

Learning minimal volume uncertainty ellipsoids

Itai Alon, David Arnon and Ami Wiesel

Abstract—We consider the problem of learning uncertainty regions for parameter estimation problems. The regions are ellipsoids that minimize the average volumes subject to a prescribed coverage probability. As expected, under the assumption of jointly Gaussian data, we prove that the optimal ellipsoid is centered around the conditional mean and shaped as the conditional covariance matrix. In more practical cases, we propose a differentiable optimization approach for approximately computing the optimal ellipsoids using a neural network with proper calibration. Compared to existing methods, our network requires less storage and less computations in inference time, leading to accurate yet smaller ellipsoids. We demonstrate these advantages on four real-world localization datasets.

Index Terms—Uncertainty ellipsoid, covariance estimation, conformal prediction.

I. INTRODUCTION

We consider uncertainty quantification for parameter estimation. Estimating unknown parameters given noisy observations is a fundamental problem in statistical signal processing, including inverse problems and localization [1]. In critical decision-making tasks, quantifying the uncertainty associated with these estimates is crucial. Accurate uncertainty regions are also necessary for developing more robust signal processing algorithms [2], [3], [4]. Traditionally, the estimation algorithms and their uncertainty regions have been designed and analyzed based on statistical models. Recently, there has been a switch to data-driven methods based on machine learning [5], [6], [7], [8]. This letter focuses on deep learning solutions to uncertainty quantification as the natural next step.

The standard approach to uncertainty quantification is to associate an uncertainty region, typically an ellipsoid, around each estimate. Ideally, the ellipsoids should have small volumes and cover the true unknown parameters. The goal is to trade off these two competing properties optimally. Classically, in model-based estimation, the shape of the ellipsoid is characterized by the likelihood, asymptotic covariance matrix, and Fisher Information matrix. The volume is traditionally computed assuming a multivariate Gaussian distribution [9], [10]. Linearization methods can address more complex uncertainty sources [11]. Other approaches rely on perturbations of the observations and an analysis of the resulting deviations [12], [13].

Recently, there has been a switch to data-driven estimation based on fitting an algorithm concerning a training set. In uncertainty quantification, the main idea is to compare the current measurement to similar examples in this set and rely

on their uncertainties. Following this line of thought, many works are variants of nearest neighbors algorithms [14], [15]. The successful CELLO approaches rely on fast estimating the covariance matrices of the errors [16], [17]. In particular, uncertainty methods typically allow imperfect coverage and a small percentage of examples outside the uncertainty region. In this sense, the problem is also related to a large body of literature on robust covariance estimation [18], [19], [20], [10], [21].

Most of the previous data-driven methods lack formal guarantees and do not always satisfy the required coverage in practice. To close this gap, conformal prediction is a distribution-free approach that uses an independent calibration phase that provides validity for any desired coverage level [22], [23], [24], [25]. Cross-conformal prediction [26] and cross-validation [27] generalize these ideas and utilize all the available data for fitting and calibration. Initially, these ideas were focused on classification problems and scalar regressions. Recent advances also address the multivariate case by marginal intervals that result in rectangular uncertainty regions [28]. Closest to our work is the extension of conformal prediction to calibration of ellipsoid uncertainty regions based on nearest neighbors [29], [30].

We present a unified Learning Minimum Volume Ellipsoids (LMVE) framework that considers the volume vs. coverage tradeoff. We demonstrate that the optimal shapes of the ellipsoids are conditional covariances in the Gaussian case. We propose a deep learning solution for more realistic settings that combines nearest-neighbor approaches, covariance estimation, and conformal prediction. The proposed LMVE neural network generates ellipsoids with minimal average volume and prescribed coverage probability. It significantly reduces memory and computation resources, improves accuracy, and is demonstrated on four real-world localization datasets.

The fully implemented code can be found in <https://github.com/ItaiAlon/LMVE/>.

II. PROBLEM FORMULATION

Let $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^n$ be a features vector and a labels vector drawn from a joint distribution $\mathcal{X} \times \mathcal{Y}$. An uncertainty ellipsoid on y given x is defined as:

$$\mathcal{E}(\mu(x), C(x)) = \{y : (y - \mu(x))^T C^{-1}(x)(y - \mu(x)) \leq 1\} \quad (1)$$

where $\mu(x)$ is the center of the ellipsoid and $C(x) \succ 0$ characterizes its shape. Generally, ellipsoids are defined by both the center and the shape. Still, in many applications, assuming that the center is fixed and given from a previous prediction stage (either model-based or data-driven) is more reasonable. Therefore, most of this letter will focus on the shape $C(x)$.

