August 2005

PICOREF[1] programme
RTPG[2]

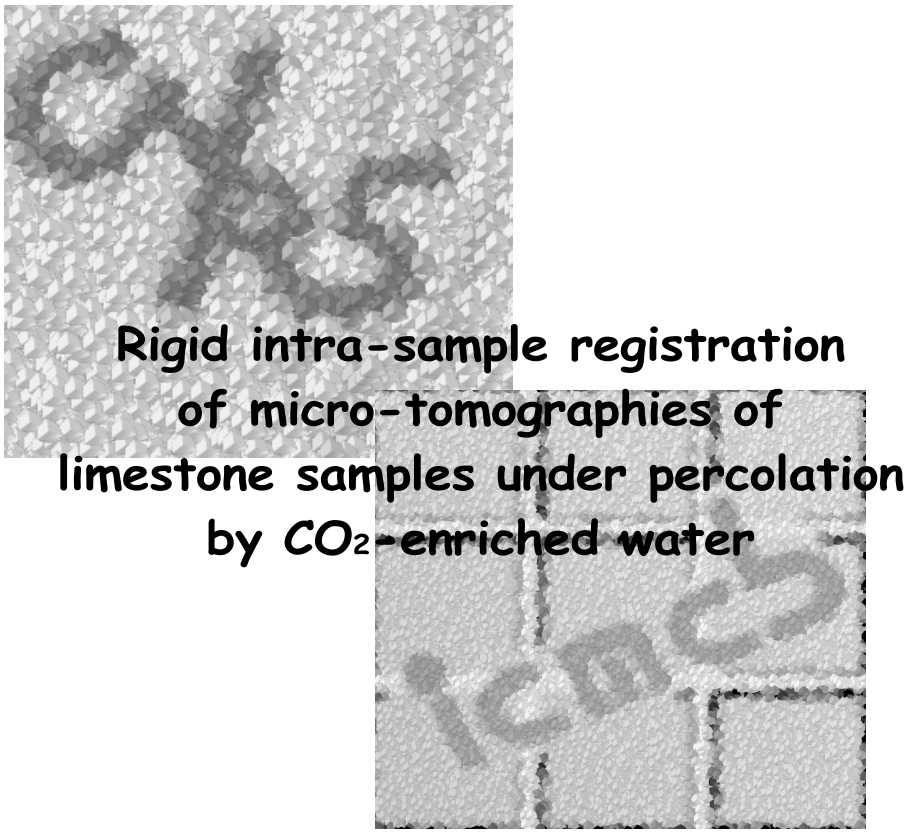

Rigid intra-sample registration of micro-tomographies of limestone samples under percolation by CO_2-enriched water

Erwan Plougonven
Dominique Bernard
ICMCB, group VIII

---

[1] **Piégeage du CO2 dans des Réservoirs géologiques en France**

[2] **Réseau des Technologies Pétrolières et Gazières**

Contents

List of Figures

1. Context

1.1. Introduction

The concept of the greenhouse effect originated in 1896 when Arrhenius established a correlation between atmospheric carbon dioxide concentration and ground temperature [1]. In the past 40 years, the CO_2 content of the atmosphere has increased from 315 to 360 parts per million, with its present world production reaching 7 billions tons per year (5.4 of which are due to the use of fossil fuels) [2, 3].

Growing concerns in climate changes pressure to reduce CO_2 emissions (as well as end-user energy demand). Immediate transition to nuclear, wind or solar energy is economically unsound, and since global hydrocarbon reserves can sustain energy demands well through the 21^{st} century [4], the possibility of CO_2 capture and sequestration is being thoroughly studied.

For such a process to significantly affect worldwide emissions, carbon sequestration must be examined at a very large scale (hundreds of billions of tons throughout the century). Since leakage rates, which are proportional to storage size and inversely proportional to storage lifetime, have to be taken into account, for the quantities mentioned, the chosen sequestration type should be very stable, with a lifetime of tens of thousands of years [5].

For this reason, certain methods seem unpropitious: terrestrial sequestration through the biomass is too marginal and oceanic sequestration has too short a lifetime (oceans have a natural turnover times of centuries). The most cost-efficient alternative seems to be geological sequestration, i.e. injection underground in porous sedimentary formations, either in depleted oil reservoirs or deep saline aquifers [6]. Because this method still presents uncertainties (seismic activity, contaminant mobilisation, ecosystem impact, ... [7]), integrity of each site should be studied beforehand. This is the main objective of the PICOREF (**PI**égeage du **C**O_2 dans les **R**éservoirs **G**éologiques **E**n **F**rance) programme: evaluation of potential French CO_2 injection sites.

Underground injection in porous alkaline mineral strata is likely to lead to carbonation, i.e. dissolution of the CO_2 in pore water, creating carbonic acid and raising the acidity of the solution, which would in turn erode the mineral (calcites or silicates) to form carbonates and bicarbonates. Although this mineral trapping permanently captures CO_2 in stable products, it can also weaken the injection site (leading to fractures in the caprock) or encourage CO_2 migration, both resulting in possible gas leakage [8]. The dynamics of the reaction must therefore be fully characterised before such solutions can be implemented. Predictive reaction-transport models for mineral sequestration are as of yet inaccurate and require more experimental measurements for their validation.

1.2. Quantification of sample evolution

1.2.1. Objectives

Another objective of the PICOREF programme is to continue the fundamental research from the previous PICOR programme, which studied the fluid-rock interaction to validate current reaction-transport models and thus the feasibility of massive CO_2 injection. The 3-year study, started in 2004 and conducted at the ISTEEM (Institut des Sciences de la Terre, de l'Environnement et de l'Espace de Montpellier), experiments on the percolation of porous or fractured reservoir rocks by supercritical fluids (fluids at pressures and temperatures above their critical values, i.e. values for vapourisation) containing CO_2, under conditions varying from 100°C and 9MPa to 150°C and 13MPa.

Evolution of the samples' properties during the percolations is recorded via X-ray microtomography (see appendix A for an overview of X-ray tomographic imaging). X-ray tomography being non-invasive, images are taken at different stages of the percolation, but before any quantification of the samples' properties' evolution can be calculated, the images must first be registered (i.e. have the position (i,j,k) in each of the images correspond to the same location in the sample).

1.2.2. Hypotheses

Because the experiments described in the previous section have not yet been completed, the images presented in this report are of porous reservoir rocks (specifically crinoidal limestone), 9 millimetres in diametre and 20 mm high (dimensions are chosen according to the field of view of the acquisition equipment, in this case the line ID-19 at the ESRF (European Synchrotron Radiation Facility)) being percolated under ambient temperature and pressure conditions with CO_2-enriched water. Specifications of this percolation can be found in [9]. The sample consists in a cylinder of limestone encased in a resin to direct the percolation flow and prevent the sample from disaggregating. This sample, for which the registration has been developed, has the same characteristics as the ones studied in the PICOREF programme (shape, dimensions, resin-encasing, and percolation procedure), which implies that it should give equally accurate results for these.

Figure 1.1.: Cross-section of one of the tomographies of the processed sample

The percolation obviously affects the geometric features inside the sample: the micrite (a sort of carbonate mud found in the porosities) can be moved or washed out, the pores become larger and their connectivity can change, the mineral is eroded in a non-uniform manner. This makes iconic or intensity-based registration (i.e. using the information of all the pixels without prior segmentation, as opposed to geometric or feature-based) unsuitable. Therefore, this process must first extract pertinent features that remain invariant throughout the percolation. This leaves several alternatives: the unseemly ones consist in using information from only one phase,

i.e. the mineral or the porosity, but directional percolation and complex pore structure might shift the phases' features in an unpredictible way. Intrinsic information from the mineral itself, such as its texture, could be used, this information is at best as important as the noise level, thus cannot be the basis of an accurate registration. The reasonable choice is to use the resin's interface with the sample, firstly because the resin is not affected by the percolation, therefore remains invariant (although sample manipulation in between the percolation stages results in minimal variations), and secondly because the interface presents conspicuous geometric features (since it moulds the polycrystalline heterogenous sample structure), i.e. provides sufficient information for precise registration[1].

