Gaussian process inference approximation for indoor pedestrian localisation

J. Medvesek✉, A. Symington, A. Trost and S. Hailes

Clutter has a complex effect on radio propagation, and limits the effectiveness of deterministic methods in wireless indoor positioning. In contrast, a Gaussian process ($\mathcal{GP}$) can be used to learn the spatially correlated measurement error directly from training samples, and build a representation from which a position can be inferred. A method of exploiting $\mathcal{GP}$ inference to obtain measurement predictions from within a pose graph optimisation framework is presented. However, $\mathcal{GP}$ inference has a run-time complexity of $\mathcal{O}(N^3)$ in the number of training samples N, which precludes it from being called in each optimiser iteration. The novel contributions of this work are a method for building an approximate $\mathcal{GP}$ inference map and an $\mathcal{O}(1)$ bi-cubic interpolation strategy for sampling this map during optimisation. Using inertial, magnetic, signal strength and time-of-flight measurements between four anchors and a single mobile sensor, it is shown empirically that the presented approach leads to decimetre precision indoor pedestrian localisation.

Introduction: It is currently possible to localise pedestrians outdoors to decimetre precision using mass-market raw global satellite navigation system receivers. However, localising pedestrians indoors to a similar accuracy is not yet possible, despite it being a key requirement for many applications including augmented reality, advertising and assisted navigation. The fundamental reasons are that environmental clutter causes complex absorption, multi-path and non-line-of-sight radio errors [1], and that dead reckoning systems drift with time.

A variety of different indoor localisation systems have been described, some of which are based on multi-sensor fusion [2] and the application of supervised machine learning techniques to train localisation models [3]. Given measurements [4], many approaches use sequential Monte Carlo methods to find a likely candidate position, but these methods suffer from scaling problems when more sensors are added to the system. Since efficient pose graph methods are commonly used in visual localisation [5], we apply this type of method, in combination with GraphSLAM implementation [6], to solve a range-based localisation problem.

Our approach casts localisation as a pose graph optimisation problem, where gradient descent is used iteratively to solve for the sensor's unknown trajectory by maximising the measurement likelihood – the weighted difference between a prediction from a measurement model and an observation. Gaussian process ($\mathcal{GP}$) regression is used to build measurement models that capture the spatial correlation of magnetic field strength, as well as anchor-specific radio signal strength and time-of-flight measurements. Inertial measurements are passed through a step and heading system (SHS) to produce bearing and distance pseudo-measurements that associate sequential state values.

Performing $\mathcal{GP}$ inference in each iteration of pose graph optimisation is computationally infeasible, as it has a complexity of $\mathcal{O}(N^3)$, where N is the number of training samples. Our contribution is an approximation strategy that assumes that there is no meaningful spatial variation in measurement values beyond a fixed resolution. We can therefore pre-compute a fixed-resolution $\mathcal{GP}$ inference map once, before the pose graph optimiser is called. During optimisation, bi-cubic interpolation – with an $\mathcal{O}(1)$ complexity – is used to sample a measurement prediction by fusing the 16 nearest map cells. We show that our approach yields approximation errors of <0.001% for a 20 cm grid size.

Proposed method: Consider a localisation system that includes m static anchors $a_i = [x_i, y_i]^T$, $i = 1, 2, \ldots, m$, and one mobile sensor $p = [x, y]^T$. The target carries a 6-degree-of-freedom inertial sensor, a triaxial magnetometer and a packet radio that not only supports communication but also provides signal strength and time-of-flight measurements. The localisation problem can be represented by Fig. 1.

The objective of localisation is to estimate values for all 'unknown' sensor positions $x_1 \ldots x_n$, given the 'predicted' initial states x_0, 'unknown' anchor positions $a_1, \ldots, a_m$, 'processes' $f(\cdot)$ that relate two sequential sensor positions, and 'corrections' $g(\cdot)$ that relate a sensor's position to a subset of anchors through magnetic and radio measurements.

