“We Know Where You Are”: Indoor WiFi Localization Using Neural Networks

Tong Mu, Tori Fujinami, Saleil Bhat

Abstract:

In this project, a neural network was trained to predict the location of a WiFi transmitter given the measured amplitude and phase (CSI data) of the transmitted signal at the WiFi receiver. Each CSI data point was mapped to a ground-truth angle of arrival (AoA) in degrees, and distance in meters. Training neural networks with this data, we were able to classify AoA from phase difference with 86% accuracy and predict distance from phase difference within around 0.5 meters.

I. Introduction:

With the ever-increasing demand for indoor localization services (e.g. location-customized app behavior, intelligent car parking, guidance for the visually impaired), accurate and fast localization techniques using WiFi have become increasingly relevant.

To this end, we trained a neural network to predict the location of a Wi-Fi transmitter unit with respect to a Wi-Fi receiver unit. Using a transmitter with one antenna and a receiver with three evenly spaced antennas, channel state information (CSI) data, (the amplitude and phase of the signal at each receive antenna), were collected. These data were fed into a neural network for training in order to predict the line-of-sight distance and angle between the transmitter and the receiver. We hypothesized that there would be a relationship between the amplitude of the received signal and distance, and a relationship between the difference in phase measured at two receive antennas and the angle. However, we found that phase difference data was better suited to predict both distance and angle of arrival. Using phase difference data as our input, we were able to predict distance with an accuracy of around 0.5m and classify angle with around 14% error.

II. Related Work:

WiFi localization recently has become an important area of research. There are a few papers detailing previous work in this area. SpotFi[3] achieved a high degree of localization accuracy (within decimeters) using only CSI data, but it used a non-machine learning algorithm. Particularly, SpotFi uses the computationally complex MUSIC algorithm which uses intensive linear algebra to decompose the received signal into a matrix of multipath components and attempts to determine the LOS component. Unfortunately, this computationally expensive algorithm makes the method unpractical for real-time systems.

Two other papers, "Decimeter-Level Localization" [4] and "WiFi Localization and Navigation for Autonomous Indoor Mobile Robots" [5] achieved decimeter-level localization accuracy in short computation times, but had to add additional hardware to obtain their results. Chronus, the algorithm and system detailed in "Decimeter-Level Localization," can compute location accurately and quickly from time of flight measurements, but the system must be changed to mimic wideband radio signals. "WiFi Localization and Navigation for Autonomous Indoor Mobile Robots" [5], details a specific application of indoor WiFi localization on a real-time system using machine learning; however, as robots generally have additional sensors, the localization algorithm used in this paper combines WiFi localization with the use of odometry. While we would like to achieve both the short computation time and accuracy seen in these paper, we only want to use the CSI data obtainable from a simple standard WiFi setup.

Finally, a last set of papers both used machine learning approaches and only CSI data. DeepFi[1], uses deep learning and fingerprinting with amplitude data only and ran their experiment in a variety of different environments such as a living room or a computer lab. Similarly, PhaseFi[2] also uses deep learning and fingerprinting, but with both CSI amplitude and phase data. While DeepFi[1] and PhaseFi[2] were significantly faster than SpotFi[1], the accuracies they achieved (around 2 meters for DeepFi[1] and 1-2 meters for PhaseFi[2]) were not as precise as those achieved in the any of the above mentioned papers.

The goal of this project was to obtain good accuracy, closer to what was achieved in SpotFi[3], at a faster rate, like in PhaseFi[2] or DeepFi[1], without the modifications to a WiFi system as were necessary in "Decimeter-Level Localization" [4] and "WiFI Localization and Navigation for Autonmous Indoor Mobile Robots" [5].

III. Experimental Setup, Features, and Data:

Our equipment for data collection consisted of a Wi-Fi transmitter with a single antenna and a Wi-Fi receiver with three antennas, spaced 2.6 cm apart (Fig. 1).

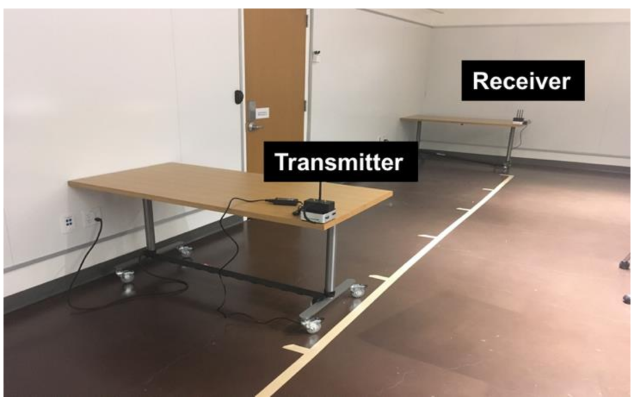

Figure 1: Experimental Setup

The Wi-Fi transmitter transmitted at 30 different subcarrier frequencies each spaced 312.5KHz apart with the center frequency at 5GHz. For distances between 1 meter to 10 meters in increments of 1 meter, and for each distance, at angles between -40 and 40 degrees (Fig. 2) in increments of 10 degrees, CSI data was collected. Note that the angle between the transmitter and receiver is the same as the angle of arrival of a line-of-sight signal propagating from the transmitter to the receiver. Therefore, predicting the angle between the devices is equivalent to predicting the line-of-sight signal's angle of arrival (AoA). Thus, the two labels we attempted to predict were distance and AoA.

