

Article

DOES: A Deep Learning-based approach to estimate roll and pitch at sea.

Fabiana Di Ciaccio ^{1,*}, Paolo Russo ² and Salvatore Troisi ¹

¹ International PhD Programme/UNESCO Chair "Environment, Resources and Sustainable Development", Department of Science and Technology, Parthenope University of Naples, Centro Direzionale Isola C4, 80143, Naples, Italy; (fabiana.diciaccio, salvatore.troisi)@uniparthenope.it

² Department of Computer, Control and Management Engineering "Antonio Ruberti", University of Rome La Sapienza, via Ariosto 25, 00185, Rome, Italy; paolo.russo@diag.uniroma1.it

* Correspondence: fabiana.diciaccio@studenti.uniparthenope.it; +39-328-0935198 (F.D.C.)

Abstract: The use of Attitude and Heading Reference Systems (AHRS) for orientation estimation is now common practice in a wide range of applications, e.g., robotics and human motion tracking, aerial vehicles and aerospace, gaming and virtual reality, indoor pedestrian navigation and maritime navigation. The integration of the high-rate measurements can provide very accurate estimates, but these can suffer from errors accumulation due to the sensors drift over longer time scales. To overcome this issue, inertial sensors are typically combined with additional sensors and techniques. As an example, camera-based solutions have drawn a large attention by the community, thanks to their low-costs and easy hardware setup; moreover, impressive results have been demonstrated in the context of Deep Learning. This work presents the preliminary results obtained by DOES, a supportive Deep Learning method specifically designed for maritime navigation, which aims at improving the roll and pitch estimations obtained by common AHRS. DOES recovers these estimations through the analysis of the frames acquired by a low-cost camera pointing the horizon at sea. The training has been performed on the novel ROPIS dataset, presented in the context of this work, acquired using the FrameWO application developed for the scope. Promising results encourage to test other network backbones and to further expand the dataset, improving the accuracy of the results and the range of applications of the method.

Keywords: AHRS, Computer Vision, Dataset Acquisition, Deep Learning, Orientation Estimation.

1. Introduction

The pose estimation problem consists in estimating the position and orientation of a vehicle, device, human or robot with respect to a reference frame, through the use of different kinds of internal or external sensors. The accurate measurement of the orientation plays in fact a critical role in a wide range of activities, e.g., robotics and human motion tracking, bio-logging for animal behaviour research, aerial vehicles and aerospace, gaming and virtual reality applications, medicine and biotechnology, indoor and outdoor pedestrian navigation, maritime and/or autonomous navigation. When Global Navigation Satellite Systems (GNSS) are not able to provide correct information about the position and attitude of a vehicle, navigation and localization operations are generally performed through the integration of different kind of sensors: inertial, odometry, laser and sonar ranging sensors, underwater positioning systems, etc. [1].

In the last years the use of low-cost technologies is becoming widely spread in numerous applications: this means that the accuracy of the pose obtained by these systems can be affected by even more disturbing factors than the traditional high-performing methods. In these circumstances,

the development of accurate and reliable orientation estimation algorithms can still be considered a very challenging task, being at the basis of the localization process and of the consequent performances of the device employed for any specific task. This finds particular application in the context of the navigation, be it aerial, maritime or pedestrian, underwater/underground or in surface, autonomous, remotely operated or traditionally performed. In the specific case of maritime navigation, the information of position and attitude of a vessel is of great interest for seafarers in different operations and scenarios (e.g., open sea, congested harbours and waterways) as it is strictly related to the safety of the navigation at any level [2]. The same goes for Unmanned Surface Vehicles (USVs), which are mainly employed in environmental monitoring, safety or navigation support and research operations. In this case, a non accurate estimation of the orientation can severely compromise the ultimate success of the mission, especially when paired to low-cost sensors and poor GNSS support.

The Inertial Measurement Unit (IMU) gives the instantaneous speed and position of the vehicle without the need for external references by integrating the measures of angular velocity and linear acceleration obtained through its three orthogonal rate-gyroscopes and –accelerometers respectively. Unfortunately, several problems are associated with these sensors; among the others, measurements are noisy and biased and the errors increase over time due to the drift of the sensors. Micro Electro-Mechanical Systems (MEMS) Attitude Heading Reference Systems (AHRS) integrate to this configuration a magnetometer which measures the variation of the Earth's magnetic field: this allows to instantly calculate an improved estimation while benefitting from lighter weight, smaller sizes and lower prices. The great potential of these devices makes them suitable for several applications exploiting the pure orientation estimation, like geomatics, surveys, augmented reality, etc.

Vision-based methods are also frequently employed for the scope: these techniques allow to understand the surrounding environment by detecting its visual features through a camera; captured color data with its high resolution contains in fact several information, and the sensors are generally low-costs and with an easy hardware setup. In this context, the detection of the horizon line is an important attribute for the maritime image processing, as it allows to estimate the camera's orientation with respect to the sea surface other than restricting the object search region when detection is performed, thus reducing the processing time and the false detection problem. Several approaches have been proposed to solve this task, however the accuracy and the processing time of the horizon line detection on high-resolution maritime image still face some issues [3].

In the last decade, Visual Odometry (VO) and Visual Simultaneous Localization and Mapping (VSLAM) techniques have been successfully developed; however, their application can be challenging too, especially when deployed in non-textured environments or poor-light conditions. Visual Inertial Odometry (VIO) systems are proposed to eliminate these limitations, combining IMU and camera to improve motion tracking performance [4]. The current VIO systems heavily rely on manual interference to analyze failure cases and refine localization results, other than requiring careful parameters tuning procedures for the specific environment they have to work in. In recent years, Deep Learning (DL) has drawn significant attentions due to its potential in learning capability and its robustness to camera parameters and challenging environments. These data-driven methods have successfully learned new features representations from images that are used to further improve the motion estimation [5]. With the aim of providing further enhancements in the orientation estimation methodologies, this paper presents DOES, Deep Orientation (of roll and pitch) Estimation at Sea, a new supportive DL low-cost model which can be combined to the actual IMU-based configuration. The idea is to train a DL architecture able to output the correct vehicle attitude (in terms of roll and pitch angles) by processing the sea horizon view recorded by a low-cost camera. In particular, the latter needs to be mounted on the bridge of traditional ships, with its axis parallel to the vehicle longitudinal axis, or on the surface of an autonomous robot to correctly frame the horizon line. A similar approach could be further tested on Unmanned Aerial Vehicles (UAVs) too. To lay the foundation for this task, preliminary intensive tests have been conducted to verify the validity of the approach. Different DL architectures have been tested for the processing of the images acquired through an Android smartphone's camera.



Figure 1. Illustration of an image from the Roll and Pitch at Sea (ROPIS) dataset.

In this context, the lack of datasets specifically designed for DL-based orientation estimation at sea has been evidenced. While tackling this issue, the need of acquisition methods assuring the synchronism of the measurements for a reliable Ground Truth (GT) has been addressed too. For this reason, this paper presents also the first release of the Roll and Pitch at Sea (ROPIS) dataset (Fig. 1), which has been created through FrameWO, an Android application developed for the scope. The choice of employing low-cost sensors meets the necessity to develop affordable and smart tools to enhance the orientation estimation; for this reason, the first deployment of the dataset has been acquired using open-source libraries and software. In this preliminary release, the operating user acquires the data in the proximity of the seashore trying to simulate the real behaviour of a ship in navigation.

The main contribution of this work stands in the provision of a supportive low-cost technology aiming at improving the accuracy of the orientation estimation results in different approaches, without the need to configure camera models or considering related issues; the obtained results are promising and strongly encourage to work for further improvements.

The paper is organized as follows: Section 2 gives a brief overview on the existing literature on the orientation estimation task exploited through different traditional, visual and DL-based methods; Section 3 gives a theoretical foundation to the subject, introducing the attitude estimation problem to further describe the DL architectures which best fit the task. In Section 4 the ROPIS dataset will be presented, highlighting the issues and solutions encountered during the app creation and the data acquisitions. Section 5 details the experiments performed on DOES while the obtained results will be presented and discussed in Section 6; final considerations and future objectives will conclude the work in Section 7.

2. Related works

The accurate measurement of the orientation plays a critical role in a wide range of activities. AHRS sensors (i.e. accelerometers, gyroscopes and magnetometers) provide reliable measurements whose integration gives accurate information about the pose (position and attitude) of any object they are rigidly attached to. In the last decade, traditional methods have seen a huge improvement due to the integration with different kind of sensors, aiming at reducing the inertial-related error accumulation and the costs whilst enhancing the robustness of the methodology. As previously mentioned, one of the most effective integration is made through visual-based method, leveraging the potential of visual features and the low-cost of the devices. The following sections give a concise review of the existing literature in the field of orientation estimation.

