Expert Systems with Applications 40 (2013) 5240–5259

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

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

DS-DPSO: A dual surrogate approach for intelligent watermarking of bi-tonal document image streams

Eduardo Vellasques *, Robert Sabourin, Eric Granger

Laboratoire d'imagerie, de vision et d'intelligence artificielle, École de Technologie Supérieure, Université du Québec, Montreal, Canada

ARTICLE INFO

Keywords:
Intelligent Watermarking
Evolutionary Computation
Particle Swarm Optimization
Surrogate-based Optimization
Gaussian Mixture Models

ABSTRACT

Intelligent watermarking (IW) techniques employ population-based evolutionary computing in order to optimize embedding parameters that trade-off between watermark robustness and image quality for digital watermarking systems. Recent advances indicate that it is possible to decrease the computational burden of IW techniques in scenarios involving long heterogeneous streams of bi-tonal document images by recalling embedding parameters (solutions) from a memory based on Gaussian Mixture Model (GMM) representation of optimization problems. This representation can provide ready-to-use solutions for similar optimization problem instances, avoiding the need for a costly re-optimization process. In this paper, a dual surrogate dynamic Particle Swarm Optimization (DS-DPSO) approach is proposed which employs a memory of GMMs in regression mode in order to decrease the cost of re-optimization for heterogeneous bi-tonal image streams. This approach is applied within a four level search for near-optimal solutions, with increasing computational burden and precision. Following previous research, the first two levels use GMM re-sampling to recall solutions for recurring problems, allowing to manage streams of heterogeneous images. Then, if embedding parameters of an image require a significant adaptation, the third level is activated. This optimization level relies on an off-line surrogate, using Gaussian Mixture Regression (GMR), in order to replace costly fitness evaluations during optimization. The final level also performs optimization, but GMR is employed as a costlier on-line surrogate in a worst-case scenario and provides a safeguard to the IW system. Experimental validation were performed on the OULU image data set, featuring heterogeneous image streams with a varying levels of attacks. In this scenario, the DS-DPSO approach has been shown to provide comparable level of watermarking performance with a 93% decline in computational cost compared to full re-optimization. Indeed, when significant parameter adaptation is required, fitness evaluations may be replaced with GMR.

1. Introduction

The decreasing costs of data transmission and storage provided numerous opportunities for sharing multimedia documents like images. This has led to the creation of a digital economy with new services that are available 24 h a day, 7 days a week, around the globe. Individuals and businesses depend more and more on sharing important documents which raises serious privacy concerns. Enforcing the security of document images is an important issue. Cryptography can solve part of this issue. However, specially with multimedia documents like images, the protection allowed by cryptography vanishes as the data has been decrypted. Digital watermarking (Cox et al., 2002) which consists of embedding image-related secret data through the manipulation of pixel values in an imperceptible manner, allows another layer of protection. Most importantly, the protection mechanism provided by digital watermarking follows the image even when it is inadvertently distributed or tampered. Enforcing the security of bi-tonal document images poses an additional challenge as bi-tonal images have lower embedding capacity and the manipulation of bi-tonal pixels is more prone to result in visual artifacts.

Digital watermarking has become an active area of research in recent years. Because of its nature, watermarking systems are subject to attacks by hackers (Voloshynovskiy et al., 2001). Robustness against attacks always comes at the cost of degradation on imperceptibility (Cox et al., 1996). Many watermarking techniques allow adjusting the trade-off between robustness and quality through the manipulation of embedding parameters. The optimal trade-off and the corresponding values vary from one image to another. To make matters worse, security requirements also vary across applications.

In a typical use case, a robust watermark is added to a document image and then a fragile watermark is added to the top of it. The robust watermark can contain, for example, information

* Corresponding author.
E-mail addresses: evellasques@livia.etsmtl.ca (E. Vellasques), robert.sabourin@etsmtl.ca (R. Sabourin), eric.granger@etsmtl.ca (E. Granger).

0957-4174/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2013.03.021

about the ownership of that document. For this reason it must survive intentional and unintentional image processing operations. The fragile watermark by its way, is easily destroyed which allows detecting that tampering has occurred. Therefore, compared to other security mechanisms such as cryptography, digital watermarking allows for a concealed type of protection. And more important, it will always follow the image (the protection provided by cryptography vanishes when the image is decrypted). The main limitation of digital watermarking is that the robustness against attacks and the corresponding image quality trade-off are defined through heuristic parameters. A poor choice of parameter can result in a robust watermark that can be easily removed. Therefore, these parameters must be adjusted for different sets of attacks, according to the embedding capacity of each image.

Manual adjustment of digital watermarking parameters is infeasible in practical applications. Evolutionary computing (EC) techniques such as Particle Swarm Optimization (PSO) (Kennedy and Eberhart, 1995) and Genetic Algorithms (Holland, 1992) have been employed in order to find embedding parameters that optimize the trade-off between image quality and watermark robustness for each image and set of attacks (Vellasques et al., 2010). In EC, a population of candidate solutions is evolved through a certain number of generations, and guided by an objective function. In intelligent watermarking (IW), objective functions are usually a combination of image quality and watermark robustness. The fitness of each candidate solution is evaluated at each generation. Each fitness evaluation requires one or more embedding, detection and attack (image processing) operations which is prohibitively expensive in industrial applications.

