



Article RIOT: Recursive Inertial Odometry Transformer for Localisation from Low-Cost IMU Measurements

James Brotchie^{1,*}, Wenchao Li¹, Andrew D. Greentree² and Allison Kealy^{1,3}

- ¹ School of Science, RMIT University, Melbourne, VIC 3001, Australia
- ² ARC Centre of Excellence for Nanoscale BioPhotonics, School of Science, RMIT University, Melbourne, VIC 3001, Australia
- ³ Victorian Department of Environment, Land, Water and Planning, Melbourne, VIC 3000, Australia
- * Correspondence: s3238455@student.rmit.edu.au

Abstract: Inertial localisation is an important technique as it enables ego-motion estimation in conditions where external observers are unavailable. However, low-cost inertial sensors are inherently corrupted by bias and noise, which lead to unbound errors, making straight integration for position intractable. Traditional mathematical approaches are reliant on prior system knowledge, geometric theories and are constrained by predefined dynamics. Recent advances in deep learning, which benefit from ever-increasing volumes of data and computational power, allow for data-driven solutions that offer more comprehensive understanding. Existing deep inertial odometry solutions rely on estimating the latent states, such as velocity, or are dependent on fixed-sensor positions and periodic motion patterns. In this work, we propose taking the traditional state estimation recursive methodology and applying it in the deep learning domain. Our approach, which incorporates the true position priors in the training process, is trained on inertial measurements and ground truth displacement data, allowing recursion and learning both motion characteristics and systemic error bias and drift. We present two end-to-end frameworks for pose invariant deep inertial odometry that utilises self-attention to capture both spatial features and long-range dependencies in inertial data. We evaluate our approaches against a custom 2-layer Gated Recurrent Unit, trained in the same manner on the same data, and tested each approach on a number of different users, devices and activities. Each network had a sequence length weighted relative trajectory error mean ≤ 0.4594 m, highlighting the effectiveness of our learning process used in the development of the models.

Keywords: inertial navigation; deep learning; sensor fusion; odometry; pose estimation; trajectory estimation; self-attention; inertial measurement unit

1. Introduction

Inertial odometry is crucial in mobile agents as it facilitates ego-motion in many applications, such as autonomous driving, health/activity monitoring, indoor navigation, human-robot interaction and augmented/virtual reality. Inertial measurement units (IMUs) are low-power, offer high privacy, and are robust in various environments. As such, they offer a cheap and completely ego-centric means of localisation. IMUs typically consist of a 3D accelerometer, 3D gyroscope and 3D magnetometer. By accurately integrating data from IMUs and other sensors, it is possible to build a reliable system for estimating the motion and position of autonomous systems and pedestrian navigation. However, low-cost inertial sensors have high levels of noise and biases, causing unbound system error growth in long-term inertial navigation [1].

Neural networks have the ability to employ continuous activation functions that inherently understand time and are capable of modelling complex non-linear system behaviours, which are typically too complex for classical mathematical approaches [2]. Recurrent neural networks (RNNs) have long been the primary choice for sequence-tosequence modelling. Most existing deep inertial navigation solutions utilise RNNs, some



Citation: Brotchie, J.; Li, W.; Greentree, A.D.; Kealy, A. RIOT: Recursive Inertial Odometry Transformer for Localisation from Low-Cost IMU Measurements. *Sensors* 2023, 23, 3217. https:// doi.org/10.3390/s23063217

Academic Editors: Rosario Schiano Lo Moriello and Giorgio De Alteriis

Received: 24 February 2023 Revised: 10 March 2023 Accepted: 14 March 2023 Published: 17 March 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). supplementing with convolutional neural networks (CNNs) (see Section 2). However, these architectures have well-documented limitations, such as an inability to capture long-term dependencies and sequential computation that cannot be parallelised [3].

These deficiencies lead to the development of new architectures. The most notable being self-attention-based Transformer models, first proposed in [4], which, since inception, have become ubiquitous in natural language processing (NLP). The success seen in NLP has proliferated its use in a number of domains. Recently, we have seen Transformers employed in computer vision (CV) [5], NLP [6], time-series forecasting [7], image recognition/production [8], text summarisation [9], speech recognition [10] and music generation [11]. These implementations have displayed the network's ability to model long dependencies between input sequence elements and can be processed in parallel, contrasting RNNs. As noted in [12], these capabilities have the advantageous property of resolving the memory bottleneck commonly found in RNNs. Additionally, a comparison to the effectiveness of a long short-term memory (LSTM) network (a RNN variant) and self-attention-based Transformer showed significant performance gains in self-attention-based techniques on datasets with long-range dependencies [13].

In contrast to CNNs, Transformers do not necessitate design specifications and are proficient in handling set functions. Additionally, their uncomplicated architecture facilitates the processing of diverse modalities through the utilisation of homogeneous processing units, which have proven to exhibit remarkable scalability to both large networks and datasets. The incorporation of a self-attention mechanism in neural networks enables the inputs to engage with one other and to be evaluated according to their correlation with the final prediction. Despite extensive investigation into this formulation, limited research has been conducted utilising self-attention and unprocessed sequential readings obtained from low-cost, noisy inertial sensors in the inertial odometry domain. The substantial achievements achieved in other sequence-to-sequence learning problems indicate that the application of self-attention-based techniques could eliminate the need for accurate dynamic models and offer a more robust solution compared to previous RNN or CNN-based methodologies.

2. Related Work

Recent work has shown that the implementation of an accurate inertial odometry algorithm can serve as a foundation for a more robust and reliable ego-motion estimation system through the fusion of inertial sensors. The body of literature in this domain includes both end-to-end solutions and work applying machine learning (ML) directly to enhance the quality of IMU measurements and error models.

In [14], the authors propose a ML-based adaptive neuro-fuzzy inference system to compensate for the errors from a low-cost IMU. Similarly, the authors of [15] look at the effectiveness of different RNN architectures for IMU measurement noise reduction. In [16], the authors successfully utilised a CNN as an accelerometer error reduction method, and the authors of [17] applied a temporal convolutional network (TCN) to construct the gyroscope measurement model.

The authors of [18] demonstrated the feasibility of using a smartphone's 6D inertial sensor (consisting of 3D accelerations and 3D angular velocities) and orientation estimation provided by its application programming interface (API) for pedestrian localisation. Termed RIDI, their approach leverages patterns in natural human movement to learn how to predict velocity and correct linear accelerations using linear least squares. Recent advances in deep learning (DL) have further accelerated data-driven-based inertial navigation. In [19], a deep learning approach called PDRNet was developed for pedestrian dead reckoning (PDR). PDRNet consists of a classification network for smartphone location recognition and a regression network for determining the change in distance and heading. These approaches rely heavily on a large number of carefully tuned parameters and users' walking habits, which leaves them fundamentally susceptible to rapid drift and a lack of generalisability. IONet [20] utilises a bidirectional long short-term memory (Bi-LSTM) and kinetic models

to regress the magnitude velocity and the changing rate of direction. RONIN [21] takes ResNet [22], a LSTM built for CV, as a backbone to again regress a velocity vector. The authors of [23] apply preintegration and an LSTM as a solution to supplement the IMU motion model for deep inertial odometry. These unified deep neural networks provide more robust solutions in highly dynamic conditions. However, they are still reliant on direct integration. Additionally, these methods rely on IMU orientation, are all limited in dimensionality and are dependent on dynamical models that require prior knowledge of the dynamics of the system.

In [24], the authors present TLIO, again using ResNet, to regress 3D displacement estimates and the uncertainty, allowing them to tightly fuse the relative state measurement into a stochastic cloning extended Kalman filter (EKF) to solve for pose, velocity and sensor biases. Owing to its reliance on an EKF, it was shown to be susceptible to a system failure during highly dynamic or unusual motion, which is in line with previous work [25]. Similarly, in [26], the authors propose a hybridised approach using an LSTM and EKF in a modular design that consists of orientation and position subsystems, termed IDOL. Having a dedicated orientation module that included 3D magnetometer readings proved advantageous in contrast to previous approaches that are reliant on the system API. The authors of [27] present a novel loss formulation for smartphone-based deep inertial odometry. The authors use a ResNet-style neural network utilising two-second inertial signals from an IMU to estimate the average velocity and direction of movement. It is noted in this work that despite the obvious benefits of incorporating the magnetometer readings, the network would not converge.