An extensive literature exists on the subject of registration (almost entirely devoted to medical imaging), and compared to most registration applications, this one seems particularly easy (for an introduction to image registration, see [10]). Using classification terms described in [11], the type of registration required here can be characterised as: rigid (it involves computing only 6 parameters: a rotation and a translation (3 parameters each), compared to cases where scaling, shearing (or worse) is needed), instrinsic (the process does not involve the use of information outside of the sample, like exterior markers) and automatic (as opposed to interactive, for which the user has to manually steers the program towards the solution). A particularity of the microtomographies on which this project focuses, which proves to be a major hindrance in the choice of the procedure, is their size: the images contain around 6 billions pixels ($2000 \times 2000 \times 1500$), therefore brute force approaches, or even algorithms of quadratic complexity are inapplicable (see [12] for an exhaustive introduction to algorithm analysis and computational complexity theory).

1.2.3. Overview of the registration

For this reason the choice of the developed algorithm, summarised by the following steps, is not only linear in time, but decomposes the process for further acceleration and potential parallelisation :

Extract pertinent and invariant features (chapter 2) No matter the specificities of this application, the first step in feature-based registration is to select the points or features of the image that contain information pertinent to the registration. In this case, the set of points $\mathcal{P}$ of the interior edge of the resin (figure 1.2(b)) is selected: the resin remains intact during the percolation (invariance), and the interior edge presents all the irregularities of the limestone (pertinence).

Find a reference tube $\mathcal{T}$ that most closely matches $\mathcal{P}$ (chapter 3) Due to the sample's cylindrical shape and the fact that the extracted points are near the sample's surface, a tubular approximation (i.e. determining the tube that best fits on these points, figure 1.2(c)) would allow for a very simple calculation of half of the solution, i.e. by registering the tomographies' reference tubes. The rest of the process would then consist of a translation along the cylinder's axis and a rotation around it. The decomposition in itself does not accelerate the registration whatsoever, but it allows a subsequent data reorganisation that will.

[1] A sensible question would be why not simply use the outside of the encasing resin. The answer is that it is not visible in the tomographies. Since reconstruction is very time-consuming, all pixels of the final image outside the region of interest are ignored and set to 0.

Reorganise the data (chapter 3) A distance map $\mathcal{D}(\mathcal{P}, \mathcal{T})$ can then be calculated, which results in a two-dimensional image (in polar coordinates) that still contains all the information pertinent to the registration without any loss of accuracy (figure 1.2(d)). The rest of the registration boils down to an alignment of the tomographies' distance maps (horizontal and vertical alignments are respectively equivalent to rotation around and translation along the tube's axis).

Align the distance maps (chapter 4) Through a search for the global minimum of a customised cross-correlation function of the distance maps, the horizontal and vertical translations that aligns the distance maps is computed. With the data reorganised in two-dimensional images, the computation time for this part is minimal.

Compute the result (chapter 5) From the transformation that registered the reference tubes and the translation that aligns the distance maps, the final transformation matrix that registers the tomographies can be calculated.

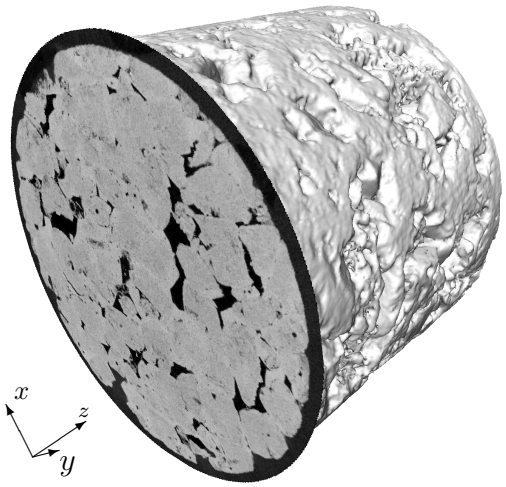

(a) Sample tomography (for the acquisition, the sample is placed vertically, i.e. the sample's axis is roughly aligned with the z-axis of the image)

(b) Set of pertinent points $\mathcal{P}$

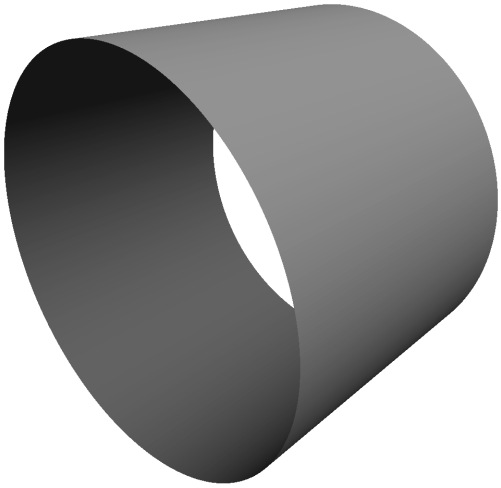

(c) Approximated cylinder $\mathcal{T}$

(d) Distance map $\mathcal{D}(\mathcal{P}, \mathcal{T})$ (since this 2D image is in polar coordinates, the horizontal and vertical axes are respectively the angle around and the distance along the sample's axis (and not the z-axis of the tomography))

Figure 1.2.: Illustration of how the registration data is extracted and reorganised for each image

2. Extracting the invariant features

2.1. Direct approach

2.1.1. Overview

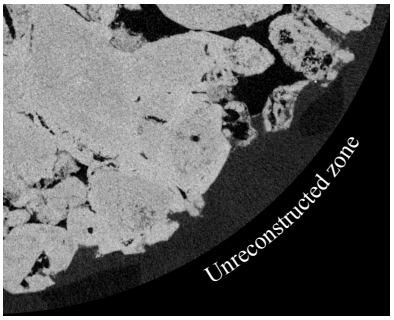

(a) Zoom on a portion of the tomography

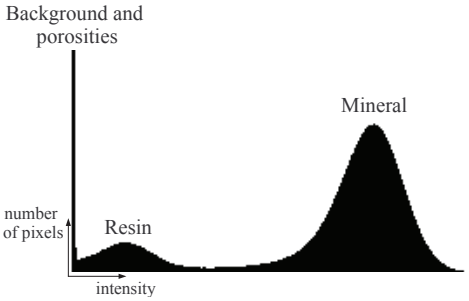

(b) Histogram of the tomography

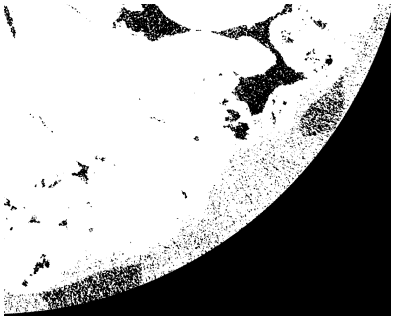

(c) Porosities and background

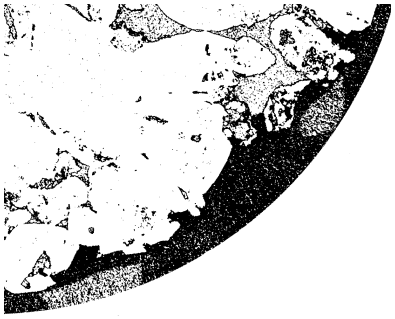

(d) Resin

(e) Mineral

Figure 2.1.: From the original image and its histogram ((a) and (b)), a simple thresholding can be used for segmentation ((c)–(e), where the points in black belong to the desired phase).

The tomographies have three distinguishable phases: the mineral, the resin and the porosity (in decreasing intensity), as shown in figure 2.1. There are indeed three distinct peaks on the histogram (c.f. 2.1(b)), which could correspond to those phases, in which case a simple thresholding filter would be reasonable means to extract the points $\mathcal{P}$ of the interior edge of the resin. Unfortunately, separation of the background and the resin is mediocre due to the noise level, so this type of filter as a basis for finding $\mathcal{P}$ is not ideal, but it is still useable: the vast majority of $\mathcal{P}$ belongs to the interface between the resin and the mineral (i.e. the interface between the resin and the porosity is only a small fraction of $\mathcal{P}$, for obvious physical reasons), so the essential segmentation would be of the mineral (as in 2.1(e)). In other words, a good approximation of $\mathcal{P}$ is the outer edge of the mineral phase. The difference between the two is

illustrated in figures 2.4(b) and 2.4(d), and can be seen in the top-right portion of figure 2.1(a). It will introduce errors on the distance map $\mathcal{D}(\mathcal{P}, \mathcal{T})$ (not on the definition of the reference tube itself), but in an negligible range. The resulting pixels extracted with this method is illustrated in figure 2.2.