We use $\mathcal{GP}$ models to describe spatially correlated radio and magnetic error. Let $z = \{z_i, i = 1, \ldots, N\}$ be a set of observed measurements drawn from a noisy process and let $X = \{x_i, i = 1, \ldots, N\}$ be a set of

corresponding coordinates. Each observation z_i can be related to a transformation function $f(x_i)$ through a Gaussian noise model as $z_i = f(x_i) + \epsilon_i$, where z_i is the measurement at position x_i, and $\{\epsilon_i\}_{i=1}^N \sim N(0, \sigma_n^2)$ is the measurement noise which is assumed to be Gaussian. A $\mathcal{GP}$ model represents a function as a non-parametric distribution specified by a set of labelled training samples and a kernel function. It requires that all function values of f are jointly Gaussian with covariance matrix K. A commonly used covariance matrix is a squared exponential covariance matrix, defined as $k(x_p, x_q) = \sigma_f^2 \exp\left(-(1/2l^2)|x_p - x_q|^2\right)$, where σ_f^2 is the signal variance and l is the length scale of the covariance function. σ_f^2, l and σ_n^2 are the hyperparameters of a $\mathcal{GP}$ model. They control the smoothness of distribution and can be learned with a standard optimisation procedure by maximising the likelihood of the training data with respect to their values, using a conjugate gradient descent.
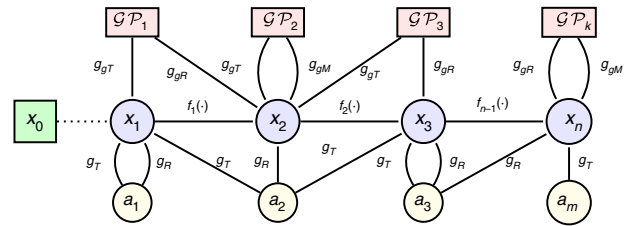

Fig. 1 *Localisation problem expressed as connected graph*
Circles x_i and a_i represent unknown sensor and landmark positions. Edges $g(\cdot)$ as corrections represent radio and magnetic measurements as well as $\mathcal{GP}$ predictions, while edges $f(\cdot)$ as processes, represent SHS pedometry. Initial states for all circles are represented as vector x_0

The main structure of our proposed localisation model is an example of a robustified, nonlinear least-squares problem of the form

$$\min_{x} \frac{1}{2} \sum_i \rho_i \left(\| f_i(x_{i_1}, \ldots, x_{i_k}) \|^2 \right) \qquad (1)$$

The expression $\rho_i(\| f_i(x_{i_1}, \ldots, x_{i_k}) \|^2)$ in (1) is known as a 'residual block', where $f_i(\cdot)$ is a 'cost function' that depends on the parameter blocks $[x_{i_1}, \ldots, x_{i_k}]$. Function ρ_i represents a 'loss function'. We have chosen the Cauchy loss function defined as $\rho_i(s) = a^2 \log(1 + (s/a^2))$, where a is a scale factor that controls influence of the loss function. The proposed loss function is a scalar function that is used to reduce the influence of outliers (e.g. multi-path) on the solution of nonlinear least squares problems. The use of a loss function is very important where many outliers are to be expected in the system, in order to prevent the solution from being pulled away from the optimal due to erroneous measurement.

Let $x \in \mathbb{R}^n$ be an n-dimensional vector of variables, and $F(x) = \begin{bmatrix} f_1(x), \ldots, f_m(x) \end{bmatrix}^\top$ be an m-dimensional function of x. Then, the general nonlinear optimisation problem that needs to be solved can be defined as $\arg \min_x (1/2) \| F(x) \|^2$. Finally, to solve the full localisation problem, we define an objective function as shown in (2)

$$\bar{X} = \arg \min_X \sum_{i \in M} \left(z_i - h_i(X) \right) R_i^{-1} \left(z_i - h_i(X) \right) \qquad (2)$$

where $X = \{x_i, i = 1, \ldots, N\}$ contains all unknown sensor positions, z_i are observations, $h_i(\cdot)$ is the measurement model and R_i is the measurement noise covariance. The processes $f(\cdot)$, as a two-dimensional displacement vector calculated from constant stride length multiplied with a bearing obtained by attitude and heading reference system, and the corrections $g(\cdot)$ are bundled into a set of measurements M. The measurement model $h_i(\cdot)$ for the ith measurement extracts the relevant states from X to 'predict' the measurement value. A residual is then calculated using the observation z_i and measurement noise covariance R_i. Corrections g represent weighted errors between predicted value $\tilde{d}$ at position x_t and the estimated value from the pre-computed map $r_{\mathcal{GP}}(t) = \sum_{i=0}^M (\tilde{d}_i - \text{map}(x_t)) \sigma_{\mathcal{GP}}^{-1}(x_t)$ for $\mathcal{GP}$ predicted magnetic and radio estimates, and $r_{RB}(t) = \sum_{i=0}^M (\tilde{d}_i - e_t) \sigma_{RB}^{-1}$ for raw radio measurements. Raw measurements are used on areas inadequately mapped with $\mathcal{GP}$, to provide a weak estimate of the absolute reference frame.

