

Localization Convolutional Neural Networks Using Angle of Arrival Images

Marcus Comiter
Harvard University
Cambridge, MA, USA
marcuscomiter@g.harvard.edu

H. T. Kung
Harvard University
Cambridge, MA, USA
kung@harvard.edu

Abstract—We introduce localization convolutional neural networks (CNNs), a data-driven time series-based angle of arrival (AOA) localization scheme capable of coping with noise and errors in AOA estimates measured at receiver nodes. Our localization CNNs enhance their robustness by using a time series of AOA measurements rather than a single-time instance measurement to localize mobile nodes. We analyze real-world noise models, and use them to generate synthetic training data that increase the CNN’s tolerance to noise. This synthetic data generation method replaces the need for expensive data collection campaigns to capture noise conditions in the field. The proposed scheme is both simple to use and also lightweight, as the mobile node to be localized solely transmits a beacon signal and requires no further processing capabilities. Our scheme is novel in its use of: (1) CNNs operating on space-time AOA images composed of AOA data from multiple receiver nodes over time, and (2) synthetically-generated perturbed training examples obtained via modeling triangulation patterns from noisy AOA measurements. We demonstrate that a relatively small CNN can achieve state-of-the-art localization accuracy that meets the 5G standard requirements even under high degrees of AOA noise. We motivate the use of our proposed localization CNNs with a tracking application for mobile nodes, and argue that our solution is advantageous due to its high localization accuracy and computational efficiency.

Index Terms—wireless sensor networks, mobile node, localization, angle of arrival (AOA), tracking, convolutional neural network (CNN)

I. INTRODUCTION

Many wireless sensor network applications require localizing mobile nodes (MNs). This need is reflected in the emerging 5G standard, which calls for sub-meter indoor localization accuracy [1]. Angle of arrival (AOA) localization is a popular localization method [2] because it imposes a light resource requirement on MNs to be localized, as the MN only needs the capability to transmit a simple beacon signal. By measuring the AOA of the arriving beacon signal, multiple receiver nodes in the network infrastructure can use the principle of *triangulation* to localize the MN [3].

In practice, however, localization with triangulation can be inaccurate due to factors such as multipath effects, environmental changes, synchronization, and sensor calibration introducing AOA measurement noise [4]. Triangulation-based localization is vulnerable to such measurement errors, especially in locating nodes in collinear regions between receiver

nodes [4]. Given the complexity of analytic modeling of these environmental and system effects, it is natural to consider data-driven approaches such as those based on neural networks [5].

In this paper, we introduce localization convolutional neural networks (CNNs) that capture and mitigate the effects of measurement errors in AOA localization. Departing from previous work in using a single AOA measurement with multilayer perceptrons (MLPs) for localization [5] [4], our localization CNNs operate over *time series* AOA data, rather than AOA data for a single-time instance. Specifically, the inputs to the CNNs are space-time matrices, which we term *AOA images*, where each row consists of data related to a set of AOA measurements taken by multiple receiver nodes at a given time instance, and subsequent rows contain AOA data measured at subsequent time instances. By working on these image-like inputs, localization CNNs are robust against AOA measurement errors occurring at a single-time instance. In addition to providing robustness against noise, the localization CNNs can be made more efficient than MLPs, and are amenable to further computation reduction via efficient CNN implementations in the literature leveraging image-like inputs [6].

To reduce the effort of acquiring AOA measurements in the field, we augment field-collected data with synthetically-generated data to train localization CNNs. We analyze the effect of noisy AOA measurements on triangulation localization and generate synthetic AOA training data accordingly using Gaussian perturbations that reflect real-world noise models [4]. By using synthetic training data, we can avoid expensive data collection campaigns that could otherwise be needed in order to obtain a diverse training data set that encompasses essential error profiles. Specifically, after the receiver nodes have acquired a set of accurate AOA measurements, we generate additional synthetic AOA samples for the location by perturbing these AOA measurements with Gaussian noise. We constrain these perturbations such that they will lead to small equilateral intersection triangles that contain the location. This data generation method provides training examples for commonly seen noise. We show that by utilizing these synthetic training data, the localization CNN performs with high accuracy even under significant noise (up to 15 degrees of measurement error per AOA measurement).

