Novelty Assessment Report

Paper: Unbalanced Soft-Matching Distance For Neural Representational Comparison With Partial Unit Correspondence

PDF URL: https://openreview.net/pdf?id=peMOI4RjmJ

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Representational similarity metrics typically force all units to be matched, making them susceptible to noise and outliers common in neural representations. We extend the soft-matching distance to a partial optimal transport setting that allows some neurons to remain unmatched, yielding rotation-sensitive but robust correspondences. This unbalanced soft-matching distance provides theoretical advantages---relaxing strict mass conservation while maintaining interpretable transport costs---and practical benefits through efficient neuron ranking in terms of cross-network alignment without costly iterative recomputation. In simulations, it preserves correct matches under outliers and reliably selects the correct model in noise-corrupted identification tasks. On fMRI data, it automatically excludes low-reliability voxels and produces voxel rankings by alignment quality that closely match computationally expensive brute-force approaches. It achieves higher alignment precision across homologous brain areas than standard soft-matching, which is forced to match all units regardless of quality. In deep networks, highly matched units exhibit similar maximally exciting images, while unmatched units show divergent patterns. This ability to partition by match quality enables focused analyses, \emph{e.g.,} testing whether networks have privileged axes even within their most aligned subpopulations. Overall, unbalanced soft-matching provides a principled and practical method for representational comparison under partial correspondence.

Core Task Landscape

This paper addresses: **Comparing Neural Representations with Partial Unit Correspondence**

A total of **19 papers** were analyzed and organized into a taxonomy with **15 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Optimal Transport and Matching Methods for Representation Comparison**
- **Representation Learning and Convergence Across Networks**
- **Neural Coding and Representational Mechanisms in Biological Systems**
- **Decoding Neural Representations Through Statistics, Intervention, and Behavior**

Complete Taxonomy Tree

- Comparing Neural Representations with Partial Unit Correspondence Survey Taxonomy
- Optimal Transport and Matching Methods for Representation Comparison
  - Partial and Unbalanced Optimal Transport for Neural Comparison ★ (1 papers)
  - [0] Unbalanced Soft-Matching Distance For Neural Representational Comparison With Partial Unit Correspondence (Anon et al., 2026) View paper
  - Partial-to-Partial Point Cloud Registration (2 papers)
  - [5] Prnet: Self-supervised learning for partial-to-partial registration (Yue Wang, 2019) View paper
  - [14] RORNet: Partial-to-Partial Registration Network With Reliable Overlapping Representations (Yue Wu, 2023) View paper
  - Hierarchical and Transfer-Based Correspondence Matching (2 papers)
  - [12] Learning to align the source code to the compiled object code (Dor Levy, 2017) View paper
  - [13] Matching neural paths: transfer from recognition to correspondence search (Nikolay Savinov, 2017) View paper
  - Neural Network Approaches to Graph and Structural Matching (2 papers)
  - [11] Issues of representation in neural networks (E Bienenstock, 1991) View paper
  - [18] Inexact matching using neural networks (J. Feng, 1994) View paper
- Representation Learning and Convergence Across Networks
  - Convergent Learning and Representation Similarity Across Initializations (2 papers)
  - [4] Towards understanding learning representations: To what extent do different neural networks learn the same representation (Wang Li-wei, 2018) View paper
  - [8] Convergent learning: Do different neural networks learn the same representations? (Li, 2015) View paper
  - Multi-View and Incomplete Data Representation Learning (1 papers)
  - [7] Contrastive and adversarial regularized multi-level representation learning for incomplete multi-view clustering. (Haiyue Wang, 2024) View paper
  - Neural Representations in Continuous and Implicit Domains (2 papers)
  - [16] EVAN: Evolutional Video Streaming Adaptation via Neural Representation (Mufan Liu, 2024) View paper
  - [19] Implicit Graphon Neural Representation (Xia Xinyue, 2022) View paper
- Neural Coding and Representational Mechanisms in Biological Systems
  - Spatial and Semantic Neural Representations in Hippocampus and Entorhinal Cortex (1 papers)
  - [1] A unified neural representation model for spatial and semantic computations (Tatsuya HAGA, 2023) View paper
  - Neural Codes for Sensory Stimuli and Written Symbols (1 papers)

