

Continuous Vehicle Localisation Using Sparse 3D Sensing, Kernelised Rényi Distance and Fast Gauss Transforms

Mark Sheehan, Alastair Harrison and Paul Newman

Abstract—This paper is about estimating a smooth, continuous-time trajectory of a vehicle relative to a prior 3D laser map. We pose the estimation problem as that of finding a sequence of Catmull-Rom splines which optimise the Kernelised Rényi Distance (KRD) between the prior map and live measurements from a 3D laser sensor. Our approach treats the laser measurements as a continual stream of data from a smoothly moving vehicle. We side-step entirely the segmentation and feature matching problems incumbent in traditional point cloud matching algorithms, relying instead on a smooth and well behaved objective function. Importantly our approach admits the exploitation of sensors with modest sampling rates - sensors that take seconds to densely sample the workspace. We show how by appropriate use of the Improved Fast Gauss Transform we can reduce the order of the estimation problem from quadratic (straight forward application of the KRD) to linear. Although in this paper we use 3D laser, our approach is also applicable to vehicles using 2D laser sensing or dense stereo. We demonstrate and evaluate the performance of our approach when estimating the full 6DOF continuous time pose of a road vehicle undertaking over 2.7km of outdoor travel.

I. INTRODUCTION

This paper addresses the problem of continuous time vehicle localisation against a 3D prior map. We envision a system where a single expensive one-shot survey of the environment can be created and subsequently leveraged and/or contributed to by many low-cost vehicles.

While there exists a substantial body of work with regards to the exploitation of a prior survey in the laser-based localisation literature, a common assumption is that of a “rigid point set”: many laser measurements are agglomerated into rigid chunks and treated as atomic entities which undergo euclidean transformations relative to each other. An example of this is treating a single rotation of the rapidly rotating Velodyne laser scanner as an atomic unit. Localisation or ego-motion estimation then becomes one of matching rigid point clouds to each other - a well understood task. This is not wildly problematic if vehicle motion is slow compared to the measurement rate - so the sensor is to a first approximation able to capture the workspace “instantaneously” - but things fall apart when this is not the case. When the vehicle is itself undergoing rapid unknown motion it is not possible to create a trustworthy and rigid workspace scan. In reality measurements are taken one by one as the vehicle undertakes an (unknown) continuous trajectory. Our formulation represents the vehicle’s trajectory as a Catmull-Rom spline with each and every laser measurement hung

The authors are with the Oxford University Mobile Robotics Group, UK.
{mcs,arh,pnewman}@robots.ox.ac.uk

off this space curve. Accordingly if we smoothly perturb the spline parameters, the back projected 3D location of all laser measurements is also smoothly perturbed. One of the contributions of this work is to suggest a continuous cost function free from any data association or feature matching complexities which can be used to optimise the trajectory parameters. This results in a necessarily smooth trajectory. We go on to show how the nature of this cost function can be exploited by a variant of the Fast Gauss Transform to reduce the complexity of the cost function from quadratic to linear.




Fig. 1. A typical survey point cloud, coloured by laser reflectance value, with vehicle poses from the GPS/INS plotted as coloured axes. We wish to deduce a continuous time representation of vehicle motion relative to this map using a slow-sampling sensor.

II. RELATED WORK

There exists a vast and rich body of related work in laser-based localisation. In particular, 2D scan matching against prior maps has a long and successful record of accomplishment in robotics. However, operation in complex, dynamic environments often necessitates the use of an environment representation far richer than a 2D scan. In [1] for example, 2D scan data is registered against a 3D survey map in order to overcome scan artefacts due to vehicle roll and pitch, ground strike and grazing effects. In [2], 3D laser and camera data are combined and feature descriptors generated and matched on a scan-by-scan basis. The authors employ Mutual Information as a “metric” and show that it offers a superior basin of convergence compared to ICP. In [3] a 3D




Fig. 2. System overview - 3D point cloud maps are created of the environment, the points in these maps are denoted by blue stars. Our robotic platform traverses the same environment again using only the 3D laser scanner, providing us only with range measurements (denoted by the green dashed lines) drawn from the same underlying environment structure. There is only one continuous trajectory which causes our live laser points to align with our survey point cloud data. This true trajectory is discovered by considering hypothesis trajectories (one of which is shown in red) which alters where the live laser points are projected, we bend the trajectory according to our cost function until our live data maximally matches our survey. Our cost function captures the alignment between our survey point cloud and the laser points projected by a hypothesis trajectory. Our cost function does not use explicit correspondences between point sets, this is highlighted by the blue dashed line, each live laser point is compared (in an efficient manner) to every single survey point allowing every laser and every survey point to contribute to the overall cost function.