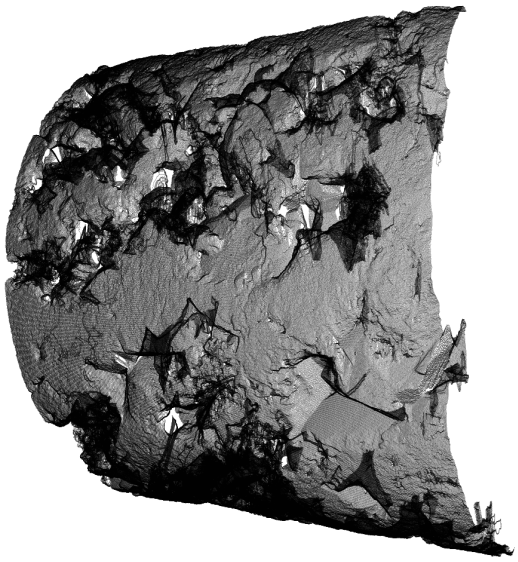

Figure 2.2.: A subset of $\mathcal{P}$

2.1.2. Technical details

Basic notions of digital topology have been put into practise in order to find $\mathcal{P}$. Understanding the technical aspects require a few prior definitions, such as neighbourhood relations and connected components. For a more thorough introduction to digital topology, consult [13].

Preliminaries

Let I be a binary image, i.e. an application $\mathbb{Z}^k \longrightarrow \{0, 1\}$ (we will introduce the notions with the case $k = 2$, but they can be intuitively extended to any $k \in \mathbb{N}$). A point (or pixel) p is defined by its indices $(p_1, p_2) \in \mathbb{Z} \times \mathbb{Z}$, and its intensity is noted $I(p)$. We consider two neighbourhood relations Γ_4 and Γ_8, defined, for every point p, by $\Gamma_4 = \{p' = (p'_1, p'_2); |p'_1 - p_1| + |p'_2 - p_2| \leqslant 1\}$ and $\Gamma_8 = \{p' = (p'_1, p'_2); \max(|p'_1 - p_1|, |p'_2 - p_2|) \leqslant 1\}$, and denote by $\Gamma^*_4(p)$ and $\Gamma^*_8(p)$ the sets $\Gamma_4(p) \setminus p$ and $\Gamma_8(p) \setminus p$. A point p' is said to be *n-adjacent* to p (with $n = 4$ or 8) if $p' \in \Gamma^*_n(p)$. Let $X \subset \mathbb{Z}^2$, $\overline{X}$ the complementary set of X, and $\{p, p'\} \subset X$ (X is often referred to as the object or the foreground, and $\overline{X}$ the background). We say that there is a *n-path between p and p'* if there exists a sequence of points $\{p_1, p_2, \ldots, p_m\}$ such that $p_1 = p$, $p_m = p'$, and $\forall i \in [1; m-1]$, p_i is n-adjacent to p_{i+1}. We say that X is *n-connected* (or an n-connected component) if $\forall \{p, p'\} \in X, \exists$ an n-path between p and p'.

Extracting the correct outer edge of the mineral phase

These notions are used on the tomographies binarised with a thresholding method, as in 2.1(e) (the black points being the foreground), where all pixels with intensity strictly lower than the threshold value belong to the background (supposably only the porosity and resin phases (and the unreconstructed portion of the image)). They are presented for two-dimensional images because the algorithm will be applied on each axial cross-section (i.e. the cross-sections perpendicular to the z-axis, which roughly corresponds to the sample's axis). This simplification is done for several reasons explained in the next section.

The process is divided in two parts: the first will identify the 4-connected components of the background which are connected to the image's border (here the image refers to an axial cross-section of the tomography), and mark the pixels of the foreground that are 4-connected to those components as the outer edge. The result of this is illustrated in figure 2.3(b). The object (or foreground) of these binary images should intuitively consist of one connected component, but it is hardly ever true, for several reasons:

- the binarisation with a simple thresholding filter is not perfect due to the noise in the image, so points from the porosity or the resin can be wrongfully classified in the foreground
- during the creation of the sample, limestone fragments can detach and sink in the resin, leading to another connected component (like the one that disappeared between 2.3(b) and 2.3(c))
- since the physical object is defined in three-dimensions (excluding string theory), there can exist discontinuities from one cross-section to another, i.e. mineral protrusions or protuberances extending along the z-axis in the resin, leading to, when looking at a horizontal cross-section, another connected component

This implies that the result contains several outer edges (which, in light of the way an outer edge was defined, are 8-connected components), therefore the correct one needs to be selected. Fortunately, the desired component has an ostentatious particularity: it's the biggest one, which is what the second part consists of, extracting the outer edge's 8-connected component containing the most number of points, as in figure 2.3(c).

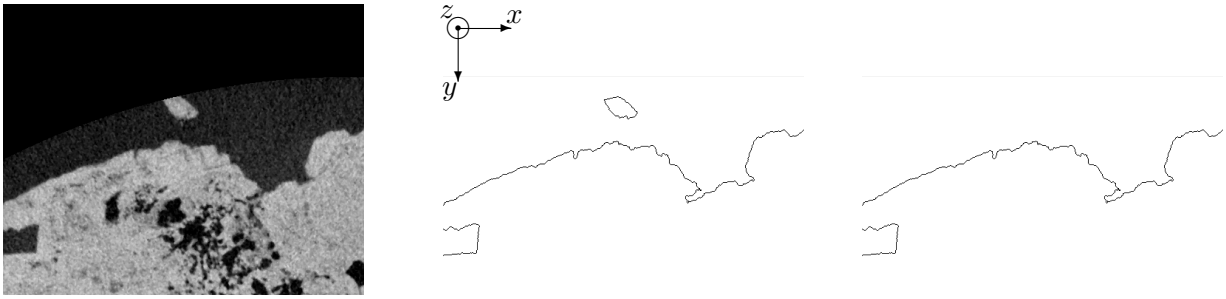

(a) Portion of a cross-section of a tomography of the sample

(b) Outer edge of the mineral phase

(c) Biggest connected component of the outer edge

Figure 2.3.: Illustration of the outer edge extraction

Method justification

Using such a thresholding technique yields acceptable results with a fixed threshold value for all tomographies of a given sample because the final post-processing of the reconstruction of the tomographies, which involves transtyping the pixels' intensities from real to integer values coded in one byte (i.e. 256 levels of grey), overlays the histograms' maxima, so that their shape and position on the $[0; 255]$ interval corresponds. This means that the mineral phase is found in the same intensity range in all tomographies, for a given sample.

Even though all the notions used for this method are extensible to three dimensions, the process is still applied on each axial cross-section independently, most importantly because finding the biggest connected component is a global operation impossible to implement for very large data structures without a time-consuming block-decomposition scheme, but also because it slightly lessens the imprecision of the method. The reason for this is explained with the figure 2.4, which shows that the result of the three-dimensional algorithm would include *all* the points of the edges of the porosities in contact with the resin (figure 2.4(d)), porosities which are likely to evolve during the percolation and thus give inaccurate information for the registration. When applied on axial cross-sections independently, (only) for porosities of similar shape and orientation as depicted in 2.4, the points of the mineral-porosity interface cannot be reached (c.f. figure 2.4(c)) so will not be used for the registration[1].