Our contributions in this paper are the following:

- Localization CNNs using space-time AOA images. As

far as we know, we are the first to use such images incorporating AOA measurements from multiple receiver nodes in space and over time in training CNN localization models robust to measurement noise.

- Model of how triangulation AOA localization is affected by noise in AOA measurements, and use of intersection triangles in constraining synthetically generated data to be used for training CNNs that are robust against errors in AOA measurements (see Section VI). This synthetic data may form entries in the space-time AOA images used in training localization CNNs.
- State-of-the-art localization accuracy achieved with computationally efficient CNNs (see Section VII-D). We achieve the sub-meter accuracy called for by 5G for indoor deployments [1].
- A motivating application scenario for tracking mobile nodes, for which the CNN-based localization approach proposed in this paper is an enabling mechanism (see Section VII-E).

II. RELATED WORK

Localization methods using AOA measurements were studied in [2] but require many antennas on each receiver to mitigate the impact of AOA measurement errors on localization accuracy, while our method uses carefully curated synthetic data to address AOA measurement errors that allows for fewer antennas to be used. Localization via other methods was proposed in [7] using frequency-modulated carrier wave (FMCW) rather than AOA, but requires the MN to be moving, while our method places no requirements on the MN's behavior.

In the context of 5G and 802.11ad that require mmWave beamforming, approaches have been proposed to localize a MN in-band. The 802.11ad Amendment II proposes Sector Level Sweep (SLS) and Beam Refinement Process (BRP) to align receiver and transmitter antennas. Improving upon this, [8] introduces Blind Beam Steering to reducing the amount of scanning required in the SLS/BRP process. Our work instead uses out-of-band beacons over lower frequency spectrum to localize MN, after which this localization information can be leveraged to perform mmWave beamforming.

Data driven localization schemes have been studied in [9]–[11], but they do not meet the required sub-meter accuracy. The work of [4], [5] show that multilayer perceptrons can be used to achieve sub-meter accuracy. Further, these works show that spatially-correlated biases and multipath effects that are present in the space are captured by the model as long as they are captured in the training data. This is a benefit of data driven methods over conventional methods: rather than trying to conventionally model the multipath effects or address them analytically, which can be to complex to carry out in many deployment scenarios, the model can instead learn these effects if they are present in the training data. Our work is similarly data driven in order to realize these benefits.

Convolutional neural networks specifically have also recently been leveraged in similar domains. The work of [12] has proposed the use of convolutional recurrent neural networks

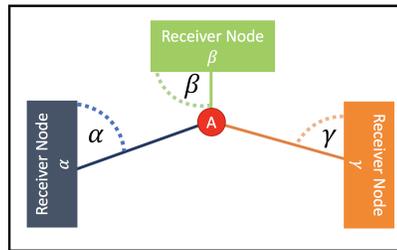


Fig. 1: The angle of arrival (AOA) measured by each of the three receiver nodes intersect at one unique point in the noiseless illustration shown here.

for predicting direction of arrival of sound waves. In that work, the input to the CNN is a discrete Fourier transform of the spectrogram of the audio signal captured over a two second interval. The temporal nature of the audio signal provides a natural medium over which the CNN may operate. In our work, we similarly leverage the time series data of AOA measurements taken over subsequent time steps to form an AOA image over which the localization CNN may operate.

III. LOCALIZATION PROBLEM STATEMENT AND MOTIVATING TRACKING APPLICATION

A. Localization Problem Statement

The localization problem we solve in this paper is shown in Figure 1. We are given a MN that continuously transmits a beacon signal within an environment and a small set of receiver nodes (three or more) that can estimate the AOA of the transmitted beacon. Given a set of AOA measurements over time, we wish to locate the MN via a learned CNN model that performs triangulation. In this paper we perform localization in a two-dimensional space, but the proposed approach can be extended to three-dimensional localization.

Complicating matters is the fact that AOA data is often noisy. As a result, there will be errors in estimated locations. To address this, we propose to use CNNs (inspired by their success in image recognition) for the localization task, as they provide a natural method for incorporating multiple measurements taken over time to combat the noise in single measurements. Beyond achieving high localization accuracy, these CNNs can be made highly efficient via model reduction [6], weight pruning [13], and quantization [14].