The most recent approach is given in [28] and iterated in [29], where the authors propose attention-driven rotation-equivariance-supervised inertial odometry based on 6D IMU readings. They adopted ResNet to show that adding a self-supervised auxiliary task based on rotation-equivariance can improve the performance of the model when it is jointly trained and can be further improved with a Test-Time Training strategy. In their follow-up, the authors propose a hybrid neural network model for inertial odometry that combines a CNN block with attention mechanisms and a Bi-LSTM network. The CNN block is used to extract spatial features from normalised 6D IMU measurements, and the attention mechanism, which include a spatial attention mechanism and a channel attention mechanism, are used to refine these features. The Bi-LSTM network is then used to capture temporal features.

The effectiveness of data-driven solutions in inertial odometry is well-documented, but these approaches share a common issue in network design. A well-designed network can improve performance in various applications [30], and IMU data, collected at high frequencies, can be challenging to process using traditional machine learning approaches, such as RNNs and CNNs. One drawback of using RNNs for IMU data processing is the issue of "washout", whereby the network's ability to remember past inputs diminishes over time [31], making it difficult to accurately process long sequences of data. On the other hand, CNNs require deep architectures to cover large enough receptive fields to effectively process IMU data, which can result in significant computational expenses during training and deployment [32]. Finally, it should be noted that any effective end-to-end solution for deep inertial odometry must contain the solutions proposed only for measurement error reduction.

3. Contributions

We investigate the efficacy of utilising self-attention for addressing the challenges in inertial navigation. IMU data from low-cost inertial sensors are inherently noisy, biased and incomplete. This can lead to inaccurate readings and make long-term tracking problematic. To address this, we incorporate a sliding window of input data and use prior network outputs as inputs to improve accuracy and robustness. By incorporating multiple readings over a short period of time, the network can average out noise and fill in gaps in the data. The window size is a hyperparameter that can be adjusted to balance the trade-off between incorporating enough information to improve accuracy while avoiding over-fitting to noise or short-term fluctuations. This formulation is designed to emulate the recursiveness leveraged in traditional mathematical approaches such as an EKF [33].

To mitigate the challenges associated with high-frequency time series data, we propose reliance entirely on the self-attention mechanism to compute representations of the inputs and outputs rather than using sequence-aligned RNNs or convolutions. As the self-attention mechanism will be the primary method for extracting information from the inputs and generating the outputs, there is the potential to be more efficient and flexible than using RNNs or CNNs, as self-attention mechanisms can capture long-range dependencies in the data and can be parallelised during the training process. Additionally, self-attention mechanisms can provide a degree of interpretability by allowing the model to identify the most important input features at each time step.

We term our approaches: Recursive Inertial Odometry Transformer (RIOT) and Attitude Recursive Inertial Odometry Transformer (ARIOT). To the best of our knowledge, our approaches are the only networks that leverage self-attention and all available IMU information (from a 3D accelerometer, 3D gyroscope and 3D magnetometer) to provide an end-to-end, 3D inertial odometry solution.

For the network design of RIOT, a number of modifications were made to the original Transformer proposed in [4]. The embedding layer is replaced by a generic linear layer to reduce the dimensionality of the input. Additionally, residual connections were added between the multi-head attention and the feed-forward layers to improve information flow through the network. Lastly, we forgo an activation function after the final linear layer to facilitate boundless position estimation.

The ARIOT model is a hierarchical Transformer; it differs from RIOT by the incorporation of an additional, internal attitude estimation network that regresses the orientation of the IMU from the sensor measurements. This subsystem benefits from the self-attentionbased framework design in [34]. However, we were able to formulate a new loss function, which, to the best of our knowledge, is absent in the literature. The angle between quaternions is a well-known quantity; however, when training a network to regress to unit quaternions, the inner product is frequently outside the range [-1, 1], resulting in numerical instability. We propose the use of Equation (18) to negate this instability. The output is used in the odometry network to further regress the accelerometer readings and prior position to give updated IMU-based localisation.

The effectiveness of RIOT and ARIOT is validated on unseen sequences in their entirety from different users, activities and smartphone IMU devices.

4. Problem Formulation

4.1. Sensor Models

First, we consider the problem of modelling measurements from a 9D IMU. It is implicit that these systems are characterised by high noise levels and time-varying additive biases. The available measurements from a typical IMU are from three-axis rate gyros, three-axis accelerometers and three-axis magnetometers. The reference frame of the IMU is termed the body frame (B), which is rotated with respect to some fixed inertial frame (I), e.g., the Earth-centered inertial (ECI) frame or the North-East-Down (NED) frame. However, for brevity, these reference frames are assumed and not incorporated into the notation.

The gyroscope measures the angular velocity of *B* relative to *I*, corrupted by a slowly varying bias and noise. Therefore, we can define the gyroscope measurements, $I_{\omega,t}$, as

$$\mathbf{I}_{\omega,t} = \boldsymbol{\omega}_t + \boldsymbol{\delta}_{\omega,t} + \boldsymbol{e}_{\omega,t} \tag{1}$$

where ω_t is the true angular velocity at each time instance t, $\delta_{\omega,t}$ denotes the time-varying bias and $e_{\omega,t}$ is the noise, typically assumed to be Gaussian, $e_{\omega,t} \sim \mathcal{N}(0, \Sigma_{\omega})$.

The accelerometer measures the linear acceleration of *B* relative to *I*. Again, with added noise and bias, the accelerometer measurements, $I_{a,t}$, are given by

$$\mathbf{I}_{a,t} = \mathbf{f}_t + \delta_{a,t} + \boldsymbol{e}_{a,t} \tag{2}$$

where \mathbf{f}_t is the specific force at each time instance t, and $\delta_{a,t}$ and $e_{a,t}$ denote the bias and noise, respectively, with $e_{a,t} \sim \mathcal{N}(0, \Sigma_{\omega})$.

Magnetometers provide information about the direction and intensity of the local magnetic field surrounding the sensor. The local magnetic field is composed of the Earth's magnetic field as well as any additional magnetic fields that arise due to the existence of magnetic materials. As magnetometer measurements are used primarily in attitude determination, we assume the magnitude of the local magnetic field vector, denoted by \mathbf{m}^l , is equal to 1—i.e., $\|\mathbf{m}^l\| = 1$. Assuming that the magnetometer only measures the local magnetic direction, its measurements, $\mathbf{I}_{m,t}$, can be modelled as

$$\mathbf{I}_{m,t} = \mathbf{R}_t^b \mathbf{m}^l + \boldsymbol{e}_{m,t},\tag{3}$$

where \mathbf{R}_t^b denotes the rotation matrix from navigation to body frame and $\mathbf{e}_{m,t} \sim \mathcal{N}(0, \Sigma_m)$ is the Gaussian noise.

By incorporating magnetometer measurements, we enable the system to determine its initial attitude. This is predicated on the principle that, given a set of two or more linearly independent vectors in two distinct reference frames, the rotation between said frames can be calculated. The underlying assumption here is that the accelerometer only measures the gravity vector, and the magnetometer only measures the local magnetic field. Hence we have four linearly independent vectors: measurements $\mathbf{I}_{a,t}$ and $\mathbf{I}_{m,t}$, the local gravity vector g^n , and the local magnetic vector, \mathbf{m}^l . Whilst this is seen as a major advantage, it does come with the drawback of requiring local magnetic field knowledge in order to transform Equation (3) into local coordinates.

4.2. Attitude and Position Estimation

Traditional attitude estimation approaches rely on gyro integration as the baseline for deriving the attitude. However, it is well documented that gyroscope measurements lack the information to give absolute attitude determination. Therefore, applying numerical gyro integration results in an accumulated error that grows boundlessly. As such, specific force measurements from an accelerometer are often used in tandem with magnetic field measurements from a magnetometer to correct the estimate, as they provide information on the absolute angular position. These methods typically involve complex mathematical models and computations and require prior state knowledge and specific sensor parameters [35].

Analogous to attitude estimation, traditional methods for position estimation are also susceptible to unbounded errors due to the lack of information for an absolute position change. To overcome this limitation, we propose using self-attention and raw data in gradient descent optimisation to analyse and retain information related to accelerometer error and bias over long sequences. The relevant features are extracted by the attention mechanism to learn the relationship between acceleration, attitude and position. In the case of ARIOT, this method also has the caveat of recognising and compensating for attitude estimation errors in the initial attitude estimation network.

