Similarity-based prediction for channel mapping and user positioning

Luc Le Magoarou

► To cite this version:

Luc Le Magoarou. Similarity-based prediction for channel mapping and user positioning. IEEE Communications Letters, In press. hal-02974852v3

HAL Id: hal-02974852

<https://hal.science/hal-02974852v3>

Submitted on 5 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Similarity-based prediction for channel mapping and user positioning

Luc Le Magoarou

Abstract—In a wireless network, gathering information at the base station about mobile users based only on uplink channel measurements is an interesting challenge. Indeed, accessing the users locations and predicting their downlink channels would be particularly useful in order to optimize the network efficiency. In this paper, a supervised machine learning approach addressing these tasks in an unified way is proposed. It relies on a labeled database that can be acquired in a simple way by the base station while operating. The proposed regression method can be seen as a computationally efficient two layers neural network initialized with a non-parametric estimator. It is illustrated on realistic channel data, both for the positioning and channel mapping tasks, achieving better results than previously proposed approaches, at a lower cost.

Index Terms—channel mapping, user positioning, neural networks, regression.

I. INTRODUCTION

WIRELESS networks are currently undergoing dramatic changes, driven by significant innovations in the physical layer. In particular, it has been recently proposed to use massive multiple input multiple output (massive MIMO) wireless systems [1], [2], [3] with a large number of antennas in the millimeter-wave band [4], [5], where a large bandwidth can be exploited.

On the other hand, machine learning techniques made possible by the growing available computing power have recently led to tremendous successes in various domains [6], [7]. Machine learning holds promise for wireless communications (see [8], [9] for exhaustive surveys). In particular, it is possible to use uplink channel data that can be acquired easily in modern MIMO systems to predict quantities of interest. In this paper, the objective is to use the channel data to predict some target vector whose content depends on the considered task. The two tasks of interest are *user positioning* (in which the user's location is to be predicted) and *channel mapping* (in which the user's downlink channel is to be predicted).

Contributions. This paper proposes a similarity-based approach to tackle these two problems. It is initialized with a non-parametric method but is fundamentally parametric. Indeed, it corresponds to a neural network whose structure at the initialization mimics nearest neighbors regression [10], [11], [12]. Fine tuning by gradient descent within this structure allows to improve the prediction results. A theoretical justification of this simple method is proposed in the context of channel data. The introduced method is then empirically

evaluated on realistic data and proves computationally efficient and accurate on the two considered tasks.

Related work. User positioning and channel mapping using channel data have both been investigated. On the user positioning side, it has been proposed to tackle the problem with a convolutional neural network (CNN) operating directly on channel data [13], or on channel data transformed to the angular domain [14]. It has also been proposed to use a classical non-parametric approach not relying on a neural network [15], operating on the received signal strengths (RSS) rather than on the channels directly. The proposed method takes advantage of the simplicity and low complexity of non-parametric approaches to guide the structure of a neural network operating on channels, which distinguishes it from all these methods. Regarding channel mapping, which is a more recent problem, it has been tackled with help of a fully connected neural network [16]. The proposed method adopts a totally different network structure, as well as a different cost function.

II. PROBLEM FORMULATION

The methods proposed in this paper apply to a wide variety of multi-user MIMO wideband systems, operating indifferently in time division duplex (TDD) or frequency division duplex (FDD), where the antennas at the base station are indifferently colocated or not (in which case it is a distributed MIMO system). Let us consider N base station antennas and S subcarriers, and denote $\mathbf{h} \in \mathbb{C}^{NS}$ the uplink channel vector between any given user and the base station and $h_{n,s} \in \mathbb{C}$ the channel for the n th antenna on the s th subcarrier. Note that no index is introduced to denote to which specific user corresponds the channel, since the proposed method treats indifferently the channels from all users.

Based on an incoming estimated uplink channel $\mathbf{h}$, the objective in this paper is to predict a target vector $\mathbf{t}$. Depending on the considered application, the target vector $\mathbf{t}$ may be the location of the user [13], its downlink channel [16] or any other quantity of interest. Let us denote f the prediction function that maps the incoming channel to the target estimate, so that the estimated target writes

$$\hat{\mathbf{t}} = f(\mathbf{h}). \quad (1)$$

In order to calibrate the prediction function f , let us assume a dataset is available containing L labeled samples made of a channel and the associated target:

$$\{\mathbf{h}_i; \mathbf{t}_i\}_{i=1}^L. \quad (2)$$

Under this very general framework, two tasks are considered in this paper.

