Acta Astronautica 98 (2014) 169–188

Contents lists available at ScienceDirect

Acta Astronautica

journal homepage: www.elsevier.com/locate/actaastro

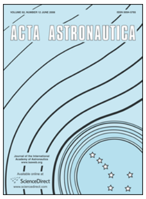

Single camera absolute motion based digital elevation mapping for a next generation planetary lander

Luke M Feetham [a,*], Nabil Aouf [a], Clement Bourdarias [b], Thomas Voirin [c]

[a] *Centre of Electronic Warfare, Cranfield University, UK*
[b] *Sensors Engineering and Image Processing, Astrium ST, France*
[c] *Control Systems Division, ESA-ESTEC, The Netherlands*

ARTICLE INFO

Article history:
Received 31 July 2013
Received in revised form
17 December 2013
Accepted 15 January 2014
Available online 10 February 2014

Keywords:
Next generation entry descent and landing
Structure from motion
Single camera absolute scale recovery
Digital elevation model
Hazard detection
Pin-point landing.

ABSTRACT

Robotic planetary surface exploration missions are becoming much more ambitious in their science goals as they attempt to answer the bigger questions relating to the possibility of life elsewhere in our solar system. Answering these questions will require scientifically rich landing sites. Such sites are unlikely to be located in relatively flat regions that are free from hazards, therefore there is a growing need for next generation entry descent and landing systems to possess highly sophisticated navigation capabilities coupled with active hazard avoidance that can enable a pin-point landing. As a first step towards achieving these goals, a multi-source, multi-rate data fusion algorithm is presented that combines single camera recursive feature-based structure from motion (SfM) estimates with measurements from an inertial measurement unit in order to overcome the scale ambiguity problem by directly estimating the unknown scale factor. This paper focuses on accurate estimation of absolute motion parameters, as well as the estimation of sparse landing site structure to provide a starting point for hazard detection. We assume no prior knowledge of the landing site terrain structure or of the landing craft motion in order to fully assess the capabilities of the proposed algorithm to allow a pin-point landing on distant solar system bodies where accurate knowledge of the desired landing site may be limited. We present results using representative synthetic images of deliberately challenging landing scenarios, which demonstrates that the proposed method has great potential.

1. Introduction

Future missions to land robotic craft on the surface of planetary bodies will require increasingly more sophisticated levels of autonomy in order to efficiently and safely carry out their mission objectives. The primary driving force behind this is an ever-increasing demand on cost effectiveness and science return. Planetary surface exploration missions are becoming much more ambitious in their science goals as they attempt to answer the bigger questions relating to the possibility of life elsewhere in our solar system. This inevitably leads to larger, more expensive robotic craft carrying large suites of complex scientific instruments. In order to take full advantage of this sophisticated equipment and maximize scientific return, a careful consideration of potential landing sites needs to be conducted and then a suitable landing site must be selected based on a compromise between its interest to scientists and the ability to land safely at that location so as to not place an unjustifiable risk of loss or damage on expensive, mission critical hardware.

The entry, descent and landing (EDL) systems used in current missions are predominantly based on technology

---

[*] Corresponding author.
E-mail address: l.feetham@cranfield.ac.uk (L. Feetham).

