Toward Robust Neural Reconstruction from Sparse Point Sets Supplementary Material

Amine Ouasfi, Shubhendu Jena, Eric Marchand, Adnane Boukhayma
Inria, Univ. Rennes, CNRS, IRISA

1. Background on Distributionally Robust Optimization

Distributionally Robust Optimization (DRO) was initially introduced by [21] and has since become a significant framework for addressing uncertainty in decision-making [4, 9] . The DRO framework operates by defining an uncertainty set $\mathcal{U}$, typically modeled as a ball of radius ϵ around an empirical distribution $\hat{Q}_n$, such that $\mathcal{U} = \{Q : d(Q, \hat{Q}_n) \leq \epsilon\}$. The specific choice of the divergence measure greatly influences both the required size of ϵ and the tractability of the resulting optimization problem. The loss function is miminized under the worst-case distribution $Q \in \mathcal{U}$ in terms of the expected loss.

In machine learning, two primary divergence measures are widely adopted: f-divergences and the Wasserstein distance. With f-divergences [3, 14, 16] convex optimization techniques are usually leveraged to define tractable uncertainty sets. Alternatively, the Wasserstein distance [5, 15] is based on a metric over the data space, enabling the inclusion of distributions with supports different from the empirical distribution, thereby offering robustness to unseen data. However, the computational complexity of Wasserstein-based DRO makes it more challenging to handle. To address these challenges, various studies have proposed tractable methods for specific uncertainty sets and loss functions. For instance, [5, 15, 22] provide practical approaches for solving DRO problems with uncertainty regions defined by Wasserstein balls. For smooth loss functions, [23] proposes an efficient formulation for certifying robustness under Wasserstein uncertainty sets. Furthermore, the Unified DRO framework (UDR) introduced by [7] establishes a connection between Wasserstein DRO and adversarial training (AT) methods, offering a novel approach where the dual variable of the DRO problem is adaptively learned during training. This contrasts with [23] , where this parameter is fixed. [1, 24] study the DRO problem using the Sinkhorn Distance instead of the Wasserstein Distance providing efficient dual formulations.

	Sparse	Dense
SPSR [10]	2.27	1.25
DIGS [2]	0.68	0.19
OG-INR [11]	0.85	0.20
NTPS [8]	0.73	-
NP [12]	0.58	0.23
SparseOcc [18]	0.49	0.20
NAP [17]	0.49	0.19
Ours (WDRO)	0.51	0.20
Ours (SDRO)	0.48	0.21

Table 1. Ablation of point cloud density

2. Additional Ablative Analysis

2.1. Varying the point cloud density

In order to assess the performance of our method under various point cloud densities we perform an ablative analysis on the SRB benchmark [25]. We present quantitative results for both 1024-sized and dense input point clouds. In the dense setting, we report results from OG-INR. Our distributionally robust training strategy outperforms competitors in the sparse case and performs on par with the state-of-the-art in the dense case. Importantly, we observe considerable improvement over our baseline (NP) in both scenarios. Fig. 1 visually supports these results, illustrating reconstructions for sparse and dense inputs. In the dense setting, our method captures finer details, emphasized by the red boxes. These results highlight the practical advantages of our approach, even for dense inputs. Interestingly, our ablative analysis reveals that for dense inputs, WDRO may exhibit slightly better performance compared to SDRO. This result is not surprising, given that WDRO is certified to effectively hedge against small perturbations [23]. Consequently, as the input becomes denser, the noise on the labels due to input sparsity diminishes, thereby favoring WDRO.

2.2. Hyperparameter Analysis

In order to determine the hyperparameters of our proposed approach (SDRO), We performed a hyperparameter search on the SRB [25] benchmark utilizing the chamfer distance between the reconstruction and the input point cloud as a