User positioning. The objective of this task is to predict the position of the user, based only on the knowledge of its uplink channel. In that case the target $\mathbf{t}$ contains the coordinates of the user's position (either in 2D or 3D depending on the context). Regarding the labeled data acquisition, it has to be obtained by an auxiliary sensor such as the GPS [15], either online or during an offline data collection phase.

Channel mapping. This task aims at eliminating the need for a downlink channel estimation phase. Indeed, in that case the target $\mathbf{t}$ contains the user's downlink channel, which has to be predicted from the knowledge of the uplink channel. Labeled data can be acquired during a phase where the base station sends downlink pilots, as explained in details in [16]. Channel mapping is especially interesting for FDD systems that use different frequencies for the uplink and downlink transmissions. In its original formulation [16], channel mapping takes as input the channel on only a subset of the base station antennas, so as to lighten the computational cost. This setting is also considered in the experimental part of this paper.

System operation. For the two considered tasks, the system operates in two phases. During the *training phase*, data are collected (using GPS or downlink pilots) and the prediction function is learned. During the *exploitation phase*, the learned prediction is used. For the proposed method, these two phases are not mutually exclusive in time. Indeed, after an initial training has been carried out, they can be intertwined and online learning can be used so as to adapt to changes in the environment or users distribution.

III. SIMILARITY-BASED PREDICTION

Let us now introduce the proposed prediction method, by unveiling the structure of the prediction function f . The method is based on the very simple rationale that similar channels lead to similar targets. The main novelty of this paper is to translate that rationale into a neural network structure instead of using a generic network structure, as was done previously [16], [14], [13].

A. The Nadaraya-Watson estimator

The similarity-based prediction method proposed here is inspired by the Nadaraya-Watson estimator [17], [18], which is a classical non-parametric regression method. It is based on an approximation of the joint density of inputs and targets by kernels located at the training points [19], [20]. Such an approximation leads to the conditional expectation of the target given the input taking the form

$$\hat{\mathbf{t}} = \mathbb{E}[\mathbf{t}|\mathbf{h}] = \frac{\sum_{i=1}^L K(\mathbf{h}, \mathbf{h}_i) \mathbf{t}_i}{\sum_{j=1}^L K(\mathbf{h}, \mathbf{h}_j)}, \quad (3)$$

where $K(\cdot, \cdot)$ is a kernel function measuring the similarity between two inputs. The estimate of the target is then a convex combination of targets corresponding to the training points. In order to perform well, such a non-parametric method requires a fine enough sampling of the channel space, which may require a very large number of training points L . However,

for the studied problem, it is argued in section III-C that the particular nature of channel data allows to overcome this issue.

B. Neural network structure

Let us now propose a neural network structure whose initialization corresponds to the Nadaraya-Watson estimator, so as to view this non-parametric method as a parametric one, whose parameters can be optimized. To do so, a kernel measuring the similarity between channels using the complex inner product is used. Indeed, denoting $\mathcal{I}_k(\mathbf{h})$ the set of indices corresponding to the k training channels that are most correlated to the current channel $\mathbf{h}$, the used kernel takes the form

$$K(\mathbf{h}, \mathbf{h}_i) \triangleq \begin{cases} |\mathbf{h}_i^H \mathbf{h}| & \text{if } i \in \mathcal{I}_k(\mathbf{h}), \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

It amounts to take into account only the k training channels most correlated to the current one. The hyperparameter k is chosen to be small (no more than a few dozens), yielding sparsity that reduces the computational cost of the method. With this kernel, the Nadaraya-Watson estimator can be viewed as an instance of k -nearest neighbors (k -NN) regression.

Instead of considering this estimator as static, it is possible to view it as the forward pass in a neural network, given in algorithm 1, where $\text{HT}_k(\cdot)$ refers to the hard thresholding operator that keeps unchanged the k entries of greatest modulus of its input and sets all the others to zero. The sequence of linear and nonlinear operation can also be visualized on figure 1.

