IADIS International Journal on Computer Science and Information Systems
Vol. 7, No.1, pp. 47-63
ISSN: 1646-3692

MULTI-MODAL MEDICAL IMAGE REGISTRATION USING NORMALIZED COMPRESSION DISTANCE

Pere-Pau Vázquez. *MOVING Group, Universitat Politècnica de Catalunya.*

Jordi Marco. *Dep. LSI, Universitat Politècnica de Catalunya.*

ABSTRACT

Image registration is an important task in medicine, especially when images have been acquired by different scanner/sensor types, since they provide information on different body structures (bones, muscles, vessels…). Several techniques have been proposed in the past, and among those, Normalized Mutual Information has been proven as successful in many cases. Recently, Normalized Compression Distance, that makes use of real-world compressors, has been proposed as a simple yet effective technique for image registration. It is especially suitable for the case of CT-MRI registration, and may improve timings over Normalized Mutual Information. However, other image modalities such as PET pose some problems and do not achieve accurate registration. In this paper we analyze and propose a valid approach for image registration using compression that works properly for different combinations of CT, MRI and PET images.

KEYWORDS

Medical Image Registration, Image Compression, Visualization.

1. INTRODUCTION

Image registration is the process of overlaying two or more images that represent the same information taken at different times, from different viewpoints, and/or by different sensors. Image registration is nowadays a central tool for image analysis, understanding, and visualization in medical and scientific applications. Aligning medical images has interest for the analysis of temporal patient evolution, the fusion of multimodal images, inter-patients comparison, and so on. In our case, we will concentrate on multimodal images, that is, images acquired with different sensors/processes, such as Computerized Tomography (CT), Positron Emission Tomography (PET), or Magnetic Resonance Imaging (MRI). Moreover, we will also consider only rigid transformations of the images. A huge number of registration algorithms have already been proposed, some of them based on intensity values or feature detection (see for instance Zitová & Flusser (2003) for a survey). One of the most prominent approaches is the maximization of mutual information (Viola & Wells (1997)). Our objective is to develop a method for robust multi modal image registration that achieves Mutual Information level results, requires little user intervention, and, if possible, is faster than previous methods. Our work relies on the use of Normalized Compression Distance to solve the image registration problem. Normalized Compression Distance is a measure that has its roots in Kolmogorov complexity and that uses compressors to determine the similarity between two sequences. It uses as input the two sequences to compare and a third sequence or file that combines both, usually the concatenation of the sequences to analyze. This measure has been previously used to address CTMRI registration Bardera et al. (2006, 2010) based on the use of *bzip2*. Unfortunately, other image modalities, such as PET yield poor results: registration point does not necessarily correspond to the correct one and, more importantly, the distance function has a high number of local minima, which will easily trap any optimization process. Our method is inspired by this previous approach though with fundamental changes. As a result, we deal with CT, MRI, and PET images, and we avoid the presence of local minima. Our system requires almost no user intervention and has the following novelties:

- Identification of the most suitable real-world compressor for image registration: We found that Prediction-by-Partial Matching family of compressors achieve high quality registration results. They are better than the ones obtained with block-based or dictionary-based compressors.
- We have proposed a new approach for image scan that provides better results than row-by-row, column-by-column, or Space Filling Curve scans.
- Data quantization: We improve the registration quality and speed up the process by quantizing

the input images.

The rest of the paper is organized as follows. In Section 2 we present some background concepts and analyze previous work. Section 3 introduces the different aspects we analyzed when addressing image registration with compression. In Section 4 we analyze the performance of real world compressors for image registration and we present our algorithm. Section 5 concludes our work with a discussion and the analysis of some lines for future study.

2. BACKGROUND

2.1 Normalized Compression Distance

Normalized Compression Distance is a universal metric of distance between sequences. It has its roots in Kolmogorov complexity. We briefy review here some background (see Li and Vitányi's book Li & Vitanyi (1993) for more details). The Kolmogorov complexity (*K(x)*) of a string x is the length of the shortest binary program to compute *x* on a universal computer (such as a universal Turing Machine). Thus, *K(x)* denotes the number of bits of information from which *x* can be computationally retrieved. As a consequence, strings presenting recurring patterns have low complexity, while random strings have a complexity that almost equals their own length. Hence, *K(x)* is the lower-bound of what a real-world compressor can possibly achieve. The conditional Kolmogorov complexity $K(x|y)$ of x relative to y is the length of a shortest program to compute x if y is provided as an auxiliary input.

Both Kolmogorov complexity and conditional Kolmogorov complexity are machine independent up to an additive constant. Bennet *et al.* Bennett et al. (1998) define the information distance between two, not necessarily equal length binary strings as the length of the shortest program that, with input x computes y, and with input y computes x. It was also shown that, up to an additive logarithmic term, it can be defined as:

$$E(x; y) = max\{K(y|x);K(x|y)\} \quad (1)$$

The information distance is a metric, up to negligible violations of the metric inequalities. Li *et al.* Li et al. (2004) present a normalized version of information distance, the similarity metric, defined as:

$$d(x; y) = max\{K(y|x);K(x|y)\} / max\{K(y);K(y)\} \quad (2)$$

The authors also prove that it is also a metric, and that this metric is universal: two files of whatever type similar with respect to a certain metric are also similar with respect to the similarity metric. Being Kolmogorov complexity not computable, it may be approximated with the use of a real-world compressor, leading to the Normalized Compression Distance (NCD):

$$NCD(x; y) = C(xy) – min\{C(x);C(y)\} / max\{C(x);C(y)\} \quad (3)$$

where function $C(F)$ is the size of the compression of a certain file F, and xy is the concatenation of files x and y. Although the similarity metric has values in [0..1], NCD values are usually in the range of [0..1.1], due to compressor imperfections. NCD has been used for applications such as language classification and handwriting recognition Cilibrasi & Vitanyi (2005). Cilibrasi and Vitányi also analyze the conditions that compressors must fulfill in order to be used for computing the Normalized Compression Distance:

- **Idempotency:** For a repetition of a string, the compressor should be able to detect the repetitions and thus compress the file to a similar size than the original string compression $C(xx) = C(x)$, and $C(\lambda) = 0$ where λ is the empty string.
- **Monotonicity:** The concatenation of two strings should yield to a less compressible file tan taking a single string alone, up to a certain precision: $C(xy) \geq C(x)$.
- **Symmetry:** $C(xy) = C(yx)$. Compression should be symmetric, that is, changing the order of the concatenated strings should no affect the length of the compression.
- **Distributivity:** $C(xy) + C(z) \leq C(xz) + C(yz)$. Real-world compressors seem to satisfy this property.

