2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN), 13-16 October 2015, Banff, Alberta, Canada

Mean Mutual Information of Probabilistic Wi-Fi Localization

Rafael Berkvens and Maarten Weyn
MOSAIC, Department of Applied Engineering
University of Antwerp – iMinds
Antwerp, Belgium
rafael.berkvens@uantwerpen.be

Herbert Peremans
ENM, Department of Applied Economics
University of Antwerp
Antwerp, Belgium

***Abstract*—This paper focusses on quantifying the localization performance of exteroceptive sensors solely by virtue of their sensor model using information theory. For the third track of the EvAAL competition, we submit our probabilistic Wi-Fi sensor model used for that research. We show the calculation of the mean mutual information for this sensor model. When applying Maximum Likelihood Estimation and k Nearest Neighbor as a localization scheme to our sensor model, we obtain an average Sample Error of 22.15 and 21.97, respectively, on the evaluation dataset. Sample Error is the metric proposed by the EvAAL competition.**

***Index Terms*—Wi-Fi, Localization, Mutual Information**

I. Introduction

This paper focusses on quantifying the localization performance of exteroceptive sensors solely by virtue of their sensor model. We define exteroceptive sensors as sensors that perceive the environment, such as GPS, camera, LIDAR, or Wi-Fi. Continuous localization, or tracking, is often aided by sensors that perceive the state of the object to be localized, like odometers, accelerometers, or gyroscopes, which we call proprioceptive sensors. A localization algorithm uses measurements from such exteroceptive or proprioceptive sensors, or a combination of those, to improve its guess of the location of the device being localized.

The localization performance of a specific sensor strongly depends on the underlying sensor model. This model is responsible for translating a sensor measurement into a likelihood over all possible poses. A localization algorithm will use this likelihood to obtain a pose estimation. The performance of a localization algorithm is generally indicated as a mean distance with a certain standard deviation between the estimated pose and the true pose. While this is a valid metric to compare localization algorithms, the additional processing and assumptions required in the localization algorithm can hide the actual performance of the underlying sensor model. It also assumes a unimodal Gaussian distribution for the pose estimate, which is seldom correct.

We are working on a method to quantify the performance of the sensor in a manner that is both more direct, *i.e.*, not associated with a particular localization scheme, and more general, *i.e.*, not assuming the location estimates to be normally distributed. Our method is based on the mean mutual information between a pose in the environment and the measurement that a sensor makes at that pose.

In order to compare our Wi-Fi sensor model with the state of the art, we join the Evaluating Ambient Assisted Living (EvAAL) competition of the Indoor Positioning and Indoor Navigation (IPIN) 2015 conference. We submit our Wi-Fi sensor model to the third, off-site track of the competition, which is a Wi-Fi localization benchmark competition. It requires competitors to evaluate their localization system on the huge UJIIndoorLoc [1] Wi-Fi fingerprint database. We will present the results of our sensor model by applying a simple Maximum Likelihood Estimation (MLE) and k Nearest Neighbor (KNN) algorithm to the localization posterior.

This paper is continued as follows. In Section II we describe our method for calculating the mean mutual information and our Wi-Fi sensor model. Next, we explain how we applied the MLE and KNN localization algorithms. Finally, we introduce the competition's evaluation metric. In Section III we present what mean mutual information the sensor model obtains in the competition dataset and our results of applying the metric to the evaluation dataset. Lastly, in Section IV we provide a short conclusion.

