
UAVs using Bayesian Optimization to Locate WiFi Devices

Mattia Carpin

University of Padova
Padova, Italy

mattiacarpin@gmail.com

Stefano Rosati*

Ecole Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

stefano.rosati@epfl.ch

Mohammad Emtiyaz Khan*

Ecole Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

emtiyaz@gmail.com

Bixio Rimoldi

Ecole Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

bixio.rimoldi@epfl.ch

Abstract

We address the problem of localizing non-collaborative WiFi devices in a large region. Our main motive is to localize humans by localizing their WiFi devices, e.g. during search-and-rescue operations after a natural disaster. We use an active sensing approach that relies on Unmanned Aerial Vehicles (UAVs) to collect signal-strength measurements at informative locations. The problem is challenging, because the measurements are received at arbitrary times and they are received only when the UAV is in close proximity to the device. For these reasons, it is extremely important to make prudent decisions with very few measurements. We use the Bayesian optimization approach based on Gaussian process (GP) regression. This approach works well for our application since GPs give reliable predictions with very few measurements while Bayesian optimization makes a judicious trade-off between exploration and exploitation. In field experiments conducted over a region of $1000 \times 1000 m^2$, we show that our approach reduces the search area to less than 100 meters around the WiFi device within 5 minutes. Overall, our approach localizes the device in less than 15 minutes with an error of less than 20 meters.

1 Introduction

We consider localization of WiFi devices in a large region using an unmanned aerial vehicle (UAV). Due to the widespread use of commercial WiFi-enabled devices (smartphones, tablets, laptops, etc.), the location of WiFi devices can be used to localize people who own those devices. Such localizations are important for many applications, e.g. to deliver a mail package for companies such as Amazon and Google, to provide internet connectivity for companies like Facebook, and perhaps most importantly, in search and rescue operations to localize victims.

Localization of people could potentially be performed using cameras mounted on a UAV [1], but this requires a clear view of the target and a great amount of image processing. Our approach is to localize people via their WiFi devices by measuring the received signal strength index (RSSI). The RSSI measurements concern the Probe Request Frame (PRF) which is frequently broadcasted by a WiFi device to actively scan the environment and discover access points (e.g. see [2]).

*Equal contributions

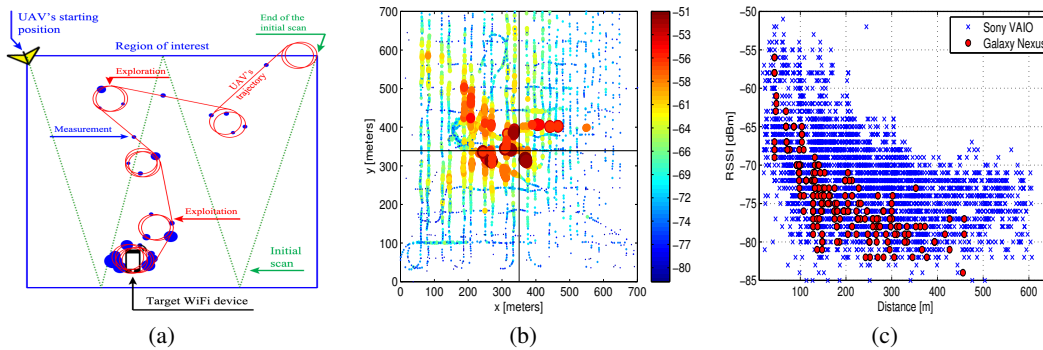


Figure 1: Figure (a) shows the typical path followed by the UAV during the localization process. Blue circles show the data, and their size is proportional to the RSSI values. Figure (b) shows the RSSI measured during an exhaustive scan of the area. Large red dots indicate high RSSI values, while small blue dots indicate low values. Figure (c) shows RSSI as a function of the distance between the UAV and device.

We show that, by actively exploring the region using a UAV, we can locate WiFi devices quickly and accurately. Our approach relies on the fact that the expected value of an RSSI measurement increases as the distance between UAV and the device decreases. The UAV therefore tries to get to a location where the expected value of the RSSI measurement is highest.

