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

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

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,

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

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

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

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

differentiability problems at $s = 0$. The last two terms in (1) integrate the point correspondences from descriptor matching:

$$E_{\text{match}}(\mathbf{u}, \mathbf{u_c}) = \int_\Omega \delta(\mathbf{x})\rho(\mathbf{x})\Psi(|\mathbf{u}(\mathbf{x}) - \mathbf{u_c}(\mathbf{x})|^2)\,\mathrm{d}\mathbf{x}, \tag{5}$$

$$E_{\text{desc}}(\mathbf{u_c}) = \int_\Omega \delta(\mathbf{x})|\mathbf{f}_2(\mathbf{x} + \mathbf{u_c}(\mathbf{x})) - \mathbf{f}_1(\mathbf{x})|^2\,\mathrm{d}\mathbf{x}, \tag{6}$$

where $\mathbf{u_c}(\mathbf{x})$ are the matched correspondences at some spatial locations determined by the delta function $\delta(\mathbf{x})$. The correspondences are matched according to the feature vectors $\mathbf{f}_1$ and $\mathbf{f}_2$ associated with the frames I_1 and I_2, respectively, and weighted by their matching score $\rho(\mathbf{x})$. The parameters α, β, $\gamma \geq 0$ balance the different energy terms in (1). Finally, the method considers pre-smoothing the images with Gaussian filter of size $\sigma \geq 0$. Brox and Malik [35] manually tuned the parameters to obtain the best average angular error on the benchmark data sets they used. We denote the parameterization of the LDOF approach by $\theta^{\text{LDOF}} = (\sigma, \alpha, \beta, \gamma)^\top$. Later on we will characterize this optical flow method by its Pareto-optimal front on the Middlebury and KITTI benchmarks by multi-objective optimization.

2.2. Anisotropic Huber-L_1 optical flow

The AHL1 approach of Werlberger et al. [36] introduces a regularization of the flow field that preserves discontinuities present in the image sequence, and penalizes deviations from that model with a robust function. An estimate of $\mathbf{u}$ is obtained as the minimizer of the energy functional

$$E_{\text{AHL1}}(\mathbf{u}) = E_{\text{int-lin}}(\mathbf{u}) + \frac{1}{\lambda}E_{\text{reg-ani}}(\mathbf{u}), \tag{7}$$

with $\lambda > 0$. The first term, which is a linearized version of the gray value constancy penalty (2), reads

$$E_{\text{int-lin}}(\mathbf{u}) = \int_\Omega \Psi(|\mathbf{u}(\mathbf{x})^\top \nabla I_2(\mathbf{x}) + I_2(\mathbf{x}) - I_1(\mathbf{x})|^2)\,\mathrm{d}\mathbf{x}, \tag{8}$$

where the authors also used the L_1 norm $\Psi(s^2) = \sqrt{s^2 + \epsilon^2}$, but with $\epsilon = 0$. They handle the differentiability at $s = 0$ in their optimization framework. The second term in (7),

$$E_{\text{reg-ani}}(\mathbf{u}) = \int_\Omega \left(\Phi_\varepsilon(|D^{1/2}\nabla u_1(\mathbf{x})|^2) + \Phi_\varepsilon(|D^{1/2}\nabla u_2(\mathbf{x})|^2)\right)\,\mathrm{d}\mathbf{x} \tag{9}$$

promotes discontinuity preserving, anisotropic regularization (cf. isotropic TV regularization (4)) of the flow field $\mathbf{u}$, where $D \in \mathbb{R}^{2\times 2}$ is a symmetric, positive definite diffusion tensor that encodes directional information of the image frames. The authors chose $D^{1/2} = \exp(-a|\nabla I|^b)\mathbf{n}\mathbf{n}^\top + \mathbf{n}^\perp\mathbf{n}^{\perp\top}$, with $a = 5.0$, $b = 0.5$, $\mathbf{n} = \frac{\nabla I}{|\nabla I|}$ and $\mathbf{n}^\perp$ the perpendicular vector to $\mathbf{n}$. The robust Huber function Φ_ε is used to penalize deviations from regularity

$$\Phi_\varepsilon(|s|^2) = \begin{cases} \dfrac{|s|^2}{2\varepsilon} & \text{if } |s| \leq \varepsilon \\ |s| - \dfrac{\varepsilon}{2} & \text{else} \end{cases} \tag{10}$$

where the authors set $\varepsilon = 0.01$. We refer to [36] for the full details of the optimization method and for an improved model that considers temporal symmetry of the flow field. It is important to mention that the linearization in (8) is only valid for small displacements. Therefore, to handle large displacements the implementation is embedded in a coarse-to-fine warping scheme.

Our preliminary experiments showed that accuracy and execution time are sensitive to scale factor (sf) for the coarse-to-fine strategy, the number of warps (nw), and the number of iterations (nit) for the optimization algorithm. We denote the parameterization of the AHL1 approach by $\theta^{\text{AHL1}} = (sf, nw, nit)^\top$. Later on we will characterize this optical flow method by its Pareto-optimal front on the Middlebury and KITTI benchmarks.

3. Proposed methodology

3.1. Evolutionary multi-objective optimization

Multi-objective optimization is targeted to look for the set of optimal solutions of a function with more than one objective. It should not be confused with the optimization of a weighted sum of its objectives. For an optical flow method parameterized by $\theta \in \Theta$ we deal with vector-valued solutions or operating points $\mathbf{v}(\theta) \in \mathbb{R}^{N_{obj}\times 1}$, where N_{obj} is the number of target objectives. Multi-objective optimization enables us to explore the whole parameter space Θ to find an operating curve describing the optimal parameter set for a specific optical flow method. Operating points are compared according to Pareto-domination, as described in Definition 1. Consequently, there is a set of points that are not dominated by any other vector, which form the *Pareto-optimal front*. 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. This curve contains much richer information than a local or global optimum obtained by single-objective optimization.

Definition 1. An operating point represented by the vector $\mathbf{v} = (v_1, \ldots, v_{N_{obj}})^\top$ with N_{obj} objectives is said to *Pareto-dominate* another solution $\mathbf{w} = (w_1, \ldots, w_{N_{obj}})^\top$ if $v_i \leq w_i$, $\forall i = 1, \ldots, N_{obj}$; and $\exists i' \in \{1, \ldots, N_{obj}\}$ such that $v_{i'} < w_{i'}$, where the inequality is fulfilled strictly.

In order to efficiently explore the parameter space we utilize *evolutionary multi-objective optimization* (EMO). A popular approach to EMO is NSGA-II [20,26], an improved non-dominated sorting genetic algorithm. The NSGA-II approach efficiently sorts non-dominated solutions and has a clever criterion for diversity preservation. It uses a simple procedure for creating a random parent population and new offspring populations based on binary tournament selection, mutation and crossover. For details we refer the reader to [26].

The NSGA-II procedure for solution sorting is as follows:

1. For every solution $\mathbf{v}(\theta^i)$, save the domination count $n_{\mathbf{v}(\theta^i)}$, the number of solutions which dominate $\mathbf{v}(\theta^i)$. Record the set $S_{\mathbf{v}(\theta^i)}$ of solutions that $\mathbf{v}(\theta^i)$ dominates. The set of solutions where $n_{\mathbf{v}(\theta^i)} = 0$ is the first non-dominated front and their non-domination rank is 1.
2. For every solution $\mathbf{v}(\theta^i)$ with $n_{\mathbf{v}(\theta^i)} = 0$, reduce the domination count of each solution $\mathbf{v}(\theta^j)$ in $S_{\mathbf{v}(\theta^i)}$ by one. Save solutions $\mathbf{v}(\theta^j)$ which reached domination count zero as a new non-dominated front. The non-domination rank of these solutions is one more than those in the previous non-dominated front.
3. Repeat Step 2 for each new non-dominated front.