nodding laser scanner is used, with gradient based features extracted from 2D slices and matched between consecutive scans, providing a localisation prior for a more complex 6D SLAM system. Bibby and Reid [4] represent vehicle trajectories using cubic splines enabling the use of a SLAM system for a dynamic nautical environment using radar data. Of particular relevance is the thread of work beginning with [5] where consecutive 3D scans are gathered and processed to extract 3D shape features which are paired using smoothness and match constraints provided by a linear system. This work is extended in [6] where continuous trajectories are created via a spring mounted laser with an IMU. Continuing the 3D theme, an offline SLAM algorithm is employed on data from a Velodyne laser scanner and a GPS/IMU to build probabilistic maps in [7]. Localisation is subsequently performed against this map using a particle filter. In [8] an ICP based approach is taken to discrete vehicle pose estimation. This problem is slightly different to our own, as they are using a Velodyne HDL-64E, which gathers data at a faster rate than our own sensor, making the motion between point acquisitions less critical. In a recent work, Tong, Furgale and Barfoot [9] have been looking at the continuous time state estimation problem using suitably chosen basis functions. This work is complementary to ours in that we choose a very simple trajectory model and focus instead on leveraging a continuous, association-free measurement model.

III. METHODOLOGY

A. Survey generation

At the outset a static survey map is constructed using GPS/INS data and measurements from a 3D laser scanner, which is comprised of three SICK LMS151 scanners configured to have a full 360° field of view¹. The laser is mounted on a vehicle and driven around a looping road. We take great care in synchronising the GPS/INS and 3D laser data streams using the TICSync algorithm [11]. We fit a Catmull-Rom spline (the details of which will be discussed in Section III-B) to the GPS/INS data. This provides a continuous representation of the vehicle trajectory which is parametrised by time. Every single laser measurement has its own time stamp. This allows us to identify a unique 6DOF pose of the vehicle for each measurement. From here it is a simple matter to back project the laser measurement to produce a 3D point in the world frame. Doing this for all measurements leads to the survey point cloud, Ω , containing N points. An example survey point cloud is shown in Figure 1 with the GPS/INS trajectory overlaid.

B. Trajectory modelling

We treat the vehicle motion as a continuous function $S(t)$, parametrised by time, allowing us to obtain a 6DOF vehicle pose for a laser point at any given time, t . Our pose

¹A more detailed account of the 3D laser scanner, including the calibration method used, can be found in [10].

representation parametrises vehicle attitude using an Euler angle representation, $[r, p, q]^T$, corresponding to roll, pitch and yaw respectively. A trajectory $\mathbf{S}(t)$ is thus composed of six spline paths – one for each element of the pose vector – such that $\mathbf{S}(t) = [S_x(t), S_y(t), S_z(t), S_r(t), S_p(t), S_q(t)]^T$.

To represent the trajectory $\mathbf{S}(t)$, we choose a piecewise cubic spline, the coefficients of which are generated using the Catmull-Rom method [12]. Catmull-Rom splines have a compact support region making them easy to compute, as well as providing C^1 continuity across control points. A Catmull-Rom spline is a special case of a Cardinal spline, which is formed using four Hermite interpolation basis functions:

$$\begin{aligned} h_1(s) &= (2s^3 - 3s^2 + 1) \\ h_2(s) &= (s^3 - 2s^2 + s) \\ h_3(s) &= (-2s^3 + 3s^2) \\ h_4(s) &= (s^3 - s^2) \end{aligned} \quad (1)$$

Given a pair of control points p_{k-1}, p_k with corresponding gradients m_{k-1}, m_k and an interpolation parameter $s \in [0, 1]$, we write the spline equation as

$$x(s) = h_1(s)p_{k-1} + h_2(s)m_{k-1} + h_3(s)p_k + h_4(s)m_k \quad (2)$$

A Cardinal spline is formed if we further specify that the gradients should be computed as a function of neighbouring control point positions,

$$m_k = (1 - c) \frac{p_{k+1} - p_{k-1}}{t_{k+1} - t_{k-1}} \quad (3)$$

where $c \in [0, 1]$ is known as the tension parameter, with $c = 0$ being the value for a Catmull-Rom spline. Since the gradients now depend on neighbouring control points, a Catmull-Rom spline segment therefore has a support of 4 control points.