Algorithm 1 Similarity-based prediction (forward pass)

Input: Input $\mathbf{h}$, dictionary matrix $\mathbf{D}$, prediction matrix $\mathbf{P}$, sparsity level k .

- 1: Correlation: $\mathbf{c} \leftarrow \mathbf{D}^H \mathbf{h}$
- 2: Hard thresholding: $\mathbf{s} \leftarrow \text{HT}_k(\mathbf{c})$
- 3: Normalization: $\mathbf{y} \leftarrow \frac{|\mathbf{s}|}{\|\mathbf{s}\|_1}$

Output: $\hat{\mathbf{t}} \leftarrow \mathbf{P} \mathbf{y}$ (estimated target)

Fig. 1: Feedforward representation of the proposed method.

The matrices $\mathbf{D}$ and $\mathbf{P}$ are the weights of the neural network that are to be optimized. Initializing them with

$$\mathbf{D} = (\mathbf{h}_1, \dots, \mathbf{h}_L), \quad \mathbf{P} = (\mathbf{t}_1, \dots, \mathbf{t}_L), \quad (5)$$

amounts to initialize the neural network as the Nadaraya-Watson estimator. These weights can then be optimized by gradient descent using back-propagation [21] to better fit the labeled dataset, according to some cost function. The precise form of the used cost functions are given in the experimental part of the paper since they depend on the considered application. This fine tuning phase helps because the distribution of the training samples may not be perfect for the ultimate goal, and the used channels can be noisy. Fine tuning is meant to correct these two potential defects.

Computational complexity. The proposed method is particularly computationally efficient. Indeed the complexity of the forward pass is dominated by the first step (multiplication by $\mathbf{D}^H$), which costs $\mathcal{O}(NSL)$ arithmetic operations. The backward pass is even more efficient since it costs only $\mathcal{O}(NSk)$ arithmetic operations, thanks to the sparsity obtained via the hard thresholding operation. Note that this complexity is much lower than the one of concurrent approaches. Moreover, if the number of training samples L is very large, it is also possible to keep only a subset of the training samples or to perform clustering to build matrices $\mathbf{D}$ and $\mathbf{P}$ in order to further reduce the complexity of the method, as is classically done for radial basis function networks (RBFN) [22]. This is one advantage of viewing the method as parametric, leading to a forward pass costing $\mathcal{O}(NS\tilde{L})$ arithmetic operations, where $\tilde{L}$ is a fixed number independent of the number of training samples. This possibility is not explored in the current paper, but an interesting avenue for future developments. The complexity of the proposed approach is compared precisely to the one of concurrent approaches in the experimental part of the paper.

C. Why should it work?

Given the simplicity of the proposed method, it is legitimate to wonder why it should work well with channel data, although in other domains such as image or audio processing, this kind of local interpolation methods are outperformed by more elaborate deep learning methods [6].

In order to answer this question, let us take a step back. Actually, data processing techniques are often based on the *manifold assumption*: Meaningful data (signals) lie near a low dimensional manifold, although their apparent dimension is much larger [23], [24] [7, Section 5.11.3] [25, Section 9.3]. The performance of local interpolation methods is heavily dependent on the dimension of the data manifold (the lower the better). In order to perform well for a reasonable number of training points, the manifold dimension should be low. Indeed, the number of required training points for a given accuracy grows exponentially with the manifold dimension [26], [27].

What is the dimension of the channel data manifold? As previous work suggests, it is reasonable to assume the existence of a *position to channel mapping* [16], i.e. a deterministic function linking the position of the user with the corresponding channel. The set of possible positions can be modeled as a two-dimensional manifold (neglecting the elevation dimension). If it is further assumed that the position to channel mapping is an homeomorphism (continuous bijection whose inverse is also continuous), then the channel vectors lie on a two-dimensional manifold (simple application of the definition of a manifold). This conclusion is very encouraging since it means that despite the high apparent dimension of channel vectors (due to the large number of antennas and subcarriers), the very low dimension of the channel manifold should allow simple local interpolation methods to perform well with a reasonable number of training points.

Note that such a simple reasoning cannot be applied to image or audio data, and the data manifold in these domains is very likely to be much higher dimensional (e.g. a few dozens

for images [28]). In summary, the channel data is inherently much less complex than image or audio data, so that local methods can work well for channel data although much more computationally complex techniques have to be used in these domains (such as deep learning).