We are interested in the set of solutions with rank 1, which are non-dominated solutions and as such they all form the Pareto front. Let us consider the parameter set

$$\Theta_{\text{P}} = \{\theta | \mathbf{v}(\theta) = (v_1(\theta), \ldots, v_{Nobj}(\theta))^\top \text{ is in the Pareto front}\} \tag{11}$$

and the lists of sorted objectives

$$\mathbf{V}_m = (v_m(\theta^1), \ldots, v_m(\theta^{|\Theta_{\text{P}}|})) \quad \text{for} \quad m = 1, \ldots, N_{obj}, \tag{12}$$

such that $v_m(\theta^i) \leq v_m(\theta^{i+1})$, $\forall i = 1, \ldots, |\Theta_P| - 1$. A small value of the crowding measure

$$d_{crowd}[i] = \sum_{m=1}^{N_{obj}} \frac{v_m(\theta^{i+1}) - v_m(\theta^{i-1})}{v_m(\theta^{|\Theta_P|}) - v_m(\theta^1)}, \quad \text{for } i = 2, \ldots, |\Theta_P| - 1, \quad (13)$$

assuming $d_{crowd}[1] = d_{crowd}[|\Theta_P|] = \infty$, indicates a high density of solutions around the ith point of the Pareto front. This measure stabilizes around a value greater than 1 as the population moves from generation to generation, which indicates that the NSGA-II algorithm preserves the diversity of solutions. In this article we choose the NSGA-II algorithm for EMO, although any other EMO algorithm could be used as well.

3.2. Optimization strategy for optical flow methods

In optical flow we estimate the motion between pairs of consecutive images. Let us consider a data set $\mathbf{I} = \{I_l^k\}_{l=1,\ldots,N_k}^{k=1,\ldots,N_{seq}}$ containing N_{seq} sequences, with the kth sequence formed of N_k images. The data set has

$$N_{\mathbf{I}} = \sum_{k=1}^{N_{seq}} (N_k - 1) \quad (14)$$

pairs of images. The performance of an optical flow method on the data set $\mathbf{I}$ given the parameter configuration $\theta \in \Theta$ will be measured by the solution vector

$$\mathbf{v}(\theta) = (AEE(\theta), T(\theta))^\top \in \mathbb{R}^2, \quad (15)$$

where the objectives *AEE* and *T* correspond to the average endpoint error and the average execution time *T* [3], respectively. The averages are computed over all $N_{\mathbf{I}}$ image pairs.

Our goal is to find the parameter configurations $\theta \in \Theta_P \subset \Theta$ which result in solution vectors $\mathbf{v}(\theta)$ at the Pareto front, i.e., solutions that are not Pareto-dominated by any other solution. For this purpose we propose the evolutionary multi-objective optimization framework shown in Fig. 1. The *Evolutionary population* block iterates over multiple generations, producing a population of optical flow parameters that is tested in the *Optical flow evaluation* block using the benchmark data set containing multiple image sequences $\mathbf{I}$ with ground truth motion $\mathbf{u}_{GT}$. The *NSGA-II* block sorts the solutions into non-dominated front sets according to the two-objective criterion, which allows the parameter population to be ranked and optimized. The algorithm outputs the optimal parameter set Θ_P and its corresponding rank-1 Pareto front $\mathbf{v}(\theta)$, $\forall \theta \in \Theta_P$.

Algorithm 1 describes the proposed methodology for parameter selection and characterization of optical flow methods in more detail. The final result of this algorithm is an optimal Pareto front for the optical flow method being studied. The user defines a variation range for each parameter and that information guides the random definition of the initial population. Afterwards, every new population is initialized using binary tournament selection, binary crossover, and binary mutation [20,26]. The parent and offspring populations are combined and sorted according to their non-domination rank and crowding measure [26]. The last step for each generation is to take the N_{ind} best individuals and keep them for the next generation. The non-domination rank and crowding measure are computed from the objective values, which are the Pareto front statistics in Fig. 1, namely *AEE* and *T* for the solutions forming the Pareto front.

Algorithm 1. Evolutionary multi-objective optimization for optical flow

Input: $\mathbf{I}$, $\mathbf{u}_{GT}$: Benchmark data set
Input: N_{gen}: Number of generations
Input: N_{ind}: Number of individuals in a generation
Output: Pareto-optimal parameter set Θ_P
Output: Pareto front statistics $\mathbf{v}(\theta) = (AEE, T)^\top$, $\forall \theta \in \Theta_P$

$P_1 \leftarrow$ Initial population with random individuals $\{\theta^j\}_{j=1,\ldots,N_{ind}}$
for $i = 1, \ldots, N_{gen}$ **do**
  $Q_i \leftarrow$ Offspring from P_i by selection, crossover and mutation
  $R_i \leftarrow P_i \cup Q_i$
  $\mathbf{v}(\theta^j) \leftarrow \text{OpticalFlowEvaluation}(\mathbf{I}, \mathbf{u}_{GT}, \theta^j)\ \forall \theta^j \in R_i$ ▸ *Eq.* (15)
  $F_1, \ldots, F_{\text{max-rank}} \leftarrow$ Split R_i into non-dominated front sets ▸ *NSGA-II*
  $P_{i+1} \leftarrow \varnothing$
  $f \leftarrow 1$
  while $|P_{i+1}| + |F_f| \leqslant N_{ind}$ **do**
    $P_{i+1} \leftarrow P_{i+1} \cup F_f$
    $f \leftarrow f + 1$
  end while
  $F_f \leftarrow$ Sort F_f by descending crowding measure ▸ *Eq.* (13)
  $P_{i+1} \leftarrow P_{i+1} \cup F_f[1 : N_{ind} - |P_{i+1}|]$
end for
$\Theta_P \leftarrow P_{N_{gen}}$

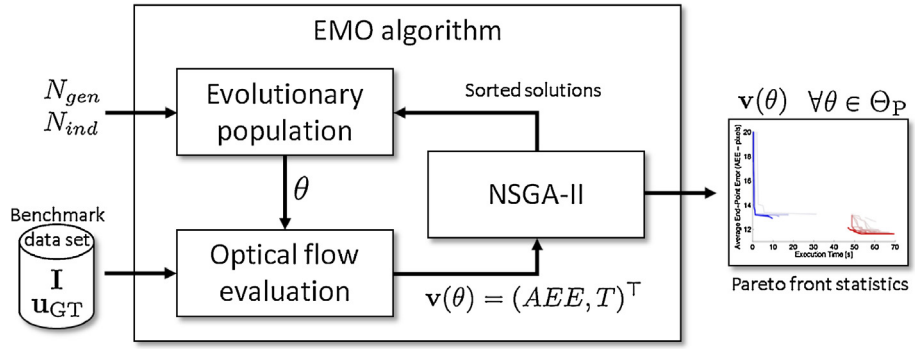

Fig. 1. Diagram of the proposed evolutionary multi-objective optimization framework for parameter selection and characterization of optical flow methods.

4. Experimental results

4.1. Data sets

We have chosen two data sets for validating the proposed methodology for multi-objective optical flow optimization: the Middlebury benchmark [3], which has become classical for evaluating optical flow methods; and the KITTI benchmark [4], a real world evaluation data set oriented to autonomous driving. Fig. 2 shows a common color scheme used to represent a motion vector at each pixel location. Figs. 3 and 4 show examples of both benchmarks alongside the ground truth motion fields. The Middlebury data set consists of eight sequences with ground truth motion obtained using hidden fluorescent textures and high resolution imaging with UV illumination. The data set also includes artificial sequences generated from 3D computer graphics models, from which the ground truth displacements are known directly. The KITTI benchmark contains 194 real training sequences with high resolution images captured from an autonomous platform while driving in the city of Karlsruhe in rural areas and on highways. Semi-dense ground truth motions were obtained by acquiring 3D points with a laser scanner and then projecting them onto the images. Some sequences contain violations to common optical flow assumptions such as brightness constancy – due to shining objects and non-uniform illumination. This affects some methods that rank high in the Middlebury benchmark, which then perform poorly on the KITTI data set [4]. These two benchmarks have another fundamental difference. The

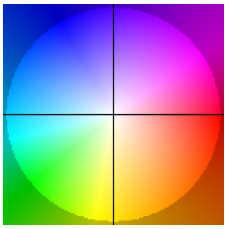

Fig. 2. Color wheel representing a motion field. The angle with the *x*-axis indicates the motion direction, and the color intensity expresses the magnitude of the displacement.

