IBLoc-UAV: Inferring On-Body Location in UAV-to-Ground Channels

Mahmoud Badi[1] and Joseph Camp[2]

[1]Qualcomm Atheros, Inc.*, Santa Clara, California, USA
[2]Southern Methodist University, Dallas, Texas, USA

***Abstract*—Drones frequently communicate with devices on the ground, often carried in different ways by users. A user might hold their device near their chest while flying a drone, walk with their phone in their pocket, or walk facing away from the drone while it is tracking them. Despite the proliferation of drones and possibility of such scenarios, the prediction of user orientation and the user equipment (UE) location on or near the user's body in Drone-to-Ground (D2G) channels has not received adequate attention in literature. In scenarios where visibility is no longer available to the drone – due to adversarial attacks or harsh weather – the wireless signal can be used to detect the user's presence, their orientation, and even the location of the device on or near their body. This is the objective of this work. We study how the baseband I/Q samples, converted into spectrogram images that span a relatively short period of time, can be used to predict the on-body device location. We leverage Convolutional Neural Networks (CNNs) to classify three different use cases of holding a device operating at two different carrier frequencies (2.5 GHz and 900 MHz). Specifically, we investigate three on-body locations: near chest while facing the drone, in-pocket while facing the drone, and near chest but facing away from the drone. We show that, using only spectrogram images as input, we can predict these use cases with an average overall accuracy of 87% and 85% at 2.5 GHz and 900 MHz, respectively. We also investigate the classification performance on a dataset that belong to a hovering location that was not seen by the model, and show that the CNN model was able to correctly classify all images belonging to user orientation and 91% of the images that belong to near-chest vs in-pocket for the same orientation. Finally, we study the application of transfer learning and CNNs on classifying the on-body location at a different carrier frequency from the one on which they were trained, and show that while one use case can be correctly predicted, more complex models and hyper-parameter tuning is needed to achieve this goal. This work could be useful for building real-time deep learning models that help drones to make intelligent decisions and adapt to changes in user postures and on-body locations in air-to-ground (A2G) channels.**

***Index Terms*—Air-to-Ground Channels, UAVs, Drones, Human Body Effects, Deep Learning, Convolutional Neural Networks.**

I. Introduction

The use of Unmanned Aerial Vehicles (UAVs) has seen a substantial increase in recent years. Drones are now being used in structural health inspection [1], object tracking, entertainment, delivery, and agriculture applications [2]. Drones and phones rely on their cameras, assisted by other sensors, to obtain an understanding of the surrounding environment and be able to track objects. For example, Google's ARCore Software Development Kit (SDK) uses Augmented Reality (AR) to perform multiple tasks such as localization, depth analysis, and motion tracking [3]. However, when visibility is poor due to obstructions, poor lighting, or harsh weather, object localization and/or tracking using a camera might no longer be an option, and relying on a different method, such as wireless signals, could be a good alternative.

A. Motivation: Impact of Human Body Blockage on Ranging

When performing ranging and localization using wireless signals, the channel and surrounding environment can significantly impact performance. For example, in Round Trip Time (RTT) based ranging protocols, such as IEEE 802.11mc, the time of arrival (ToA) of the line-of-sight (LOS) component is critical to the accuracy of the ranging measurement. The accuracy of the estimated ToA depends on the bandwidth used and the nature of the multipath channel [4]. The richer the channel in multipath reflections, the more unreliable and inaccurate the ToA and RTT ranging are, respectively. This is due to confusion in estimating the direct path due to multiple reflected paths. Furthermore, when an object is blocking the path between two devices that are performing ranging, the direct path might be severely attenuated compared to diffracted/reflected paths. This results in an incorrect multipath component being identified as the direct path leading to an erroneous ranging estimate. Therefore, when performing ranging, it is desirable to know if the target is blocked or not to better evaluate the reliability of the ranging measurement.