We use the accelerometer and gyroscope measurements as inputs to the dynamics for the purpose of estimating the position. The state vector includes the position and a quaternion parametrisation of the attitude (detailed in Section 5.2). We use the inertial measurements along with prior positions to estimate the attitude and position.

The dynamics of the position for an interval of time Δt are given by the equation

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \Delta t \boldsymbol{v}_t + \frac{\Delta t^2}{2} \left(\mathbf{R}_t^n (\mathbf{I}_{a,t} - \delta_{a,t}) + \boldsymbol{g}^n + \boldsymbol{e}_{a,t} \right)$$
(4)

where v_t is the velocity, time is denoted as t and \mathbf{R}_t^n is the rotation matrix from the body to the navigation frame. We switched the sign on the noise term for convenience. The dynamics of the attitude is then given by

$$\mathbf{q}_{t+1} = \mathbf{q}_t \odot \exp_q \left(\frac{\Delta t}{2} (\mathbf{I}_{\omega,t} - \boldsymbol{\delta}_{\omega,t} - \boldsymbol{e}_{\omega,t}) \right) \odot \exp_q \left(\frac{\Delta t}{2} f(\mathbf{I}_{a,t}, \mathbf{I}_{m,t}) \right)$$
(5)

where $f(\cdot)$ is a function of the accelerometer and magnetometer measurements to calculate the correction term for the quaternion, and the notation \odot denotes the quaternion multiplication given by

$$\begin{pmatrix} j_1 \\ j_2 \\ j_3 \\ j_4 \end{pmatrix} \odot \begin{pmatrix} k_1 \\ k_2 \\ k_3 \\ k_4 \end{pmatrix} = \begin{pmatrix} j_1k_1 - j_2k_2 - j_3k_3 - j_4k_4 \\ j_1k_2 + j_2k_1 + j_3k_4 - j_4k_3 \\ j_1k_3 - j_2k_4 + j_3k_1 + j_4k_2 \\ j_1k_4 + j_2k_3 - j_3k_2 + j_4k_1 \end{pmatrix}.$$
 (6)

A significant advantage of traditional state estimation algorithms over the neural networks is the retention of prior state estimates to update subsequent states. This allows for the algorithm to use past information to correct or refine the current estimate, which improves accuracy and reliability. In contrast, neural networks are typically trained to make predictions based on input data without explicit retention of past estimates.

We recognise that RNNs are somewhat the exception here as they can be designed to have internal state memory that can retain past state information; however, this does not actually give the network the desired recursive property. Instead, it acts more as a pseudo-recursion in which the hidden states store some information from previous steps and use it to influence future steps. This distinction is crucial in understanding the nature and limitations of RNNs in terms of recursive behaviour. Furthermore, networks that adopt this come with the aforementioned memory bottlenecks and vanishing gradient drawbacks.

5. Proposed Solution

5.1. Network Components

ARIOT and RIOT models are presented in Sections 5.2 and 5.3, respectively. Here we will introduce the common components found in both networks. The modular-specific adaptations will follow. In model design, we follow the original NLP Transformer proposed in [4], comprising encoder-decoder blocks and multi-head attention (MHA). An advantage of this intentionally straightforward system design is that it is efficient to implement and provides an out of the box solution. The input of the standard Transformer is a 1D sequence of token embeddings. To handle IMU data, the sequence embeddings are expanded to *N*-dimensions corresponding to feature inputs, each with a set of additional position embeddings, which represent the temporal information. The network produces a sequence of representations for each input time step, which is then used as feature vectors in downstream tasks.

5.1.1. Positional Encoding

In the Transformer model, as described in [4], the relative sequential position is not explicitly encoded. To incorporate relative sequential position information, we add sinusoidal position encoding functions over the inputs before the first layer. The values of the encoding are calculated using the trigonometric functions sin and cos. The argument of these functions is the product of the sequence position (*pos*) and a scaling factor $10000^{2i/d_{model}}$, where *i* is an index variable, $0 \le i \le \frac{d_{model} - 1}{2}$, used to calculate different dimensions of the positional encoding vector. For each value of *i*, two dimensions of the positional encoding the sine function ($PE_{(pos,2i+1)}$). The idea behind using both the sine and cosine functions is to capture both the magnitude and phase of the sequence. This, in turn, allows the model to

attend different parts of the input sequence at different stages of processing. Additionally, this approach allows the model to generalise to different sequence lengths and attend to elements based on their relative positions rather than their absolute positions. This can be useful when the model needs to handle sequences of different lengths or in cases of mismatched sampling [36]. The positional encoding in both networks is defined by

$$PE_{(pos,2i)} = \sin\left(pos/10000^{2i/d_{\text{model}}}\right)$$

$$PE_{(pos,2i+1)} = \cos\left(pos/10000^{2i/d_{\text{model}}}\right)$$
(7)

where *pos* denotes the position, *i* the dimension and d_{model} is the model dimensionality; in this work d_{model} is 64 and 224 for the attitude and position networks (see Section 7), respectively.

5.1.2. Self-Attention

Self-attention sublayers in these networks employ h = 2 heads. The self-attention mechanism works by first projecting the IMU measurements into a higher-dimensional space using a linear transformation parameterised by a set of weights $W^Q, W^K, W^V \in \mathbb{R}^{d_1 \times d_{model}}$. This projection is parameterised by a set of weights, which are learned during training. These parameter matrices are unique per layer and attention head. The transformed input data is then passed through a function (often called the "attention function"), which produces a set of attention weights for each input element, representing the importance of each input element in regressing to an attitude or position estimation.

Each attention head operates on an input sequence $\mathbf{I} = (I_1, \ldots, I_n)$ of *n* elements, where $I_i \in \mathbb{R}^{d_1}$. A new sequence of the same length in computed as $\mathbf{z} = (z_1, \ldots, z_n)$, where $z_i \in \mathbb{R}^{d_{\text{model}}}$, and each output element is computed as a weighted sum of a linearly transformed input per

$$z_i = \sum_{j=1}^{l} \alpha_{i,j} \left(x_j \mathbf{W}^V \right) \tag{8}$$

where each weight coefficient, $\alpha_{i,j}$, is computed through a softmax function, which normalises the compatibility scores for each element to produce a probability distribution over the input sequence

$$\alpha_{i,j} = \frac{\exp e_{i,j}}{\sum_{k=1}^{n} \exp e_{i,k}}$$
(9)

where e_{ii} is computed using a compatibility function that compares two input elements,

$$e_{ij} = \frac{\left(x_i \mathbf{W}^Q\right) \left(x_j \mathbf{W}^K\right)^T}{\sqrt{d_z}}.$$
(10)

Scaled dot product is used as the compatibility function to enable efficient computation. Linear transformations of the inputs add sufficient expressive power. The self-attention layer is implemented using MHA.

In the context of using IMU information to estimate an attitude quaternion or position, the self-attention mechanism is used to weigh the different sensor measurements differently, depending on how relevant they are to the position estimate. For example, the gyroscope measurements may be given a higher weight than the accelerometer measurements when estimating rotational motion, while the accelerometer measurements may be given a higher weight when estimating linear acceleration.

5.1.3. Encoder

The element-wise addition of the input vector and positional encoding vector is fed into two identical encoder layers. Each encoding layer is made up of two sub-layers: a MHA sub-layer and a fully connected feed-forward (FF) sub-layer. In the case of the ARIOT Our encoder follows the Query–Key–Value model, proposed in [4], where the scaled dot-product attention used is given by

Attention(
$$\mathbf{Q}, \mathbf{K}, \mathbf{V}$$
) = softmax $\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{D_k}}\right)\mathbf{V}$ (11)

where the input, **I**, is used to obtain the queries $\mathbf{Q} = \mathbf{I}(k)\mathbf{W}^Q \in \mathbb{R}^{N \times D_k}$, keys $\mathbf{K} = \mathbf{I}(k)\mathbf{W}^K \in \mathbb{R}^{M \times D_k}$ and values $\mathbf{V} = \mathbf{I}(k)\mathbf{W}^V \in \mathbb{R}^{M \times D_v}$; each **W** is the respective weight matrices updated during training, N, M denote the lengths of queries and keys (or values) and D_k, D_v denote the dimensions of keys (or queries) and values. The MHA consists of H different sets of learned projections instead of a single attention function as

MultiHeadAttn($\mathbf{Q}, \mathbf{K}, \mathbf{V}$) = Concat(head₁,...,head_H) \mathbf{W}^{O}

where head_i = Attention $(\mathbf{QW}_{i}^{Q}, \mathbf{KW}_{i}^{K}, \mathbf{V}_{i}^{V})$.