Localisation is carried out by assigning values to X in such a way that the sum of the weighted square residuals is minimised. The $\mathcal{GP}$ predictions from the radio and magnetic models are approximated from pre-computed maps, since the fundamental computational complexity of $\mathcal{GP}$ inference is $\mathcal{O}(N^3)$, which prevents $\mathcal{GP}$ predictions from being made inside of a solver

with millions of requested calls. Moreover, the solving strategy requires that x be updated using $x \leftarrow x + \Delta x$, where Δx may be an infinitesimally small quantity. We addressed this problem by performing bi-cubic interpolation (3) on the pre-computed map, defined as

$$p(x, y) = \sum_{i=1}^{4} \sum_{j=1}^{4} a_{ij} x^i y^j \qquad (3)$$

where $p(x, y)$ is the interpolated value at an arbitrary position (x, y), depending on the 16 nearest cells in the map. With this procedure, the runtime complexity of $\mathcal{GP}$ prediction decreases from $\mathcal{O}(N^3)$ to the $\mathcal{O}(1)$ operation required to interpolate a value from the pre-computed map.

Since, in most cases, the measurement model $h_i(\cdot)$ includes nonlinear functions, we have to settle finding for a local minimum. The general strategy of solving nonlinear optimisation problems is to solve a sequence of approximations to the original problem. At each iteration, the approximation is solved to determine a correction Δx to the vector x. For nonlinear least squares, an approximation can be constructed by using the first-order linearisation $F(x + \Delta x) \approx F(x) + J(x)\Delta x$. The result of this approximation is that solvers are non-optimal and may therefore converge to local minima. This effect can be mitigated by initialising the solver with a reasonable approximation of the true state values.

Experimental results: A set of experiments was conducted in a 20×20 m office space containing several large pillars and other environmental clutter. Four anchors were deployed at the corners of the space and 20 min worth of raw inertial and radio measurements were recorded by a small bespoke mobile sensor. Both the sensors and the anchors were based around an IEEE 802.15.4 compatible Freescale MC13224 microcontroller with a low-cost single chip InvenSense MPU-9150 IMU, and a nanoPAN radio module for time-of-flight measurements. The sensor was carried on the waist of a pedestrian, who was simultaneously tracked by an electronic theodolite, to obtain millimetre-accuracy ground truth.

The data set was then split into a training and a test set of 10 min each. The training set was used to generate nine (one per anchor for radio and one for magnetic measurements) $\mathcal{GP}$ approximation maps. An example of one such map is shown in Fig. 2a. To evaluate how much error is introduced by our bi-cubic interpolation approximation, we carried out exact inference and approximations for 10 000 sample points. Table 1 shows the mean value of relative error. From this data, we chose a grid step size of 0.2 m, for which the relative error drops below 0.001%.
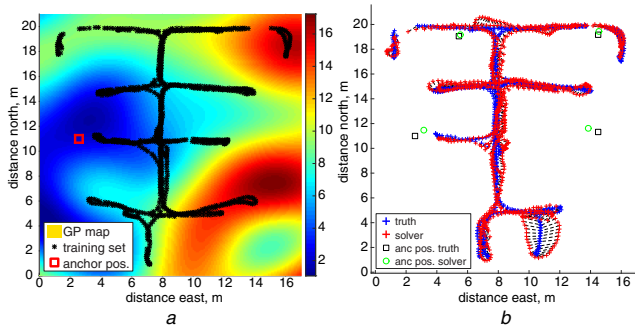

Fig. 2 *$\mathcal{GP}$ ToF map of one anchor node and localisation solver result fusing ALL sensors including $\mathcal{GP}$ over 10 min period*

a $\mathcal{GP}$ ToF map of one anchor node
b Localisation solver result fusing ALL sensors

Table 1: LUTinterpolation error

Grid step size (m)	0.05	0.1	0.2	0.5	1.0	2.0
Error – $\mathcal{GP}$ value (%)	1.7×10^{-5}	1.1×10^{-4}	7.8×10^{-4}	0.01	0.1	2.0
Error – $\mathcal{GP}$ variance (%)	4.6×10^{-5}	1.3×10^{-4}	2.9×10^{-3}	0.05	0.5	5.7

The training set was also used to estimate measurement uncertainties. Variances used in $\mathcal{GP}$ correction blocks are provided by the $\mathcal{GP}$ itself, where constant variance for range-based measurements and for process f was estimated by running the proposed solver on a test trajectory included in the simulated annealing algorithm.

