On the Detection of Unknown Locally Repeating Patterns in Images

Diogo Pratas* and Armando J. Pinho

Signal Processing Laboratory, IEETA/DETI
University of Aveiro, 3810-193 Aveiro, Portugal

Abstract. Detecting unknown repeated patterns that appear multiple times in a digital image is a great challenge. We have addressed this problem in a recent work and we have shown that, using a compression based approach, it is possible to find exact repetitions. In this work, we continue this study, introducing a procedure for detecting unknown repeated patterns that occur in a close vicinity. We use finite-context modelling to pinpoint the possible locations of the repetitions, by exploring the connection between lossless image compression and image complexity. Since repetitions are associated to low complexity regions, the repeating patterns are revealed and easily detected. The experimental results show that the proposed approach provides increased ability to eliminate false positives.

1 Introduction

Solomonoff, Kolmogorov, Chaitin and others have contributed in a decisive way to the theory of complexity and to its application in a number of fields [11, 12, 6, 2, 14, 10]. The Kolmogorov complexity of A, denoted by $K(A)$, is defined as the size of the smallest program that produces A and stops. Unfortunately, it is not computable and to overcome this limitation, we have to rely on approximations provided by computable measures, that provide upper bounds on the Kolmogorov complexity. Lempel-Ziv based complexity measures [7], linguistic complexity measures [5] or compression-based complexity measures [3], are some of them.

Lossless compression is an obvious choice for approximating the Kolmogorov complexity. In fact, together with the decoder, a bitstream produced by a lossless compression algorithm can be used to reconstruct the original data, and the sum of the number of bits required for representing both the decoder and the bitstream can be viewed as an estimate of the Kolmogorov complexity. In order to be suitable to approximate the Kolmogorov complexity, a compression algorithm needs to gather knowledge of the data while the compression is performed, i.e., it needs to find dependencies and to create an internal model of the

* This work was partially funded by FEDER through the Operational Program Competitiveness Factors - COMPETE and by National Funds through FCT - Foundation for Science and Technology in the context of the project FCOMP-01-0124-FEDER-022682 (FCT reference PEst-C/EEI/UI0127/2011).

data. Lempel-Ziv compression algorithms are within this class. They are also the most often used compression algorithms in compression-based complexity applications, including those reported in the imaging field [13, 9, 4]. Unfortunately, although the Lempel-Ziv compression techniques are quite effective in data sequences, they do not perform as well in the case of images.

A finite-context model provides, on a symbol by symbol basis, an information measure that corresponds in essence to the number of bits required to represent the symbol, conditioned by the accumulated knowledge of all past symbols. In a recent work [15], we have developed a method for locating exact repetitions in images using models that gather information from the whole past. In this paper, we propose an approach that relies only on local past occurrences, exploring the repeated patterns of certain regions. We use this information to build complexity surfaces, i.e., images in which the intensity of the pixels indicate how complex is the corresponding region of the original image. Patterns that occur more than once in the image tend to require less bits to encode as repetitions of these patterns are found. Since the repetitions are associated with low complexity regions in the complexity surface, they can be more easily pinpointed.

This paper is organized as follows. In Section 2, we describe the method. In Section 3, we provide experimental results, namely using the whole past occurrences and the local past occurrences. Finally, in Section 4, we draw some conclusions.

2 The proposed method

The main idea is to use finite-context modelling in order to explore the multiple approaches of the accumulated knowledge of the past, namely the whole past occurrences and the local past occurrences.

A finite-context model collects statistical information of a source. For every outcome, it assigns probability estimates to the symbols of the alphabet $\mathcal{A} = \{s_1, s_2, \ldots, s_{|\mathcal{A}|}\}$, where $|\mathcal{A}|$ denotes the size of $\mathcal{A}$. The estimates are calculated taking into account a conditioning context computed over a finite and fixed number, $k > 0$, of past outcomes $x_{n-k+1..n} = x_{n-k+1} \ldots x_{n-1}x_n$ (order-k finite-context model) [1]. In the case of image data, the notion of recent past usually refers to spatial proximity. Therefore, $x_{n-k+1..n}$ may refer to the set of the k spatially closest samples and not necessarily to the k most recently processed samples, although causality is always imposed.