The data compressors with these properties are named *normal compressors*. Most real-world compressors do fulfill those properties, at least to a point where they are usable for NCD computation. Some of the candidates are: stream-based (*gzip*), block-based (*bzip2*), and statistical (*paq8px*) compressors. Note that we are not searching for the best compression, as reducing the file size does not imply a reduction in NCD computation (Cilibrasi & Vitanyi (2005)). Even with tested compressors, their behavior may change with the size of the file. As a consequence, stream-based compressors (such as the based on Lempel-Ziv) will probably improve their behavior with respect to symmetry, because they adapt to the regularities of the string throughout the compression process. However, large files may affect the efficiency in NCD computation for block-based compressors. As studied by Cebrián *et al.* Cebrián et al. (2007), in the case of *bzip2*, the best option works properly for files up to 900KB before being compressed. Larger sizes make the comparison processes less effective. Other predictive compressors, such as the ones belonging to the PPM (Prediction by Partial Matching) family are not precisely symmetric, but for long strings, they are close to symmetry. If symmetry is not fulfilled to a great extent, the compressor should not be used at all.

Normalized compression distance has been used for music clustering Cilibrasi et al. (2004) and music style modeling Dubnov et al. (2003), automatic construction of the phylogeny tree based on whole mitochondrial genomes Li et al. (2001), the automatic construction of a language trees Li et al. (2004), Benedetto et al. (2002). However, its use is not always suitable, or, at least, using NCD on the raw data may not yield good results, as demonstrated by Tran Tran (2007) or Rocha *et al.* Rocha et al. (2006).

2.2 Image Analysis using Normalized Compression Distance

Normalized Compression Distance has been used for image classification Cilibrasi & Vitanyi (2005) with gray-scale images. Lan and Harvey Lan & Harvey (2005) show that the measure performs better than histogram-based approaches in object recognition, using PPM-based compression. A further work by Li and Zu (Li & Zhu (2006)) improves the optimization task by using a Lempel-Ziv encoding of the data and using either the dictionary, or the compressed patterns for measuring image similarity.

Bardera *et al.* Bardera et al. (2006) use Normalized Compression Distance for image registration. In order to do this, they select a window of pixels in one image and another one in the other reference image. Then, pixels are interleaved forming a new image where the red channel holds the pixels of reference image 1 and green channel the ones of reference image 2. These images are then compressed using JPEG 2000 and the compressed size is used as C(xy) in equation 3. They also present a second approach where the gray-scale values are treated as elements of a string, and *bzip2* is used to compress the resulting string. Again, the values of both images are interleaved. Although this approach works for the CT-MRI registration, it has problems for the PET-MRI image pair.

Vázquez and Marco, Vázquez and Marco (2012) analyzed image similarity using NCD using different compressors and image formats. Their study showed that, among the many compressors that can be chosen, for image similarity, some algorithms were superior: Prediction-by-Partial Matching (implemented in compressors such as *paq8px* or *7-zip*), and block-based encoding (implemented in programs such as *bzip2*). However, this study has not been performed with a different comparison approach: image correlation, that is what image alignment seek for. The study also found interesting results that will be used here: the selection of the file format is also vital. Their outcomes show that, for image comparison using compressors, *simple* formats that do not have extensive headers are superior to compressed formats or other formats that store a large amount of information in the headers. We will therefore restrict our file format to a simple one, the greyscale mode (magic number ‘P5’) of PPM (Portable Pixel Map).

Image scan directions and concatenation building also have an important task in image registration, and have also not been analyzed previously in this context further than pointing out different possibilities (Bardera et al. (2006), Macedonas et al. (2008)). This paper intends to provide some answers to these questions for the concrete problem of medical image registration. The following section will introduce the compressor families and the different kinds of tests we carried out before arriving to our proposed solution.

3. COMPRESSION-BASED REGISTRATION

3.1 Data Compression Algorithms

In general, a data compression algorithm focuses in identifying and extracting data redundancy. There are several ways to do this. In short, some of the main ideas under the standard compressors:

- **Huffman coding:** The core idea is to assign a fixed-length code to each symbol. The most frequent the symbol, the lower number of bits it is assigned.
- **Arithmetic coding:** The general idea is to replace a stream of input symbols with a single floating-point number. By codifying a stream into a single number, instead of replacing each of its symbols by different codes, a lower amount of bits are wasted.
- **Dictionary-based compression:** These methods encode variable-length strings of symbols as single tokens. These tokens are indices to a phrase dictionary. Note that this makes the compressor to look for, not only redundant symbols, but redundant repetitions of different symbols, if existing.
- **Block-sorting compression:** It is based on a transformation that permutes the order of the characters. As a consequence, after the transformation, repeated characters will be grouped together. Thus, other techniques that work on repeated characters such as run-length encoding can be applied to the transformed string to reduce data size.
- **Prediction by Partial Matching:** It is an adaptive statistical data compression technique that uses a set of previous symbols in the uncompressed symbol stream to predict the next symbol in the stream Cleary et al. (1995).

The main issue with multimodal image registration is that we are not looking for similarities but for correlation. This is due to the fact that images from different acquisition methods do not show the same information, on the contrary, it is somewhat complementary (bones versus soft tissues...). Therefore, the same gray level may indicate different information and be placed at different positions. That is why Mutual Information achieves so good results. Thus, the selection of the most appropriate compressor may not be straightforward. The first approach for compression-based registration with multi-modal images used either *JPEG2000* (a compression algorithm specially tailored to deal with images, and that uses wavelets and entropy coding), or block-based compression. In some cases, good results were achieved, but for some image source combinations, such as MRI-PET, results are inferior. Our objective is to identify the best compression scheme (if any) for image registration, and to determine which image sources combinations are feasible.