http://dx.doi.org/10.1016/j.actaastro.2014.01.011

developed during the 1960s to 1970s [1], which, although proven to be reliable on many occasions, is severely limited in terms of sophistication and autonomous capability. For U.S. missions in particular, the backbone EDL system design is that of the Viking Mars Landers of 1976 [1]. The sensors typically employed in this type of system consist of radar altimeters to determine height and vertical velocity, through determining the temporal derivative of the height; an inertial measurement unit (IMU) to propagate the spacecraft's position and velocity; and perhaps also Doppler radar to determine horizontal velocity. This type of system has a number of limitations that affect the achievable landing precision. Firstly, an initial position determination is carried out just prior to atmospheric entry, often using Earth-based measurements, therefore potentially containing significant error. Following this, the position and velocity are propagated on-board the spacecraft throughout the entry phase using only IMU measurements, which, due to the nature of inertial measurements, leads to significant error accumulation. During the latter stages of the descent phase the on-board radar systems are activated to measure height and velocity. These measurements can be used to prevent further error accumulation and correct for errors in motion estimates, but they cannot be used to correct for the already accumulated errors in position. Consequently, past and present planetary landers have very large landing uncertainty ellipses for the expected touchdown spot, where the major axis length can range from 300 to 20 km (Viking-Curiosity) [2]. In addition to the huge uncertainty, the current sensors are incapable of providing any information on potential surface hazards that may cause critical damage to the spacecraft. Consequently, mission designers must select landing sites where the entire uncertainty ellipse lies in a region of relatively flat terrain that is free from hazards; therefore, currently, the above mentioned compromise is skewed heavily in favour of safety and against scientific interest. Indeed, the most scientifically interesting regions are rarely located in such hospitable places [2,3], leading to many landing sites being ruled out or otherwise requiring long distances to be traversed in order to reach the primary target location of the mission, which itself carries significant risk and reduces cost effectiveness. There is therefore a need for providing the capability of a precise and safe pin-point landing at a site of significant scientific interest (which may be surrounded by numerous hazards) in order to dramatically improve cost effectiveness.

For the sake of argument, throughout this paper we assume that a desirable landing site has been selected in a manner similar to how landing sites are currently chosen – i.e. through an extensive survey based on data previously acquired through orbital observations, in which accessibility and science return are traded off based on the capabilities of the landing craft. This could be done either by using data from other spacecraft, e.g. for a Mars mission, observations from existing orbiters may be used to make a selection early on in the mission design phase; or data could be acquired from an orbital phase of the mission in question prior to landing that may be used by the ground segment to select a suitable landing site. However, it can be pointed out that the presented method does not actually require any prior knowledge of the landing site terrain or of the motion parameters in order to estimate the structure and motion of the spacecraft during decent. Therefore this technique could also be applied to enable a safer landing than is possible with current EDL technology in situations where there is very little prior knowledge, such as on distant solar system bodies where it may not be possible to carry out detailed orbital observations prior to entry.

The focus of this work is to develop the foundations of a next-generation EDL system that has the ability of making a pin-point landing at a predetermined site of interest, as well as the capability of detecting and avoiding potential hazards. As a starting point towards achieving this aim we propose a method for accurately estimating the motion parameters of a descending spacecraft using a single camera as well as measurements from an IMU. The specific class of technique that is applied in this work is known as a structure from motion (SfM) algorithm, which, as well as providing estimates of the spacecraft’s motion, also allows the structure in the scene to be estimated. Accurate recovery of the 3D scene structure would provide a starting point for a hazard detection system, therefore this type of technique has great potential in EDL.

Single camera SfM algorithms suffer from an inability to determine the motion and structure parameters as absolute quantities. This arises as a consequence of the 3D–2D image formation process, in which a small object close to the camera may produce an identical image to that of a larger object further away. Therefore, the translational motion and the scene structure can only be recovered up to an unknown scale factor in the absence of any further information due to the inability to determine depth in the scene. In this work we make use of measurements from an IMU in order to obtain the additional information needed to determine the unknown scale factor, using a technique inspired by the work of Nutzi et al. [4].

In the remaining sections of this paper a multi-source, multi-rate data fusion algorithm will be presented that aims at combining a single camera, recursive, feature-based structure from motion (SfM) algorithm with measurements from an on-board inertial measurement unit (IMU), using the Extended Kalman filter (EKF). The paper describes how the fusion of these measurements enables direct estimation of the unknown scale factor, and results are presented showing that this leads to an accurate estimation of the motion parameters and structure of the scene. The algorithm is tested using sequences of representative artificial images that are generated from an artificial terrain model of a planetary surface, using the Planet and Asteroid Natural Scene Generation Utility (PANGU). These tests are carried out from three different initial heights above the surface (2000 m, 1000 m, and 500 m), in which no knowledge of the landing site terrain structure is assumed in advance (this is deliberately in contrast to the approaches in [5–12], all of which make use of a priori mapped landmarks, i.e. absolute 3D locations of certain surface features are known in advance and used to assist motion estimation by tracking these features throughout descent image sequences) and where only very limited knowledge of the motion parameters is