Since collecting measurements is costly, both in terms of time and battery consumption, we use Bayesian optimization which allows the UAV to collect measurements in informative locations to find the WiFi device quickly and with high accuracy. More importantly, in our application, the environment is unknown and a “black-box” approach, such as Bayesian optimization, is appropriate since it avoids making assumptions about the environment. Our approach therefore is completely data-driven.

Fig. 1(a) summarizes our approach for locating a single WiFi device. In the beginning, the UAV does a quick scan of the area to get a few initial measurements. Based on these measurements, the UAV computes an estimate of the distribution of the RSSI over the whole region. The UAV then decides to fly to a new location either to “exploit”, i.e., move to a location which is more likely to give higher RSSI measurements, or “explore”, i.e., go to a previously unexplored location. We use a Bayesian optimization based approach which automatically trades-off between exploitation and exploration by using an *acquisition* function. This process is repeated after a new measurement is received. The algorithm terminates when the confidence in the location estimate is reasonably low.

2 Problem Definition, Data Characteristics, and Challenges

Our goal is to localize a WiFi device placed in an unknown location \mathbf{x}^{true} inside a two-dimensional region \mathcal{X} . The RSSI measurements are received at random times. Let us suppose that by the time t , we have received n_t measurements with time-stamps t_1, t_2, \dots, t_{n_t} such that $t_1 < t_2 < \dots < t_{n_t} < t$. Denote by \mathbf{x}_{t_i} the location of the UAV when the i 'th measurement was made and by y_{t_i} the corresponding RSSI measurement. Note that the RSSI measurement is a random process Y that depends on many factors such as the current position of UAV, the true location of the device and the time of measurement, but for notation simplicity we only show dependency w.r.t. the position of the UAV at time t , i.e. \mathbf{x}_t . Therefore, at any time t , we have a sequence of triplets: $\{t_i, \mathbf{x}_{t_i}, y_{t_i}\}$ for $i = 1, 2, \dots, n_t$ such that $t_i < t_{i+1} < t$. Denote the set of triplets at time t by \mathcal{D}_t .

Fig. 1(b) shows the measurements gathered from a laptop (Sony Vaio Pro-13 ultrabook) during an exhaustive scan over the region that took more than 60 minutes. The device is located in the center of the figure, and the size and color of points vary according to signal strengths. We clearly see that the signal strength is on average highest around the device, justifying our search for the maximum value for the expected RSSI to localize the device.

Fig. 1(c) shows the measured RSSI as a function of the distance between the UAV and device. Obviously, the measurements are stochastic due to the noise introduced during propagation through

the channel that includes reflection, multipath scattering and shadowing. Therefore, at different distances, the distribution of the RSSI is different, making the measurements *heteroscedastic*. Interestingly, different devices show different behaviour.

Furthermore, the RSSI measurements are collected at random times since the original time of broadcast is unknown. The probability of dropping a packet, and consequently of missing a measurement also increases with the distance, violating the *missing-at-random* assumption [3].

Finally, in our scenario, \mathcal{X} could potentially be large and therefore it is crucial to develop an algorithm that is capable of collecting measurements where they really matter. We need to make good decisions with very few measurements, otherwise we may end up performing an exhaustive scan of \mathcal{X} which will defeat the whole purpose. As a result of these issues, naive methods such as hill-climbing do not work for our problem. Such methods need a very good estimate of the gradient, and such estimate would require far more measurements than the ones we have.

3 Our Approach

In our approach, the UAV first performs a coarse scan over the region. For a region of around $1000 \times 1000 m^2$, this takes around 3 mins. Afterwards, the UAV repeatedly performs the following tasks whenever a new measurement is received.