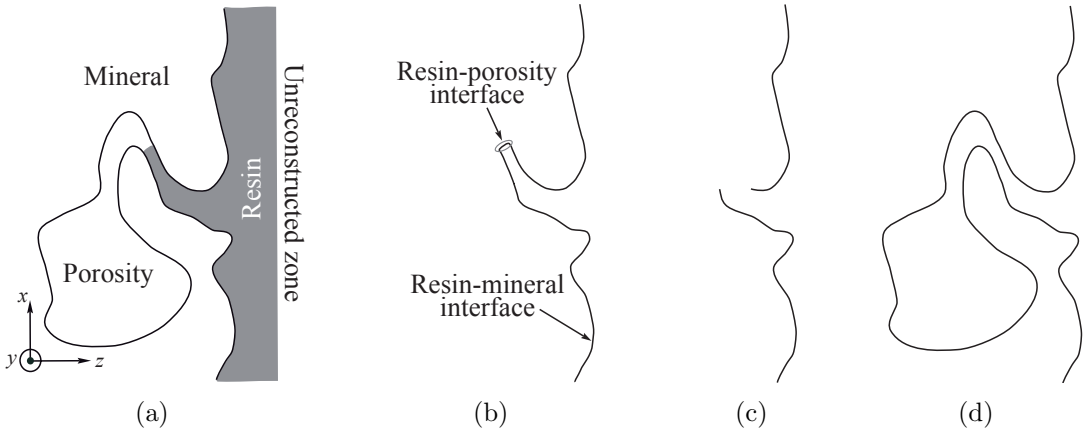

Figure 2.4.: A situation where independent processing of the axial cross-sections is less inaccurate. Illustrations of coronal cross-sections of : (a), a portion of the original tomography, (b), the corresponding ideal set $\mathcal{P}$, (c), the outer edge of the mineral phase for every axial cross-section (the method used), and (d), the outer edge of the mineral phase.

2.2. Class of alternatives

2.2.1. More preliminaries

The most important flaw of the method described is that, by construction, it will miss all the pixels from the resin-porosity interface. All the attempts made to solve this needed four

[1]It creates a discontinuity in the extracted surface, but is in no way relevant to the process's accuracy

more digital topology definitions. We call an *dilation* of an image I (using n-connectivity) the application: $I(p) \longmapsto \Gamma(I(p)) = \max_{p' \in \Gamma_n(p)} I(p')$ The complementary definition, a *erosion*, is defined in a dual way: $I(p) \longmapsto \Gamma^{-1}(I(p)) = \min_{p' \in \Gamma_n(p)} I(p')$. Dilation and erosion are defined in such a way as to be naturally applicable to greyscale images.

We call an *opening* of an image I (using n-connectivity) the application consisting of an erosion followed by a dilation, i.e. $\Gamma^{-1} \circ \Gamma$. The dual application, a *closing*, is the inverted sequence $\Gamma \circ \Gamma^{-1}$. We define an opening (or closing) of size i as i erosions followed by i dilations (or vice-versa).

2.2.2. Failed alternatives

All tested alternatives were similar in design and failed for the same reason. One example of such a method is depicted in figure 2.5 which describes the process and shows two different situations, the upper one leading to the desired result, and the lower one demonstrating its uselessness. The main problem that arises when trying to segment only the resin phase, which has an intensity in between that of the pores and the mineral, is that a thin layer of pixels in between the mineral and the pores will be considered as part of the resin. If a pore is in contact with the resin, then this layer of pixels will be connected to the ones from the resin, as in figure 2.5(c). The easiest way to remove this connection is with an opening, which removes thin features. Unfortunately, because air bubbles appeared in the resin during the fabrication of the sample, some areas in the resin show up darker in the tomography. This leads to a fragile segmentation of that phase, as most of the pixels in those areas are considered as not belonging to it[2]. Thus while the opening disconnects the undesired mineral-pore interface pixels from the resin, it also creates breaches in the resin (lower half of figure 2.5(d)), which results in taking pixels from the inside of the resin phase and even pixels from the boundary of the unreconstructed zone (figure 2.5(f)) instead of the desired ones.

2.2.3. Other paths to investigate

Other possibilities for extracting the ideal set $\mathcal{P}$, which are theoretically promising, can be described in two distinct steps: first, remove as much noise as possible in order to have clearly separate intensities for each phase (good results seemed to be achieved by combining a strong gaussian blur, a median filter, and a conditional diffusion[3]); secondly, by using notions of cross-section topology such as lower levelling kernels of upper regularisation on the gradient of the image [15] (the use of filtered leveling might bypass the need for the first step) to conserve the crests, from which $\mathcal{P}$ could easily be extracted. Nevertheless, supposing it does extract the ideal set of points sought, its effect on the accuracy of the registration is not at all certain.

[2]lowering the lower threshold of this segmentation leads to incorporating points of the porosities
[3]inappropriately named anisotropic diffusion by [14]

