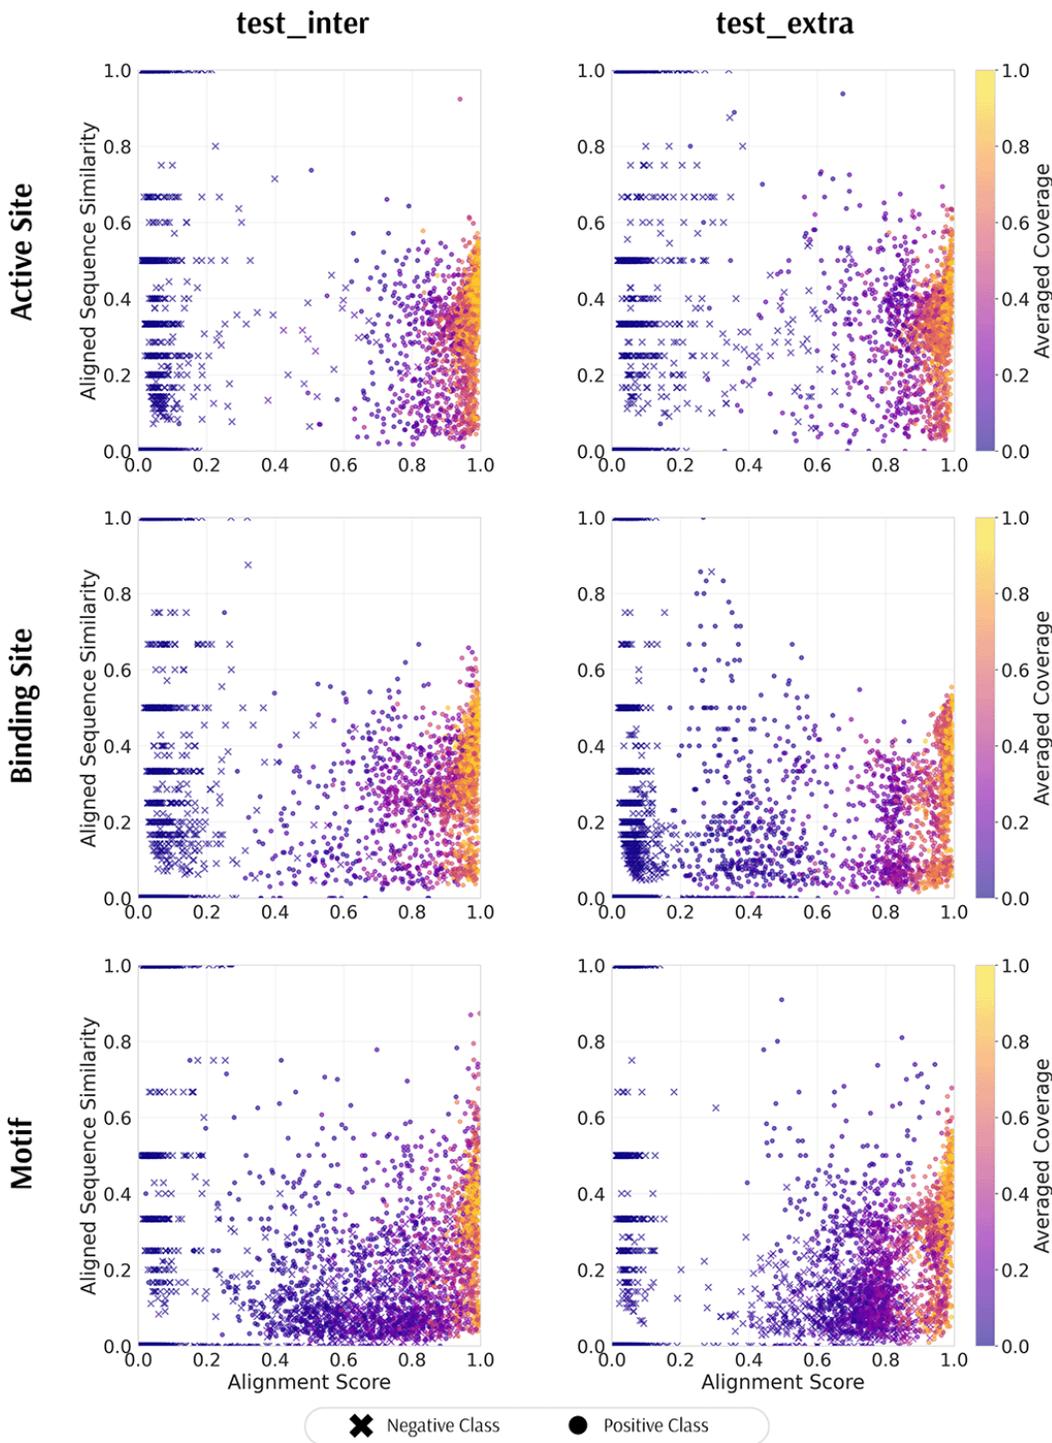
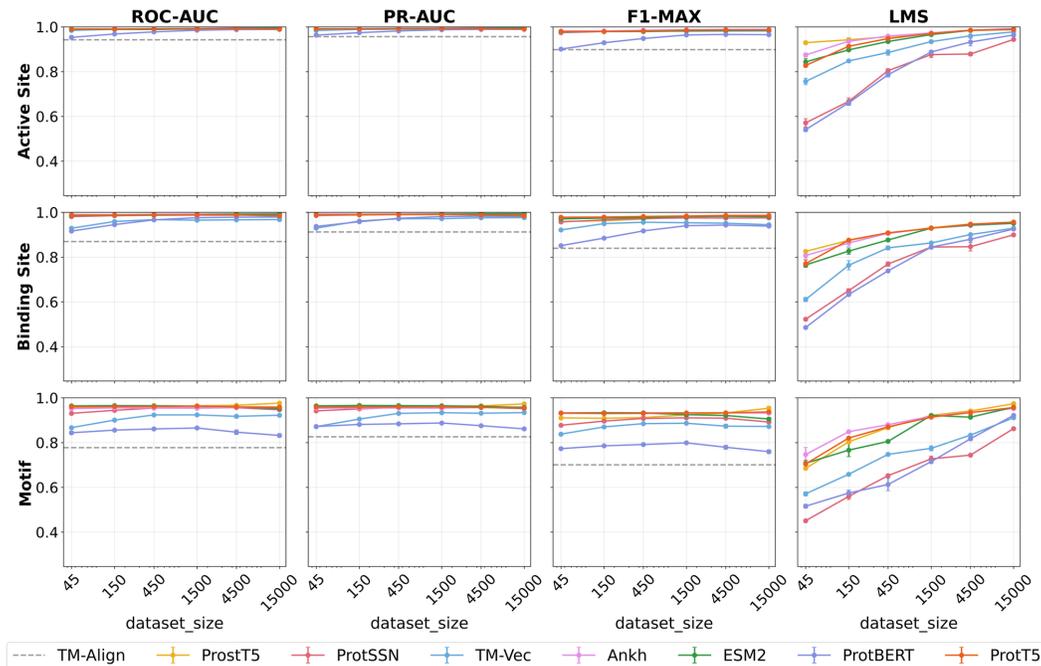
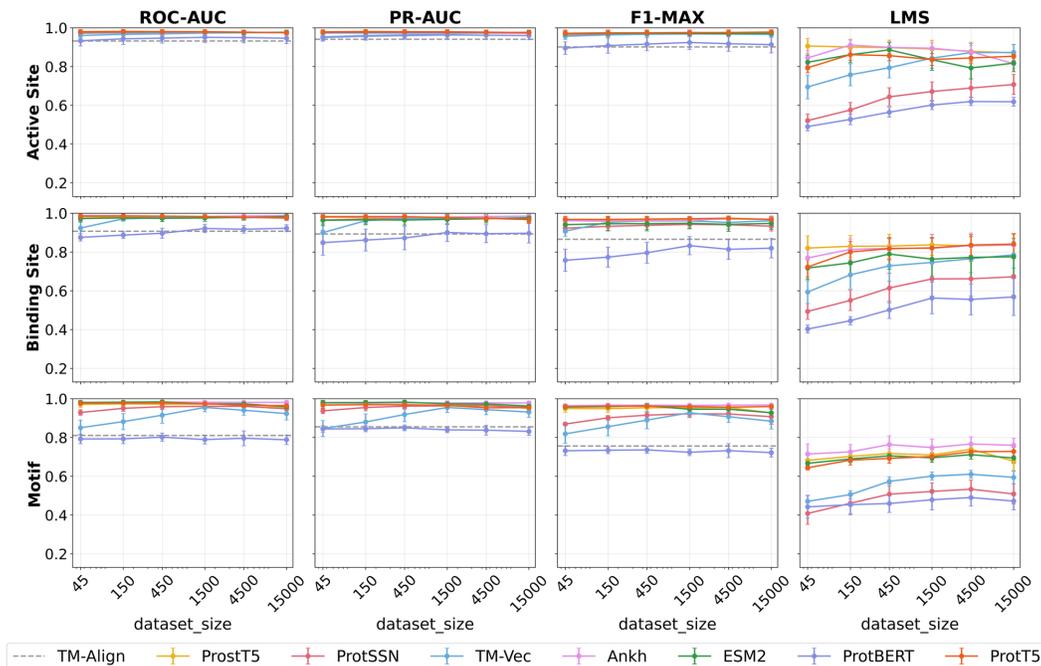