Middlebury library shows relatively small pixel displacements compared to the KITTI library. The histograms of Fig. 5 show that in the Middlebury benchmark most displacements are less than 5 pixels and no greater than 25 pixels. However, the KITTI data set includes quite large displacements. Fig. 5 shows that some are larger than 100 pixels. Classic optical flow algorithms are based on functionals with linearized data terms, which is only suitable for small displacements [40]. Image sequences containing large displacements are a clear challenge for most optical flow methods, which report large errors when evaluated on the KITTI benchmark. It will be interesting to see later the performance of the LDOF (Section 2.1) and the AHL1 (Section 2.2) methods on both libraries, considering that the first method is designed to handle large displacements, while the second method is based on a linearized functional.

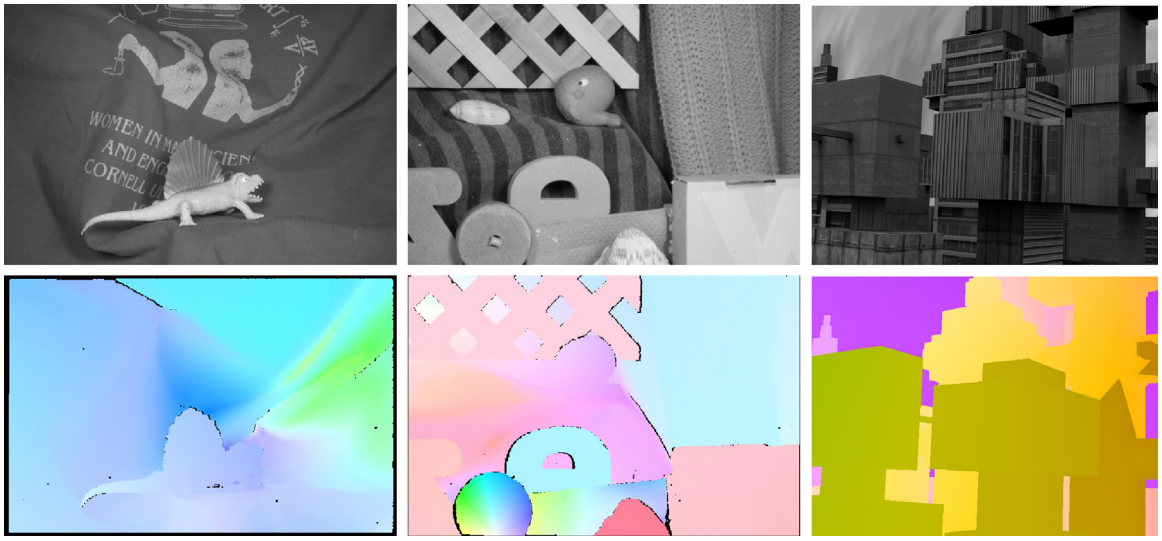

Fig. 3. Three sample frames from the Middlebury benchmark [3] and their corresponding ground truth motion towards the next frame.

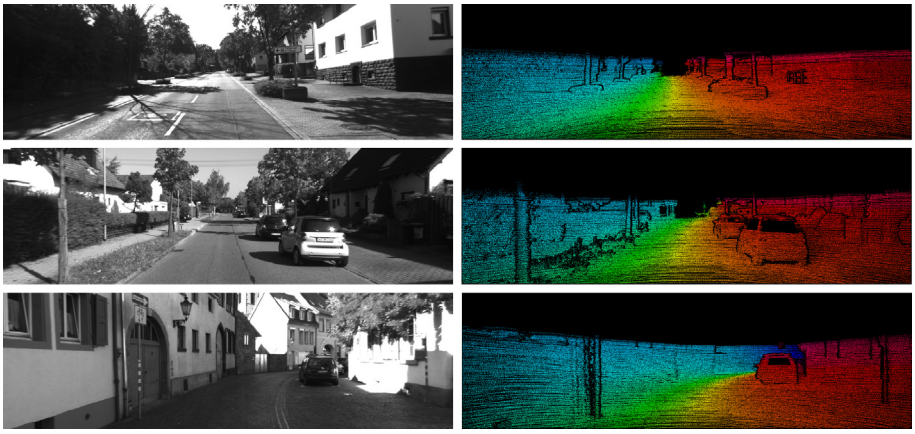

Fig. 4. Three sample frames from the KITTI benchmark [4] and their corresponding ground truth motion towards the next frame.

4.2. Experimental setup

In our multi-objective optimization framework we use the code[2] provided by the authors of both optical flow methods. Table 1 shows the default parameter setting in the provided code as well as the range of values we tested in our experiments. We refer to Section 2 for more details about these parameters. It is important to note that the authors do not provide clear indications about to how the parameters can be chosen. Therefore, we aim at studying the effect of different parameterizations on the multi-objective performance of both the LDOF and the AHL1 optical flow methods, information that is integrated in the Pareto front curve of each method.

4.3. Evaluation

We now evaluate the proposed evolutionary multi-objective optimization Algorithm 1 using the two chosen benchmark data sets. We use the parameter settings for the optical flow methods described in the previous section. Our approach efficiently examines the whole parameter space Θ, with $\theta \in \Theta$, to find the Pareto fronts described by the solution vectors $\mathbf{v}(\theta) = (AEE(\theta), T(\theta))^{\top}$, where *AEE* is the average end-point error and *T* the average execution time.

4.3.1. Repeatability test

Fig. 6 reports two realizations of our methodology for a small example, where one variable (σ) of the LDOF method was optimized to look for the best values of *AEE* and *T* for one sequence (RubberWhale) in the Middlebury data set. Fig. 6(a) shows that both final Pareto fronts are close to each other, covering almost the same ranges in both objectives. Fig. 6(b) shows that the values of the crowding measure for both runs tend to get closer to each other. Fig. 6(c–f) shows that the minimum and average values of both objectives tend to converge to very close values for both executions of our methodology. We can conclude that the results of multi-objective optimization of optical flow are not highly sensitive to the random components of the methodology.

4.3.2. Middlebury benchmark

Fig. 7(a, b) shows the evolution of the Pareto fronts in the objective space for both optical flow methods LDOF and AHL1, respectively. The first noticeable result is that the *x*–*y* axes differ in both methods, meaning that they have different performance on the Middlebury library, which will be further discussed in the next section. Fig. 7(c, d) shows the evolution of the *AEE* for both methods. It can be seen in the average curves that the *AEE* was reduced in more than 10% and 40%, respectively. Similarly, Fig. 7(e, f) shows a reduction of the minimum execution time of about 25% for the LDOF approach and a reduction in the average execution time of about 70% for the AHL1 approach. We can conclude that the proposed EMO strategy can effectively optimize the joint accuracy – speed objective for both optical flow methods when tested on the Middlebury data set.