B. Motivating Tracking Application

As an illustrative application of this work, we consider tracking a moving MN (e.g., a robot) to ensure that it follows a prescribed path (e.g., for security concerns or path correction purposes). This is depicted in Figure 2, where the MN moving from location A to B is expected to be at location i at time $t_i = i \cdot \Delta$, where Δ is some time interval, for $i = 1, 2, \dots$. We denote the *permissible region*, or the area that the MN is allowed to be in at time $t_i = i$, with a small grey region around location i . Enabled by localization CNNs, at a given time, the MN may momentarily stop moving and allow the system to check if it has substantially deviated from the permissible

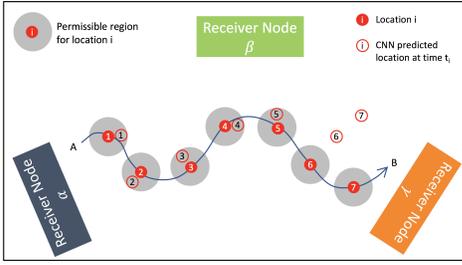


Fig. 2: A tracking application scenario. The localization CNN detects when the mobile node is off the path (e.g., $t_i = 6, 7$).

regions for some number of consecutive time instances, in which case the system can trigger the necessary alarms and take the proper corrective actions.

This application is complicated by noise in AOA measurements. If this noise leads to inaccurate location estimations, the system alarm will experience potentially costly false positives and false negatives when the MN is near the boundary of a permissible region. As the level of ambiguity triggering false alarms relates to the magnitude of the localization errors, we strive to train localization CNNs that minimize localization errors in the presence of AOA measurement noise.

The localization CNN introduced in this paper addresses this challenge. Specifically, localization CNNs can be used directly to localize a MN to see if it is within the permissible region at a given time. Alternatively, for this particular application, we also present a simplified variant of the localization CNN that is exactly the same as a localization CNN, but only outputs if the MN is in the correct permissible region at time i rather than an exact localization. See details in Section VII-E.

IV. PROPOSED LOCALIZATION CNNs WITH AOA IMAGES

A. Localization CNNs Using Time Series Input

In this section we describe localization CNNs, a novel application of CNNs that accept “angle of arrival (AOA) images” as input and output the localization of the MN. In contrast to previous methods that localize a MN from a single set of AOA measurements, the localization CNN of this paper incorporates a time series of measurements (for example, multiple measurements taken at a given location over time during inference), providing increased robustness against noise. An overview of this time series-based localization pipeline is shown in Figure 3. Additionally, as our experiments show, localization CNNs can be modest in size (e.g., three convolutional layers and one linear layer) and therefore efficient, yet still achieve high localization accuracy.

Figure 8 shows an example structure of a modest localization CNN used for experiments in this paper. This localization CNN consists of a number of convolutional layers, each followed by a maximum pooling operation, followed by a number of linear layers. These and other hyperparameters (such as convolutional filter size) can be set based on application needs (e.g., desired localization accuracy and computational

efficiency). Localization CNNs are trained in a traditional fashion via backpropagation.

B. AOA Image

An AOA image is formed by collecting multiple AOA measurement sets at discrete time intervals. Each AOA image is associated with a particular location of the MN (the label that the localization CNN tries to predict). At each time interval, a set of AOA measurements, consisting of one AOA measurement from each receiver node, is obtained to form a row of the image. Because of noise in the system, each of the measurements taken will be different even though the MN has not moved. As more measurement sets are acquired, they are stacked on top of one another, forming the AOA image.

Figure 4 on the left depicts an AOA image, produced from the AOA measurements collected at a sequence of discrete time intervals. Formally, the AOA image is a matrix where each row is the set of AOA samples measured by the receiver nodes at a given time, where sequential rows correspond to additional samples taken at subsequent times. As such, if S samples over time are taken, then we obtain a $\{S \times \# \text{ receiver nodes}\}$ -sized matrix as the AOA image. The localization CNN will take an AOA image as input to predict the location of the MN.

C. Reshaping AOA Image

Given that the raw AOA image is tall and skinny as there are more samples (rows) than receiver nodes (columns), the area over which convolution filters can operate is limited. To make effective use of the information captured from the time series data via repeated convolution operations, prior to being fed into the CNN, the AOA image is *reshaped* to be approximately square, as shown in Figure 4. There is no requirement for the AOA image to be exactly square, as the reshaping is used just to create a larger area over which convolutional filters may operate. After the AOA image is reshaped, it is fed into the localization CNN, a CNN trained to predict the location of a MN from its reshaped AOA image. At inference time, input AOA images are reshaped in an identical manner.

