UNPAIRED POINT CLOUD COMPLETION USING UNBALANCED OPTIMAL TRANSPORT MAP

Taekyung Lee
IPAI,* Seoul National University
dlxorud1231@snu.ac.kr

Jaemoo Choi
Georgia Institute of Technology
jchoi843@gatech.edu

Jaewoong Choi†
Sungkyunkwan University
jaewoongchoi@skku.edu

Myungjoo Kang†
Dept. of Mathematical Sciences & RIM, Seoul National University
mkang@snu.ac.kr

ABSTRACT

Real-world point cloud is incomplete and lacks corresponding complete point cloud pairs. To address the challenge of unpaired point cloud completion, we introduce UOT-UPC, a novel approach grounded in the (Unbalanced) Optimal Transport (OT) problem. Specifically, we demonstrate that solving OT-based framework provides an effective approach to unpaired point cloud completion; the optimal transport map T^* can serve as an unpaired point cloud completion model. Furthermore, we extend this formulation to incorporate the Unbalanced Optimal Transport (UOT) Map, achieving competitive performance in unpaired point cloud completion. This extension also addresses the class imbalance, a phenomenon in which differences in the composition ratios of two distributions cause undesired distribution shifts in the OT problem. In this paper, we provide theoretical evidence supporting the UOT-based framework's competency to manage the class imbalance, and we validate its effectiveness through experiments.

1 INTRODUCTION

The point cloud serves as a fundamental representation for 3D object. In real-world, the sensor-acquired point cloud is often incomplete and lacks the corresponding complete representation. Consequently, numerous methods have been proposed for unpaired point cloud completion, which is aimed at reconstructing complete point cloud from incomplete input without paired training data.

However, most existing methods rely on heuristic approaches. To address this limitation, we adopt the Optimal Transport (OT) problem (Monge, 1781). Specifically, we propose why OT-based approach is particularly suited to unpaired point cloud completion and demonstrate that the solution of the OT problem becomes one of unpaired point completion models.

Building on this perspective, we propose a novel method, UOT-UPC, leveraging the Unbalanced Optimal Transport Map. While the classical OT framework cannot tackle the class imbalance, the Unbalanced OT framework overcomes the limitation (Eyring et al., 2024). The experiments demonstrate that our model not only outperforms existing unpaired point cloud completion methods but also effectively handles the class imbalance. Now, we summarize our contributions as follows:

- UOT-UPC is the first to employ UOTM for unpaired point cloud completion.
- UOT-UPC demonstrates a competitive performance for point cloud completion and class imbalance.

Notations and Assumptions Let $\mathcal{X}$, $\mathcal{Y}$ be compact complete metric spaces with probability distributions μ and ν (absolutely continuous with respect to the Lebesgue measure), where μ (source) represents the incomplete point clouds and ν (target) represents the complete point clouds. For any measurable map T, $T_{\#}\mu$ denotes the pushforward of μ and $\Pi(\mu, \nu)$ denote the set of joint

*Interdisciplinary Program in Artificial Intelligence
†Corresponding Author

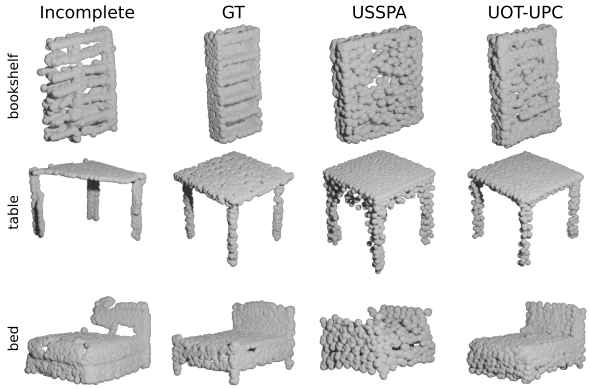

Figure 1: **Comparison of generated samples** from UOT-UPC and USSPA in the single-category.

distributions on $\mathcal{X} \times \mathcal{Y}$ with marginals μ and ν. Additionally, for a function $f : \mathbb{R} \to [-\infty, \infty]$, its convex conjugate is given by $f^*(y) = \sup_{x \in \mathbb{R}}\{\langle x, y\rangle - f(x)\}$.

2 Background

2.1 Optimal Transport

The Optimal Transport(OT) problem is finding the cost-minimizing way to transport source distribution μ to target distribution ν for the given cost function $c(\cdot, \cdot)$. Monge (1781) first introduced the exploration of OT problem assuming deterministic transport map $T : \mathcal{X} \to \mathcal{Y}$ such that $T_{\#}\mu = \nu$.

$$C(\mu, \nu) := \inf_{T_{\#}\mu = \nu} \left[\int_{\mathcal{X}} c(x, T(x)) d\mu(x) \right]. \tag{1}$$

where the minimum is taken over transport map T. However, Monge's OT problem has limitations. The formulation is non-convex, and solution T^* may not exist depending on μ and ν (Villani et al., 2009). To overcome this limitation, Kantorovich (1948) formulated the following relaxation:

$$C_{ot}(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) \right]. \tag{2}$$

where the minimum is taken over transport plan π. In distinction to Monge's OT problem, the minimizer transport plan π^* always exists on some mild assumptions on $(\mathcal{X}, \mu), (\mathcal{Y}, \nu)$, and cost function c (Villani et al., 2009). Under the assumption that both μ and ν are absolutely continuous with respect to the Lebesgue measure, the deterministic optimal transport map T^* exists and the optimal coupling is characterized by $\pi^* = (I \times T^*)_{\#\mu}$ (Villani et al., 2009).