Figure 9: Sequence-similarity patterns of aligned substructures. Each panel shows how the aligned-sequence similarity varies with PLASMA’s alignment score, colored by the averaged coverage between the query and candidate proteins. These plots illustrate that high alignment scores do not simply arise from high sequence similarity; the alignment quality is driven by structural correspondence rather than sequence identity. All results use embeddings from Ankh.

M HYPERPARAMETER ANALYSIS



(a) test_inter



(b) test_extra

Figure 10: Performance vs dataset fraction. PLASMA demonstrates high performance in predicting the existence of substructure similarities even with minimal training data (45 samples), and, in most cases, this ability remains stable when the dataset size increases. However, the LMS of PLASMA noticeably improves as dataset size increases, indicating that training is important for predicting the precise local of similar substructures.

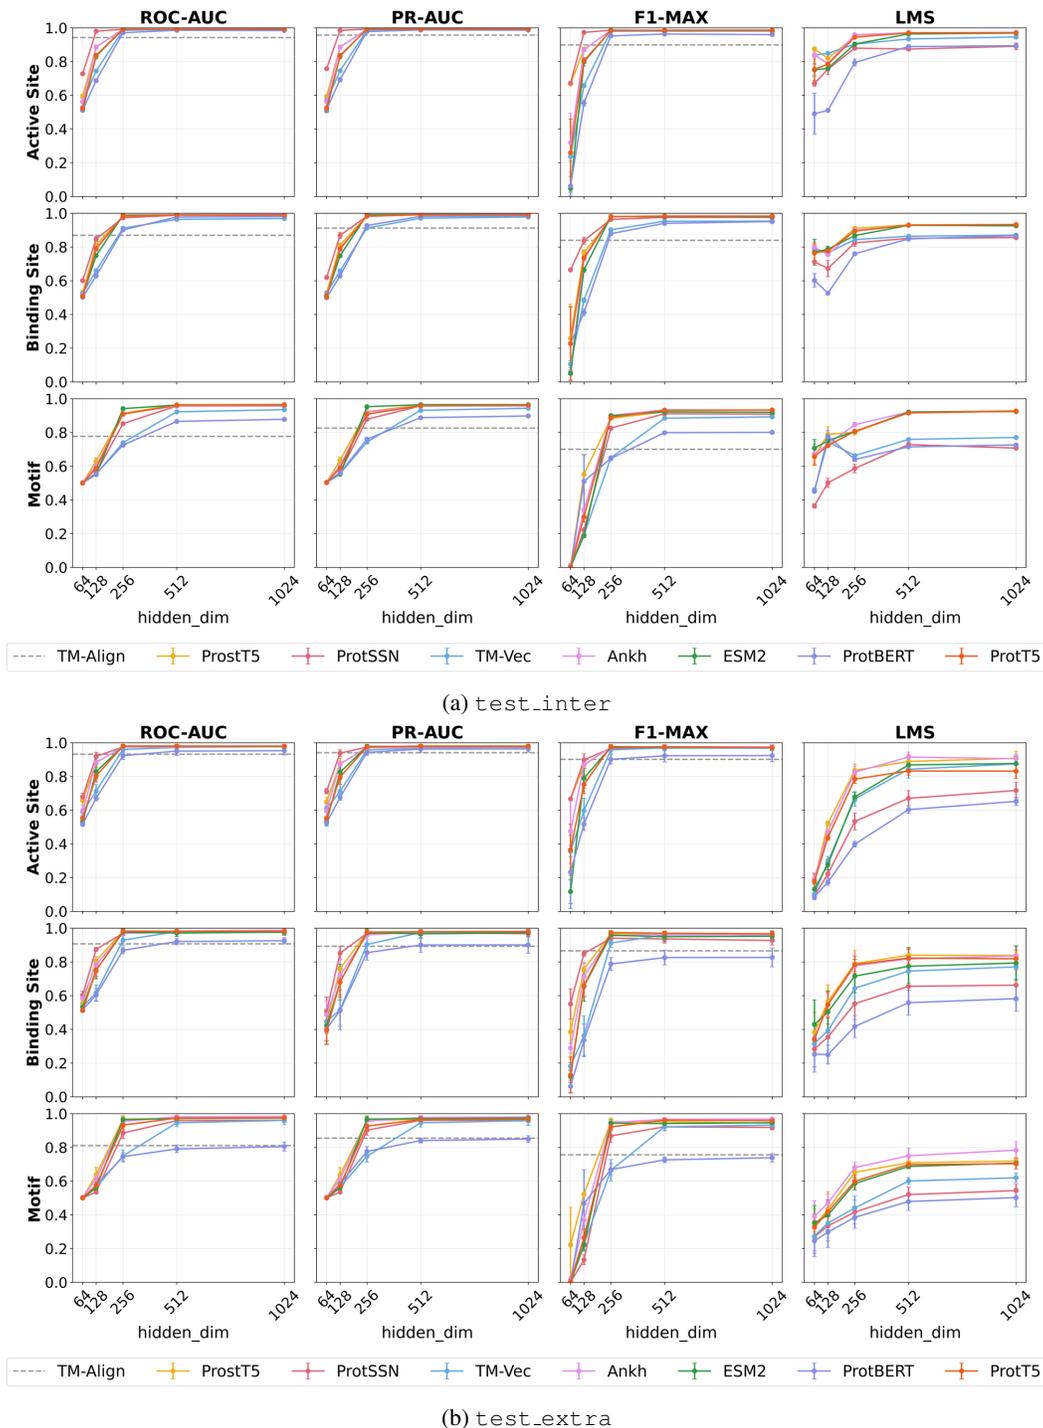


Figure 11: Performance vs hidden dimension size of the siamese network. While PLASMA’s performance remains stable when the hidden dimension size is greater than 256, it would significantly drop when the hidden dimension size is less than this number.