4.3.3. KITTI benchmark

Fig. 8(a, b) shows the evolution of the Pareto front for the LDOF and AHL1 methods, where the latter reveals a narrower evolution indicating more stable generations, although the former method

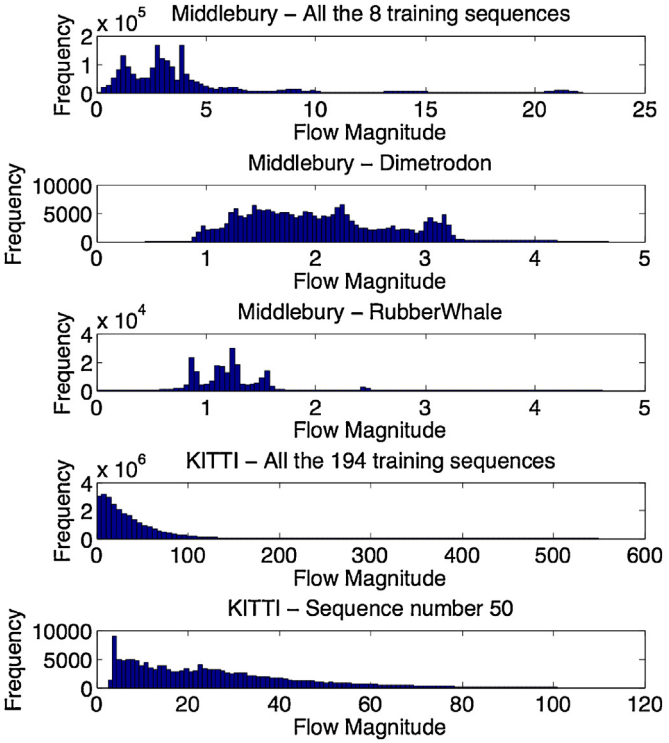

Fig. 5. Displacements histograms. The KITTI benchmark has much larger motions that the Middlebury benchmark, which poses a significant challenge to optical flow methods.

[2] LDOF: http://lmb.informatik.uni-freiburg.de/resources/software.php, AHL1: http://gpu4vision.icg.tugraz.at/index.php?content=subsites/flowlib/flowlib.php.

Table 1
Parameter settings for the optical flow methods to optimize using the proposed evolutionary multi-objective approach.

Method	Parameter	Default	Min	Max	Encoding	No. of bits
LDOF	σ	0.8	0.4	1.6	Real	–
	α	30	15	60	Real	–
	β	300	150	600	Real	–
	γ	5	2.5	10	Real	–
AHL1	*sf*	0.8	0.01	0.99	Real	–
	nw	10	0	20	Binary	8
	nit	100	0	200	Binary	8

performs better in terms of *AEE*. Fig. 8(c) shows a reduction close to 8% in both the minimum and average *AEE* curves for the LDOF method. Fig. 8(d) displays a slight reduction in the minimum *AEE* and almost no change in the final generation for the average curve. Fig. 8(e) shows a reduction of 18% and 10% in the minimum and average execution time of the LDOF method. Fig. 8(f) shows a relatively small reduction in the minimum execution time, and a more significant reduction of 76% in the average execution time of the AHL1 method. Here we can conclude that the KITTI data set results more challenging for the AHL1 optical flow approach, while the LDOF approach, which was designed to handle large displacements, performs as expected.

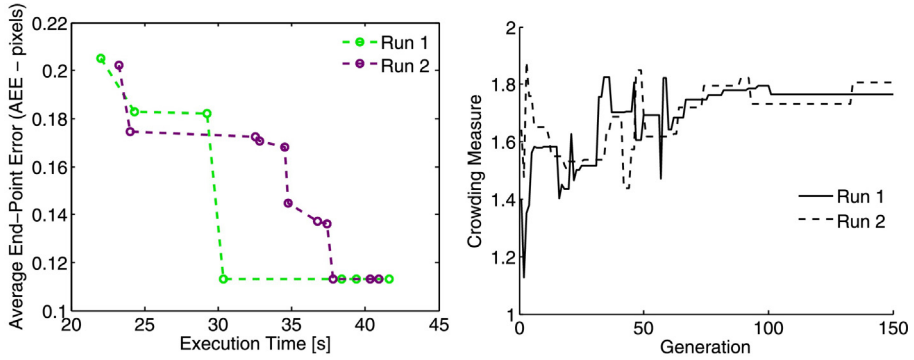

(a) Final Pareto front

(b) Crowding measure evolution

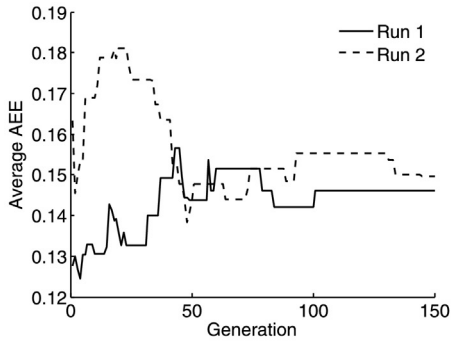

(c) Avg. AEE evolution

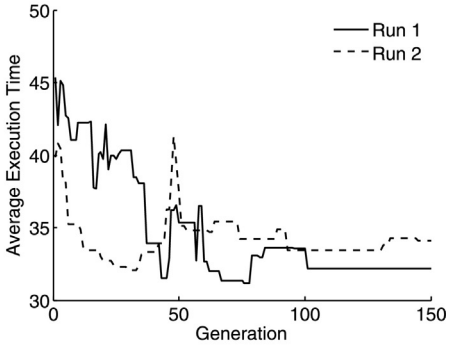

(d) Avg. execution time evolution

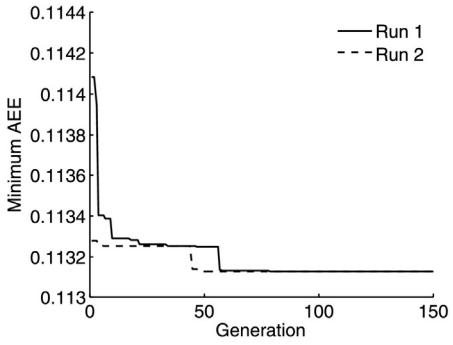

(e) Min. AEE evolution

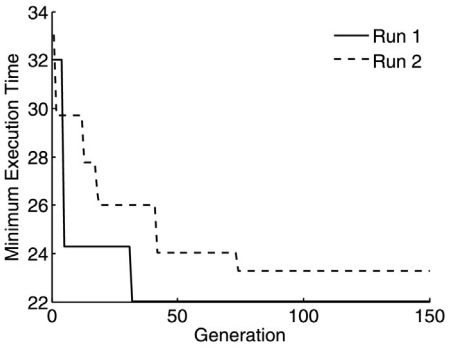

(f) Min. execution time evolution

Fig. 6. Test of repeatability for multi-objective optimization of optical flow. Results given for two executions, considering the LDOF method, one variable (σ) and one sequence of the Middlebury data set (RubberWhale). (a) Final Pareto fronts for both runs. (b) Evolution of the crowding measure in dependence of the generation number. (c) Evolution of the average AEE. (d) Evolution of the average execution time. (e) Evolution of the minimum AEE. (f) Evolution of the minimum execution time.