II. Methods

To quantify the localization performance of a sensor model, we generalize the approach discussed by Steckel and Peremans [2], where the mean mutual information is calculated between a pose and the measurements that can be performed at a certain pose. In other words, what does a sensor measurement tell us about our pose. We start with the mutual information between two random variables as defined by Cover [3], and use the notation from MacKay [4]:

$$I(X;Y) \equiv \sum_{xy \in \mathcal{A}_X \mathcal{A}_Y} P(x,y) \log \frac{P(x,y)}{P(x)P(y)}, \qquad (1)$$

where X and Y are ensembles which are defined as triples consisting of an outcome x and y, respectively, an alphabet $\mathcal{A}_X$ and $\mathcal{A}_Y$, and the probabilities $\mathcal{P}_X$ and $\mathcal{P}_Y$. An outcome x is a random variable that has a value from $\mathcal{A}_X$, say a_o, with a probability from $\mathcal{P}_X$, p_o. The number of elements in both $\mathcal{A}_X$ and $\mathcal{P}_X$ is O, with a_o having a probability of p_o. The function

$P(x)$ indicates the probability of the outcome x, which is p_o in this situation. In the following, we define the ensemble $P = (p, \mathcal{A}_P, \mathcal{P}_P)$ for all possible poses in the environment, assumed to be finite. A pose is defined as both a location and a heading, hence our choice for P. To avoid confusion, the function $P(x)$ will always indicate the probability of x and we will not use the shorthand p for the probability of $x = a$, but use p to indicate a pose. We also define the ensemble $M_{p_i} = (\vec{m}_{p_i}, \mathcal{A}_{M_{p_i}}, \mathcal{P}_{M_{p_i}})$ for all possible measurements at the pose p_i. We emphasize with our notation that a measurement is a vector. Inserting our ensembles into Equation (1) gives:

$$I(P; M_{p_i}) \equiv \sum_{p\vec{m}_{p_i} \in \mathcal{A}_P \mathcal{A}_{M_{p_i}}} P(p, \vec{m}_{p_i}) \log \frac{P(p, \vec{m}_{p_i})}{P(p)P(\vec{m}_{p_i})}, \quad (2)$$

where we assume that a logarithm is a base two logarithm, thus the resultant mutual information is expressed in bits. A measurement is sampled from the Gaussian distribution $\mathcal{N}$ of the actual measurements at pose p_i:

$$\vec{m}_{k,p_i} = \mathcal{N}\left(\overline{M_{p_i}^{+}}, \Sigma_{M_{p_i}^{+}}\right), \quad (3)$$

where $M_{p_i}^{+}$ are the actual sensor measurements performed at the pose p_i; $\overline{M_{p_i}^{+}}$ is the sample mean of $M_{p_i}^{+}$; and $\Sigma_{M_{p_i}^{+}}$ is the covariance matrix of $M_{p_i}^{+}$.

The sensor model, *i.e.*, the probability density function of the sensor measurement $\vec{m}$ given the pose p_j is described by:

$$P(\vec{m} \mid p_j) = \frac{1}{\sqrt{(2\pi)^{|M_{p_j}^{+}|} \det \Sigma_m}} \quad (4)$$
$$\exp\left(-\frac{1}{2}\left(\vec{m} - \overline{M_{p_j}^{+}}\right)^T \Sigma_m^{-1} \left(\vec{m} - \overline{M_{p_j}^{+}}\right)\right),$$

where p_j is an element of all possible poses P; $|M_{p_j}^{+}|$ is the dimension of the measurements in $M_{p_j}^{+}$; and Σ_m is the covariance matrix of the sensor model of the given sensor.

The posterior probability of the pose p_j given a particular sensor measurement $\vec{m}$ is calculated using Bayes:

$$P(p_j|\vec{m}) = \frac{P(\vec{m} \mid p_j)P(p_j)}{\sum_{p\in\mathcal{A}_P} P(\vec{m} \mid p)P(p)}, \quad (5)$$

where we assume a uniform distribution for $P(p_j)$ since we do not have any prior information about the pose. The marginal probability of a measurement $\vec{m}$ performed at pose p_i is not known, so we first calculate the mutual information [3], [4] between all possible poses P and the measurement $\vec{m}_{k,p_i}$:

$$I(P; \vec{m}_{k,p_i}) = \sum_{p\in\mathcal{A}_P} P(p|\vec{m}_{k,p_i}) \log \frac{P(p \mid \vec{m}_{k,p_i})}{P(p)}, \quad (6)$$

where $\vec{m}_{k,p_i}$ is a measurement sampled from the distribution of the actual measurements at pose p_i, see Equation (3).

