# Real-time 3-D Motion and Structure of Point Features: a Front-end System for Vision-based Control and Interaction 

Hailin Jin Paolo Favaro Stefano Soatto<br>Washington University, One Brookings Dr. 1127, St. Louis - MO 63130<br>\{hljin,fava, soatto\}@essrl.wustl.edu


#### Abstract

We present a system that consists of one camera connected to a Personal Computer that can (a) select and track a number of high-contrast point features on a sequence of images, (b) estimate their three-dimensional motion and position relative to an inertial reference frame, assuming rigidity, (c) handle occlusions that cause pointfeatures to disappear as well as new features to appear. The system can also (d) perform partial self-calibration and (e) check for consistency of the rigidity assumption, although these features are not implemented in the current release. All of this is done automatically and in realtime ( 30 Hz ) for $40-50$ point-features using commercial off-the-shelf hardware.

The system is based on an algorithm presented by Chiuso et al. [3] whose properties have been analyzed in [2]. In particular, the algorithm is provably observable, provably minimal and provably stable - under suitable conditions. The core of the system, consisting of C++ code ready to interface with a frame grabber as well as Matlab code for development, is available at http://ee.wustl.edu/ ~ soatto/research.html.

We demonstrate the system by showing its use as (1) an ego-motion estimator, (2) an object tracker, (3) an interactive input device. All without any modification of the system settings.


## 1 Introduction

We are interested in using vision as a sensor for machines to interact with the environment by moving, tracking, manipulating objects etc. In order to do so, a machine must be able to estimate its three-dimensional (3-D) motion relative to the scene and - to an extent that depends upon the application - the shape of the scene. The requirement of real-time operation imposes constraints on the algorithm to be used (it needs to be causal, for we cannot measure the future) as well as on the representation of shape (via point features; fancier representations cannot be handled by current commercial hardware). We consider, therefore, an $N$-tuple of points in the threedimensional Euclidean space, represented by their image on a reference plane, $\mathbf{y}_{0}$ and their depth along the pro-
jection rays $\rho$. Let the points move under the action of a rigid motion represented by a translation vector $T$ and a rotation matrix $R$, with linear velocity $V$ and rotational velocity $\widehat{\omega}$, represented by a skew-symmetric $3 \times 3$ matrix using the "hat" notation (see [2] for details). We assume that we can measure the (noisy) projection $\mathbf{y}^{i}(t)=$ $\pi\left(R(t) \mathbf{y}_{0}^{i} \rho^{i}+T(t)\right)+\mathbf{n}^{i}(t) \in \mathbb{R}^{2} \forall i=1 \ldots N$ which corresponds to an ideal perspective projection model. This choice is not crucial and the discussion can be easily extended to other projection models (e.g. spherical, orthographic, para-perspective, etc.). By organizing the timeevolution of the configuration of points and their motion, we end up with a discrete-time, non-linear dynamical system. In [2] it is proven that the following model is minimal, and that the Extended Kalman Filter (EKF) based on it is stable (in mean-square and with probability one):

| $\left(\mathbf{y}_{0}^{i}(t+1)=\mathbf{y}_{0}^{i}(t) \quad i=4 \ldots N\right.$ | $\mathbf{y}_{0}^{i}(0)=\mathbf{y}_{0}^{i}$ |
| :---: | :---: |
| $\rho^{i}(t+1)=\rho^{i}(t) \quad i=2 \ldots N$ | $\rho^{i}(0)=\rho_{0}^{i}$ |
| $T(t+1)=\exp (\widehat{\omega}(t)) T(t)+V(t)$ | $T(0)=T_{0}$ |
| $\Omega(t+1)=\log _{S O(3)}(\exp (\widehat{\omega}(t)) \exp (\widehat{\Omega}(t)))$ | $\Omega(0)=\Omega_{0}$ |
| $\begin{aligned} & V(t+1)=V(t)+\alpha_{V}(t) \\ & \omega(t+1)=\omega(t)+\alpha_{\omega}(t) \end{aligned}$ | $\begin{aligned} V(0) & =V_{0} \\ \omega(0) & =\omega_{0} \end{aligned}$ |
| $\mathbf{y}^{i}(t)=\pi\left(\exp (\widehat{\Omega}(t)) \mathbf{y}_{0}^{i}(t) \rho^{i}(t)+T(t)\right)+\mathbf{n}^{i}(t)$ |  |
| where $\mathbf{n}^{i} \sim \mathcal{N}\left(0, \Sigma_{n}\right)$ | $=1 \ldots N$. |

The notation $\log _{S O(3)}(R)$ stands for $\Omega$ such that $R=e^{\widehat{\Omega}}$ and is computed by inverting Rodrigues' formula ${ }^{11}$. Note that the indices of $\mathbf{y}_{0}$ and $\rho$ in the state start from 4 and 2 respectively. This is to guarantee that the model is minimal and stable ${ }^{2}$.