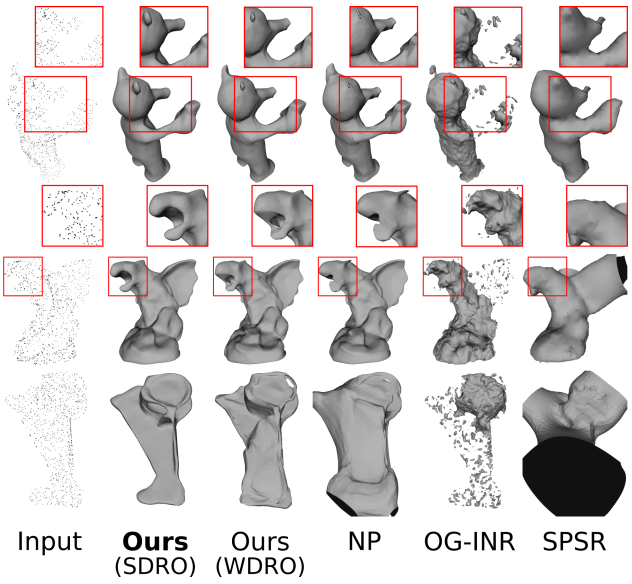

Figure 1. SRB [25] unsupervised reconstructions from sparse (1024 pts) unoriented point clouds without data priors.

validation metric. For the remaining datasets, we employed the same hyperparameters.

We carry out here an ablation study where we vary each one of the hyperparameters λ and ρ while fixing the remaining ones in order to better understand the behavior of our approach (SDRO) and its sensitivity to the choice of these hyperparameters.

Regularization parameter λ. This parameter controls how close the worst-case distribution Q' is to the nominal distribution. Fig. 3 illustrates how a very high value for this parameter minimizes the regularization impacts of SDRO by maintaining the worst-case samples around the nominal samples. Conversely, excessively low values lead to overly pessimistic estimations over-smoothing the results, despite greatly improving over the NP baseline.

Regularization parameter ρ. This parameter is responsible for the strength of the entropic regularization: it controls how the SDRO worst case distribution is concentrated around the support points of WDRO worst case distribution [24]. Consequently, it has to be defined such that it facilitates finding challenging distributions around the surface while maintaining a useful supervision signal. According to Fig. 2, it is important to utilize a sufficiently high ρ value in order to hedge against the right family of distributions. Contrastively, very high values can result in increased variance. Notice that ρ_{avg} here corresponds to average σ_p over the input points $\mathbf{P}$.

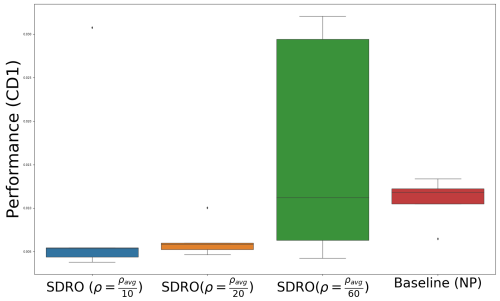

Figure 2. Ablation of the regularization parameter ρ.

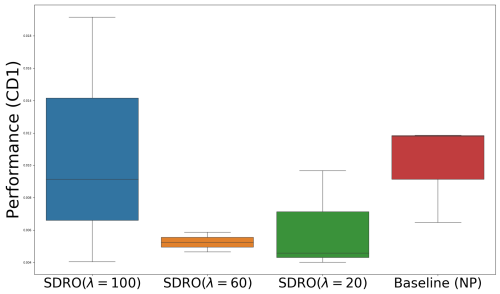

Figure 3. Ablation of the regularization parameter λ.

Algorithm 1 The training procedure of our method with WDRO.

Input: Point cloud $\mathbf{P}$, learning rate α, number of iterations N_{it}, batch size N_b.
WDRO hyperparameters: ϵ, σ_0, α_{wdro}, N_{it}^{wdro}, η_λ.
Output: Optimal parameters θ^*.
Compute local st. devs. $\{\sigma_p\}$ ($\sigma_p = \max_{t \in Knn(p,\mathbf{P})} ||t - p||_2$).
$\mathfrak{Q} \leftarrow$ sample($\mathbf{P}$,$\{\sigma_p\}$). (Equ. **??**)
Compute nearest points in $\mathbf{P}$ for all samples in $\mathfrak{Q}$.
Initialize $\lambda_1 = \lambda_2 = 1$.
Initialize λ.
for N_{it} times **do**
  Sample N_b query points $\{q, q \sim Q\}$.
  Initialize N_b points $\{q'\}$, ($q' \sim \mathcal{N}(q, \sigma_0 \mathbf{I}_3)$).
  for N_{it}^{wdro} times **do**
    $q' \leftarrow q' + \alpha_{wdro} \nabla_{q'} [\mathcal{L}(\theta, q') - \lambda c(q, q')]$
  end for
  $\lambda \leftarrow \lambda - \eta_\lambda \left(\epsilon - \frac{1}{N_b} \sum_{i=1}^{N_b} c\,(q_i', q_i)\right)$
  Compute WDRO losses $\{\mathcal{L}_{\text{WDRO}}(\theta, q)\}$ (Equ. **??**)
  Compute combined losses $\{\mathfrak{L}(\theta, q)\}$ (Equ. **??**)
  $(\theta, \lambda_1, \lambda_2) \leftarrow (\theta, \lambda_1, \lambda_2) - \alpha \nabla_{\theta,\lambda_1,\lambda_2} \Sigma_q \mathfrak{L}(\theta, q)$
end for

3. Training algorithm for WDRO

We provide in Algorithm 1 the detailed training procedure for WDRO.