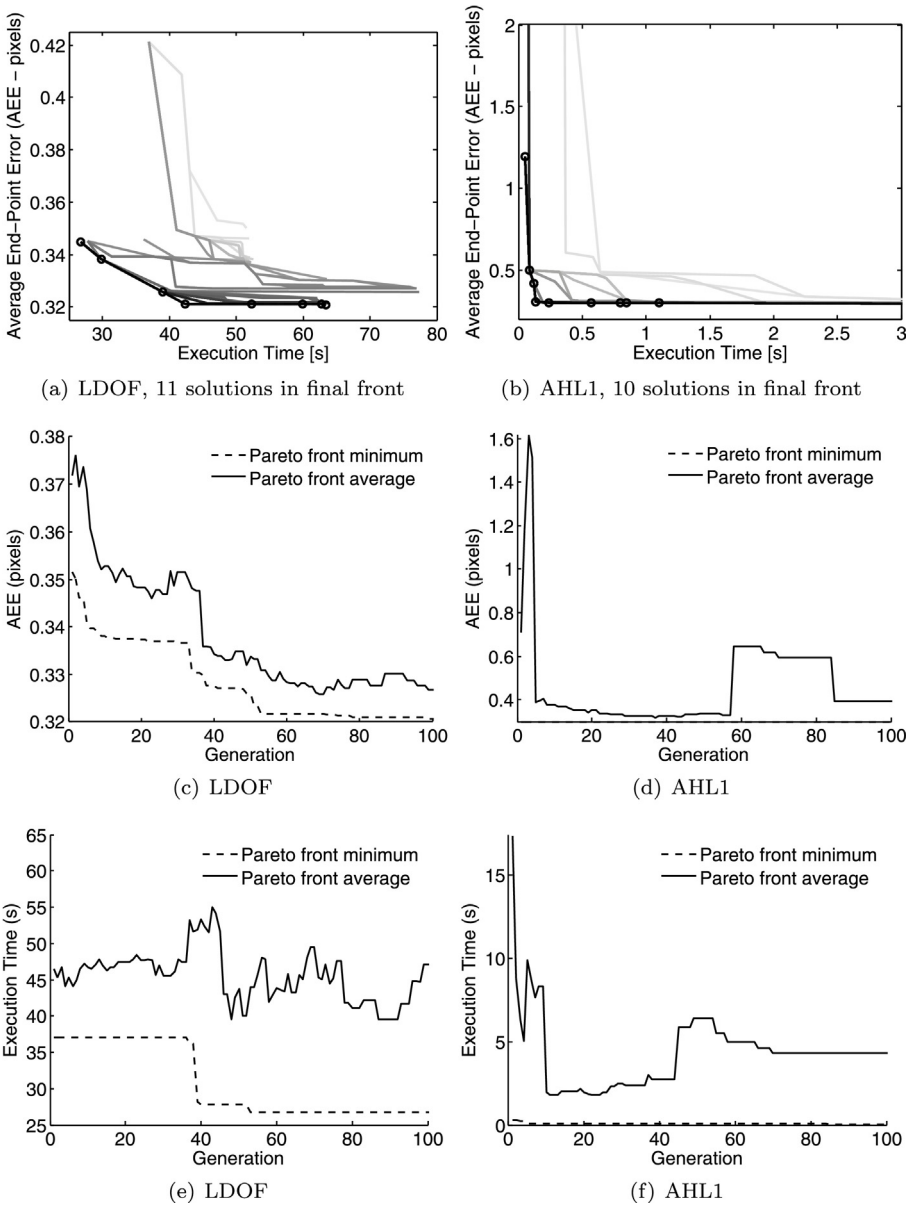

(a) LDOF, 11 solutions in final front (b) AHL1, 10 solutions in final front

(c) LDOF (d) AHL1

(e) LDOF (f) AHL1

Fig. 7. Multi-objective optical flow optimization using the Middlebury benchmark. Evolution of the Pareto front for the (a) LDOF and the (b) AHL1 optical flow methods. The lightest gray corresponds to the Pareto front for the first generation. Darker curves represent successive generations, and the dashed black curve is the result for the latest generation. Obtained Pareto solutions are shown as circles. (c and d) Evolution of the AEE. (e and f) Evolution of execution time.

4.4. Comparison and characterization of optical flow methods

Fig. 9 compares the Pareto fronts of the optical flow methods tested on both benchmarks. Fig. 9(a) shows that in the Middlebury benchmark the LDOF approach is Pareto-dominated (recall Definition 1) by the AHL1 approach. That is, every point on the final LDOF's operating curve is improved by at least one point on the final AHL1's operating curve. Then, the Pareto front of the AHL1 approach can be considered as the optimal operating curve for the Middlebury benchmark. Note, however, that the difference between the lowest errors (*AEE*) of both methods is not very significant, while the difference in execution time (*T*) is considerable. Fig. 9(b) shows that none of the methods dominates the other in the KITTI benchmark. Thus, the optimal operating curve for the KITTI benchmark would include points from the Pareto fronts of the two methods. We can then characterize the different motion estimation methods for a given application (or benchmark) by the coverage of their Pareto fronts in the *AEE*–*T* objective space. The proposed framework enables combining operating curves from several optical flow methods to cover the whole objective space. Then, attending to the accuracy-speed constraints of a given application, a point in the optimal operating curve can be chosen, which specifies the utilization of an optical flow method and its parameter configuration.

The selection of one specific solution from this set of non-dominated solutions is a very interesting topic. Multi-criteria decision making could be applied for that goal. In that sense, we chose to work for *a posteriori* articulation of preferences [41], that is, the decision makers will define their preferences after the optimization. One decision making process for that should take into account:

- Considering the cost structure of the specific application field, if available. Is it possible to assign costs to optical flow error and execution time? In that case, a weighted sum of objectives can be used to score solutions in the Pareto-front, without running a new optimization.

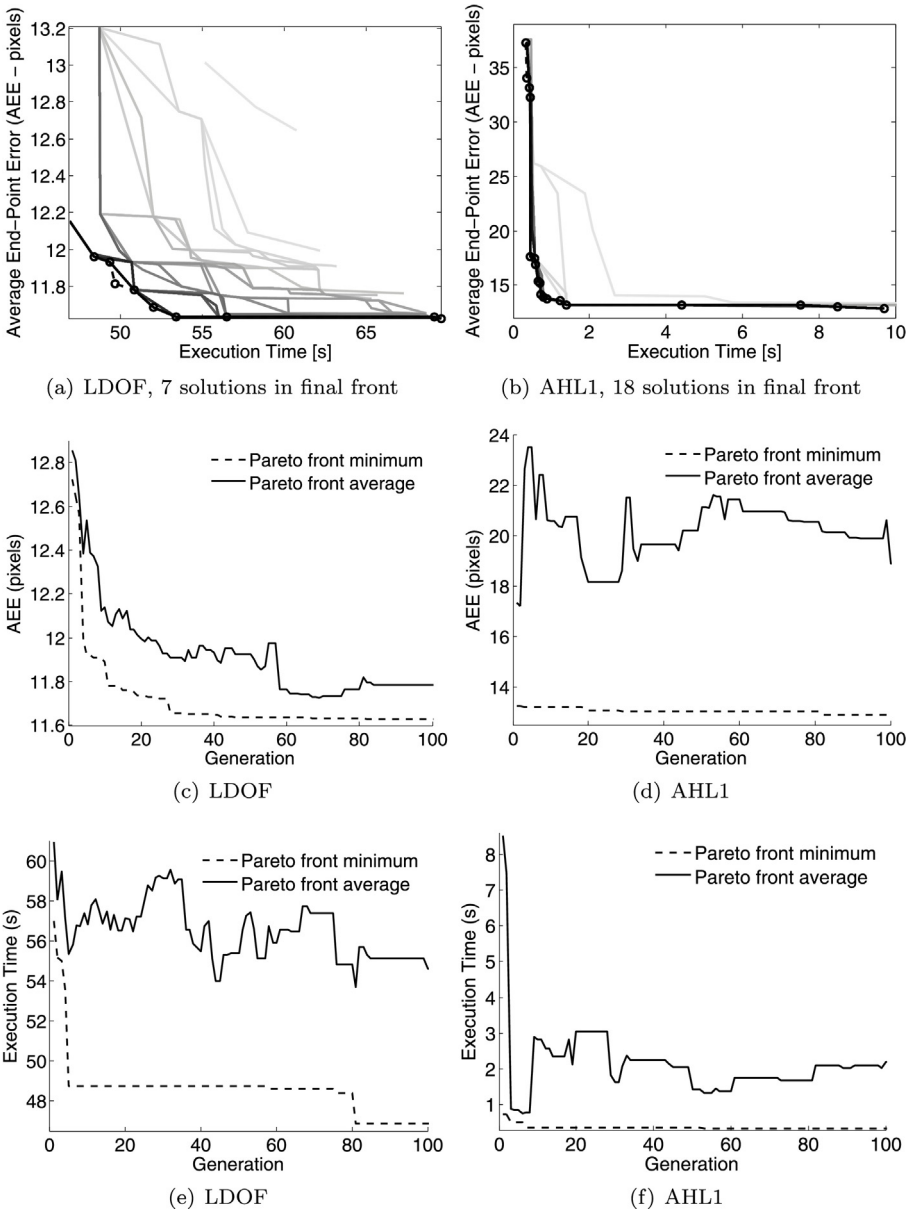

(a) LDOF, 7 solutions in final front (b) AHL1, 18 solutions in final front

(c) LDOF (d) AHL1

(e) LDOF (f) AHL1

Fig. 8. Multi-objective optical flow optimization using the KITTI benchmark. Evolution of the Pareto front for the (a) LDOF and the (b) AHL1 optical flow methods. The lightest gray corresponds to the Pareto front for the first generation. Darker curves represent successive generations, and the black curve is the result for the latest generation. Obtained Pareto solutions are shown as small circles. (c and d) Evolution of the AEE. (e and f) Evolution of execution time.