2.2 Unbalanced Optimal Transport

The standard OT problem assumes $\pi_0 = \mu$ and $\pi_1 = \nu$. However, this rigid constraint induces sensitivity to outliers (Balaji et al., 2020) and is unsuited to the class imbalance (Eyring et al., 2024). To address these issues, a new type of OT problem called Unbalanced Optimal Transport (UOT) was introduced (Chizat et al., 2018) (Liero et al., 2018) defined as follows:

$$C_{uot}(\mu, \nu) = \inf_{\pi \in \mathcal{M}_+(\mathcal{X} \times \mathcal{Y})} \left[\int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) + D_{\Psi_1}(\pi_0 | \mu) + D_{\Psi_2}(\pi_1 | \nu) \right], \tag{3}$$

where $\mathcal{M}_+(\mathcal{X} \times \mathcal{Y})$ denotes the set of positive Radon measures on $\mathcal{X} \times \mathcal{Y}$. Each D_{Ψ_i} is an f-divergence induced by a convex entropy function Ψ_i, which is defined by $D_{\Psi_i}(\pi_j | \eta) = \int \Psi_i \left(\frac{d\pi_j(x)}{d\eta(x)} \right) d\eta(x)$. The UOT problem relaxes the hard marginal constraints of the OT problem through f-divergence (Vacher & Vialard, 2023). The UOT problem can be seen as the OT problem between $\pi_0 \approx \mu$ and $\pi_1 \approx \nu$ (Choi et al., 2023). This relaxation imparts robustness to outliers (Balaji et al., 2020) and addresses the class imbalance (Eyring et al., 2024).

Table 1: **Point cloud completion comparison in the single-category setting**, assessed by L1-Chamfer distance $cd^{l1} \times 10^2$ (↓). The boldface denotes the best performance among the unpaired.

	Method	AVG	chair	table	trash bin	TV	cabinet	bookshelf	sofa	lamp	bed	tub
Paired	PoinTr (Yu et al., 2021)	14.37	13.65	12.52	15.26	12.69	17.32	13.99	12.36	17.05	15.13	13.77
	Disp3D (Wang et al., 2022)	7.78	6.24	8.20	7.12	7.12	10.36	6.94	5.60	14.03	6.90	5.32
	TopNet (Tchapmi et al., 2019)	7.07	6.39	5.79	7.40	6.26	8.37	7.02	5.94	8.50	7.81	7.25
Unpaired	ShapeInv (Zhang et al., 2021)	21.39	17.97	17.28	33.51	15.69	26.26	25.51	14.28	16.69	32.33	14.43
	Unpaired (Chen et al., 2020)	10.47	8.41	7.52	12.08	6.72	17.45	9.95	6.92	19.36	10.04	6.22
	Cycle4 (Wen et al., 2021)	11.53	9.11	11.35	11.93	8.40	15.47	12.51	10.63	12.25	15.73	7.92
	USSPA (Ma et al., 2023)	8.56	8.22	7.68	10.36	7.66	**10.77**	7.84	**6.14**	11.93	**8.20**	6.75
	UOT-UPC (Ours)	**7.60**	**7.51**	**6.33**	**8.83**	**6.07**	11.54	**7.32**	6.61	**7.30**	9.00	**5.45**

3 METHODS

3.1 MOTIVATIONS

Task formulation as (Unbalanced) Optimal Transport Map Unpaired point cloud completion transfers point cloud from incomplete point cloud distribution to a complete point cloud distribution. Consider two point cloud sets $X = \{x_i | x_i \in \mathcal{X}, i = 1, 2, ..., N\}$ and $Y = \{y_i | y_i \in \mathcal{Y}, i = 1, 2, ..., M\}$. X and Y are independently sampled from μ and ν, i.e., X and Y are unpaired. We define an unpaired point cloud completion model T as follows:

$$T : \mathcal{X} \to \mathcal{Y}, \quad x \text{ (incomplete point cloud)} \to T(x) \text{ (complete point cloud)} \tag{4}$$

This completion model T must satisfy two conditions:

(i) $T(x)$ should generate a complete point cloud ($T(x) \in \nu$).

(ii) T should align each incomplete x with its corresponding complete y.

In light of this, the optimal transport map T^* defined in Eq. 1 is well-suited for unpaired point cloud completion. By definition, T^* transports μ to ν while minimizing the cost function $c(x, T^*(x))$. This guarantees condition (i). If an appropriate cost function is chosen such that T^* also satisfies condition (ii), then T^* serves as an effective model for unpaired point completion. More specifically, the cost function's value quantifies the alignment between x and T(x), remaining low for well-aligned pairs and high for incompatible pairs.

Class Imbalance In the unpaired setting, the absence of corresponding complete point clouds for incomplete ones often results in the class imbalance. For instance, while the incomplete point cloud distribution μ exhibits a chair-to-table ratio of 50% to 50%, the complete point cloud distribution ν demonstrates a ratio of 70% to 30%. Under the classical OT framework, the OT map transfers this unintended distribution shift, aligning an extraneous 20%p of tables in μ with 20%p of chair in ν. To tackle this, we integrate the Unbalanced OT framework into our method (See Sec 3.2).