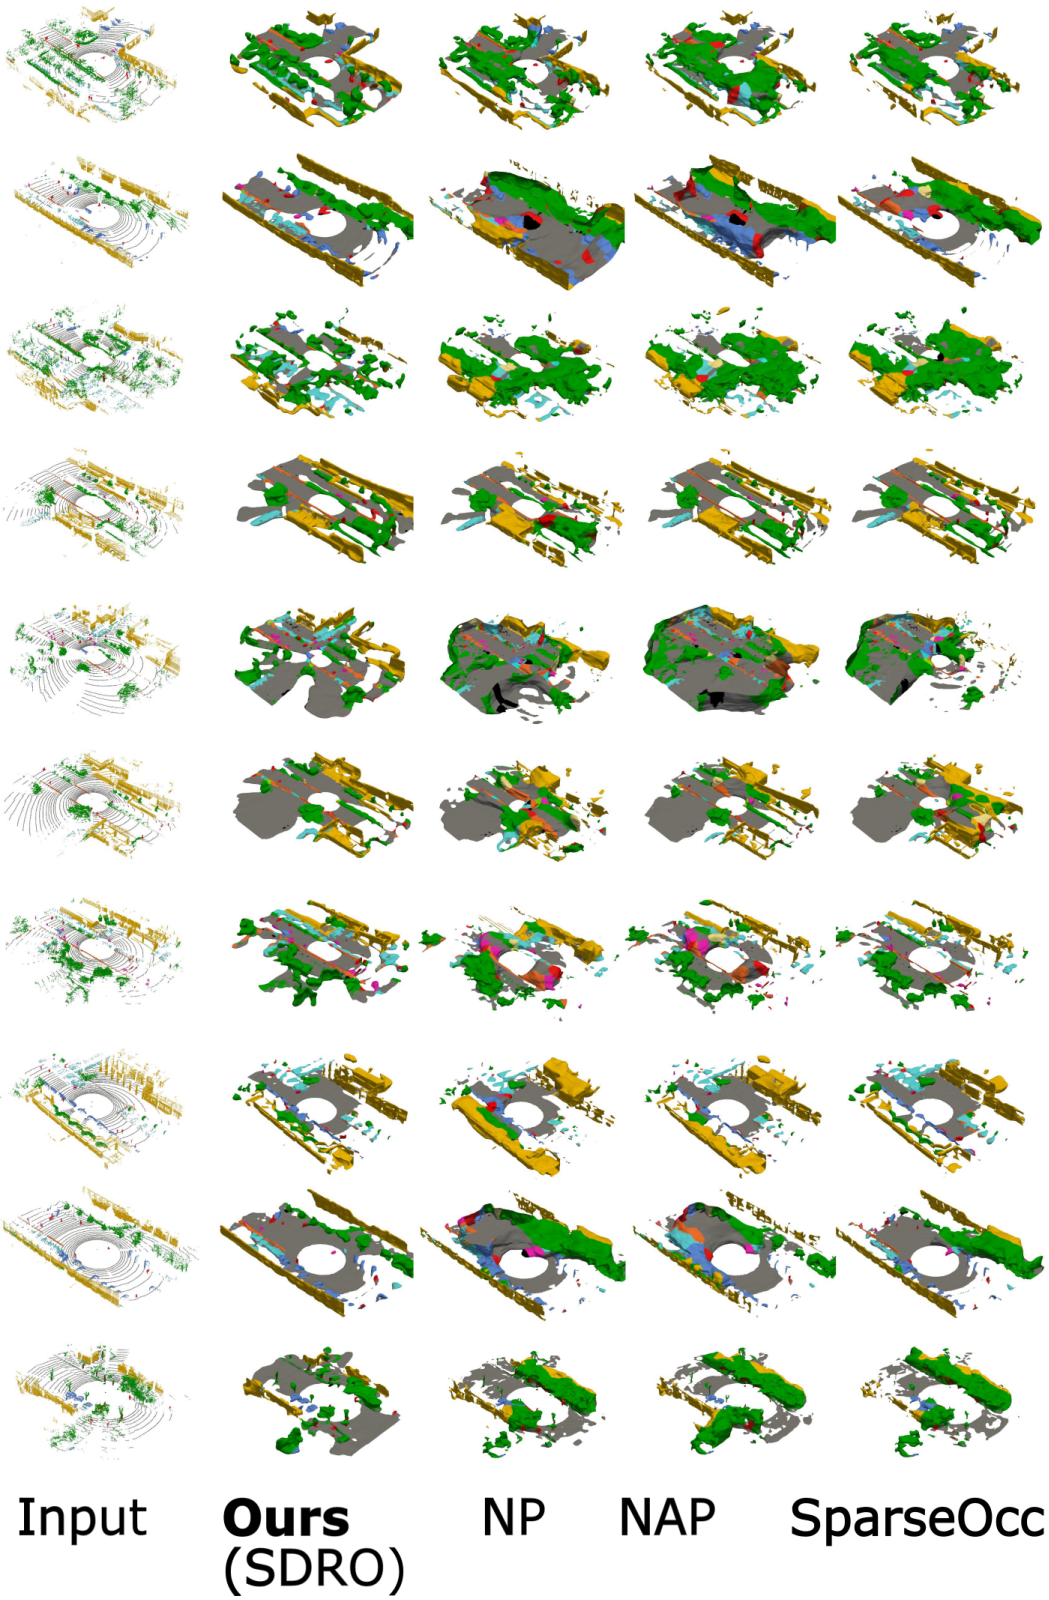

Figure 4. SemanticPOSS [19] reconstructions from road scene LiDAR data.

4. Additional Qualitative Results

We provide additional qualitative comparisons using SemanticPOSS road scene LiDAR data. Fig. 4 highlights the supe-

riority of our method in this challenging scenario compared to NAP and SparseOcc. This is particularly evident in highly noisy regions, such as trees, where these methods struggle, whereas our SDRO approach demonstrates robust performance.

5. Evaluation Metrics

Building on the definitions provided in [6] and [25], we present the formal definitions of the metrics used for evaluation in the main submission. Let $\mathcal{S}$ and $\hat{\mathcal{S}}$ denote the ground truth and predicted meshes, respectively. Following [8], all metrics are approximated using 100k samples drawn from $\mathcal{S}$ and $\hat{\mathcal{S}}$ for ShapeNet and Faust, and 1M samples for 3DScene. For SRB, we also utilize 1M samples, as suggested by [2] and [11].

Chamfer Distance (CD1) The L_1 Chamfer Distance is computed using the two-way nearest-neighbor distance::