We define a vehicle pose as a tuple $\mathbf{P}_i = ([x, y, z, r, p, q], [t])$. We seek a continuous trajectory in the form of a Catmull-Rom spline which is uniquely determined by four vehicle poses which are themselves control points for the spline. We shall invoke an iterative and recursive strategy. At time step k we shall calculate an updated trajectory $S(t)$ by optimising over a new control point \mathbf{P}_k in isolation. Given $S(t)$, a stream of laser points $L_i = \{x_i, y_i, z_i\}$ where each point i has a time stamp t_i , we can obtain an estimate of vehicle pose for every measured laser point. By application of a kinematic chain this yields a back projected point cloud \mathbf{L} in the world frame. The following sections will set out in detail how this optimisation proceeds. Before progressing we should be explicit about why this paper is not about ICP. Firstly we use a continuum based approach using a slow sensor, significant motion occurs between individual laser measurements. Secondly our approach is not limited by point to point correspondences, we know in reality that laser data is sampled from surfaces. Matching explicitly to other measurements is not sensible.

C. The cost function

We wish to find the vehicle trajectory which maximally aligns our live point stream with our survey map. Quantifying the quality of this match allows us to form an optimisation problem yielding the current vehicle trajectory. The method used to quantify the overlap between the deformable \mathbf{L} and the rigid survey $\mathbf{\Omega}$ is an entropy measure called the Kernelised Rényi Distance (KRD) [13]. Entropy expresses the amount of disorder in a system, therefore we expect our system's lowest entropy to correspond to a maximal overlap between the survey map and the live laser stream, yielding the transform of the current vehicle location relative to the prior map.

We now outline the derivation of the Kernelised Rényi Distance from an information theoretic quantity known as Rényi Quadratic Entropy. At this point we draw on our previous work [14] to recapitulate the provenance of this quantity.

Let us begin with a model that all measurements - both at survey and runtime

$$\hat{\mathbf{X}} = \{\mathbf{\Omega}, \mathbf{L}\} \quad (4)$$

are samples of a latent distribution, $p(\mathbf{x})$ which represents the chance of a laser measurement originating from a 3D location, \mathbf{x} . We approximate $p(\mathbf{x})$ with Parzen Window [15] density estimation. With a Gaussian centred on each data point, we can use a Gaussian Mixture Model (GMM) to model this distribution

$$p(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N G(\mathbf{x} - \hat{\mathbf{x}}_i, \sigma^2 \mathbf{I}), \quad (5)$$

where $G(\mu, \Sigma)$ is a Gaussian with mean μ and isotropic covariance $\Sigma = \sigma^2 \mathbf{I}$.

Rényi [16] proposed an entropy measure which has a closed-form solution for Gaussian Mixture Models [17]. While we use this to quantify the compactness of a point cloud, it has also found application in the field of point cloud registration [18]). The Rényi entropy H_R of a random variable X with pdf $p(x)$ is:

$$H_R[X] = \frac{1}{1 - \alpha} \log \int p(x)^\alpha dx \quad \alpha > 0, \alpha \neq 1 \quad (6)$$

where α determines the weighting of events: high values of α approaching infinity consider only the highest probability events, whereas lower values of α weight high and low probability events more equally. For the case where $\alpha \rightarrow 1$, Equation 6 becomes the familiar Shannon Entropy measure [19]. For $\alpha = 2$ we obtain:

$$H_{RQE}[X] = -\log \int p(x)^2 dx, \quad (7)$$

which is known as the Rényi Quadratic Entropy (RQE).

Substituting Equation 5 into Equation 7 yields:

$$\begin{aligned}
H_{RQE}[\hat{\mathbf{X}}] &= -\log \int \left(\frac{1}{N} \sum_{i=1}^N G(\mathbf{x} - \hat{\mathbf{x}}_i, \sigma^2 \mathbf{I}) \right)^2 d\mathbf{x} \\
&= -\log \left(\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \int G(\mathbf{x} - \hat{\mathbf{x}}_i, \sigma^2 \mathbf{I}) \right. \\
&\quad \left. G(\mathbf{x} - \hat{\mathbf{x}}_j, \sigma^2 \mathbf{I}) d\mathbf{x} \right) \quad (8)
\end{aligned}$$

which can be simplified using,

$$\int G(x - x_i, \Sigma_1) G(x - x_j, \Sigma_2) dx = G(x_i - x_j, \Sigma_1 + \Sigma_2), \quad (9)$$

leading to a closed-form expression for the Rényi Quadratic Entropy of the mixture model,

$$H_{RQE}[\hat{\mathbf{X}}] = -\log \left(\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N G(\hat{\mathbf{x}}_i - \hat{\mathbf{x}}_j, 2\sigma^2 \mathbf{I}) \right). \quad (10)$$