2.1. Inertial-based methods

There exists a large amount of literature on the use of inertial sensors for position and orientation estimation. The reason for this is related to their robust algorithms and their accurate solutions which makes them suitable for being used in several fields. Interestingly, relatively simple position and orientation estimation algorithms work quite well in practice, even if the model choice can sensibly affect the accuracy of the estimates [6].

There is a large and ever-growing number of application areas for inertial sensors, as for example robotics and human motion tracking [7,8], bio-logging for animal behavior research [9], aerial vehicles and aerospace [10,11], gaming, virtual reality and indoor pedestrian navigation [12–14], etc. In fact, the use of accurate inertial sensors and magnetic compasses was first introduced in the navigation field, but along with the development of MEMS technology, low-cost and small-size inertial and magnetic compass sensors appeared in various kinds of consumer electronics, game consoles, virtual reality applications and so on. The orientation representations and sensor fusion still remain the challenges to overcome [15]. Real-time orientation estimation algorithms based on low-cost IMU are analyzed in [16], where the approach is based on the relationships between the quaternion representing the platform orientation and the measurements of the sensors and the integration is performed through an Extended Kalman Filter (EKF). Researchers in [17] developed a low-cost and low-weight attitude estimator for autonomous helicopter based on an inclinometer and a gyroscope, while fusing the data coming from the sensors through a classic complementary filter; in [18] a gyro-free, quaternion-based attitude determination system which exploits low cost sensors is proposed. [19] implemented a complementary filter able to infer Micro Aerial Vehicle (MAV) attitude from observations of gravity and magnetic field, with the final algorithm able to work with both IMU and MARG sensors. Authors in [11] exploited an AHRS device together with a Unscented Kalman Filter algorithm to perform attitude estimation on UAVs. The same filter has been used in [20], which developed a novel navigation system for autonomous underwater vehicles that works without the presence of a GPS device, not available in underwater scenarios. Researchers in [21] proposed an Adaptive Kalman Filter which is able to provide pose estimations based on low-cost AHRS devices, while [22] and [23] investigated the use of AHRS in smartphones as cheap but reliable devices for angles estimation. A novel error-state Kalman filter is presented in [24], which yields highly accurate estimates of IMU orientation that are robust to poor measurement updates from fluctuations in the local magnetic field and/or highly dynamic movements. An indoor pedestrian navigation method based on shoe-mounted MEMS IMU and ultra-wideband is discussed in [25], in which a quaternion-based Kalman Filter is used to integrate the data and to reduce the complexity of the method. In [26] a new orientation estimation strategy for a non-accelerated platform is presented: it is based on a low-cost IMU and the orientation angles are obtained through a nonlinear Luenberger observer, while the common offset issues on the magnetometer are calibrated by a recursive least-square algorithm. Authors in [27] utilize common bicycling motions to calibrate the 2D- and 3D-mounting orientation of a MEMS IMU on an electric bicycle. The method is independent of sensor biases and requires only a very low computation expense, so the estimation can be realized in real-time.

2.2. Vision-based methods

The possibility to employ visual data to perform orientation and in general pose estimation has been widely deepened in the past decades. Many researches have been focused on the horizon line detection, due to its relevance for visual geo-localization, port security, etc. However, some special features in real marine environments (e.g., clouds clutter, sea glint and weather conditions) frequently result in different kinds of interference in optical images. Authors in [28] propose a Sea-Sky Line (SSL) detection method for USVs based on the computation of the gradient saliency, through which the line features of the SSL are effectively enhanced while other disturbances are attenuated. The SSL identification is achieved according to regions contrast, line segment length and orientation features, and optimal state estimation of SSL detection is implemented by a cubature Kalman filter. In [29] a fast

method for detecting the horizon line in maritime scenarios is proposed. It combines a multi-scale approach and a region-of-interest (ROI) detection, which is an efficient way to reduce the amount of required processing information. The results are then combined to produce a single edge map on which the Hough transform and a least-square method are sequentially applied to accurately estimate the horizon line. The Hough transform is also used in [30], which proposed a sea-sky line detection system based on the local Otsu segmentation; similarly, authors in [31] recognize the horizon line in maritime images through a two-phase, coarse-fine detection algorithm which increases the overall method robustness. Another quick horizon line detection method is proposed in [32], which extracts the horizon line in real maritime image with improved reliability and faster execution with respect to other competitors. The horizon detection through vision sensors is also frequently exploited to obtain redundant orientation information in the field of unmanned aerial navigation. For example, authors in [33] proposed two attitude estimation methods: the first one searches for the best line fitting the horizon in thermal images, which allows to further estimate the pitch and roll angles using an infinite horizon line model. The second method exploits a Convolutional Neural Network (CNN) which predicts the angles on the basis of the raw pixel intensities from the same kind of images.

However, these methods alone cannot be considered totally robust and reliable, since the position and slope of the horizon are strictly related to the camera intrinsic (i.e., focal length, optical center, pixel aspect ratio and skew) and extrinsic (rotation and translation) parameters and to the model used to parametrize them. In [34] the authors surveyed a plethora of methods which perform pose estimation by fusing visual, inertial and magnetic measurements, integrating them through the use of an EKF. The combined use of IMU and vision information has been explored by [1], which exploits SURF visual features together with accelerometer and gyroscope data to retrieve the robot pose in an indoor setting. A comprehensive analysis of the behaviour of these features when used for visual odometry can be found in [35]. VO, VIO and SLAM algorithms have recently received much attention for their efficient and accurate ego-motion estimation in robotics. A VIO algorithm for the estimation of the motional state of UAVs with high accuracy is presented in [36]. It is based on the fusion of visual data and pre-integrated inertial measurements in a joint optimization framework and the on a stable initialization of scale and gravity using relative pose constraints. To account for the ambiguity and uncertainty of VIO initialization, a local scale parameter is adopted in the online optimization.

The use of stereo camera sensors for VO is a low-cost and effective way to estimate attitude, but may encounter problems in underwater setting due to poor imaging condition and inconsistent motion caused by water flow. Authors in [37] proposed a robust and effective stereo underwater VO system that can overcome the aforementioned difficulties and accurately localize the AUV. In the context of underwater robotics, another VO method designed to be robust to these visual perturbations is proposed in [38]: it demonstrated to outperform state-of-the-art SLAM methods under many of the most challenging conditions. A novel keyframe-based SLAM system with loop-closing and relocalization capabilities targeted for the underwater domain is presented in [39]. This paper addresses drift and loss of localization by providing a robust initialization method to refine scale using depth measurements and a fast preprocessing step to enhance the image quality. Authors in [40] presented a tightly coupled monocular VI-SLAM algorithm, which provides accurate and robust motion tracking at high frame rates on a standard CPU. A visual-inertial EKF is exploited to track the motion, then a globally consistent map is constructed to feed it back to the EKF state vector and reduce the drift. In a parallel thread, a global map is constructed to perform a keyframe-based visual-inertial bundle adjustment to optimize the map, together with a correction module to further eliminate the accumulated drift. ORB-SLAM3 [41] is another worth mentioning method, as it is the first system able to perform visual, visual-inertial and multi-map SLAM with monocular, stereo and RGB-D cameras, using pin-hole and fisheye lens models. It uses a feature-based tightly-integrated VI-SLAM system that fully relies on Maximum-a-Posteriori estimation, even during the IMU initialization phase, resulting in a system that operates robustly in real time, in small and large, indoor and outdoor environments, which is 2 to 5 times more accurate than previous approaches.

The rise of Deep Learning, with powerful architectures able to tackle complex tasks such as classification [42], detection [43], segmentation [44], denoising [45], super resolution [46], has definitely changed the way vision data is exploited for pose estimation. Instead of relying on engineered, fixed features (e.g. SIFT [47], SURF [48]), recent algorithms exploit deep networks as powerful features extractors or by directly estimating the pose vector in an end-to-end model, from input images to the output prediction. For example, in order to estimate camera orientation, authors of [49] exploited a LSTM deep network together with a linear Kalman Filter to combine IMU and camera data, while in DeepVIO [5] the authors fused 2D optical flow features together with standard inertial data, obtaining state of the art results on KITTI [50] and EuRoC [51] datasets. The combination of a traditional IMU with a LIDAR laser scan has been proposed by [52], which built a recurrent CNN to perform this aggregation on a scan-to-scan basis. In [53] researchers proposed a method to estimate a camera six degrees of freedom and absolute scale by exploiting unsupervised data, getting good results in terms of pose accuracy on KITTI benchmark. In the more recent work of Almalioglu et al [54], the authors developed a generative framework able to exploit a GAN [55] model on unlabelled RGB images for 6-DoF pose camera motion prediction, demonstrating the efficacy of their approach both on KITTI and Cityscapes [56] datasets. The former method has been improved by Feng et al. [57] with a stack of GAN layers which demonstrated to be effective on ego-motion estimation tasks. A comprehensive review of the state of the art deep models for pose estimation can be found in [58].