In order to illustrate this, we may perform a naïve experiment: We create two files, one containing 16 *a*'s followed by 16 *b*'s, and another one containing the same symbols, but interleaved. If we compress using the command `gzip -9` (maximum compression), we will find that the size of the second file is smaller than the previous one. A different compression strategy, such as run-length encoding, would yield a totally different result. This behavior will be useful for registration, and it inspired us to explore with different compressors such as Prediction-by-Partial Matching schemes.

3.2 Image Scan and Concatenation

The Normalized Compression Distance works on strings. However, images contain 2D information that is not trivially changed to a 1D string Macedonas et al. (2008). Therefore, some possibilities arise, such as building a 1D sequence from an image by taking pixels row by row, or column by column.

Macedonas *et al.* Macedonas et al. (2008) have tested both cases and stated that the row-by-row approach yields slightly better results. This is the approach also followed by Li and Zhu Li & Zhu (2006). Bardera *et al.* Bardera et al. (2006) use zig-zag scan, and point as a future work the analysis of other 2D to 1D sequence conversion such as the use of space filling curves. One of the objectives of this paper is to determine if the 2D to 1D transformation has any impact when using compressors for image registration. Thus, we will test different 2D to 1D sequence transformation methods. This will be our first task. We evaluated the following 2D to 1D sequence transformation methods:

- **Row-by-row:** This is the most typical approach, as tested by Macedonas *et al.* Macedonas et al.

(2008), it seems to obtain superior results to column-by-column.

- **Column-by-column:** Previously tested in Macedonas et al. (2008), the second typical option of image transform to 1D sequence. The authors claim it yields slightly worse results tan row-by-row.
- **Space Filling Curves:** Transformation used in environments where spatial coherence must be taken into account. We tested this approach using the Hilbert curve Lawder & King (2000).
- **Random sampling:** The idea is to create a string by (pseudo-)randomly selecting the pixels from the input image.

For images containing similar information, such as when comparing for image equality, the first three approaches may take advantage of the spatial coherence. However, for image registration, the two images may contain greyscale levels that do not necessarily need to match, but the correlation between both images must be maximized. We will see that random sampling of the images may improve the image registration, especially for compressor algorithms that make use of dictionaries in some way, such as Prediction by Partial Matching algorithms and Lempel-Ziv-based systems. Our intuition is that this strategy leads to a wider dictionary with smaller words. Prediction by Partial Matching and Lempel-Ziv methods seem to take advantage of this fact in order to detect the pixel-by-pixel correlation. Moreover, as already stated in the introduction, there are different ways to combine the pixels of the input images in order to build the concatenated one. If we follow the spirit of Cilibrasi and Vitányi Cilibrasi & Vitanyi (2005), we should concatenate the two images. However, Bardera *et al.* Bardera et al. (2006) showed that interleaving pixel by pixel yielded good results. To the authors' knowledge, there has been no previous study on the effect of different ways to construct the concatenated string. We have therefore evaluated several methods for image combination:

- **Image concatenation:** This would be the classical approach.
- **Pixel interleaving:** The concatenated image is built by combining one pixel from each image, each taken from the same position of both images.
- **N-Pixel interleaving:** The concatenated image is created by building groups of n pixels from each image. We used values of $n = 2$ and $n = 4$.

Note that the pixel interleaving, with 1, 2, or 4 pixels, does not change the complexity of algorithm for string construction but up to a constant (the length of the trivial program that may interleave the pixels), and therefore, we may keep on using Normalized Compression Distance as a distance metric. The same happens with the (pseudo-)random sampling commented above, as long as we have the pseudo-random sampling sequence computed in advance, as usually happens in most cases (note that the random function of most programming languages actually returns numbers from a set of predefined pseudo-random sequences). Again, pixel interleaving obtains the best results because Prediction by Partial Matching-based compressors and the ones based on Lempel-Ziv are able to gather the correlation from the images when the concatenated version is encoded this way.

4. IMAGE REGISTRATION USING REAL WORLD COMPRESSORS

4.1 Experimental Setup

In this section we will analyze the behavior of different real world compressors when using them for multimodal Image Registration. Concretely, we will analyze the suitability of Normalized Compression Distance for CT-MRI-PET registration. The experiments are performed by taking a pair of images that, otherwise indicated, are almost aligned. From those images, one is taken as the source, and the other one as the destination. In order to make the plots more intuitive, we only show transformations involving translations in both X and Y directions. For the tests, we move the source image over the destination one by applying translations from (-10;-10) to (10; 10). Of course, the destination image has an extra frame of 10 pixels width (in background color) around the original image, and we clip the destination image according to the position of the moving one. The results are shown as a 3D chart where we plot the Normalized Compression Distance using different compressors and parameterizations. Since the important point is the minimum value, and the function shape, due to the lack of space, we clipped the plots in order to show only the informative parts of the distance function. The minimum distance is the registration point. In order to compare the results, we will analyze the same image pair using Normalized Mutual Information. The desired result is a function that decreases as we approach the matching point and this one coincides with the one found with Normalized Mutual Information. Note that, for several reasons, the images coming from two different capture devices do not exactly simple the same regions in space. As a consequence, some small translation from the matching point given by NMI may be visually acceptable. A very important key issue is that the function should not be plagued with local minima, as this may challenge the optimization process.

Although we analyzed a high number of compressors, we will mainly show the results from the ones that show better behavior: a block-based compressor: *bzip2*, a compressor based on the Prediction by Partial Matching scheme: *paq8px*, one of the most effective data compressors, in terms of compression rate, according to Bergmans (n.d.), but very time consuming. In Section 5 we will summarize the results obtained with other compressors. Since some of the approaches are orthogonal, such as image scan versus image combination, throughout the paper we will analyze different configurations and incrementally incorporate them in the following sections, in order to reduce space and provide more informative results.

4.2 Image Concatenation versus Pixel-by-Pixel Interleaving

As already stated, Bardera *et al.*, Bardera et al. (2006, 2010) showed that pixel-by-pixel interleaving was a good means to achieve CT-MRI registration with a block-based compressor. Despite that, the image combination possibilities were not deeply discussed. As a consequence, we first analyze the performance of different schemes using the same compressor. In contrast to the original approach, which analyzed the image using zig-zag, we perform the image scan row-by-row, as we obtain equivalent results. An analysis on different image scans is presented next. The first experiment is CT-MRI registration. The images used are shown in Figure 1-left.