Manuscript received October, 2023.

The authors are with the School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel (e-mail: itai.alon1@mail.huji.ac.il; amiw@cs.huji.ac.il).

The second author is with Rafael Advanced Defence Systems - Research, Development and Engineering, Israel (e-mail: david.arnon@gmail.com).

This research was partially supported by ISF grant number 2672/21.

Ideally, the ellipsoids should satisfy two competing properties. First, to ensure validity, each ellipsoid should satisfy the coverage constraint defined as

$$\Pr[y \in \mathcal{E}(\mu(x), C(x))] \geq \eta \quad (2)$$

and $\eta \in [0, 1)$ a prescribed coverage probability. Second, it should minimize the uncertainty defined by the volume

$$\text{Vol}_{\mathcal{E}}(x) = \text{const} \cdot \det(C(x))^{\frac{1}{2}} \quad (3)$$

where the constant is the volume of the unit ball in that dimension. To formulate this tradeoff, we use a constrained optimization and minimize the expected volume subject to a coverage constraint

$$\begin{aligned} (MVE) : \quad & \min_{\mathcal{E}(\mu(\cdot), C(\cdot))} \mathbb{E}[\text{Vol}_{\mathcal{E}}(x)] \\ & \text{s.t.} \quad \Pr[y \in \mathcal{E}] \geq \eta. \end{aligned} \quad (4)$$

For completeness, our uncertainty ellipsoids assume a Bayesian setting in which both x and y are random, and the expectations are computed. This differs from the non-Bayesian definition of confidence regions, which assumes that y are deterministic unknown and will be addressed in future work.

The main goal of this paper is to develop a method for approximately solving (MVE) using a neural network. At inference time, it will take x as an input and output a center $\mu(x)$ and a shape matrix $C(x)$ that defines the uncertainty ellipsoid associated with the corresponding unknown y . We assume access to independent and identically distributed (i.i.d.) pairs of examples (x_i, y_i) for $i = 1 \dots m$ to train this network.

III. THEORY

In this section, we theoretically analyze the optimal solution to (MVE) in the simplistic Gaussian setting.

We begin by considering a single ellipsoid (without any features x). This result embodies a widely recognized conventional wisdom, yet we are unaware of a previous formal proof.

Lemma 1. *Let $\eta \in [0, 1]$ and $y \sim \mathcal{N}(\mu, \Sigma)$ with dimension n . Consider the optimization*

$$\begin{aligned} \min_{\mathcal{E}} \quad & \text{Vol}(\mathcal{E}) \\ \text{s.t.} \quad & \Pr[y \in \mathcal{E}] \geq \eta \end{aligned} \quad (5)$$

The optimal argument is

$$\mathcal{E}(\mu, F_{\chi_n^2}^{-1}(\eta) \cdot \Sigma) \quad (6)$$

where $F_{\chi_n^2}^{-1}(\cdot)$ is the inverse chi-square cdf with n degrees of freedom, and the optimal value is $F_{\chi_n^2}^{-1}(\eta)^n \cdot \text{Vol}(\Sigma)$.

Proof. Assume an optimal solution $\mathcal{E}(\hat{\mu}, \hat{\Sigma})$ and let κ satisfy $\text{Vol}(\kappa\Sigma) = \text{Vol}(\hat{\Sigma})$. We only care about the symmetric difference, so denote:

$$\begin{aligned} S_- &:= \mathcal{E}(\mu, \kappa\Sigma) \setminus (\mathcal{E}(\hat{\mu}, \hat{\Sigma}) \cap \mathcal{E}(\mu, \kappa\Sigma)) \\ S_+ &:= \mathcal{E}(\hat{\mu}, \hat{\Sigma}) \setminus (\mathcal{E}(\hat{\mu}, \hat{\Sigma}) \cap \mathcal{E}(\mu, \kappa\Sigma)) \end{aligned} \quad (7)$$