V. LOCALIZATION NOISE MODELING WITH INTERSECTION TRIANGLES

In this section we discuss how noise in AOA measurements affects the localization task. Following this, in Section VI we discuss a synthetic data generation method that leverages these insights to improve the accuracy of localization CNNs operating on noisy data.

We first examine the intersection patterns of AOA measurements with and without noise, shown in Figure 5. At the top of the figure, we show that without noise, the AOAs intersect at one single point (the MN’s true location). However, in the presence of noise, the intersection exhibits a more complex pattern. In the bottom of the figure, we show the intersection pattern of “noisy” AOA measurements (i.e., perturbations of the true AOA measurements). Under noise, the AOAs intersect at three distinct points, forming an *intersection triangle* around

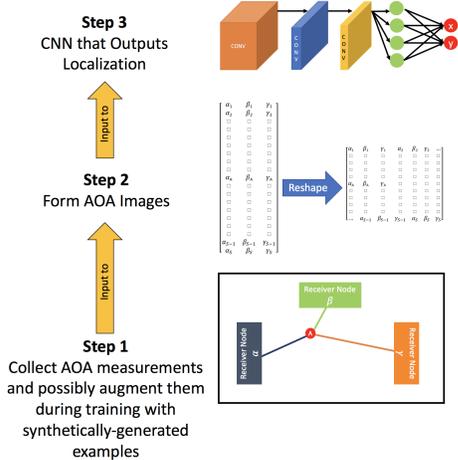


Fig. 3: The pipeline for producing a localization of a mobile node, denoted as “A”, from a time series of AOA measurements. A number of AOA measurements are taken sequentially or generated synthetically, producing an AOA image, which is then reshaped and fed into the localization CNN, which produces the final localization of the target point.

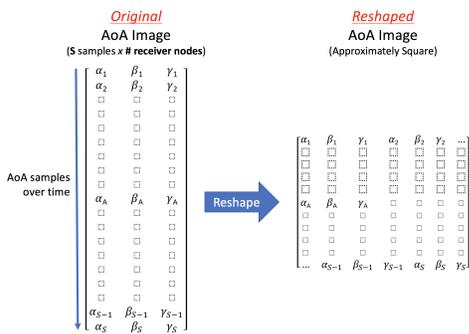


Fig. 4: Left: the angle of arrival (AOA) image that is used as input to the localization CNN. Each row consists of AOA measurements at a given time for all receiver nodes (in this illustration, there are three receiver nodes). The time series of measurements are stacked sequentially to form the AOA image. Right: the AOA image is reshaped to be as square (equal number of rows and columns) as possible to create a larger area over which the convolutional filters can operate.

the MN’s true location. We also show that the triangle’s *midpoint* is the average localization point under noise. Depending on the shape of triangle (determined by the amount and nature of the noise in the AOA measurements), this midpoint may or may not be close to the MN’s true location.

We use the size and shape of the intersection triangle to assess the effect of the noise on the localization task. Figure 6 shows different triangle patterns. We use two metrics to characterize intersection triangles: 1) triangle perimeter, and 2) the ratio of the length of the longest side to the shortest side. We use these metrics to categorize the size (small or large) and shape (equilateral or skewed) of the triangle. Larger amounts of noise correspond to larger triangles, and smaller amounts

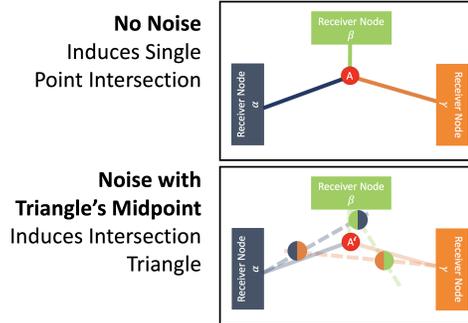


Fig. 5: The different intersection patterns created by AOA measurements with and without noise. Without noise (top), the three AOA measurements intersect at a single point (red anchor circle A). With noise (bottom), the three AOA measurements have three intersection points (the dual colored circles), which together form the *intersection triangle*. The *midpoint* of this triangle is the average localization point from the noisy measurements (red anchor circle A’).