In Figure 2 we compare the different behavior in registration when using image appending and image interleaving, respectively. Note that, independently on the quality of the matching point, pure image concatenation leads to a high number of local minima. This poses problems for the registration process, as an optimization function may be easily trapped in any of them.

We also tried different pixel range sizes, that is, instead of concatenating one pixel from each image, we used groups of 2 and 4 pixels from each image. These N-pixel interleaving strategies do not provide good results. Actually, the bigger the number of pixels we group, the worse the result. We will show this in the following section. From now on, the experiments shown will incorporate pixel-by-pixel interleaving.

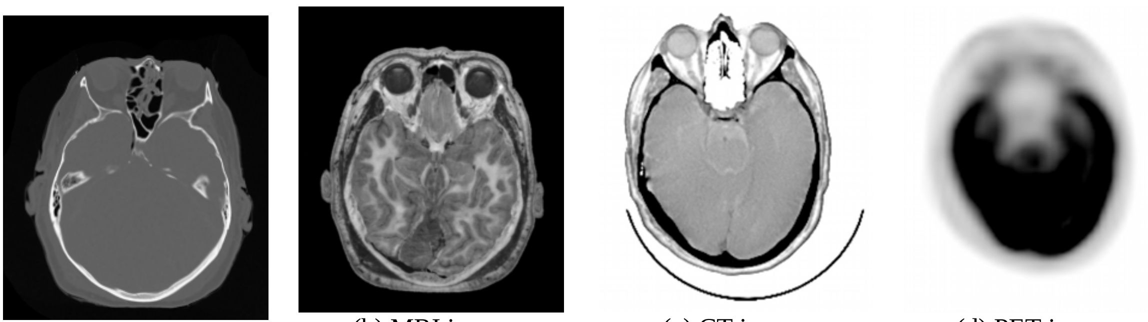

(a) CT image (b) MRI image (c) CT image (d) PET image

Figure 1. Left: CT and MRI pair to register. Right: CT-PET image pair.

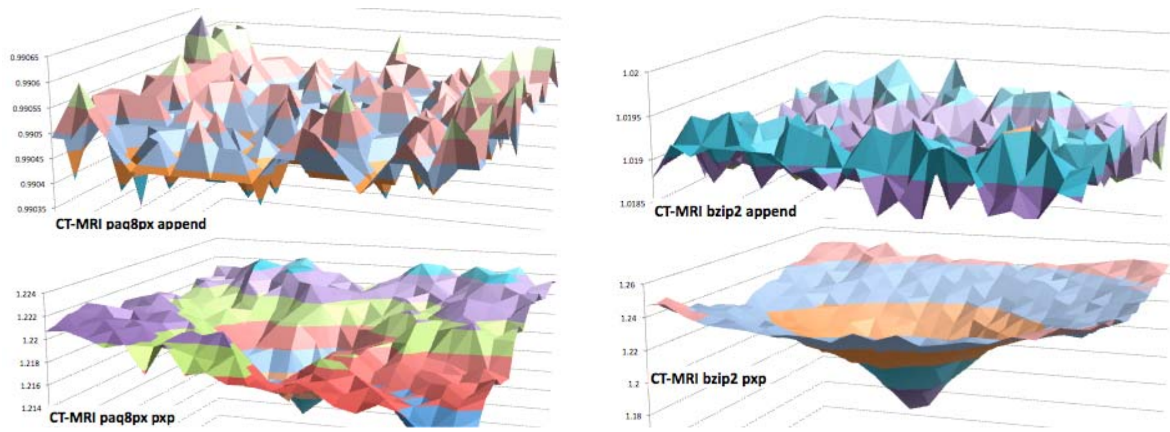

Figure 2. CT-MRI registration using different image combination methods. We compare regular file append (first and third for *paq8px* and *bzip2*, respectively) and pixel-by-pixel interleaving (second and fourth for *bzip2* and *paq8px*, respectively).

4.3 2D to 1D Image Transformation

An image is a 2D data structure. For file writing, this information is transformed to a 1D array. This transformation can be done in different ways. So far, we have scanned the input images row-by-row. However, some other alternatives have been pointed out in literature. Since no experimental evidence on which method would yield better results for image registration, we have tested several approaches: row-by-row, column-by-column, space filling curves, and (pseudo-)random sampling.

Row-by-row and column-by-column have already been used for image comparison, with little advantage for row-by-row according to Macedonas *et al.* Macedonas et al. (2008). Space Filling Curves were suggested as a possibility by Bardera *et al.* Bardera et al. (2006). The rationale behind the use of Space Filling Curves is the fact that such curves take advantage of spatial coherence. From our experiments, it turned out that none of the previous approaches was optimal for image registration, as in most cases the function still contains a high number of local minima (Figure 3). Then, we came up with a totally different solution: random sampling the input images. We build the combined image by pixels selected from pseudo-random positions in the input images (the randomly selected position is the same for both input images). The reasoning behind is, when we are addressing image registration, we are not looking for local coherence, but pixel correlation between the two input images. Space Filling Curves may worsen the results because the information contents may vary greatly from one image to the other, and therefore, we are, somehow, counteracting the compressor task. For dictionary-based compressors, the dictionary construction might benefit from having a larger set of smaller words, than a smaller set of larger words. This is what may be induced by random image sampling. Actually, this technique improves image registration with dictionary-based compressors such as *paq8px* and *gzip*, and even block-based compressors such as *bzip2*.