By the choice of κ we get:

$$\text{Vol}(\kappa\Sigma) = \text{Vol}(\hat{\Sigma}) \implies S := \text{Vol}(S_-) = \text{Vol}(S_+) \quad (8)$$

The Gaussian density has an equal value on the edge of $\mathcal{E}(\mu, \kappa\Sigma)$, has a larger value inside, and has a smaller value outside. Therefore,

$$p(y) \geq p \quad \forall y \in S_- \quad , \quad p(y) < p \quad \forall y \in S_+ \quad (9)$$

where

$$p := \min_{y \in \mathcal{E}(\mu, \kappa\Sigma)} p(y). \quad (10)$$

Thus

$$\begin{aligned} & \int_{y \in \mathcal{E}(\mu, \kappa\Sigma)} p(y) \, dy - \int_{y \in \mathcal{E}(\hat{\mu}, \hat{\Sigma})} p(y) \, dy \\ &= \int_{y \in S_-} p(y) \, dy - \int_{y \in S_+} p(y) \, dy \\ &\geq S \cdot \left(\inf_{y \in S_-} p(y) - \sup_{y \in S_+} p(y) \right) \\ &\geq S \cdot (p - p) = 0 \end{aligned} \quad (11)$$

with equality if and only if all parameters are equal. This means that $S_- = S_+ = \emptyset$ and the optimal parameters are $\mu = \hat{\mu}$ and $\kappa\Sigma = \hat{\Sigma}$.

Finally, it remains to choose the minimum scaling that satisfies the constraint, namely $\kappa = F_{\chi_n^2}^{-1}(\eta)$. \square

Next, we turn to a more practical estimation problem with features x and unknown labels y .

Theorem 1. *Let x and y be jointly Gaussian, then the optimal solution to (MVE) is*

$$\begin{aligned} \mu(x) &= \mathbb{E}[y|x] \\ C(x) &= \kappa(x) \cdot \mathbb{E}[(y - \mu(x))(y - \mu(x))^T] \end{aligned} \quad (12)$$

$\kappa(x) > 0$ is a scaling factor that satisfies the MVE coverage constraint.

Proof. The conditional Gaussian distribution is given by

$$p(y|x) = \mathcal{N}(y; \mu(x), C_{y|x}) \quad (13)$$

where $C_{y|x}$ does not depend on x . Problem (MVE) can be expressed as

$$\begin{aligned} \min_{\mathcal{E}(x)} \quad & \int_{\mathcal{X}} p(x) \cdot \text{Vol}_{\mathcal{E}}((x)) \, dx \\ \text{s.t.} \quad & \int_{\mathcal{X}} p(x) \cdot \int_{y \in \mathcal{E}(x)} p(y|x) \, dy \, dx \geq \eta \end{aligned} \quad (14)$$

Introducing auxiliary variables, we can rewrite the constraint

$$\begin{cases} \Pr[y \in \mathcal{E}|x] \geq \bar{\eta}(x) \quad \forall x \in \mathcal{X} \\ \int_{\mathcal{X}} p(x) \cdot \bar{\eta}(x) \, dx \geq \eta \end{cases} \quad (15)$$

We can apply Lemma 1 for each x separately, and its solution will always be the same $C_{y|x}$ times a scalar factor that depends on x . Together, the optimization problem is equal to

$$\begin{aligned} \min_{\bar{\eta}(x)} \quad & \int_{\mathcal{X}} p(x) \cdot F_{\chi_n^2}^{-1}(\bar{\eta}(x))^n \, dx \cdot \text{Vol}(C_{y|x}) \\ \text{s.t.} \quad & \int_{\mathcal{X}} p(x) \cdot \bar{\eta}(x) \, dx \geq \eta. \end{aligned} \quad (16)$$

\square

The theorem formally shows the relation between uncertainty regions and covariance estimation in the joint Gaussian distribution. In this simplistic case, the shape of the ellipsoids is defined by the conditional covariances. These are independent of the conditioning, and the dependence on x is only through the calibration factor. This is not true in general. In other distributions, there is no clear relation between the shapes and the covariances of the error and the shapes typically depend on the values of x . The next section proposes a neural network solution that addresses these more practical settings.

IV. LEARNING CONFIDENCE ELLIPSOID NETS

In this section, we propose LMVE, a neural network-based solution to (MVE). We rely on a standard neural architecture with fully connected layers to ensure rich expressibility, generalization, and low computational complexity at inference time.