1. Using the data \mathcal{D}_t at time t , the UAV computes the (predictive) distribution $p(y|\mathbf{x}_*, \mathcal{D}_t)$ using Gaussian process (GP) regression. This is the probability density function of the measurements y at candidate locations $\mathbf{x}_* \in \mathcal{X}$ given \mathcal{D}_t .
2. Given the distribution, it goes to a new location that maximizes the *expected improvement* function. This step is based on the Bayesian-optimization framework [4] and automatically makes a trade-off between *exploitation* and *exploration*.
3. After moving to the new location, the UAV waits for a new measurement. When a new measurement is received, it updates the estimate of the true location and decides whether to terminate the search or to repeat the steps. If no measurement is received, we repeat the step 2 and go to the next best location.

For prediction, we use a GP regression model with a Gaussian likelihood. For the GP prior, we set mean to 0 and use squared-exponential covariance function. We set the hyperparameter using maximum (marginal) likelihood. For computational reasons, we limit the number of candidate locations to 350. Our GP implementation runs on the computer in the UAV. For Bayesian optimization, we use expected-improvement (EI) as our acquisition function. The type of UAV we used in our experiments is described in [5].

4 Experiments and Results

We use real data shown in Fig. 1 (b) to simulate the behavior of the UAV during the localization. In order to simulate the RSSI measurement we build a synthetic model of the channel using the dataset of real observations. The simulator also models the random times at which signals are emitted, and the probability of missing a measurement as function of the distance. Fig. 2(a) shows the average localization error as a function of time, obtained with 1000 independent runs of the algorithm. It is clear that, as the time increases and more packets are received, the UAV improves the accuracy reducing the error, thus confirming convergence of our algorithm. Fig. 2(b) shows the trajectories of multiple runs where we see that the UAV finds the device in every run.

These results suggest that within 5 minutes the UAV is able to reach within 100 meters of the device. This is a huge improvement over the time taken to do a full scan of the region, e.g. the scan of Fig. 1(b) took around an hour. In 5 minutes, one would typically get 10 to 30 RSSI measurements only. Our method therefore uses the available data efficiently to reduce the average distance to around 50 meters.

We also conducted a few field experiments. We show results to locate a Samsung Galaxy Nexus i9250 smartphone, positioned in an area of $1000 \times 1000 m^2$. The UAV followed the trajectory shown in Fig. 3(a). Fig. 2(c) shows the corresponding localization error with time. We see that within 6 minutes the error is reduced to 50 meters, validating the simulation results. The whole flight took around 15 minutes, after which the localization error was 18 meters.

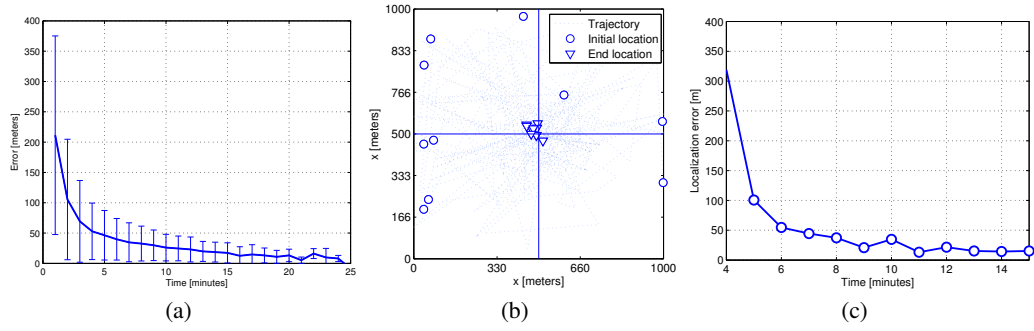


Figure 2: Figure (a) shows simulations results using real data. We plot the mean localization error plus and minus standard deviation as a function of the time. Figure (b) shows the trajectories followed by the UAV during 10 runs of the localization process. The source is at the intersection of the horizontal and vertical line. Figure (c) shows the results obtained from a real-world field experiment where we plot localization error vs time.

