Image and Vision Computing 29 (2011) 484–498

Contents lists available at ScienceDirect

Image and Vision Computing

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

Evolutionary-computer-assisted design of image operators that detect interest points using genetic programming☆

Gustavo Olague [a,*], Leonardo Trujillo [a,b]

[a] Proyecto Evovisión, Departamento de Ciencias de la Computación, División de Física Aplicada, Centro de Investigación Científica y de Educación Superior de Ensenada, Km. 107 Carretera Tijuana-Ensenada, 22860, Ensenada, BC, Mexico

[b] Instituto Tecnológico de Tijuana, Av. Tecnológico S/N, Fracc. Tomás Aquino, Tijuana, B.C., Mexico

ARTICLE INFO

Article history:
Received 3 May 2010
Received in revised form 29 January 2011
Accepted 22 March 2011

Keyword:
Interest points
Computer assisted design
Evolutionary computation
Genetic programming
Evolutionary computer vision

ABSTRACT

This work describes a way of designing interest point detectors using an evolutionary-computer-assisted design approach. Nowadays, feature extraction is performed through the paradigm of interest point detection due to its simplicity and robustness for practical applications such as: image matching and view-based object recognition. Genetic programming is used as the core functionality of the proposed human-computer framework that significantly augments the scope of interest point design through a computer assisted learning process. Indeed, genetic programming has produced numerous interest point operators, many with unique or unorthodox designs. The analysis of those best detectors gives us an advantage to achieve a new level of creative design that improves the perspective for human-machine innovation. In particular, we present two novel interest point detectors produced through the analysis of multiple solutions that were obtained through single and multi-objective searches. Experimental results using a well-known testbed are provided to illustrate the performance of the operators and hence the effectiveness of the proposal.

1. Introduction

Computer vision (CV) is concerned with the development of artificial systems that can automatically analyze and interpret visual information, and current systems have obtained impressive performance on many high-level tasks that include object recognition [15], object detection [2], image classification [7], image retrieval [36], object categorization [12], and 3D reconstruction [17,57]. Nevertheless, the answer to the conundrum of artificial vision from the point of view of machine intelligence presents fundamental problems with underlying difficulties that attract continuously the interest of researchers from diverse fields of research, such as pattern recognition, artificial intelligence, and cognitive science, to name but a few. In particular, researchers from evolutionary computation are working actively in numerous theoretical and practical CV problems, see [3,4,39].

In this way, for many domains of technological and scientific endeavor, solutions to particular problems are traditionally the outcome of a detailed design process undertaken by a group of human experts, and in this respect CV is not an exception. Problem solving requires that a scientist or engineer makes a series of design choices in order to produce a final solution [42]. Therefore, the attributes of a particular solution will depend upon the initial assumptions that are made, and on the overall understanding that the human expert possesses regarding the nature of the problem domain. One shortcoming for this approach to problem solving is that sometimes, when a different type of solution is desired, or required, then the design process must be changed and executed once more, using a different set of assumptions and analytical perspectives. As a result, a large number of competing proposals can exist for what appear to be very basic and simple problems; for example, the problem of interest point detection [64]. The design of candidate solutions can also be understood as an informed search process. For instance, in the scenario described above the search is guided by human expertise and operates within a domain-specific space of possible solutions. When the domain of a problem is well-known, then some useful properties of the search space could conceivably be inferred in order to improve the search process. However, even if this is the case, and for difficult problems it cannot be assumed, the space is normally very large, complex and non-linear [42].

In this work, we follow the genetic programming framework to extend the traditional approach for problem solving just described. Indeed, from all evolutionary-based methodologies inspired by biological evolution genetic programming provides a framework to find computer programs that perform a user-defined task. This technique is a powerful machine learning approach that is still largely unknown in the computer vision literature. This paper takes a further step in the traditional way of designing CV programs using what we call an evolutionary computer assisted design (E-CAD) concept. The

☆ This paper has been recommended for acceptance by Maja Pantic, PhD.

* Corresponding author. Tel.: +526461750500 x 23429.
E-mail addresses: olague@cicese.mx (G. Olague), leonardo.trujillo.ttl@gmail.com (L. Trujillo).

0262-8856/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.imavis.2011.03.004

idea is to explore the design space of a very specific and useful task found in many machine vision systems, such as the interest point detection, using the genetic programming technique. Later, when numerous design solutions are available the decision maker could create novel human designs from the set of human-competitive machine intelligence solutions. Thus, the human designer is given the possibility to create new designs that are far away of the abilities of a designer using more traditional approaches. Indeed, the advantage is clear because as we will show in the experiments the designer starts a new design stage from a whole set of competitive designs; thus, closing the human-machine invention cycle within a repeating process that can be easily reconfigured to produce new suitable designs according to the requirements of the task. This is possible because a priori knowledge is easily incorporated within the genetic programming framework that uses an adaptive learning paradigm to approach the curse of dimensionality.