From the results shown in Fig. 3, where we fused together different measurement sources with the main process f, it is evident that the overall localisation accuracy is significantly improved, where $\mathcal{GP}$ predictions were included into the localisation solver.

Since it may not be possible to build $\mathcal{GP}$ maps for the entire area of an environment, it is important to include raw (or filtered) range measurements to prevent the solution from drifting away at unmapped areas. Note that we used the same initial state x_0 for all tests, to exclude preliminary errors due to different initial conditions. Initial states were predicted by a multi-level algorithm including maximising measurement likelihood of range measurements. In obtaining the full solver solution on a 10 min test set that contains 20 000 unknown states, it takes 2 s to predict initial states and 4 s to solve the localisation problem, giving a 0.5 m mean localisation error. This is shown in Fig. 2b.
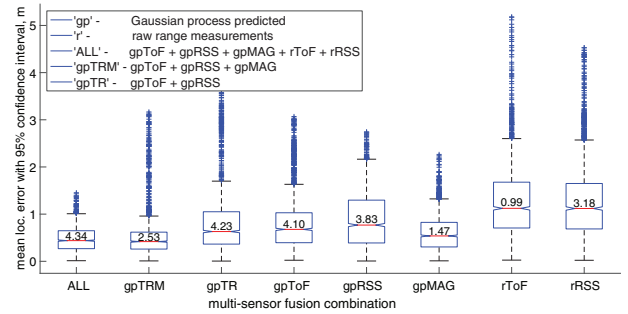

Fig. 3 *Mean localisation error for different multi-sensor fusion combinations (values inside boxes represent solving time in seconds)*

Conclusion: In this Letter, we have shown how $\mathcal{GP}$ measurement models can be integrated with pose graph optimisation in order to localise pedestrians indoors. Notably, we present a novel bi-cubic interpolation method for inference approximation, which enables solutions to be calculated in reasonable time. We also show empirically that this approximation does not add significant error, and that our system enables decimetre tracking from inertial, radio and magnetic measurements obtained from mass-market sensors.

One or more of the Figures in this Letter are available in colour online.

J. Medvesek, A. Symington and S. Hailes (*Department of Computer Science, University College London, London, United Kingdom*)

✉ E-mail: jan.medvesek@gmail.com

A. Trost (*Department of Electronics, Faculty of Electrical Engineering, Ljubljana, Slovenia*)

References

1 Rohrig, C., and Muller, M.: 'Indoor location tracking in non-line-of-sight environments using a IEEE 802.15. 4a wireless network'. IEEE Intelligent Robots and Systems Conf., St. Louis, USA, October 2009
2 DSouza, M., Wark, T., Karunanithi, M., and Ros, M.: 'Evaluation of real-time people tracking for indoor environments using ubiquitous motion sensors and limited wireless network infrastructure', *Pervasive Mob. Comput.*, 2013, **9**, pp. 498–515
3 Yen, H., and Wang, C.: 'Adapting Gaussian processes for cross-device Wi-Fi localization'. IEEE Indoor Positioning and Indoor Navigation Conf., Montbeliard-Belfort, France, October 2013
4 Galov, A., and Moschevikin, A.: 'Bayesian filters for ToF and RSS measurements for indoor positioning of a mobile object'. IEEE Indoor Positioning and Indoor Navigation Conf., Montbeliard-Belfort, France, October 2013
5 McDonald, J., Kaess, M., Cadena, C., Neira, J., and Leonard, J.: 'Real-time 6-DOF multi-session visual SLAM over large-scale environments', *Robot. Auton. Syst.*, 2013, **61**, pp. 1144–1158
6 Huang, J., and Millman, D.: 'Efficient, generalized indoor WiFi GraphSLAM'. IEEE Int. Conf. on Robotics and Automation, Shanghai, China, May 2011