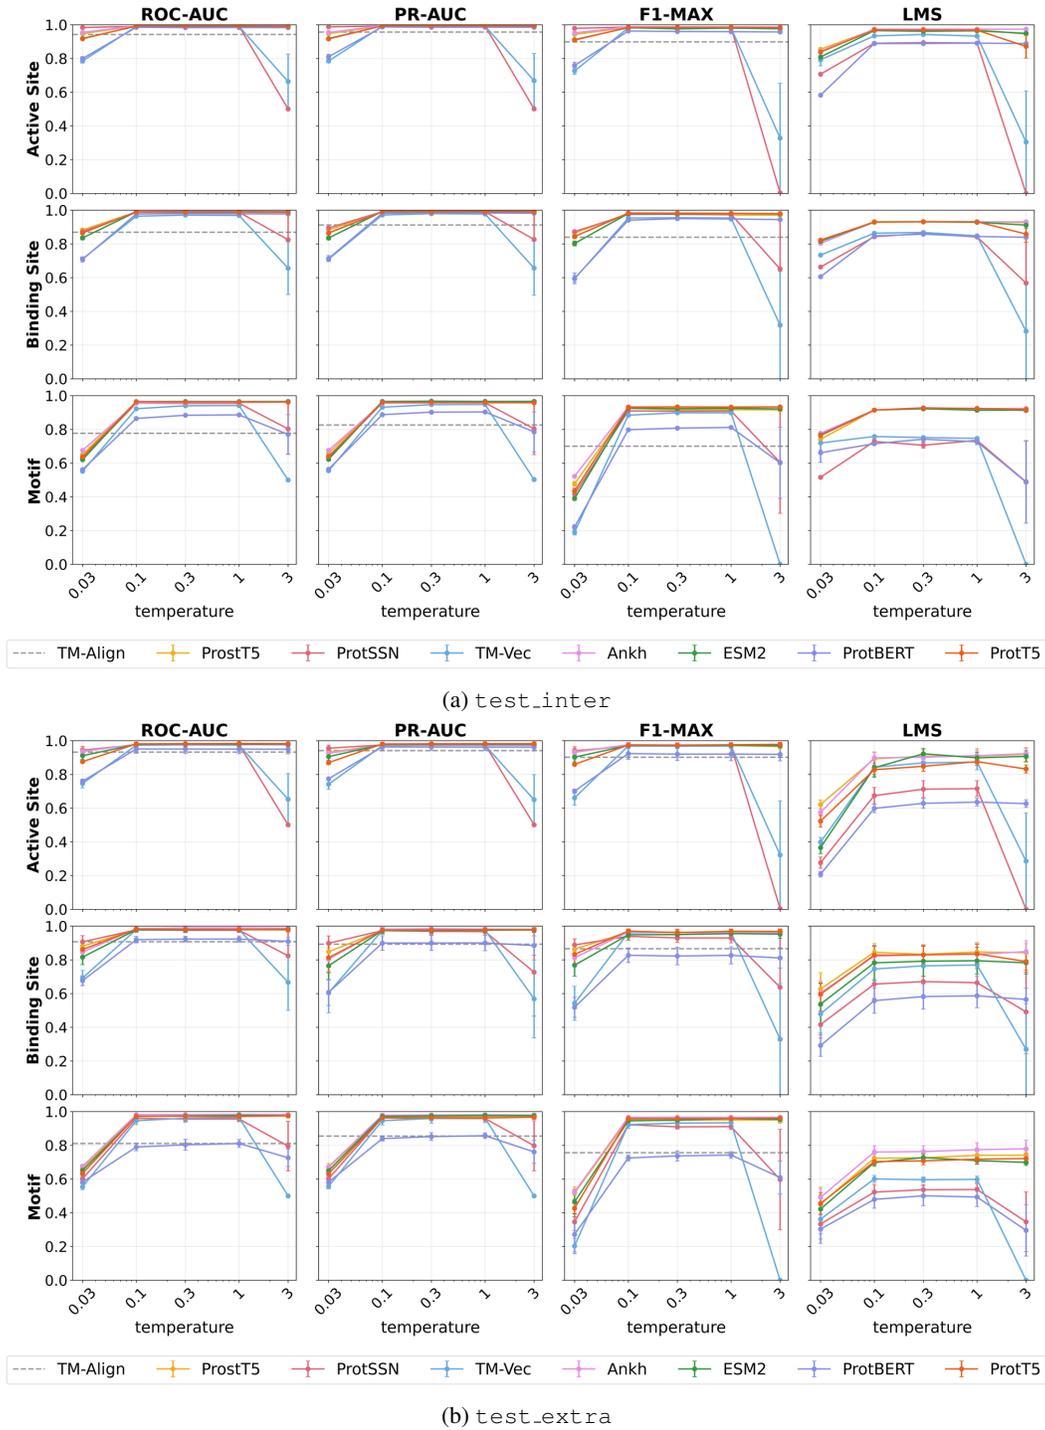


Figure 12: Performance vs Sinkhorn temperature (τ). PLASMA’s performance remains stably high within the 0.1–1 range, but when out of this range, PLASMA’s performance noticeably drops.