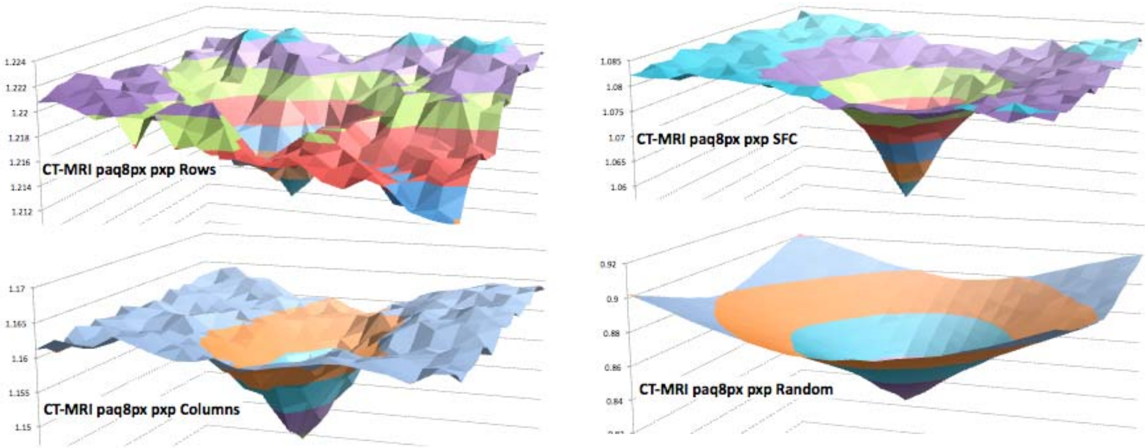

Figure 3. CT-MRI registration using different image scan strategies and *paq8px*. In this case, column by column yields better results than row-by-row. However, the best matching is always found using random scan (bottom right). In this case the matching point is good, and the function shape is soft, thus minimizing the probabilities of any optimization algorithm to get trapped in a local mínimum.

We show the results of these different configurations in Figure 3. We used again the CT-MRI pair in Figure 1-left. Note that random scan (bottom) produces higher quality results than with other techniques, both in terms of function shape, and a more accurate matching of the registration point. Row-by-row and column-by-column often obtain different, contrary results, but none is always better than the other. Sometimes row-by-row scan generates a softer shape, and sometimes column-by-column generates a better shape. The results are unpredictable, and we have not been able to find a consistent behavior throughout the different tests. Moreover, except for the random sampling, none of the previous image scan methods improves the registration point.

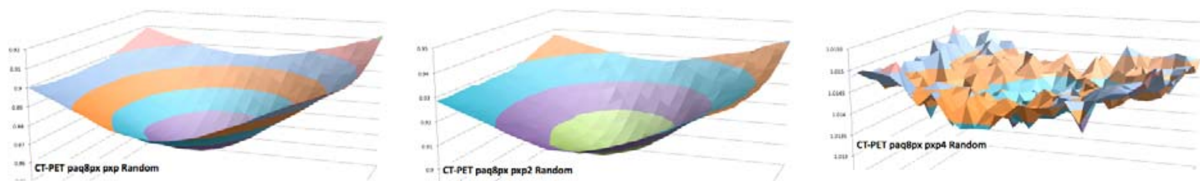

Figure 4. CT-PET image registration using *paq8px* and different ways to combine image pixel interleaving: single pixel interleaving (left), 2-pixel interleaving (center), and 4-pixel interleaving (right). Note how the best results are obtained with single pixel interleaving.

As already commented, we previously experimented with different interleaving strategies with no positive results. In order to assess our intuition supporting the random sampling strategy, we experimented again with N-pixel interleaving. If our intuition is truth, incrementing the number of pixels taken into account in the interleaving process should worsen the registration function. We used the CT-PET image pair in Figure 1-right. This is shown in Figure 4 where pixel interleaving sizes are compared: one (pxp), two (pxp2) and four (pxp4) pixels. The method that obtains better registration is always 1 pixel wide image combination. This enforces our idea that correlation is better captured if we randomly sample the input images.

4.4 Improving Registration using Image Quantization

When addressing image registration using Normalized Mutual Information, the number of bins selected for the histograms does influence the registration results. More concretely, the registration may improve if we select a lower number of bins, say 100 for instance. Furthermore, a lower number of bins also accelerates the registration algorithm because the joint histogram is sensitively simpler. This bin reduction also makes the algorithm more robust to the presence of noise.

In our experiments we follow the same idea. We quantize the input images to 16 bins (as shown in Figure 5). This reduces the noise and the amount of information, but the important details are not removed. As a consequence, registration results are improved.

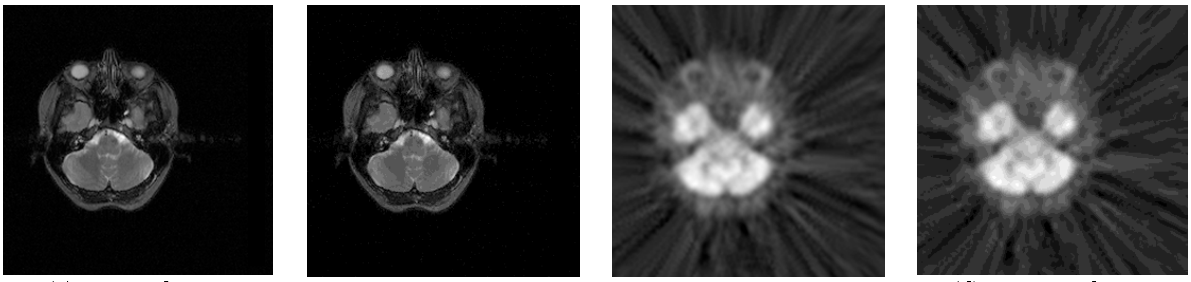

(a) Original MRI (b) Original PET (c) Quantized MRI (d) Quantized PET

Figure 5. Quantization of the input images to 16 values. Left column shows the original images to register, and the right column shows the same images after the quantization.

Once the images are quantized, we compare both using the Normalized Compression Distance. We set one of the two images as the source, and the other one as the destination. Again, we move the source image over the destination one by applying translations from (-10;-10) to (10; 10). The minimum distance is the registration point. As shown in Figure 6, we are able to correctly register the MRI-PET pair, a difficult example, since the information in the PET image is not very detailed.