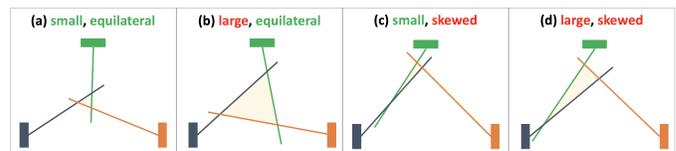


Fig. 6: Depending on AOA measurement noise, differently sized (small/large) and shaped (equilateral/skewed) intersection triangles are produced. Each receiver’s incident line indicates a perturbed AOA measurement. When generating synthetic training data or gathering test data, we retain only perturbed AOA measurements whose intersection triangles are both small and equilateral in (a) and discard synthetic AOA data with intersection triangles as in (b), (c), and (d).

of noise correspond to smaller triangles. The shape of the triangle determines if its midpoint is close to the true location of MN (yes in the small equilateral case, and no otherwise). As we discuss in Section VI, we leverage this insight and use small (e.g., perimeter is less than 1.5 meters for a given $4\text{m} \times 4\text{m}$ area) equilateral triangles to constrain the generation of synthetic data as well as qualify test data in order to improve the performance of the localization CNN.

VI. DATA GENERATION PROCESS VIA CONSTRAINED PERTURBATIONS ON AOA MEASUREMENTS

We now discuss how synthetic training data can be generated to improve localization CNN performance. After a set of noiseless training AOA measurements are acquired for a location (i.e., they lead to an intersection triangle so small that it is effectively an intersection point), we perturb the measurements to generate synthetic training examples for the location as follows. We add Gaussian noise to the measurements, inducing a number of triangle intersection patterns to be formed as illustrated in Figure 6. For each triangle intersection pattern, two statistics are calculated: 1) the perimeter (proxy

for size), and 2) the ratio of largest to smallest side length (proxy for skewness). Perturbations leading to triangles that are too large or too skewed are removed from consideration, and the rest form an AOA image for the location.

We now explain why constraints on the type of triangles (size and skewness) are used. If the triangles are highly skewed or extremely large, this is a result of a large amount of irregular perturbations in the AOA measurements. As the corresponding noisy AOA measurements are by definition a rare event, and because we desire that the training data to be representative of the “typical” type of noisy AOA measurements seen, we do not wish to introduce these rare cases into the training set. At test time, the same constraint is used to remove noisy input samples with triangular intersection patterns not meeting these constraints (in practice done by continuing to sample AOAs until the test AOA image is formed).

VII. PERFORMANCE RESULTS

A. Setup of Experiments

We use a setup with three receiver nodes locating a single MN in a $4\text{m} \times 4\text{m}$ area for all experiments. AOA measurements are created using the real-world validated simulator introduced in [4]. This simulator creates the $4\text{m} \times 4\text{m}$ virtual space in which the receiver nodes and MN can be placed at arbitrary positions. For each MN location, the simulator calculates the precise AOA measurement between the MN and each of the receiver nodes. This AOA reflects the relative geometry of the MN and each of the receiver nodes, and is therefore the ground-truth AOA subject to no noise, corresponding to real-world field measurements under perfect (no noise) conditions.

For training, we use 200 (AOA image, location) training points corresponding to 200 random locations within the area. Each AOA image contains 225 sets of AOA measurements generated by the perturbation method of Section VI. This original AOA image is of size 225×3 and is reshaped to be an approximately square AOA image of size 27×25 .

For testing, we use 625 previously unseen and randomly chosen MN locations. Following this, because in real situations there is noise in the measurements, Gaussian noise is added independently to these AOA measurements. As noted earlier, the Gaussian noise added is a model that has been validated as the same noise seen in real world experiments in [4]. That is, this Gaussian noise model accurately captures the noise seen in real-world field measurements in various environments (including secluded outdoor environments and populated indoor environments). The field experiments performed to demonstrate the real-world validity of the noise model used in these experiments are depicted in Figure 7.

We use the CNN shown in Figure 8 for the localization experiments. The CNN has a modest size, consisting of three convolutional layers with 200, 30, and 16 filters of size 2×2 , respectively, each followed by a maximum pooling. The final convolutional layer is followed by a linear layer with 50 neurons, which is followed by the output layer that produces the two-dimensional coordinates of the predicted localization.