This is a measure of the compactness of the points in \mathbf{X} , dependent on only the parameter σ , and the pairwise distances between points in $\hat{\mathbf{X}}$.

The Kernelised Rényi Distance [13], can be obtained from Equation 10 if $\hat{\mathbf{X}}$ is treated as a partition of two \mathbb{R}^3 point sets, in our case this partition is defined in Equation 4. Considering only the cross components of the entropy contribution we obtain our cost function as:

$$E(\mathbf{L}, \Omega) = -\log \left(\frac{1}{NM} \sum_j^M \sum_i^N \frac{e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}}}{8\pi^{\frac{3}{2}} \sigma^3} \right), \quad (11)$$

where $\mathbf{d} = L_j - \Omega_i$, the vector between the i^{th} and j^{th} points in Ω and \mathbf{L} , and N and M denote the number of points in Ω and \mathbf{L} , respectively.

As we seek to optimise $E(\mathbf{L}, \Omega)$ we drop the constant scaling factors and log, as it is a monotonic operator simplifying our cost function to:

$$E(\mathbf{L}, \Omega) = -\sum_j^M \sum_i^N e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}}, \quad (12)$$

which depends only on pairwise distances between the survey and laser point stream.

Equation 11 will return the lowest entropy when \mathbf{L} and Ω are maximally aligned. Because we have constructed the problem such that \mathbf{L} is a function of the current vehicle pose, $\hat{\mathbf{P}}_k$, we can express the problem as a minimisation:

$$\hat{\mathbf{P}}_k = \arg \min_{\hat{\mathbf{P}}_k} E(\mathbf{L}, \Omega) \quad (13)$$

providing us with our best estimate of the current vehicle pose $\hat{\mathbf{P}}_k$ and by construction the vehicle trajectory.

IV. EFFICIENCY VIA THE FAST GAUSS TRANSFORM

Naive calculation of Equation 11 can be achieved in $O(NM)$ time. However, this can be reduced to $O(N + M)$ time, using a method known as the Improved Fast Gauss Transform (IFGT) [20], [21], [22]. We are working with two sets of points in \mathbb{R}^3 , $\Omega_1, \dots, \Omega_N \in \Omega$ and $L_1, \dots, L_M \in \mathbf{L}$. For consistency with the literature we refer to Ω as the source points and \mathbf{L} as the target points, where the source points have associated weightings $q_1, \dots, q_N \in \mathbf{Q}$. For these the IFGT efficiently evaluates equations of the form

$$g(L_j) = \sum_{i=1}^N q_i e^{-\frac{\|\Omega_i - L_j\|^2}{h^2}}, \quad (14)$$

where the system bandwidth is h . We note that our cost function of Equation 12 may be written as a summation of terms in this form, with $h = 2\sigma$ and $q_i = 1, \forall i$.

The IFGT is guaranteed to satisfy the error bound $\frac{\|\hat{g}(L_j) - g(L_j)\|}{\sum_{i=1}^N q_i} \leq \epsilon$ for a specified ϵ . The decrease in computational complexity is achieved through the combination of a truncated Taylor series factorisation and a space subdivision scheme, which allows insignificant contributions to be ignored.