- [6] The neural code for written words: a proposal (S. Dehaene, 2005) View paper
- Conditioned and Unconditioned Stimulus Representations in Amygdala (1 papers)
- [2] Neural representations of unconditioned stimuli in basolateral amygdala mediate innate and learned responses (Felicity Gore, 2015) View paper
- Representational Geometry and Incomplete Stimulus Encoding (1 papers)
- [10] Representational geometry of incomplete faces in macaque face patches. (Dongyuan Li, 2023) View paper

- Decoding Neural Representations Through Statistics, Intervention, and Behavior
  - Multisensory Integration and Receptive Field Alignment (1 papers)
  - [9] Multisensory integration: current issues from the perspective of the single neuron (Barry E. Stein, 2008) View paper
  - Sensory Perception Decoding via Statistics and Intervention (1 papers)
  - [3] Cracking the neural code for sensory perception by combining statistics, intervention, and behavior (Stefano Panzeri, 2017) View paper
  - Memory Suppression Through Partial Neural Activation (1 papers)
  - [15] Briefly cuing memories leads to suppression of their neural representations (Jordan Poppenk, 2014) View paper
  - Foundations of Neuronal Representations in Motor Control (1 papers)
  - [17] Foundations of Neuronal Representations (Justin C. Sanchez, 2022) View paper

Narrative

Core task: comparing neural representations with partial unit correspondence. This field addresses the challenge of aligning and comparing learned or biological neural representations when units (neurons or features) do not correspond one-to-one across systems. The taxonomy organizes work into four main branches. Optimal Transport and Matching Methods develop algorithmic frameworks—often rooted in optimal transport or graph matching—to find correspondences between representations even when they are incomplete or unbalanced. Representation Learning and Convergence Across Networks examines how different architectures or training regimes produce similar or divergent internal codes, exploring questions of convergence and generalization. Neural Coding and Representational Mechanisms in Biological Systems investigates how real neurons encode information, including population codes and sensory integration. Decoding Neural Representations Through Statistics, Intervention, and Behavior focuses on extracting interpretable structure from neural activity via statistical analysis, causal interventions, or behavioral readouts. Together, these branches span computational, theoretical, and empirical perspectives on understanding and comparing neural codes.

A particularly active line of work within Optimal Transport and Matching Methods tackles partial and unbalanced scenarios where not all units have counterparts, a setting that arises naturally when comparing networks of different sizes or biological recordings with missing data. Unbalanced Soft Matching[0] contributes to this direction by proposing methods that relax strict one-to-one constraints, allowing flexible alignment even when correspondence is incomplete. This contrasts with earlier exact matching approaches like Inexact Neural Matching[18], which assumed more rigid structure, and complements recent geometric methods such as PRNet Partial Registration[5] and RORNet Partial Registration[14] that handle partial overlaps in spatial or feature domains. Meanwhile, works like Understanding Learning Representations[4] and Convergent Learning[8] explore whether different training procedures yield aligned codes, raising questions about when and why representations converge. Situating the original paper within this landscape, Unbalanced Soft Matching[0] fits squarely in the optimal transport branch, emphasizing robustness to imbalance and offering a principled framework that bridges classical matching theory with modern representation comparison challenges.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

The original leaf focuses on optimal transport extensions (partial, unbalanced) for aligning neural representations when units may not fully correspond. Siblings address partial correspondence in different domains: hierarchical semantic matching for object parts, neural architectures for graph/structural alignment, and point cloud registration. All share the challenge of incomplete or asymmetric matching, but differ in mathematical formulation (optimal transport vs. graph matching vs. geometric registration) and application domain (neural units vs. semantic parts vs. 3D geometry).