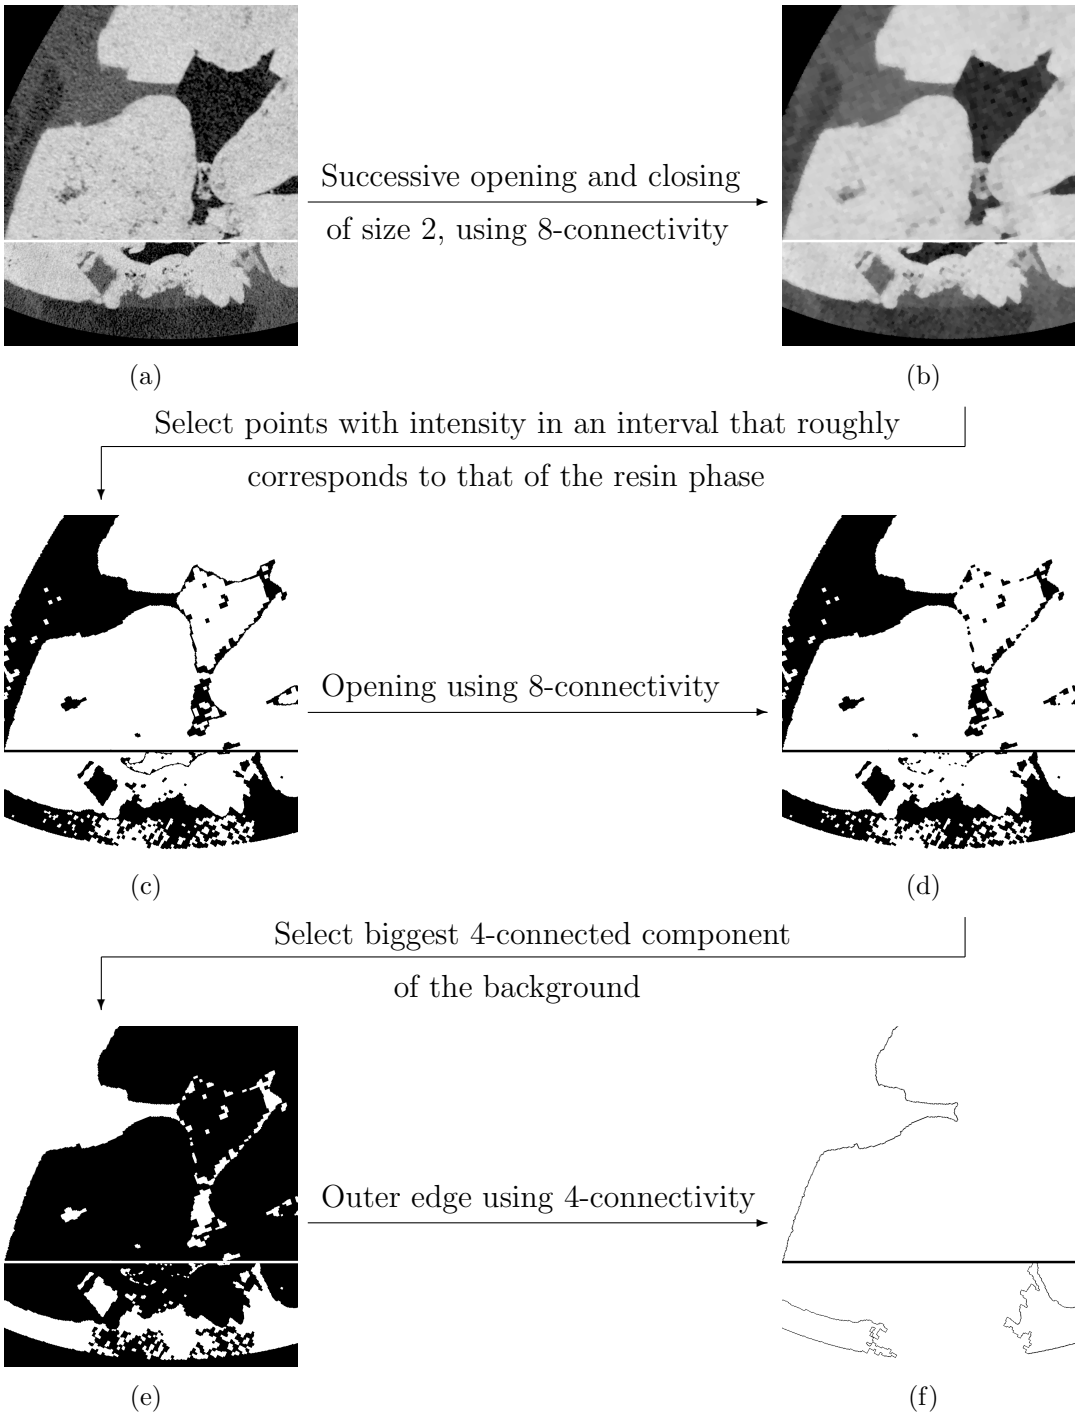

Figure 2.5.: A failed alternative method for selecting $\mathcal{P}$, illustrated by two extracts from the tomography, the top extract showing what was intended with this method. The first stage removes some of the noise that corrupts the following thresholding step. The mineral segmentation step (a double thresholding) that gives (c) also extracts a thin layer of pixels between the pores and the mineral, so the following opening, which gives (d), is used to disconnect that layer from the set that actually represents the resin. The biggest connected component of the background of (d) is supposed to correspond to the inside of the sample, i.e. only the pore and the mineral, giving (e) whose outer edge should be the inner edge of the resin phase. Except the bottom extract shows where it fails.

3. Reorganising the data

The set $\mathcal{P}$ of points, considered as containing the necessary information to compute the registration, is now used in two interrelated ways: firstly, the set is approximated by a reference tube $\mathcal{T}$ in the manner of a best-fit approach, and secondly, the set and the tube are used to compute a distance map $\mathcal{D}(\mathcal{P}, \mathcal{T})$ (represented as a two-dimensional image) of one from the other. Using both $\mathcal{T}$ and $\mathcal{D}(\mathcal{P}, \mathcal{T})$ from each tomography of a given sample, the complete registration transformation can be determined.

3.1. Description

Three essential hypotheses are used to prove the validity the registration process and justify its efficacy in terms of computational complexity:

1. the sample is almost cylindrical
2. $\mathcal{P}$ represents the edge of the sample
3. the sample's axis is almost aligned with the z-axis of the tomography

The first suggests the help of the shape in the registration: if the sample's image in each tomography was approximated by a reference cylinder, then the registration problem can be reduced to two simpler subproblems: determining the transformation from one cylinder to another, and finding the rotation around and translation along one's axis. The second hypothesis implies that $\mathcal{P}$ and only $\mathcal{P}$ is necessary for determining the reference tube $\mathcal{T}$ and thus a cylindrical approximation of the sample.

The tube $\mathcal{T}$ and the distance map $\mathcal{D}(\mathcal{P}, \mathcal{T})$ can be used to represent the positions of the points of $\mathcal{P}$. Instead of using a set of triplets (i, j, k) (defined in the Euclidean space of the tomography) of all the points of $\mathcal{P}$, an image defined in the tubular plane (O, θ, z') (z' being the height along the cylinder's axis and θ the angle around it) containing the distances of the points of $\mathcal{P}$ from $\mathcal{T}$ (a simple example is shown in figure 3.1, where the sample is a hexagonal cylinder) can be used for the latter subproblem, so that it is reduced to an alignment of distance maps, the horizontal translation (on θ) from this alignment corresponding to the rotational part, and the vertical translation (on z') to the translational part of the subproblem.

The approach first requires that the reference tube be determined. From the first hypothesis and from the fact that the outer edge of the coat of resin is not visible on the tomography, the best definition of $\mathcal{T}$ is the tube with the smallest radius and containing $\mathcal{P}$[1]. In light of the amount of data to use to determine such a cylinder, a local method (one allowing the parallelisation of the problem) should be used. Knowing that the intersection of a tube and a

[1]Note that only the axis of the reference tube is essential for the registration, its radius just changes a global offset on the projection. Nevertheless, it is impossible to determine an axis without having previously found the tube, since its axis is corollary.

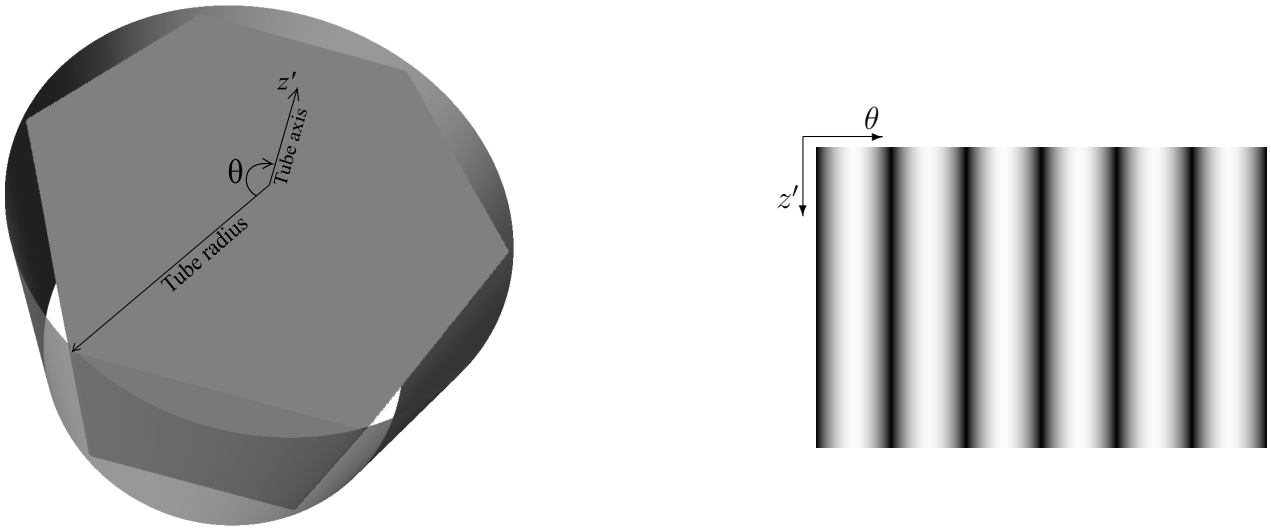

(a) Determination of the reference tube

(b) Projection on the tube

Figure 3.1.: Example illustrating the use of a reference tube and polar coordinates: the sample is a hexagonal cylinder and the pixels of $\mathcal{P}$ are on its rectangular facets. The determined tube surrounds the sample, in contact with its six edges. Figure (b) shows the resulting projection on the tube, starting (on the left), arbitrarily, from the edge joined by the radius drawn in (a), and moving around clockwise (again, chosen arbitrarily) until the angle θ reaches 2π. The light zones of (b) correspond to the facets of the hexagonal cylinder, whose centres are farthest from the tube (thus whitest in the image), and the dark areas correspond to the edges (there are seven vertical dark zones while the hexagon has six edges because the first and last zones in (b) originate from the same edge).

plane is an ellipse (supposing that tube and plane are infinite and not parallel), a local method would consist firstly in finding the smallest ellipse surrounding the points of $\mathcal{P}$ in each plane, and secondly in finding the cylinder that best fits this set of ellipses.

The third hypothesis not only means that the reference tube is not parallel to an axial plane but also allows a further simplification with a negligeable loss of accuracy: since the sample's axis is almost perpendicular to the axial plane, and knowing the shape of the intersection between a tube and a perpendicular plane, a set of circles would be determined instead of ellipses.

3.2. Finding the circles

Let $\mathcal{P}_i \in \mathcal{P}$ be the points contained in the i^{th} axial cross-section (this means that $\bigcup_{i\in[1,d]} \mathcal{P}_i = \mathcal{P}$, d being the number of axial cross-sections in the tomography).

3.2.1. Geometric optimisation: the discrete 1-centre problem

Let $\mathcal{P}_i$ be a finite set of points in a Euclidean plane.

What is the smallest circle containing all points in $\mathcal{P}_i$?

This seemingly easy problem first appeared in 1857 [16], and has been thoroughly studied, solved and generalised since.