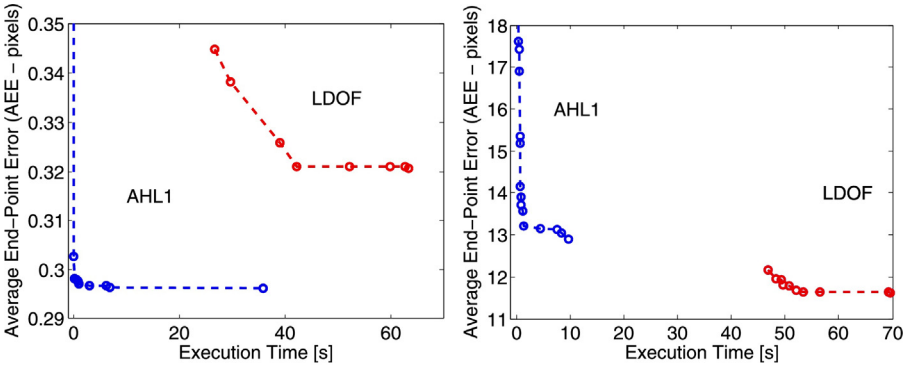

(a) Middlebury benchmark (b) KITTI benchmark

Fig. 9. Comparison of the obtained Pareto fronts for optical flow methods LDOF and AHL1 in the (a) Middlebury and the (b) KITTI benchmarks. Note that while in the Middlebury dataset AHL1 dominates LDOF, in the KITTI dataset no method is dominated.

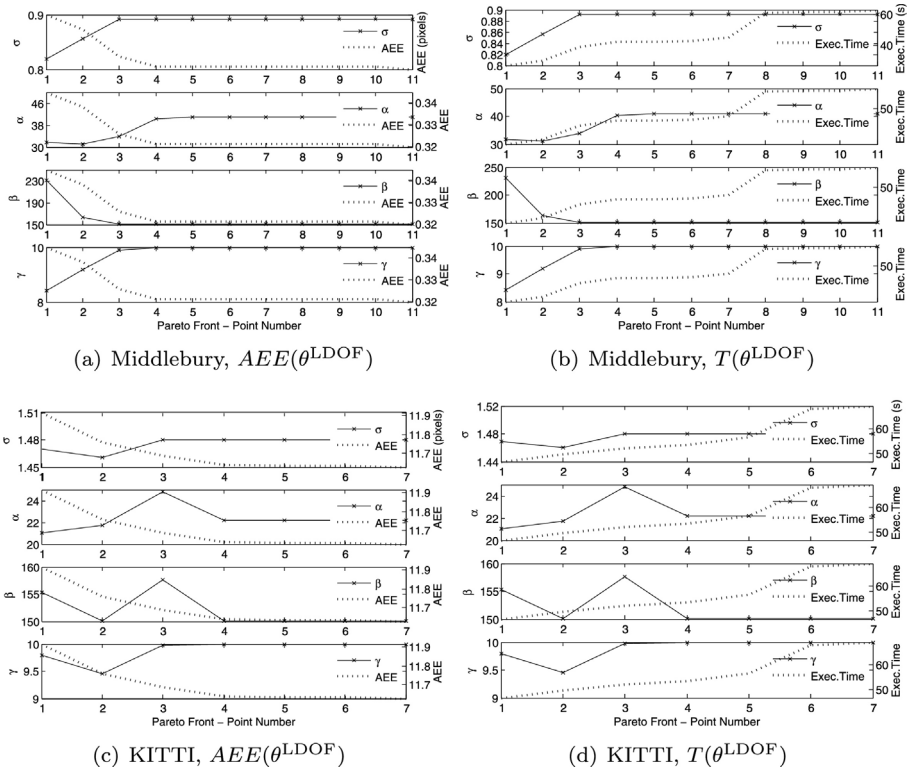

(a) Middlebury, $AEE(\theta^{\mathrm{LDOF}})$ (b) Middlebury, $T(\theta^{\mathrm{LDOF}})$

(c) KITTI, $AEE(\theta^{\mathrm{LDOF}})$ (d) KITTI, $T(\theta^{\mathrm{LDOF}})$

Fig. 10. Parameter variation along the final Pareto front of the LDOF method parameterized by $\theta^{\mathrm{LDOF}} = (\sigma, \alpha, \beta, \gamma)^{\top}$ tested on the (a, b) Middlebury and the (c, d) KITTI benchmarks. The operating points on the Pareto curve were sorted according to execution time in ascending order and then assigned a point number.

- Considering the accuracy and speed constraints for the specific application. Is there a constraint for optical flow error or execution time?

It is also interesting to observe in Fig. 9 that the shape of the Pareto front of each motion estimation approach is quite similar in both benchmarks. However, the relative position of the juxtaposed operating curves in the objective space depends on the performance of each method on the specific benchmark. This is directly related to the model assumptions and implementation of the optical flow methods. The LDOF approach models large displacements with special constraints and refrains from explicitly linearizing the energy term. Despite the accuracy of this method for large motions, it is considerably slower than the AHL1 approach. The latter does linearize its model assumptions, which disables its capability to deal with large motions. However, that enables parallelization of the implementation along with the use of graphics processing units (GPUs) to accelerate the computations.

4.5. Parameter analysis

Figs. 10 and 11 plot the parameter variation along the final Pareto front of both optical flow methods, respectively, tested on the Middlebury and the KITTI benchmarks. Note that the operating points on the Pareto curve are sorted according to execution time in ascending order and assigned a point number. We can then inspect how the different optical flow parameters change along the Pareto front. Table 2 shows the relative range

$$\Delta_x = \frac{x_{\max} - x_{\min}}{x_{\max}} \times 100\% \tag{16}$$

for the parameters $\theta^{\mathrm{LDOF}} = (\sigma, \alpha, \beta, \gamma)^{\top}$ and $\theta^{\mathrm{AHL1}} = (sf, nw, nit)^{\top}$. Similarly, Table 3 shows the relative range for the objectives $AEE(\theta)$ and $T(\theta)$. We observe that the parameters of the LDOF method present more variability in the Middlebury benchmark than in the KITTI benchmark, which also occurs with the objectives *AEE* and *T*. On the other hand, the parameters of the AHL1 method fluctuate much more in the KITTI benchmark, and presents similar variability in both objectives.

The LDOF method is less susceptible to changes in the parameter σ, which adjusts the amount of pre-smoothing of the frame sequence. In terms of *AEE*, it performs better for the large displacements in the KITTI benchmark, where the parameter α – controlling the regularity of the estimated motion field – presents the largest variability (15%). In the Middlebury benchmark the parameter β is the most volatile (34.9%), which regulates the influence of matched correspondences in the frame sequence. Overall, it is interesting to note that for increasing (σ, α, γ) and decreasing β, the end-point error goes down while the execution time goes up.

Table 2
Relative range of the optical flow parameters along the Pareto front.

Method	Benchmark	Δ_σ	Δ_α	Δ_β	Δ_γ
LDOF	Middlebury	8.1%	24.3%	34.9%	15.9%
	KITTI	1.3%	15.0%	4.9%	5.4%
		Δ_{sf}		Δ_{nw}	Δ_{nit}
AHL1	Middlebury	10.7%		14.3%	71.9%
	KITTI	92.0%		93.8%	81.0%

Table 3
Relative range of the objectives *AEE* and *T* along the Pareto front.