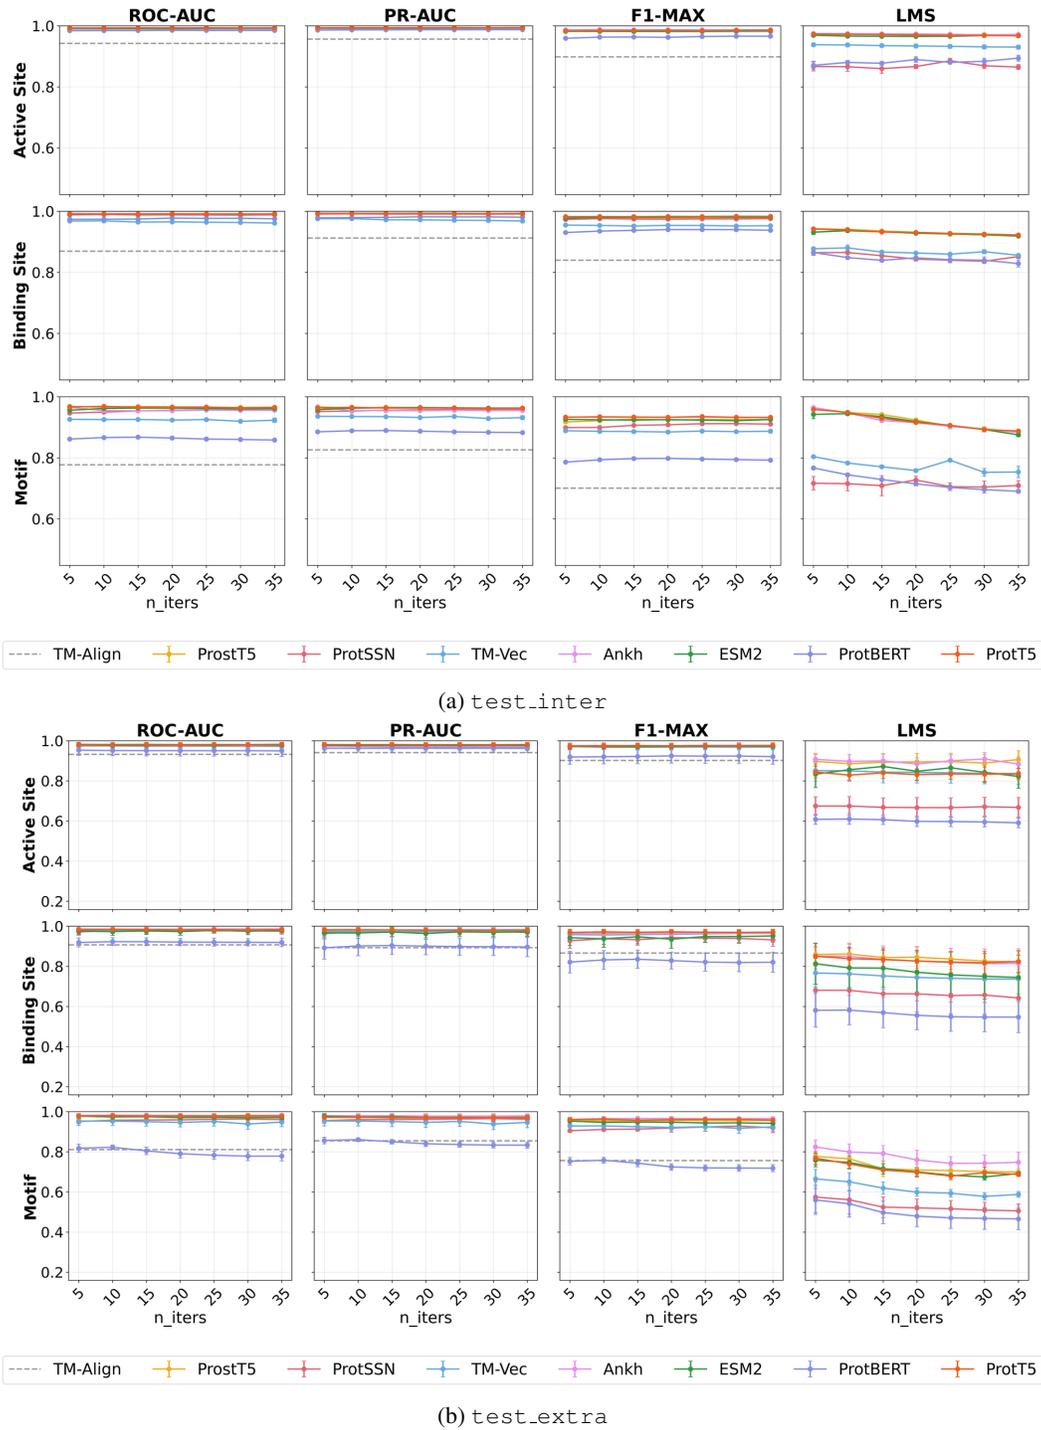
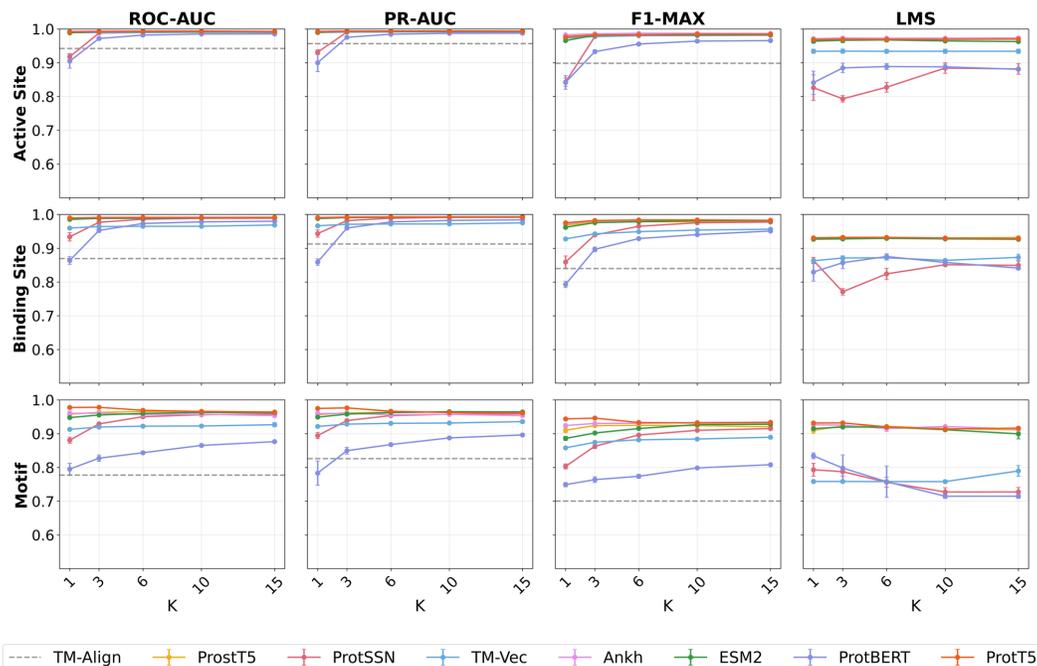
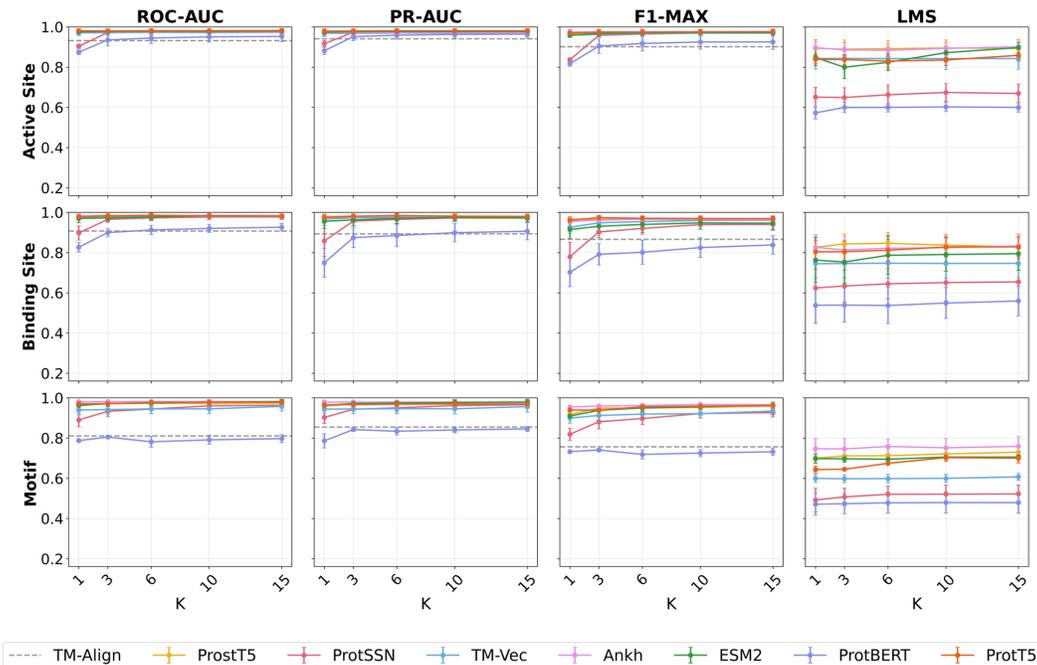


Figure 13: Performance vs number of Sinkhorn iterations T . In most cases, PLASMA’s performance is insensitive of the setting of T , but for analyzing motifs, we can see a subtle decreasing trend as the number of iteration increases.



