

Full Scaled 3D Visual Odometry from a Single Wearable Omnidirectional Camera

Daniel Gutiérrez-Gómez, Luis Puig, J.J. Guerrero

*Aragon Institute of Engineering Research (I3A)
University of Zaragoza, Mariano Esquillor s/n, 50018, Zaragoza, Spain.
Tel. +34-976842472, Fax +34-976761914, e-mail: danielgg@unizar.es*

Abstract

In the last years monocular SLAM has been widely used and tested to obtain highly accurate maps and trajectory estimations of a moving camera. However one of the issues of this methodology is that, due to the impossibility of the depth being measured in a single image, global scale is not observable and scene and camera motion can only be recovered up to scale. This problem gets aggravated as we deal with larger scenes since it is more likely that scale drift arises between different map portions and their corresponding motion estimates. To compute the absolute scale we need to know some kind of dimension of the scene (*e.g.* real size of an element of the scene, velocity of the camera or baseline between two frames) and somehow integrate it in the SLAM estimation. In this paper, we present a method to recover the scale of the scene using an omnidirectional camera mounted on a helmet carried by a person. The high precision of visual SLAM allows the head vertical oscillation during walking to be perceived in the trajectory estimation. By performing a spectral analysis on the camera vertical displacement, we can measure the step frequency. We relate the step frequency to the speed of the camera by an empirical formula based on biomedical experiments on human walking. This velocity measurement is integrated in a Particle Filter to estimate the current scale factor and the 3D motion estimation in real size. We tested our approach on real world image sequences and it has shown to provide a very accurate estimate of the real scaled trajectory being able to cope with scale drift.

Keywords: Omnidirectional vision, catadioptric camera, Visual SLAM, feature matching, Extended Kalman Filter (EKF), spectral analysis, Discrete Fourier Transform (DFT).

1 Introduction

Visual SLAM has become one of the most trending research fields in the last years and has been addressed both by using stereo and monocular systems. The main feature of stereo systems is that, knowing the baseline of the cameras, detected landmarks of the scene can be triangulated and the visual odometry and landmark positions can be completely estimated. SLAM approaches using stereo systems have been presented in (Nistér, Naroditsky, and Bergen, 2006; Paz et al., 2008; Mei et al., 2010).

On the other hand, due to the impossibility to extract the depth of a landmark just from one single image, monocular systems only allow the camera motion and scene to be estimated up to an unknown scale. With this in mind, stereo systems may seem more appropriate than monocular ones to perform visual SLAM. However the use of single cameras for visual SLAM is still appealing since they are cheaper, more compact and easier to calibrate than stereo systems.

One of the most important and successful works on monocular SLAM is the one developed by Davison (Davison, 2003). This approach uses a pattern of known size to initialize some feature locations allowing the SLAM to start. Thus the scale of the map is fixed by the size of this initial pattern. Later, the inverse depth parametrization for the map points (Civera, Davison, and Montiel, 2008) allowed the SLAM to start automatically without the need of using an initialization pattern. In this case the scale is arbitrarily fixed by a depth prior for the map features and the acceleration noise setup parameter.

Although the scale can be initially set up (*e.g.* either by a pattern of known size or some kind of prior) it is likely that scale drift arises between different portions of the scene as the size of the map gets larger. This drift acts as a source of incremental error in the SLAM estimation, which leads to a deformation of the final map even after applying loop closing techniques. In (Strasdat, Montiel, and Davison, 2010), Strasdat *et al.* propose a loop closing method which corrects the map deformation due to scale drift.

Visual SLAM using omnidirectional cameras has been proposed in (Corke, Strelow, and Singh, 2004; Tardif, Pavlidis, and Daniilidis, 2008; Mei, 2007; Rituerto, Puig, and Guerrero, 2010). Due to the 360° field of view (FoV) of omnidirectional cameras, features last longer on the image than in the case of conventional cameras, specially during big camera rotations. The increased lifespan of the features on the image translates in a better estimation of the position of the features on the map, a lower need to initialise new features and an increased robustness .

In this work we extend our SLAM approach for catadioptric cameras (Gutierrez et al., 2011) which derives from state of the art EKF monocular SLAM for conventional cameras (Civera et al., 2010). This approach is used to compute the visual odometry from sequences of images acquired with a catadioptric camera mounted on a helmet (Fig. 1).

