Applied Soft Computing 46 (2016) 1067–1078

Contents lists available at ScienceDirect

Applied Soft Computing

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

Multi-objective optimization for parameter selection and characterization of optical flow methods

Jose Delpiano [a,*], Luis Pizarro [b], Rodrigo Verschae [c,1], Javier Ruiz-del-Solar [c,d]

[a] School of Engineering and Applied Sciences, Universidad de los Andes, Santiago, Chile
[b] Department of Computer Science, University College London, London, UK
[c] Advanced Mining Technology Center, Santiago, Chile
[d] Department of Electrical Engineering, Universidad de Chile, Santiago, Chile

ARTICLE INFO

Article history:
Received 16 June 2015
Received in revised form 20 January 2016
Accepted 22 January 2016
Available online 15 February 2016

Keywords:
Multi-objective optimization
Optical flow
Parameter selection

ABSTRACT

Optical flow methods are among the most accurate techniques for estimating displacement and velocity fields in a number of applications that range from neuroscience to robotics. The performance of any optical flow method will naturally depend on the configuration of its parameters, and for different applications there are different trade-offs between the corresponding evaluation criteria (e.g. the accuracy and the processing speed of the estimated optical flow). Beyond the standard practice of manual selection of parameters for a specific application, in this article we propose a framework for automatic parameter setting that allows searching for an approximated Pareto-optimal set of configurations in the whole parameter space. This final Pareto-front characterizes each specific method, enabling proper method comparison and proper parameter selection. Using the proposed methodology and two open benchmark databases, we study two recent variational optical flow methods. The obtained results clearly indicate that the method to be selected is application dependent, that in general method comparison and parameter selection should not be done using a single evaluation measure, and that the proposed approach allows to successfully perform the desired method comparison and parameter selection.

1. Introduction

Following the seminal works of Lucas and Kanade [1] and Horn and Schunck [2] on local and global optical flow estimation, respectively, numerous variants of these and other sophisticated approaches have been proposed to solve the ill-posed problem of motion recovery. Optical flow benchmarks such as Middlebury [3], KITTI [4] and MPI-Sintel [5] list most modern methods and evaluate them on image sequences with known motion under varying conditions. The importance of benchmarking to evaluate and rank the different optical flow approaches is relevant as low-level motion cues are the cornerstone of many high-level machine vision and pattern recognition systems [6,7], where applications may impose certain constraints on the accuracy of the estimated motion and/or the speed at which such estimations can be obtained.

The parameter space has a direct influence on the performance of optical flow methods. Searching for an accurate and/or fast estimation often leads to different parameter settings of the same method. The literature on optimization of the hyper-parameters mainly focuses on accuracy, while speed is improved with multigrid solvers and GPU implementations [8–11]. Some works employ statistical methods to learn the parameter space. For example, simultaneous perturbation stochastic approximation can be used for learning model parameters from training data [12], and Bayesian inference for estimating the regularization parameter [13] and other model parameters [14]. Some works have also compared the execution times of several approaches [15,16,8,17]. Nevertheless, most articles do not consider the problem of parameter estimation in a multi-objective sense as there is little research regarding the method's accuracy and speed simultaneously as figures of merit for model selection. A recent example is [18], where the authors compare several meta-heuristics in the Middlebury and Sintel databases, but only taking a single objective (the accuracy) into account.

In this article, we argue that a parameter setting that compromises between both criteria, accuracy and speed, might be the right operating point for a given application. In fact, such a compromise has been observed in certain species in the animal kingdom, which exhibit a behavioral trade-off between accuracy and speed in completing specific tasks [19]. By using (vector-valued) multi-objective optimization we can explore the whole parameter space to find an operating curve describing the optimal parameter set that best describes the accuracy-speed compromise for a specific optical flow approach. Moreover, the operating curves of different methods can be juxtaposed in order to select the most suitable method at specific operating points given the joint objective of minimizing the alignment error and the execution time. The proposed framework is tested considering two objectives, though it can deal with multiple objectives without any modification.