Recent efforts to decrease the computational cost of IW techniques for streams of document images is promising. In the Dynamic PSO (DPSO) system proposed in Vellasques et al. (2011), IW of homogeneous streams of bi-tonal document images (or problems) was formulated as a special type of dynamic optimization problem (DOP[1]). In this special formulation of DOP, a stream of document images corresponds to a stream of recurring optimization problems. A change detection mechanism assigns case-based solutions of previously-seen problem instances (associated with previous document images) to new similar problem instances (associated with new images). This significantly reduced the number of costly re-optimization operations, allowing for a significant decrease in computational burden. In the DPSO system proposed in Vellasques et al. (in press), Gaussian mixture modeling (GMM) of optimization history was employed in order to model previous optimization problems.

In that approach, which is based on the combined use of PSO and GMM, a memory of GMMs is incrementally built using fitness (phenotypic) and parameter (genotypic) information of all intermediary solutions found during each case of optimization. This memory is split in two levels: Short Term Memory (STM) and Long Term Memory (LTM). For every new image, solutions are sampled from each GMM in the memory and re-evaluated. The distributions of sampled and re-evaluated fitness values are compared with the use of a statistical test and if they are considered to be similar, the best re-evaluated solution is employed directly for the new image avoiding a costly re-optimization operation. Otherwise, re-optimization is triggered. After that, a GMM is trained using fitness and parameter values of all intermediary solutions (from all generations). The new GMM plus the global best solution will form a probe. The STM probe is replaced with the new probe. For the LTM, the new GMM is merged with the most similar GMM in the memory if their distance is below a threshold estimated on the last cases of update. Otherwise it is simply added to the LTM, replacing the least recalled probe if a memory limit has been reached. The main motivation in using a GMM is that it allows modeling multimodal fitness landscapes. Moreover, new models can be easily merged into existing models, which provides an incremental learning capability to the system. This is crucial in scenarios involving long streams of heterogeneous document images as the memory needs to adapt automatically to variations in the stream. Readers interested in understanding better how is GMM employed in the optimization of long streams of document images are referred to our previous article (Vellasques et al., in press). This approach allowed for a significant decrease in the cost for IW of **heterogeneous** streams of document images compared to the case-based approach. In both approaches, when a new optimization problem is similar to a previously-solved one, solutions in memory corresponding to that previous problem should be recalled, avoiding a costly re-optimization process.

The basic assumption behind that approach is that after a learning phase, most new problem instances will result in recall rather than re-optimization operations. However, a significant variation in the stream of optimization problems such as that caused by a new attack, will result in an increase in the number of re-optimization operations. The time complexity of re-optimization is orders of magnitude higher than that of recall. Each attempt of recalling a memory element has a time complexity comparable to a single iteration in the optimization phase, and optimization generally requires generally 50 plus iterations. Decreasing this cost is an important issue. It has been demonstrated in literature that optimization strategies based on the use of an associative memory (Yang and Yao, 2008) outperform other dynamic optimization strategies in cyclic/recurrent problems. These techniques rely on storage of high performance solutions, as well as information about their fitness landscape using a density estimate. The most common approach to associative memory optimization is to inject memory solutions in the initial population, in a strategy named memory-based immigrants (Wang et al., 2007).

One limitation of approaches based on associative memory is that for a case of re-optimization, the density estimates will only provide an initial set of candidate solutions. After that, these solutions are evolved with the use of EC and the knowledge of previous problems provided by that estimate is not explored during the optimization process whatsoever. It has been observed in the Estimation of Distribution Algorithms (EDA) literature that probabilistic models can be employed in order to guide the optimization process (Pelikan et al., 2002). A second limitation is that although memory-based immigrants can reduce the number of generations needed for convergence, still, each generation involves re-evaluating the fitness of each solution.

In surrogate-based optimization, costly fitness evaluation operations are replaced by a regression model. Sampling, model update and optimization are applied in an iterative manner. The advantage of such approach is that most of the fitness evaluations required for optimizations are performed using a regression model at a fraction of the cost of an exact fitness evaluation. There are two schools of thought: a first one that sees a surrogate as an oracle that will replace the objective function (Queipo et al., 2005) and a second one that sees a surrogate as a compact database employed in order to forecast good solutions during optimization, accelerating convergence (Parno et al., in press). Both provide different ways of addressing the trade-off between model precision and fitness evaluation cost. The first one favors decreasing fitness evaluation over precision and is preferred in situations where the model provides a precise representation of the fitness landscape and/or the computational cost of fitness evaluation is too high. The second one favors precision over decreasing fitness evaluations and is preferred in situations where the model does not provide a precise representation of the fitness landscape and/or the cost of optimization is not too

[1] In a DOP, the optimum location and/or fitness value change over time.