Similarities: - All four subtopics handle scenarios where complete one-to-one correspondence is unavailable or inappropriate - Each addresses alignment or matching problems with partial, incomplete, or unbalanced data structures - Methods must reason about which elements should match and which remain unmatched

Differences: - Original leaf uses optimal transport theory with mass constraints; siblings use hierarchical features, neural graph networks, or geometric registration frameworks - Original targets neural network unit comparison; hierarchical methods target semantic object parts; graph matching targets abstract structures; point cloud methods target 3D geometric data - Original explicitly models unbalanced mass and transport costs; siblings focus on structural consistency (graphs), semantic hierarchy (parts), or geometric overlap (point clouds) - Hierarchical methods leverage pretrained recognition models; original and point cloud methods are often self-supervised or unsupervised; graph matching uses specialized neural architectures

Suggested Search Directions: - Hybrid methods combining optimal transport with hierarchical semantic features for multi-level neural comparison - Graph-structured optimal transport for neural architectures with explicit connectivity patterns - Unbalanced transport formulations adapted to geometric point cloud alignment problems

Sibling Subtopics

- **Hierarchical and Transfer-Based Correspondence Matching** (leaves: 1, papers: 2)
- Scope: Methods using hierarchical semantic features or transfer learning from recognition tasks to establish low-level part correspondences.
- Exclude: Excludes direct supervised correspondence learning or non-hierarchical matching; see other matching or representation learning categories.
- **Neural Network Approaches to Graph and Structural Matching** (leaves: 1, papers: 2)
- Scope: Neural network architectures designed for inexact graph matching or structural alignment problems with partial consistency.
- Exclude: Excludes point cloud or image-based matching; see partial registration or hierarchical correspondence categories.
- **Partial-to-Partial Point Cloud Registration** (leaves: 1, papers: 2)
- Scope: Self-supervised or learning-based methods for aligning partially overlapping point clouds with incomplete correspondence.
- Exclude: Excludes full-shape registration or non-geometric alignment tasks; see neural representation comparison or semantic alignment.

Contributions Analysis

Overall novelty summary. The paper proposes an unbalanced soft-matching distance that extends optimal transport to allow partial neuron correspondences, addressing robustness to outliers and noise in neural representation comparison. It resides in the 'Partial and

Unbalanced Optimal Transport for Neural Comparison' leaf, which currently contains only this paper as a sibling. This indicates a relatively sparse research direction within the broader optimal transport branch, suggesting the work occupies a niche position in the taxonomy where explicit unbalanced transport formulations for neural alignment are not yet densely populated.

The taxonomy reveals neighboring leaves focused on partial point cloud registration (PRNet, RORNet) and hierarchical correspondence matching, which address partial overlap in geometric or semantic domains but do not explicitly formulate unbalanced transport for neural units. The broader 'Representation Learning and Convergence' branch explores whether networks learn similar codes but lacks the algorithmic machinery for partial matching. The paper's contribution bridges classical optimal transport theory with practical neural comparison challenges, diverging from rigid one-to-one matching (e.g., Inexact Neural Matching) and complementing geometric registration methods by targeting neuron-level alignment with explicit mass relaxation.

Among the three contributions analyzed, the core unbalanced soft-matching distance examined ten candidates with zero refutable prior work, suggesting novelty within the limited search scope. The L-curve heuristic for regularization selection examined four candidates with one refutable match, indicating some overlap with existing parameter selection methods. Efficient neuron ranking examined two candidates with one refutable match, pointing to prior work on alignment-based ranking. The statistics reflect a modest search scale (sixteen total candidates), so these findings characterize novelty relative to top semantic matches rather than exhaustive coverage of the field.

Given the limited search scope and sparse taxonomy leaf, the work appears to introduce a principled extension of soft-matching to unbalanced settings, a direction not densely explored in the examined literature. The core transport formulation shows novelty among the candidates reviewed, while auxiliary contributions (L-curve heuristic, ranking) have more substantial prior work. The analysis covers top semantic matches and does not claim exhaustive field coverage, leaving open the possibility of related work outside the examined set.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Unbalanced soft-matching distance for partial neural correspondence

Description: The authors extend the soft-matching distance to a partial optimal transport framework that permits some neurons to remain unmatched rather than forcing all units into correspondence. This relaxes strict mass conservation constraints while maintaining interpretable transport costs and enables rotation-sensitive but robust alignments between neural populations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Partially Aligned Cross-modal Retrieval via Optimal Transport-based Prototype Alignment Learning

URL: View paper

Brief Assessment

Partial Cross Modal[24] applies optimal transport to cross-modal retrieval (aligning images and texts via prototypes), not to neural network unit correspondence or representational similarity analysis between neural populations.