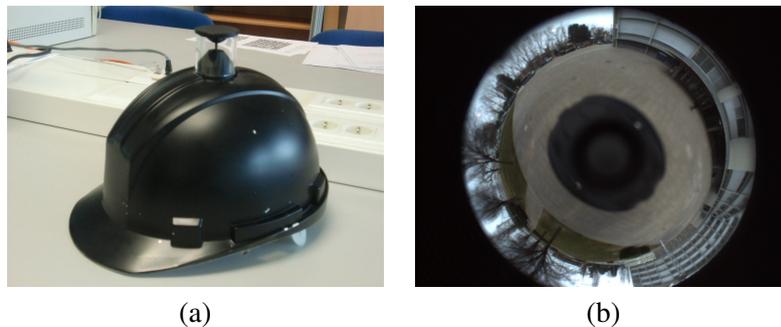


Figure 1: (a) Hemlet-camera device used in our experiments. (b) Omnidirectional image captured with our device.

One induced effect of walking on humans is a head vertical oscillation whose frequency matches up with the step frequency (Hirasaki et al., 1999). Our work is based on the premise that if the 6 d.o.f. visual SLAM is accurate enough , this vertical oscillatory motion of the head should be observed. Fig. 2a depicts an example of this behaviour, where the camera trajectory was obtained by performing a Visual SLAM algorithm. Hence the step frequency of the camera carrier could be measured by estimating the power spectra of the vertical component of the camera trajectory (Fig. 2b).

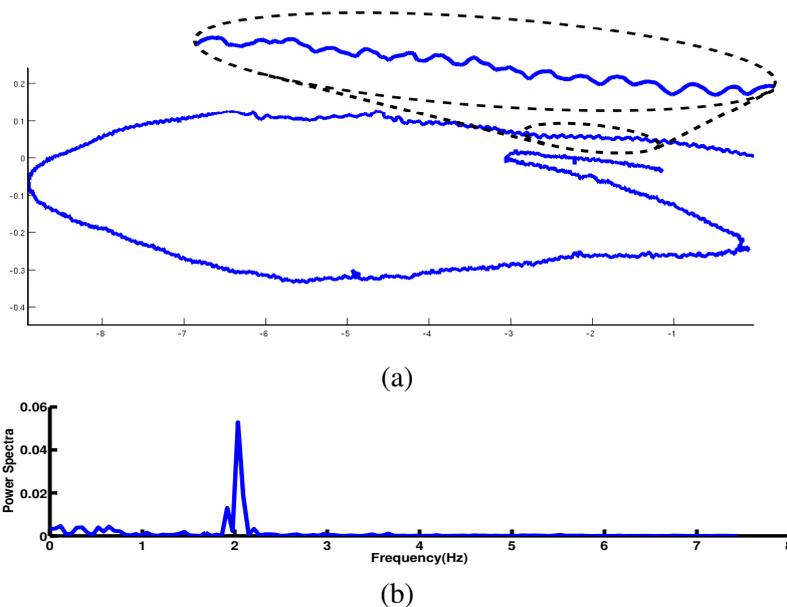


Figure 2: (a) Trajectory estimation of Visual SLAM from a head-mounted catadioptric camera. (b) Power spectra of the vertical component.

Walking speed is strictly calculated as the product of step frequency and stride length. However there exists biomedical studies like the one lead by Grieve (Grieve, 1968), which show an empirical relation between step frequency and the walking speed with no dependence on the stride length. Further studies explain this relation as the result of a human tendency to choose a step frequency that minimizes metabolic cost of locomotion at a given walking speed (Zarrugh, Todd, and Ralston, 1974; Kuo, 2001).

Based on this, we propose an approach to calculate the scale of the visual odometry from a single omnidirectional camera carried on the head of a person. This is done by first performing spectral analysis on short sections of the trajectory to extract the step frequency. Then we compute the estimated walking velocity using the relation between step frequency and walking speed, and finally this estimation is integrated into a particle filter which recursively computes the scale factor.

The paper is structured as follows. Related work on scale estimation with monocular vision systems is described in Section 2. In Section 3 we explain in detail our approach for computing the scale factor. Section 4 present the experimental results and finally in Section 5 we extract the conclusions.

2 Related work

The literature on scale estimation in monocular SLAM shows two ways to address this problem.

In some works additional sensors are used to compute the real scale of the SLAM trajectory and map. In (Civera et al., 2010), Civera *et al.* use GPS information to align and scale the SLAM estimation by a rigid transformation which minimizes the distance between corresponding tra-

jectory points. However this approach was designed for a benchmarking purpose and has little practical utility. Lupton *et al.* (Lupton and Sukkarieh, 2008) and Nützi *et al.* (Nützi et al., 2010) use an IMU to resolve the scale. The former makes the true map scale observable by integrating the visual data and the IMU data within an information filter. This allows the computation of the true map and trajectory estimations with no bias due to acceleration noise and feature depth priors. The latter fuses the SLAM estimation and IMU data in an EKF framework to compute the scale factor. In (Eudes et al., 2010), Eudes *et al.* integrate the odometry and visual information to obtain the scaled map. In (Scaramuzza, Fraundorfer, and Siegwart, 2009a) Scaramuzza *et al.* use the vehicle speed measurement to compute the distance between the last two frames and recover the 3D structure by triangulation of the common image points.