Fig. 1 shows the results of a ranging dataset that we captured using the IEEE802.11mc ranging protocol with and without the blockage of a human body along the path. Our experiments were carefully designed and conducted in an anechoic chamber with minimal reflections to emphasize the impact of human body on RTT ranging error. A phone was placed on a fixed platform while a person was holding the other phone and performing RTT ranging. The ranging error is then calculated as the difference between the true distance and what is reported by the RTT ranging application. From the figure, we can see that when a person holds the phone while facing away from the other device, the ranging error significantly increases, reaching up to 4 m in some instances. This confirms that the attenuation

*During the experiments, data collection, and the conceptualizing phases of this work, Mahmoud Badi was with the Department of Electrical and Computer Engineering, Southern Methodist University. He is currently a Staff Engineer with Qualcomm Atheros, Inc., Santa Clara, USA. Email :{mbadi}@qti.qualcomm.com

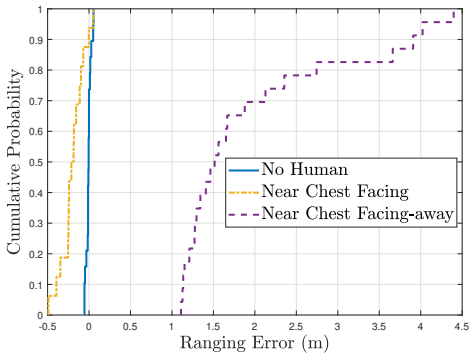

Fig. 1. Impact of a human body on ranging error in an anechoic chamber.

of the direct path and the reflections and diffraction caused by the human body contribute to the uncertainty in ranging measurements. This result is one example that motivates the need for models that can accurately capture the orientation of the user and location of the device on their body.

Since electromagnetic models that aim to describe such scenarios are cumbersome to derive, computationally expensive, and difficult to integrate with UAV-based channels, it is desired to leverage relatively simple and lightweight [5] deep learning models to predict different user postures, orientations, and on-body device locations, to enable drones to dynamically and intelligently adapt to channel changes on the ground.

B. Contributions

In this work, we leverage an in-field dataset of UAV-to-Ground channels that we collected and analyzed in [6] and build (offline) CNN deep learning models to predict the orientation of the user and the on-body location of the communication device. We investigate the effectiveness of using only spectrograms as an input to CNNs in achieving this task. The hope is for this work to help researchers build lightweight, real-time CNN-based models that can perform UAV-to-User predictions to enable drones to make intelligent decisions regarding their connectivity, energy consumption, or path planning. For example, once the drone determines that the user is facing away, it can better evaluate the reliability of its ranging measurement before deciding to get closer or further from them. Such decision, if successful, might result in better energy consumption by leveraging a more appropriate power level at the new drone hovering position. While other works, such as [5], [12], [14], have extensively studied the use of CNNs for wireless signal and channel classification, this work – to the best of our knowledge – is the first to use only I/Q samples and CNNs to predict the on-body device location in UAV-to-Ground channels.

Our contributions can be described as follows:

- We propose a framework for investigating the use of CNN-based models to classify UAV-to-Ground channels. We build and validate the trained model (offline) against a dataset that spans two different frequencies and many hovering positions for three different use cases[1] of a person holding a device on the ground.
- We show that lightweight CNNs[2] using only spectrogram images as input, can predict the user orientation and device location in drone-to-ground (D2G) channels with reasonable accuracy (85% to 87%). We also demonstrate the CNN model's ability to predict the use case on datasets that were not seen by the model (not included in the training/validation phase) at the furthest drone hovering locations.
- We explore transfer learning and test the ability of one model, which was built for one carrier frequency, to predict user orientation and device location at another frequency. We show that while one use case was correctly predicted, more hyper-parameter tuning and potentially more complex models are needed to achieve this task across all use cases.
- Finally, we provide access to the code and datasets used in our study, enabling further exploration and experimentation with UAV-to-Ground classification using CNNs.[3].