LMVE is defined as

$$\begin{aligned} C_\theta(x) &= R_\theta(x)^T R_\theta(x) + \epsilon I \\ R_\theta(x) &= L_3 \cdot \sigma(L_2 \cdot \sigma(L_1 \cdot x + b_1) + b_2) + b_3 \end{aligned} \quad (17)$$

where $L_1 \in \mathbb{R}^{4d \times d}$, $L_2 \in \mathbb{R}^{d \times 4d}$, $L_3 \in \mathbb{R}^{n^2 \times d}$, $b_1 \in \mathbb{R}^{4d}$, $b_2 \in \mathbb{R}^d$, $b_3 \in \mathbb{R}^{n^2}$ and $\theta = (L_1, L_2, L_3, b_1, b_2, b_3)$. The operator σ is the ReLU activation function with dropout, and $\epsilon > 0$ is a regularization parameter that ensures positive definiteness.

Problem (MVE) is challenging for learning due to its non-differentiability and strict constraint. Therefore, LMVE is based on three phases: initialization, training, and calibration. We divide the available data into two subsets, samples z_1, \dots, z_{m_t} for the initialization and training (that is, training set) and samples $z_{m_t+1}, \dots, z_{m_t+m_c}$ for the calibration (namely validation set), where $m = m_t + m_c$.

First, in the initialization phase, we imitate an existing baseline. We use the baseline outputs as approximate labels and train our network to approximate them using a standard Mean Squared Error (MSE) loss. The baseline and the initialization phase only use samples from the training set.

Second, in the training phase, we consider (MVE) using its Lagrange penalized form directly

$$\min_{\theta} \Pr[y \notin \mathcal{E}(\mu(x), C_\theta(x))] + \lambda \mathbb{E}[\text{Vol}_\theta(x)] \quad (18)$$

We replace the probability and expectation with empirical averages over the training set. We tried different smooth surrogates for the probability function but found that the best results were obtained by minimizing the average Mahalanobis distance. Together, the training loss is

$$\begin{aligned} \min_{\theta} \frac{1}{T} \sum_{i=1}^T \ell(x_i, y_i; \theta) \\ \ell(x, y; \theta) := M_\theta(x, y) + \lambda \det(C_\theta(x))^{\frac{1}{2}} \\ M_\theta(x, y) := (y - \mu(x))^T C_\theta^{-1}(x)(y - \mu(x)) \end{aligned} \quad (19)$$

where λ is a hyperparameter that scales between the volume's importance and accuracy. To tune it, we rely again on the

existing calibrated baseline and choose lambda to balance the two terms:

$$\lambda = \frac{\frac{1}{T} \sum_{j=1}^T M_B(x_j, y_j)}{\frac{1}{T} \sum_{j=1}^T \det(C_B(x_j))} \quad (20)$$

where the ellipsoids are obtained from the calibrated baseline (denoted by subscript B). We emphasize that the baseline is trained and calibrated only using the training set to avoid dependency on the validation set.

Third, to ensure the strict coverage constraints, we use a calibration step using the theory of conformal prediction [30]. We evaluate our network on the validation samples and compute their Mahalanobis distances

$$\alpha_i = M_\theta(x_i, y_i) \quad i = m_t, \dots, m_t + m_c \quad (21)$$

We sort these distances and define α_q as the q 'th largest value where

$$q = \frac{\lceil (m_c + 1)\eta \rceil}{m_c} \quad (22)$$

We then rescale the ellipsoids by α_q so that the final LMVE shapes are defined as

$$C_{LMVE}(x) = \alpha_q \cdot C_\theta(x) \quad (23)$$

This ensures that

$$\Pr[y \in \mathcal{E}_{LMVE}(x)] \geq \eta \quad (24)$$

as required.

V. EXPERIMENTS

In this section, we compare the algorithms on different real-world datasets.

Algorithms: In all the methods, the centers of the ellipsoids μ are computed identically using a multi-output Support Vector Regressor (SVR) fitted on the training set. We compare three algorithms for generating ellipsoids. The shapes $C(x)$ are fitted on the training set and then rescaled to satisfy the coverage constraint on the validation set as defined in 23. The algorithms are:

- Gaussian Ellipsoid (GE): the sample version of (12) defined as

$$\hat{C}_{GE}(x) = \frac{1}{T} \sum_{i=1}^T (y_i - \mu_i)(y_i - \mu_i)^T \quad (25)$$

We approximate the conditional mean using SVR for consistency with the other algorithms.