In other works the scale of the scene is estimated without additional sensors. As commented before, it is possible to fix the scale with the size of the pattern used to initialize the landmarks needed to start the SLAM. However it does not avoid scale drift in large maps. Loethe *et al.* (Loethe et al., 2010) use the prior knowledge of the distance of the camera to the ground plane to compute the scale factor of the scene, which is well suited for camera mounted on vehicles. Scaramuzza *et al.* (Scaramuzza et al., 2009b) exploit non-holonomic motion constraints of wheeled vehicles to resolve the absolute scale, although this is only possible when the vehicle turns.

The approach presented in this paper fits in the second category, since no additional sensors are used. It takes advantage of the vertical oscillatory motion of a head-mounted camera to extract the step frequency. The prior information in this case is an empirical formula which allows us to compute the walking speed from the extracted step frequency. The computed walking speed is integrated in a particle filter framework used to smoothly estimate the scale factor rejecting possible spurious estimations of the walking speed.

3 Problem formulation

The V-SLAM approach used in this work is based on the Extended Kalman Filter (EKF) which is divided in two parts. In the first part, *Prediction*, the new state of the system is estimated from the previous time step state through the motion model. The second part of the algorithm, *Update*, uses the measurements of the environment to improve the new state prediction. The full state vector, composed of both the map and last camera location, is modelled as a multidimensional Gaussian distribution coded by its mean vector and covariance matrix.

The state of the system is given by the state vector \mathbf{x}

$$\mathbf{x} = \left(\underbrace{\mathbf{r}, \mathbf{q}, \mathbf{V}, \boldsymbol{\omega}}_{\text{Camera state}}, \underbrace{x_i, y_i, z_i, \theta_i, \phi_i, \rho_i, \dots}_{\text{3D points (IDP)}}, \underbrace{X_j, Y_j, Z_j, \dots}_{\text{3D points}} \right) \quad (1)$$

where $\mathbf{r}_{(3 \times 1)}$ is the camera pose, $\mathbf{q}_{(4 \times 1)}$ is the quaternion of its orientation and $\mathbf{V}_{(3 \times 1)}$ and $\boldsymbol{\omega}_{(3 \times 1)}$ are its linear and angular velocities, respectively. The size of the part of the state vector encoding the location of the landmarks depends on their parameterisation, which in turn depends on the depth uncertainty. Features with large depth uncertainty are parameterised using inverse depth parameterisation (IDP) (Civera, Davison, and Montiel, 2008). As the depth of the landmarks cannot

be measured from one single image they are initialised with an arbitrary depth prior of ρ_{0i} with large uncertainty. In successive observations of the feature, depth estimation is gradually refined. If the depth uncertainty of a feature decreases under a certain threshold then the state of that feature is given by its cartesian coordinates in the world reference frame.

The inability to measure the initial depth of the features involves the unobservability of the absolute scale of the scene. The scale of the SLAM reconstruction is biased due to the difference between the arbitrary depth prior ρ_{0i} and the true depth of the first measured landmarks. Another source of the drift is that usually old landmarks are lost and new ones are initialized.

In this work we perform SLAM with an omnidirectional camera on a hemlet carried by a human operator. Since the head moves vertically with the step frequency during walking, SLAM visual odometry can be used to estimate the step frequency by spectral analysis.

To adress the problem of determining the real scale dynamically, we propose a method which is performed iteratively on trajectory sections and can be subdivided in four steps:

- Spectral analysis on the SLAM visual odometry for step frequency estimation.
- Empirical estimation of walking speed from step frequency.
- Integration of walking speed in a particle filter for recursive scale factor estimation.
- Scaling of the final visual odometry.

3.1 Spectral analysis on SLAM visual odometry

In the case of our omnidirectional camera, the camera frame is oriented with its z -axis pointing approximately to the direction of the normal to the ground plane, so the head vertical oscillation is given by the z -component of the camera position vector. To estimate the step frequency we apply spectral analysis to the z -component of the trajectory. The trajectory is divided in sections of N points and for each section the data sequence (z_1, z_2, \dots, z_N) is considered, where z_n is the z -component of the n -th trajectory point of the current analyzed section.

Prior to computing the power spectra, the data sequence is filtered to eliminate the low frequency harmonics introduced by the change of the ground height along the path.