The probabilities, $P(X_{n+1} = s|x_{n-k+1..n}), \forall_{s \in \mathcal{A}}$, are calculated using symbol counts that are accumulated while the image is processed. Here, the symbol counts are accumulated according to a fixed-size window of the most recent symbols, discarding the rest. Thus, when the window size is small, the model explores the local statistics. On the other hand, when the window size has the size of the past, the model explores the whole image statistics. We estimate the probabilities using

$$P(X_{n+1} = s|x_{n-k+1..n}) = \frac{C(s|x_{n-k+1..n}) + \alpha}{C(x_{n-k+1..n}) + \alpha|\mathcal{A}|}, \tag{1}$$

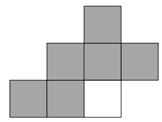

Fig. 1: Context template of order six used by the finite-context model.

where $C(s|x_{n-k+1..n})$ represents the number of times that, in the past, the information source generated symbol s having $x_{n-k+1..n}$ as the conditioning context and where

$$C(x_{n-k+1..n}) = \sum_{a \in \mathcal{A}} C(a|x_{n-k+1..n}) \tag{2}$$

is the total number of events that has occurred so far in association with context $x_{n-k+1..n}$. Parameter α allows balancing between the maximum likelihood estimator and an uniform distribution, preventing the estimator from generating zero probabilities. It can be seen that when the total number of events, n, is large, this estimator behaves essentially as a maximum likelihood estimator. For $\alpha = 1$ is the well-known Laplace estimator.

We use finite-context modelling to build complexity surfaces of the images. A complexity surface is an image ϕ, with the same geometry as the original image f, where each pixel $\phi_{i,j}$ contains the code length required to encode $f_{i,j}$ estimated by the finite-context model, i.e.,

$$\phi_{i,j} = -\log_2 P(F_{i,j} = f_{i,j}|c_{k,i,j}), \tag{3}$$

where $c_{k,i,j}$ denotes the (usually) two-dimensional order-k context and i, j the pixel coordinates. In the experiments reported in this paper, we used the context configuration depicted in Fig. 1.

The image data are scanned pixel by pixel, in raster-scan order. If a certain pattern A is found for the first time, it will require a certain number of bits to be represented, corresponding to a certain complexity. When that same pattern is seen again, the number of bits needed to encode that occurrence will be smaller, because meanwhile the encoder has constructed an internal representation of the pattern. Using this approach, the first seen occurrence of the pattern could be masked. This dependency is easy to remove, if desired: it is enough to scan the image in several directions and retain the minimum complexity found in all scans. In the examples presented in this paper, we have scanned the image in four different directions, resulting from rotating the image in steps of ninety degrees.

Finite-context modelling requires memory resources that grow exponentially with the size of the alphabet. Therefore, usually, even gray-level images use alphabets that render these models almost useless. To cope with this problem, before computing the complexity surfaces, we perform a reduction in the number of intensities to a maximum of twenty, using Lloyd-Max quantization [17, 16].

2.1 Experimental results

Using the method described in Section 2, we have studied its ability for finding both exact and approximate repetitions, assuming the whole past occurrences. In the first image of the first row of Figure 2, we can see that the method performs very well on exact repetitions, perfectly detecting the two equal squares. Looking at approximate repetitions, the models seems to provide good results until a level image disparity of twenty percent. For a higher level of disparity, the method performance reduces substantially.

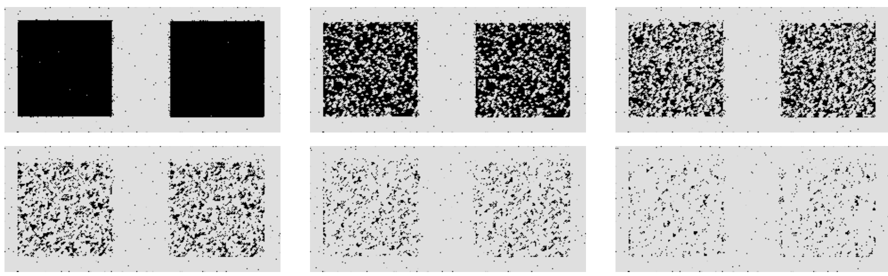

Fig. 2: Complexity surfaces using the whole image distribution, where the darker regions correspond to low complexity occurrences motivated by the repeating patterns. In the first row, from left to right, the pair of images differ in zero, ten and twenty percentage of the pixels. In the second row, from left to right, the pair of images differ in thirty, forty and fifty percentage of the pixels.

Fig. 3: On the left, the image was changed by inserting the same small textured region into a different place. On the middle, the logarithm of the complexity surface, using the whole image distribution. On the right, the logarithm of the complexity surface of the local distribution (200 pixels). The darker regions correspond to low complexity occurrences, motivated by the repeating patterns.

