Image and Vision Computing 52 (2016) 97–113

Contents lists available at ScienceDirect

Image and Vision Computing

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

A robust similarity measure for volumetric image registration with outliers☆

Patrick Snape[a,*], Stefan Pszczolkowski[c], Stefanos Zafeiriou[a], Georgios Tzimiropoulos[b], Christian Ledig[a], Daniel Rueckert[a]

[a] Imperial College London, Department of Computing, London SW7 2AZ, UK
[b] University of Nottingham, School of Computer Science, Nottingham NG8 1BB, UK
[c] University of Nottingham, School of Medicine, Nottingham NG7 2UH, UK

ARTICLE INFO

Article history:
Received 9 February 2015
Received in revised form 30 April 2016
Accepted 5 May 2016
Available online 29 May 2016

Keywords:
Image registration
Lucas–Kanade
Normalised gradient
Free-form deformation

ABSTRACT

Image registration under challenging realistic conditions is a very important area of research. In this paper, we focus on algorithms that seek to densely align two volumetric images according to a global similarity measure. Despite intensive research in this area, there is still a need for similarity measures that are robust to outliers common to many different types of images. For example, medical image data is often corrupted by intensity inhomogeneities and may contain outliers in the form of pathologies. In this paper we propose a global similarity measure that is robust to both intensity inhomogeneities and outliers without requiring prior knowledge of the type of outliers. We combine the normalised gradients of images with the cosine function and show that it is theoretically robust against a very general class of outliers. Experimentally, we verify the robustness of our measures within two distinct algorithms. Firstly, we embed our similarity measures within a proof-of-concept extension of the Lucas–Kanade algorithm for volumetric data. Finally, we embed our measures within a popular non-rigid alignment framework based on free-form deformations and show it to be robust against both simulated tumours and intensity inhomogeneities.

1. Introduction

Image registration is an important procedure in many areas of computer vision for both 2D and volumetric 3D images. Given its relevance, there is a large body of prior work concentrating on methodologies for performing accurate registration. In this paper, we are most interested in techniques that attempt to densely align two images according to a global similarity measure. Many global similarity measures have been proposed, yet only a few focus on being robust to the presence of outliers and systematic errors. We consider a similarity measure to be robust if it is not biased by the presence of noise and/or occlusions within the image to be registered. Generally, if a similarity measure is biased by noise or occlusions, this will manifest as a failure to accurately register the images.

In the case of 2D images, systematic errors and outliers are common in the form of illumination variance and occlusions, which naturally occur in so called "in-the-wild" images. Although 2D image alignment is a broad field, much of the existing work focuses on augmenting existing efficient algorithms with improved robust properties. For example, one of the first algorithms to describe a 2D image alignment approach was the Lucas–Kanade (LK) algorithm [1]. The LK algorithm concentrates on recovering a warp that best maximises a similarity measure between two images. Numerous extensions to the LK algorithm have been proposed [2,3,4] and most are based on ℓ_2 norm minimisation [3,5,6,7]. Most notably, the inverse compositional framework proposed by Baker and Matthews [5] provides a computationally efficient framework for solving the least squares problem.

For volumetric, or 3D images, outliers can occur in the form of pathologies, and systematic errors are commonly seen as intensity inhomogeneities caused by image acquisition artefacts such as bias fields [8]. Several methods have been proposed for registration of medical images with mismatches, focusing on robustness [9], tumour models [10] or Bayesian models [11]. However, previous methods [12] all require prior knowledge of what constitutes a mismatch in order to detect and ignore them. Additionally, a number of

☆ This paper has been recommended for acceptance by Ioannis A. Kakadiaris, PhD.

* Corresponding author.
E-mail addresses: p.snape@imperial.ac.uk (P. Snape), stefan.pszczolkowskiparraguez@nottingham.ac.uk (S. Pszczolkowski), s.zafeiriou@imperial.ac.uk (S. Zafeiriou), yorgos.tzimiropoulos@nottingham.ac.uk (G. Tzimiropoulos), christian.ledig@imperial.ac.uk (C. Ledig), d.rueckert@imperial.ac.uk (D. Rueckert).

http://dx.doi.org/10.1016/j.imavis.2016.05.006
0262-8856/

methods based on mutual information have been proposed to reduce the effect of intensity inhomogeneities in the registration [13,14,15].

To the best of our knowledge, no existing similarity measure provides robustness against both outliers and intensity inhomogeneities for registration of 3D images *without prior knowledge of the type of dissimilarity*. However, the 2D similarity measure recently proposed in Ref. [16], has been shown to be robust against both general occlusions and illumination variance. This measure is formulated as the cosine of normalised gradient orientations and is simple and efficient to compute. For 3D images, we seek to provide a similarity measure that can utilise the robust properties of the cosine function. This requires calculating a similarity between the two images being aligned that can be represented as an angle. In this work, we provide two separate methodologies of measuring angular similarity between 3D images.

We clarify that when we state 3D images, we are referring to volumetric data where it is valid to compute gradients along all 3 of the principle axes. This may be in contrast to other usages of the term "3D" when referring to data sources such as range images or mesh data. In this case, we would refer to these data sources as 2.5D data, as the computation of the gradient for the depth axis is non-trivial. In fact, as mentioned by Baker et al. [17], the inverse compositional algorithm is no longer valid when extended to 2.5D data due to the representation of the data as a surface. Therefore, given that the treatment of 2.5D data is totally different from the 3D volumetric images that we use here, we do not further consider it.