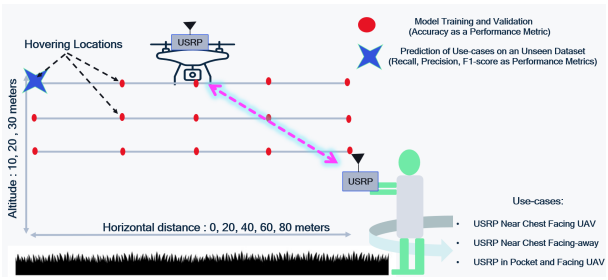

Fig. 2. An illustration of the UAV Air-to-Ground experiments that we conducted in [6]

The rest of the paper is organized as follows. In Section II we describe the signal model. In Section III we briefly introduce the experiments and data collection process. The data preparation and processing is described in Section IV and the deep learning CNN model training and validation is discussed in Section V. Related work is discussed in Section VI and the concluding remarks are presented in Section VII.

II. Signal Model

We can describe the transmitted signal as $x(t) = m(t)\cos(2\pi f_c t + \phi_o)$, where $m(t)$ is the message signal, f_c, is the carrier frequency which is generated by the local oscillator (2.5 GHz or 900 MHz) with a random phase offset of ϕ_o. Due to reflections and scattering caused by the UAV body [7], ground induced multipath reflections, and the human body if it exists in the path of the signal, the received signal

[1]We use the terms use cases and on-body locations interchangeably throughout this work.

[2]The work in [5] builds a real-time CNN model and compares computation resources showing the lightweight nature of their CNN model. Our work uses a simpler CNN architecture than [5]. Hence, our claim of the presented CNN as lightweight seems valid.

[3]The code and dataset for the proposed framework can be found here: http://muddi.lyle.smu.edu/Repository_human.html

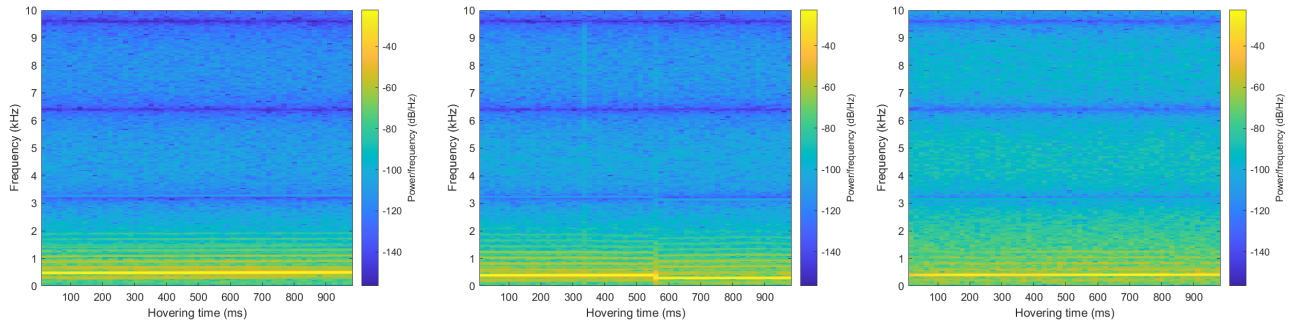

(a) (b) (c)

Fig. 3. Spectrogram of the received baseband I/Q samples from a hovering time of 1 second. Samples shown are from the three different use cases: (a) Near-Chest Facing, (b) In-pocket Facing, and (c) Near-Chest Facing Away.

will be a combination of multiple paths. If, $m(t)$ = 1, then the received signal can be written as:

$$r(t) = \Re\{[\sum_{n=0}^{N(t)} \alpha_n(t) e^{-j\Phi_n(t)}] e^{j2\pi f_c t}\} \quad (1)$$

Here, $N(t)$ is the number of multipath components, and the phase term Φ is given by $\Phi_n(t) = 2\pi f_c \tau_n(t) - \phi_{D_n} - \phi_o$. The propagation delay of the n^{th} component is denoted by τ_n, and it is equal to $d_n(t)/c$ with $d_n(t)$ being the separation distance in meters and c being the speed of light. $\alpha_n(t)$ is the amplitude of the n^{th} multipath component, and ϕ_{D_n} is the Doppler shift. In terms of in-phase and quadrature components, we can rewrite the received signal as [6]:

$$r(t) = r_I(t)\cos(2\pi f_c t) + r_Q(t)\sin(2\pi f_c t) \quad (2)$$

Here, $r_I(t) = \sum_{n=0}^{N(t)} \alpha_n(t)\cos\Phi_n(t)$ and $r_Q(t) = \sum_{n=0}^{N(t)} \alpha_n(t)\sin\Phi_n(t)$ are the in-phase and quadrature components, respectively. We can see that the received signal will vary over time. Fluctuations due to hovering, blockage caused by the human body, and reflections from the ground will all contribute to the time-varying nature of the signal.