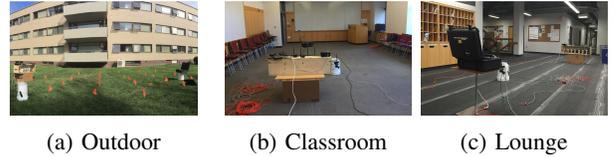


Fig. 7: Pictures of outdoor and indoor real world experiments.

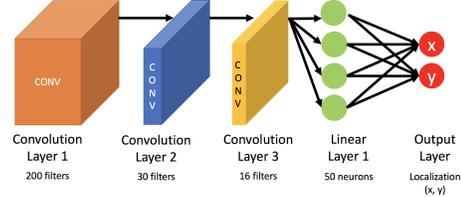


Fig. 8: The CNN used for experiments in this paper.

All results on localization accuracy are reported in terms of Median Squared Error (Median SE).

B. Increasing Number of Samples over Time for Improved Localization Accuracy

We first show that taking additional input samples over time reduces localization error at test time, justifying the use of time series-based AOA images for localization CNNs. To allow for a clean conclusion, in this experiment the models are trained on data without noise, and the training and test sets do not use the synthetic data generation process of Section VI.

Figure 9 shows Median SE as the number of samples in each AOA image is increased (increasing the number of rows in the AOA image before reshaping) for test data with and without noise. While without noise increasing the number of samples taken does not affect error, with Gaussian noise $\mathcal{N}(0, 0.05)$, increasing the number of samples reduces error. This is explained in that for the noiseless case, additional samples do not add additional information, but under noise, additional samples reveal further statistics about the noisy signal. This provides justification for the use of time series data for the localization task, as additional measurements decrease localization error in the presence of noise.

C. Distribution of Triangle Sizes and Median-based Selection of Size Threshold

We next examine the distribution of the sizes of the intersection triangles induced by Gaussian noise on AOA measurements. These results will allow us to choose appropriate parameters in constraining perturbations on AOA measurements when generating synthetic training data.

We create a set of triangular intersection patterns induced by noisy measurements for two levels of noise ($\sigma = 0.01, 0.05$), and examine the relative frequency of the sizes of triangles formed, shown in Figure 10. Under both, the majority of the triangles have small size, but there are some outliers, as reflected in the exponential decay. Together, these observations enable us to pick a proper size threshold. A *mean-based* threshold should not be used, as it is not robust to the large

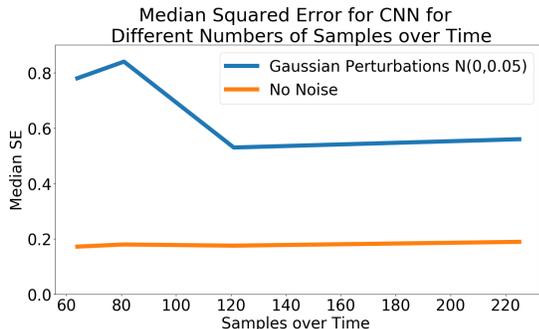


Fig. 9: Median SE for different amounts of input samples used in forming AOA images at test time with and without Gaussian noise. With noise, additional samples reduce error.

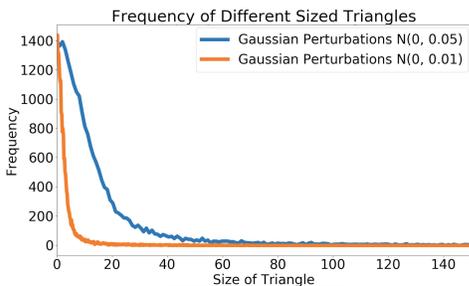


Fig. 10: The frequency of different sizes of triangles.

outliers. Therefore, we use a *median-based* threshold that provides robustness to outliers. Further, because there are many small triangles, there will be a sufficient number of triangles to generate synthetic noisy data whose triangular intersection patterns meet the constraints, even with a small perimeter threshold. For all experiments, we use the median as a threshold for perimeter and skewness. While empirically this median-based threshold performs the best, when it is not practical to calculate this, such as when field conditions are rapidly changing or there is not sufficient time to collect data for estimating the amount of noise present, we find that setting a single universal small threshold on the perimeter (e.g., 1.5) for different noise levels does not greatly impact accuracy.