3.2 PROPOSED METHOD: ESTIMATION OF UNBALANCED OPTIMAL TRANSPORT MAP

In this paper, we propose **UOT-UPC**, a point completion model based on the Unbalanced OT (UOT) framework. Our method is to learn the UOT Map T^* from incomplete point cloud distribution μ to complete point cloud distribution ν. We adopt the UOTM (Choi et al., 2023), which utilizes semi-dual formulation of the UOT problem:

$$C_{uot}(\mu, \nu) = \sup_{v \in \mathcal{C}} \left[\int_{\mathcal{X}} -\Psi_1^* \left(-v^c(x)\right)) \, d\mu(x) + \int_{\mathcal{Y}} -\Psi_2^*(-v(y)) d\nu(y) \right], \tag{5}$$

where the c-transform of v is defined as $v^c(x) = \inf_{y \in \mathcal{Y}} [c(x, y) - v(y)]$. Following earlier works (Korotin et al., 2021; Fan et al., 2022), we approximate T^* with a parametrized map T_θ, such that:

$$T_\theta(x) \in \underset{y \in \mathcal{Y}}{argin}f [c(x, y) - v(y)] \Leftrightarrow v^c = c(x, T_\theta(x)) - v(T_\theta(x)). \tag{6}$$

Now, we term the optimal maximizer v^* of Eq. 5 as the optimal potential function in the UOT problem. We parameterize v^* by v_ϕ and replace v^c with the right-hand side of Eq 6, yielding the following learning objective:

$$\mathcal{L}_{v_\phi, T_\theta} = \inf_{v_\phi} \left[\int_{\mathcal{X}} \Psi_1^* \left(-\inf_{T_\theta} \left[c\left(x, T_\theta(x)\right) - v_\phi\left(T_\theta(x)\right) \right] \right) d\mu(x) + \int_{\mathcal{Y}} \Psi_2^* \left(-v_\phi(y)\right) d\nu(y) \right]. \tag{7}$$

Table 2: **Point cloud completion comparison** in the single-category setting and the multi-category setting, assessed by L1 Chamfer Distance $cd^{l1} \times 10^2$ (↓) and F-scores $F^{0.1\%}_{\text{score}} \times 10^2$, $F^{1\%}_{\text{score}} \times 10^2$ (↑).

Method		Single-category		Multi-category		
		$F^{0.1\%}_{\text{score}}$ ↑	$F^{1\%}_{\text{score}}$ ↑	cd^{l1} ↓	$F^{0.1\%}_{\text{score}}$ ↑	$F^{1\%}_{\text{score}}$ ↑
Paired	PoinTr (Yu et al., 2021)	-	-	14.37	18.35	80.41
	Disp3D (Wang et al., 2022)	-	-	7.78	30.29	78.26
	TopNet (Tchapmi et al., 2019)	-	-	7.07	12.33	80.37
Unpaired	ShapeInv (Zhang et al., 2021)	15.58	66.53	19.35	16.98	69.66
	Unpaired (Chen et al., 2020)	12.20	64.33	10.12	10.86	66.68
	Cycle4 (Wen et al., 2021)	9.98	60.14	12.00	8.61	56.57
	USSPA (Ma et al., 2023)	17.49	73.41	**8.96**	16.88	**72.31**
	UOT-UPC (Ours)	**18.43**	**75.59**	**8.96**	**19.25**	71.52

Note when the optimal potential v^* is given, the UOT Map $T*$ can be interpreted as the OT Map between $\pi_0(x) = \Psi_1^{*\prime}(-v^{*c}(x))\mu(x)$ and $\pi_1(x) = \Psi_2^{*\prime}(-v^*(x))\nu(x)$ (Choi et al., 2023). These rescaling factors $\Psi_i^{*\prime}(\cdot)$ provide the flexibility to tackle the class imbalance (Eyring et al., 2024). For the UOT-based formulation, we employ InfoCD (Lin et al., 2024) as a cost function. InfoCD is a regularized Chamfer distance robust to outliers, enabling more accurate cost estimation between incomplete and complete point clouds. In Sec 4, our model exhibits competitive performance in unpaired point completion and effectively addresses the class imbalance (See Appendix A for details).

4 EXPERIMENTS

Unpaired point cloud completion requires learning from two distributions: incomplete and complete point cloud distribution. The USSPA benchmark (Ma et al., 2023) is a suitable for unpaired point cloud completion, containing ten categories (e.g., chair, table). We use this dataset to train and evaluate our model with respect to (1) Unpaired Point Cloud Completion and (2) Class Imbalance.

4.1 UNPAIRED POINT CLOUD COMPLETION

We evaluate our model in both single- and multi-category settings. Table 1 and Table 2 report the results for each. We measure completion quality using L1 Chamfer distance(cd^{l1}), and F-score($F^{0.1\%}_{score}$, $F^{1\%}_{score}$). In both settings, our model achieves state-of-the-art performance among unpaired models for most categories. Fig 1 shows our generated outputs in the single-category setting (Appendix B.1).

4.2 CLASS IMBALANCE

We assess our model under the class imbalance, which is prevalent in unpaired point cloud completion. Specifically, we select two categories and adjust their ratios in μ and ν. For instance, we choose 'TV' and 'Table' and modify the ratio in ν by a factor r, TV : Table $= 6.4 : 21.3$ in $\mu \rightarrow$ TV : Table $= 6.4 : 21.3 \times r$ in ν (See Table 3).