(a) test_inter



(b) test_extra

Figure 14: Performance vs the kernel size of the diagonal convolution (k). For interpolation tasks and in particular when using PROTSSN, PROTBERT, or TM-VEC as the backbone, there is a trade-off between detecting the existence of substructure similarities and predicting the precise location of similar regions—the former prefers higher k while the latter prefers lower k . However, for other cases, PLASMA demonstrates stable performance regardless the choice of k .

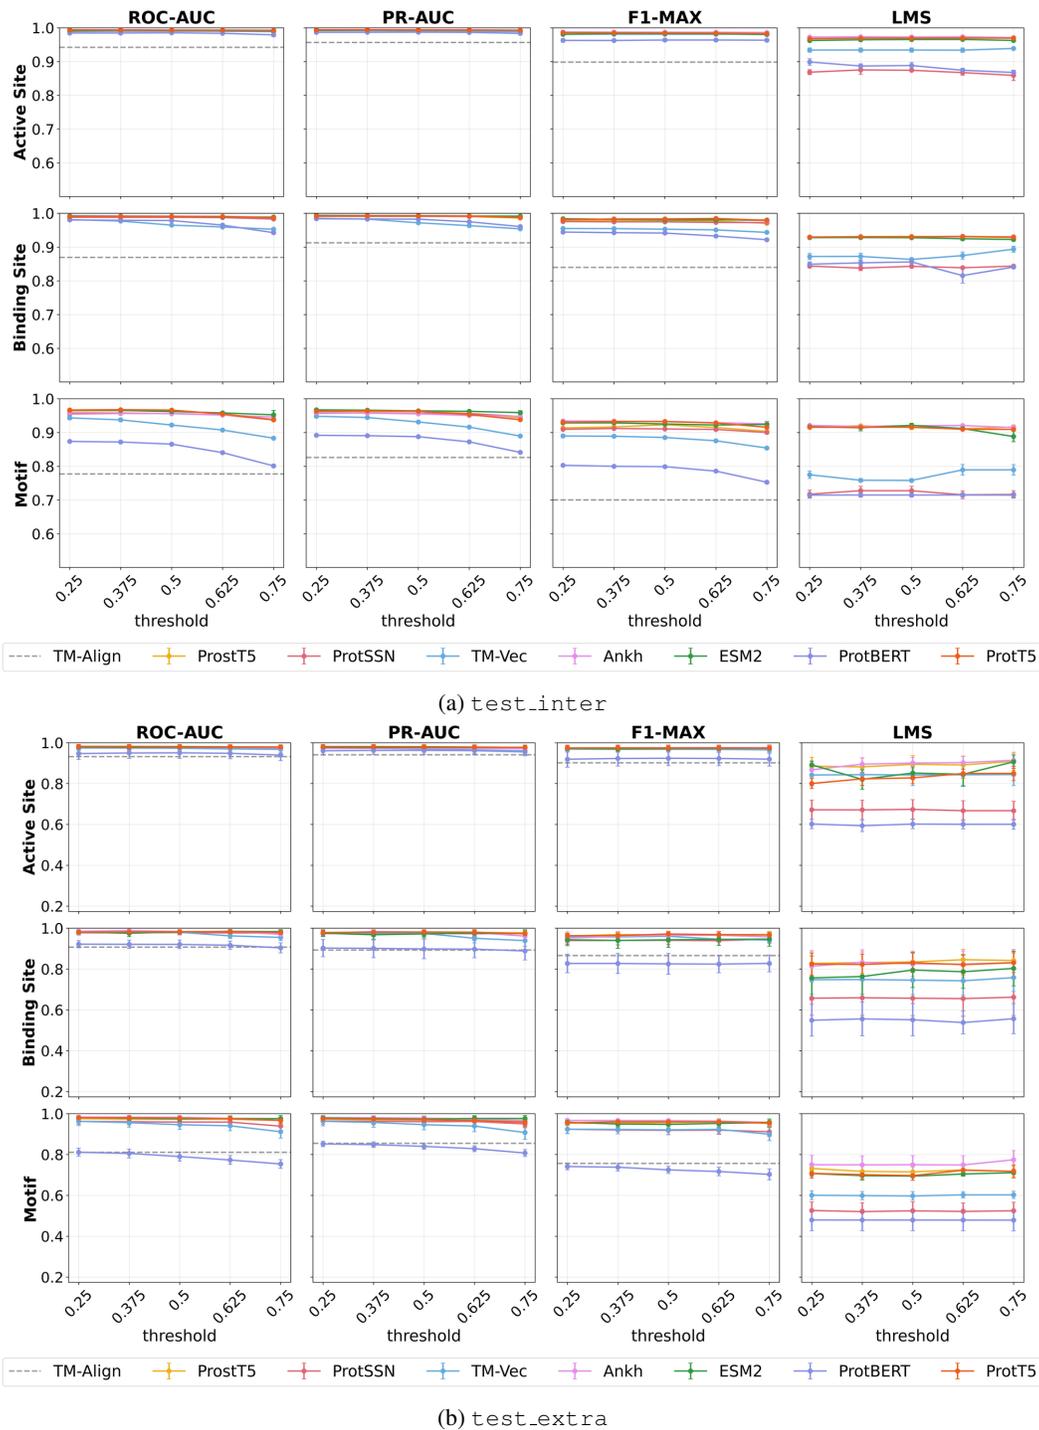


Figure 15: Performance vs residue matching threshold (ρ). PLASMA’s performance remains stable overall when choosing different ρ values, but for some backbone choices, such as TM-VEC and PROTBERT, PLASMA shows a slight preference over lower ρ values.