D. Performance of Localization CNN Trained on Synthetically Generated Data

We now examine the performance of our localization CNN trained with data synthetically-generated by the method described in Section VI. The localization error of models trained with synthetically-generated data is shown in Figure 11. We note that the localization CNN easily meets the sub-meter localization accuracy goal of the 5G standard [1], as even under the most amount of noise (corresponding to 15 degrees of standard deviation in the additive white Gaussian noise), the Median SE is only 0.29, corresponding to 0.54m error.

These results can be explained by understanding the type of perturbations contributed by the synthetic data generation method. If we did not constrain the type of perturbations used in generating the training AOA images, we would be including

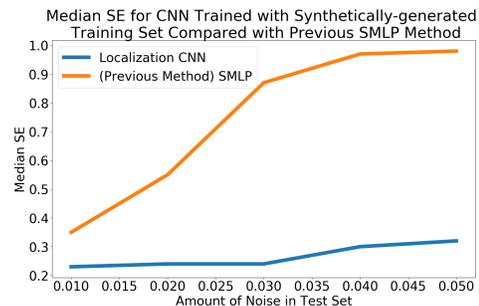


Fig. 11: Median SE of our localization CNN (blue) trained on synthetically generated data for different noise levels meets the sub-meter localization accuracy goal and improves significantly over the Structured MLP (SMLP) method (orange).

unconstrained noise that is very rare, which would have the risk of 1) not being useful at test time if similar outliers are not seen, and 2) potentially causing the model to overfit to these rare examples. In contrast, synthetic training examples generated with constrained perturbations represent data that is more likely to be seen at test time, and does not contain the outliers that may cause the model to potentially overfit.

In Figure 11 we also compare against the localization results of [4] using a Structured MLP (SMLP) that accepts a single AOA measurement (unlike the localization CNN capable of accepting multiple AOA measurements via an AOA image). The localization CNN significantly outperforms the previous SMLP model at all noise levels, especially at high noise levels (SMLP has 0.98 Median SE and localization CNN has 0.32 Median SE). Further, the computational cost in terms of multiply-accumulate operations (MACs) is significantly less for the localization CNN than the SMLP. The localization CNN used in the experiments uses 440K MACs, while the Structured MLP uses 725K MACs. Therefore, the localization CNN significantly improves over the performance of the Structured MLP while reducing computational cost.

Finally, we examine the effect of using fewer training points in training the localization CNN. Figure 12 shows localization CNN accuracy for different amounts of training points (i.e., the number of AOA images used in training the localization CNN) both with and without noise. Even with a 50% or 75% reduction in training points from 200 AOA images to 100 and 50 AOA images, respectively, Median SE increases only from 0.24 to 0.28 and 0.33, respectively, in the noiseless case.

E. Application Tracking Results

We present results for the tracking scenario introduced in Section III-B. We use the same $4\text{m} \times 4\text{m}$ area used in the other experiments and the scenario of Figure 2, but with four $1\text{m} \times 1\text{m}$ permissible regions rather than seven (shaped as squares rather than circles for ease of implementation). To show that fewer training AOA images can be used, for each region i , a classifier CNN_i is trained to output “yes” if the MN is within permissible region i at time i , and “no” otherwise (this is exactly the same as the localization CNN described in

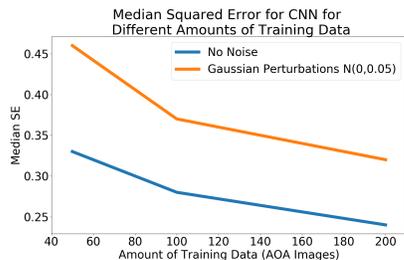


Fig. 12: The localization CNN can still perform with minimal accuracy loss when relatively less training data is used.

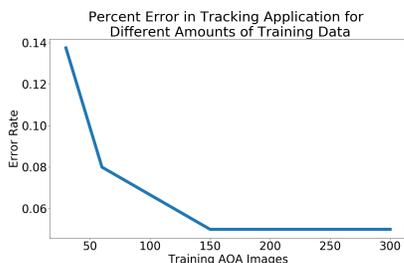


Fig. 13: Tracking error as function of training data amount.

this paper except it outputs this binary classification rather than location). At test time, test data is generated from locations within the four permissible regions and an additional area completely outside these four permissible regions (as would occur if the MN completely leaves the prescribed path). Each AOA image consists of 121 measurement sets. We assume that the MN is momentarily stationary as it acquires the AOA measurements to form the AOA image.