The power spectra Γ is calculated by applying the Discrete Fourier Transform (DFT) to the filtered data sequence as follows:

$$\Gamma(f_m) = \frac{1}{F_s N} \left\| \sum_{n=1}^N z_n \exp \left(-j \frac{2\pi f_m (n-1)}{F_s} \right) \right\|^2 \quad (2)$$

$$f_m = \frac{m F_s}{N} \quad m = 0, 1, \dots, \frac{N}{2} \quad (3)$$

where F_s is the sampling frequency, which in our case is the number of frames per second (fps) of the camera and f_m are the frequencies for which the spectrogram is sampled.

We extract from the power spectra the maximum peak in the interval of feasible human step frequencies, which are assumed to be in the range between 1 and 3 Hz. This frequency is taken as the estimated step frequency f_{step} .

3.2 Walking speed estimation

To estimate the walking speed we consider the biomedical work by Grieve (Grieve, 1968) where a relation between the step frequency (f_{step}) and the walking speed (V_{walk}) normalized with height (H) is presented:

$$V_{walk} = \alpha f_{step}^\beta H \quad (4)$$

where V_{walk} is in m/s, f_{step} in Hz, H in meters and α and β are characteristic parameters which differ from one individual to another.

For higher accuracy, we have computed our own α and β parameters for the camera operator. We measured the time t_i it took the operator to walk a distance of $s = 100$ meters at the times per step ΔT_i given by a metronome ranging from $0.45s$ to $0.80s$ in intervals of $0.05s$ (see Table 1).

ΔT_i [s]	t_i [s]	$f_i = \frac{1}{T_i}$ [Hz]	$V_i' = \frac{s}{t_i H}$ [$\frac{1}{s}$]
0.45	48.18	2.22	2.08
0.50	55.60	2	1.80
0.55	61.63	1.82	1.62
0.60	74.54	1.67	1.34
0.65	84.42	1.54	1.19
0.70	94.63	1.43	1.06
0.75	104.42	1.33	0.96
0.80	116.06	1.25	0.86

Table 1: Experimental data used to compute the empirical *Step frequency-Walking velocity* relationship for the camera operator.

Normalized walking speeds V_i' and the step frequencies f_i were computed from the raw experimental data and then a power fitting was applied to obtain the values of $\alpha = 0.329$ and $\beta = 1.534$ (Fig. 3).

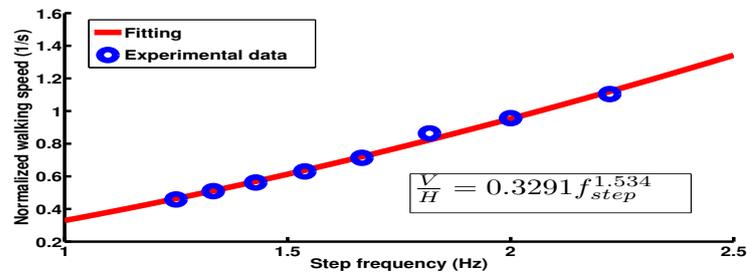


Figure 3: Power fitting of the experimental data to compute the relation between walking speed and step frequency ($\mu_{err} = 0.018$, $max_{err} = 0.04$).

3.3 Particle Filter for scale factor tracking

Having the walking velocity estimate, the scale factor for each section k could be easily computed by $d_k = \frac{V_{walk,k}}{V_{SLAM,k}}$, where $V_{SLAM,k}$ would be the mean of the speed at the N trajectory points used in the computation of the power spectra. However, given the empirical method for the walking speed estimation and the possible high variability of the SLAM velocity along N frames, we decided to use a probabilistic filter for the computation of the scale factor. This allows us to introduce an uncertainty to the scale factor and at the same time the rejection of spurious estimations of the walking speed.

Thus, we consider a dynamic system whose state \mathbf{x}_k is composed by the magnitude of the SLAM velocity $V_{SLAM,k}$ and the logarithm of the scale factor $\lambda_k = \log_{10}(d_k)$. We take the logarithm instead of directly the scale factor to avoid negative scale values.

To track the scale factor a particle filter with Sampling Importance Resampling is designed (Gordon, Salmond, and Smith, 1993). We use a particle filter rather than an EKF so that it can deal with high uncertainty priors of the scale factor which would involve a large linearization error in an EKF approach.

Hence the state of the system is approximated by a set of particles

$$S_k = \left\{ (\mathbf{x}_k^{(L)}, w_k^{(L)}) \mid L = 1, 2, \dots, P \right\} \quad (5)$$

where P is the number of particles and $\mathbf{x}_k^{(L)}$ and $w_k^{(L)}$ are respectively the state vector and the resampling weight of particle L . Each iteration of the particle filter is denoted by the index k and integrates the last N positions of the camera used for walking velocity estimation.

$$\mathbf{x}_k^{(L)} = \begin{bmatrix} V_{SLAM,k}^{(L)} \\ \lambda_k^{(L)} \end{bmatrix} \quad (6)$$

The particles are initialized such that the initial values of $\lambda_0^{(L)}$ are drawn from a Gaussian distribution $\lambda_0 \sim \mathcal{N}(0, \sigma_0)$, where σ_0 is a parameter related to the orders of magnitude being scoped out.