3. Method

This section aims at providing a theoretical background to fully understand the fundamentals of the proposed work. In particular, a general overview on the orientation estimation process is given in subsection 3.1, with some details on the sensors embedded in an AHRS and on the coordinate frame to which the smartphone device (and the related measures) is referred. Subsection 3.2 presents in a concise but detailed way the deep architecture models analysed and tested during the work.

3.1. Orientation estimation overview

The orientation of a rigid body is usually expressed by a transformation matrix in which the elements are generally parameterized in terms of Euler angles, rotation vectors, rotation matrices, and unit quaternions [59]. The Euler angles are the most intuitive expression as they allow a simple analysis of the body orientation in the 3D space. These angles are defined as follows:

- ϕ represents the rotation around the x axis (*roll* angle);
- θ defines the rotation around the y axis (*pitch* angle);
- ψ is related to the rotation around the z axis (*yaw* angle).

The integration of high-rate raw data acquired by the IMU sensors or of the more cost-effective AHRS is at the basis of the orientation estimation process. The accelerometer measures the acceleration in m/s^2 applied to a device, including the force of gravity: velocity is determined if the linear acceleration component is integrated once and position if the integration is performed twice. The results can be of poor accuracy due to the extensive noise and accumulated drift from which it suffers. The gyroscope measures the device rate of rotation (i.e. the angular velocity) in rad/s , from which the rotation angle can be calculated by integration. Gyroscopes run at a high rate, allowing them to track fast and abrupt movements, but they suffer from serious drift problems caused by the accumulation of measurement errors over long periods. Therefore, the fusion of both an accelerometer and gyroscope data is suitable to determine the pose of an object and to make up for the weakness of one over the other. The magnetometer measures the Earth's magnetic field in μT , which is helpful in heading determination; the drawback is that the presence of metallic objects within the environment could influence data collected through measurements. The drift introduced by the sensors causes errors accumulation: this means that the navigation information provided by the INS can be considered reliable and accurate only within short times, while it is still impossible for a pure inertial navigation

system to maintain the high-precision level throughout a mission. For this reason, the integration of the measurements provided by the three sensors aims at reducing the errors accumulation caused by the single one; this is generally made through filtering techniques and fusion methods. Moreover, information provided by external devices can considerably improve the accuracy of the estimations, especially when low-cost sensors could facilitate the process and make it more practical.

In this context, the objective of the present work is to provide a supportive mean to improve the attitude estimations obtained by common AHRS: DOES is a low-cost DL architecture developed to recover orientation information from the view of a camera pointing the horizon at sea, which will be placed on the bow of a navigating vehicle in future experiments. The training has been performed on the ROPIS dataset, acquired using an application developed for the scope on an Android smartphone which simultaneously collects the frames and calculates the corresponding Ground Truth data using the AHRS sensors.

The IMU-AHRS measurements of the smartphones are generally expressed in a custom body reference frame. The Android developer website defines its frame relative to the device's screen when the device is held in its default orientation (see Figure 2) [60]. In particular, the frame originates in the center of the device with the horizontal x axis pointing to the right, the vertical y axis pointing up and the z axis points toward the outside of the screen face, so that the the coordinates behind the screen have negative Z values. The related attitude information is then referred to the same coordinates.




Figure 2. Device coordinate system used by the Android Sensor API [60].

During the ROPIS dataset acquisition the smartphone has been kept in landscape mode, recording the horizon view. It has to be noticed that the coordinate frame does not change its definition, so in this setting the z axis points in the user direction, the y axis to his/her left and the x upwards.

3.2. Deep Learning architectures

DOES model is composed of a pre-trained backbone CNN and two additional Fully Connected (FC) layers to output the roll and pitch estimates. Several, well established architectures have been tested as backbone for the final network, as for example the VGG16-19 [61] and ResNet18-50-152 [62]; the resulting numerical comparison will be reported in Section 6, Tab. 3.

The VGG-16 and VGG-19 networks are based on the popular VGG architecture. They are composed of several convolutional layers followed by a Rectified Linear Unit (ReLU) activation function and interspersed by max pooling layers. Two FC layers are concatenated in order to produce the final features which are fed to a classification layer. These two networks differ only by the quantity and dimension of the convolutional layers employed, with a total number of parameters equal to 138M and 144M respectively. Despite being among the first developed deep architectures, with a huge amount of trainable parameters making them prone to overfitting, VGG models are still incredibly widespread, thanks to their ease of use for fine-tuning purposes on different tasks [63,64].




Figure 3. DOES architecture with default ResNet18 backbone network.

ResNet is a family of deep models based on the *residual* architecture. Differently from the VGG, the ResNet is made of a series of residual blocks in which the feature maps calculated by the convolutional layers are added to the input, so that each residual block calculates an *update* (hence residual) of the input feature maps. This approach makes the network resilient to the vanish gradient problem [65], improving convergence speed and the final accuracy result. Moreover, all the ResNet models avoid the use of the FC layers after the convolutional blocks, reducing the total number of trainable parameters and thus lessening the overfitting effect on training data. Authors of ResNet developed three versions with different number of layers (18, 50, 152) and with different number of visual features before the classification step (512 for the former, 2048 for the others). The number of free parameters for the 18, 50 and 152 layers models are 11M, 23M and 60M respectively.

In the experiments presented in this work, all the networks have been fine-tuned on the proposed ROPIS dataset starting from the ImageNet [66] pre-trained weights. The ResNet18 has been chosen among the others as the default DOES backbone since it produced the best accuracy while keeping at the same time a fast inference speed. Figure 3 reports the DOES network with the default ResNet18 backbone.

Two additional FC layers have been added as additional branches on top of the highest set of visual features in the backbone network to separately estimate the roll and pitch angles; for example, in the case of the ResNet models, this correspond to the global average pooling layer. Some different estimation procedures have been experimented, as the one described in [67]: it proposes to map the float angle value to a set of fixed bins, which then undergo a standard classification procedure with a final mapping back to the float value. However, in this work it has been experimentally found that this approach adds a layer of complexity without increasing the overall performances; this led to the decision to add a FC layer for each angle, which is able to accomplish the regression task with a good accuracy. Both the backbone network and the additional FC layers are jointly trained by back-propagation with the use of a standard Mean Square Error Loss (squared L2 norm). Two separated losses are calculated for each of the two angles, as reported in Eq. 1 for roll (L_{roll}) and Eq. 2 for pitch (L_{pitch}), where y and \hat{y} are the GT and predicted values respectively. The final loss L_{final} is then obtained as a simple addition of the aforementioned quantities, as shown in Eq. 3. The GT roll and pitch values have undergone a prior normalization process, which subtracts to each of them the mean and divides by the variance, both calculated over the entire dataset.

$$L_{roll}(y_{roll}, \hat{y}_{roll}) = \frac{1}{n} \sum_{i=1}^n (y_{roll} - \hat{y}_{roll}^i)^2 \quad (1)$$

$$L_{pitch}(y_{pitch}, \hat{y}_{pitch}) = \frac{1}{n} \sum_{i=1}^n (y_{pitch} - \hat{y}_{pitch}^i)^2 \quad (2)$$

$$L_{final} = L_{roll}(y_{roll}, \hat{y}_{roll}) + L_{pitch}(y_{pitch}, \hat{y}_{pitch}) \quad (3)$$

4. ROPIS data acquisition process

The lack of datasets designed for DL-based orientation estimation at sea lead to the necessity of searching for methods to acquire a set of data for the scope. In the following section, the development of the Android application and the obtained ROPIS dataset will be described in detail.

4.1. Device internal sensors and characteristics

In order to train the model, the dataset needs to contain a large amount of images showing the horizon and the corresponding GT data in terms of roll and pitch angles. The latter needs to be given with the best possible accuracy, as the learning process results will depend on it, which is strictly related to the instrumentation employed for the acquisition. With the aim of producing a low-cost and flexible solution, in this work the authors avoided the use of costly, high-end IMU devices and developed an Android application to acquire the dataset through a common smartphone. The presented ROPIS dataset in its first release has been totally collected through a OnePlus Nord smartphone, characterized by an average price and equipped with the most common sensors as reported in Table 1.