Figure 13 shows the total percent of wrong classifications for different amounts of training data, where the test data is perturbed with Gaussian noise drawn from $\mathcal{N}(0, 0.2)$. With 150 training AOA images, the error rate is under 5%, and with substantially fewer training samples (30 AOA images), the error rate remains low (14%).

Figure 14 shows the error rate for different amounts of Gaussian noise added to the test data. With moderate amounts of noise, the error rate is extremely low (less than 1%). Even under exceedingly large amounts of noise, the error rate is still low (13%). When the amount of training data is reduced 5-fold (300 to 60), the error rate increases only about 2%.

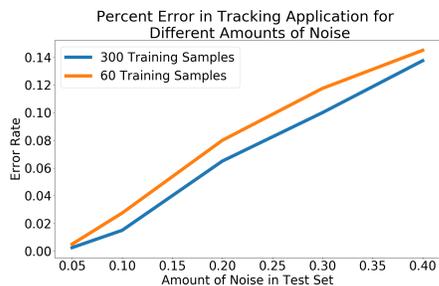


Fig. 14: Tracking error as function of Gaussian noise.

VIII. CONCLUSION

We have introduced localization CNNs, a novel application of CNNs that take an AOA image representing a time series of AOA measurements as input, and output a localization for a mobile node. Localization CNNs achieve high accuracy, easily meeting the sub-meter accuracy goal of the emerging 5G standard. The formulation of the CNN to take image-like input and the notion of an AOA image mapping tuples of AOA measurements into an image-like form are novel and allow for a natural learning model to incorporate time series data to improve localization performance. We have validated the presented methods extensively in terms of absolute localization accuracy and with a novel application tracking MN compliance to a prescribed path. We have further demonstrated that the methods are highly robust to the Gaussian noise model that we have shown in previous work to exist in real-world environments. Finally, we have presented synthetic data generation methods leveraging insights from analyzing noise patterns to constrain generated data to the most useful forms based on triangular intersection patterns of noisy AOA measurements.

REFERENCES

- [1] N. Alliance, "5g white paper," *Next Generation Mobile Networks, White paper*, 2015.
- [2] J. Xiong and K. Jamieson, "Arraytrack: a fine-grained indoor location system." Usenix, 2013.
- [3] C. Savarese, J. M. Rabaey, and J. Beutel, "Location in distributed ad-hoc wireless sensor networks," in *Acoustics, Speech, and Signal Processing, 2001. Proceedings.(ICASSP'01). 2001 IEEE International Conference on*, vol. 4. IEEE, 2001.
- [4] M. Z. Comiter, M. B. Crouse, and H. Kung, "A data-driven approach to localization for high frequency wireless mobile networks," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, 2017.
- [5] M. Z. Comiter, M. B. Crouse, and H. Kung, "A structured deep neural network for data-driven localization in high frequency wireless networks," *International Journal of Computer Networks & Communications (IJNCN)*, vol. 9, 2017.
- [6] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [7] F. Adib, Z. Kabelac, and D. Katabi, "Multi-person localization via rf body reflections," in *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, 2015.
- [8] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015.
- [9] X. Wang, L. Gao, S. Mao, and S. Pandey, "Csi-based fingerprinting for indoor localization: A deep learning approach," *IEEE Transactions on Vehicular Technology*, vol. 66, 2017.
- [10] G. Félix, M. Siller, and E. N. Álvarez, "A fingerprinting indoor localization algorithm based deep learning," in *Ubiquitous and Future Networks (ICUFN), 2016 Eighth International Conference on*. IEEE, 2016.
- [11] J. Luo and H. Gao, "Deep belief networks for fingerprinting indoor localization using ultrawideband technology," *International Journal of Distributed Sensor Networks*, 2016.
- [12] S. Adavanne, A. Politis, and T. Virtanen, "Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network," *arXiv preprint arXiv:1710.10059*, 2017.
- [13] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," *arXiv preprint arXiv:1510.00149*, 2015.
- [14] D. Lin, S. Talathi, and S. Annapureddy, "Fixed point quantization of deep convolutional networks," in *International Conference on Machine Learning*, 2016.