In the first step of the particle filter, particles are sampled down by a proposal distribution $p(\mathbf{x}_k | \mathbf{x}_{k-1})$:

$$\mathbf{x}_k^{(L)} \sim p(\mathbf{x}_k | \mathbf{x}_{k-1}^{(L)}) \quad (7)$$

In our system the sampling of the proposal distribution includes both the update of the SLAM velocity, which is taken as a control input coming from the visual odometry, and the possible drift in the scale. This is encoded in the following equations:

$$V_{SLAM,k}^{(L)} = \mu_{V,k} + \nu^{(L)} \quad (8)$$

$$\lambda_k^{(L)} = \lambda_{k-1}^{(L)} + \alpha^{(L)} \quad (9)$$

with $\nu^{(L)} \sim \mathcal{N}(0, \sigma_{V,k})$ and $\alpha^{(L)} \sim \mathcal{N}(0, \sigma_{drift})$, and where $\mu_{V,k}$ and $\sigma_{V,k}$ are the averaged speed and the corresponding standard deviation of the last set of N SLAM trajectory points used for spectral analysis, and σ_{drift} is the standard deviation prior of the scale drift between two consecutive sections, which is modelled as Gaussian noise.

Once the prediction has been computed, the samples are weighted by integrating the estimation of the walking velocity $V_{walk,k}$ from spectral analysis as follows:

$$w_k^{(L)} = p(V_{walk,k} | \mathbf{x}_{k-1}^{(L)}) \quad (10)$$

where $p(V_{walk,k} | \mathbf{x}_{k-1}^{(L)})$ is the probability density function defined by the measurement model $h(\mathbf{x}_k)$ and the statistics of the sensor noise. By assuming that the velocity estimation is affected by Gaussian noise of zero mean and standard deviation σ_{Vwalk} to be set up empirically, weights are computed as:

$$\omega_k^{(L)} = p(V_{walk,k} | \mathbf{x}_k^{(L)}) = \phi \left(\frac{V_{walk,k} - h(\mathbf{x}_k^{(L)})}{\sigma_{Vwalk}} \right) \quad (11)$$

where $\phi(z)$ is the probability density function of the standard normal distribution and the measurement function $h(\mathbf{x}_k^{(L)})$ is given by:

$$h(\mathbf{x}_k^{(L)}) = V_{SLAM,k}^{(L)} 10^{\lambda_k^{(L)}} \quad (12)$$

Finally, the weights are normalized for each particle and the next set of particles S_{k+1} is resampled by drawing P particles from the current set with probability equal to the corresponding weights.

3.4 Scaling of the trajectory

To track the scale factor, the original SLAM trajectory has been splitted in sections of N points, being each of these sections denoted by their corresponding time-step k . The scale factor to be applied to each section k is obtained by averaging the logarithmic scale values of the particle set S_k and undoing the logarithmic change as follows:

$$\bar{\lambda}_k = \frac{\sum_{i=1}^P \lambda_k^{(i)}}{P} \quad (13)$$

$$\bar{d}_k = 10^{\bar{\lambda}_k} \quad (14)$$

This scale factor must be applied to the position and velocity of the N camera states of section k . To simplify these state variables are encapsulated in a vector:

$$\mathbf{X}_k(n) = (r_x^n, r_y^n, r_z^n, v_x^n, v_y^n, v_z^n) \quad n = 1, 2, \dots, N \quad (15)$$

Finally, to ensure the continuity in position and velocity, each vector $\mathbf{X}_k(n)$ is scaled recursively as follows:

$$\hat{\mathbf{X}}_k(n) = \hat{\mathbf{X}}_{k-1}(N) + \bar{d}_k[\mathbf{X}_k(n) - \mathbf{X}_{k-1}(N)] \quad k = 2, 3, \dots \quad (16)$$

$$\hat{\mathbf{X}}_1(n) = \bar{d}_1 \mathbf{X}_1(n) \quad (17)$$

where $\hat{\mathbf{X}}_k(n)$ is the vector which includes the scaled position and velocity of the trajectory points contained in section k .

4 Experiments

In these experiments we use a catadioptric omnidirectional camera with a resolution of 1024x768 and a frame rate of 15 fps. This camera is mounted on a helmet carried by a human operator. We acquired 3 sequences of images walking along the same path of 230 meters with 3 different step frequencies. The ground truth is obtained using a metronome with a resolution 0.01 seconds. The metronome was set up to 0.70, 0.60 and 0.50 seconds per beat for each sequence, which translates in step frequencies of 1.43 Hz, 1.67 Hz and 2 Hz, respectively. The trajectory and the map corresponding to each image sequence are constructed with our SLAM approach for catadioptric cameras presented in (Gutierrez et al., 2011).