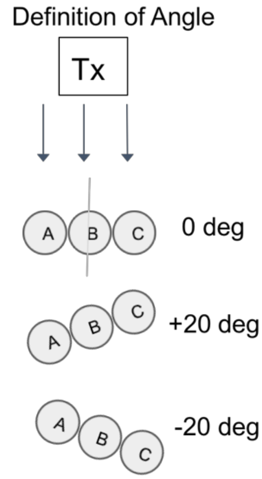

Figure 2: Definition of Angle of Arrival

Once the data was collected, we performed pre-processing. First, we linearized the phase data to remove its periodic properties. Second, for every combination of two antennas, we calculated the difference between the linearized phase measured at each antenna for each subcarrier. That is, we obtained the linearized phase difference between antennas A and B, antennas A and C, and antennas B and C on each of the 30 subcarriers.

With this setup, one training example in our data consisted of the received amplitude at each antenna on each of the 30 subcarriers, the difference in the received phase for every set of two antennas on each of the 30 subcarriers, and the ground truth labels of the parameters we wanted to predict: AoA and distance. Initially, this meant we had a total of 180 features (30 subcarriers x 3 antenna x 1 distance + 30 subcarriers x 3 antenna pairs x 1 phase difference). However, we discovered that the amplitude features were not useful in predicting distance. Therefore, for our final results, we had 90 features and two ground-truth labels for each training sample. We collected around 1900 samples of data for each distance-angle pair resulting in 171891 samples total.

IV. Methods:

For a line-of-sight angle of arrival θ, a signal travels an additional distance of $dsin(\theta)$

to the second antenna in the array compared to the first antenna (Fig 3). This results in a difference between the phase measured at each antenna of $2\pi d\ sin(\theta) f/c$.

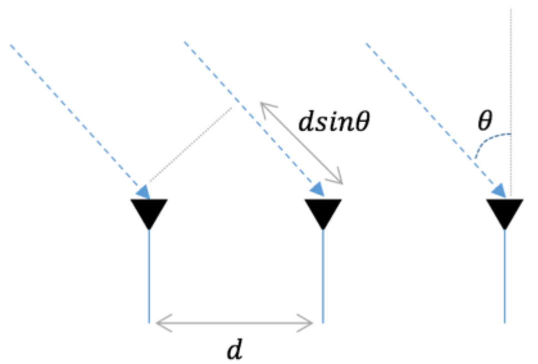

Figure 3: Phase difference based on angle and distance

However, for each signal, there is more than one propagation path (it can bounce off walls, objects, etc. before reaching the receiver). This is called multipath. As such, the phase difference is actually given by:

$$\angle\left(\sum_{i=1}^{M} a_i e^{j\varphi_i}\right) - \angle\left(\sum_{j=1}^{N} a_j e^{j\varphi_j}\right)$$

where the phase difference between the two line-of-sight components is given by $2\pi d\ sin(\theta) f/c$. The amplitudes, a_i, a_j, and phases φ_i, φ_j are dependent on both the distance and AoA (Fig. 4). Amplitude decreases due to path loss as distance (r) increases, and the angle of the multipath components (φ) with respect to the line of sight component changes with distance. Additionally, most antennas have high directivity, meaning they have high gain (g) in one direction which falls off quickly as the signal arrives off angle. Therefore, as the AoA is changed from zero degrees, the amplitude decreases as well.

We hypothesized that a neural network could learn a pattern for the multipath terms, and thus, be able to draw a relationship between line-of-sight angle of arrival and phase difference. Too see this, we cast the problem as a classification problem with 9 angle of arrival classes corresponding to the angles at which we made our measurements and used MATLAB's neural network tool with 30 layers and the default configuration: a sigmoidal activation function for each layer and scaled conjugate gradient backpropagation as the update weight/bias update procedure.