The goal of finding the optimal solution for real-world problems is usually a quite arduous task; in particular, our work follows the approach of artificial evolution. Thus, one of the main keys to set up such a proposal is to begin with a clearly specified problem where it is often much easier to propose a performance or evaluation criteria that could help to define a measure of optimality. In this case, the search could then be carried out automatically by a computer algorithm that exploits the information that such measures of performance can return. In this way, the search could be carried out using evolutionary computation (EC), a population based meta-heuristic based on the principles of artificial Darwinism that have shown great success at exploring large search spaces; thus, producing solutions that are well adapted to the prescribed objectives [9,19,24]. Indeed, EC is based on the core principles of biological evolution, a natural process that exhibits an adaptive power that by far outstrips that of any human-engineered system [54]. Currently, a large amount of experimental evidence exists that confirms the ability of EC to outperform man-made solutions in many domains, such as antenna design, mathematical proofs, and even CV [4,25,26,44,56,62]. In many cases, the stochastic nature of EC allows it to sample large portions of the search space, and sometimes produce solutions that might not be evident to a human expert. For example, the case of network design in photogrammetry in which the design of a specific network, considered as not atypical, was rediscovered by means of evolutionary computing [38,41]. Moreover, in our previous work on the design of interest point detectors genetic programming was able to rediscover Beaudet's detector [60]. However, we do not suggest that the use of machine learning in general, or EC in particular, should completely substitute the design work that a human expert could perform. On the contrary, we agree with the argument that a more complete strategy would cooperatively include both methodologies, thereby blending the complimentary skills of each [55], in what others have called a computer assisted design (CAD) process [34,42].

In this paper, we employ an E-CAD based approach in the search for optimal image operators that detect low-level features known as interest points [64]. We use genetic programming, one of the more advanced forms of EC, to automatically synthesize candidate solutions that can be represented using tree structures. The evolutionary search is guided by two performance criteria, the geometric and photometric stability of detected points given by the repeatability rate [50], and a measure of how disperse the set of detected points are over the image plane [64]. In order to achieve a design these objectives are concurrently considered using two different techniques: first, both criteria are included into a single objective function and the search returns the best single solution found; and second, we pose a multi-objective problem that searches for a diverse set of Pareto optimal solutions. In each case, we use E-CAD to propose novel interest point detectors using the operators that the evolutionary algorithm generates. The first one is characterized by its simplicity and the high performance it achieves on standard tests; we call it the *Gaussian Intensity Neighborhood* (GIN) interest point detector. On the other hand, the other is a parameterized operator for interest point detection that allows for fine grained control of the amount of point dispersion without sacrificing the geometric stability; we call it the *Multi-Objective Parameterized* (MOP) interest point detector, and to our knowledge it is unique in CV literature.

The remainder of this paper is organized as follows. Section 2 presents the basic problem of interest point detection, reviews previous work and defines the performance criteria. Genetic programming is introduced in Section 3, and a review of applications to CV is outlined. Section 4 describes the single objective approach to evolutionary-computer-assisted design of interest point operators and introduces the Gaussian Intensity Neighborhood detector. Then, the multi-objective approach is presented in Section 5 and the Multi-Objective Parameterized interest point detector is explained. Finally, Section 7 contains a brief summary and outlines possible lines of future research.

2. Interest point detection

Currently, many CV systems employ a local approach to feature extraction and description, by focusing on small and highly invariant features called interest points [31,33,49,50,64]. It is also important to understand that the performance of these systems directly depends on the quality of the underlying detection and description algorithms that are used. Keeping to the former, there are dozens of proposed interest point detectors available in CV literature, most of which are the direct product of a human-based approach to problem solving and/or design.

Using the taxonomy of local features given in [50,64], we can say that interest points are detected by algorithms that focus on image intensity values and only make weak assumptions regarding the underlying structure of the observed scene [64]. Interest points are salient image pixels that are unique and distinctive; i.e., they are quantitatively and qualitatively different from other points, and they normally represent only a small fraction of the total image area [32,50].

2.1. Problem definition

A measure of how salient or interesting each pixel is can be obtained using a mapping of the form $K(x):\mathbb{R}^{+}\rightarrow\mathbb{R}$ which we call an interest point operator. Each interest point detector will employ a different operator K; in this way, a *detector* refers to the complete algorithmic process that extracts interest points, while an *operator* only computes the corresponding interest measure. Applying K to an image I produces what can be called an *interest image* I^{*}, see Fig. 1. Afterwards, most detectors follow the same basic process: non-maxima suppression that eliminates pixels that are not local maxima, and a thresholding step that obtains the final set of points. Therefore, a pixel $\mathbf{x}$ is tagged as an interest point if the following conditions hold,

$$K(\mathbf{x}) > max\{K(\mathbf{x_W})\,|\,\forall \mathbf{x_W}\in\mathbf{W}, \mathbf{x_W}\neq\mathbf{x}\}\wedge K(\mathbf{x})>h, \tag{1}$$

where $\mathbf{W}$ is a square neighborhood of size $n\times n$ around $\mathbf{x}$, and h is an empirically defined threshold. The first condition in Eq. (1) accounts for non-maximum suppression and the second is the thresholding step, the process is shown in Fig. 1. Experiments in the current work use $n=5$, while h depends on the operator.

2.2. Previous proposals

In this section only a brief overview of previous work is given, a thorough discussion is beyond the scope of this paper. For example, there exists a group of interest point detectors that employ operators that are based on the *auto-correlation* or *second-moment matrix*. This