3.2.2. The implemented solution

The chosen solution implements the Elzinga-Hearn algorithm [17]:

1. Choose two points $\{p_1, p_2\} \in \mathcal{P}_i$
2. Construct the circle $\mathcal{C}$ with centre c and radius r whose diameter is $[p_1p_2]$

   If $\forall p \in \mathcal{P}_i$, $\|\overrightarrow{pc}\| \leqslant r$, then $\mathcal{C}$ is the smallest circle ■
   Otherwise, choose a point $p_3 \in \mathcal{P}_i$, $\|\overrightarrow{p_3c}\| > r$
3. **If** the triangle $p_1p_2p_3$ is right or obtuse, then rename the two points opposite the right/obtuse angle as p_1 and p_2, ang go to step 2.
   Otherwise, construct the circle $\mathcal{C}'$ with centre c' and radius r' circumscribing the acute triangle $p_1p_2p_3$

   If $\forall p \in \mathcal{P}_i, \|\overrightarrow{pc'}\| \leqslant r'$, then $\mathcal{C}'$ is the smallest circle ■
   Otherwise, choose a point $p_4 \in \mathcal{P}_i, \|\overrightarrow{p_4c'}\| > r'$
4. Let $s = \max_{p \in \{p_1,p_2,p_3\}} \|\overrightarrow{p_4p}\|$ and $t \in \{p_1, p_2, p_3\}$ such as $]tp_4[\cap (sc') \neq \emptyset$ (e.g. figure 3.2).

   Rename the points $\{s, t, p_4\}$ as $\{p_1, p_2, p_3\}$ and go to step 3.

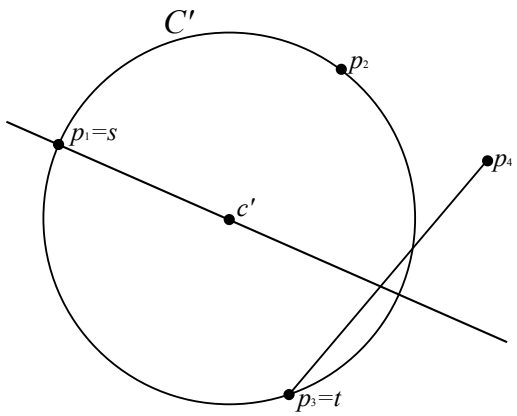

Figure 3.2.: Illustration of the fourth step of the Elzinga-Hearn algorithm: p_1 is farthest from p_4 and is therefore s, p_3 is the point on (sc')'s side opposite to p_4 and is therefore t.

3.2.3. Alternatives

The worst-case scenario can give the Elzinga-Hearn algorithm a complexity of $O(n^2)$ [18]. Megiddo presented in [19] an linear-time algorithm for the unweighted 1-centre problem, but apparently, in practical problems, the Elzinga-Hearn algorithm remains faster [20], although this would have to be verified.

Avoiding the last simplification and determining a set of ellipses instead of circles is a questionable improvement because not only will it hinder computation speed, but its effect will be negligible on the approximation and even more imperceptible on the registration.

3.3. Approximating parallel circles by a tube

3.3.1. Description

We suppose that we now have a set of circles parallel to the axial plane. From these, a tube must be determined (see figure 3.3). Concerning the radius of this cylinder, an average of the circles' radii is taken, although it is unimportant, as explained in the footnote page 16. Axis-wise, the problem is to find a line from a set of points, wich are the circles' centres, problem which is commonly called a linear regression.

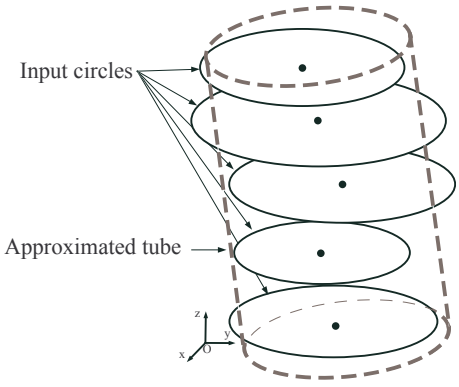

Figure 3.3.: Approximation of a set of circles with a tube

3.3.2. Linear regression

Let $\mathcal{S} \subset \mathbb{R}^3$ be a set of d points that represent the circles' centres. We note $\mathcal{S} = \{p_1, p_2, \ldots, p_d\}$ and $\forall i \in [1; d], p_i = \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix}$.

The problem is to project $\mathcal{S}$ on the one-dimensional Euclidean subspace (i.e. the line) that minimises its distortion: if we define a line as a pair $(a, \overrightarrow{b}) \in \mathbb{R}^3 \times \mathbb{R}^3$, where a is a point on the line and $\overrightarrow{b}$ the vector which defines its direction, the line that best fits the points is defined by $\min\limits_{(a,\overrightarrow{b})} \sum\limits_{p_i \in \mathcal{S}} \|p_i(a + (\overrightarrow{Oa}.\overrightarrow{b}).\overrightarrow{b})\|$.

The solution that was implemented decomposes the problem: project $\mathcal{S}$ on two orthogonal planes (coronal and sagittal), perform a linear regression of the two projected sets separately in their respective planes, and back-project those lines in 3D space so as to obtain an intersection of two planes, which is the result.

Let $(\mathcal{S}', \mathcal{S}'') \subset \mathbb{R}^2 \times \mathbb{R}^2$ be the projections of $\mathcal{S}$ on the coronal and sagittal planes (Oxz) and (Oyz) respectively[2], as shown in figure 3.4, and note $\mathcal{S}' = \{p'_1, p'_2, \ldots, p'_d\}$ and $\mathcal{S}'' = \{p''_1, p''_2, \ldots, p''_d\}$, i.e. $\forall i \in [1; d], p'_i = \begin{pmatrix} x_i \\ z_i \end{pmatrix}$ and $p''_i = \begin{pmatrix} y_i \\ z_i \end{pmatrix}$.

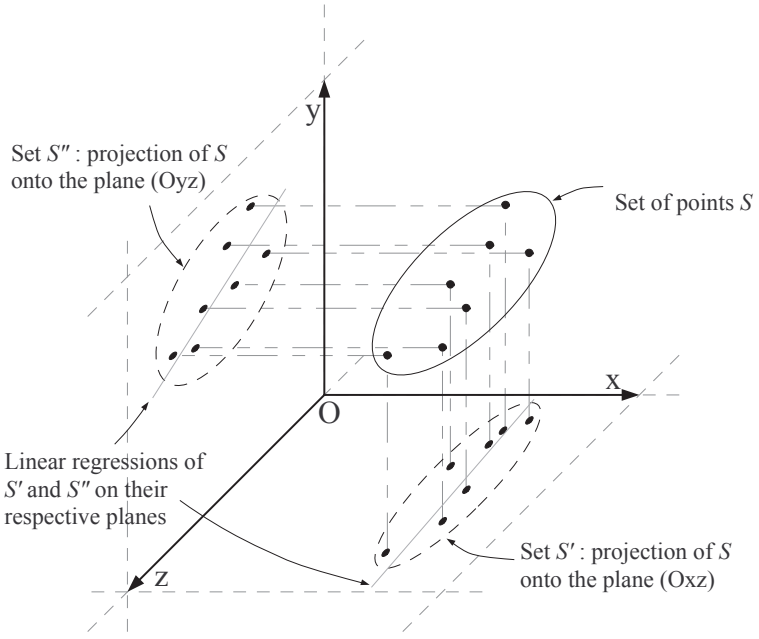

Figure 3.4.: Decomposition of the linear regression

Let Δ, Δ' and Δ'' be the respective linear regressions of $\mathcal{S}$, $\mathcal{S}'$ and $\mathcal{S}''$. Δ' and Δ'' are defined by the equations $x = a_1.z + b_1$ and $y = a_2.z + b_2$, both in $\mathbb{R}^2$, but if they are extended to $\mathbb{R}^3$, these equations describe planes. If we now suppose that $(\Delta', \Delta'') \subset \mathbb{R}^3 \times \mathbb{R}^3$, then $(\Delta' \perp (Oxz)) \wedge (\Delta'' \perp (Oyz))$, and since $(Oxz) \perp (Oyz)$, then $\Delta' \perp \Delta''$. This is to affirm that the planes Δ' and Δ'' are not parallel, or that they are bound to intersect, and their intersection is Δ, which is obviously defined by $\begin{cases} x = a_1.z + b_1 \\ y = a_2.z + b_2 \end{cases}$.