The projections are parameter matrices $W_i^Q \in \mathbb{R}^{D_{\text{model}} \times D_k}, W_i^K \in \mathbb{R}^{D_{\text{model}} \times D_k}, W_i^V \in \mathbb{R}^{D_{\text{model}} \times D_v}$ and $W^O \in \mathbb{R}^{hD_v \times D_{\text{model}}}$. In this work, we employ h = 2 parallel attention layers or heads. For each, we use $D_k = D_v = D_{\text{model}} / h$.

In addition to the attention sub-layers, each encoder/decoder layer consists of a fully connected FF network consisting of linear transformation and activation functions. We use a LeakyReLU [37] activation in the FF network as follows

LeakyReLU(x) =
$$\begin{cases} x, & \text{if } x \ge 0\\ 1 \times 10^{-3} \cdot x, & \text{otherwise} \end{cases}$$
 (12)

The point-wise FF network is a fully connected module

$$FFN(\mathbf{H}') = LeakyReLU(\mathbf{H}'\mathbf{W}^{1} + \mathbf{b}^{1})\mathbf{W}^{2} + \mathbf{b}^{2}$$
(13)

where \mathbf{H}' is the output of the previous layer, $\mathbf{W}^1 \in \mathbb{R}^{D_m \times D_f}$, $\mathbf{W}^2 \in \mathbb{R}^{D_f \times D_m}$, $\mathbf{b}^1 \in \mathbb{R}^{D_f}$ and $\mathbf{b}^2 \in \mathbb{R}^{D_m}$ are trainable parameters and D_f denotes the inner-layer dimensionality. Each sub-layer has a Layer Normalisation Module inserted around each module. That is,

$$\mathbf{H}' = \text{LayerNorm}\big(\text{SelfAttn}(\mathbf{X}) + \mathbf{X}\big) \tag{14}$$

where SelfAttn(\cdot) denotes self-attention module and LayerNorm(\cdot) the normal layer operation. The resultant vector is then fed into the decoder.

5.1.4. Decoder

The decoder is composed of two identical layers. The decoder contains the sub-layers found in the encoder, with the addition of a third sub-layer that performs MHA over the output vector from the encoder. The MHA mechanism allows the model to attend to multiple parts of the input sequence in parallel, allowing it to capture a more detailed and nuanced representation of the input. This is achieved by dividing the attention mechanism into multiple "heads". Each head performs attention with a different linear projection. Additionally, the self-attention mechanism in the decoder stacks prevents positions from influencing subsequent positions to ensure that predictions at *k* can depend only on the known outputs at or before k - 1. In our attitude network, the output maps the final layer into the estimated quaternion through a hyperbolic tangent. For the position, no outbound function is used past linearisation.

5.2. Attitude Recursive Inertial Odometry Transformer

The Attitude Recursive Inertial Odometry Transformer is a hierarchical framework composed of two self-attention-based encoder-decoder networks, depicted in Figures 1 and 2. The foundation for the initial network is based on previous work [34] and functions to regress attitude estimation from 9D inertial measurements (from Equations (1)–(3)). This allows for the componential estimation of both attitude and position estimation in a single framework, providing a robust solution for inertial odometry. The use of self-attention mechanisms within both modules allows for the modelling of long-term dependencies in the data, effectively handling the high dynamic motion present over long sequences.



Figure 1. Information flow in the ARIOT model training during the first cycle (**left**) and second cycle (**right**). The $_Ax$ and $_px$ denote the attitude, and following positional neural network that is described in Section 5.2 and and visualised in Figure 2. The network inputs are given in Section 4.1, specifically Equations (1)–(3). Inference is performed in a similar depiction as the second cycle, without minimisation of the loss, detailed in Section 6.3.



Figure 2. ARIOT model architecture, where the networks, ${}_{A}x$ and ${}_{p}x$ depicted in Figure 1, are visualised left to right, respectively. The network inputs are given in Section 4.1, specifically Equations (1)–(3).

We follow by parameterising the attitude in quaternions. Quaternions, which are a type of representation of attitudes in \mathbb{R}^4 , have several advantages over representations in \mathbb{R}^3 . They are free of discontinuities and singularities and are more computationally efficient and numerically stable. To be a valid representation of an attitude, a quaternion must be a unit quaternion. Unit quaternions have a one-to-one correspondence with rotation matrices, and they double cover the group SO(3), meaning that both **q** and $-\mathbf{q}$ represent the same attitude. However, by requiring that $\mathbf{q}_0 \ge 0$, we can ensure that there is a unique correspondence between quaternions and rotation matrices [38].

We propose to use the self-attention mechanism and raw 9D IMU data in gradient descent optimisation to analyse and retain information related to gyroscope error and bias over long sequences. This minimises the complexity by forgoing preintegration. Additionally, the solution is unconstrained by not forcing the network into predefined dynamic models. These features and the inclusion of magnetometer measurements also have the advantage of the network being an out-of-the-box solution where the local magnetic field is known.

We propose a new loss function for quaternions, which we call the Quaternion Loss. To define the Quaternion Loss, we first introduce some quaternion background and notation. A quaternion is a 4-tuple (x, y, z, w), where x, y, z, w are real numbers. Quaternions can be represented in the form

$$\mathbf{q} = w + xi + yj + zk \tag{15}$$

where *i*, *j*, *k* are the imaginary units, satisfying $i^2 = j^2 = k^2 = ijk = -1$.

Quaternions can be used to represent rotations in three-dimensional space by setting w to the cosine of the rotation angle and x, y, z to the sine of the rotation angle, multiplied by the rotation axis [39]. Given a pair of quaternions ($\mathbf{q}_1, \mathbf{q}_2$), we can measure the similarity between them using the inner product as

$$\langle \mathbf{q}_1, \mathbf{q}_2 \rangle = x_1 x_2 + y_1 y_2 + z_1 z_2 + w_1 w_2$$
 (16)

This product is related to the angle between the quaternions by the following,

$$\cos(\theta) = \frac{\langle \mathbf{q}_1, \mathbf{q}_2 \rangle}{|\mathbf{q}_1| \cdot |\mathbf{q}_2|} \tag{17}$$

where θ is the angle between the quaternions and $|\cdot|$ denotes the L2 norm.

We then define the Quaternion Loss function as

$$\mathcal{L}(\mathbf{q}_1, \mathbf{q}_2) = \cos^{-1}(\operatorname{clamp}(\langle \mathbf{q}_1, \mathbf{q}_2 \rangle, -1 + \epsilon, 1 - \epsilon))$$
(18)

where

$$\operatorname{clamp}(x, a, b) = \begin{cases} a & \text{if } x < a \\ x & \text{if } a \le x \le b \\ b & \text{if } x > b \end{cases}$$

and ϵ is a small positive constant used to avoid numerical instability when the inner product is outside the range [-1, 1].

The mean angle across the batch is then returned as

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} \theta_i \tag{19}$$

where *N* is the batch size.

5.3. Recursive Inertial Odometry Transformer

The Recursive Inertial Odometry Transformer is a self-attention-based encoder-decoder network. Forgoing the attitude module to directly apply self-attention to raw 9D IMU data (from Equations (1)–(3)) in gradient descent optimisation for 3D displacement regression, depicted in Figure 3.



Figure 3. RIOT architecture and information flow during training cycles. (**a**) RIOT model architecture, where the network, $_{R}x$ depicted in (**b**), is visualised. The network inputs are given in Section 4.1, specifically Equations (1)–(3). (**b**) Information flow of the RIOT model training during the first cycle (**top**) and second cycle (**bottom**). The $_{R}x$ denotes the neural network described in Section 5.3 and shown in (**a**). Inference is performed in a similar depiction as the second cycle, without minimisation of the loss, detailed in Section 6.3.