As shown in Table 3, our model outperforms the others and maintains low variance. while USSPA exhibits high variance. In contrast, OT-UPC (the OT Map based approach) performs poorly and unstably. The stable training dynamics in learning the transport map underscore another advantage of the UOT over OT (Choi et al., 2024). In summary, these results confirm the strong robustness of UOT-UPC against the class imbalance. See Appendix B.2 for additional categories.

Table 3: **Comparison of the class imbalance robustness** ($cd^{l1} \times 10^2$ (↓)) on TV and Table.

r	0.3	0.5	0.7	1
USSPA	7.60	6.97	8.08	7.97
OT-UPC	23.77	23.74	29.79	27.21
Ours	**6.86**	**6.75**	**6.75**	**6.94**

5 CONCLUSION

We propose UOT-UPC which is an unpaired point completion model based on the UOT Map. We reformulate the OT map for unpaired point completion and extend it to the Unbalanced OT Map to solve the class imbalance. Our model shows great performance and strong robustness to the class imbalance. A limitation is the use of InfoCD as a cost function; future work will explore alternatives.

ACKNOWLEDGEMENTS

This work was supported by the NRF grant [RS-2021-NR059802, RS-2024-00421203, RS-2024-00496127] and MSIT/IITP grant[RS-2021-II211343-GSAI]. JW was partially supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) [RS-2024-00349646].

REFERENCES

Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. *Advances in Neural Information Processing Systems*, 33:12934–12944, 2020.

Xuelin Chen, Baoquan Chen, and Niloy J. Mitra. Unpaired point cloud completion on real scans using adversarial training. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=HkgrZ0EYwB`.

Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11): 3090–3123, 2018.

Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Generative modeling through the semi-dual formulation of unbalanced optimal transport. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Analyzing and improving optimal-transport-based adversarial networks. In *The Twelfth International Conference on Learning Representations*, 2024.

Luca Eyring, Dominik Klein, Théo Uscidda, Giovanni Palla, Niki Kilbertus, Zeynep Akata, and Fabian J Theis. Unbalancedness in neural monge maps improves unpaired domain translation. In *The Twelfth International Conference on Learning Representations*, 2024. URL `https://openreview.net/forum?id=2UnCj3jeao`.

Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 605–613, 2017.

Jiaojiao Fan, Shu Liu, Shaojun Ma, Yongxin Chen, and Hao-Min Zhou. Scalable computation of monge maps with general costs. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.

Leonid Vitalevich Kantorovich. On a problem of monge. *Uspekhi Mat. Nauk*, pp. 225–226, 1948.

Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, Alexander Filippov, and Evgeny Burnaev. Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark. *Advances in Neural Information Processing Systems*, 34:14593–14605, 2021.

Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures. *Inventiones mathematicae*, 211(3): 969–1117, 2018.

Fangzhou Lin, Yun Yue, Ziming Zhang, Songlin Hou, Kazunori Yamada, Vijaya Kolachalama, and Venkatesh Saligrama. Infocd: a contrastive chamfer distance loss for point cloud completion. *Advances in Neural Information Processing Systems*, 36, 2024.

Changfeng Ma, Yinuo Chen, Pengxiao Guo, Jie Guo, Chongjun Wang, and Yanwen Guo. Symmetric shape-preserving autoencoder for unsupervised real scene point cloud completion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13560–13569, 2023.

Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pp. 3481–3490. PMLR, 2018.

Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pp. 666–704, 1781.

Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. Stabilizing training of generative adversarial networks through regularization. *Advances in neural information processing systems*, 30, 2017.

Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3405–3414, 2019.

Lyne P Tchapmi, Vineet Kosaraju, Hamid Rezatofighi, Ian Reid, and Silvio Savarese. Topnet: Structural point cloud decoder. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 383–392, 2019.

Adrien Vacher and François-Xavier Vialard. Semi-dual unbalanced quadratic optimal transport: fast statistical rates and convergent algorithm. In *International Conference on Machine Learning*, pp. 34734–34758. PMLR, 2023.

Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

Yida Wang, David Joseph Tan, Nassir Navab, and Federico Tombari. Learning local displacements for point cloud completion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1568–1577, 2022.

Xin Wen, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. Cycle4completion: Unpaired point cloud completion using cycle transformation with missing region coding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13080–13089, 2021.

Rundi Wu, Xuelin Chen, Yixin Zhuang, and Baoquan Chen. Multimodal shape completion via conditional generative adversarial networks. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pp. 281–296. Springer, 2020.

Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(3): 3121–3138, 2022.

Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. Pointr: Diverse point cloud completion with geometry-aware transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12498–12507, 2021.

Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 international conference on 3D vision (3DV)*, pp. 728–737. IEEE, 2018.

Junzhe Zhang, Xinyi Chen, Zhongang Cai, Liang Pan, Haiyu Zhao, Shuai Yi, Chai Kiat Yeo, Bo Dai, and Chen Change Loy. Unsupervised 3d shape completion through gan inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1768–1777, 2021.

A IMPLEMENTATION DETAILS

Unless otherwise stated, our implementation follows the experimental settings and hyperparameters of USSPA Ma et al. (2023).