N FURTHER ALIGNMENT MATRIX VISUALIZATIONS

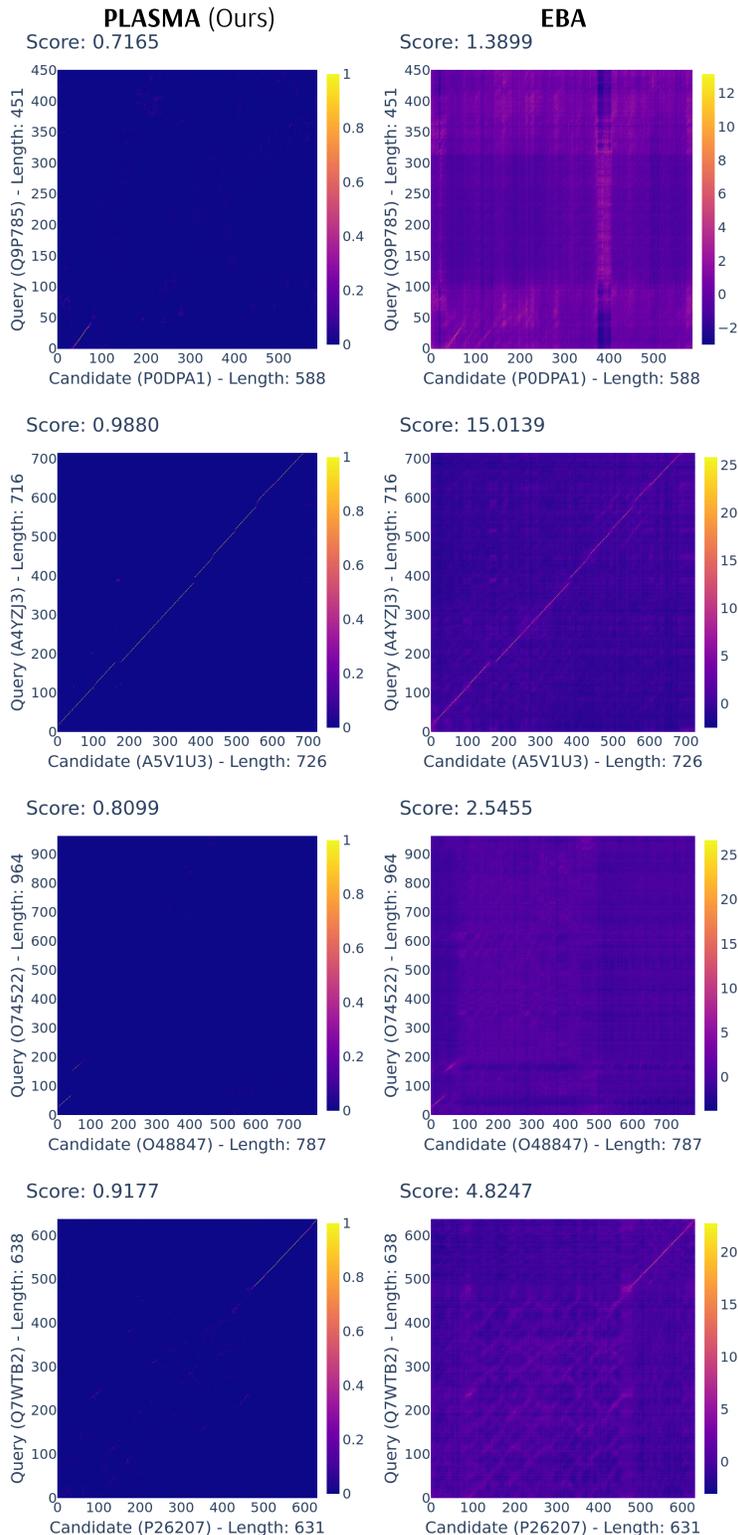


Figure 16: Alignment matrix visualizations of random positive pairs from `test_inter`. (Part 1)