To set some notation, let us define a two-objective operating point of an optical flow method parameterized by θ as $\mathbf{v}(\theta) = (AEE(\theta), T(\theta))^\top \in \mathbb{R}^2$, described in terms of its average end-point error *AAE* and execution time *T*. It is clear that the performance is a function of the model parameters θ. Naturally, we can only compute the *AEE* if we know the ground truth motion. Therefore, we work with benchmark data sets to train our approach. Nevertheless, we later discuss how to use our framework in scenarios where no ground truth is available. Given a parameter configuration θ^i

* Corresponding author. Tel.: +56 962380898.
E-mail addresses: jd@miuandes.cl (J. Delpiano), l.pizarro@ucl.ac.uk (L. Pizarro), rodrigo@verschae.org (R. Verschae), jruizd@ing.uchile.cl (J. Ruiz-del-Solar).
[1] Now with the Graduate School of Informatics, Kyoto University, Kyoto, Japan.

http://dx.doi.org/10.1016/j.asoc.2016.01.037
1568-4946/

and its corresponding solution $\mathbf{v}(\theta^i)$ in the objective space, the aim of multi-objective optimization is to find a set of points $\{\mathbf{v}(\theta^j)\}_j$ that have a better value for at least one of the objectives and equal or better values for the remaining objectives. This results in a set of solutions that corresponds to an approximation of the *Pareto-optimal front* [20] for the tested optical flow method. This front contains much richer information than a local or global optimum obtained by single-objective optimization. The Pareto front can be considered as a receiver operating characteristic (ROC) curve that characterizes the optimized method by its most distinctive parameter settings (operating conditions) in one curve. Having the Pareto fronts for various methods in the same axes permits a quantitative comparison of them that is inherently multi-objective and therefore suited to a problem with possibly conflicting objectives.

Searching for the optimal parameter setting represents a challenging problem, especially when dealing with multi-objective optimization [21,22] and when there is no analytical form for the multi-objective function being optimized. The most common approaches in this case are random search, grid search, weighted sum of the objectives, and evolutionary algorithms. We use evolutionary algorithms [23]. In particular, we employ genetic algorithms [24] for this task. Genetic algorithms can solve problems with multiple solutions. They do not require objective function derivatives, thus they are easy to implement and can cope with non-continuous problems. Standard genetic algorithms search the parameter space in an evolutive manner, considering only one objective. To optimize several objectives concurrently we utilize an evolutionary multi-objective optimization (EMO) strategy [20]. Among the many existing alternatives [25], we use NSGA-II (an improved Non-dominated Sorting Genetic Algorithm) [20,26], a successful approach for EMO that has a fast approach for non-dominated solution sorting and a smart criterion for diversity preservation. Two particular reasons for using this method in the present work are that it is well understood and that open source implementations exist (we use the implementation made available by the authors). Other nature-inspired derivative-free optimization algorithms include the usage of fuzzy logic to combine solutions of particle swarm optimization and genetic algorithms for a single objective [27], techniques inspired in gravitational forces for single-objective optimization [28], and a hybrid approach based on NSGA-II and neural networks for the optimization of time-intensive simulations [29]. Multi-objective optimization has been applied before to other computer vision tasks, such as segmentation [30], face detection [31], tracking [32] and 3D vision [33]. To the best of our knowledge, there is very limited research related to multi-objective optimization of optical flow methods. One exception is the work of Salmen et al. [34], where the authors look for highly accurate and efficient optical flow algorithms. However, they work with non-dense optical flow methods and define efficiency as the number of flow vectors found per frame. Thus, they are not considering algorithm speed.

The main contributions of our work are as follows:

- We propose a methodology for parameter selection and characterization of optical flow methods based on multi-objective optimization considering the joint accuracy-speed compromise.
- Our multi-objective optimization strategy can be applied to tune the parameters of any optical flow method optimally (variational or not). In general, the parameter space of an optical flow method can be very large, which makes the optimization task very challenging.
- We use the proposed method to analyze two methods, namely the *large displacement optical flow* method of Brox and Malik [35] and the *anisotropic Huber-L_1 optical flow* approach of Werlberger et al. [36], in two recent databases (Middlebury and KITTI).

The current article is an extended version of [37], with the main differences being:

- In [37], we test multi-objective optimization of optical flow in a classical method, namely combined local global (CLG). For our new article, we describe two recent optical flow methods, the large displacement optical flow by Brox and Malik [35] and the Huber-L_1 optical flow by Werlberger et al. [36], in a common mathematical framework. We use those two methods as a proof-of-concept for our methodology, as it can cope with any optical flow method, variational or not. For the experiments in [37], we worked with our own implementation of the CLG method. Now, we constrained ourselves to using only optical flow implementations made available by their authors in order to have fair characterization and comparison of the optical flow methods.
- In [37], we chose to use the Middlebury data set. In the present article, we test our methodology with the classical Middlebury data set, but also in the current and more challenging KITTI data set. This time we added an analysis of distribution of displacements. This kind of analysis is needed to understand the differences in performance between the two selected optical flow methods.
- The aim of our previous article [37] was the characterization of optical flow methods, but did not consider the optimal selection of parameters for them. In [37], our experimental results end up characterizing three variants of a classical optical flow method, the CLG optical flow by Bruhn et al. [38]. However, we did not study the optimum parameters related to the solutions in those final Pareto fronts. In this article, we analyzed the parameter settings of the non-dominated solutions, particularly their variation range along the Pareto front.
- The comparison among optical flow methods is just suggested in our conference paper [37]. In our present manuscript, we characterize optical flow methods and make use of that characterization to compare two optical flow methods when applied to two relevant data sets.