The exponential component of Equation 14 can be expressed as a factorisation around any arbitrary point in space, c , as follows:

$$\begin{aligned}
g(L_j) &= \sum_{i=1}^N q_i e^{-\frac{\|\Omega_i - L_j\|^2}{h^2}}, \\
&= \sum_{i=1}^N q_i e^{-\frac{\|(\Omega_i - c) - (L_j - c)\|^2}{h^2}}, \\
&= \sum_{i=1}^N q_i e^{-\frac{\|L_j - c\|^2}{h^2}} e^{-\frac{\|\Omega_i - c\|^2}{h^2}} e^{\frac{2(\Omega_i - c) \cdot (L_j - c)}{h^2}} \quad (15)
\end{aligned}$$

This results in three exponential terms. The first and second exponentials depend only on points in sets Ω and \mathbf{L} respectively. The separation of Ω and \mathbf{L} in the third exponential can be achieved through the use of a Taylor series expansion:

$$e^{\frac{2(\Omega_i - c) \cdot (L_j - c)}{h^2}} = \sum_{n=0}^{p-1} \frac{2^n}{n!} \left[\frac{(\Omega_i - c)}{h} \right]^n \left[\frac{(L_j - c)}{h} \right]^n + \text{err}(p) \quad (16)$$

Where p is the number of terms used in the Taylor expansion. This allows us to rewrite Equation 15 in terms of a set of coefficients which can be evaluated independently of points in \mathbf{L} by evaluating Ω around c :

$$C_\alpha = \sum_{n=0}^{p-1} \sum_{i=1}^N q_i \frac{2^n}{n!} \left[\frac{\|\Omega_i - c\|^2}{h} \right]^n \quad (17)$$

yielding the full equation:

$$g(L_j) = \sum_{n=0}^{p-1} C_\alpha e^{-\|L_j - c\|^2} \left[\frac{\|L_j - c\|^2}{h} \right]^n \quad (18)$$

This formulation allows points Ω in a region around c to contribute to a single Taylor Series expansion, which can later be evaluated for points \mathbf{L} around the same cluster centre c . This results in the approach having linear complexity with respect to the number of points in \mathbf{L} and Ω rather than quadratic.

On a housekeeping note - at low bandwidths efficient clusters cannot be formed without expanding the Taylor series to a large number of terms. In this scenario, the problem could be addressed simply by using a KD-Tree [23] to search the space, with a radius proportional to a combination of the bandwidth and ϵ . A detailed analysis of the trade-off between these two methods is provided in [20].

In our work we use the IFGT and KD-Tree implementations provided by [24]. The toolbox automatically selects the best method by efficiently estimating the approximate number of floating point operations required based on the data and the desired bandwidth.

A. Jacobian formulation

With knowledge that we can reduce our system from quadratic to linear complexity, we demonstrate both how we compute the Jacobian of Equation 13 and how we make the calculation computationally tractable. Let $\mathbf{T}(t)$ be the transformation mapping a laser measurement at time t , when the vehicle is at pose Γ , to \mathbf{L}^w in the world frame. We seek the variation of Equation 12 with respect to the control point / current vehicle pose \mathbf{P}_k . This is readily accessed via application of the chain rule.

$$\frac{dE}{d\mathbf{P}_k} = \frac{dE}{d\mathbf{L}^w} \frac{d\mathbf{L}^w}{d\mathbf{T}} \frac{d\mathbf{T}}{d\Gamma} \frac{d\Gamma}{d\mathbf{S}} \frac{d\mathbf{S}}{d\mathbf{P}_k} \quad (19)$$

The last four of these terms can be evaluated easily by simple differentiation. The first term however, $\frac{dE}{d\mathbf{L}^w}$, deserves closer attention.

Remembering our cost function as:

$$E(\mathbf{L}, \Omega) = - \sum_j^M \sum_i^N e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}} \quad (20)$$

We differentiate with respect to the x, y, z components separately. But only the differentiation for the x component is discussed here as the processes for y and z are similar. We use subscript notation L_{xi} and Ω_{xj} to denote the x component of laser points L_i and Ω_j respectively:

$$\frac{dE(\mathbf{L}, \Omega)}{dL_x} = - \sum_j^M \sum_i^N \frac{(L_{xi} - \Omega_{xj})e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}}}{2\sigma^2} \quad (21)$$

We can split this result into two summations in the form

of Equation 14

$$\frac{dE(\mathbf{L}, \Omega)}{dL_x} = \frac{1}{2\sigma^2} \left(\sum_j^M \sum_i^N L_{xi} e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}} - \sum_j^M \sum_i^N \Omega_{xj} e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}} \right) \quad (22)$$

and swapping the ordering of the summations for the first part of the Equation:

$$\frac{dE(\mathbf{L}, \Omega)}{dL_x} = \frac{1}{2\sigma^2} \left(\sum_i^N L_{xi} \sum_j^M e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}} - \sum_j^M \Omega_{xj} \sum_i^N e^{-\frac{(\mathbf{d}^T \mathbf{d})}{4\sigma^2}} \right) \quad (23)$$

each summation maps to the form of Equation 14 again using $h = \sqrt{2}\sigma$ and $\mathbf{Q} = \mathbf{1}$ as:

$$\frac{dE(\mathbf{L}, \Omega)}{dL_x} = \frac{1}{2\sigma^2} \left(\sum_i^N L_{xi} g(\Omega_i) - \sum_j^M \Omega_{xj} g(L_j) \right) \quad (24)$$

This is the most computationally expensive part of the Jacobian, condensed to two IFGT calculations. Which have been numerically verified.