III. Experiment Setup and Procedure

Here, we briefly discuss the experiments that were carried out to collect the datasets. For more details, we refer the reader to [6]. We used two Universal Software Radio Peripherals (USRP) E312s from Ettus Research™. The transmitting radio and antenna are mounted on the UAV, while the receiver USRP is held by the user on the ground. The transmitter and receiver are sampling at 64k samples/s and measurements are recorded for a period of 20 seconds per hovering position. In this work, we assume channel reciprocity [8], and use the dataset as if it were collected at the drone, not by the user on the ground.

We investigate three on-body locations: (*i*) Near chest and facing (NCF) towards the Tx UAV, (*ii*) Near chest and facing away (NCFA) from the Tx UAV, and (*iii*) In-pocket while facing (IPF) the Tx UAV. For each use case, we perform drone-to-ground experiments at carrier frequencies of 900 MHz and 2.5 GHz, totaling 6 experiment sets. In each experiment, three altitudes and 5 horizontal distances are visited by the drone, where it hovers for 20 seconds. The hovering positions are shown in Fig. 2.

IV. Data Pre-processing

The received I/Q baseband samples are processed as follows. Due to rapid fluctuations in the signal caused by drone hovering, we apply a moving average with a length of $W = 10$. We experimented with and without using a moving average and observed better prediction results after applying the moving average. The selection of a moving average of a length of 10 was chosen after seeing about 10% increase in validation accuracy compared to not using a moving average. In [5] the authors saw a similar improvement (6%) in performance after using a smoothing filter. We also tested the model's performance with I only as input to the spectrogram function, Q only, and the envelope (I+Q), and there was little to no difference in results. Hence, we opted to use only one component of the I/Q samples to consume less computation time and resources.

The smoothed samples are then fed to the MATLAB function *Spectrogram* with a Hamming window and a Fast Fourier Transform (FFT) of size N=1024. The function computes the short-term Fourier Transform on the windowed data over 1 second and creates a spectrogram plot. The spectrogram is then saved as a PNG file. The process is repeated for all 3 scenarios across all the hovering locations. The total number of spectrogram images per frequency band is 3 (on-body locations) x 5 (distances) x 3 (altitudes) x 20 (images/second per hovering location) = 900 images. The total number of images is 1800 images spanning 2.5 GHz and 900 MHz data. A sample of the spectrogram of the three on-body locations is shown in Fig. 3 and the data processing flow is shown in Fig. 4. We can see from Fig. 3 that, to the naked eye, there seems to be little to no difference between the spectrogram images of the three use cases. Hence, we rely on CNNs due to their proven ability to effectively extract discriminating features from images [14], [15].

Since this work performs the classification offline (on a previously collected dataset), we are able to label the images according to the use-cases. In a real-time scenario of CNN-based prediction of user-cases, the labeling of the images

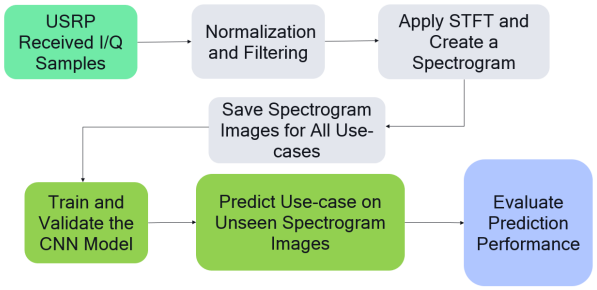

Fig. 4. Our framework for processing the received baseband I/Q samples and turning the problem into a deep learning CNN model based prediction.

can happen in a previous stage using sensors or cameras prior to the wireless-based prediction stage. As previously mentioned, in scenarios where vision is no longer available due to adversarial attacks or harsh weather/environment, relying on wireless signals to detect user orientation and device location can be a good alternative.