Method	Benchmark	Δ_{AEE}	$\Delta_{\Delta T}$
LDOF	Middlebury	7.3%	57.8%
	KITTI	2.4%	33.2%
AHL1	Middlebury	83.7%	99.2%
	KITTI	65.4%	96.6%

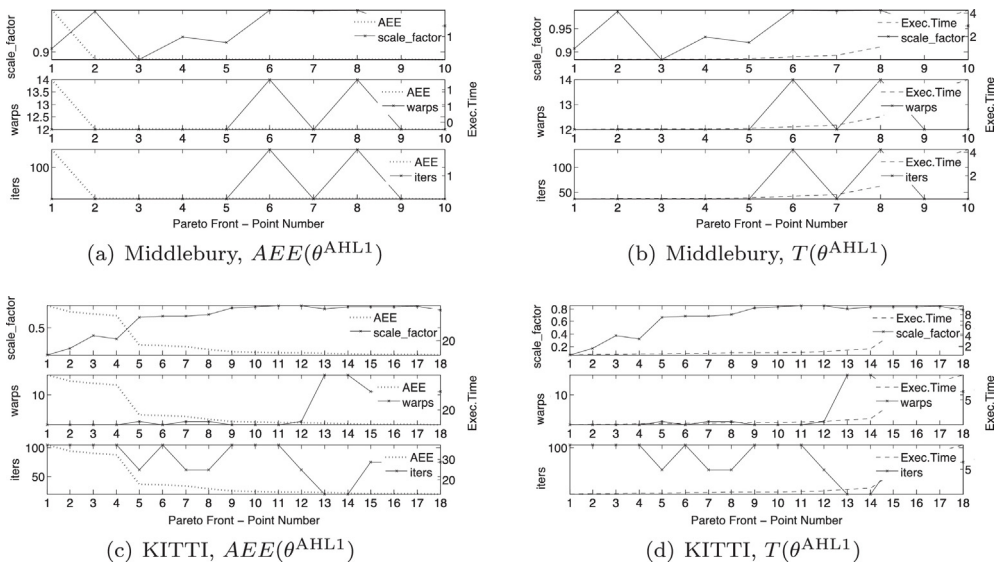

(a) Middlebury, $AEE(\theta^{\mathrm{AHL1}})$ (b) Middlebury, $T(\theta^{\mathrm{AHL1}})$

(c) KITTI, $AEE(\theta^{\mathrm{AHL1}})$ (d) KITTI, $T(\theta^{\mathrm{AHL1}})$

Fig. 11. Parameter variation along the final Pareto front of the AHL1 method parameterized by $\theta^{\mathrm{AHL1}} = (sf, nw, nit)^{\top}$ tested on the (a, b) Middlebury and the (c, d) KITTI benchmarks. The operating points on the Pareto curve were sorted according to execution time in ascending order and then assigned a point number.

In the AHL1 method the number of iterations (*nit*) has a large volatility in both benchmarks. The influence of the other two parameters (*sf*, *nw*) varies: they are more stable in Middlebury where good performance is achieved in terms of *AEE*, but they fluctuate considerably more in KITTI – parameter space exhaustively explored – in an attempt to improve performance on these challenging sequences. We observed that increasing (*sf*, *nw*, *nit*) leads to higher accuracy and higher execution times. However, the strength of this method lies in its ability to run very fast using GPU hardware.

As we have seen, the optical flow methods here studied have strengths and weaknesses. Understanding the relationships amongst the different model parameters can aid in choosing an adequate parameterization for these methods, attending to specific accuracy-speed constraints and to the type of frame sequences dealing with.

5. Conclusions and future perspectives

The proposed multi-objective optimization framework is useful to obtain the Pareto fronts for different optical flow methods. These fronts are formed of operating points in the accuracy-speed objective space, which are related to their corresponding parameter configurations. The fronts can be used as ROC curves to easily compare the methods and to characterize their performance by the extension of the Pareto fronts in the objective space. Moreover, the Pareto-dominance concept allows the combination of fronts from multiple optical flow methods into a unified optimal operating curve. Then, for a motion estimation task similar to those studied in this article, it would be possible to select the most suitable configurations from the optimal operating curve, which indicate the method and parameters to be used for that specific task.

We have shown the benefits of using our optical flow optimization approach with benchmark data to evaluate the degree to which the estimated motion field matched the ground truth motion. A natural extension of this work is to consider the general case when no ground truth is available, which would impede computing error metrics such as the *AEE* (average end-point error). To solve this problem, we need to define an error or distance metric that would work with any image sequence independently of ground truth availability. One possibility is to consider the *displaced frame difference* [42], which regards constancy assumptions such as Eq. (2) or Eq. (3) as global modelling error. A more robust alternative is the confidence measure proposed in [43], defined inversely proportional to the local energy contribution considering all terms in the energy functional. It would also be interesting to explore the 17 confidence measures recently examined in [44] in the context of stereo matching. Similarly, the proposed methodology could be applied to other computer vision problems where two or more performance measures can be defined and their joint optimization results advantageous.

Another extension of this work would be the development of novel multi-objective rankings for optical flow methods. Currently, most available rankings such as in Middlebury and KITTI juxtapose methods based on the fulfillment of single objectives independently, which does not provide any guidance toward the goal of choosing the best performing method – and its parameter configuration – in terms of the joint fulfillment of multiple objectives. The availability of multi-objective rankings would make possible to discard optical flow methods which are known to be Pareto-dominated by other ones, so that the focus is on those that fulfill multiple application-specific requirements and whose Pareto fronts provide recommendations on the most suitable parameterizations. Moreover, even before the development of multi-objective rankings, the concepts of Pareto-dominance and Pareto front can be applied to available data from particular problems.

Acknowledgements

This research was partially funded by Universidad de los Andes FAI grant #05/2013, and the FONDECYT Projects 1130153 and 3120218 (CONICYT, Chile).

References

[1] B.D. Lucas, T. Kanade, An iterative image registration technique with an application to stereo vision (DARPA), in: Proceedings of the 1981 DARPA Image Understanding Workshop, 1981, pp. 121–130.
[2] B.K. Horn, B.G. Schunck, Determining optical flow, Artif. Intell. 17 (1981) 185–203.
[3] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. Black, R. Szeliski, A database and evaluation methodology for optical flow, Int. J. Comput. Vis. 92 (2011) 1–31.
[4] A. Geiger, P. Lenz, R. Urtasun, Are we ready for autonomous driving? The KITTI vision benchmark suite, in: Computer Vision and Pattern Recognition (CVPR), Providence, USA, 2012, pp. 3354–3361.
[5] D. Butler, J. Wulff, G. Stanley, M. Black, A naturalistic open source movie for optical flow evaluation, in: ECCV, 2012, pp. 611–625.