V. RESULTS




Fig. 3. An overhead satellite image of the Begbroke site overlaid with vehicle trajectories obtained from our localisation system for four loops of the site. The results from our localisation system are shown in yellow, while those from the GPS/INS data are shown in Red and are scaled according to the uncertainty values provided by the GPS/INS system. GPS data is particularly bad for the bottom left section of the image due to tree cover. In this region we see that our systems trajectory remains consistent over multiple runs. On the final traversal, where the survey vehicle was driven on the opposite side of the road (where possible) our localisation system remains in agreement with the GPS.

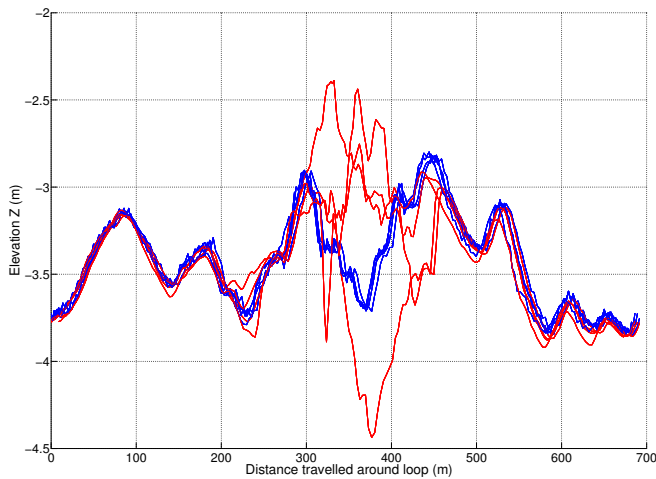


Fig. 4. Raw elevation data in metres plotted for multiple passes of the same region for both the GPS/INS (red) and for our system (blue). Indexed by the distance the vehicle has travelled. We see that two systems generally agree with the exception of the region from 200M to 450M around the loop, which corresponds to the tree covered region of Figure 3, where the line of sight to the GPS satellites is blocked, for this region our system provides a consistent elevation estimate for multiple passes whereas the GPS/INS system provides elevations different by up to 2 Metres for the same region.

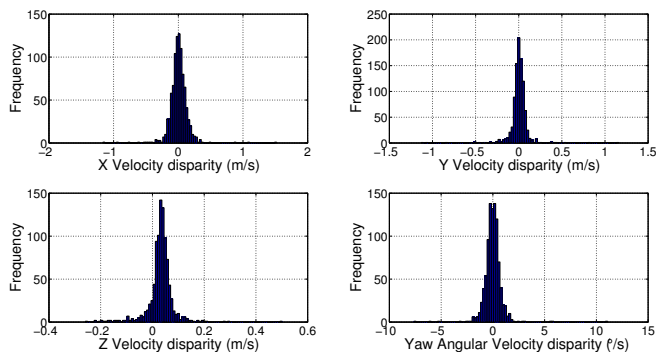


Fig. 5. Histograms showing the velocity disparities between the GPS/INS and our system of Figure 6 for X,Y,Z and Yaw. All of the histograms have a low variance around zero velocity disparity with the exception of the Z histogram. Part of the reason the variance isn't lower, is because the GPS signal suffers for a region of our route, this is also likely the reason for the slight offset in the Z velocity disparity see Figure 4 for further details.

This section provides a ground truth evaluation of our continuum method over 2.7km of data.