The problem is therefore to find a_1, b_1, a_2 and b_2. Using the method of least squares, if we note $\overline{u} = \frac{1}{d}\sum_{i=1}^{d} u_i$ the average of u_i, $C_{uv} = \frac{1}{d}\sum_{i=1}^{d}(u_i - \overline{u}).(v_i - \overline{v})$ the covariance between u and v, and $\sigma_u^2 = C_{uu}$ the variance of u, then $a_1 = \dfrac{C_{xz}}{\sigma_z^2}$, $b_1 = \overline{x} - a_1.\overline{z}$, $a_2 = \dfrac{C_{yz}}{\sigma_z^2}$ and $b_2 = \overline{y} - a_1.\overline{z}$.

If the computation of (a_1, b_1, a_2, b_2) is implemented literally from the previous results, two passes on $\mathcal{S}$ are necessary: a first pass to find $(\overline{x}, \overline{y}, \overline{z})$, and a second pass for $(\sigma_z, C_{xz}, C_{yz})$, dependent on the first pass's result. The formulas for variance and covariance can be changed

[2]Regression, by definition, is the estimation of a conditional expected value of an output variable (also called dependent or endogenous variable), given an input (or independent, or exogenous) variable. This concept fits well for our problem, since the z-coordinates of the points of $\mathcal{S}$ are known in advance, the points being evenly distributed along the z-axis. Therefore, it feels natural for the z-coordinates to serve as input variable for the linear regressions.

so as to need only one pass on the points of $\mathcal{S}$. The result (only shown for a_1 and b_1) is written:

$$a_1 = \frac{d\sum_{i=1}^{d} x_i.z_i - \sum_{i=1}^{d} x_i.\sum_{i=1}^{d} z_i}{d\sum_{i=1}^{d} z_i^2 - \left(\sum_{i=1}^{d} z_i\right)^2} \qquad b_1 = \frac{\sum_{i=1}^{d} z_i^2.\sum_{i=1}^{d} x_i - \sum_{i=1}^{d} z_i.\sum_{i=1}^{d} x_i.z_i}{d\sum_{i=1}^{d} z_i^2 - \left(\sum_{i=1}^{d} z_i\right)^2}$$

3.4. Obtaining the distance map

In order to simplify the remaining processing, the reference cylinders will be defined as a triplet $(Q, R, \mathfrak{r}) \in \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}$, where Q and R are the intersections of the cylinder's and the planes $z = 0$ and $z = d$ respectively, and $\mathfrak{r}$ is its radius.

3.4.1. Overview

Once the reference tube $\mathcal{T}$ has been determined, $\mathcal{D}(\mathcal{P}, \mathcal{T})$ can be calculated. The intensities of the pixels in this distance map, which represent the respective distances from the reference tube of the points of $\mathcal{P}$, are actually signed values, the sign indicating whether the pixel is inside or outside $\mathcal{T}$ (negative was chosen arbitrarily to denote interior points)[3]. In the example depicted in figure 3.5(a), M and M' belong to $\mathcal{P}$, P and P' are their respective projections on $\mathcal{T}$ and the grid drawn on the tube represents the pixels of the resulting 2D image. As is shown, each point's projection falls inside a pixel of the result, whose intensity will equal the signed distance between the point and its projection (for the two pixels considered in the figure, their intensities will be respectively $-\|\overrightarrow{PM}\|$ and $\|\overrightarrow{P'M'}\|$).

Depending on the resolution of the resulting image, it is possible that more than one point is projected in a given pixel, in which case an averaging of the distances is done.

3.4.2. Conditions for the consistency of the ensuing alignment

In order to be able to perform a meaningful alignment of the distance maps, two necessary conditions must be met.

Distance coherency

This concerns the resolution of the projections. For example, if the image of an object in one picture is twice the size of the same object's image in another picture, then registering them will require scaling. Dimensional coherency means that no scaling is needed. This implies that for two points M_1 and M_1' (belonging to $\mathcal{P}$) in one tomography which correspond to two other points M_2 and M_2' in another, the distance $\|\overrightarrow{P_1P_1'}\|$ of the first pair's projection on $\mathcal{T}$ must

[3]Using the set of circles which surround all the points of $\mathcal{P}$ to find the tube does not imply that the tube contains all the points of $\mathcal{P}$.

coincide with the distance $\|\overrightarrow{P_2P_2'}\|$ of the second. This condition is implicitely met along the θ-axis of the projection[4], but asks for a judicious choice of vertical resolution of $\mathcal{D}(\mathcal{P}, \mathcal{T})$.

Let $\{M_1, M_1'\} \subset \mathcal{P}_1$ be two points of one tomography, $\{P_1, P_1\}$ their projections, $\{M_2, M_2'\} \subset \mathcal{P}_2$ two points in a second tomography (not necessarily corresponding to the first pair) and $\{P_2, P_2'\}$ their projections. The condition can be written as follows:

$$\frac{\|\overrightarrow{M_1M_1'}\|}{\|\overrightarrow{P_1P_1'}\|} = \frac{\|\overrightarrow{M_2M_2'}\|}{\|\overrightarrow{P_2P_2'}\|}$$

Position coherency

This point underlines the importance of keeping track of the relation between z'-coordinates in the distance map and the position along the tube's axis. For instance, if the boundaries of the distance map were chosen in such a way as to include, in the range they define, all projections of the points of the sample's image, then such boundaries would depend on the orientation of the axis, in which case the relation would be lost. The simplest solution would be that Q (or rather the plane perpendicular to $\mathcal{T}$ and that contains Q) corresponds to $z' = 0$ and that R corresponds to $z' = h$, the height of the distance map, as illustrated in figure 3.5(b).

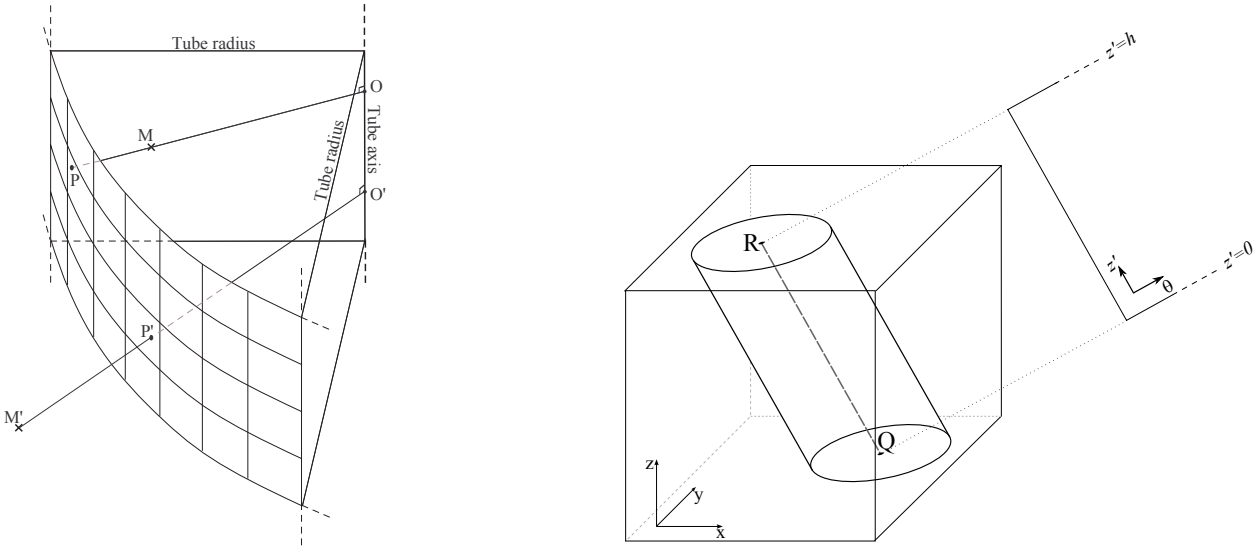

(a) Construction

(b) Condition for alignment consistency

Figure 3.5.: Definition of $\mathcal{D}(\mathcal{P}, \mathcal{T})$

[4]If the reference tubes of the tomographies to register are accurate, they have the same radius, and if the resolution along the θ axis of the projections were chosen in the same way, then the widths of the projections are equal, therefore no horizontal scaling is required. In practical cases, the width of the images varied by almost 0.05%, in which cases they should be resized to have the same width before finding correlations

4. Aligning the distance maps

4.1. Image alignment

Let the reference image be $I_r : [0; w] \times [0; h_r] \longrightarrow \mathbb{R}$ and the model image, i.e. the one on which the alignment transformation will be applied, $I_m : [0; w] \times [0; h_m] \longrightarrow \mathbb{R}$.