V. CNN Model Training and Validation for Air-to-Ground Channel Classification

In this section, we describe the deep learning CNN model we use. Throughout this work, we use Keras library with TensorFlow (TF) [9].

A. CNN Model Description

Splitting the dataset: The only input to the CNN model is the spectrogram images we described in the previous section. The data was split into two subsets, one for training and one for validation with a 80/20 split. This ratio was chosen after experimenting with 70/30 and 85/15 and found out that the former resulted in poor performance while the latter resulted in overfitting.

Normalization: Images are normalized to have their values between 0 and 1 instead of 0 to 255 which are the default RGB channel range values.

CNN architecture: The neural network follows the sequential model mentioned here [9]. It comprises three convolutional blocks with a maxpooling layer in each one of these blocks. At the end of these blocks there is a fully connected (FC) layer that is activated by a ReLU function. It is worth mentioning that we tried to use other activation layers such as Sigmoid but the prediction performance significantly deteriorated. Hence, ReLU was the activation layer of choice. The CNN architecture we chose is similar to [9] due to its validated performance in recognizing images. Further details are available as comments in the code we shared.

The optimizer algorithm: we chose Adam as the optimizer algorithm due to its success in classification tasks and computational efficiency [10]. The default values for the Adam optimizer were chosen. The learning rate is $\alpha = 0.001$; the exponential decay rate of the first and second moment estimates are $\beta_1 = 0.9$, $\beta_2 = 0.999$, respectively. The numerical stability constant, ϵ, is set to default value of 10^{-7}.

B. Training and Validation Accuracy

We use the holdout method with 80% of the data used for training and 20% for validation. We use 40 epochs for both training and validation. We evaluate the model's performance via an accuracy metric, which is defined as the number of correct predictions across *all* three use-cases divided by the number of all predictions. In other words, $Accuracy = \frac{Number of Correct Predictions}{Total Number of Predictions}$. For example, if we correctly classified the three use-cases 830 times, then, $Accuracy = \frac{830}{900} = 92.2\%$.

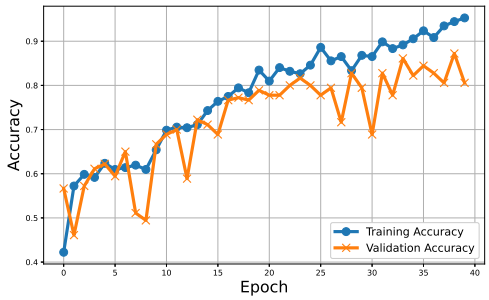

Fig. 5. Training and validation accuracy of the 2.5 GHz dataset spanning the three use cases.

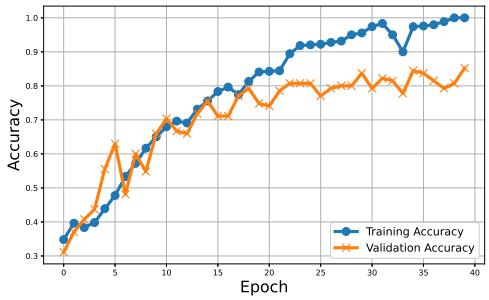

Fig. 6. Training and validation accuracy of the 900 MHz dataset spanning the three use cases.

The prediction accuracy of the model on 2.5 GHz are shown in Fig. 5. We can see that the 2.5 GHz prediction accuracy can reach about 87% at the 39th epoch. In the 900 MHz dataset, the validation prediction accuracy is about 85%. This can be seen in Fig. 6. It is worth noting that after investigating the model's performance, we found that most of the false predictions come from confusing the NCF with the IPF scenario or vice versa and not with the NCFA test case. Note that the accuracy results span the whole dataset, *i.e.*, all three use cases and does not address potential issues such as class imbalance or the weight and impact of misclassifying specific use-cases.