The experiments are divided in two parts. In the first part we test the accuracy of the spectral analysis carried out on the three trajectories using different setups and we select the best one among them. In the second part we evaluate our approach to estimate the scale factor with which we scale the trajectories provided by the SLAM algorithm.

4.1 Spectral analysis for step frequency estimation

First, we evaluate the feasibility of using spectral analysis on the SLAM trajectory to measure the step frequency. As stated in Section 3 visual odometry is divided in sections of N frames and the DFT is carried out on each section.

To compute the DFT we use the FFTW (Fast Fourier Transform West) C library (Frigo and Johnson, 2005). We compare different section dimensions of $N = 100$ and $N = 200$ frames. As the routines of this library perform faster when the length of the data sequence is a power of 2, data

sequences are padded with zeros to a length of N_p to fill this condition. A greater padding involves an increased resolution of the spectrogram. Thus, we compare two zero-padding instances ZP1 and ZP2. ZP1 corresponds to a padding being N_p the power of 2 closest to N . ZP2 corresponds to a padding with $N_p = 1024$.

In Fig. 4 the results of the measured step frequency of the three trajectories using different setups for the DFT are shown.

It can be observed that taking $N = 200$ points provides more accurate estimations and less outliers. This is done at the expense of increasing the interval between two consecutive estimations. It is also shown that a greater zero-padding does not provide any improvement in accuracy. Thus we select a setup of $N = 200$ data points and the ZP1 padding instance to compute the DFT for spectral analysis.

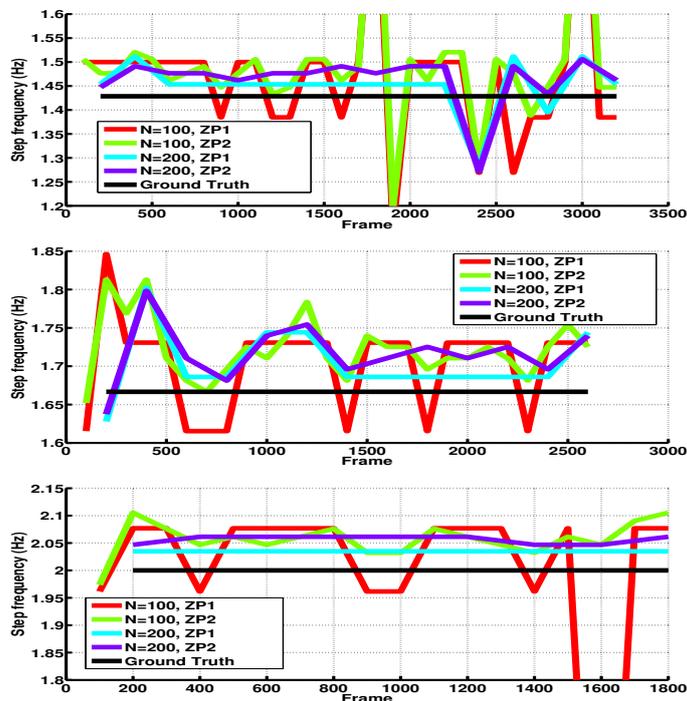


Figure 4: Spectral analysis along the same path at the three step frequencies of 1.43 (top), 1.67 (center) and 2 Hz (bottom) with different setups for the computation of the DFT.

4.2 Scaling of trajectories

In this section we scale the visual odometry from each sequence of images using the approach described in Section 3. The number of particles used in the particle filter is $P = 5000$. Standard deviations of the gaussians used to model probability distributions have been adjusted.

For the gaussian distribution used to initialize the values of the logarithmic scale factor for the initial samples, the standard deviation was set up to $\sigma_0 = 1$. This setup allows us to consider an initial uncertainty interval for the scale factor between 10^{-2} and 10^2 with a 95% of confidence. The standard deviation of the gaussian noise modeling scale drift was fixed to $\sigma_{drift} = 0.05$. The

standard deviation of the walking measurement estimation was set up to a high value of 0.2 m/s to avoid a particle depletion of the particle filter due to a low sensor noise.

After adjusting the priors we tested our approach on the three SLAM trajectories walked at different step frequencies. In Fig. 5 we show the final reconstruction of the trajectory compared to the Ground Truth over a satellite view from Google Maps. It can be observed the great improvement respect to the raw visual odometry estimation from the SLAM algorithm. Notice also that our approach provides a better estimation than applying a uniform scale factor. The reason is that the dynamic estimation of the scale factor every N frames (Fig. 6) allows us to correct the scale drift of the raw visual odometry.