The OnePlus Nord mounts a BMI260 IMU, which contains a 16-bit tri-axial gyroscope (G) and accelerometer (A) providing fast, precise inertial sensing in smartphones and Human-Machine Interface (HMI) applications (i.e., advanced gesture, activity and context recognition, etc.). The IMU is characterized by a noise density of $160\mu g/\sqrt{Hz}$ (A) and $0.008dps/\sqrt{Hz}$ (G), a Zero-g/Zero-rate offset of $\pm 20mg$ (A) and $\pm 0.5dps$ (G) and an output data rate up to $1.6kHz$ (A) and $6.4kHz$ (G). Moreover, it mounts the industry’s first self-calibrating gyroscope with motionless Component Re-Trimming (CRT) functionality, which compensates MEMS typical soldering drifts, ensuring post-soldering sensitivity errors down to $\pm 0.4\%$ [68].

The MMC5603 is a monolithic complete 3-axis Anisotropic Magnetoresistance Effect (AMR) magnetic sensor with on-chip signal processing. It has an on-chip automatic degaussing with built-in SET/RESET function, allowing to eliminate thermal variation-induced offset error (Null field output) and to clear the residual magnetization resulting from strong external fields. It has a true frequency response up to $1KHz$ and can measure magnetic fields within the full scale range of $\pm 30Gauss$ (G) with $2mG$ total Root Mean Square (RMS) noise level, enabling heading accuracy of $\pm 1deg$ in electronic compass applications [69].

Table 1. OnePlus Nord smartphone general specifics [70].

General	Main Sensors	Rear Camera - Main
OS: OxygenOS Android™ 10	IMU: Bosch BMI260	Megapixels: 48
CPU: Qualcomm® Snapdragon™ 765G	Magn: MEMSIC MMC5603	Pixel Size: 0.8 μm/48M; 1.6 μm (4 in 1)/12M
GPU: Adreno 620	Camera: Sony IMX586	Lens Quantity: 6P
RAM: 8GB/12GB LPDDR4X	Proximity sensor	Aperture: f/1.75
Storage: 256GB UFS2.1	Ambient light sensor	OIS, EIS: Yes

The Sony IMX586 stacked CMOS image sensor is mounted as the main camera of the OnePlus Nord, and features 48 effective megapixels with an ultra-compact pixel size of $0.8\mu m$. The sensor uses the Quad Bayer color filter array, where adjacent 2×2 pixels come in the same color, making high-sensitivity shooting possible. During low light shooting, the signals from the four adjacent pixels are added, raising the sensitivity to a level equivalent to that of $1.6\mu m$ pixels (12 megapixels), resulting in bright, low noise images [71].

4.2. FrameWO application development

The FrameWO app has been developed in a free Open Source environment, the B4X suite, which supports the majority of PC, smartphones and embedding operating systems [72] (e.g., Android, iOS, Windows, MacOS, Linux, Arduino, RaspberryPI) and uses a modern version of Visual Basic as programming language. The Android version (B4A) allows to wrap existing Java code as an external

library and then to reference it from the B4A IDE, obtaining in release mode performances similar to those of Java. The size of a simple app is generally around 100 KB.

As previously mentioned, the necessary prerequisite for the dataset to meet the scope of this study is to associate to each frame the corresponding GT; however, the images size is much more larger than that of the IMU data, thus introducing a delay in their storage which affected their simultaneity. For this reason, the app captures the frames in YUV format (allowing for a better compression of the image) and converts them in JPEG only at the end of the process; this also avoids to run out of memory during the acquisition. A detailed overview on the YUV model can be found in [73]. Furthermore, several tests have been performed to determine an acquisition frequency value suitable for both the high-rate IMU data and the low-rate camera frames: the application offers in fact the possibility to set the camera acquisition frequency in *msec* to choose the best option for the needs.

As regards the GT, the API of Android [60] has been used to work on the raw measures read by the sensors and to obtain the Euler angles of interest. The *getRotationMatrix* function takes as input the gravity and geomagnetic field in vector form to compute the inclination matrix *I* and the rotation matrix *R*, transforming a vector from the device coordinate system to the world coordinate system (defined as a direct orthonormal basis). By definition, *R* is the identity matrix when the device is aligned with the world coordinate system (i.e., when the device *X* axis points toward East, the *Y* axis points to the North Pole and the device is facing the sky) and *I* is a simple rotation around the *X* axis transforming the geomagnetic vector into the same coordinate space as gravity, i.e., the world coordinate space (see Eq. 4, where *g* is the magnitude of gravity and *m* is the magnitude of the geomagnetic field).

$$\begin{aligned} \begin{bmatrix} 0 & 0 & g \end{bmatrix} &= R * gravity \\ \begin{bmatrix} 0 & m & 0 \end{bmatrix} &= I * R * geomagneticfield \end{aligned} \quad (4)$$

In order to isolate the gravity vector, a discrete-time low-pass filter with a smoothing factor $\alpha = 0.2$ has been applied to the accelerometer measurements. The Euler angles are recovered through the *getOrientation* function, which calculates them from the elements of the rotation matrix *R* [60,74].

The measurements are updated at the fastest rate provided by the Android API, which is in the order of few milliseconds. The time sampling has been set equal to 100msec, that means that 10 times in a second the device simultaneously registers the orientation and the corresponding image. As a final result, the data is saved in a directory named with the date and time of the specific acquisition, which is further renamed to specify the scenario characteristics of the moment. This directory contains all the frames, saved as *n_YYYY-MM-DD_HHMMSS.jpg*, and a *data.txt* file which lists the frame name, its index *n*, and the related GT.

4.3. Dataset structure

The ROPIS dataset in its first release has been mainly acquired in Italy, in the cities of Gaeta (Lazio) and Racale (Puglia). It consists of 22173 sRGB TrueColor JPEG images, with resolution set to 2592x1168, for a total dimension of 42.3 GB. Six different subsets have been acquired in as many locations, each presenting different characteristics in terms of scenarios and meteo-marine conditions; five of them have been chosen for the training set, from which a total of 100 frames has been separated for the validation set, and the last acquisition has been used as test set. The use of a dedicated test set with images coming from a separate location allows to verify the ability of DOES to generalize to new, different scenes with respect to the training and validation set. More in the specific, in each place eight different acquisitions have been made trying to simulate the behaviour of a ship in navigation in both static and dynamic conditions: this aims at emulating the induced oscillations which resemble the true motion of the ship. To improve the generalization ability of the model, the data has been acquired at