The input to the network is a concatenation of the inertial measurements and true position priors in the first cycle of training; then, true position priors are replaced by estimated position priors. The input is passed through an embedding layer to generate embedded representations. The encoder then applies self-attention to compute a weighted sum of the embedded representations for each time step, which is used to compute a context vector. The context vector is then passed through a decoder to estimate the 3D position at each time step. The equations for the input sequence, the embedding function and the self-attention mechanism are provided in Section 5.1. The model is then trained to minimise the Mean Square Error (MSE) loss function in Equation (20) using the ADAM optimisation algorithm [40].

$$\mathcal{L}(\hat{\mathbf{p}}, \mathbf{p}) = \frac{1}{NT} \sum_{n=1}^{N} \sum_{t=1}^{T} \|\hat{\mathbf{p}}_{n,t} - \mathbf{p}_{n,t}\|^2$$
(20)

where $\|\cdot\|^2$ represents the squared Euclidean norm, *N* is the batch size, *T* the sequence length and $\hat{\mathbf{p}}$ and \mathbf{p} are the estimated and true positions, respectively.

6. Evaluation

Despite numerous proposed solutions in the literature attempting to solve inertial navigation, these approaches evaluate their algorithms using their datasets with various preprocessing and alignment techniques, such as the Umeyama algorithm [41]. Under these conditions, it is difficult to compare directly to these different algorithms. Additionally, to the best of our knowledge, no other approach leverages all available IMU information. However, the inclusion of magnetometer measurements comes with the drawback of our solutions being dependent on the local magnetic field, as the magnetometer readings are

used to disambiguate the orientation of the IMU. This results in our network calibrations being regional-specific and not generalisable to other datasets without local magnetic field knowledge. To this end, we build on our own implementation of a RNN as a means of comparison.

6.1. Gated Recurrent Unit

Recent work on RNNs has shown that a Gated Recurrent Unit (GRU) surpasses the preferred LSTM in a number of scenarios [42–44]. Additionally, GRUs have fewer parameters, making it more computationally efficient, and has been shown to be more robust to noise and missing data [45].

We have added our own implementation of a two-layer GRU as a means of comparison. GRU has been shown to be effective in inertial attitude estimation [46]; however, to the best of our knowledge, GRUs are untested in the inertial odometry domain. The network is formulated with the hidden state h_t at time step t as follows

$$r_t = \sigma(\mathbf{W}_{ir}x_t + \mathbf{b}_{ir} + \mathbf{W}_{hr}h_{t-1} + \mathbf{b}_{hr})$$
(21)

$$z_t = \sigma(\mathbf{W}_{iz}x_t + \mathbf{b}_{iz} + \mathbf{W}_{hz}h_{t-1} + \mathbf{b}_{hz})$$
(22)

$$\tilde{h}_t = \text{LeakyReLU}(\mathbf{W}_{ix}x_t + \mathbf{b}_{ix} + r_t * (\mathbf{W}_{hx}h_{t-1} + \mathbf{b}_{hx}))$$
(23)

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h_t}$$
(24)

where x_t is the input at time step t, \mathbf{W}_{i*} and \mathbf{b}_{i*} are the input-to-hidden weights and biases, \mathbf{W}_{h*} and b_{h*} are the hidden-to-hidden weights and biases and σ is the sigmoid function. Equations (21) and (22) compute the reset gate, r_t , and update gate, z_t , respectively. These gates control the amount of information that is passed through to the next time step. Equation (23) computes the candidate hidden state \tilde{h}_t , and Equation (24) updates the hidden state h_t by combining the previous hidden state h_{t-1} and the candidate hidden state \tilde{h}_t .

We implement this network in the same manner as RIOT, depicted in Figure 4, where a stack of two GRU layers transforms the 9D IMU input at sampling instant *t*, concatenated with the 3D position vector at time t - 1, to an N_n -dimensional feature vector h_t , with $N_n = 200$ being the number of neurons per layer.



Figure 4. GRU model training during the first cycle (left) and second cycle (right). Our model is made up of two GRU cells, described in Section 6.1, with 200 neurons per layer. The network inputs are given in Section 4.1, specifically Equations (1)–(3). Inference is performed in a similar depiction as the second cycle, without minimisation of the loss, detailed in Section 6.3.

6.2. Training and Dataset

This approach was trained and tested on publicly available smartphone data published by Chen et al. [47]. The dataset contains 158 sequences, totalling more than 42 km in total distance and incorporates a variety of attachments, activities and users to best reflect the broad use cases seen in real life. The data were captured via five different users and four different types of off-the-shelf consumer smartphones. The IMU data was collected and synchronised with a frequency of 100 Hz, which is generally accepted in various applications and research [48–50]. A high-precision optical motion capture system (Vicon) was used to capture full pose ground truth at 0.01 m location and 0.1 degree attitude accuracy [51]. The dataset was randomly divided into training, validation and test sets, following [52]. A single sequence was left out for each of the variables as a means of complete, unseen comparison with other techniques. To avoid overfitting and to improve the computing efficiency, we used a sliding window to capture 100 measurements every 50 to feed into the encoder and used random search to tune the hyperparameters. This gave us 63,614 training samples, 18,175 validation samples and 9089 test samples. The implementation of all adaptations was carried out with PyTorch. The attitude network converges after 300 epochs. The position network converges after 120 and 30 epochs, using true and recursive inputs, respectively. A learning rate of 0.001, an ADAM optimiser and a dropout of 0.2 were used across each implementation. The training was conducted in parallel on $4 \times$ Nvidia V100 GPUs.

6.3. Inference

The inference procedure for each model closely resembles the second training cycle, as illustrated in Figures 1, 3b and 4 for ARIOT, RIOT and GRU, respectively. The initial window of each sequence is pre-padded with zeroes, followed by a given initial position. The position window is then iteratively updated by processing the subsequent windows of data. This inference approach is designed to reflect both recurrent architectures and recursive mathematical models whilst leveraging the benefits of self-attention. This process is visualised in Figure 5.





6.4. Evaluation Metrics

In order to quantitatively assess the performance of each approach on each unobserved sequence of length *K*, the following three metrics were employed:

• Absolute Trajectory Error (ATE) (m)

$$ATE = \sqrt{\frac{1}{K} \sum_{k=1}^{K} \left\| \left(\hat{\mathbf{p}}_{k} - \mathbf{p}_{k} \right) \right\|^{2}}$$
(25)

The ATE is commonly used to assess the performance of a guidance or navigation system and represents the global accuracy of the estimated position.

• Relative Trajectory Error (RTE) (m)

$$\sqrt{\frac{1}{K}\sum_{k=1}^{K} \left\| \left(\hat{\mathbf{p}}_{k+\Delta t} + \mathbf{p}_{k+\Delta t} \right) - \left(\hat{\mathbf{p}}_{k} + \mathbf{p}_{k} \right) \right\|^{2}}$$
(26)

The RTE is a measure of the difference between the estimated and true position at a given time relative to the distance between the two positions. It is often used to quantify the location position consistency over a predefined duration Δt ; $\Delta t = 1$ s in

this work.

Cumulative Distribution Function (CDF)

$$\int_0^e f(x)dx, \quad \boldsymbol{e} = \left[\frac{1}{K}\sum_{k=1}^K \sqrt{\left\|\left(\hat{\mathbf{p}}_k - \mathbf{p}_k\right)\right\|^2}\right]$$
(27)

The CDF is the distribution function f(x), used to characterise the distribution of a variable. In this context, it is used to describe the probability that the error in the estimated position will be less than or equal to a certain value. f(x) is the probability density function of the localisation error *e*.

ATE and RTE are used in deep inertial odometry papers [21,29], and CDF is a common metric in indoor localisation research [53].

6.5. Discussion

This work presents three approaches for evaluation of unseen sequences from different users, devices and activities. The ATE and RTE evaluation results are quantified in Table 1, with the best-performing approach for each sequence and metric highlighted in bold. In addition, a qualitative analysis was conducted on the model's output, which revealed a close correspondence between the predicted trajectories and the ground truth trajectories. This is depicted in Appendix A, which provides visualisations of the position estimates for each approach during the first and last minute of data.