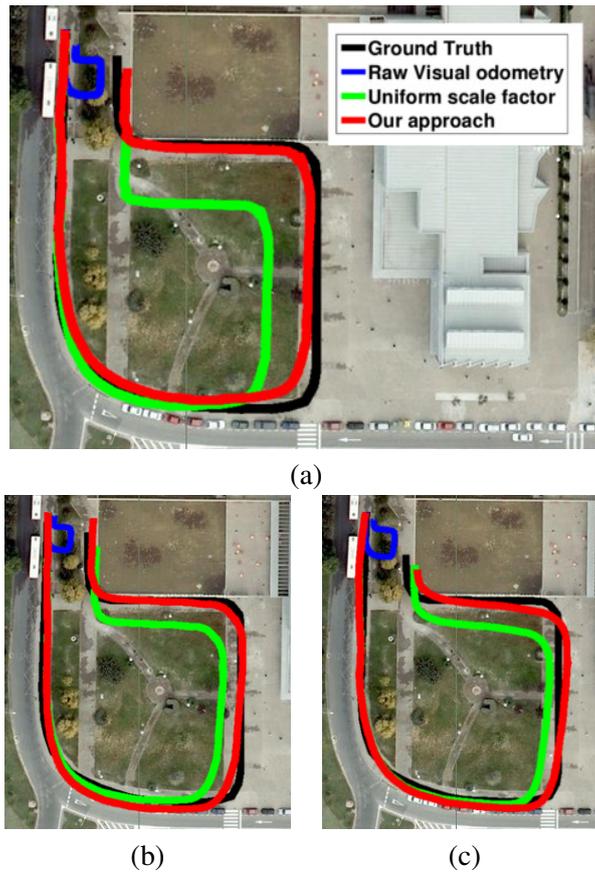


Figure 5: Visual odometry estimations using different approaches on the three trajectories walked at different step frequencies of about 1.43 Hz (a), 1.67 Hz (b) and 2 Hz (c).

The implementation of our approach has been done in MATLAB and it has been performed offline on the estimates provided by a real time monoSLAM application. In Fig. 7 we show the computation time to perform the spectral analysis and to scale each section of the trajectory. It can be observed that after an initial computation time of 0.3 seconds for the first section it steadily converges to a time of 0.015 seconds. Thus our approach is able to be used in a real time application taking advantage of the time that takes to fill the data sequence vector with the N elements for next spectral analysis.

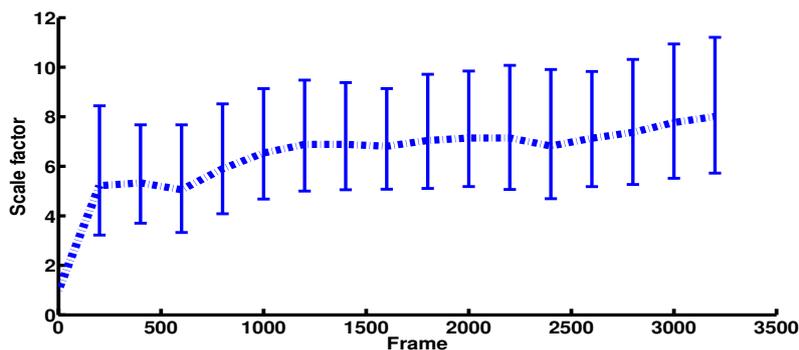


Figure 6: Scale factor estimation corresponding to the trajectory taken at a step frequency of 1.43 Hz. Vertical lines represent the uncertainty with a confidence of 95%. Variation of the scale along the trajectory implies that scale drift is observed.

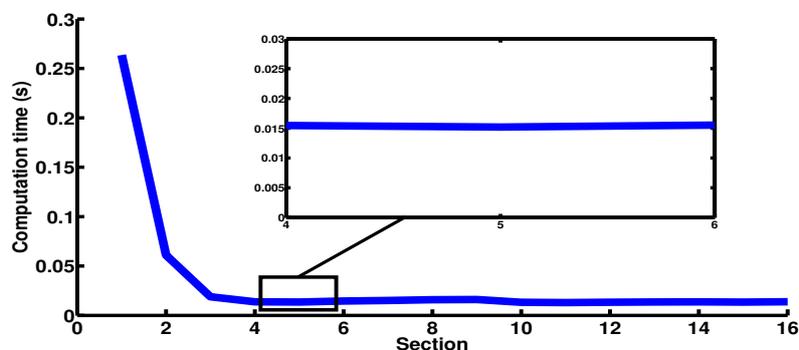


Figure 7: Computation time used to perform our approach for the different sections.