C. Testing the Model on Unseen UAV Hovering Datasets

In the previous section, we trained the model and validated its prediction performance over the entire dataset and across the 3 use cases (NCF, NCFA, and IPF). Here, we will train the model on a portion of the dataset and then test it on another unseen portion of I/Q images. The unseen images belong to

I/Q samples from the three use cases when the drone was hovering at the furthest distance of 80 m horizontal distance and 30 m altitude (a distance of $d_x = \sqrt{80^2 + 30^3}$=85.4 m). This hovering position is denoted by the star in Fig. 2.

The purpose of this exercise it to address the validity of predicting UAV A2G channels at new (and further) flying locations after a period of training in previous and closer hovering locations. This investigation is useful for applications when the drone had enough time to label some portion of the dataset using knowledge obtained from cameras and sensors, but then loses its vision due to adversarial conditions (attacks, weather, etc.). The question we pose is this: if we had enough time to train the drone at relatively close locations to understand the user orientation and device location on/near their body, can it still predict the orientation/location of the person/device at further locations relying solely on the generated spectrogram images that span a short time interval?

Prediction metric: We test the prediction performance of the model on unseen spectrogram images belonging to each scenario/test-case when the flying drone is hovering at 80 m horizontal distance and 30 m altitude. We test the prediction performance on 20 images at 2.5 GHz and 15 images at 900 MHz. While accuracy is a good performance metric, it fails to address class imbalance issues and it does not provide insight on the per-class/test-case prediction capability of the model. For these reasons, we use Recall, Precision, and F1-score as metrics to judge the prediction capability of the CNN model. The prediction decision has to have a confidence score of 50% or more to be considered. We then calculate Precision, Recall, and the F1-score.

Precision is calculated as the number of times the model correctly predicted the class/test-case divided by the *total* number of times it made that prediction across all classes, *i.e.,* both correct and false predictions.

Recall is calculated as the number of times the model correctly predicted the class divided by the total number of tested images in that class.

F1-score combines both recall and precision according to the formula $F1 = 2 \times \frac{Recall \times Precision}{(Recall + Precision)}$. It has both contributions from recall and precision, *i.e.,* it includes the number of times we falsely predicted a certain class and the number of times we falsely missed it.

The results of the 2.5 GHz and 900 MHz drone-to-ground classifications are summarized in Table I and Table II, respectively. We make the following observations: the near-chest facing-away (NCFA) scenario was perfectly predicted by the CNN model across all unseen images. This is evident by the perfect recall, precision, and F1-score results. This finding means if a UAV air-to-ground channel was trained on I/Q samples at short distances with labels pertaining to the orientation of the user, we can completely rely on the I/Q baseband samples to accurately predict the orientation of the user and determine whether they are facing or facing-away from the drone at further distances.

When the person was facing the drone while the device was placed near their chest (NCF) or in their pocket (IPF), the model's average F1-score was 0.91. This means, on average, the predicted location of the device was correct 91% of the time across the two frequencies. While these are the prediction results at one hovering location, we expect to see similar results at other hovering positions as long as the surrounding environment remains the same.

TABLE I
CNN MODEL PREDICTION ACCURACY OF SPECIFIC USE-CASES ON UNSEEN DATA AT 2.5 GHZ

Metric/Class	Near Chest Facing	Near Chest Facing-away	In-pocket Facing
Precision	0.94	1	0.86
Recall	0.85	1	0.95
F1-Score	0.89	1	0.90

TABLE II
CNN MODEL PREDICTION ACCURACY OF SPECIFIC USE-CASES ON UNSEEN DATA AT 900 MHZ

Metric/Class	Near Chest Facing	Near Chest Facing-away	In-pocket Facing
Precision	0.88	1	1
Recall	1	1	0.87
F1-Score	0.93	1	0.93

D. Transfer Learning: Using a Dataset of One Frequency to Predict A2G Channels at Other Frequencies

In this section we try to answer the following question: If we train the CNN model to predict an A2G channel using a labeled dataset from one frequency, can it accurately predict the A2G channel at other frequencies? We believe that such a question is worth investigating due to the potential of reducing the training phase of multi-frequency channels.