- NLE [30]: local sample covariance matrices of the 5% nearest neighbors $N(x)$ of each x with a small GE regularization

$$\begin{aligned} \hat{C}_{NLE}(x) = 0.95 \cdot \frac{1}{|N(x)|} \sum_{i \in N(x)} (y_i - \mu_i)(y_i - \mu_i)^T \\ + 0.05 \cdot \hat{C}_{GE} \end{aligned} \quad (26)$$

- Copula [28]: We also compare to the empirical copula method. It provides α_s that promises $\Pr \left[\frac{|y_i - \hat{y}_i|}{\sigma_i} \leq \alpha_s^{(i)} \forall i \in [d] \right] \geq \eta$ for a new sample y . We

TABLE I
RESULTS OVER 90% CALIBRATION.

	GE	NLE	LMVE	Copula
ble_rssi	89.8 (±4.8)	90.1 (±3.1)	89.6 (±4.3)	88.8 (±3.7)
enb	88.2 (±5.2)	89.8 (±4.2)	89.3 (±4.8)	88.9 (±5.4)
indoor_localization	89.8 (±1.0)	89.7 (±1.2)	89.6 (±1.0)	89.8 (±1.0)
residential_building	88.6 (±7.0)	88.5 (±6.5)	88.7 (±7.6)	83.9 (±6.6)
ble_rssi	10.2 (±1.5)	9.6 (±1.0)	7.8 (±1.8)	54.8 (±18.3)
enb	34.0 (±7.6)	18.3 (±3.4)	17.90 (±1.1)	263.9 (±118.4)
indoor_localization	$7.0 \cdot 10^3$ (±178.6)	$4.3 \cdot 10^3$ (±136.5)	$2.0 \cdot 10^3$ (±51.6)	$22.4 \cdot 10^3$ (±1355)
residential_building	$2.2 \cdot 10^5$ (±15.4 · 10 ⁴)	$1.3 \cdot 10^5$ (±7.8 · 10 ⁴)	$1.1 \cdot 10^5$ (±6.8 · 10⁴)	$69.9 \cdot 10^5$ (±6603 · 10 ⁴)

used the same net structure that was suggested in the article.

- LMVE: $C_{LMVE}(x)$ is the neural network output where hyperparameters are chosen to maximize performance on the validation set. Its initialization is based on the NLE baseline.

Datasets: We compare four real-world localization datasets, as detailed in Table II. The full data process, which explains the dimension difference, can be found on GitHub.

TABLE II
DATASETS

Name	Size	Input	Output	Source
ble_rssi	1,420	13	2	[31]
enb	768	8	2	[32]
indoor_localization	19,937	519	2	[33]
residential_building	372	103	2	[34]

Each dataset was divided into three parts: a training set (81%), a validation set (9%), and a test set (10%). We conducted 50 experiments and reported the average performance and standard deviation.

Results for 90% coverage are in Table I. The top section shows preferred coverage values close to 90%. The bottom section displays mean volume values, with smaller values indicating better performance. Strong calibration makes the methods accurate, but LMVE improves average volume. GE is the cheapest for storage and computation. NLE is expensive for large datasets. LMVE is efficient with a simple three-layer neural network. Further comparison is in Table III.

TABLE III
COMPUTATIONAL COMPLEXITY IN \mathcal{O} NOTION. N IS THE NUMBER OF NEIGHBORS IN THE NLE. L IS THE HIDDEN LAYER SIZE.

	Memory	Inference Time
GE	d^2	const
NLE	$m_t(n + d)$	$N \cdot (d \log(m_t) + n^2)$
LMVE	$d^2 + n^2$	$d^2 + n^2$
Copula	$dL + L^2 + Ln$	$dL + L^2 + Ln$

We tested more complex losses, such as hinge loss and robust generalized Gaussian-based losses, that smooth the probability function. However, they did not lead to significant improvements in some datasets and were therefore omitted.