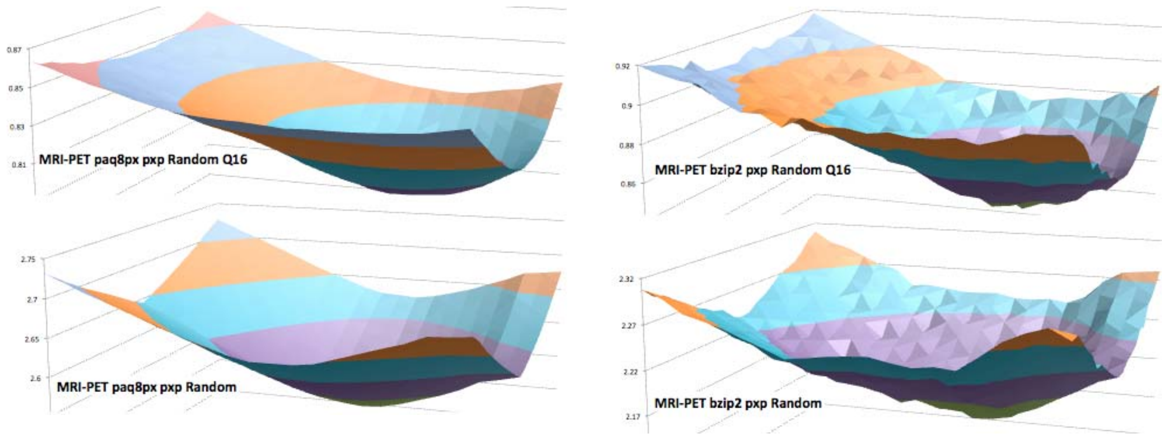

Figure 6. Comparison of MRI-PET registration using *paq8px* and *bzip2* using quantization and with unmodified image sources. In this case, quantization does not especially improve the function shape, but it improves the matching point. Actually, we match the result obtained by Normalized Mutual Information when using NCD and *paq8px* compressor.

4.5 Results

We have tested our registration scheme with several multimodal image pairs, and some of them already appeared in this paper. So far, we have only analyzed the adequacy of the shape of the distance function. In this section we will further analyze the obtained registration points with a pair of CT-MRI, another one consisting in CT-PET, and finally, we also test an MRI-PET pair.

From now on, the experiments are carried out using all improvements: image quantization, image interleaving, and random scan. For the sake of clarity, we plot the NCD values obtained by *paq8px* compressor, since it is the one that showed most robustness in the experiments. Over the plot, which is shown as a contour chart, we indicate the registration point obtained with this compressor. We also add the registration point obtained by Normalized Mutual Information and the results we obtain with our algorithm and other compressors such as *bzip2* and 7z. We also added 7z since it is a PPM-based compressor that achieves almost as robust results as *paq8px* with a very low time cost.

Our method performs CT-MRI registration very efficiently. It even yields very good results with other compressors such as *gzip*. In Figure 7 (left) we can see the results, that belong to the images shown in Figure 1. In this case, all the registration points computed by our method with different compressors (*paq8px*, *bzip2*, and 7z) correspond to the value computed using Normalized Mutual Information. In this case, *gzip* also correctly finds the registration point.

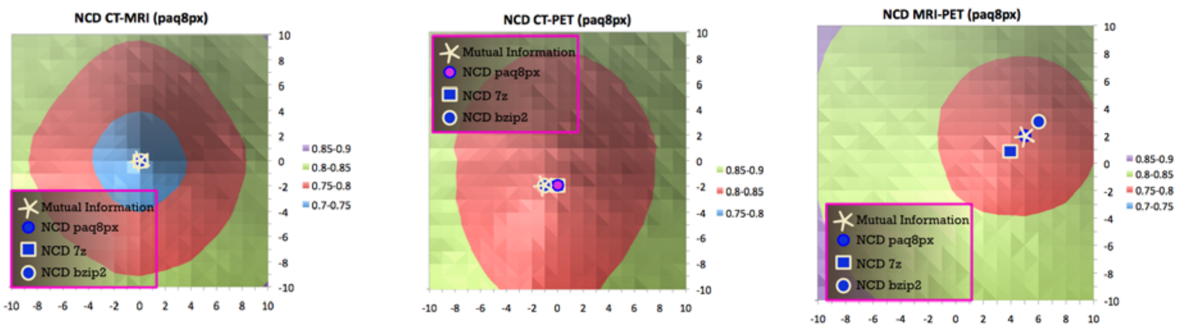

Figure 7. Comparison of the different registration methods with different image modalities. Left: CT-MRI registration. Al registration points coincide with the NMI method. Center: CT-PET registration. NMI and *bzip2* achieve the same result, while using *paq8px* or 7z obtains a point shifted only one pixel away. Right: MRI-PET registration.

As already said, CT-PET registration is usually difficult because of the lack of details of PET images (see Figure 1-right). If we do not use our improvements, *bzip2*, is unable to match the registration point. However, using pixel interleaving and random scan, results are highly improved. In Figure 7 (center) we show the results obtained with our algorithm. In this case, the block-based compressor obtains the same result than NMI, while *paq8px* and 7z are shifted a pixel in the X direction. This is not a bad result, because as noted previously, the images do not exactly sample the same region of the body, and both matchings are visually acceptable. This is shown in Figure 8, where we fused the images at their registered positions using two different channels (red and green). Left image shows the registration position indicated by our algorithm, and the right one shows the registration position as calculated using NMI. Notice that both registration points are perfectly acceptable. Furthermore, it is important to take into account that a change in the number of bins used for the registration process with NMI usually shifts the registration point a couple or three pixels away. The MRI-PET registration is also difficult due to the lack of details in PET images. However, our method correctly finds a registration point for these kind of source data. We show the results in Figure 7 (right). Note that *paq8px* finds the correct point according to NMI, but both *bzip2* and 7z achieve very good results. In all cases, the matching points are visually acceptable.

5. SILHOUETTE-BASED REGISTRATION

The use of silhouettes for image registration is not new, see for example Betting & Feldmar, Betting & Feldmar (1995). However we have addressed this issue in a different way: our objective is to determine if accurate registration can be achieved using NCD over medical images that have been processed for silhouette detection. In our case, we use a Laplacian filter to process the images, as shown in Figure 8.