Finally, we establish the mean mutual information by Monte Carlo approximation [5] from the mutual information between all possible poses P and the measurement $\vec{m}_{k,p_i}$ as follows:

$$< I(P; M_{p_i}) > \cong \frac{1}{K} \sum^{K} I(P; \vec{m}_{k,p_i}), \quad (7)$$

where K is the number of Monte Carlo samples. However, since the UJIIndoorLoc database provides a validation dataset, we will use this dataset to calculate the mutual information.

Our Wi-Fi sensor model is a probabilistic version of the sensor model used in [6]. We define a measurement vector containing the received signal strength (RSS) values of every access point in the environment, $\vec{m} = \vec{w}$, where $\vec{w} = \{w_1, w_2, \ldots, w_A\}$ with A the set of access points in the environment, so w_a is the RSS value of access point a in measurement $\vec{w}$. If the signal strength of an access point is very low at a certain pose, our active RFID reader or Wi-Fi scanner might not pick up the signal, and we say that we did not see the access point at that pose. The vectors are compared by access point to calculate the likelihood of pose p_j. Assuming the access points to be independently received, the sensor model becomes:

$$P(\vec{w} \mid p_j) = \prod_{a}^{A} P\left(w_a \mid p_j\right), \quad (8)$$

where $\vec{w}$ is a Wi-Fi measurement. We define four mutually exclusive events when comparing the RSS values of an access point in the set of measurements collected at p_j denoted by $W_{p_j}^{+}$ with our new measurement denoted by w: either a *hit*, a *miss*, an *extra*, or a *none*. A *hit* occurs when the access point has an RSS value both in the collection of measurements and in the new measurement. A *miss* occurs when the access point has an RSS value in the collection, but was not seen in the new measurement. An *extra* occurs when the access point has an RSS value in the new measurement, but was not seen in the collection of measurements at p_j. Lastly, a *none* occurs when the access point was not seen both in the collection and in the new measurement. The probability of w_a in (8) is defined as:

$$P(w_a \mid p_j) = \begin{cases} P(w_a, hit|p_j) & \textit{hit}, \quad (9a) \\ P(w_a, miss|p_j) & \textit{miss}, \quad (9b) \\ P(w_a, extra|p_j) & \textit{extra}, \quad (9c) \\ P(w_a, none|p_j) & \textit{none}, \quad (9d) \end{cases}$$

denoting with W_{a,p_j}^{+} the set of measurements collected at pose p_j that included access point a; w_t a threshold RSS value, under which our hardware will no longer detect an access point—if we do not detect an access point, we say that we are under this threshold. We define the four conditions as: *hit* is $W_{a,p_j}^{+} \neq \varnothing$ and $w_a > w_t$, *miss* is $W_{a,p_j}^{+} \neq \varnothing$ and $w_a < w_t$, *extra* is $W_{a,p_j}^{+} = \varnothing$ and $w_a > w_t$, and *none* is $W_{a,p_j}^{+} = \varnothing$ and $w_a < w_t$. The expressions in (9) can then be calculated as:

$$P(w_a, hit|p_j) \quad (10a)$$
$$= \frac{1}{\sqrt{2\pi}\sigma_w} \exp - \frac{\left(w_a - \overline{W_{a,p_j}^{+}}\right)^2}{2\sigma_w^2} P(W_{a,p_j}^{+} \neq \varnothing|p_j),$$

$$P(w_a, miss|p_j) \quad (10b)$$
$$= \frac{1}{\sqrt{2\pi}\sigma_w} \int_{-\infty}^{w_t} \exp - \frac{\left(x - \overline{W_{a,p_j}^{+}}\right)^2}{2\sigma_w^2} dx P(W_{a,p_j}^{+} \neq \varnothing|p_j),$$