We conducted ablation studies on a few datasets to comprehend the contribution of each component to the overall system. However, we only report some results in Table IV due to space limitations. The key conclusions are as follows:

TABLE IV

RESULTS OVER 90% CALIBRATION ON ENB [32] DATASET. THE RESULTS ARE COMPARED TO ACCURACY OF 89.325 (±4.883) AND EXPECTED VOLUME OF 17.90 (±1.184).

	LMVE
Without smart init	88.3 (±5.1)
Fixed $\lambda = 1$	88.8 (±4.7)
2 Layers	87.7 (±4.7)
50% Validation	89 (±4.1)
Without smart init	18.23 (±51.5)
Fixed $\lambda = 1$	11.40 (±4.3)
2 Layers	16.58 (±7.7)
50% Validation	17.37 (±10.5)

Smart initialization enhances the performance. Reducing the number of layers in architecture leads to a decline in coverage. The effect of lambda value is less significant and depends on the dataset. Finally, as expected, using a larger validation set improves the coverage but results in inaccuracies in the volume.

VI. DISCUSSION

This paper proposes a scalable deep learning framework called the LMVE neural network. This framework outputs uncertainty ellipsoids around predictions. LMVE combines CELLO and NLE and improves upon them by providing modern and efficient data modeling.

LMVE is an important deep-learning technique for estimating uncertainty ellipsoids, but many questions remain regarding its usage. Regarding problem formulation, LMVE only considers the average coverage of random labels. Extending its application to confidence regions around deterministic unknown labels would be interesting. Training LMVE is challenging, so future work should focus on improving initialization methods, optimization algorithms, and hyperparameter tuning approaches.

APPENDIX: TECHNICAL DETAILS ON LMVE

We fitted the models in PyTorch using dropout parameters of 0.1 or 0.5, and 200,000 iterations of an Adam optimizer (half for initialization and half for training) with rates $\{10^{-3}, 10^{-5}\}$.

ACKNOWLEDGE

We want to thank Erez Peterfreund, Roy Friedman, and Tzvi Diskin for the consulting.