In this investigation, we train the CNN model on a dataset that belongs to one frequency (*e.g.,* 2.5 GHz) and try to classify unseen images that belong to the three use-cases at one drone hovering location at another frequency (*e.g.,* 900 MHz).

Training the model on 900 MHz to predict on-body locations at 2.5 GHz: We followed the same training and validation procedure mentioned in sec V-B. Then, the model, after reaching the 85% validation accuracy, was used to classify the 2.5 GHz images that belonged to the three use-cases. We note the following: in predicting the NCF (facing) use-case, the model did very well resulting in correctly classifying the 15 images with an average confidence of 93%. In other words, the recall was 100% with a confidence of 93%. On predicting the class of NCFA (facing-away) images, it also correctly predicted all images (recall of 100%) with a confidence of almost 100%. However, when it was tasked to predict the class of the IPF (in-pocket) use-cases, it failed. It only correctly predicted 1 image with the rest of images being falsely predicted as NCF with high confidence.

Training the model on 2.5 GHz to predict on-body locations at 900 MHz: The results of this prediction were worse than the previous one. The model was only able to achieve a good recall of 93% on the NCFA class. Unfortunately, it made the false prediction of NCFA on almost all images resulting in a large number of false positives and poor prediction performance for the other two on-body locations (IPF and NCF).

From the above result, we conclude that for CNN-based transfer learning to be successful in predicting the on-body device location in A2G channels, further investigation is required. This investigation includes studying the potential need for more complex models and/or hyper-parameter tuning, more sophisticated signal processing techniques, or more measurements to find patterns and get better insight.

VI. Related Work

The use of machine learning and deep learning methods to predict and/or classify wireless channels has seen a significant increase in recent years. For example, the work in [11] used support vectors and neural networks to predict ground-to-UAV channels in smart farming applications; it showed that these methods can outperform ray tracing methods. However, only the prediction of the received signal strength (RSS) or path loss of user-free channels was studied. When it comes to using deep learning models that leverage CNNs to predict or classify wireless channels, the literature has shown that CNNs can perform this task with great accuracy. For example, the work in [5] deploys a lightweight CNN model on an SDR platform and performs real-time classification of different WiFi signals resulting in comparable accuracy to other methods that are more resource-intensive. The work in [12] shows that it is possible to train CNNs to distinguish between *simulated* LOS and NLOS ultra-wideband channels. The work in [14] has shown that using CNNs, combined with a Walsh-Hadamard Transform (WHT), twelve *emulated* wireless channel classes – that were determined by the Rician K-factor – were accurately classified even in relatively low SNR regions. When wireless channels involve humans performing certain activities, CNNs have been shown to perform well in predicting those activities. For example, the work in [13] shows that using CNNs along with the transmission and reflection coefficients, six different activities were predicted with an accuracy of more than 97%. In [15], the authors show that using Mel-spectrogram and CNNs, they were able to recognize coughs with an accuracy of 98% outperforming other machine learning methods such as K-means and support vector machines (SVMs).

While the above mentioned works offer great insight regarding the ability of CNNs in learning and performing predictions with impressive accuracy, none have investigated CNNs with *real* UAV A2G channels that involve humans as part of the channel. This study offers the first initial insight regarding the use of spectrograms and CNNs to classify UAV A2G channels at different drone hovering positions, frequencies, and for various human use-cases of holding a device.

VII. Conclusion

In this work, we leveraged a previously collected dataset [6] that spans different frequencies and use-cases of A2G channels, and built (offline) deep learning CNN models with the objective of classifying these channels using *only* spectrogram images as an input. We showed that using lightweight CNNs, we can classify three different use-cases of holding a communication device with an accuracy of 85% to 87%. We also tested the CNN model on unseen data that belonged to the furthest UAV hovering position and demonstrated that the CNN can still accurately predict the orientation of the user and location of the device. Finally, we investigated the use of transfer learning to classify the channels at other frequencies and showed that more complex hyper-parameter tuning and/or models are needed to achieve this task. This work serves as a preliminary investigation to open the door for the design and implementation of real-time CNN models that can make drones communicate with users on the ground more intelligently by dynamically adapting to the changing environment around them. These design and implementation efforts will involve studying computational resources, ML model generalization to unseen scenarios, and applying transfer learning to other frequencies.