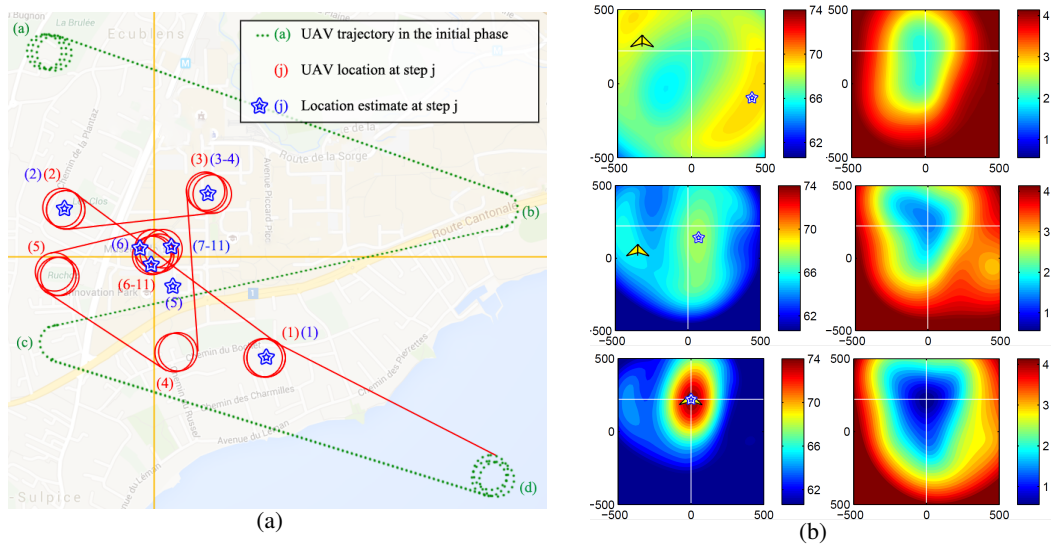


Figure 3: Figure (a) shows the trajectory followed by the UAV during the localization process. The source is at the intersection of the horizontal and vertical line. Figure (b) shows the posterior mean (left column) and standard deviation (right column) at different times during the localization process. The first row shows the estimate around 5 mins when $n_t = 7$ measurements are received. The second and third are around 9 and 15 minutes with 16 and 47 measurements, respectively.

Fig. 3(b) shows the evolution of predictive means and variances. From the same figure, we can also see whether the UAV is in the exploration or exploitation mode. After 5 minutes, the UAV had collected only 7 measurements. The predicted mean at that time was almost flat with a high variance almost everywhere (see the first row Fig. 3(b)). The star indicates the estimated device position, while the other marker shows where the UAV is going to collect the next measurement. As we can see, the UAV decides to do an exploration of the region moving away from the peak (shown with star). The subsequent rows in Fig. 3(b) show predictions around 9 and 15 minutes, respectively. We observe that as time progresses, the mean concentrates around the true location of the device and the variance at the same location reduces. We also see that the UAV keeps exploring at 9 mins and then does exploitation at 15 minutes. The algorithm terminates at 15 minutes, with a localization error of approximately 18 meters.

References

- [1] A. Symington, S. Waharte, S. Julier, and N. Trigoni. Probabilistic target detection by camera-equipped UAVs. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 4076–4081, May 2010.
- [2] K. Benkic, M. Malajner, P. Planinsic, and Z. Cucej. Using RSSI value for distance estimation in wireless sensor networks based on ZigBee. In *Systems, Signals and Image Processing, 2008. IWSSIP 2008. 15th International Conference on*, pages 303–306, June 2008.
- [3] Roderick J A Little and Donald B Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, Inc., New York, NY, USA, 1986.
- [4] Eric Brochu, Vlad M Cora, and Nando de Freitas. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. eprint arXiv:1012.2599, arXiv.org, December 2010.
- [5] Stefano Rosati, Karol Kruzelecki, Louis Traynard, and Bixio Rimoldi. Speed-Aware Routing for UAV Ad-Hoc Networks. In *4th International IEEE Workshop on Wireless Networking & Control for Unmanned Autonomous Vehicles: Architectures, Protocols and Applications*, 2013.