TOWARDS GENERAL-PURPOSE IMAGE REGISTRATION

By

Gehua Yang

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Computer Science

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Examining Committee:

Charles V. Stewart, Thesis Adviser
Badrinath Roysam, Member
Daniel Freedman, Member
Richard Radke, Member

Rensselaer Polytechnic Institute
Troy, New York

June 2007
(For Graduation August 2007)

ABSTRACT

This thesis presents four image-registration-related methods. The first two methods aid the construction of a fully-automatic retinal image registration system, called RIVERS.

The first method, automatic masking, is a preprocessing technique for image registration that separates the retinal surface from the background in each image. It is especially important when aligning retinal images scanned from slides or images acquired from different cameras.

The second method, a covariance-driven refinement technique, is developed to handle pairs with extreme low overlap (e.g. 5%) in the context of aligning a set of sparsely-overlap images.

The third method, named GDB-ICP, is an automated 2d-image-pair registration algorithm using a hypothesis-and-test strategy and an extension of the Dual-Bootstrap method for refinement. This algorithm is capable of aligning images taken of a wide variety of natural and man-made scenes as well as many medical images, tolerating low overlap, substantial orientation and scale differences, large illumination variations, and physical changes in the scene. An important component of this is the ability to automatically reject pairs that have no overlap or have too many differences to be aligned well. Experimental results on a data set of 22 challenging image pairs show that the algorithm effectively aligns 19 of the 22 pairs and rejects 99.8% of the misalignments that occur when all possible pairs are tried.

An extension of the algorithmic principle of GDB-ICP, the fourth method, a 3d-to-2d alignment algorithm, is applied to estimate the location of a hand-held camera with respect to a 3d model augmented with texture information. Little prior knowledge is assumed about the camera location. A key issue is that initially the model-to-image mapping is well-approximated by a simple 2d-to-2d transformation based on a local model surface approximation. However, the algorithm must transition to the 3d-to-2d projection necessary to solve the position estimation problem. The experiments are conducted on a collection of 9 range scans and 60 image cover-

ing approximately a 100m x 100m region of RPI campus. The algorithm successfully and correctly determines the camera location of 52 images, while indicates it cannot find an alignment for the remaining 8.