$$\mathrm{CD}_1 = \frac{1}{2|\mathcal{S}|}\sum_{v\in\mathcal{S}}\min_{\hat{v}\in\hat{\mathcal{S}}}\|v-\hat{v}\|_2 + \frac{1}{2|\hat{\mathcal{S}}|}\sum_{\hat{v}\in\hat{\mathcal{S}}}\min_{v\in\mathcal{S}}\|\hat{v}-v\|_2.$$

Chamfer Distance (CD2) The L_2 Chamfer Distance is computed using the two-way nearest-neighborr squared distance:

$$\mathrm{CD}_2 = \frac{1}{2|\mathcal{S}|}\sum_{v\in\mathcal{S}}\min_{\hat{v}\in\hat{\mathcal{S}}}\|v-\hat{v}\|_2^2 + \frac{1}{2|\hat{\mathcal{S}}|}\sum_{\hat{v}\in\hat{\mathcal{S}}}\min_{v\in\mathcal{S}}\|\hat{v}-v\|_2^2.$$

F-Score (FS) For a given threshold τ, the F-Score between the ground truth mesh $\mathcal{S}$ and the predicted mesh $\hat{\mathcal{S}}$ is defined as:

$$\mathrm{FS}\left(\tau,\mathcal{S},\hat{\mathcal{S}}\right) = \frac{2\ \mathrm{Recall}\cdot\mathrm{Precision}}{\mathrm{Recall}\ +\ \mathrm{Precision}},$$

where

$$\mathrm{Recall}\left(\tau,\mathcal{S},\hat{\mathcal{S}}\right) = \left|\ \left\{v\in\mathcal{S},\ \text{s.t.}\ \min_{\hat{v}\in\hat{\mathcal{S}}}\|v-\hat{v}\|_2\langle\tau\right\}\ \right|,$$
$$\mathrm{Precision}\left(\tau,\mathcal{S},\hat{\mathcal{S}}\right) = \left|\ \left\{\hat{v}\in\hat{\mathcal{S}},\ \text{s.t.}\ \min_{v\in\mathcal{S}}\|v-\hat{v}\|_2\langle\tau\right\}\ \right|.$$

Following [13] and [20], we set τ to 0.01.