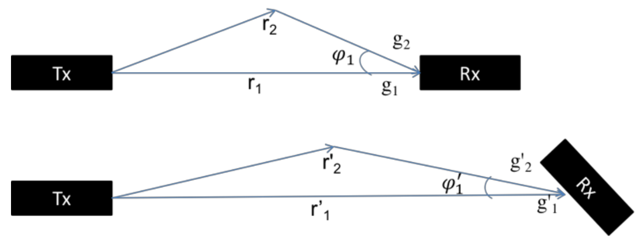

Figure 4: Parameter change with distance and AoA

Similarly, we tried to establish a relationship between phase difference and distance. The different multipath components can sometimes destructively interfere. In this case, the receiver will not observe a signal component with that phase. These patterns can be qualitatively seen in Fig. 5 below. Since the phase of each multipath component is a function of distance and room geometry, we hypothesized that each distance would have a characteristic set of "missing" phases which the neural network could learn.

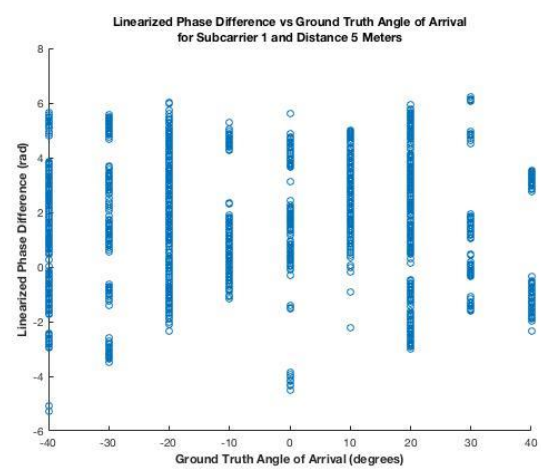

Figure 5: Linearized phase difference vs AoA at d=5m

For this case, we ran a regression with the same hidden layer configuration as used for the AoA classification case.

V. Results and Discussion:

There was no observed dependence between the distance and amplitude of the received signal. As shown below (Fig. 6), it can be qualitatively seen that the variance in amplitudes for every distance in our experiment was very high. As a result, it was impossible to recover any distance information from amplitude alone.

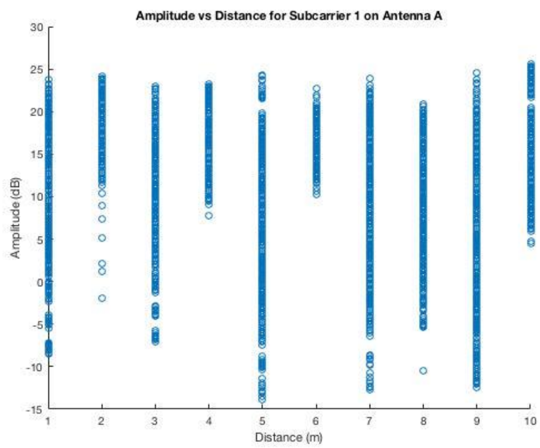

Figure 6: Amplitude vs Distance

However, the neural networks were successfully trained to learn a dependence between both phase difference and distance, and phase difference and AoA.

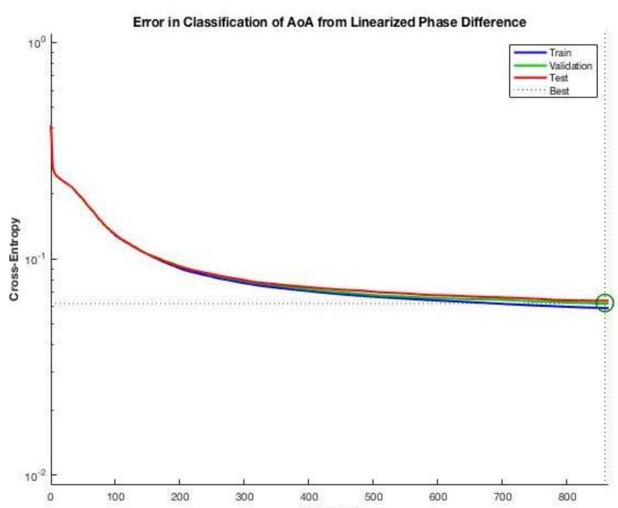

Figure 7: The Learning Curve of the Neural Network for Classifying AoA

Classification of AoA	Samples	Cross-Entropy
Training	125923	13.5%
Validation	26984	14.1%
Test	26984	14.3%

Figure 8: Results for AoA Classification

For AoA, we were able to classify the angles between -40˚, -30˚,…, 20˚, 30˚, 40˚ with 14.3% accuracy on the test data. The results are detailed in the plot and table above (Fig. 7, Fig. 8).