A.1 NETWORK

We adopt the generator and discriminator architectures from the USSPA framework as completion model T_θ and potential v_ϕ. For the potential v_ϕ, the final sigmoid layer of the discriminator is omitted to allow for the parameterization of the potential function, enabling outputs to assume any real values. Additionally, we remove the feature discriminator to streamline the architecture. In the potential v_ϕ, we implement the encoder proposed by Yuan et al. (2018) in their Point Cloud Networks (PCN). Following the encoder, we employ an MLPConv layer specified as $\text{MLPConv}(C_{in}, [C_1, \ldots, C_n]) = \text{MLPConv}(1024, [256, 256, 128, 128, 1])$, which indicates that the output y is computed as follows:

$$y = \text{Conv1D}_{C_4=128, C_5=1}(\text{ReLU}(\ldots \text{ReLU}(\text{Conv1D}_{C_{in}=1024, C_1=256}(x))\ldots)) \tag{8}$$

Here, $\text{Conv1D}_{C_{\text{in}}, C_{\text{out}}}$ represents a 1D convolutional layer with C_{in} input channels and C_{out} output channels.

The completion model T_θ receives as input a concatenation of the incomplete point cloud and a complete point cloud. These inputs are processed independently to generate distinct complete point cloud samples. The completion model T_θ follows an Encoder-Decoder architecture, augmented by an upsampling refinement module (upsampling module) in sequence. The upsampling module is implemented using a 4-layer MLPConv network, where the final MLPConv layer is responsible for refining and adding detailed structures to the output (Ma et al., 2023). Specifically, the inputs to the last MLPConv layer are composed of the skeleton point cloud produced by the Encoder-Decoder structure and the features extracted from the third MLPConv layer.

A.2 IMPLEMENTATION DETAIL

Training The functions Ψ_1^* and Ψ_2^* are defined using the Softplus activation, $\text{SP}(x) = 2\log(1 + e^x) - 2\log 2$.[1] As a regularization term, we incorporate the density loss dl proposed by Ma et al. (2023), and we designate a coordinate value of 10.5 for dl. The objective of Potential v_ϕ is to assign high value to target sample y while assigning lower values to generated sample $\hat{y}$. We utilize the Adam optimizer with $\beta_1 = 0.95, \beta_2 = 0.999$ and learning rates of 1.0×10^{-5} for both the potential v_ϕ and completion model T_θ, respectively. The training is conducted with a batch size 4. The maximum epoch of training is 480. We report the final results based on the epoch that yields the best performance.

In practice, we adopt a mixture-based strategy that merges target distribution with the source distribution. Specifically, we constructa mixed distribution $\tilde{\mu}$ by combining 50% of the incomplete distribution μ with 50% of the complete distribution ν, i.e., $\tilde{\mu} = 0.5 \times \mu + 0.5 \times \nu$. Hence, our approach learns the UOTM that maps $\tilde{\mu}$ to ν. The detailed training procedure is provided in Algorithm 1

cost function Concerning the loss function $L_{v,T}$, we employ InfoCD(Lin et al., 2024) as the cost function c with a coordinate value of $\tau = 0.044$. For the hyperparameters of InfoCD, we set τ_{infocd} to 2 and λ_{InfoCD} to 1.0×10^{-7}.

- $\text{InfoCD}(x_i, y_j) = \ell_{\text{InfoCD}}(x_i, y_j) + \ell_{\text{InfoCD}}(y_j, x_i)$.
  where $\ell_{\text{InfoCD}}\ (x_i, y_i) = -\frac{1}{|y_i|}\sum_n \log\left\{\frac{\exp\left\{-\frac{1}{\tau'}\min_m d(x_{im}, y_{in})\right\}}{\sum_n \exp\left\{-\frac{1}{\tau}\min_m d(x_{im}, y_{in})\right\}}\right\}, \quad \lambda = \frac{\tau'}{\tau}$

[1]The softplus function is translated and scaled to satisfy $\text{SP}(0) = 0$ and $\text{SP}'(0) = 1$.

Algorithm 1 Training algorithm of UOT-UPC

Require: The mixture of the incomplete and complete point cloud distribution μ. The complete point cloud distribution ν. $\Psi_i^*(x) = \mathrm{Softplus}(x)$. Generator network T_θ and the discriminator network v_ϕ. dl is density loss. Total iteration number K.

for $k = 0, 1, 2, \ldots, K$ **do**
  Sample a batch $X \sim \mu$, $Y \sim \nu$.
  $\mathcal{L}_T = \frac{1}{|X|} \sum_{x \in X} c\left(x, T_\theta(x)\right) - v_\phi\left(T_\theta(x)\right) + dl\left(T_\theta(x)\right)$
  Update θ by minimizing the loss $\mathcal{L}_T$.
  $\mathcal{L}_v = \frac{1}{|X|} \sum_{x \in X} \Psi_1^*\left(-c\left(x, T_\theta(x)\right) + v_\phi\left(T_\theta(x)\right)\right) + \frac{1}{|Y|} \sum_{y \in Y} \Psi_2^*(-v_\phi(y))$
  Update ϕ by minimizing the loss $\mathcal{L}_v$.
end for

Evaluation Metrics

- $L1$-Chamfer Distance cd^{l1} (Fan et al., 2017)

$$cd^{l1}(x_i, y_j) = \frac{1}{2}\left(\frac{1}{|x_i|}\sum_m \min_n \|x_{im} - y_{jn}\|_2 + \frac{1}{|y_j|}\sum_n \min_m \|x_{im} - y_{jn}\|_2.\right) \tag{9}$$

  where each of x_i, y_j is point cloud

- F score F^{α}_{score} (Tatarchenko et al., 2019)

$$F^{\alpha}_{score} = \frac{2 \times P(\alpha) \times R(\alpha)}{P(\alpha) + R(\alpha)} \tag{10}$$

  where $P(\alpha) = \frac{|\{x_{im} \in x_i | \min_n(\|x_{im} - y_{jn}\|_2) < \alpha\}|}{|x_i|}$ measures the accuracy of x_i,
  and $R(\alpha) = \frac{|\{y_{jn} \in y_j | \min_m(\|x_{im} - y_{jn}\|_2) < \alpha\}|}{|y_j|}$ measures the completeness of x_i.