## 2 Implementation

The implementation of the EKF based on the above model is discussed in detail in [3], including partial self-calibration and handling of occlusions. Here we only describe the

[^0]"recipe" algorithm in a level of detail that should be sufficient to follow the implementation made available on the Web, or to implement the algorithm independently.
Initialization Choose the initial conditions $\mathbf{y}_{0}^{i}=\mathbf{y}^{i}(0)$; $\rho_{0}^{i}=1 ; T_{0}=0 ; \Omega_{0}=0 ; V_{0}=0 ; \omega_{0}=0$ for $i=1 \ldots N$ where $N$ is chosen depending upon the performance of the hardware ( $N=50$ in our case, a dual PIII 733 Mhz system). For the initial variance $P_{0}$, choose it to be block diagonal with blocks $\Sigma_{n^{i}}(0)$ given from the analysis of the feature-tracking algorithm ${ }^{3}$ corresponding to $\mathbf{y}_{0}^{i}$, a large positive number $M$ (typically 100-1000 units of focal length) corresponding to $\rho^{i}$, zeros corresponding to $T_{0}$ and $\Omega_{0}$ (fixing the inertial frame to coincide with the initial reference frame). We also choose a large positive number $W$ for the blocks corresponding to $V_{0}$ and $\omega_{0}$. The variance $\Sigma_{w}(t)$ is a design parameter that is available for tuning. Finally, set $P(0 \mid 0)=P_{0}$ and $\hat{\xi}(0 \mid 0) \doteq\left[\mathbf{y}_{0}^{4^{T}}, \ldots, \mathbf{y}^{N_{0}^{T}}, \rho_{0}^{2}, \ldots, \rho_{0}^{N}, T_{0}^{T}, \Omega_{0}^{T}, V_{0}^{T}, \omega_{0}^{T}\right]^{T}$ Transient During the first transient of the filter, we do not allow for new features to be acquired. Whenever a feature is lost, its state is removed from the model and its best current estimate is placed in a storage vector. The transient can be tested, for instance, by a threshold on the innovation, a threshold on the variance of the estimates, or by a fixed time interval. We choose a combination of them, with the time interval set to 30 frames, corresponding to the first second of video. The recursion to update the state $\xi$ and the variance $P$ proceed as follows: Let $f$ and $h$ denote the state and measurement model, so that equation (1) can be written in concise form as
\[

\left\{$$
\begin{array}{l}
\xi(t+1)=f(\xi(t))+w(t) \quad w(t) \sim \mathcal{N}\left(0, \Sigma_{w}\right)  \tag{2}\\
y(t)=h(\xi(t))+n(t) \quad n(t) \sim \mathcal{N}\left(0, \Sigma_{n}\right)
\end{array}
$$\right.
\]

Prediction:


[^1]$\hat{\Theta}^{i}=e^{\widehat{\Omega}}\left[\begin{array}{c}\mathbf{y}_{0}^{i} \\ 1\end{array}\right] \rho_{i}$.
The matrix $\Pi$ is block-diagonal with blocks $\frac{1}{\rho^{i}}\left[\begin{array}{cc}I_{2} & -\hat{y}^{i}\end{array}\right]$ where $\hat{y}^{i}=\pi\left(e^{\Omega} \mathbf{y}_{0}^{i} \rho^{i}+T\right)$.
Regime Whenever a feature disappears, we remove it from the state as during the transient. However, during regime operation a feature selection module works in parallel with the filter to select new features so as to maintain roughly a constant number $N$ and a distribution as uniform as possible across the image plane. We implement this by randomly sampling points on the plane, then searching around that location for a feature that passes a "sum of square difference"-type test. Once a new pointfeature is found, a "subfilter" is initialized. Its evolution is given by

and $P_{\tau}$ is updated according to the usual Riccati equation. After a probation period the feature is inserted into the state.
Tuning The variance $\Sigma_{w}(t)$ is a design parameter which we choose to be block diagonal with the blocks corresponding to $T(t)$ and $\Omega(t)$ equal to zero (a deterministic integrator). We choose the remaining parameters using the Periodogram test. In practice, we choose the blocks corresponding to $\mathbf{y}_{0}^{i}$ equal to the variance of the measurements, and the elements corresponding to $\rho^{i}$ all equal to $\sigma_{\rho}$. We then choose the blocks corresponding to $V$ and $\omega$ to be diagonal with element $\sigma_{v}$, and then we change $\sigma_{v}$ relative to $\sigma_{\rho}$ depending on whether we want to allow for more or less regular motions. We then change both, relative to the variance of the measurement noise, depending on the level of desired smoothness in the estimates.

## Acknowledgments

Supported by NSF grant IIS-9876145 and ARO grant DAAD19-99-1-0139. The authors wish to thank Xiaolin Feng and Pietro Perona for their generous sharing of their code for feature tracking.

## References

[1] A. Azarbayejani and A. Pentland. Recursive estimation of motion, structure and focal length. IEEE Trans. Pattern Anal. Mach. Intell., 17(6):562-575, 1995.
[2] A. Chiuso and S. Soatto. 3-D Motion and structure from 2-D motion causally integrated over time. Part I: theory. Tutorial lecture notes. IEEE Conf. on Robotics and Automation, April 2000.
[3] A. Chiuso and P. Favaro and H. Jin and S. Soatto. 3-D motion and structure causally integrated over time. Part 2: experiments. In Proc. of the Eur. Conf. on Computer Vision, June 2000 (in press).


[^0]:    ${ }^{1}$ A Matlab implementation of $\log _{S O(3)}$ is included in the software distribution.
    ${ }^{2}$ There are in the literature similar models that do not include $\mathbf{y}_{0}$ in the state, for instance [1]. It is shown in 2] that such models are subminimal, and therefore have a non-zero-mean measurement error that results in a highly biased estimate.

[^1]:    ${ }^{3}$ We assume that the tracking error is independent in each point, and therefore $\Sigma_{n}$ is block diagonal with diagonal equal to 1 pixel std. in the current implementation.