$$P(w_a, extra|p_j) \tag{10c}$$
$$= \int_{-\infty}^{\infty} \alpha \left[\frac{1}{\sqrt{2\pi}\sigma_w} \int_{-\infty}^{w_t} \exp - \frac{\left(x - \overline{W^+_{a,p_j}}\right)^2}{2\sigma_w^2} dx \right]^N$$
$$\times \frac{1}{\sqrt{2\pi}\sigma_w} \exp - \frac{\left(w_a - \overline{W^+_{a,p_j}}\right)^2}{2\sigma_w^2} d\overline{W^+_{a,p_j}}$$
$$\times P(W^+_{a,p_j} = \varnothing | p_j),$$

$$P(w_a, none|p_j) \tag{10d}$$
$$= \int_{-\infty}^{\infty} \alpha \left[\frac{1}{\sqrt{2\pi}\sigma_w} \int_{-\infty}^{w_t} \exp - \frac{\left(x - \overline{W^+_{a,p_j}}\right)^2}{2\sigma_w^2} dx \right]^N$$
$$\times \left[\frac{1}{\sqrt{2\pi}\sigma_w} \int_{-\infty}^{w_t} \exp - \frac{\left(x - \overline{W^+_{a,p_j}}\right)^2}{2\sigma_w^2} dx \right] d\overline{W^+_{a,p_j}}$$
$$\times P(W^+_{a,p_j} = \varnothing | p_j),$$

where σ_w is the sensor model's kernel width; and $P(W^+_{a,p_j} \neq \varnothing | p_j)$ is the probability of measuring an access point at pose p_j. For Wi-Fi, the threshold w_t is set to $-89\,\mathrm{dBm}$, based on our hardware specifications. The kernel width is set to $3.46\,\mathrm{dB}$, based on [7].

In order to calculate our results in the EvAAL competition, we must estimate a location based on the posterior distribution after applying the sensor model. We assume a uniform distribution over all poses, so the posterior distribution is in fact the normalized likelihood. We calculated the result for two simple localization algorithms, MLE and KNN with $k = 4$. The MLE algorithm selects the location with the highest likelihood as its estimate. KNN selects the k locations with highest posterior probability. Then, it calculates a weighted average location based on the posterior probability values. In fact, one could say that we applied KNN twice, once with $k = 1$ and once with $k = 4$. To deal with the floors and buildings, we also calculated the weighted average floor and building and rounded these averages to whole numbers. We recognize that more complex algorithms are likely to produce better results, yet we are interested in the mean mutual information of sensor models before these algorithms are applied.

For the EvAAL competition, a performance metric is created based on spatial error. The accuracy on which the competitors will be ranked is the average sample error (SE) for each sample in the validation dataset that will be provided during the competition. The SE for a single sample is calculated as:

$$D = \sqrt{\sum_{x,y} (E - R)^2} \tag{11}$$

$$SE = D + pen_1 + pen_2 \tag{12}$$

where E is the pose estimated by the localization algorithm; R is the real pose; pen_1 is a penalty of 50, applied if the localization algorithm does not predict the building correctly; and pen_2 is a penalty of 4, applied if the localization algorithm does not predict the floor correctly.

The UJIIndoorLoc training dataset that is used in the EvAAL competition has 933 unique locations with an average of 20 samples per location. The covered area is $108\,703\,\mathrm{m}^2$, divided over four to five floors. For comparison, our own typical testing environment, CPM.E2, has 660 unique Wi-Fi sample poses, using a grid of $(0.3\,\mathrm{m}, 0.3\,\mathrm{m}, 0.25\pi\,\mathrm{rad})$, with usually one or two, but up to 32 samples per pose. The covered area is $59.4\,\mathrm{m}^2$ of office floor, divided over eight different orientations. This means that the UJIIndoorLoc training dataset has many more samples, but that its sample locations are much sparser than the sample poses in our environment. Additionally, the UJIIndoorLoc training dataset has 520 different Wi-Fi access points, of which there are 18 visible on average in a single sample. In our environment, there are 67 different Wi-Fi access points, of which 10 are visible on average in a single sample. Finally, the training dataset has 19937 samples, and the validation dataset has 1111 samples.