A.3 OT-UPC

OT-UPC is made by modifying the Φ_i^* of 7. Let Φ_i be an indicator function at $x = 1$. Then, Φ_i^* would be a linear function: $\Phi_i^*(x) = x$. For the potential v_ϕ, we implement MLPConv$(512, [128, 128, 1])$ following the PCN encoder (Yuan et al., 2018). We incorporate R1 regularization (Roth et al., 2017) and R2 regularization (Mescheder et al., 2018) to the loss function $L_{v,T}$. Both regularization terms are assigned coordinate values $r1 = r2 = 0.2$. The density loss dl is excluded from the $L_{v,T}$. A gradient clipping value of 1.0 is applied. We use Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$ and a learning rate $lr_{T_\theta} = 5.0 \times 10^{-5}$ for the completion model T_θ. In addition, we use Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$ and learning rate $lr_{v_\phi} = 1.0 \times 10^{-7}$ for the potential v_ϕ. All other settings not explicitly mentioned follow those of our model, UOT-UPC.

B Additional Results

B.1 Additional Qualitative Results

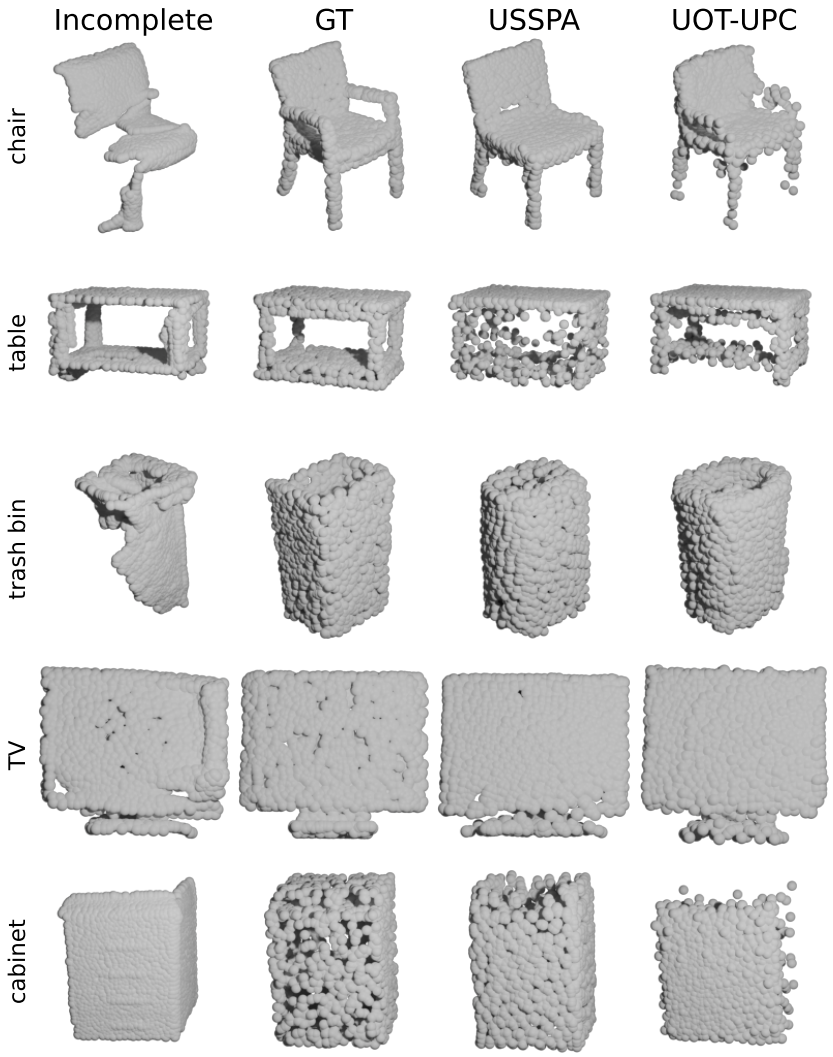

Figure 2: **Comparison of generated samples** from our UOT-UPC and USSPA in the single-category setting.