[6] C.-h. Chen, L.-F. Pau, P.S.-P. Wang, Handbook of Pattern Recognition and Computer Vision, World Scientific, 2010.
[7] E.R. Davies, Computer and Machine Vision: Theory, Algorithms, Practicalities, Academic Press, 2012.
[8] A. Bruhn, J. Weickert, C. Feddern, T. Kohlberger, C. Schnörr, Variational optical flow computation in real time, IEEE Trans. Image Process. 14 (5) (2005) 608–615.
[9] M. Werlberger, T. Pock, H. Bischof, Motion estimation with non-local total variation regularization, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2010, pp. 2464–2471.
[10] P. Gwosdek, H. Zimmer, S. Grewenig, A. Bruhn, J. Weickert, A highly efficient GPU implementation for variational optic flow based on the Euler-Lagrange framework, in: K. Kutulakos (Ed.), Trends and Topics in Computer Vision, Vol. 6554 of Lecture Notes in Computer Science, Springer, Berlin/Heidelberg, 2012, pp. 372–383, http://dx.doi.org/10.1007/978-3-642-35740-4_29.
[11] R. Ranftl, K. Bredies, T. Pock, Non-local total generalized variation for optical flow estimation, in: Computer Vision-ECCV 2014, Springer, 2014, pp. 439–454.
[12] Y. Li, D. Huttenlocher, Learning for optical flow using stochastic optimization, in: D. Forsyth, P. Torr, A. Zisserman (Eds.), European Conference on Computer Vision – ECCV 2008, Vol. 5303 of Lecture Notes in Computer Science, Springer, Berlin/Heidelberg, 2008, pp. 379–391.
[13] K. Krajsek, R. Mester, A maximum likelihood estimator for choosing the regularization parameters in global optical flow methods, in: 2006 IEEE International Conference on Image Processing, 2006, pp. 1081–1084, http://dx.doi.org/10.1109/ICIP.2006.312743.
[14] P. Heas, C. Herzet, E. Memin, Bayesian inference of models and hyperparameters for robust optical-flow estimation, IEEE Trans. Image Process. 21 (4) (2012) 1437–1451, http://dx.doi.org/10.1109/TIP.2011.2179053.
[15] C. Sun, Fast optical flow using 3D shortest path techniques, Image Vis. Comput. 20 (13–14) (2002) 981–991, http://dx.doi.org/10.1016/S0262-8856(02)00112-9.
[16] D. Gibson, M. Spann, Robust optical flow estimation based on a sparse motion trajectory set, IEEE Trans. Image Process. 12 (4) (2003) 431–445.
[17] D. Sun, S. Roth, M.J. Black, A quantitative analysis of current practices in optical flow estimation and the principles behind them, Int. J. Comput. Vis. 106 (2) (2014) 115–137.
[18] D.R. Pereira, J. Delpiano, J.P. Papa, On the optical flow model selection through metaheuristics, EURASIP J. Image Video Process. 1 (2015), http://dx.doi.org/10.1186/s13640-015-0066-5.
[19] L. Chittka, A.G. Dyer, F. Bock, A. Dornhaus, Psychophysics: bees trade off foraging speed for accuracy, Nature 424 (6947) (2003) 388, http://dx.doi.org/10.1038/424388a.
[20] K. Deb, A. Kumar, Real-coded genetic algorithms with simulated binary crossover: studies on multimodal and multiobjective problems, Complex Syst. 9 (6) (1995) 431–454.
[21] R. Marler, J. Arora, Survey of multi-objective optimization methods for engineering, Struct. Multidiscip. Optim. 26 (6) (2004) 369–395, http://dx.doi.org/10.1007/s00158-003-0368-6.
[22] T. Weise, Global Optimization Algorithms – Theory and Application, it-weise.de (self-published), Germany, 2009.
[23] T. Bäck, Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms, Oxford University Press, 1996.
[24] D.E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning, Addison-Wesley, 1989.
[25] A. Zhou, B.-Y. Qu, H. Li, S.-Z. Zhao, P.N. Suganthan, Q. Zhang, Multiobjective evolutionary algorithms: a survey of the state of the art, Swarm Evolut. Comput. 1 (1) (2011) 32–49, http://dx.doi.org/10.1016/j.swevo.2011.03.001.
[26] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Trans. Evolut. Comput. 6 (2) (2002) 182–197, http://dx.doi.org/10.1109/4235.996017.
[27] F. Valdez, P. Melin, O. Castillo, An improved evolutionary method with fuzzy logic for combining particle swarm optimization and genetic algorithms, Appl. Soft Comput. 11 (2) (2011) 2625–2632, http://dx.doi.org/10.1016/j.asoc.2010.10.010.
[28] R.-C. David, R.-E. Precup, E.M. Petriu, M.-B. R?dac, S. Preitl, Gravitational search algorithm-based design of fuzzy control systems with a reduced parametric sensitivity, Inf. Sci. 247 (2013) 154–173, http://dx.doi.org/10.1016/j.ins.2013.05.035.
[29] A.-C. Z?voianu, G. Bramerdorfer, E. Lughofer, S. Silber, W. Amrhein, E. Peter Klement, Hybridization of multi-objective evolutionary algorithms and artificial neural networks for optimizing the performance of electrical drives, Eng. Appl. Artif. Intell. 26 (8) (2013) 1781–1794, http://dx.doi.org/10.1016/j.engappai.2013.06.002.
[30] M. Everingham, H. Muller, B. Thomas, Evaluating image segmentation algorithms using the Pareto front, in: A. Heyden, G. Sparr, M. Nielsen, P. Johansen (Eds.), Computer Vision – ECCV 2002, Vol. 2353 of Lecture Notes in Computer Science, Springer, Berlin/Heidelberg, 2002, pp. 255–259.
[31] R. Verschae, J. Ruiz-del-Solar, M. Köppen, R.V. Garcia, Improvement of a face detection system by evolutionary multi-objective optimization, in: Proceedings of the Fifth International Conference on Hybrid Intelligent Systems, IEEE Computer Society, Washington, DC, USA, 2005, pp. 361–366, http://dx.doi.org/10.1109/ICHIS.2005.63.
[32] X. Benlian, W. Zhiquan, A multi-objective-ACO-based data association method for bearings-only multi-target tracking, Commun. Nonlinear Sci. Numer. Simul. 12 (8) (2007) 1360–1369, http://dx.doi.org/10.1016/j.cnsns.2006.02.007.
[33] I. Vite-Silva, N. Cruz-Cortés, G. Toscano-Pulido, L.G. Fraga, Optimal triangulation in 3D computer vision using a multi-objective evolutionary algorithm, in: Proceedings of the EvoWorkshops 2007: Applications of Evolutionary Computing, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 330–339.
[34] J. Salmen, L. Caup, C. Igel, Real-time estimation of optical flow based on optimized Haar wavelet features, in: Evolutionary MultiCriterion Optimization, Springer, 2011, pp. 448–461.
[35] T. Brox, J. Malik, Large displacement optical flow: descriptor matching in variational motion estimation, IEEE Trans. Pattern Anal. Mach. Intell. 33 (3) (2011) 500–513.
[36] M. Werlberger, W. Trobin, T. Pock, H. Bischof, A. Wedel, D. Cremers, Anisotropic Huber-L1 optical flow, in: Proceedings of the British Machine Vision Conference, BMVC 2009, 2009, pp. 1–11.
[37] J. Delpiano, L. Pizarro, R. Verschae, J. Ruiz-del-Solar, Multi-objective optimization for characterization of optical flow methods, in: Proceedings of the 9th International Conference on Computer Vision Theory and Applications – VISAPP, 2014, pp. 566–573.
[38] A. Bruhn, J. Weickert, C. Schnörr, Lucas/Kanade meets Horn/Schunck: combining local and global optic flow methods, Int. J. Comput. Vis. 61 (3) (2005) 1–21.
[39] T. Brox, A. Bruhn, N. Papenberg, J. Weickert, High accuracy optical flow estimation based on a theory for warping Proceedings of the European Conference on Computer Vision, ECCV 2004, vol. 4, 2004, pp. 25–36.
[40] J. Weickert, A. Bruhn, T. Brox, N. Papenberg, A survey on variational optic flow methods for small displacements, in: O. Scherzer (Ed.), Mathematical Models for Registration and Applications to Medical Imaging, Springer, Berlin, 2006, pp. 103–136.
[41] K. Deb, M. Koksalan, Guest editorial special issue on preference-based multiobjective evolutionary algorithms, IEEE Trans. Evolut. Comput. 14 (5) (2010) 669–670, http://dx.doi.org/10.1109/TEVC.2010.2070371.
[42] G. Le Besnerais, F. Champagnat, Dense optical flow by iterative local window registration Proceedings of the International Conference on Image Processing, vol. 1, 2005, pp. 137–140.
[43] A. Bruhn, J. Weickert, A confidence measure for variational optic flow methods, in: R. Klette, R. Kozera, L. Noakes, J. Weickert (Eds.), in: Geometric Properties for Incomplete Data, vol. 31, Springer, Netherlands, 2006, pp. 283–298, http://dx.doi.org/10.1007/1-4020-3858-8_15.
[44] X. Hu, P. Mordohai, A quantitative evaluation of confidence measures for stereo vision, IEEE Trans. Pattern Anal. Mach. Intell. 34 (11) (2012) 2121–2133, http://dx.doi.org/10.1109/TPAMI.2012.46.