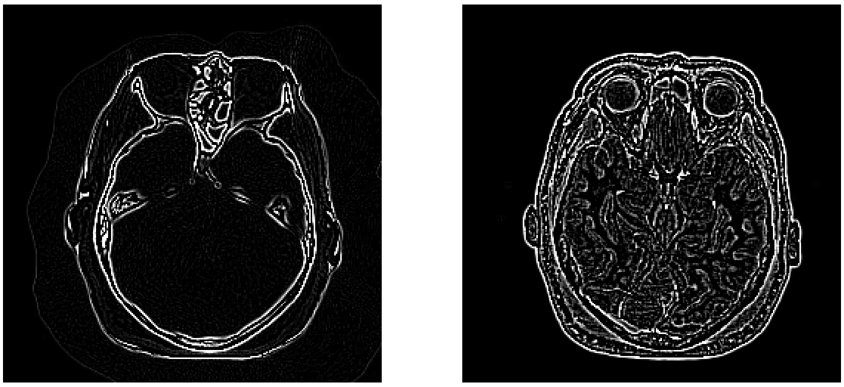

Figure 8. Filtered images. Left image shows the filtered CT image while the right image shows the MRI filtered image.

We tested some models using this approach and the results are promising. For instance, for the CT-MRI pair, we obtain good matching points, as shown in Figure 9. However, the function still presents local minima that might generate errors on optimization. Moreover, there is still a small mismatch when using *bzip2* since the registration point is displaced one pixel away from the theoretical one.

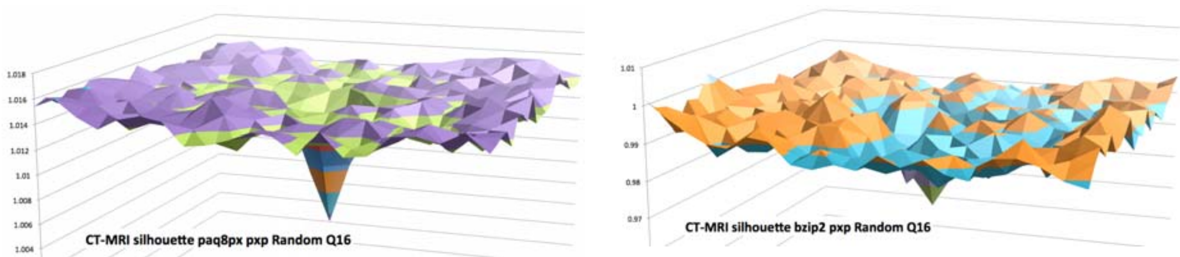

Figure 9. Registration using the silhouette-based approach. This method incorporates all previous optimizations, including random scan, pixel interleaving, and image quantization. Registration point obtained with *paq8px* coincides with the previous methods, but the point found using *bzip2* is one pixel far.

6. CONCLUSIONS AND FUTURE WORK

6.1 Conclusions

In this paper we presented a novel approach to image registration using compression that is able to cope with images that are not rich in details such as PET images. Accurate registration is obtained by modifying the classical Normalized Compression Distance measurement in three ways: a) We quantize the input images, b) we scan the images in a pseudo-random manner, and c) we combine the images pixel by pixel to form the concatenated file. All these three modifications lead to a robust algorithm that is able to align medical images of different sources with little or none user intervention in the definition of parameters.

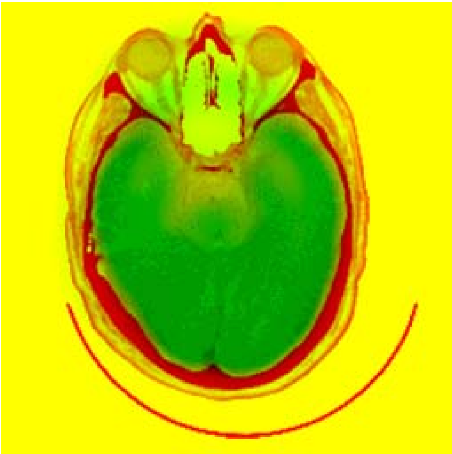

(a) NCD registration (*paq8px*)

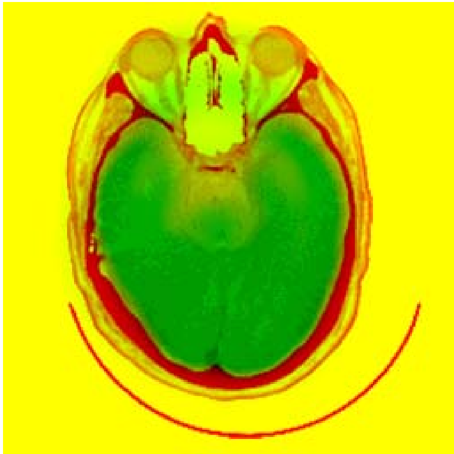

(b) Registration using NMI

Figure 10. Comparison on the registration points found by our method using *paq8px* and Normalized Mutual Information, respectively. Note that there are not big differences in the result and that both are acceptable.

Throughout the process, we also make other interesting findings: a) For image scan, Space Filling Curves and row-by-row or column-by-column did not give good results. b) Regular image concatenation is also unsuitable for image registration. c) PPM-based compressors are more robust for image registration than other schemes. Even though there is still room for improvements, the results we obtain are comparable to the ones obtain by Normalized Mutual Information (see Figure 10). Therefore, we believe that the use of compressors from the Prediction by Partial Matching family is a promising line of investigation. Their performance is superior than with block-based compressors. Moreover, though *paq8px* is very costly, another PPM-based compressor, 7z works very fast, and has proven useful for image registration. Its results outperform NMI times even when computing NMI with a reduced number of bins such as 100.

6.2 Compressor Selection

Throughout our tests we also tested other compressors, such as jpeg2000, *rzip* (tailored to find redundancies placed at high distances), lzma (an algorithm that uses Lempel-Ziv and Markov Chain coding), *hffzip* (based on Huffman coding), and *gzip* (based on Lempel-Ziv). Except for *gzip*, that gives good results for CT-MRI registration, the other ones had not satisfactory results in any case: neither the best value was close to the optimal, nor the shape of the distance function indicated robustness of the measure.

We show a couple of these results using *rzip* and *hffzip* in Figure 11. Note that the results are not optimal. In one case, the number of local minima is very high, and it does not find the correct matching point for a relatively simple model. For the second case, Huffmann coding does not take into account groups of symbols, and therefore, the compression is unable to find redundancies, and it is shown in Figure 11-right as a complete planar function, where no maxima nor minima are found.