IV. EXPERIMENTS

In this section, the proposed similarity-based approach is illustrated on the channel mapping and user positioning tasks. **Implementation details.** The method is implemented with help of the PyTorch library [29], so that gradients are computed automatically. The optimization is done with minibatch gradient descent (size of the minibatches depending on the application) using the Adam optimization algorithm [30]. Note that complex weights and inputs are handled classically by stacking the real and imaginary parts so that the neural network treats only real numbers. Moreover, for the fine tuning, the column of $\mathbf{D}$ corresponding to the current training sample is excluded in order to avoid trivial solutions during training (using for prediction a channel which is exactly the current training sample).

A. Channel mapping

The channel mapping application is investigated with help of the DeepMIMO dataset [31], which is itself based on the ray-tracing simulator Wireless InSite by Remcom [32]. In order to ease comparisons with prior art, the considered setting is exactly the same as in [16], namely multipath channels (5 paths) obtained from an indoor environment consisting in a $10\text{ m} \times 10\text{ m}$ room, with $N = 64$ antennas on the ceiling and $S = 16$ subcarriers, an uplink frequency of 2.4 GHz and a downlink frequency of 2.5 GHz (see [16] for a more detailed description). The labeled dataset consists of L uplink channels for which only 8 randomly picked antennas are considered paired with the associated downlink channel (on all 64 antennas) corresponding to users randomly located in the room, with L varying in our experiment. The cost used in order to fine tune the neural network of figure 1 (for 100 epochs) is the opposite of the downlink spectral efficiency averaged over subcarriers obtained when using the estimated downlink channel as precoder, taking the form

$$-\mathbb{E} \left[\frac{1}{K} \sum_{k=1}^K \log_2 \left(1 + \frac{|\mathbf{h}_{D,k}^H \hat{\mathbf{h}}_{D,k}|^2}{\|\hat{\mathbf{h}}_{D,k}\|_2^2} \right) \right], \quad (6)$$

where $\mathbf{h}_{D,k}$ is the downlink channel on the k th subcarrier, $\hat{\mathbf{h}}_{D,k}$ is its estimation given by the neural network, and the expectation being in practice estimated by averaging over minibatches of 1000 channels. The hard thresholding parameter is fixed to $k = 5$.

Results. Figure 2 summarizes the obtained results for $L \in \{1000, 3000, 6000, 12000, 30000, 60000\}$, showing the spectral efficiency (opposite of the cost (6)) averaged over 1000 randomly picked test channels. Results are shown for the initialization (blue dots) and after fine tuning by gradient descent for $L \in \{3000, 6000, 12000\}$ (red crosses). An upper bound (obtained using the true channel as precoder) is also shown (black dashed line). It is interesting to notice that

Fig. 2: Channel mapping results.

only a reasonable amount of training data is required to get very close to the upper bound with the proposed method. Indeed, the obtained results are within 6% of the upper bound without fine tuning and within 4% with fine tuning, with only 12000 training points. This performance is superior to the previously proposed method [16, Figure 7], which obtained results within 9% of the upper bound with 120000 training samples (ten times more). Moreover, the proposed method yields a much lighter network (forward pass costing around $1e6$ arithmetic operations) compared to the fully connected neural network of [16] (around $3e7$ arithmetic operations). In summary, the proposed method yields better results with fewer training samples, at a lower cost. This experiment shows the great potential of the similarity-based approach for the channel mapping problem. However, the number of required training samples may still seem large in regard of the quite small area considered in this experiment. This issue should be taken into account in future work. Note that this issue is less significant in the next experiment regarding the user positioning task.

B. User positioning

The user positioning task is investigated with help of a dataset originally intended for the IEEE Communication Theory Workshop (CTW) 2020 data competition [33], built with a channel sounder described in [13]. The considered base station comprises $N = 56$ antennas and communicates over $K = 924$ subcarriers at the central frequency of 1.27 GHz. The labeled dataset consists of $L = 4979$ channels $\mathbf{h}$ with the corresponding location $\mathbf{p}$ that are measured in an area of approximately 1 km^2 in the streets of Stuttgart, Germany. It is divided into 4096 channels used for training and 883 channels used as validation data.