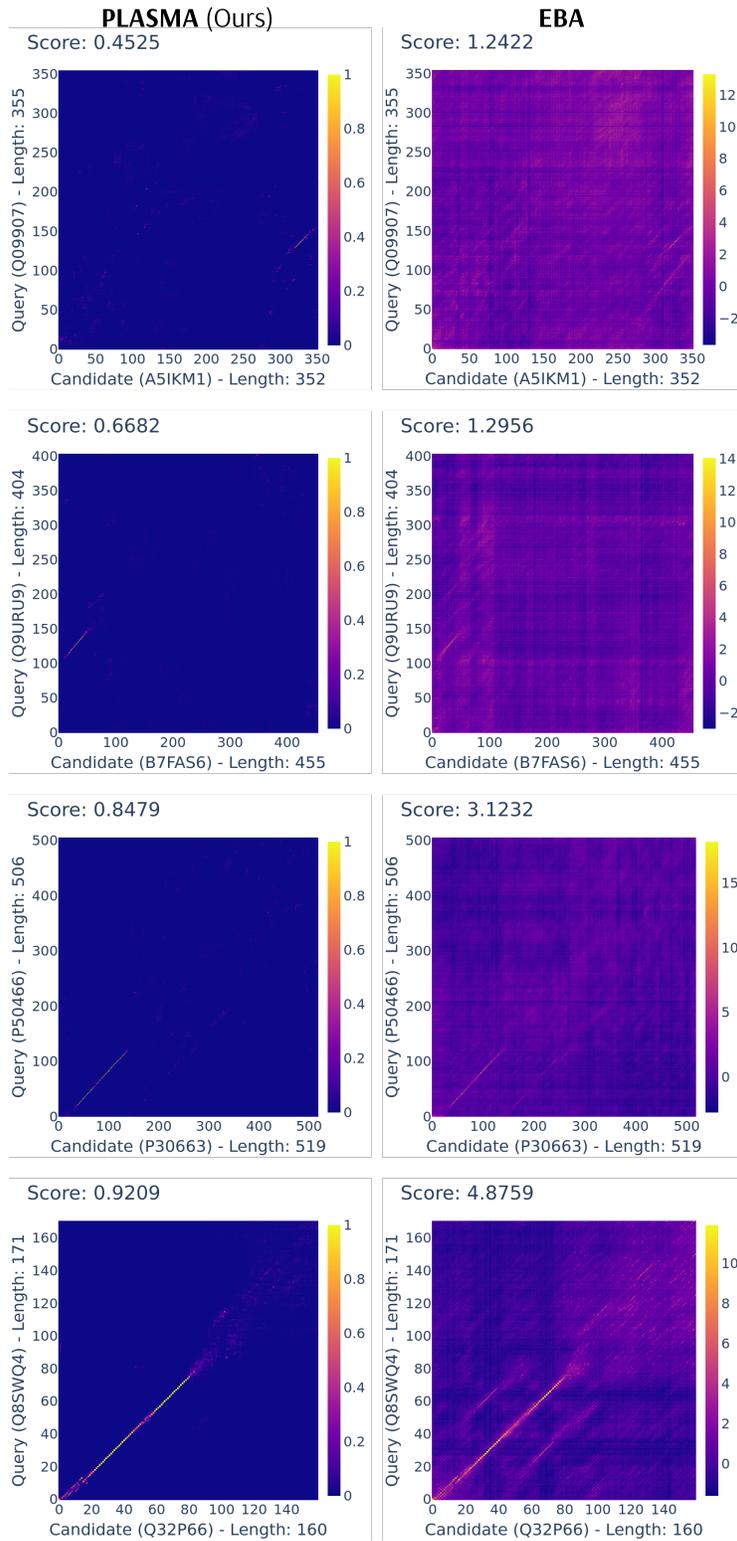


Figure 17: Alignment matrix visualizations of random positive pairs from `test_inter`. (Part 2)

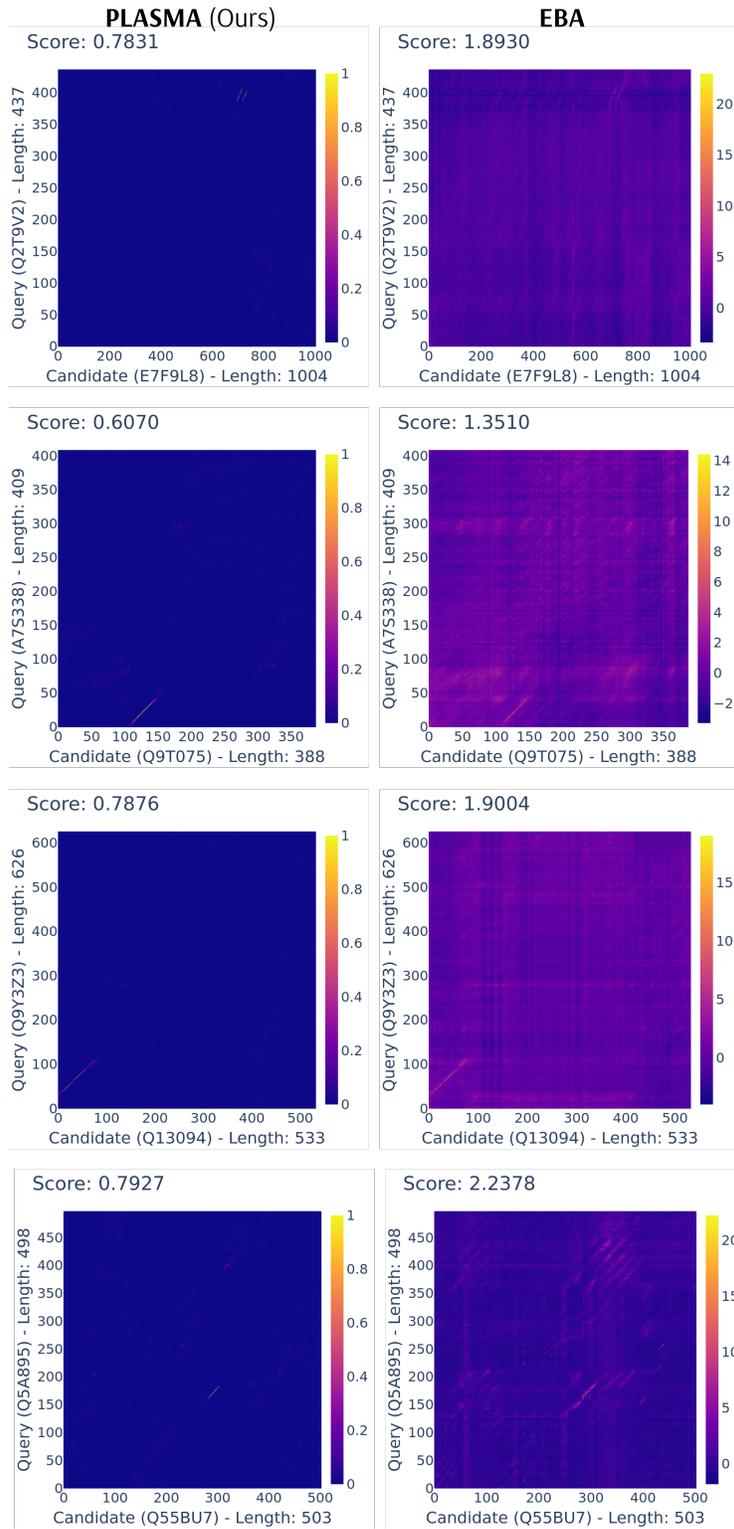


Figure 18: Alignment matrix visualizations of random positive pairs from `test_inter`. (Part 3)