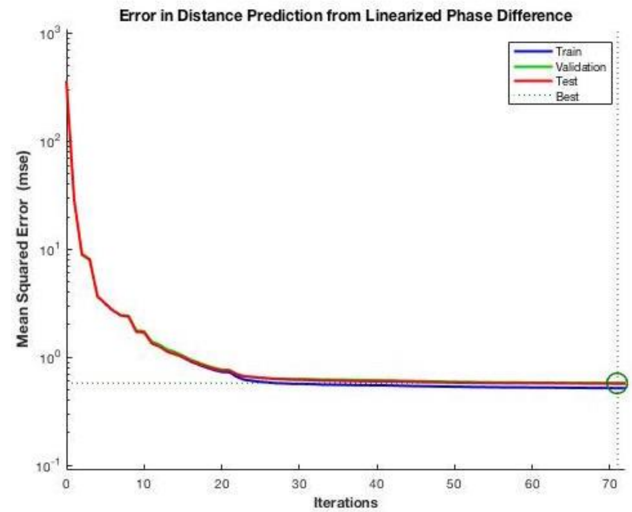

Figure 9: The Learning Curve of the Neural Network for Predicting Distance

Predicting Distance	Samples	MSE
Training	125923	0.518m
Validation	26984	0.578m
Test	26984	0.570m

Figure 10: Results for Predicting Distance

For distance, through regression, we were able to predict the distance within 0.57 meter accuracy on the test data. The results are detailed in the plot and table above (Fig. 8, Fig. 10).

We originally hypothesized a correlation between amplitude and distance, and phase differences and AoA, but initial results were inconclusive. We next thought to try a variation of this by filtering or transforming the data. We thought that taking the FFT and using a match filter which would give us the amplitudes of the various multipath components and thus the amplitude of the line-of-sight component. However, we discovered that the sampling time is only 10ms which would only allow us to recover frequencies below 50 Hz. This is not nearly fast enough to recover our signal which has a center frequency of 5 GHz. Additionally the distances were too small to see any significant amplitude fall-off due to path loss. As a result, we did not find any correlation between amplitude and distance as we initially predicted.

On the other hand, we found that phase difference data was sufficient to make both AoA and distance predictions. Intuitively, our neural network learns a phase difference "fingerprint" or specific pattern for each AoA and distance and makes predictions based on which "fingerprint" the new data point matches best. While our predictions for both AoA and distance were less accurate than the MUSIC algorithm used in SpotFi [3], our method makes predictions much faster than MUSIC does. Our predictions also have much better accuracy than the AoA and phase predictions found in the DeepFi [1] and PhaseFi[2] papers which was another one of our original goals. Thus, we have taken a small hit in accuracy relative to the SpotFi results [3], in return for a large increase in speed.

Furthermore, varying the number of layers used in the neural network gives us further control in the tradeoff between accuracy and speed. More layers usually increases accuracy, but decreases training and prediction speed. This was a trend we noticed as we increased the layers in our neural network, suggesting that we did not overfit our training set as increasing training was still leading to test error decreases.

VI. Conclusion and Future Work:

In conclusion, using neural networks, we were able to classify angle with an accuracy of around 86% and predict distance within around 0.5 meters (~10% of mean distance).

There are a few future ideas that could be implemented to extend this project. A different unknown phase offset is introduced by the receiver circuitry at each of the three receive antennas. Furthermore, the value of this offset changes every time the receiver is powered off and on again. Thus, our phase difference data has “noise” built into it that we did not account for. To improve our results, we would try removing this phase offset by measuring the phase offset at each antenna, and then subtracting this value from the raw data before training our neural network. This would most likely increase accuracy for both the distance and angle of arrival prediction.

Another consideration is that the phase difference “fingerprint” depends on the geometry of, and the arrangement of reflectors in, the room. As such, it is possible that even a small change in the environment could cause a large change in the accuracy of the learned model. Therefore, an important next step is change the arrangement of objects in the room without re-training the neural network, and observe whether the accuracy of the predictions change drastically.

Finally, if more space were available, we would try collecting data at longer distances. We believe that the signal attenuation rate was too small for our receiver to detect a noticeable change in amplitude over the distance range of our experiment. Thus, different set of data over long enough distances might provide predictions to be made from amplitude data.

VII. References:

[1] X. Wang, L. Gao, S. Mao, S. Pandey, “CSI-based Fingerprinting for Indoor Localization: A Deep Learning Approach,” IEEE Transactions on Vehicular Technology, 2016.

[2] X. Wang, L. Gao, and S. Mao, “PhaseFi: Phase fingerprinting for indoor localization with a deep learning approach,” IEEE Internet of Things Journal, Dec, 2015.

[3] M. Kotaru, K. Joshi, D. Bharadia and S. Katti, “SpotFi: Decimeter Level Localization Using WiFi,” SIGCOMM ’15, August 17 - 21, 2015

[4] D. Vasisht, S. Kumar, D. Katabi. "Decimeter-Level Localization with a Single WiFi Access Point," Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation, March 16-18, 2016.

[5] J. Biswas, M. Veloso. "WiFi Localization and Navigation for Autonomous Indoor Mobile Robots," IEEE International Conference on Robotics and Automation, 2010.