Table 1. Two-dimensional position error metric comparison. A complete sequence was left out of the training data for each variable in the dataset. This was performed as a means of unseen comparison over full sequences, allowing for different user, activity and device evaluations as well an overlook at the generalisability of each approach. Note that each network is capable of producing a 3D position estimate; however, as the data was largely taken on a level plane where the discrepancy in the *z*-axis is far smaller than the *x-y* plane, the addition of the vertical dimension would skew the error metrics. The best-performing model over each sequence and for each metric has been made bold. RIOT performs best under most conditions. However, ARIOT tracks better during highly dynamic motion.

| | User 2 | | User 3 | | User 4 | | User 5 | | Pocket | | Running | |
|-------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| Model | ATE | RTE |
| | (m) |
| GRU | 0.0796 | 0.0110 | 0.0692 | 0.0100 | 0.0757 | 0.0121 | 0.0856 | 0.0114 | 0.1013 | 0.125 | 0.1589 | 0.0171 |
| ARIOT | 0.0994 | 0.0093 | 0.0934 | 0.0088 | 0.0960 | 0.0094 | 0.1027 | 0.0100 | 0.1059 | 0.0088 | 0.1279 | 0.0144 |
| RIOT | 0.0681 | 0.0090 | 0.0655 | 0.0085 | 0.0654 | 0.0091 | 0.0721 | 0.0096 | 0.0676 | 0.0085 | 0.0990 | 0.0140 |
| | Slow Walking | | Trolley | | Handbag | | Handheld | | iPhone 5 | | iPhone 6 | |
| Model | ATE | RTE |
| | (m) |
| GRU | 0.2634 | 0.0077 | 0.0881 | 0.0116 | 0.2021 | 0.0112 | 7.352 | 0.0357 | 0.1172 | 0.0138 | 0.1133 | 0.0110 |
| ARIOT | 0.1082 | 0.0060 | 0.1033 | 0.0099 | 0.1096 | 0.0091 | 0.3196 | 0.0129 | 0.1046 | 0.0090 | 0.1036 | 0.0089 |

A closer look at Figures A1 and A2 in Appendix A gives a clear visualisation of the advantages of the self-attention-based models in maintaining smooth, life-like path trajectories that almost mirror the true path. In contrast, the GRU path estimates are seemingly noisy, consistently creating a far greater total distance length than RIOT or ARIOT. Whilst all models are implemented recursively, the GRU's inability to attend to the entire sequence in updating the current position greatly affects the overall performance.

RIOT performed best overall with the lowest ATE and RTE values, with the exception of when the IMU was handheld. When the IMU is handheld and has implied consistent dynamic motion, the attitude estimation module in ARIOT is beneficial as it can help to disambiguate the accelerometer measurements that are affected by both linear acceleration and gravity. This led to a more accurate trajectory estimate. However, when the IMU is mostly stable or cyclic through the motion, the additional complexity of the attitude estimation module is redundant and actually hinders performance. We hypothesise that the model overly leans on attitude representation, which is only beneficial in highly dynamic scenarios.

When analysing the performance of our models, it is important to consider the characteristics of the data and the specific scenario. We theorise that the reason for the superior performance of RIOT is due to the simpler architecture, which is seemingly better suited for scenarios where the IMU is less dynamic. On the other hand, the additional complexity of ARIOT's attitude estimation module allows for improved handling of dynamic motion.

It is evident that the GRU performed considerably worse than both RIOT and ARIOT in all of the sequences. This is likely due to the fundamentally inferior design of the RNN model, leading to its inability to effectively process the complex motion present over long sequences. However, it is important to note that the GRU model still performed relatively well, which highlights the effectiveness of our learning process used in the development of the models.

This analysis is further evidenced in Figure 6, which depicts the mean CDF of the localisation error over the total set of test sequences. RIOT performs almost consistently, indicated by the steep gradient of the CDF in the lower error range, whereas for ARIOT and GRU, the errors are more spread out over a wider range of values.



Figure 6. CDF of localisation error for each approach totalled over all test sequences. The CDF plots the percentage of localisation error values that fall below a certain percentage threshold, allowing evaluation for each approach at different levels of localisation error (m). RIOT performs best with a higher concentration of low errors in comparison to ARIOT and GRU.

Our models utilise multi-headed self-attention, which is achieved through multiple parallel attention mechanisms. By allowing the model to attend to different parts of the input sequence dynamically, self-attention can capture complex relationships and dependencies. Each self-attention mechanism calculates an attention matrix A of size $T \times T$, where T is the sequence length, by utilising the softmax operation, as described in Equation (11). The attention scores determine the influence of the input time features on the higher-level output time features.

The matrix visualisations in Figures 7a,c and 8a,c, provide a glimpse into how the model is weighing and combining multiple inputs to make a prediction. The values of the attention matrix depict two attention heads from the first self-attention layer from each

encoder as an adjacency matrix between input nodes and output nodes. The matrix can also be represented as a bipartite graph, as shown in Figures 7b,d and 8b,d. The edge weights represent the strength of the attention, and the opacity of the edges indicates the magnitude. The input time series is shown above the attention graph as a reference, and the attention scores are depicted as vertical lines corresponding to the values in *A*.

From the visual representations of the attention matrices, we can directly observe the distinctions between the different self-attention heads and encoders. The heads in the first encoder appear to be highly concentrated on the latter part of the sequence, whereas the heads in the second encoder concentrate on the beginning but have greater overall attention. From Figures 7 and 8, we observe that the model considers both short and longer-term temporal dependencies in the data when making predictions rather than just focusing on the prior time step, as seen in traditional methodologies. This is largely the reason for the accurate and stable position estimates, especially in situations where the motion is complex or noisy.



Figure 7. Visualisations of the self-attention scores from the first encoder in RIOT on an arbitrary sequence of input data. (**a**,**b**) (blue) depict the attention scores from the first head of the first encoder as a matrix and bipartite graph, respectively. (**c**,**d**) (red) depict the attention scores from the second head of the second encoder as a matrix and bipartite graph, respectively. (**Left**): The heat matrix displays the attention scores assigned to each input element in a sequence. The darker the colour, the higher the attention weight given to that element, indicating that it has a greater impact on the final output. (**Right**): The graph represents each input element as a node on one side of the graph, while the attention scores assigned by the model are represented as nodes on the other side. Edges connecting the nodes represent the attention weights or the degree to which the model is considering each input element. The thickness of the edges represents the magnitude of the attention weights, with thicker edges indicating higher attention scores.



Figure 8. Visualisations of the self-attention scores from the second encoder in RIOT on an arbitrary sequence of input data. (**a**,**b**) (blue) depict the attention scores from the first head of the second encoder as a matrix and bipartite graph, respectively. (**c**,**d**) (red) depict the attention scores from the second head of the second encoder as a matrix and bipartite graph, respectively. (**c**,**d**) (red) depict the attention scores from the second head of the second encoder as a matrix and bipartite graph, respectively. (**Left**): The heat matrix displays the attention scores assigned to each input element in a sequence. The darker the colour, the higher the attention weight given to that element, indicating that it has a greater impact on the final output. (**Right**:) The graph represents each input element as a node on one side of the graph, while the attention scores assigned by the model are represented as nodes on the other side. Edges connecting the nodes represent the attention weights or the degree to which the model is considering each input element. The thickness of the edges represents the magnitude of the attention weights, with thicker edges indicating higher attention scores.

In summary, we evaluated the performance of three novel recursive deep inertial odometry frameworks. Our results show that self-attention-based networks have superior performance over a GRU-based RNN, with RIOT performing best overall, with a sequence length weighted mean ATE of 0.0865 m and RTE of 0.0091 m. The mean RTE and ATE of ARIOT and GRU were 0.1134 and 0.0095 m and 0.4594 and 0.0130 m, respectively. Our results also revealed that a simpler architecture could generally yield better results; however, having an attitude module dramatically improves performance in specific scenarios where the IMU experiences highly dynamic motion, highlighting the importance of evaluating solutions on diverse datasets.

7. RIOT Ablations

Model Dimensionality: We trained the model with differing dimensionality vector sizes from 56 to 896. Increasing the dimensionality of the model makes a small but measurable improvement up to 224. This finding aligns with the general principle in deep learning in which complexity reaches a point where passing it leads to overfitting, resulting in degraded performance on new data.

Encoder-Decoder Blocks: Increasing the number of encoder-decoder blocks did not result in a decrease in the model's perplexity. We trained three different models, with two, four and six blocks.