To measure the performance of our system, we first performed a laser survey of our test site, with vehicle poses provided by a GPS/INS system with differential corrections from a local base station. We then drove four more laps of the site (amounting to just over 2.7km), using our laser based localisation system to estimate the trajectory of the vehicle. On the fourth lap, the car was deliberately driven on the wrong side of the road, to ensure that the vehicle trajectory was offset from the survey trajectory. During these runs, we also recorded GPS/INS data to provide a ground truth trajectory. The test route has one area which is significantly covered by trees, which can cause problems with lost signals and multipath returns both from the GPS

satellites and the differential base station. For that reason, the GPS/INS trajectory unfortunately cannot be regarded as 'ground truth' at all times (and is thus one of the major motivations behind the work in this paper). In the absence of higher quality localisation information, we provide results based on the disparity between the trajectory reported by our own system and that reported by the GPS/INS. Figure 3 shows the trajectories driven, along with the GPS uncertainty around the site. Figure 4 demonstrates the inconsistency of the GPS/INS result in the problem area. Note that while the GPS elevation trace shows considerable variance over multiple runs, our laser-based localisation maintains low variance throughout.

Figure 6 shows the x,y,z velocity disparity (in the vehicle frame) between our system and the GPS/INS over the course of the four loops. Apart from the four short segments when the vehicle was in the poor signal area, the disparities remain small, with a zero mean distribution, showing that the recovered trajectory remains closely tied to the GPS trajectory.

The difference between the localisation in X,Y,Z and Yaw obtained by our system, and that obtained by the GPS/INS system is shown in Figure 7 for a single loop, for which the GPS/INS system reported the lowest uncertainty. These results exclude the tree covered region where the GPS/INS system reported high uncertainty in its measurements. We see the disagreement between the two systems is in the order of centimetres, a large amount of this disagreement is likely due to inaccuracies in the GPS/INS measurements.

VI. CONCLUSIONS

This paper has addressed the problem of performing vehicle localisation using a continuous trajectory with a Kernelised Rényi Distance based cost function. The goal of this work was to operate on a continuum (so scan matching rigid chunks of laser data is out of the question) and to estimate the pose of the vehicle not as a discrete set of poses but as a continuous smooth function. Moreover we wished to operate with laser sample densities far below that suitable for the application of ICP-like matching and crucially, to do so without needing to make explicit point to point associations.

Our results indicate this is a good method to match laser data streams as the estimated trajectory closely matches the ground truth INS/GPS data with mean absolute velocity disparities of less than 8 cm/s in x, 5.5 cm/s in y and 4.5 cm/s in z over a 2743 metre dataset. The increased calculation speed provided by the IFGT has allowed us to calculate these results at a speed just behind real time. The use of a continuous and smooth motion model is mandated by the use of a sensor which cannot acquire snap-shots of the scene and one for which vehicle motion while sensing must be modelled. In the context of laser based navigation, it remains interesting to understand what can be done in the absence of a Velodyne. What we propose here is applicable to high and low bandwidth sensing alike.




Fig. 6. Plots showing the velocity disparity between our system and the GPS/INS system. We see that the velocity disparities are of a very small scale with the exception of four parts, each of which corresponds to the region where the GPS signal is weakened by foliage cover. The histograms of these results are shown in Figure 5 including all outliers.




Fig. 7. Box plots, showing the difference between the trajectory obtained from our localisation system and pose values from the GPS/INS system for X, Y, Z and Yaw in the vehicles reference frame. The top and the bottom of each box represents the 25th and 75th percentiles of the data, while the line inside the box represents the median of the data, the tails of the box show the range of the data excluding outliers, which are displayed as red crosses. This plot is generated for a single loop of the science park. Regions where the GPS/INS reported high uncertainty, corresponding to the tree covered region have been excluded from these results.

VII. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of Guidance Ltd. This work has been generously supported by Guidance Ltd. As well as the Office of Naval Research grant N00014-08-1-0337 under Bezaad Kamgar-Parsi. Paul Newman was supported by an EPSRC Leadership Fellowship, EPSRC Grant EP/I005021/1. We thank Rohan Paul for helpful conversations and give a special thank you to Dr Benjamin Davis for his insights, helping with data collection and for maintaining the research platform so well.