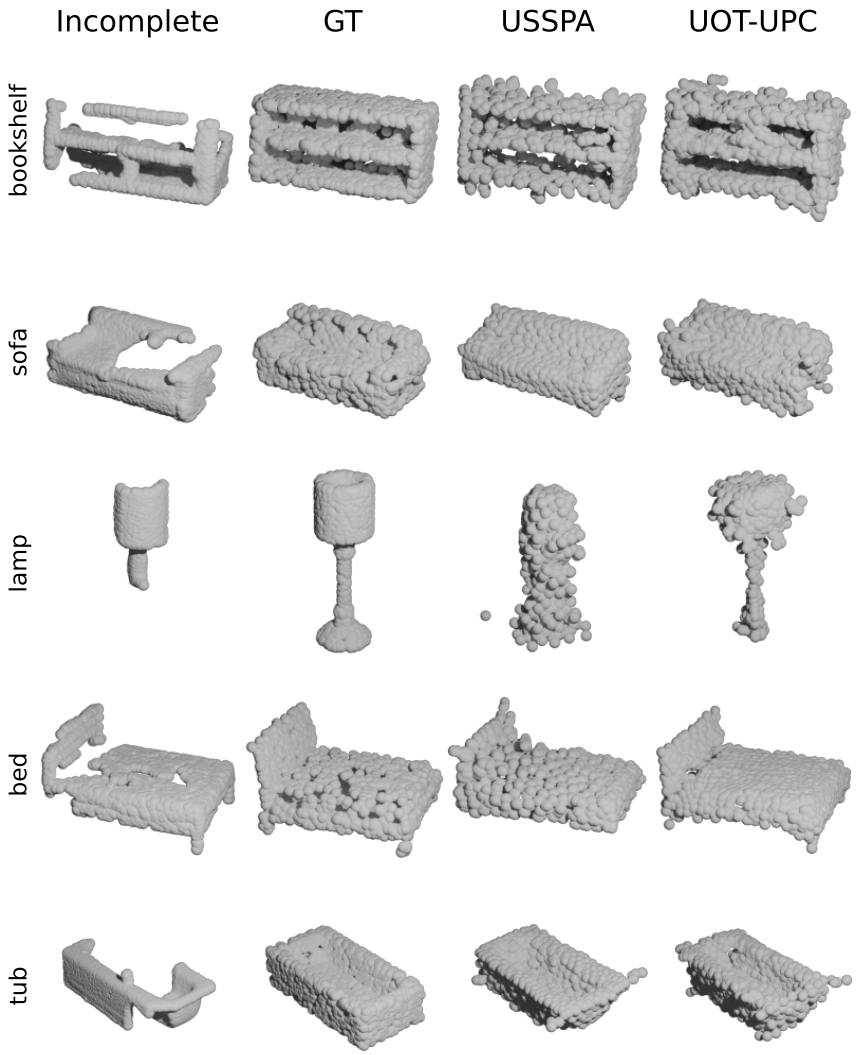

Figure 3: **Comparison of generated samples** from our UOT-UPC and USSPA in the single-category setting.

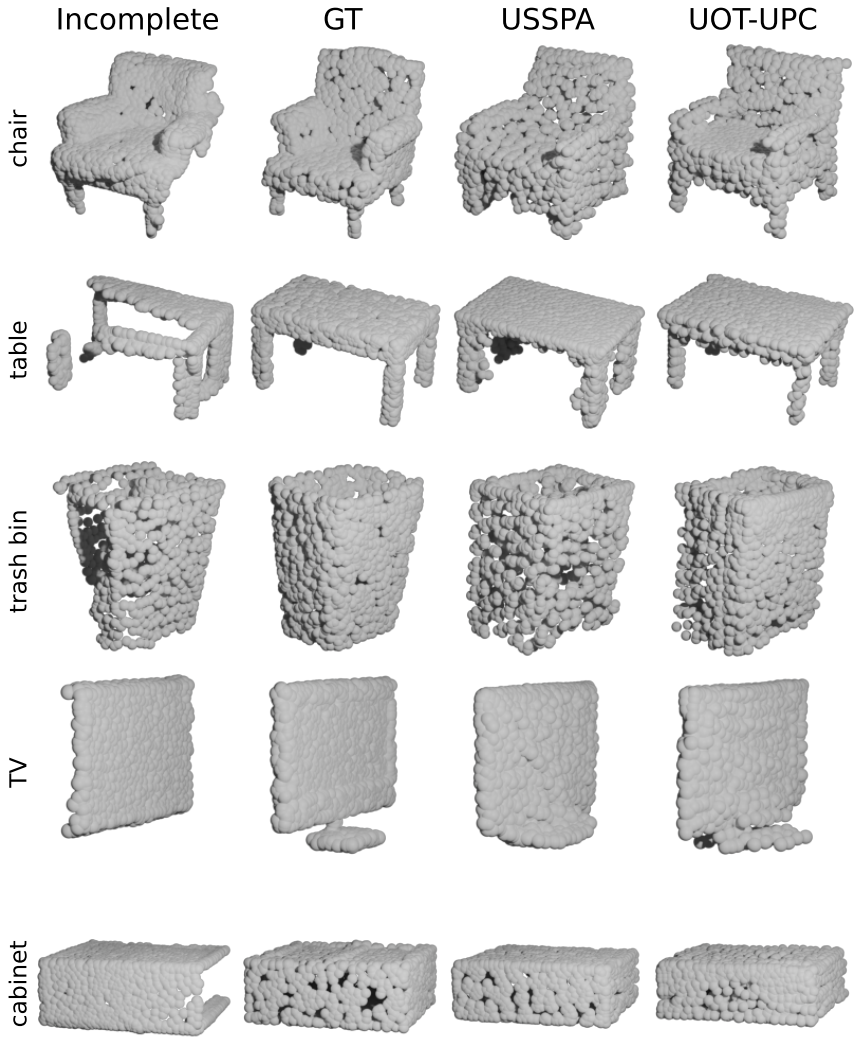

Figure 4: **Comparison of generated samples** from our UOT-UPC and USSPA in the multi-category setting.