In Figure 3, we show the performance of the method in more difficult shapes (non-rectangular), in this case a person. As it can be seen, in the second image the method performs well when the whole past occurrences are used. However, it fails if only the local past occurrences are used (third image), protruding the sky and the shadow, since they share nearly the same intensity.

There are some situations where groups having nearly the same intensity compose an image, such as the rocks and shells in Figure 4. In the second image, using the model with the whole past occurrences, we were able to identify the shell an their three copies, although with a light increase of the evidence that make up the rest of the image. Using the local past occurrences (third image), we were also able to attain the same result, but in this case with reverse grays.

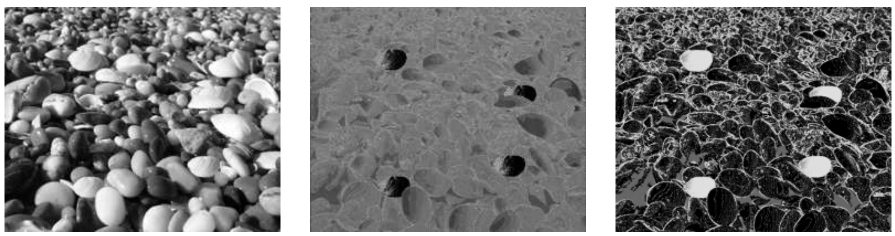

Fig. 4: On the left, the image was changed by inserting the same small textured region into three different places. On the middle, the logarithm of the complexity surface using the the whole image distribution. On the right, the logarithm of the complexity surface of the local distribution (200 pixels). The darker regions correspond to low complexity occurrences, motivated by the repeating patterns.

Now we consider a different case, where we have two equal squares and the rest of the image is a background with a single colour, as can be seen in Figure 5. In the second image, the method based on the whole past occurrences identify the two squares and the background together, i.e., it was unable to separate the objects. On the other hand, the method based on the local past occurrences, identified only the background, showing the outline of the objects.

Fig. 5: On the left, the image contains two equal squares and a background with a single color. The second image shows the logarithm of the complexity surface relatively to the first image, using the whole image distribution. The third image shows the logarithm of the complexity surface relatively to the first image, using the local distribution (50 pixels). The darker regions correspond to low complexity occurrences, motivated by the repeating patterns.

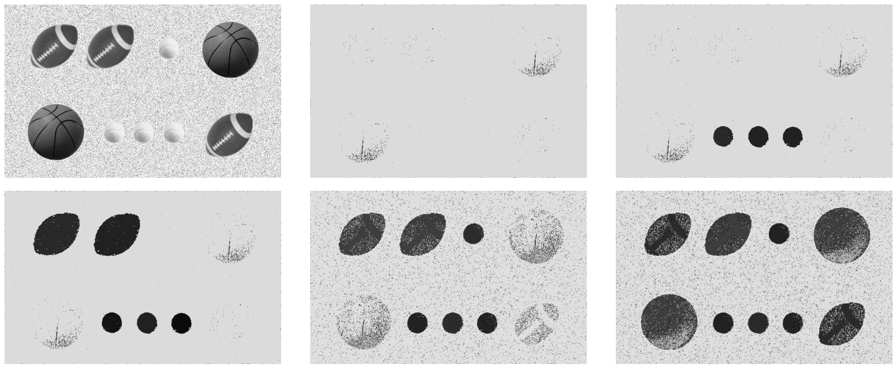

Fig. 6: Complexity surfaces using several depth values of the local past. In the first row, the first image is the original, characterized by containing several copies of tennis, rugby and basketball balls. The second and third images correspond to the complexity surfaces using the last 50 and 100 pixels of past information, respectively. In the second row, from left to right, the images corresponds to the complexity surface using, respectively, 200 pixels, 100,000 pixels and all pixels available (whole past occurrences).

In order to understand the influence of the *magic value* corresponding to the local past occurrences, we have generated the complexity surfaces using several values, as it can be seen in Figure 6. Accordingly, using the last 50 pixels of history, the approach is unable to find any pattern. For a past depth of 100 pixels, the approach detects a repeated pattern containing three tennis balls. As the history depth increases, additional repeating patterns are detected, showing that this value plays an important role in this method, allowing being adjusted according to the need.