---

2. Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling

URL: View paper

Brief Assessment

Velocity Growth Matching[29] addresses unbalanced optimal transport for single-cell dynamics modeling, not neural network unit correspondence. The candidate focuses on biological cell state transitions and mass growth in RNA sequencing data, while the original paper develops methods for comparing neural representations across networks by allowing partial neuron matching.

---

3. Learning Partial Graph Matching via Optimal Partial Transport

URL: View paper

Brief Assessment

Partial Graph Matching[31] addresses partial graph matching between discrete node sets using optimal partial transport, while the original paper focuses on comparing continuous neural representations (tuning curves, voxel responses) across networks. The candidate's framework is designed for graph structure matching with binary assignments, not for measuring representational similarity in neural populations where rotation-sensitivity and continuous transport costs are central concerns.

---

4. From one to all: Learning to match heterogeneous and partially overlapped graphs

URL: View paper

Brief Assessment

Heterogeneous Graph Matching[26] addresses graph matching with partial overlap using partial optimal transport, but focuses on matching nodes across heterogeneous graphs (different node types) rather than neural network units. The original paper's contribution is specific to neural representational comparison with rotation-sensitive correspondences between neural populations, which is a fundamentally different application domain.

---

5. Jointly aligning cells and genomic features of single-cell multi-omics data with co-optimal transport

URL: View paper

Brief Assessment

Co Optimal Transport[27] focuses on aligning cells and genomic features in single-cell multi-omic datasets, not neural network representations. The application domain (biological cells vs. neural networks) and technical objectives (multi-omic integration vs. neural correspondence) are fundamentally different.

---

6. Learning to rematch mismatched pairs for robust cross-modal retrieval

URL: View paper

Brief Assessment

Rematch Mismatched Pairs[23] addresses cross-modal retrieval with mismatched image-text pairs using partial optimal transport, not neural network unit correspondence. The candidate focuses on rematching multimedia data pairs across modalities, while the original develops metrics for comparing neural representations within networks.

---

7. Neural Optimal Transport for Dynamical Systems: Methods and Applications in Biomedicine

URL: View paper

Brief Assessment

Neural Optimal Transport[25] focuses on optimal transport methods for dynamical systems in biomedicine. The extremely limited context provided shows no discussion of neural network unit correspondence, soft-matching distances, or representational similarity metrics that are central to the original contribution.

---

8. Enhancing robust semi-supervised graph alignment via adaptive optimal transport

URL: View paper

Brief Assessment

Robust Graph Alignment[22] focuses on semi-supervised graph alignment across different networks using optimal transport, not neural network unit correspondence. The candidate addresses cross-domain graph matching with anchor links, while the original develops partial matching for neural populations with unmatched neurons.

---

9. Alpine: Partial Unlabeled Graph Alignment

URL: View paper

Brief Assessment

Alpine Graph Alignment[28] addresses partial graph alignment using adjacency matrices for structural matching, while the original paper develops partial optimal transport for neural tuning curve correspondence. These are fundamentally different domains with distinct technical approaches.

---

10. REALIGN: Regularized Procedure Alignment with Matching Video Embeddings via Partial Gromov-Wasserstein Optimal Transport

URL: View paper

Brief Assessment

REALIGN Procedure Alignment[30] focuses on aligning instructional video frames for procedure learning, not neural network unit correspondence. The partial transport formulation serves a different purpose (handling background frames in videos) rather than comparing neural representations across networks or brain recordings.