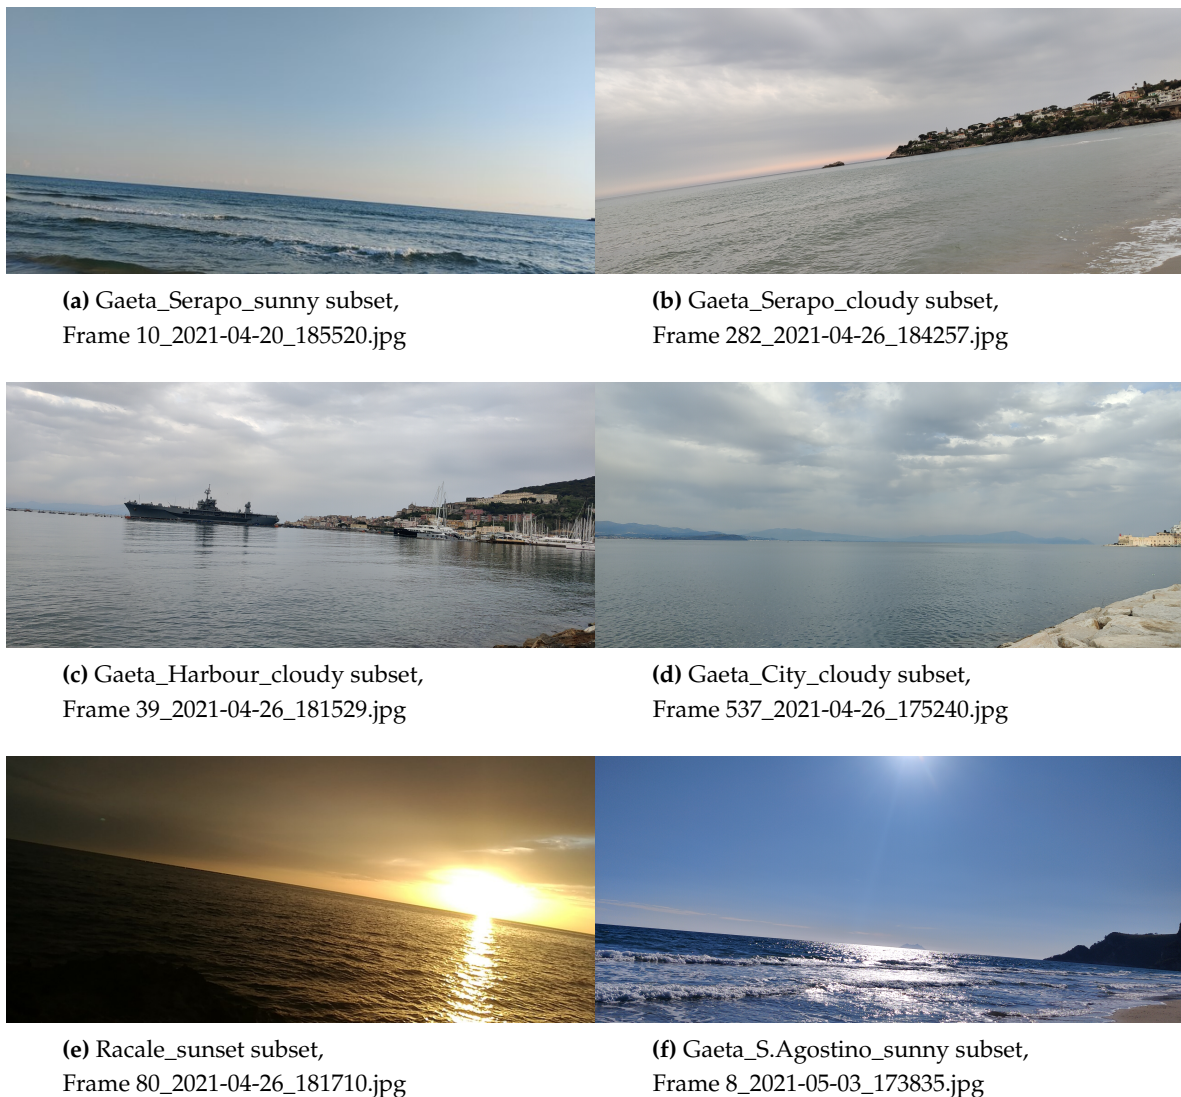


Figure 4. ROPIS dataset samples. Figures 4a to 4e belong to the training set, Figure 4f to the test set.

different day times and with sunny and cloudy sky; Figure 4 shows different samples of the ROPIS dataset.

Two aspects of this data need to be discussed:

- The point of view of the ROPIS images presents some differences with respect to the acquisitions taken on board the ship, since it adds parts of the land in the image foreground, such as sand, rocks, etc. However, this does not affect the learning procedure as the DL networks are able to recognize useful and useless image features, discarding the latter.
- A frame representing the real view from a navigating vehicle should depict some elements in the scene, such as the bow structures and some part of the bridge floor from a ship, or some of the USV sections. Although these specific features do not appear in ROPIS, DOES demonstrated its robustness to similar images cluttering present in the frames. Further experiments will be made to precisely assess their impact on the learning process.

For this reason, many other acquisitions will be made to improve the dataset structure to enhance the DOES capability to generalize over more complex and realistic data.

5. Experimental setup

In this section some details on the training process will be given, together with a brief overview of the evaluation metrics used to appraise the performance of DOES. Finally, the problem related to the comparison of DOES with other methods will be discussed.

5.1. Training details

DOES has been developed in Python programming language with the use of Pytorch framework; the code is publicly available¹. DOES has been trained using a standard fine-tuning procedure: the convolutional kernels were pre-trained on ImageNet while the additional FC layers have been initialized with random values drawn upon Pytorch default uniform distribution. Both convolutional and FC layers have been trained using the Adam [75] optimizer and a fixed learning rate set to 0.001. DOES has been trained on the ROPIS training set for a total of 10 epochs: it has in fact been noticed that a larger number of epochs led to an increase of the overfitting without any improvement of the accuracy. The images have been resized to a preliminary 512x288 resolution and then centrally cropped to obtain a final resolution of 224x224; a zero mean-unit variance normalization has been applied to both the images and the GT sets, with the corresponding mean and variance calculated over the specific training data.

The data augmentation process consisted of random changes in the colours of the images, using the *ColorJitter* transformation function of Pytorch which allows to set different values of brightness, contrast, saturation and hue: this resulted in an increase of the training dataset which further enhanced the generalization abilities of DOES. No random cropping nor image flipping have been applied during this process: in fact, the former would have caused the neglecting of the relative sea height information given by the images while the latter could have changed the correct roll angle perception by the network. The data augmentation procedure has naturally been deactivated during the testing phase, while the resize and center crop processes have been applied also to the test images; furthermore, the predicted roll and pitch values have been de-normalized before calculating the evaluation metrics presented in the following section 5.2. The selected data augmentation values (brightness and hue equal to 0.5, contrast and saturation equal to 5), as well as all the other training hyper-parameters, have been tuned on the validation set.

5.2. Evaluation metrics

DOES has been evaluated on the basis of the regression metrics implemented by the Scikit library in the *sklearn.metrics* module, which contains the most common utility functions to measure the regression performance.

The Mean Absolute Error (MAE) computes a risk metric corresponding to the expected value of the absolute error (Eq. 5); it is the average absolute difference between the predicted and the true value, expressed in the same scale as the data being measured. Each error contributes to MAE in proportion to its absolute value.

$$MAE(y, \hat{y}) = \frac{1}{n} \sum_{i=0}^{n-1} |y_i - \hat{y}_i| \quad (5)$$

The Root Mean Square Error (RMSE) represents the square root of the second sample moment of the differences between predicted values and the observed values (or the quadratic mean of these differences, also called residuals). It is a measure of accuracy and it is sensitive to outliers (Eq. 6). In fact, since the errors are squared before they are averaged, the RMSE gives a relatively high weight to large errors, making it more useful when large errors are particularly undesirable. RMSE does not

¹ <https://github.com/fabidicia/does>

necessarily increase with the variance of the errors, growing instead with the variance of the frequency distribution of error magnitudes.

$$RMSE(y, \hat{y}) = \frac{1}{n} \sqrt{\sum_{i=0}^{n-1} (y_i - \hat{y}_i)^2} \quad (6)$$

The Standard Deviation (STD) is a measure of the amount of dispersion (or variation) of the samples. A low standard deviation indicates that the values tend to be close to the mean μ (also called the expected value) of the set, while a high standard deviation indicates that the values are spread out over a wider range (Eq. 7).

$$\sigma(\hat{y}) = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \mu)^2} \quad (7)$$

Finally, the Median Absolute Error (MedAE) is calculated by taking the median of all the absolute differences between the GT and the prediction (Eq. 8). It is a non-negative floating point with best value of 0.0, robust to outliers since the median is not affected by values at the tails.

$$MedAE(y, \hat{y}) = median(|y_i - \hat{y}_i|, \dots, |y_n - \hat{y}_n|) \quad (8)$$

5.3. Methodology comparison

The comparison between DOES and other state of the art methods turned out to be a non trivial task for several reasons; among the others, the Deep Learning based solutions currently developed for the estimation of roll and pitch are either released without source code (as for example in [33]) or employed for very different tasks (e.g., head pose estimation [67]), thus making the comparison not properly correct or practically impossible. Generally speaking, traditional Horizon Line Detection (HLD) algorithms can be used as a proxy for this kind of estimations; the roll and pitch angles can in fact be correlated to the slope and position of the horizon line. However, as previously mentioned, this would require the correct knowledge of the intrinsic and extrinsic camera parameters and of the transformation matrix between the camera and the smartphone reference systems. To address this problem, a Linear Least Squares method has been applied to calibrate the HLD algorithms on the basis of the minimization of the squared error calculated between their output predictions and the GT values.

More in detail, given a set of measurements $M = [m_1, m_2 \dots m_n]$ and the corresponding set of ground truth values $[G = g_1, g_2 \dots g_n]$, the aim is to approximate the solution for the over-determined linear system (Eq. 9).

$$\begin{bmatrix} g_1 = x_2 + x_1 * m_1 \\ g_2 = x_2 + x_1 * m_2 \\ \vdots \\ g_n = x_2 + x_1 * m_n \end{bmatrix} \quad (9)$$

This system can be expressed in matrix form as in Eq. 10, where A is the known *design matrix* defined as $A = [M^T, 1^T]$, $B = G^T$ is the known target vector and $X = [x_1, x_2]$ is the solution of the Linear Least Square method. It represents the linear transformation (Eq. 11) which better minimizes the squared norm (Eq. 12).

$$AX - B = 0 \quad (10)$$

$$g = x_2 + x_1 * m \quad (11)$$

$$\frac{\|Ax - b\|^2}{2} \quad (12)$$

Two of the most renowned HLD algorithms by the scientific community have been selected to perform this comparison and are briefly described in the following lines.

The **Otsu** method [76] is a popular technique used to threshold the image between sky and non-sky regions. It is a reasonable fast and simple algorithm which performs fairly well on heterogeneous sets of data. The threshold value T is automatically computed by the algorithm through the assumption that the grayscale histogram of the image pixels intensities is bi-modal; the threshold is set so that the distance between the two histogram peaks is maximized.