Attention Heads: We trained the model with two, three and four attention heads in each encoder-decoder block. There were small, almost immeasurable improvements in the

networks' performance on the test set. However, when applied to the unseen sequences, the models with three or four heads performed considerably worse. We believe increasing the number of heads past two forces the network into overfitting.

Window Size: We trained the model with differing window sizes from 50 to 500 (0.5 to 5 s). As the window size incrementally increased over 100, we saw better test set results but worse results on the unseen sequences. Increasing the window size of the input data exponentially increases the model complexity, as RIOT has 12 input features. The added complexity forces the network into overfitting.

8. Conclusions

This work proposes novel self-attention-based recursive neural network models, RIOT and ARIOT, for pose invariant inertial odometry. The proposed approaches utilise a sliding window as a hyperparameter to mitigate noise spikes and missing measurements. True position priors are included in the training process in conjunction with raw inertial measurements and ground truth displacement data, allowing for recursion and the ability to learn both motion characteristics and systemic error bias and drift. The evaluation results demonstrate that RIOT outperforms ARIOT and a GRU in terms of position error metrics, with a sequence length weighted mean Absolute Trajectory Error (ATE) of 0.0865 m and sequence length weighted mean Relative Trajectory Error (RTE) of 0.0091 m. These results are significantly better than the existing deep-learning inertial odometry methods in the literature, highlighting the effectiveness of the proposed approaches and learning methodology. Future work will consider the scalability of these approaches and make them local magnetic field agnostic.

Author Contributions: Conceptualisation, methodology, validation, investigation, software and writing—original draft preparation, J.B.; formal analysis, software and writing—review and editing, W.L.; formal analysis, writing—review and editing, A.D.G.; supervision, A.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research was undertaken with the assistance of resources from the National Computational Infrastructure (NCI Australia), an NCRIS enabled capability supported by the Australian Government (Grant No. LE160100051).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The publicly available dataset analysed in this work can be found here: [47].

Acknowledgments: The training was conducted in parallel on $4 \times$ Nvidia V100 GPUs, made possible with the assistance of resources and services from the National Computational Infrastructure (NCI), which is supported by the Australian Government.

Conflicts of Interest: The authors declare no conflict of interest.



Appendix A. Position Estimate Visulations from the First and Last Minute of Each Network

Figure A1. Position estimates from the first and last minute of each network as well as the true path for the first half of the unseen sequences, detailed in Section 6.5. The sequence run time is given in addition to each network's estimated total path length. The initial position is given to each network and is emphasised with individual markers. The final network estimate for each approach is also emphasised to easily visualise the positional drift over each sequence time period. (**a**) User 2: Total Time: 263.5 s, TRUE Distance: 220.9 m, ARIOT Distance: 234.1 m, RIOT Distance: 228.1 m, GRU Distance: 260.7 m. (**b**) User 3: Total Time: 258.0 s, TRUE Distance: 203.3 m, ARIOT Distance: 215.5 m, RIOT Distance: 209.7 m, GRU Distance: 229.3 m. (**c**) User 4: Total Time: 434.7 s, TRUE Distance: 364.3 m, ARIOT Distance: 130.3 m, ARIOT Distance: 138.6 m, RIOT Distance: 134.0 m, GRU Distance: 153.7 m. (**e**) Pocket: Total Time: 622.9 s, TRUE Distance: 493.7 m, ARIOT Distance: 526.3 m, RIOT Distance: 512.5 m, GRU Distance: 512.5 m, GRU Distance: 512.5 m, GRU Distance: 512.5 m, GRU Distance: 493.7 m, ARIOT Distance: 526.3 m, RIOT Distance: 493.7 m, ARIOT Distance: 526.3 m, RIOT Distance: 512.5 m, GRU Distance: 650.7 m. (**f**) Running: Total Time: 302.9 s, TRUE Distance: 393.1 m, ARIOT Distance: 422.2 m, RIOT Distance: 410.5 m, GRU Distance: 447.6 m.



Figure A2. Position estimates from the first and last minute of each network as well as the true path for the second half of the unseen sequences, detailed in Section 6.5. The sequence run time is given in addition to each network's estimated total path length. The initial position is given to each network and is emphasised with individual markers. The final network estimate for each approach is also emphasised to easily visualise the positional drift over each sequence time period. (**a**) Slow Walking: Total Time: 303.7 s, TRUE Distance: 166.6 m, ARIOT Distance: 174.7 m, RIOT Distance: 171.1 m, GRU Distance: 185.2 m. (**b**) Trolley: Total Time: 370.0 s, TRUE Distance: 328.4 m, ARIOT Distance: 346.5 m, RIOT Distance: 337.4 m, GRU Distance: 376.8 m. (**c**) Handbag: Total Time: 365.8 s, TRUE Distance: 305.9 m, ARIOT Distance: 321.9 m, RIOT Distance: 317.1 m, GRU Distance: 169.7 m, RIOT Distance: 159.8 m, GRU Distance: 358.6 m. (**e**) iPhone 5: Total Time: 183.8 s, TRUE Distance: 142.4 m, ARIOT Distance: 149.8 m, RIOT Distance: 145.3 m, GRU Distance: 158.3 m. (**f**) iPhone 6: Total Time: 173.4 s, TRUE Distance: 133.4 m, ARIOT Distance: 141.6 m, RIOT Distance: 137.4 m, GRU Distance: 150.3 m.

References

- El-Sheimy, N.; Hou, H.; Niu, X. Analysis and modeling of inertial sensors using Allan variance. *IEEE Trans. Instrum. Meas.* 2007, 57, 140–149. [CrossRef]
- Chen, S.A.B.S.; Billings, S.A. Neural networks for nonlinear dynamic system modelling and identification. Int. J. Control. 1992, 56, 319–346. [CrossRef]
- 3. Pascanu, R.; Mikolov, T.; Bengio, Y. On the difficulty of training recurrent neural networks. In Proceedings of the International Conference on Machine Learning, PMLR, Atlanta, GE, USA, 17–19 June 2013; pp. 1310–1318.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
- Khan, S.; Naseer, M.; Hayat, M.; Zamir, S.W.; Khan, F.S.; Shah, M. Transformers in vision: A survey. ACM Comput. Surv. (CSUR) 2022, 54, 1–41. [CrossRef]
- Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv 2018, arXiv:1810.04805.
- Mohammdi Farsani, R.; Pazouki, E. A transformer self-attention model for time series forecasting. J. Electr. Comput. Eng. Innov. (JECEI) 2021, 9, 1–10.
- Parmar, N.; Vaswani, A.; Uszkoreit, J.; Kaiser, L.; Shazeer, N.; Ku, A.; Tran, D. Image transformer. In Proceedings of the International Conference on Machine Learning. PMLR, Stockholm, Sweden, 10–15 July 2018; pp. 4055–4064.
- 9. Liu, P.J.; Saleh, M.; Pot, E.; Goodrich, B.; Sepassi, R.; Kaiser, L.; Shazeer, N. Generating wikipedia by summarizing long sequences. *arXiv* 2018, arXiv:1801.10198.
- Povey, D.; Hadian, H.; Ghahremani, P.; Li, K.; Khudanpur, S. A time-restricted self-attention layer for ASR. In Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Calgary, AB, Canada, 15–20 April 2018; pp. 5874–5878.
- 11. Huang, C.Z.A.; Vaswani, A.; Uszkoreit, J.; Shazeer, N.; Simon, I.; Hawthorne, C.; Dai, A.M.; Hoffman, M.D.; Dinculescu, M.; Eck, D. Music transformer. *arXiv* **2018**, arXiv:1809.04281.
- 12. Merkx, D.; Frank, S.L. Comparing Transformers and RNNs on predicting human sentence processing data. *arXiv* 2020, arXiv:2005.09471.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving Language Understanding by Generative Pre-Training. 2018. Available online: https://www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf (accessed on 10 November 2022).
- 14. Mahdi, A.E.; Azouz, A.; Abdalla, A.E.; Abosekeen, A. A machine learning approach for an improved inertial navigation system solution. *Sensors* **2022**, 22, 1687. [CrossRef] [PubMed]
- 15. Han, S.; Meng, Z.; Zhang, X.; Yan, Y. Hybrid deep recurrent neural networks for noise reduction of MEMS-IMU with static and dynamic conditions. *Micromachines* **2021**, *12*, 214. [CrossRef]
- Chen, H.; Aggarwal, P.; Taha, T.M.; Chodavarapu, V.P. Improving inertial sensor by reducing errors using deep learning methodology. In Proceedings of the NAECON 2018-IEEE National Aerospace and Electronics Conference, Dayton, OH, USA, 23–26 July 2018; pp. 197–202.
- Huang, F.; Wang, Z.; Xing, L.; Gao, C. A MEMS IMU gyroscope calibration method based on deep learning. *IEEE Trans. Instrum. Meas.* 2022, 71, 1–9. [CrossRef]
- Yan, H.; Shan, Q.; Furukawa, Y. RIDI: Robust IMU double integration. In Proceedings of the European Conference on Computer Vision (ECCV), Munich, Germany, 8–14 September 2018; pp. 621–636.
- Asraf, O.; Shama, F.; Klein, I. PDRNet: A deep-learning pedestrian dead reckoning framework. *IEEE Sens. J.* 2021, 22, 4932–4939. [CrossRef]
- Chen, C.; Lu, X.; Markham, A.; Trigoni, N. Ionet: Learning to cure the curse of drift in inertial odometry. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; Volume 32.
- 21. Yan, H.; Herath, S.; Furukawa, Y. Ronin: Robust neural inertial navigation in the wild: Benchmark, evaluations, and new methods. *arXiv* 2019, arXiv:1905.12853
- 22. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
- 23. Khorrambakht, R.; Lu, C.X.; Damirchi, H.; Chen, Z.; Li, Z. Deep Inertial Odometry with Accurate IMU Preintegration. *arXiv* 2021, arXiv:2101.07061.
- 24. Liu, W.; Caruso, D.; Ilg, E.; Dong, J.; Mourikis, A.I.; Daniilidis, K.; Kumar, V.; Engel, J. Tlio: Tight learned inertial odometry. *IEEE Robot. Autom. Lett.* 2020, *5*, 5653–5660. [CrossRef]
- Brotchie, J.; Li, W.; Kealy, A.; Moran, B. Evaluating Tracking Rotations Using Maximal Entropy Distributions for Smartphone Applications. *IEEE Access* 2021, 9, 168806–168815. [CrossRef]
- Sun, S.; Melamed, D.; Kitani, K. IDOL: Inertial deep orientation-estimation and localization. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtually, 2–9 February 2021; Volume 35, pp. 6128–6137.
- Wang, Y.; Cheng, H.; Wang, C.; Meng, M.Q.H. Pose-invariant inertial odometry for pedestrian localization. *IEEE Trans. Instrum. Meas.* 2021, 70, 1–12. [CrossRef]