Appendix

Symbol	Meaning
Ω	The global point cloud survey
L	Point data gathered from the laser scanner alone
N	The number of survey point cloud measurements Ω
M	The number of live laser measurements L
E	Kernelised Rényi Distance based cost function
σ	The size of the isotropic Gaussian window
P	Vehicle poses
S	Cubic spline coefficients
t	Time

REFERENCES

- [1] I. Baldwin and P. Newman, "Road vehicle localization with 2D push-broom lidar and 3D priors," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2012.
- [2] G. Pandey, J. R. McBride, S. Savarese, and R. M. Eustice, "Toward mutual information based automatic registration of 3D point clouds," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, (Vilamoura, Algarve [Portugal]), October 2012. Accepted, To Appear.
- [3] K. Lingemann, A. Nüchter, J. Hertzberg, and H. Surmann, "High-speed laser localization for mobile robots," *Robotics and Autonomous Systems*, vol. 51, no. 4, pp. 275–296, 2005.
- [4] C. Bibby and I. Reid, "A hybrid slam representation for dynamic marine environments," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 257–264, IEEE, 2010.
- [5] M. Bosse and R. Zlot, "Continuous 3D scan-matching with a spinning 2D laser," in *ICRA'09. IEEE International Conference on Robotics and Automation*, pp. 4312–4319, IEEE, 2009.
- [6] M. Bosse, R. Zlot, and P. Flick, "Zebedee: Design of a spring-mounted 3-d range sensor with application to mobile mapping," *Robotics, IEEE Transactions on*, vol. PP, no. 99, pp. 1–16, 2012.
- [7] J. Levinson and S. Thrun, "Robust vehicle localization in urban environments using probabilistic maps," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 4372–4378, IEEE, 2010.
- [8] F. Moosmann and C. Stiller, "Velodyne slam," in *Intelligent Vehicles Symposium (IV), 2011 IEEE*, pp. 393–398, IEEE, 2011.
- [9] C. Tong, P. Furgale, and T. Barfoot, "Gaussian process gauss-newton: Non-parametric state estimation," in *Ninth Conference on Computer and Robot Vision*, pp. 206–213, IEEE, 2012.
- [10] M. Sheehan, A. Harrison, and P. Newman, "Automatic self-calibration of a full field-of-view 3D n-laser scanner," in *In Proceedings of the International Symposium on Experimental Robotics (ISER2010)*, (New Delhi and Agra, India), December 2010.
- [11] A. Harrison and P. Newman, "Ticsync: Knowing when things happened," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 356–363, IEEE, 2011.
- [12] E. Catmull and R. Rom, "A class of local interpolating splines," *Computer aided geometric design*, vol. 74, pp. 317–326, 1974.
- [13] B. Srinivasan and R. Duraiswami, "Efficient subset selection via the kernelized rényi distance," in *Computer Vision, 2009 IEEE 12th International Conference on*, pp. 1081–1088, IEEE, 2009.
- [14] M. Sheehan, A. Harrison, and P. Newman, "Self-calibration for a 3D laser," *The International Journal of Robotics Research*, vol. 31, pp. 675–687, 2012.
- [15] E. Parzen, "On estimation of a probability density function and mode," *The annals of mathematical statistics*, vol. 33, pp. 1065–1076, Jan 1962.
- [16] A. Rényi, "On measures of entropy and information," in *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pp. 547–561, 1961.
- [17] J. Principe and D. Xu, "Learning from examples with rényi's information criterion," in *Signals, Systems, and Computers, 1999. Conference Record of the Thirty-Third Asilomar Conference on*, vol. 2, pp. 966–970 vol.2, 1999.
- [18] Y. Tsin and T. Kanade, "A correlation-based approach to robust point set registration," in *Computer Vision-ECCV 2004*, pp. 558–569, Springer, 2004.
- [19] C. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, July,October 1948.
- [20] V. Raykar, C. Yang, R. Duraiswami, and N. Gumerov, "Fast computation of sums of gaussians in high dimensions," tech. rep., Department of Computer Science and Institute for Advanced Computer Studies University of Maryland, 2005.
- [21] C. Yang, R. Duraiswami, N. Gumerov, and L. Davis, "Improved fast gauss transform and efficient kernel density estimation," in *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pp. 664–671, Ieee, 2003.
- [22] C. Yang, R. Duraiswami, and L. Davis, "Efficient kernel machines using the improved fast gauss transform," *Advances in neural information processing systems*, vol. 17, pp. 1561–1568, 2005.
- [23] J. Friedman, J. Bentley, and R. Finkel, "An algorithm for finding best matches in logarithmic expected time," *ACM Transactions on Mathematical Software (TOMS)*, vol. 3, no. 3, pp. 209–226, 1977.
- [24] V. Morariu, B. Srinivasan, V. Raykar, R. Duraiswami, and L. Davis, "Automatic online tuning for fast gaussian summation," in *Advances in Neural Information Processing Systems*, 2008.