Dimensionality reduction. Channels described in the previous paragraph are of very high dimension. Indeed, they can be seen as 56×924 complex matrices. Using directly the method of algorithm 1 on these channels would result in a very high computational complexity. In order to reduce the computational burden, the dimensionality of input channels is

Fig. 3: User positioning results.

reduced by computing the left singular vector corresponding to the largest singular value and using it as input for algorithm 1. This results in inputs of size $N = 56$ complex numbers (112 real numbers) only. The cost used to fine tune the network of figure 1 (for 50 epochs) corresponds to the average localization error, taking the form

$$\mathbb{E} [\|\mathbf{p} - \hat{\mathbf{p}}\|_2] \quad (7)$$

where $\mathbf{p}$ is the true position and $\hat{\mathbf{p}}$ its estimation given by the neural network, and the expectation being in practice estimated by averaging over minibatches of 100 channels. The hard thresholding parameter k is varied between 2 and 16.

Baselines. The proposed method is compared to two baselines taking the same inputs. The first one is an instance of extreme learning machine (ELM) as was proposed for user positioning [34], with a random layer of 50000 neurons and a rectified linear activation (ReLU) [7, Chapter 6]. The second one is a simple multilayer perceptron (MLP) [7, Chapter 6] comprising two layers made of 112 neurons with a ReLU non-linearity followed by a last linear one with three neurons (outputting location estimates). It is trained for 200 epochs with the Adam optimization algorithm [30].

Results. Figure 3 summarizes the obtained results for a varying number of neighbors (corresponding to the parameter k of the hard thresholding operator) and for the two baselines. First of all, note that the proposed method leads to better positioning results than the baselines (the MLP leading to approximately 70 m and the ELM to approximately 80 m average errors). Then, it is interesting to notice that the number of neighbors has only a modest influence on the localization error, provided it is chosen not too small (larger than 4). Second, fine tuning shows beneficial for the localization task (except for $k = 1$, where the fine tuning decreases a lot the positioning accuracy). Indeed, the initialization gives errors around 52 m while fine tuning allows to attain errors around 42 m (approximately 20% better). Keeping in mind that the area on which localization is sought is of 1 km^2 , such results are pretty accurate and encouraging. Moreover, note that most of the time the localization error is much smaller than the average (the median is

around 20 m for $k = 10$). Note that these results are the first to be reported for this outdoor localization task. Concurrent parametric methods such as convolutional neural networks were only applied to indoor localization for now [13].

V. CONCLUSION

In this paper, a generic similarity-based neural network was introduced in order to operate on channel data, for the channel mapping and user positioning tasks. It was theoretically motivated relying on the manifold hypothesis. Moreover, the proposed method was empirically validated on realistic data for the two aforementioned tasks. In the future, it would be very interesting to investigate further the sample complexity of the method, in order to optimize the training of such systems. Moreover, one can also envision using several heuristics aimed at improving the training of such similarity-based neural networks [22].