Contribution 2: L-curve heuristic for automatic regularization selection

Description: The authors introduce an L-curve method to automatically determine the optimal fraction of mass to transport between neural populations. This heuristic identifies the point of maximal positive curvature in the cost-regularization tradeoff curve, enabling principled selection of how many units should be matched without requiring prior knowledge of noise levels.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Prestack waveform inversion by using an optimized linear inversion scheme

URL: View paper

Prior Art Analysis

Prestack Waveform Inversion[34] demonstrates prior use of the L-curve method for automatic regularization parameter selection in an optimization context. The candidate paper explicitly states they 'use l-curve method to acquire the optimal regularization weight adaptively' in their inversion scheme. This directly refutes the novelty claim of introducing the L-curve heuristic for automatic regularization selection, as the method was already established and applied in seismic inversion problems before the original paper's submission.

Evidence

Evidence 1 - **Rationale**: Both papers use the L-curve method to automatically determine optimal regularization parameters. The candidate paper explicitly states using 'l-curve method to acquire the optimal regularization weight adaptively,' demonstrating that this approach was already established in optimization problems prior to the original paper's work. - **Original**: To address this, we adopt an l-curve heuristic (cultrera & callegaro, 2020), inspired by classical regularization methods for ill-posed problems (e.g., tikhonov regularization). the l-curve captures the tradeoff between transport distance and regularization strength, with the 'elbow' typically indic... - **Candidate**: to avoid falling into local extrema in the process of solving the objective function, we introduce an optimal transport method into the objective function to improve its convexity and use l-curve method to acquire the optimal regularization weight adaptively.

Evidence 2 - **Rationale**: Both papers describe using the L-curve to identify optimal regularization by balancing competing objectives. The candidate's 'adaptively determined regularization weight in each iterative step' parallels the original's identification of optimal regularization at the curve's elbow point, showing this adaptive selection strategy existed prior to the original work. - **Original**: the optimal regularizations0 is identified at the curve's point of maximal positive curvature (the elbow), which balances low transport cost against aggressive regularization. - **Candidate**: in this paper, we have developed a model-based prestack waveform inversion (pwi) with a generalized propagation matrix scheme as forward operator, where a regularized function is minimized with the limited memory-broydenâfletcherâgoldfarbâshanno technique to determine a model update corresponding to...

---

2. Model Error Covariance Estimation for Weak Constraint Data Assimilation

URL: View paper

Brief Assessment

Model Error Covariance[35] applies the L-curve method to estimate model error variance in data assimilation, not to determine optimal transport fractions between neural populations. The technical domains and applications are fundamentally different.

---

3. Zero-shot physics-guided deep learning for subject-specific MRI reconstruction

URL: View paper

Brief Assessment

Zero Shot MRI[33] applies L-curve methods to MRI reconstruction for determining early stopping in neural network training, not for selecting regularization in optimal transport problems between neural populations as in the original paper.

---

4. Reducing errors in the GRACE gravity solutions using regularization

URL: View paper

Brief Assessment

GRACE Regularization[32] applies L-curve to geophysical gravity solutions for mass transport signals, not to optimal transport problems in neural network comparison or neuroscience applications.

Contribution 3: Efficient neuron ranking by alignment quality

Description: The method provides a computationally efficient approach to rank neurons by their cross-population alignment quality. A single optimization at an appropriate regularization value achieves results nearly identical to exhaustive brute-force ranking while requiring substantially fewer operations, making it practical for identifying highly-aligned or poorly-aligned neural subpopulations.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Safe Crossover of Neural Networks Through Neuron Alignment

URL: View paper

Brief Assessment

Safe Crossover Alignment[21] focuses on aligning neurons between parent networks for genetic algorithm crossover operations, not on ranking neurons by cross-population alignment quality or avoiding iterative recomputation for representational comparison tasks.

---

2. Zeroth-Order Adaptive Neuron Alignment Based Pruning without Re-Training

URL: View paper

Prior Art Analysis

Zeroth Order Alignment[20] demonstrates prior work that achieves efficient neuron ranking by alignment quality without iterative recomputation. The candidate paper presents a method that ranks neurons by their cross-population alignment using a single optimization at an appropriate regularization value, achieving results nearly identical to brute-force ranking while requiring substantially fewer operations ($O(n^3 \log n)$ vs $O(n^4 \log n)$). This directly addresses the same problem of efficiently identifying highly-aligned or poorly-aligned neural subpopulations without costly iterative recomputation, which the original paper claims as novel.