Therefore, looking at the previous examples, a repeated pattern can also be considered a local region. To overcome this issue, we explore the relation of using the method based on the whole past occurrences and the method based on the local past occurrences. In Figure 7, the complexity surface (second image) successfully identified a repeated pattern in the three rocks and in the reflection of the light in the water. This last one is a repeating pattern that we do not want to evidence (false-positive). The reflection of the light in the water is a local pattern, as it can be seen in the third image. Thus, we subtract the values from the model based on the local past occurrences to the model based on the whole past occurrences (fourth image), showing only the desired repetitive patterns (true-positives). To improve the visualization of the repeated patterns, the image can been modified using a false-colour in the low complexity zones.

3 Conclusions

Detecting unknown repeated patterns that appear multiple times in a digital image is an interesting problem but also a challenging one. The absence of *a priori* knowledge about the pattern makes the task difficult.

Fig. 7: The first image was changed by inserting the same small textured region into two different places. The second image shows the logarithm of the complexity surface, using the the whole image distribution. The third image shows the logarithm of the complexity surface of the local distribution (200 pixels). The last image shows the output of the subtraction of the third image to the second image. The darker regions correspond to low complexity occurrences, motivated by the repeating patterns.

By exploring the connection between image compression and image complexity, we took an information-theoretic approach based on finite-context modelling, to pinpoint the possible locations of the repeated patterns. We build a complexity surface of the image containing the repeated patterns, and associate them to areas of low complexity. This approach has two ways of collecting statistical information. One is based on the whole past occurrences and the other is based on the local past occurrences. The complexity surfaces resulting from the approach based on the local past occurrences can be subtracted to the complexity surfaces resulting from the one based on the whole past occurrences. We found that this procedure helps in eliminating false positives.

The proposed method shows good performance in a number of different situations, such as in non-rectangular shapes, but it has shortcomings. It is unable to deal well with repetitions that share a mutation above twenty percent and cannot deal with operations such as rotations. However, in this work, we have taken an additional step in the direction of being able to identify unknown repeated patterns in images.

References

[1] T. C. Bell, J. G. Cleary, and I. H. Witten. *Text compression.* Prentice Hall, 1990.

[2] G. J. Chaitin. On the length of programs for computing finite binary sequences. *Journal of the ACM*, 13:547–569, 1966.

[3] T. I. Dix, D. R. Powell, L. Allison, J. Bernal, S. Jaeger, and L. Stern. Comparative analysis of long DNA sequences by per element information content using different contexts. *BMC Bioinformatics*, 8(Suppl. 2):S10, 2007.

[4] I. Gondra and D. R. Heisterkamp. Content-based image retrieval with the normalized information distance. *Computer Vision and Image Understanding*, 111:219–228, 2008.

[5] G. Gordon. Multi-dimensional linguistic complexity. *Journal of Biomolecular Structure & Dynamics*, 20(6):747–750, 2003.

[6] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[7] A. Lempel and J. Ziv. On the complexity of finite sequences. *IEEE Trans. on Information Theory*, 22(1):75–81, Jan. 1976.

[8] M. Li, X. Chen, X. Li, B. Ma, and P. M. B. Vitányi. The similarity metric. *IEEE Trans. on Information Theory*, 50(12):3250–3264, Dec. 2004.

[9] A. Mallet, L. Gueguen, and M. Datcu. Complexity based image artifact detection. In *Proc. of the Data Compression Conf., DCC-2008*, page 534, Snowbird, Utah, 2008.

[10] J. Rissanen. Modeling by shortest data description. *Automatica*, 14:465–471, 1978.

[11] R. J. Solomonoff. A formal theory of inductive inference. Part I. *Information and Control*, 7(1):1–22, Mar. 1964.

[12] R. J. Solomonoff. A formal theory of inductive inference. Part II. *Information and Control*, 7(2):224–254, June 1964.

[13] N. Tran. The normalized compression distance and image distinguishability. In *Human Vision and Electronic Imaging XII — Proc. of the SPIE*, page 64921D, Jan. 2007.

[14] C. S. Wallace and D. M. Boulton. An information measure for classification. *The Computer Journal*, 11(2):185–194, Aug. 1968.

[15] A. J. Pinho and P. J. S. G. Ferreira Finding unknown repeated patterns in images In *Proc. of the European Signal Processing Conference, EUSIPCO 2011*, pages 584–588 , Barcelona, Spain Aug. 2011.

[16] J. Max. Quantizing for minimum distortion. *IRE Trans. on Information Theory*, 6:7–12, Jan. 1960.

[17] S. P. Lloyd. Least squares quantization in PCM. *IEEE Trans. on Information Theory*, 28(2):129–137, Mar. 1982.