REFERENCES

- [1] F. Rusek, D. Persson, B. K. Lau, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up mimo: Opportunities and challenges with very large arrays," *IEEE Signal Processing Magazine*, vol. 30, no. 1, pp. 40–60, 2013.
- [2] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive mimo for next generation wireless systems," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 186–195, 2014.
- [3] L. Lu, G. Y. Li, A. L. Swindlehurst, A. Ashikhmin, and R. Zhang, "An overview of massive mimo: Benefits and challenges," *IEEE journal of selected topics in signal processing*, vol. 8, no. 5, pp. 742–758, 2014.
- [4] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5g cellular: It will work!" *IEEE access*, vol. 1, pp. 335–349, 2013.
- [5] A. L. Swindlehurst, E. Ayanoglu, P. Heydari, and F. Capolino, "Millimeter-wave massive mimo: the next wireless revolution?" *IEEE Communications Magazine*, vol. 52, no. 9, pp. 56–62, 2014.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [7] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [8] T. O'Shea and J. Hoydis, "An introduction to deep learning for the physical layer," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 4, pp. 563–575, 2017.
- [9] T. Wang, C.-K. Wen, H. Wang, F. Gao, T. Jiang, and S. Jin, "Deep learning for wireless physical layer: Opportunities and challenges," *China Communications*, vol. 14, no. 11, pp. 92–111, 2017.
- [10] C. J. Stone, "Consistent nonparametric regression," *The annals of statistics*, pp. 595–620, 1977.
- [11] J. K. Benedetti, "On the nonparametric estimation of regression functions," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 2, pp. 248–253, 1977. [Online]. Available: <http://www.jstor.org/stable/2984802>
- [12] J. W. Tukey, *Exploratory data analysis*. Reading, MA, 1977, vol. 2.
- [13] M. Arnold, J. Hoydis, and S. ten Brink, "Novel massive mimo channel sounding data applied to deep learning-based indoor positioning," in *SCC 2019; 12th International ITG Conference on Systems, Communications and Coding*. VDE, 2019, pp. 1–6.
- [14] J. Vieira, E. Leiting, M. Sarajlic, X. Li, and F. Tufvesson, "Deep convolutional neural networks for massive mimo fingerprint-based positioning," in *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE, 2017, pp. 1–6.
- [15] V. Savić and E. G. Larsson, "Fingerprinting-based positioning in distributed massive mimo systems," in *2015 IEEE 82nd vehicular technology conference (VTC2015-Fall)*. IEEE, 2015, pp. 1–5.
- [16] M. Alrabeiah and A. Alkhateeb, "Deep learning for tdd and fdd massive mimo: Mapping channels in space and frequency," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*. IEEE, 2019, pp. 1465–1470.
- [17] E. A. Nadaraya, "On estimating regression," *Theory of Probability & Its Applications*, vol. 9, no. 1, pp. 141–142, 1964.
- [18] G. S. Watson, "Smooth regression analysis," *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 359–372, 1964.
- [19] M. Rosenblatt, "Remarks on some nonparametric estimates of a density function," *Ann. Math. Statist.*, vol. 27, no. 3, pp. 832–837, 09 1956. [Online]. Available: <https://doi.org/10.1214/aoms/1177728190>
- [20] E. Parzen, "On estimation of a probability density function and mode," *The annals of mathematical statistics*, vol. 33, no. 3, pp. 1065–1076, 1962.
- [21] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning Internal Representations by Error Propagation*. Cambridge, MA, USA: MIT Press, 1986, p. 318–362.
- [22] F. Schwenker, H. A. Kestler, and G. Palm, "Three learning phases for radial-basis-function networks," *Neural networks*, vol. 14, no. 4-5, pp. 439–458, 2001.
- [23] G. Carlsson, "Topology and data," *Bulletin of the American Mathematical Society*, vol. 46, no. 2, pp. 255–308, 2009.
- [24] G. Peyré, "Manifold models for signals and images," *Computer Vision and Image Understanding*, vol. 113, no. 2, pp. 249–260, 2009.
- [25] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer Science & Business Media, 2010.
- [26] P. J. Bickel, B. Li *et al.*, "Local polynomial regression on unknown manifolds," in *Complex datasets and inverse problems*. Institute of Mathematical Statistics, 2007, pp. 177–186.
- [27] S. Kpotufe, "k-nn regression adapts to local intrinsic dimension," in *Advances in neural information processing systems*, 2011, pp. 729–737.
- [28] H. minn Lu, Y. Fainman, and R. Hecht-Nielsen, "Image manifolds," in *Applications of Artificial Neural Networks in Image Processing III*, N. M. Nasrabadi and A. K. Katsaggelos, Eds., vol. 3307, International Society for Optics and Photonics. SPIE, 1998, pp. 52 – 63. [Online]. Available: <https://doi.org/10.1117/12.304659>
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in neural information processing systems*, 2019, pp. 8026–8037.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] A. Alkhateeb, "DeepMIMO: A generic deep learning dataset for millimeter wave and massive MIMO applications," in *Proc. of Information Theory and Applications Workshop (ITA)*, San Diego, CA, Feb 2019, pp. 1–8.
- [32] "Remcom, wireless insite," <https://www.remcom.com/wireless-insite/>.
- [33] "Ctw 2020 data competition," <https://ctw2020.ieee-ctw.org/data-competition/>.
- [34] A. Decurninge, L. G. Ordóñez, P. Ferrand, H. Gaoning, L. Bojie, Z. Wei, and M. Guillaud, "Csi-based outdoor localization for massive mimo: Experiments with a learning approach," in *2018 15th International Symposium on Wireless Communication Systems (ISWCS)*. IEEE, 2018, pp. 1–6.