It is important to note that there is a large amount of existing work on using gradient information for image registration on volumetric data. The concept of normalised image gradients was introduced to the field of medical image registration by Pluim et al. [18]. In Ref. [18], normalised mutual information (NMI) [19] is weighted voxelwise by the normalised image gradients in order to incorporate spatial information. After this initial work, the first similarity based solely on normalised gradients was proposed by Haber and Modersitzki [20]. This similarity measure is based on the *squared cosine* of the normalised gradients and is equivalent to minimising the *squared inner product*. In contrast, our proposal is to use the *cosine* of the normalised gradient orientations and is equivalent to minimising the *inner product*. This seemingly small difference, the squaring of the cosine, causes outliers to bias the similarity measure and thus affects the robustness. Despite these properties of the squared cosine measure, it has been successfully utilised in the literature [21,22,23] for registering images that do not contain outliers.

Preliminary work on the cosine of orientations has been shown in our previous work [24,25]. In Ref. [24], we gave preliminary results that show that the cosine of normalised gradient orientations represents a robust similarity measure in the presence of both occlusions and intensity inhomogeneities. We extend this work in two major areas.

Firstly, we note that that there are two separate angular measures that can be defined in order to compute the cosine of normalised gradient orientations between two images. These orientations are based on the spherical coordinates of the gradients and the inner product between the gradients. As a proof of concept, we directly extend the methodology of Ref. [16] to provide evidence that our similarity measures are robust to occlusions and intensity inhomogeneities. Although the extension of LK-type algorithms to 3D is simple and was proposed in Ref. [17], no previous investigation has been done on how similarity measures perform when extended to 3D. Since no previous work has investigated the use of 3D Lucas–Kanade for robust registration, we chose to extend existing robust 2D methods into 3D. These extended state-of-the-art 2D methods are then compared against our proposed methods using a synthetic dataset.

Secondly, in order to show that our similarity measures compete with state-of-the-art techniques, we embed them within a widely and successfully used non-rigid registration framework based on free-form deformations (FFD) [26]. This FFD registration framework differs from our LK example as it is a local deformation model and thus contains many thousands of parameters. This large parameter space makes Gauss–Newton optimisation infeasible due to the memory requirements of inverting the Hessian matrix. For this reason, registration techniques that focus on local deformations are generally solved using gradient descent algorithms that incorporate line searches [26]. We compare against the NMI measure [19], DRAMMS [27] and the cosine squared measure [20] and confirm the robustness of the proposed similarity measures on simulated pathological data from a tumour database. **Secondly**, we provide more extensive evaluation into the robust properties of the cosine of normalised gradient orientations when applied to volumetric data.

The remainder of the paper is organised as follows. In Section 2 we discuss relevant related works and consider the advantages of our proposed measures. In Section 3 we give a thorough explanation of the theory behind our work and empirically verify its robustness on our chosen data. In Section 4 we describe the application of our similarity measure within the Lucas–Kanade algorithm and in Section 5 we show its application within the non-rigid framework of Rueckert et al. [26]. Experimental results within both rigid and non-rigid alignment are given in Section 6. Finally, we draw conclusions in Section 7.

2. Related work

In this paper we are most interested in similarity measures that have been proposed for parametric image alignment. This is primarily due to the fact that parametric image alignment methods are often very computationally efficient. In particular, we are interested in approaches that show robustness to both intensity inhomogeneities and occlusions.

Within 2D image registration there has been a large body of work that focuses on illumination invariance. Within the most popular parametric image alignment framework, the Lucas–Kanade algorithm (LK), one of the the earliest attempts was by Hager and Belhumeur [28]. In Ref. [28], the authors incorporate a linear appearance basis that models illumination variance and excelled in the area of object tracking. However, this requires prior knowledge of the type of object under consideration so that the appearance basis can be precomputed. There is a large amount of existing literature about incorporating prior knowledge via linear bases within the LK algorithm. However, we do not consider them here, as we are most interested in modelling unseen structured variation via robust similarity measures and thus prior knowledge is generally not available.

Within the LK framework, a number of robust measures have been proposed. Black and Jepson [29] proposed incorporating robust statistics into the LK framework and showed their effectiveness in the presence of illumination variance. Dowson and Bowden [3] incorporated the mutual information (MI) measure within the LK framework and found it to be superior to an SSD based measure for illumination variance. However, using mutual information as a similarity measure requires updating the Jacobian and Hessian matrices at each iteration. For this reason, the MI measure is unable to make use of more computationally efficient LK algorithms such as the inverse compositional (IC) method proposed in Ref. [5]. Although the authors do give an ad hoc solution that involves fixing the values of the Jacobian and Hessian matrices, it is still a more complex implementation than the original inverse compositional algorithm. Evangelidis and Psarakis [6] provide a correlation measure between images and a computational framework that is invariant to illumination differences. This is a great strength of the algorithm, however, the correlation measure is still grossly affected by outliers. Lucey et al. [4] propose a method of weighting the LK fitting with a large bank of filters in a computationally efficient manner. This is a very