- Cao, X.; Zhou, C.; Zeng, D.; Wang, Y. RIO: Rotation-equivariance supervised learning of robust inertial odometry. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 18–24 June 2022; pp. 6614–6623.
- Wang, Y.; Cheng, H.; Meng, M.Q.H. A2DIO: Attention-Driven Deep Inertial Odometry for Pedestrian Localization based on 6D IMU. In Proceedings of the 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23–27 May 2022; pp. 819–825.
- 30. Abiodun, O.I.; Jantan, A.; Omolara, A.E.; Dada, K.V.; Mohamed, N.A.; Arshad, H. State-of-the-art in artificial neural network applications: A survey. *Heliyon* **2018**, *4*, e00938. [CrossRef] [PubMed]
- Mohajerin, N.; Waslander, S.L. Multistep prediction of dynamic systems with recurrent neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* 2019, 30, 3370–3383. [CrossRef]
- Alzubaidi, L.; Zhang, J.; Humaidi, A.J.; Al-Dujaili, A.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M.A.; Al-Amidie, M.; Farhan, L. Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *J. Big Data* 2021, 8, 53. [CrossRef]
- Jiménez, A.R.; Seco, F.; Prieto, J.C.; Guevara, J. Indoor pedestrian navigation using an INS/EKF framework for yaw drift reduction and a foot-mounted IMU. In Proceedings of the 2010 7th Workshop on Positioning, Navigation and Communication, Dresden, Germany, 11–12 March 2010; pp. 135–143.
- Brotchie, J.; Shao, W.; Li, W.; Kealy, A. Leveraging Self-Attention Mechanism for Attitude Estimation in Smartphones. Sensors 2022, 22, 9011. [CrossRef] [PubMed]
- Crassidis, J.L.; Markley, F.L.; Cheng, Y. Survey of nonlinear attitude estimation methods. J. Guid. Control Dyn. 2007, 30, 12–28. [CrossRef]
- Coviello, G.; Avitabile, G.; Florio, A.; Talarico, C. A study on IMU sampling rate mismatch for a wireless synchronized platform. In Proceedings of the 2020 IEEE 63rd International Midwest Symposium on Circuits and Systems (MWSCAS), Springfield, MA, USA, 9–12 August 2020; pp. 229–232.
- Maas, A.L.; Hannun, A.Y.; Ng, A.Y. Rectifier nonlinearities improve neural network acoustic models. In Proceedings of the ICML Citeseer, Atlanta, GA, USA, 16–21 June 2013; Volume 30, p. 3.
- 38. Kok, M.; Hol, J.D.; Schön, T.B. Using inertial sensors for position and orientation estimation. arXiv 2017, arXiv:1704.06053.
- 39. Huynh, D.Q. Metrics for 3D rotations: Comparison and analysis. J. Math. Imaging Vis. 2009, 35, 155–164. [CrossRef]
- 40. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. arXiv 2014, arXiv:1412.6980.
- 41. Umeyama, S. Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. Pattern Anal. Mach. Intell.* **1991**, *13*, 376–380. [CrossRef]
- Yang, S.; Yu, X.; Zhou, Y. Lstm and gru neural network performance comparison study: Taking yelp review dataset as an example. In Proceedings of the 2020 International Workshop on Electronic Communication and Artificial Intelligence (IWECAI), Shanghai, China, 12–14 June 2020; pp. 98–101.
- Gruber, N.; Jockisch, A. Are GRU cells more specific and LSTM cells more sensitive in motive classification of text? *Front. Artif. Intell.* 2020, 3, 40. [CrossRef] [PubMed]
- 44. Cahuantzi, R.; Chen, X.; Güttel, S. A comparison of LSTM and GRU networks for learning symbolic sequences. *arXiv* 2021, arXiv:2107.02248.
- 45. Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; Liu, Y. Recurrent neural networks for multivariate time series with missing values. *Sci. Rep.* **2018**, *8*, 6085. [CrossRef]
- 46. Weber, D.; Gühmann, C.; Seel, T. RIANN—A robust neural network outperforms attitude estimation filters. *AI* 2021, *2*, 444–463. [CrossRef]
- 47. Chen, C.; Zhao, P.; Lu, C.X.; Wang, W.; Markham, A.; Trigoni, N. Oxiod: The dataset for deep inertial odometry. *arXiv* 2018, arXiv:1809.07491.
- 48. Vleugels, R.; Van Herbruggen, B.; Fontaine, J.; De Poorter, E. Ultra-Wideband Indoor Positioning and IMU-Based Activity Recognition for Ice Hockey Analytics. *Sensors* 2021, 21, 4650. [CrossRef] [PubMed]
- 49. Girbés-Juan, V.; Armesto, L.; Hernández-Ferrándiz, D.; Dols, J.F.; Sala, A. Asynchronous Sensor Fusion of GPS, IMU and CAN-Based Odometry for Heavy-Duty Vehicles. *IEEE Trans. Veh. Technol.* **2021**, *70*, 8617–8626. [CrossRef]
- Dey, S.; Schilling, A. A Function Approximator Model for Robust Online Foot Angle Trajectory Prediction Using a Single IMU Sensor: Implication for Controlling Active Prosthetic Feet. *IEEE Trans. Ind. Inform.* 2022, 19, 1467–1475. [CrossRef]
- 51. Vicon. Vicon Motion Capture Systems. 2017. Available online: https://www.vicon.com/?s=Motion%20Capture%20Systems (accessed on 1 November 2021).
- 52. Goodfellow, I.; Bengio, Y.; Courville, A. Deep Learning; MIT Press: Cambridge, MA, USA, 2016.
- Xue, W.; Qiu, W.; Hua, X.; Yu, K. Improved Wi-Fi RSSI Measurement for Indoor Localization. IEEE Sens. J. 2017, 17, 2224–2230. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.