REFERENCES

- [1] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.
- [2] A. M. Elbir, K. V. Mishra, S. A. Vorobyov, and R. W. Heath, "Twenty-five years of advances in beamforming: From convex and nonconvex optimization to learning techniques," *IEEE Signal Processing Magazine*, vol. 40, no. 4, pp. 118–131, 2023.
- [3] H. Luo, Q. Li, L. Yang, and J. Qin, "An efficient algorithm for robust fractional qcqp and its applications to multiuser beamforming with bounded channel uncertainties," *IEEE Transactions on Signal Processing*, vol. 70, pp. 6096–6111, 2022.
- [4] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust statistics for signal processing*. Cambridge University Press, 2018.
- [5] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, "Model-based deep learning," *Proceedings of the IEEE*, 2023.
- [6] S. Khobahi, N. Shlezinger, M. Soltanalian, and Y. C. Eldar, "Lord-net: Unfolded deep detection network with low-resolution receivers," *IEEE Transactions on Signal Processing*, vol. 69, pp. 5651–5664, 2021.
- [7] T. Diskin, Y. C. Eldar, and A. Wiesel, "Learning to estimate without bias," *IEEE Transactions on Signal Processing*, 2023.
- [8] Y. Dahan, G. Revach, J. Dunik, and N. Shlezinger, "Uncertainty quantification in deep learning based kalman filters," *arXiv preprint arXiv:2309.03058*, 2023.
- [9] A. Censi, "An accurate closed-form estimate of icp's covariance," in *Proceedings 2007 IEEE international conference on robotics and automation*. IEEE, 2007, pp. 3167–3172.
- [10] F. Pascal, P. Forster, J.-P. Ovarlez, and P. Larzabal, "Performance analysis of covariance matrix estimates in impulsive noise," *IEEE Transactions on signal processing*, vol. 56, no. 6, pp. 2206–2217, 2008.
- [11] M. Brossard, S. Bonnabel, and A. Barrau, "A new approach to 3d icp covariance estimation," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 744–751, 2020.
- [12] A. G. Buch, D. Kraft *et al.*, "Prediction of icp pose uncertainties using monte carlo simulation with synthetic depth images," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 4640–4647.
- [13] V. Zambianchi, F. Bassi, A. Calisti, D. Dardari, M. Kieffer, and G. Pasolini, "Distributed nonasymptotic confidence region computation over sensor networks," *IEEE transactions on Signal and Information Processing over Networks*, vol. 4, no. 2, pp. 308–324, 2017.
- [14] N. Wang and A. E. Raftery, "Nearest-neighbor variance estimation (nnve) robust covariance estimation via nearest-neighbor cleaning," *Journal of the American Statistical Association*, vol. 97, no. 460, pp. 994–1019, 2002.
- [15] C. Johnstone and B. Cox, "Conformal uncertainty sets for robust optimization," in *Conformal and Probabilistic Prediction and Applications*. PMLR, 2021, pp. 72–90.
- [16] W. Vega-Brown, A. Bachrach, A. Bry, J. Kelly, and N. Roy, "Cello: A fast algorithm for covariance estimation," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 3160–3167.
- [17] D. Landry, F. Pomerleau, and P. Giguere, "Cello-3d: Estimating the covariance of icp in the real world," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8190–8196.
- [18] E. Ollila, D. P. Palomar, and F. Pascal, "Shrinking the eigenvalues of m-estimators of covariance matrix," *IEEE Transactions on Signal Processing*, vol. 69, pp. 256–269, 2020.
- [19] E. Ollila and D. E. Tyler, "Regularized m-estimators of scatter matrix," *IEEE Transactions on Signal Processing*, vol. 62, no. 22, pp. 6059–6070, 2014.
- [20] Y. Sun, P. Babu, and D. P. Palomar, "Robust estimation of structured covariance matrix for heavy-tailed elliptical distributions," *IEEE Transactions on Signal Processing*, vol. 64, no. 14, pp. 3576–3590, 2016.
- [21] J. A. Fishbone and L. Mili, "Highly robust complex covariance estimators with applications to sensor array processing," *IEEE Open Journal of Signal Processing*, 2023.
- [22] V. Vovk, A. Gammerman, and C. Saunders, "Machine-learning applications of algorithmic randomness," 1999.
- [23] C. Saunders, A. Gammerman, and V. Vovk, "Transduction with confidence and credibility," 1999.
- [24] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.
- [25] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," *arXiv preprint arXiv:2107.07511*, 2021.
- [26] V. Vovk, "Cross-conformal predictors," *Annals of Mathematics and Artificial Intelligence*, vol. 74, no. 1, pp. 9–28, 2015.
- [27] R. F. Barber, "Is distribution-free inference possible for binary regression?" *Electronic Journal of Statistics*, vol. 14, no. 2, pp. 3487–3524, 2020.
- [28] S. Messoudi, S. Destercke, and S. Rousseau, "Copula-based conformal prediction for multi-target regression," *Pattern Recognition*, vol. 120, p. 108101, 2021.
- [29] S. Feldman, S. Bates, and Y. Romano, "Calibrated multiple-output quantile regression with representation learning," *arXiv preprint arXiv:2110.00816*, 2021.
- [30] S. Messoudi, S. Destercke, and S. Rousseau, "Ellipsoidal conformal inference for multi-target regression," in *Conformal and Probabilistic Prediction with Applications*. PMLR, 2022, pp. 294–306.
- [31] M. Mohammadi, A. Al-Fuqaha, M. Guizani, and J.-S. Oh, "Semisupervised deep reinforcement learning in support of iot and smart city services," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 624–635, 2017.
- [32] G. Tsoumakas, E. Spyromitros-Xioufis, J. Vilcek, and I. Vlahavas, "Mulan: A java library for multi-label learning," *The Journal of Machine Learning Research*, vol. 12, pp. 2411–2414, 2011.
- [33] J. Torres-Sospedra, R. Montoliu, A. Martínez-Usó, J. P. Avariento, T. J. Arnau, M. Bénédicto-Bordonau, and J. Huerta, "Ujiindoorloc: A new multi-building and multi-floor database for wlan fingerprint-based indoor localization problems," in *2014 international conference on indoor positioning and indoor navigation (IPIN)*. IEEE, 2014, pp. 261–270.
- [34] M. H. Rafiei and H. Adeli, "A novel machine learning model for estimation of sale prices of real estate units," *Journal of Construction Engineering and Management*, vol. 142, no. 2, p. 04015066, 2016.