Ettinger et al. [77] is a computer vision-based HLD algorithm that performs exhaustive search in the 2D line parameters space over the whole image looking at the best values which separate sky from terrain. However, being a slow algorithm on high resolution images, a modified version has been implemented that uses a two-stage objective: the *global* one searches for a narrow range of combinations of the pitch and roll horizon line angles corresponding to a half-plane that likely subdivides the sky from the rest of the image. The *local* one aims at searching exhaustively through these combinations to find the half-plane that maximizes the difference (in average intensity) of the two half-planes in their immediate vicinity. This method assumes that the sky pixels have higher intensity values than the ground pixels (higher mean), and that the sky has higher consistency of representation (lower variance).

6. Results and discussion

This section contains an assessment of the results provided by DOES. Table 2 shows DOES performances in terms of evaluation metrics with respect to the selected horizon line detection algorithms. DOES is able to achieve sensible better results both on roll and pitch angles, with a Mean Absolute Error lower than 2.0 degrees, as opposed to the other methods which exhibit worse performance with respect to all the indicators.

Table 2. DOES performances compared to those of the two HLD methods.

	DOES		Otsu [76]		Ettinger [77]	
	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>
MAE [deg]	1.84	1.89	4.48	3.76	4.04	3.77
RMSE [deg]	2.49	2.50	5.44	4.75	5.01	4.78
STD [deg]	1.68	1.64	3.09	2.90	2.97	2.93
MedAE [deg]	1.31	1.34	4.04	3.19	3.44	3.15

The MAE and the RMSE can be used together to diagnose the variation in the errors in a set of predictions. The RMSE is generally higher than the MAE, and the greater is the difference between them, the greater will be the variance in the individual errors of the samples; moreover, if the RMSE is close to the MAE, then all the errors are of the same magnitude. In the case of the current comparison, the small gap between RMSE and MAE demonstrates the ability of DOES to produce fewer outliers than Otsu and Ettinger. In addition, the STD values of the three methods show that the results obtained by DOES are significantly more clustered than the others, meaning that they are closer to the mean value and as such can be considered more reliable. The good performances of DOES are further confirmed by the MedAE value, which is sensibly lower than the counterparts. These findings can be summarized in Figure 5, which shows the MAE behaviour analysing the outputs percentage belonging to different MAE intervals (Fig. 5a) together with the empirical cumulative distribution (Fig. 5b) for the roll angle. The same evaluation can be made for the pitch angle (Fig. 6), which exhibits similar performances with respect to the roll. Another important consideration related to this comparison regards the inference time of DOES; the average estimation time on a single image is 100-150msec with




Figure 5. Graphical distribution of the errors for the estimation of the Roll angle.

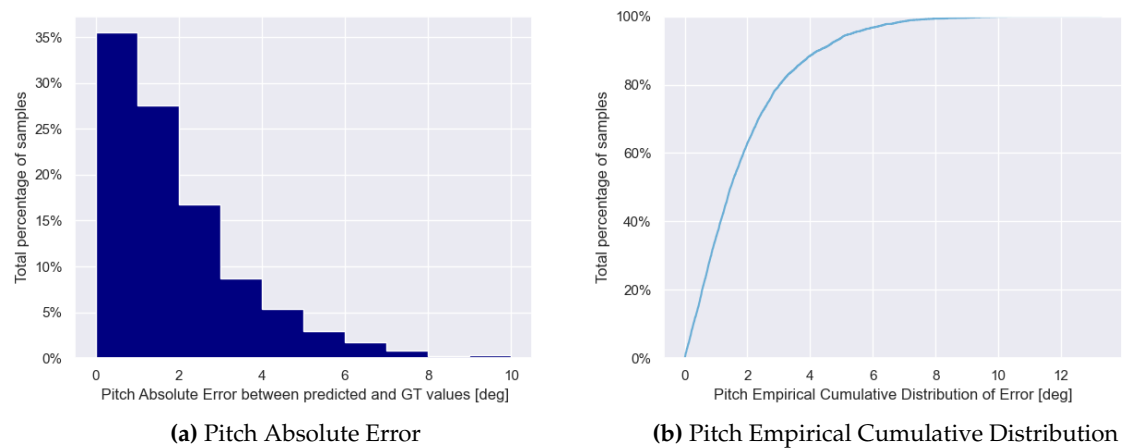


Figure 6. Graphical distribution of the errors for the estimation of the pitch angle.

any of the tested backbones, while Otsu and Ettinger inference time is comprised between 100 and 11000 msec, making them unsuitable for real-time applications on high-resolution images.

Table 3 shows a detailed comparison between DOES with its default proposed network and some alternative backbones: DOES is able to produce good performances with all the residual networks, while both VGG-19 and VGG-19bn struggle to produce reasonable results. More in detail, the MAE and RMSE results of ResNet18 are slightly better than the 50- and 152-layers versions, with the powerful DenseNet161 model able to produce a similar accuracy only on the roll angle. The performing results obtained by the ResNet18, together with the fastest training and inference speed due to the smaller number of trainable parameters (TP) with respect to the other architectures, make ResNet18 the first choice for the deployment of DOES as long as new models specifically developed for the scope will be released.

Finally, it is worth mention that a separated test has been made on a small subset of noisy images acquired rightly after the sunset and as such characterized by low-light conditions (see Fig. 7). Despite the absence of similar kind of samples in the training set, DOES obtained impressive results, with a MAE equal to 2.43° and a RMSE of 2.97° , confirming its capability to successfully work even in more challenging scenarios.

Table 3. Comparative results on different DOES backbones. *TP* indicates the number of trainable parameters.

	ResNet18 <i>TP = 11M</i>		ResNet50 <i>TP = 23M</i>		ResNet152 <i>TP = 58M</i>		VGG19 <i>TP = 139M</i>		VGG19bn <i>TP = 139M</i>		DenseNet161 <i>TP = 26M</i>	
	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>	<i>roll</i>	<i>pitch</i>
MAE [deg]	1.84	1.89	1.91	1.96	1.94	2.17	4.56	4.72	2.57	2.31	1.87	2.00
RMSE [deg]	2.49	2.50	2.58	2.60	2.68	2.79	5.50	5.70	3.30	3.01	2.56	2.61
STD [deg]	1.68	1.64	1.77	1.71	1.85	1.75	3.07	3.19	2.08	1.93	1.75	1.67
MedAE [deg]	1.31	1.34	1.35	1.49	1.37	1.77	4.11	4.31	2.09	1.87	1.35	1.58




Figure 7. A frame from the low-light condition separated set.

7. Conclusions

This paper presents a novel Deep Learning-based approach to the orientation estimation problem, which has been developed and intensively tested on a new dataset (the ROPIS dataset) specifically built for the scope and released in the context of this work. Deep Orientation (of roll and pitch) Estimation at Sea (DOES) is able to predict the attitude of the device in terms of roll and pitch angles by analysing the frames recorded by the camera pointing towards the sea horizon. DOES has been tested using several known architectures (e.g., ResNet152, ResNet18, VGG19) and with different configurations and hyper-parameters, obtaining excellent results. Unlike other visual-based methods, DOES is able to produce the output without the explicit knowledge of the camera intrinsic and extrinsic parameters or the distortions introduced by the camera lens. There is in fact no necessity to make any assumption on the use of specific models to parametrize the camera, since the model training only depends on the dataset given as input; the latter generally provides different sampling characteristics, thus making the network able to learn and then estimate the orientation regardless of the camera specifics. The ROPIS dataset has been created for this particular task and is here presented in its first release; the lack of public datasets suitable for DL applications made it necessary to search for a valid alternative for the experiments conduction. For this reason, an Android application has been developed using the Open Source B4A platform which will be made publicly available online. This app allows to simultaneously acquire the frames to be fed to the model as input, and the orientation estimations measured through the internal sensors of the smartphone, which will be used as Ground Truth in the training/testing phases.

ROPIS dataset is intended to be further improved by the introduction of more subsets of data collected in different scenarios (i.e., during the dusk/dawn, rainy days, etc) and environments (e.g., different cities coastlines, onboard of a vessels), using different acquisition devices. This will improve the DOES ability to generalize over heterogeneous data, making it even more invariant to the camera configurations, the acquisition condition and cluttering factors, thus providing better results in any kind of situation in which the vehicle will be navigating. In this regard, the authors wish to encourage

the users to download and test the FrameWO application with the aim of enhancing the ROPIS and its usage among the scientific community, to give a concrete contribution to this task.

The objective of this project is to develop a supportive technology to be integrated to the existing low-cost methodologies employed for the orientation estimation task. In fact, it has to be noticed that this approach has been specifically designed using affordable devices and applications and, as such, its results are not intended (at least in its preliminary version) to reach the accuracy provided by high-precision modern sensors. Further experiments will be made to test other light-weight DL architectures, which could be deployed on low-resources embedded hardware with the aim of providing better accuracy results in real-time applications on autonomous vehicles.