Evidence

Evidence 1 - **Rationale**: Both papers claim to provide efficient methods for neuron ranking/selection without requiring extensive recomputation, establishing that this approach existed prior to the original paper's submission. - **Original**: this unbalanced soft-matching distance provides theoretical advantages-relaxing strict mass conservation while maintaining interpretable transport costs-and practical benefits through efficient neuron ranking in terms of cross-network alignment without costly iterative recomputation. - **Candidate**: we introduceneuronal, a noveltop-uppruning algorithm that outperforms, in most cases, the previous state-of-the-art approaches over both language modeling datasets and zero-shot tasks, while providing hyperparameter adaptation and reduced runtime to obtain the non-uniform sparsity allocation.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Unbalanced Soft-Matching Distance For Neural Representational Comparison With Partial Unit Correspondence View paper
- [1] A unified neural representation model for spatial and semantic computations View paper
- [2] Neural representations of unconditioned stimuli in basolateral amygdala mediate innate and learned responses View paper
- [3] Cracking the neural code for sensory perception by combining statistics, intervention, and behavior View paper
- [4] Towards understanding learning representations: To what extent do different neural networks learn the same representation View paper
- [5] Prnet: Self-supervised learning for partial-to-partial registration View paper
- [6] The neural code for written words: a proposal View paper
- [7] Contrastive and adversarial regularized multi-level representation learning for incomplete multi-view clustering. View paper
- [8] Convergent learning: Do different neural networks learn the same representations? View paper
- [9] Multisensory integration: current issues from the perspective of the single neuron View paper
- [10] Representational geometry of incomplete faces in macaque face patches. View paper
- [11] Issues of representation in neural networks View paper
- [12] Learning to align the source code to the compiled object code View paper
- [13] Matching neural paths: transfer from recognition to correspondence search View paper
- [14] RORNet: Partial-to-Partial Registration Network With Reliable Overlapping Representations View paper
- [15] Briefly cuing memories leads to suppression of their neural representations View paper
- [16] EVAN: Evolutional Video Streaming Adaptation via Neural Representation View paper
- [17] Foundations of Neuronal Representations View paper
- [18] Inexact matching using neural networks View paper
- [19] Implicit Graphon Neural Representation View paper
- [20] Zeroth-Order Adaptive Neuron Alignment Based Pruning without Re-Training View paper
- [21] Safe Crossover of Neural Networks Through Neuron Alignment View paper
- [22] Enhancing robust semi-supervised graph alignment via adaptive optimal transport View paper
- [23] Learning to rematch mismatched pairs for robust cross-modal retrieval View paper
- [24] Partially Aligned Cross-modal Retrieval via Optimal Transport-based Prototype Alignment Learning View paper
- [25] Neural Optimal Transport for Dynamical Systems: Methods and Applications in Biomedicine View paper
- [26] From one to all: Learning to match heterogeneous and partially overlapped graphs View paper
- [27] Jointly aligning cells and genomic features of single-cell multi-omics data with co-optimal transport View paper
- [28] Alpine: Partial Unlabeled Graph Alignment View paper
- [29] Joint Velocity-Growth Flow Matching for Single-Cell Dynamics Modeling View paper
- [30] REALIGN: Regularized Procedure Alignment with Matching Video Embeddings via Partial Gromov-Wasserstein Optimal Transport View paper
- [31] Learning Partial Graph Matching via Optimal Partial Transport View paper
- [32] Reducing errors in the GRACE gravity solutions using regularization View paper
- [33] Zero-shot physics-guided deep learning for subject-specific MRI reconstruction View paper
- [34] Prestack waveform inversion by using an optimized linear inversion scheme View paper
- [35] Model Error Covariance Estimation for Weak Constraint Data Assimilation View paper