Normal consistency (NC) measures the alignment of surface normals between two meshes $\mathcal{S}$ (ground truth) and $\hat{\mathcal{S}}$ (prediction). Denoting the normal at a point v in $\mathcal{S}$ by n_v, it is defined as

$$\mathrm{NC} = \frac{1}{2|\mathcal{S}|}\sum_{v\in\mathcal{S}} n_v\cdot n_{\mathrm{closest}(v,\hat{\mathcal{S}})} + \frac{1}{2|\hat{\mathcal{S}}|}\sum_{\hat{v}\in\hat{\mathcal{S}}} n_{\hat{v}}\cdot n_{\mathrm{closest}(\hat{v},\mathcal{S})},$$

where

$$\mathrm{closest}(v,\hat{\mathcal{S}}) = \mathrm{argmin}_{\hat{v}\in\hat{\mathcal{S}}}\|v-\hat{v}\|_2.$$

References

[1] Waïss Azizian, Franck Iutzeler, and Jérôme Malick. Regularization for wasserstein distributionally robust optimization. *ESAIM: Control, Optimisation and Calculus of Variations*, 29:33, 2023. 1

[2] Yizhak Ben-Shabat, Chamin Hewa Koneputugodage, and Stephen Gould. Digs: Divergence guided shape implicit neural representation for unoriented point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19323–19332, 2022. 1, 4

[3] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013. 1

[4] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, 167: 235–292, 2018. 1

[5] Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019. 1

[6] Alexandre Boulch and Renaud Marlet. Poco: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6302–6314, 2022. 4

[7] Tuan Anh Bui, Trung Le, Quan Tran, He Zhao, and Dinh Phung. A unified wasserstein distributional robustness framework for adversarial training. *arXiv preprint arXiv:2202.13437*, 2022. 1

[8] Chao Chen, Zhizhong Han, and Yu-Shen Liu. Unsupervised inference of signed distance functions from single sparse point clouds without learning priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 4

[9] Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010. 1

[10] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *TOG*, 2013. 1

[11] Chamin Hewa Koneputugodage, Yizhak Ben-Shabat, and Stephen Gould. Octree guided unoriented surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16717–16726, 2023. 1, 4

[12] Baorui Ma, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Neural-pull: Learning signed distance functions from point clouds by learning to pull space onto surfaces. In *ICML*, 2021. 1

[13] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 4

[14] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015. 1

[15] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018. 1

[16] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *NIPS*, pages 2208–2216, 2016. 1

[17] Amine Ouasfi and Adnane Boukhayma. Few-shot unsupervised implicit neural shape representation learning with spatial adversaries. *arXiv preprint arXiv:2408.15114*, 2024. 1

[18] Amine Ouasfi and Adnane Boukhayma. Unsupervised occupancy learning from sparse point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21729–21739, 2024. 1

[19] Yancheng Pan, Biao Gao, Jilin Mei, Sibo Geng, Chengkun Li, and Huijing Zhao. Semanticposs: A point cloud dataset with large quantity of dynamic instances. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 687–693, 2020. 3

[20] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision*, pages 523–540. Springer, 2020. 4

[21] Herbert E Scarf, KJ Arrow, and S Karlin. *A min-max solution of an inventory problem*. Rand Corporation Santa Monica, 1957. 1

[22] Soroosh Shafieezadeh Abadeh, Peyman M Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. *Advances in neural information processing systems*, 28, 2015. 1

[23] Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017. 1

[24] Jie Wang, Rui Gao, and Yao Xie. Sinkhorn distributionally robust optimization. *arXiv preprint arXiv:2109.11926*, 2021. 1, 2

[25] Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, and Daniele Panozzo. Deep geometric prior for surface reconstruction. In *CVPR*, 2019. 1, 2, 4