Author Contributions: Conceptualization, F.D.C and P.R.; methodology, F.D.C and P.R.; software, F.D.C and P.R.; validation, F.D.C, P.R. and S.T.; formal analysis, F.D.C and P.R.; investigation, F.D.C and P.R.; resources, F.D.C, P.R. and S.T.; data curation, F.D.C, P.R. and S.T.; writing—original draft preparation, F.D.C and P.R.; writing—review and editing, F.D.C, P.R. and S.T.; visualization, F.D.C, P.R. and S.T.; supervision, F.D.C, P.R. and S.T.; project administration, F.D.C, P.R. and S.T.; funding acquisition, F.D.C, P.R. and S.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Acknowledgments: The Authors would like to give a special thanks to Mr. Alberto Greco, which has been of fundamental importance in the development stage of the app employed to acquire the ROPIS dataset.

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

Abbreviations

The following abbreviations are used in this manuscript:

AHRS	Attitude and Heading Reference Systems
AMR	Anisotropic Magnetoresistance Effect
API	Application Programming Interface
CNN	Convolutional Neural Network
DL	Deep Learning
DOES	Deep Orientation (of roll and pitch) Estimation at Sea
EKF	Extended Kalman Filter
FC	Fully Connected
GNSS	Global Navigation Satellite Systems
GT	Ground Truth
HLD	Horizon Line Detection
HMI	Human-Machine Interface
IMU	Inertial Measurement Unit
INS	Inertial Navigation System
JPEG	Joint Photographic Experts Group
KF	Kalman Filter
MAE	Mean Absolute Error
MAV	Micro Aerial Vehicle
MedAE	Median Absolute Error
MEMS	Micro Electro-Mechanical Systems
ReLU	Rectified Linear Unit
RMS	Root Mean Square
RMSE	Root Mean Square Error
ROI	Region Of Interest
ROPIS	Roll and Pitch at Sea
SLAM	Simultaneous Localization and Mapping
SSL	Sea-Sky Line
STD	Standard Deviation
TP	Trainable Parameters
UAV	Unmanned Aerial Vehicles
USV	Unmanned Surface Vehicle
VIO	Visual Inertial Odometry
VO	Visual Odometry
VSLAM	Visual Simultaneous Localization and Mapping

References

1. Alatisse, M.B.; Hancke, G.P. Pose estimation of a mobile robot based on fusion of IMU data and vision data using an extended Kalman filter. *Sensors* **2017**, *17*, 2164.
2. Del Pizzo, S.; Gaglione, S.; Angrisano, A.; Salvi, G.; Troisi, S. Reliable vessel attitude estimation by wide angle camera. *Measurement* **2018**, *127*, 314–324.
3. Ganbold, U.; Akashi, T. The real-time reliable detection of the horizon line on high-resolution maritime images for unmanned surface-vehicle. 2020 International Conference on Cyberworlds (CW). IEEE, 2020, pp. 204–210.
4. Huang, G. Visual-inertial navigation: A concise review. 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019, pp. 9572–9582.
5. Han, L.; Lin, Y.; Du, G.; Lian, S. DeepVIO: Self-supervised deep learning of monocular visual inertial odometry using 3D geometric constraints. *arXiv preprint arXiv:1906.11435* **2019**.
6. Kok, M.; Hol, J.D.; Schön, T.B. Using Inertial Sensors for Position and Orientation Estimation. *Found. Trends Signal Process.* **2017**, *11*, 1–153. doi:10.1561/20000000094.

7. Avci, A.; Bosch, S.; Marin-Perianu, M.; Marin-Perianu, R.; Havinga, P. Activity recognition using inertial sensing for healthcare, wellbeing and sports applications: A survey. 23th International conference on architecture of computing systems 2010. VDE, 2010, pp. 1–10.
8. Luinge, H.J.; Veltink, P.H. Measuring orientation of human body segments using miniature gyroscopes and accelerometers. *Medical and Biological Engineering and computing* **2005**, *43*, 273–282.
9. Fourati, H.; Manamanni, N.; Afilal, L.; Handrich, Y. A nonlinear filtering approach for the attitude and dynamic body acceleration estimation based on inertial and magnetic sensors: Bio-logging application. *IEEE Sensors Journal* **2010**, *11*, 233–244.
10. Adler, S.; Schmitt, S.; Wolter, K.; Kyas, M. A survey of experimental evaluation in indoor localization research. 2015 International Conference on Indoor Positioning and Indoor Navigation (IPIN). IEEE, 2015, pp. 1–10.
11. De Marina, H.G.; Pereda, F.J.; Giron-Sierra, J.M.; Espinosa, F. UAV attitude estimation using unscented Kalman filter and TRIAD. *IEEE Transactions on Industrial Electronics* **2011**, *59*, 4465–4474.
12. Vertzberger, E.; Klein, I. Attitude adaptive estimation with smartphone classification for pedestrian navigation. *IEEE Sensors Journal* **2021**, *21*, 9341–9348.
13. Renaudin, V.; Combettes, C. Magnetic, acceleration fields and gyroscope quaternion (MAGYQ)-based attitude estimation with smartphone sensors for indoor pedestrian navigation. *Sensors* **2014**, *14*, 22864–22890.
14. Harle, R. A survey of indoor inertial positioning systems for pedestrians. *IEEE Communications Surveys & Tutorials* **2013**, *15*, 1281–1293.
15. Phuong, N.H.Q.; Kang, H.J.; Suh, Y.S.; Ro, Y.S. A DCM based orientation estimation algorithm with an inertial measurement unit and a magnetic compass. *Journal of Universal Computer Science* **2009**, *15*, 859–876.
16. Kim, A.; Golnaraghi, M. A quaternion-based orientation estimation algorithm using an inertial measurement unit. PLANS 2004. Position Location and Navigation Symposium (IEEE Cat. No. 04CH37556). IEEE, 2004, pp. 268–272.
17. Baerveldt, A.J.; Klang, R. A low-cost and low-weight attitude estimation system for an autonomous helicopter. Proceedings of IEEE International Conference on Intelligent Engineering Systems. IEEE, 1997, pp. 391–395.
18. Gebre-Egziabher, D.; Elkaim, G.H.; Powell, J.; Parkinson, B.W. A gyro-free quaternion-based attitude determination system suitable for implementation using low cost sensors. IEEE 2000. Position Location and Navigation Symposium (Cat. No. 00CH37062). IEEE, 2000, pp. 185–192.
19. Valenti, R.G.; Dryanovski, I.; Xiao, J. Keeping a good attitude: A quaternion-based orientation filter for IMUs and MARGs. *Sensors* **2015**, *15*, 19302–19330.
20. Allotta, B.; Caiti, A.; Costanzi, R.; Fanelli, F.; Fenucci, D.; Meli, E.; Ridolfi, A. A new AUV navigation system exploiting unscented Kalman filter. *Ocean Engineering* **2016**, *113*, 121–132.
21. Li, W.; Wang, J. Effective adaptive Kalman filter for MEMS-IMU/magnetometers integrated attitude and heading reference systems. *The Journal of Navigation* **2013**, *66*, 99–113.
22. Di Ciaccio, F.; Gaglione, S.; Troisi, S. A Preliminary Study on Attitude Measurement Systems Based on Low Cost Sensors. International Workshop on R3 in Geomatics: Research, Results and Review. Springer, 2019, pp. 103–115.
23. Michel, T.; Geneves, P.; Fourati, H.; Layaïda, N. On attitude estimation with smartphones. 2017 IEEE International Conference on Pervasive Computing and Communications (PerCom). IEEE, 2017, pp. 267–275.
24. Vitali, R.V.; McGinnis, R.S.; Perkins, N.C. Robust error-state Kalman filter for estimating IMU orientation. *IEEE Sensors Journal* **2020**, *21*, 3561–3569.
25. Wen, K.; Yu, K.; Li, Y.; Zhang, S.; Zhang, W. A new quaternion Kalman filter based foot-mounted IMU and UWB tightly-coupled method for indoor pedestrian navigation. *IEEE Transactions on Vehicular Technology* **2020**, *69*, 4340–4352.
26. Aligia, D.A.; Rocca, B.A.; De Angelo, C.H.; Magallán, G.A.; González, G.N. An orientation estimation strategy for low cost IMU using a nonlinear Luenberger observer. *Measurement* **2021**, *173*, 108664.
27. Schnee, J.; Stegmaier, J.; Lipowsky, T.; Li, P. Auto-correction of 3D-orientation of IMUs on electric bicycles. *Sensors* **2020**, *20*, 589.