References

[1] Meira, G.d.S.; Guedes, J.V.F.; Bias, E.d.S., " UAV-Embedded Sensors and Deep Learning for Pathology Identification in Building Façades: A Review ," Drones 2024, 8, 341. https://doi.org/10.3390/drones8070341

[2] Yang, S.; Li, L.; Fei, S.; Yang, M.; Tao, Z.; Meng, Y.; Xiao, Y., " Wheat Yield Prediction Using Machine Learning Method Based on UAV Remote Sensing Data ," Drones 2024, 8, 284. https://doi.org/10.3390/drones8070284

[3] Google AR Core website: https://developers.google.com/ar

[4] C. Ma, B. Wu, S. Poslad and D. R. Selviah, "Wi-Fi RTT Ranging Performance Characterization and Positioning System Design," in IEEE Transactions on Mobile Computing, vol. 21, no. 2, pp. 740-756, 1 Feb. 2022, doi: 10.1109/TMC.2020.3012563.

[5] Cetin, R.; Gecgel, S.; Kurt, G.K.; Baskaya, F., " Convolutional Neural Network-Based Signal Classification in Real Time ," IEEE Embed. Syst. Lett. 2021, 13, 186–189

[6] M. Badi, S. Gupta, D. Rajan and J. Camp, "Characterization of the Human Body Impact on UAV-to-Ground Channels at Ultra-Low Altitudes," in IEEE Transactions on Vehicular Technology, vol. 71, no. 1, pp. 339-353, Jan. 2022, doi: 10.1109/TVT.2021.3122413.

[7] M. Badi, J. Wensowitch, D. Rajan and J. Camp, "Experimentally Analyzing Diverse Antenna Placements and Orientations for UAV Communications," in IEEE Transactions on Vehicular Technology, vol. 69, no. 12, pp. 14989-15004, Dec. 2020, doi: 10.1109/TVT.2020.3031872

[8] Y. Shi, M. Badi, D. Rajan and J. Camp, "Channel Reciprocity Analysis and Feedback Mechanism Design for Mobile Beamforming Systems," in IEEE Transactions on Vehicular Technology, vol. 70, no. 6, pp. 6029-6043, June 2021, doi: 10.1109/TVT.2021.3079837.

[9] Image Classification using Tensor Flow. Website: https://www.tensorflow.org/tutorials/images/classification

[10] Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization, " arXiv preprint arXiv:1412.6980 (2014).

[11] Duangsuwan, S.; Juengkittikul, P.; Myint Maw, M., " Path loss characterization using machine learning models for GS-to-UAV enabled communication in smart farming scenarios, " Int. J. Antennas Propag. 2021, 2021, 5524709.

[12] P. ShirinAbadi and A. Abbasi,"UWB Channel Classification Using Convolutional Neural Networks," 2019 IEEE 10th Annual Ubiquitous Computing, Electronics Mobile Communication Conference (UEMCON), 2019, pp. 1064-1068

[13] Y. Kim and Y. Li, "Human Activity Classification With Transmission and Reflection Coefficients of On-Body Antennas Through Deep Convolutional Neural Networks," IEEE Trans. Antennas Propag., Vol. 65, no. 5, pp. 2764-2768, May 2017

[14] G. Baldini, F. Bonavitacola and J. -M. Chareau, "Fading Channel Classification with Walsh-Hadamard Transform and Convolutional Neural Network," 2023 International Conference on Smart Applications, Communications and Networking (SmartNets), Istanbul, Turkiye, 2023, pp. 1-6, doi: 10.1109/SmartNets58706.2023.10215941.

[15] Zhou Q, Shan J, Ding W, Wang C, Yuan S, Sun F, Li H, Fang B. , " Cough Recognition Based on Mel-Spectrogram and Convolutional Neural Network, " Front Robot AI. 2021 May 7;8:580080. doi: 10.3389/frobt.2021.580080. PMID: 34026854; PMCID: PMC8138471.