5 Conclusions

In this paper we have presented a new approach to estimate the absolute scale of the visual odometry of a head-mounted omnidirectional camera provided by a monocular SLAM application. Our approach combines spectral analysis on the visual odometry to estimate the walking velocity with a particle filter framework. With this information we are able to track the scale factor. The experiments show that our approach provides an accurate dynamic estimation of the absolute scale which is robust to scale drift. Although it was implemented to operate offline, low computation times show that this approach is suitable to be integrated in a real time SLAM application. This possibility will be considered in future work.

6 Acknowledgments

This work has been supported by the project DPI2009-14664-C02-01. Thanks to the I3A Fellowship Program.

References

- Civera, J., Davison, A. J. and Montiel, J. M. M. (2008). Inverse depth parametrization for monocular slam. *IEEE Transactions on Robotics*, 24(5):932–945.
- Civera, J. et al. (2010). 1-Point RANSAC for EKF Filtering: application to real-time structure from motion and visual odometry. *Journal of Field Robotics*, 27(5):609–631.
- Corke, P., Strelow, D. and Singh, S. (2004). Omnidirectional visual odometry for a planetary rover. In: *Intelligent Robots and Systems*, 4:pp. 4007–4012.
- Davison, A. J. (2003). Real-time simultaneous localisation and mapping with a single camera. In: *International Conference on Computer Vision*, 2:pp. 1403–1410.
- Eudes, A. et al. (2010). Fast odometry integration in local bundle adjustment-based visual slam. In: *International Conference on Pattern Recognition*, pp. 290–293.
- Frigo, M. and Johnson, S. G. (2005). The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231. Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- Gordon, N. J., Salmond, D. J. and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F Radar and Signal Processing*, 140(2):107–113.
- Grieve, D. W. (1968). Gait patterns and the speed of walking. *Biomedical Engineering*, 3(March):119–122.
- Gutierrez, D. et al. (2011). Adapting a real-time monocular visual slam from conventional to omnidirectional cameras. In: *11th OMNIVIS, held with International Conference on Computer Vision*, pp. 343–350.
- Hirasaki, E. et al. (1999). Effects of walking velocity on vertical head and body movements during locomotion. *Experimental Brain Research*, 127(2):117–130.
- Kuo, A. D. (2001). A simple model of bipedal walking predicts the preferred speed-step length relationship. *Journal of Biomechanical Engineering*, 0:264–269.
- Lothe, P. et al. (2010). Real-time vehicle global localisation with a single camera in dense urban areas: Exploitation of coarse 3d city models. In: *Computer Vision and Pattern Recognition*, pp. 863–870.
- Lupton, T. and Sukkarieh, S. (2008). Removing scale biases and ambiguity from 6dof monocular slam using inertial. In *International Conference on Robotics and Automation*, pp. 3698–3703.
- Mei, C. (2007). *Laser-Augmented Omnidirectional Vision for 3D Localisation and Mapping*. PhD thesis, INRIA Sophia Antipolis, Project-team ARobAS.
- Mei, C. et al. (2010). RSLAM: A system for large-scale mapping in constant-time using stereo. *International Journal of Computer Vision*, 0:1–17. Special issue of BMVC.

- Nistér, D., Naroditsky, O. and Bergen, J. (2006). Visual odometry for ground vehicle applications. *Journal of Field Robotics*, 23(1):3–26.
- Nützi, G. et al. (2010). Fusion of imu and vision for absolute scale estimation in monocular slam. *Journal of Intelligent Robotic Systems*, 61(1-4):287–299.
- Paz, L. M. et al. (2008). Large scale 6dof slam with stereo-in-hand. *IEEE Transactions on Robotics*, 24(5):946–957.
- Rituerto, A., Puig, L. and Guerrero, J. J. (2010). Visual slam with an omnidirectional camera. In: *20th International Conference on Pattern Recognition*, pp. 348–351.
- Scaramuzza, D., Fraundorfer, F. and Siegwart, R. (2009). Real-time monocular visual odometry for on- road vehicles with 1-point ransac. In: *International Conference on Robotics and Automation*, pp. 4293–4299.
- Scaramuzza, D. et al. (2009). Absolute scale in structure from motion from a single vehicle mounted camera by exploiting nonholonomic constraints. In: *12th International Conference on Computer Vision*, pp. 1413–1419.
- Strasdat, H., Montiel, J. M. M. and Davison, A. (2010). Scale drift-aware large scale monocular slam. In: *Proceedings of Robotics: Science and Systems*.
- Tardif, J.-P., Pavlidis, Y. and Daniilidis, K. (2008). Monocular visual odometry in urban environments using an omnidirectional camera. In: *Intelligent Robots and Systems*, pp. 2531-2538.
- Zarrugh, M., Todd, F. and Ralston, H. (1974). Optimization of energy expenditure during level walking. *European Journal of Applied Physiology and Occupational Physiology*, 33:293–306.