III. Results

We will now present the mean mutual information of our sensor model and the results of applying the MLE and KNN localization algorithms on our Wi-Fi sensor model in the UJIIndoorLoc evaluation database.

The cumulative distribution of mean mutual information provided by our sensor model in the UJIIndoorLoc database, calculated using the validation dataset, is shown in Fig. 1. The maximum amount of mutual information in the environment, $9.87\,\mathrm{bit}$, is reached in about $60\,\%$ of the samples, and an additional $10\,\%$ with a value close to the maximum. This can be seen by the steep increase at the 0.3 ratio of positions mark. At these locations, the sensor model can indicates an exact location at which it thinks the measurement to be sampled. The localization posterior will behave like a Dirac delta function. This location does not need to be correct, as can be seen in Fig. 2, which will be discussed shortly. For the bottom $20\,\%$ of samples the mean mutual information drops. The bottom 5 samples have $0\,\mathrm{bit}$ mutual information, which means that their posterior location distribution is uniform.

In our usual environment, CPM.E2, the maximum amount of mutual information is $9.03\,\mathrm{bit}$. This is reached in only very few poses. Since the environment is much more densely sampled, as discussed before, there is less signal variation between nearby poses than with the UJIIndoorLoc database. The localization posterior will decrease more smoothly around the estimated location, which causes a decrease in mutual information.

As the large difference between average and median SE in Table I suggests, the high average SE value is due to large outliers. These seem to be mainly caused by floor and building fails. The distribution of SE for MLE and KNN can also be seen in Fig. 2. At first sight, it is rather striking that using KNN but slightly improves the result, compared with MLE. However, since our measurement model results in a localization posterior that is very specific, and because we use

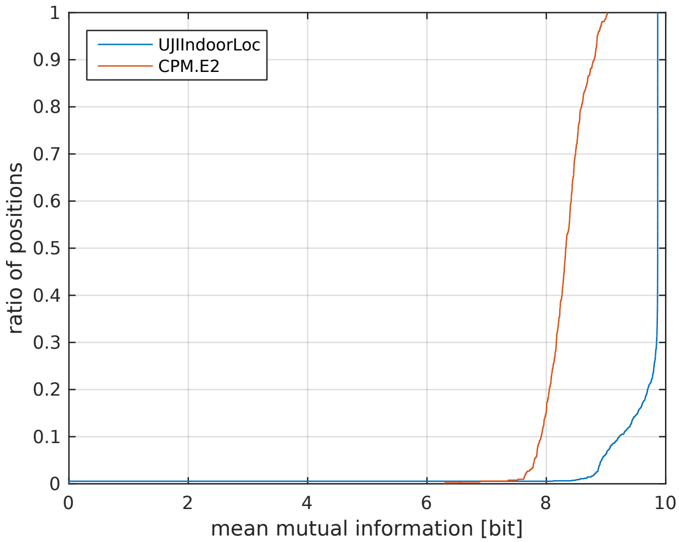

Fig. 1. The mean mutual information between a location in the UJIIndoorLoc validation database and measurements from the validation database, according to our sensor model. For comparison, we also included this result for our usual environment, CPM.E2, which are calculated using Monte Carlo approximation.