28. Wang, B.; Su, Y.; Wan, L. A sea-sky line detection method for unmanned surface vehicles based on gradient saliency. *Sensors* **2016**, *16*, 543.
29. Jeong, C.Y.; Yang, H.S.; Moon, K. Fast horizon detection in maritime images using region-of-interest. *International Journal of Distributed Sensor Networks* **2018**, *14*, 1550147718790753.
30. Yongshou, D.; Bowen, L.; Ligang, L.; Jiucui, J.; Weifeng, S.; Feng, S. Sea-sky-line detection based on local Otsu segmentation and Hough transform. *Opto-Electronic Engineering* **2018**, *45*, 180039–180039.
31. Sun, Y.; Fu, L. Coarse-fine-stitched: A robust maritime horizon line detection method for unmanned surface vehicle applications. *Sensors* **2018**, *18*, 2825.
32. Praczyk, T. A quick algorithm for horizon line detection in marine images. *Journal of Marine Science and Technology* **2018**, *23*, 164–177.
33. Carrio, A.; Bavle, H.; Campoy, P. Attitude estimation using horizon detection in thermal images. *International Journal of Micro Air Vehicles* **2018**, *10*, 352–361.
34. Ligorio, G.; Sabatini, A.M. Extended Kalman filter-based methods for pose estimation using visual, inertial and magnetic sensors: Comparative analysis and performance evaluation. *Sensors* **2013**, *13*, 1919–1941.
35. Chien, H.J.; Chuang, C.C.; Chen, C.Y.; Klette, R. When to use what feature? SIFT, SURF, ORB, or A-KAZE features for monocular visual odometry. 2016 International Conference on Image and Vision Computing New Zealand (IVCNZ). IEEE, 2016, pp. 1–6.
36. Hong, E.; Lim, J. Visual-inertial odometry with robust initialization and online scale estimation. *Sensors* **2018**, *18*, 4287.
37. Zhang, J.; Ila, V.; Kneip, L. Robust visual odometry in underwater environment. 2018 OCEANS-MTS/IEEE Kobe Techno-Oceans (OTO). IEEE, 2018, pp. 1–9.
38. Ferrera, M.; Moras, J.; Trouvé-Peloux, P.; Creuze, V. Real-time monocular visual odometry for turbid and dynamic underwater environments. *Sensors* **2019**, *19*, 687.
39. Rahman, S.; Li, A.Q.; Rekleitis, I. Svin2: an underwater slam system using sonar, visual, inertial, and depth sensor. 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2019, pp. 1861–1868.
40. Quan, M.; Piao, S.; Tan, M.; Huang, S.S. Accurate monocular visual-inertial SLAM using a map-assisted EKF approach. *IEEE Access* **2019**, *7*, 34289–34300.
41. Campos, C.; Elvira, R.; Rodríguez, J.J.G.; Montiel, J.M.; Tardós, J.D. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multimap SLAM. *IEEE Transactions on Robotics* **2021**.
42. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely connected convolutional networks. Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 4700–4708.
43. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask r-cnn. Proceedings of the IEEE international conference on computer vision, 2017, pp. 2961–2969.
44. Russo, P.; Tommasi, T.; Caputo, B. Towards Multi-source Adaptive Semantic Segmentation. International Conference on Image Analysis and Processing. Springer, 2019, pp. 292–301.
45. Russo, P.; Di Ciaccio, F.; Troisi, S. DANA++: A Smart Approach for Denoising Underwater Attitude Estimation. *Sensors* **2021**, *21*, 1526.
46. Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; Change Loy, C. Esrgan: Enhanced super-resolution generative adversarial networks. Proceedings of the European conference on computer vision (ECCV) workshops, 2018, pp. 0–0.
47. Lowe, D.G. Object recognition from local scale-invariant features. Proceedings of the seventh IEEE international conference on computer vision. Ieee, 1999, Vol. 2, pp. 1150–1157.
48. Bay, H.; Tuytelaars, T.; Van Gool, L. Surf: Speeded up robust features. European conference on computer vision. Springer, 2006, pp. 404–417.
49. Rambach, J.R.; Tewari, A.; Pagani, A.; Stricker, D. Learning to fuse: A deep learning approach to visual-inertial camera pose estimation. 2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR). IEEE, 2016, pp. 71–76.
50. Geiger, A.; Lenz, P.; Stiller, C.; Urtasun, R. Vision meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)* **2013**.
51. Burri, M.; Nikolic, J.; Gohl, P.; Schneider, T.; Rehder, J.; Omari, S.; Achtelik, M.W.; Siegwart, R. The EuRoC micro aerial vehicle datasets. *The International Journal of Robotics*

- Research **2016**, [<http://ijr.sagepub.com/content/early/2016/01/21/0278364915620033.full.pdf+html>]. doi:10.1177/0278364915620033.
52. Li, C.; Wang, S.; Zhuang, Y.; Yan, F. Deep sensor fusion between 2d laser scanner and imu for mobile robot localization. *IEEE Sensors Journal* **2019**.
 53. Li, R.; Wang, S.; Long, Z.; Gu, D. Undeepvo: Monocular visual odometry through unsupervised deep learning. 2018 IEEE international conference on robotics and automation (ICRA). IEEE, 2018, pp. 7286–7291.
 54. Almalioglu, Y.; Saputra, M.R.U.; de Gusmao, P.P.; Markham, A.; Trigoni, N. Ganvo: Unsupervised deep monocular visual odometry and depth estimation with generative adversarial networks. 2019 International conference on robotics and automation (ICRA). IEEE, 2019, pp. 5474–5480.
 55. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems* **2014**, 27.
 56. Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; Schiele, B. The cityscapes dataset for semantic urban scene understanding. Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 3213–3223.
 57. Feng, T.; Gu, D. Sganvo: Unsupervised deep visual odometry and depth estimation with stacked generative adversarial networks. *IEEE Robotics and Automation Letters* **2019**, 4, 4431–4437.
 58. Zhao, C.; Sun, Q.; Zhang, C.; Tang, Y.; Qian, F. Monocular depth estimation based on deep learning: An overview. *Science China Technological Sciences* **2020**, pp. 1–16.
 59. Bernal-Polo, P.; Barberá, H. Orientation Estimation by Means of Extended Kalman Filter, Quaternions, and Charts. *Journal of Physical Agents* **2017**, 8. doi:10.14198/JoPha.2017.8.1.03.
 60. SensorManager.java. https://developer.android.com/guide/topics/sensors/sensors_overview. Accessed: 2021-11-06.
 61. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* **2014**.
 62. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
 63. He, Y.; Zhu, C.; Wang, J.; Savvides, M.; Zhang, X. Bounding box regression with uncertainty for accurate object detection. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 2888–2897.
 64. Long, J.; Shelhamer, E.; Darrell, T. Fully convolutional networks for semantic segmentation. Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 3431–3440.
 65. Veit, A.; Wilber, M.J.; Belongie, S. Residual networks behave like ensembles of relatively shallow networks. *Advances in neural information processing systems* **2016**, 29, 550–558.
 66. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. 2009 IEEE conference on computer vision and pattern recognition. Ieee, 2009, pp. 248–255.
 67. Ruiz, N.; Chong, E.; Rehg, J.M. Fine-grained head pose estimation without keypoints. Proceedings of the IEEE conference on computer vision and pattern recognition workshops, 2018, pp. 2074–2083.
 68. BMI260: IMU combining accelerometer and gyroscope. <https://www.bosch-sensortec.com/products/motion-sensors/imus/bmi260/>. Accessed: 2021-10-01.
 69. Monolithic, High Performance, Low Cost 3-axis Magnetic Sensor. <http://www.memsic.com/uploadfiles/2020/08/20200827165137254.pdf>. Accessed: 2021-10-01.
 70. OnePlus Nord - Specs. <https://www.oneplus.com/uk/nord-specs>. Accessed: 2021-10-01.
 71. Sony Releases Stacked CMOS Image Sensor for Smartphones. <https://www.sony.com/en/SonyInfo/News/Press/201807/18-060E/>. Accessed: 2021-10-01.
 72. Simple, powerful and modern development tools. <https://www.b4x.com/>. Accessed: 2021-10-04.
 73. Podpora, M.; Korbas, G.P.; Kawala-Janik, A. YUV vs RGB-Choosing a Color Space for Human-Machine Interaction. FedCSIS (Position Papers), 2014, pp. 29–34.
 74. Sensors Overview. https://github.com/aosp-mirror/platform_frameworks_base/blob/master/core/java/android/hardware/SensorManager.java. Accessed: 2021-11-06.
 75. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* **2014**.
 76. Otsu, N. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics* **1979**, 9, 62–66.

77. Ettinger, S.M.; Nechyba, M.C.; Ifju, P.G.; Waszak, M. Vision-guided flight stability and control for micro air vehicles. *Advanced Robotics* **2003**, *17*, 617–640.

Sample Availability: The dataset is available from the authors.