In the recent related work of Pereira et al. [18], several meta-heuristics were taken into account when evaluating the *large displacement optical flow* [35] in the Middlebury and Sintel databases, but with the main differences with our work being that [18] only deals with single-objective optimization, that we work with the more recent and challenging KITTI dataset, that in addition to *large displacement optical flow* we consider the *anisotropic Huber-L_1 optical flow*, and that we perform a quantitative analysis of the obtained optimal parameters for each analyzed method.

1.1. Organization of the paper

Section 2 describes the two advanced optical flow methods that were chosen as sample methods to evaluate multi-objective optimization. Section 3 presents a novel evolutionary multi-objective methodology for parameter selection and characterization of optical flow methods. Section 4 analyzes the optimization results, comparing both methods by their accuracy-speed operating curves. It also studies the resulting parameter settings. Section 5 concludes the paper suggesting the development of multi-objective rankings in modern benchmarks of optical flow methods.

2. Advanced optical flow methods

We briefly discuss the parameter space of two recent variational methods for optical flow estimation, the *Large Displacement Optical Flow* (LDOF) approach of Brox and Malik [35] and the *Anisotropic Huber-L_1 Optical Flow* (AHL1) approach of Werlberger et al. [36]. As mentioned above, these methods are chosen as proof-of-concept since the proposed multi-objective optimization framework presented in Section 3 will work similarly for any other method, variational or not.

Using a common notation to describe both methods, let $I_1, I_2 : \Omega \subset \mathbb{R}^2 \to \mathbb{R}$ be two consecutive grayscale images defined on the rectangular grid Ω. The optical flow field aligning both frames is a function $\mathbf{u} : \Omega \to \mathbb{R}^2$. That is, $\mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), u_2(\mathbf{x}))^\top$, $\forall \mathbf{x} = (x_1, x_2)^\top \in \Omega$.

2.1. Large displacement optical flow

Traditional variational methods for optical flow computation fail to estimate the motion of small scale structures moving fast. The LDOF approach overcomes this problem by incorporating point correspondences from descriptor matching into the variational setting [35]. The optical flow field $\mathbf{u}$ is obtained as the minimizer of the energy functional

$$E_{\mathrm{LDOF}}(\mathbf{u}) = E_{\mathrm{int}}(\mathbf{u}) + \gamma E_{\mathrm{grad}}(\mathbf{u}) + \alpha E_{\mathrm{reg}}(\mathbf{u}) + \beta E_{\mathrm{match}}(\mathbf{u}, \mathbf{u_c}) + E_{\mathrm{desc}}(\mathbf{u_c}) \,. \quad (1)$$

The first three terms in (1) originate in the classic variational optical flow formulation [39] which penalize model deviations from gray value constancy, gradient constancy, and regularization (smoothness) of the solution, respectively,

$$E_{\mathrm{int}}(\mathbf{u}) = \int_\Omega \Psi(|I_2(\mathbf{x}+\mathbf{u}(\mathbf{x})) - I_1(\mathbf{x})|^2)\,\mathrm{d}\mathbf{x}\,, \quad (2)$$

$$E_{\mathrm{grad}}(\mathbf{u}) = \int_\Omega \Psi(|\nabla I_2(\mathbf{x}+\mathbf{u}(\mathbf{x})) - \nabla I_1(\mathbf{x})|^2)\,\mathrm{d}\mathbf{x}, \quad (3)$$

$$E_{\mathrm{reg}}(\mathbf{u}) = \int_\Omega \Psi(|\nabla u_1(\mathbf{x})|^2 + |\nabla u_2(\mathbf{x})|^2)\,\mathrm{d}\mathbf{x}. \quad (4)$$

The regularized L_1 norm (also called TV) $\Psi(s^2) = \sqrt{s^2+\epsilon^2}$ is used to penalize model deviations. The authors set $\epsilon = 0.001$ to avoid