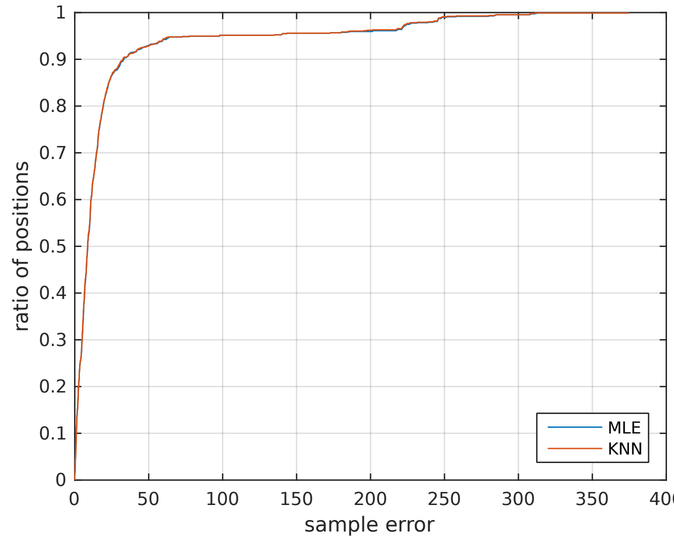

Fig. 2. Distribution of SE for the EvAAL UJIIndoorLoc validation database.

TABLE I
EVAAL LOCALIZATION RESULTS.

Property	MLE	KNN
Average SE	22.15	21.97
Median SE	8.70	8.67
Average D	19.13	18.96
Median D	8.29	8.29
Building fail %	4.86	4.86
Floor fail %	14.76	14.58

a weighted version of the KNN algorithm, the additional poses used to calculate the pose estimation are usually cancelled because of their much lower posterior value.

IV. Conclusion

We showed the mean mutual information between locations in the environment and measurements that could be performed at those locations, interpreted by our probabilistic Wi-Fi sensor model. We saw that for most samples in the UJIIndoorLoc validation dataset, the mutual information between a sample and the locations in the training database is very high. This indicates a very selective posterior distribution, where only one location has a higher location probability. As the sample error suggests otherwise, this unwarranted certainty is likely caused by our measurement model not handling the sparse samples in the UJIIndoorLoc database very well, as compared with the more densely sampled database that we built ourselves.

Accordingly, our localization algorithms do not seem to be performing with a very high accuracy in the UJIIndoorLoc database. The more sparsely sampled environment, compared to our own environment, might be better incorporated in a different localization algorithm. Such an algorithm might simulate additional samples based on the training dataset to create a better model of the environment, or might determine the building or floor on a hierarchical manner, to prevent building and floor fails. Also note that our main focus is on calculating the mean mutual information between a pose in the environment and measurements performed at that pose, using any probabilistic sensor model. It is likely that other localization algorithms could incorporate the information provided by the Wi-Fi sensor model to obtain a more accurate location estimation.

References

[1] J. Torres-Sospedra, R. Montoliu, A. Martnez-Usó, J. P. Avariento, T. J. Arnau, M. Benedito-Bordonau, and J. Huerta, "UJIIndoorLoc: A new multi-building and multi-floor database for WLAN fingerprint-based indoor localization problems," in *Proc. Fifth Int. Conf. Indoor Position. Indoor Navig.*, Busan, Korea, 2014.
[2] J. Steckel and H. Peremans, "Biomimetic sonar for biomimetic SLAM," in *Sensors, 2012 IEEE*. Taipei, Taiwan: IEEE, Oct. 2012, pp. 1–4.
[3] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York, NY: John Wiley & Sons, 1991.
[4] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
[5] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge: The MIT Press, 2006.
[6] R. Berkvens, A. Jacobson, M. Milford, H. Peremans, and M. Weyn, "Biologically Inspired SLAM Using Wi-Fi," in *Intell. Robot. Syst. (IROS), 2014 IEEE/RSJ Int. Conf.* Chicago, IL, USA: IEEE, 2014, pp. 1804–1811.
[7] Y.-S. Chiou, C.-L. Wang, S.-C. Yeh, and M.-Y. Su, "Design of an adaptive positioning system based on WiFi radio signals," *Comput. Commun.*, vol. 32, no. 7-10, pp. 1245–1254, May 2009.