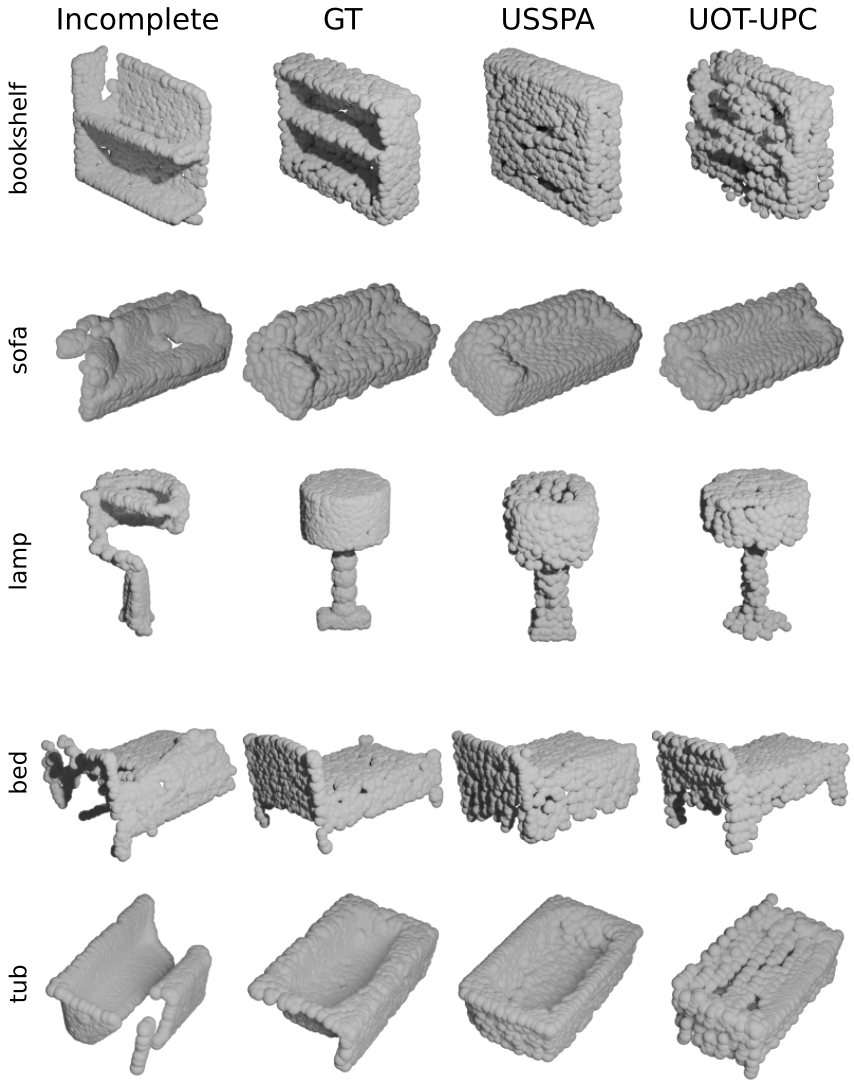

Figure 5: **Comparison of generated samples** from our UOT-UPC and USSPA in the multi-category setting.

B.2 COMPARISON OF CLASS IMBALANCE ROBUSTNESS FOR DIVERSE CLASS COMBINATIONS.

Table 4: **Comparison of class imbalance robustness** ($cd^{l1} \times 10^2$ (↓)) between UOT-UPC (ours), USSPA, and OT-UPC on diverse class combinations (Data1, Data2). Our UOT-UPC consistently outperforms other models across a wide range of class imbalance ratios in both additional class settings.

(a) (Data1, Data2) = (Lamp, Trash bin) with sample count = (1.1 : 8.0 * r).

r	0.3	0.5	0.7	1
USSPA	10.16	9.49	10.21	10.21
OT	25.68	21.95	28.41	25.36
Ours	**9.42**	**9.48**	**9.57**	**9.44**

(b) (Data1, Data2) = (Lamp, Bed) with sample count = (1.1 : 2.9 * r).

r	0.3	0.5	0.7	1
USSPA	9.64	9.78	9.27	9.79
OT	18.99	21.23	19.27	22.12
Ours	**8.95**	**8.91**	**8.98**	**8.73**

C RELATED WORKS

Unpaired point completion models have developed following recent advancements in unsupervised learning. Unpaired (Chen et al., 2020) introduces a GAN-based framework that aligns latent representations of incomplete point cloud with those of complete point cloud. Wu et al. (2020) proposes a conditional GAN approach for generating multiple plausible completions conditioned on the incomplete data. ShapeInv (Zhang et al., 2021) employs an optimization-based GAN-inversion technique (Xia et al., 2022) that identifies the optimal generator input noise yielding a complete point cloud from an incomplete point cloud. Cycle4 (Wen et al., 2021) employs dual cyclic transformations between the latent spaces of incomplete and complete point cloud through missing region encoding. USSPA (Ma et al., 2023) proposes a symmetric shape-preserving method based on GAN.

In this paper, we propose an unbalanced optimal transport approach to unpaired point cloud completion. To the best of our knowledge, this constitutes the first attempt to leverage the unbalanced optimal transport map within unpaired point cloud completion.