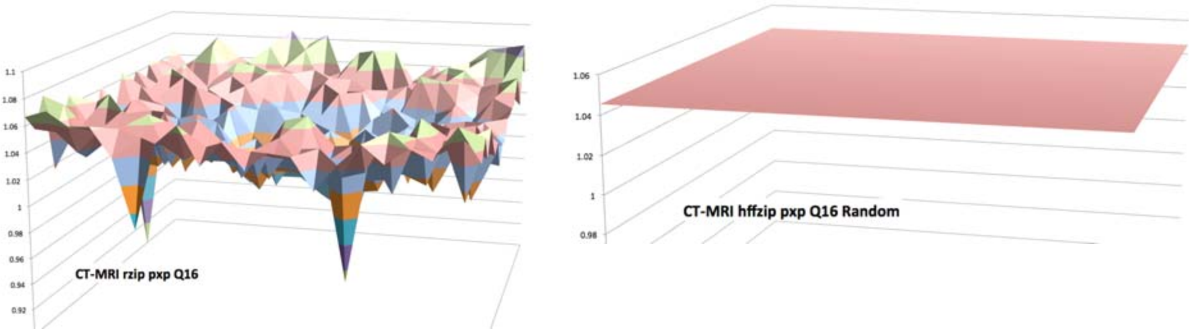

Figure 11. Peformance of the rzip and hffzip compressors. The first one is designed to find redundancies at large distances of the file, while the second implements Huffman coding. The experiment was carried out using the CT-MRI pair. Note that the results are not optimal.

6.3 Future Work

In future we want to continue working on the approach using silhouettes, since it has shown promising results. We also plan to use our method for other image modalities such as SPECT datasets. Another line that might find interesting results is related to the automated detection of features in the images. This has also been used previously for image registration, but not using NCD. We plan to analyze some of these algorithms, such as SIFT, SURF, and so on, in order to determine if they can be used for medical image registration using NCD.

ACKNOWLEDGEMENT

This project has been supported by TIN2010-20590-C02-01 Project of the Spanish Government.

REFERENCES

Bardera, A., Feixas, M., Boada, I. & Sbert, M. (2006), Compression-based image registration, in *IEEE International Conference on Information Theory*, IEEE, pp. 436-440.

Bardera, A., Feixas, M., Boada, I. & Sbert, M. (2010), Image registration by compression, *Information Sciences*, 180(7), 1121-1133.

Benedetto, D., Caglioti, E. & Loreto, V. (2002), Language trees and zipping, *Physical Review Letters*, 88(4).

Bennett, C., Gacs, P., Li, M., Vitanyi, P. & Zurek, W. (1998), Information distance, *IEEETIT: IEEE Transactions on Information Theory*, 44.

Bergmans, W. (n.d.), Maximum compression (lossless compression software), *http://www.maximumcompression.com/*. [Online; accessed 23th February 2010].

Betting, F. and Feldmar, J. (1995) 3D–2D projective registration of anatomical surfaces with their projections. In Bizais, Y. and Barillot, C. (eds), *Information Processing in Medical Imaging*, 275–286. Kluwer Academic Publishers, Dordrecht.

Cebrián, M., Alfonseca, M. & Ortega, A. (2007), The normalized compression distance is resistant to noise, *IEEE Trans. on Information Theory*, 53(5), 1895-1900.

Cilibrasi, R. & Vitanyi, P. (2005), Clustering by compression, *IEEE Trans. Information Theory*, 51(4), 1523-1545.

Cilibrasi, R., Vitanyi, P. & de Wolf, R. (2004), Algorithmic clustering of music based on string compression, *Journal of Computer Music*, 28(4), 49-67.

Cleary, J. G., Teahan, W. J. & Witten, I. H. (1995), Unbounded length contexts for ppm, in *DCC'95: Proc. of the Conference on Data Compression*, IEEE Computer Society.

Dubnov, S., Assayag, G., Lartillot, O. & Bejerano, G. (2003), Using machine-learning methods for musical style modeling, *Computer*, 36(10), 73-80.

Lan, Y. & Harvey, R. (2005), Image classification using compression distance, in *Proceedings of the 2nd International Conference on Vision, Video and Graphics*, pp. 173-180.

Lawder, J. K. & King, P. J. H. (2000), Using space-filling curves for multi-dimensional indexing, in *BNCOD 17: Proceedings of the 17th British National Conference on Databases*, Springer-Verlag, London, UK, pp. 20-35.

Li, M., Badger, J., Chen, X., Kwong, S., Kearney, P. & Zhang, H. (2001), An information-based sequence distance and its application to the whole mitochondrial genome phylogeny, *Bioinformatics*, 17(2), 149-154.

Li, M., Chen, X., Li, X., Ma, B. & Vitanyi, P. (2004), The similarity metric, *IEEE Transactions Information Theory*, 50(12), 3250-3264.

Li, M. & Vitanyi, P. M. (1993), *An Introduction to Kolmogorov Complexity and Its Applications*, Springer-Verlag, Berlin.

Li, M. & Zhu, Y. (2006), Image classification via lz78 based string kernel: A comparative study, in *Pacific Asia Conference on Knowledge Discovery and Data Mining*, pp. 704-712.

Macedonas, A., Besiris, D., Economou, G. & Fotopoulos, S. (2008), Dictionary based color image retrieval, *J. Visual Comunication and Image Representation*, 19(7), 464-470.

Rocha, J., Rosselló, F. & Segura, J. (2006), Compression ratios based on the universal similarity metric still yield protein distances far from cath distances, *CoRR: Computer Research Repository*, abs/q-bio/0603007.

Tran, N. (2007), The normalized compression distance and image distinguishability, in *Proceedings 19th IS&T/SPIE Symposium on Electronic Imaging Science and Technology*, San José, USA.

Vázquez, P., & Marco, J. (2012), Using Normalized Compression Distance for image similarity measurement: An experimental study, *The Visual Computer*, 28, 1-22, *to appear*, 2012.

Viola, P. & Wells, III, W. M. (1997), Alignment by maximization of mutual information, *Int. J. Computer Vision*, 24(2), 137-154.

Zitová, B. & Flusser, J. (2003), Image registration methods: a survey, *Image and Vision Computing*, 21, 977-1000.