4.1.1. Cross-correlation

The basic problem, which is in this case an intensity-based alignment is to find the maximum correlation coefficient[1] of the images when varying the translation, i.e.

$$\max_{(T_x,T_y)} \frac{\sum_{x,y} I_r(x,y).T(I_m(x,y)) - \sum_{x,y} I_r(x,y) \sum_{x,y} T(I_m(x,y))}{\sigma_{I_r}.\sigma_{T(I_m)}}$$

where σ_I is the standard deviation of the image I, and T is the application:

$$\begin{array}{rccl} T: & \mathbb{R}^2 & \longrightarrow & \mathbb{R}^2 \\ & I(x,y) & \longmapsto & T(I(x,y)) = \begin{cases} I(x-T_x, y-T_y) & \textbf{if } (0 < x-T_x < w) \wedge (0 < y-T_y < h_r) \\ I(w+(x-T_x), y-T_y) & \textbf{if } (-w < x-T_x < 0) \wedge (0 < y-T_y < h_r) \\ I(x-T_x-w, y-T_y) & \textbf{if } (w < x-T_x < 2w) \wedge (0 < y-T_y < h_r) \\ \emptyset & \textbf{otherwise} \end{cases} \end{array}$$

This special translation application has two traits specific to this registration process: it has no output outside the vertical bounds of I_r (there is no correlation where the two images do not overlap), and the translation repeats I_m along the x-axis, since the x-coordinate represents the angle around the reference tube's axis, which is periodic.

The implementation of the cross-correlation function calculates the correlation coefficient for integral translations, and locates the maximum, which is not optimal since only integral translations were tested. To refine this maximum, a non-linear interpolation is performed in its 4-neighbourhood. More precisely, a natural cubic spline interpolation was chosen.

4.1.2. Natural cubic splines

A spline is a piecewise polynomial function that locally has a simple form and presents globally a good flexibility and smoothness. A cubic spline has n third-order polynomial functions which pass through $n + 1$ control points. Since a third-order plynomial is defined by four coefficients $\{a, b, c, d\}$ such that $Y(t) = at^3 + bt^2 + ct + d$, the interpolation consists in finding the $4n$ coefficients necessary to define all the piecewise functions.

[1] The Pearson product-moment correlation coefficient was used

To find the maximum correlation, two splines are used, along each axis. We describe the process along the x-axis, but the same is implicitely applied along the other direction. Let ρ_1, ρ_2 and ρ_3 be the correlations between I_r and respectively $T_{(T_x-1,T_y)}(I_m)$, $T_{(T_x,T_y)}(I_m)$ and $T_{(T_x+1,T_y)}(I_m)$. They are the control points which define the two third-order polynomials $Y_1(t) = a_1t^3 + b_1t^2 + c_1t + d_1$ and $Y_2(t) = a_2t^3 + b_2t^2 + c_2t + d_2$ (c.f. figure 4.1). To find eight unknowns, eight equations are needed:

$$\begin{cases} Y_1(0) = \rho_1 \\ Y_1(1) = \rho_2 \\ Y_2(0) = \rho_2 \\ Y_2(1) = \rho_3 \\ Y_1'(1) = Y_2'(0) \\ Y_1''(1) = Y_2''(0) \\ Y_1''(0) = 0 \\ Y_2''(1) = 0 \end{cases} \Longrightarrow \begin{cases} d_1 = \rho_1 \\ a_1 + b_1 + c_1 + d_1 = \rho_2 \\ d_2 = \rho_2 \\ a_2 + b_2 + c_2 + d_2 = \rho_3 \\ c_1 + 2b_1 + 3a_1 = c_2 \\ 6a_1 + 2b_1 = 2b_2 \\ 2b_1 = 0 \\ 6a_2 + 2b_2 = 0 \end{cases}$$

The first four lines of the system indicates that the curves must contain the control points, the two next lines forces the first and second derivative of the curves at the middle control point to be equal, and the two last lines are the constraints that make these cubic splines natural, i.e. having a null second derivative at the exterior control points. Solving this system of equation gives the curves' coefficients:

$$(a_1, b_1, c_1, d_1) = (\tfrac{\rho_1 - 2\rho_2 + \rho_3}{4}, 0, \tfrac{-5\rho_1 + 6\rho_2 - \rho_3}{4}, \rho_1)$$
$$(a_2, b_2, c_2, d_2) = (\tfrac{\rho_1 + 2\rho_2 - \rho_3}{4}, \tfrac{3\rho_1 - 6\rho_2 + 3\rho_3}{4}, \tfrac{-\rho_1 + \rho_3}{2}, \rho_2)$$

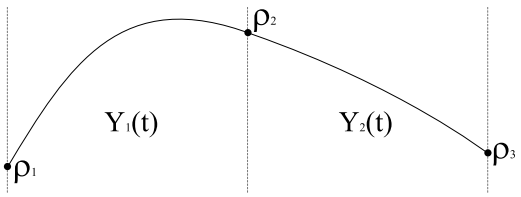

Figure 4.1.: The simplest practical case of the natural cubic splines interpolation: three control points define the two polynomials to find. If the centre control point is the maximum integral correlation coefficient found, then it does not correspond with the maximum of the interpolated function.

The coefficients define the curves but do not give the maximum on the interval $[T_x - 1; T_x + 1]$. To find it, it must first be determined on which side of T_x it lies. If the signs of the derivative of one polynomial at its control points does not change, then the maximum is on the other polynomial[2]. Once the correct polynomial Y has been chosen, the equation $Y'(t) = 0$ is solved, i.e. $t = \frac{-b_i \pm \sqrt{b_i^2 - 4a_ic_i}}{2a_i}$, and the maximum sought is the solution t_c that lies in the interval $[0; 1]$[3]. The best translation will thus equal either $T_x - 1 + t_c$ or $T_x + t_c$, depending on which side of T_x the maximum lies.

[2]The other polynomial's signs of the derivative at its control points change, therefore the intermediate value theorem states that there exists a point in between where the derivative equals 0, i.e. where the extremum lies

[3]By construction, there is a solution and the extrema in $[0; 1]$ cannot be a minimum.

4.1.3. Improvements

The computation of the cross-correlation uses a brute force approach, that is it calculates the correlation coefficient independently for every translation possible. This is very time-consuming, but can be accelerated using the cross-correlation theorem, which states, for $f(t)$ and $g(t)$ the functions to correlate : $(f \bigotimes g)(t) = FT[\overline{FT[f]}FT[g]](t)$, where $FT[]$ is the Fourier transform, and $\overline{x}$ is the conjugate of x.

Of course, this modification does not influence the result. On the other hand, the choice of the interpolation does: the result can gain in accuracy if a bigger neighbourhood was inspected, i.e. having more control points to define the spline. Another limitation is the fact that the correlation coefficient is a value between -1 and 1, but interpolations tend to ignore such constraints. For example, in our case, if $(\rho_1, \rho_2, \rho_3) = (-0.9, 0.9, 0.8)$, then the interpolation finds a maximum correlation coefficient of 1.19, which is an impossible value. Fortunately this is an extreme case which never occurs in practical examples, considering the resolution of the images compared to the geometric features of the sample.

5. Solution and results

5.1. Obtaining the transformation

The transformation that registers a model tomgoraphy to a reference uses the previous results, which consist of two items:

- $A_m = (Q_m, R_m)$ and $A_r = (Q_r, R_r)$ the axes of $\mathcal{T}_m$ and $\mathcal{T}_r$, defined by their extremities.
- $(T_x, T_y) \in \mathbb{R}^2$ the optimal translation that aligns $\mathcal{D}_m(\mathcal{P}_m, \mathcal{T}_m)$ with $\mathcal{D}_r(\mathcal{P}_r, \mathcal{T}_r)$, which are to be interpreted as a rotation around and a translation along $\mathcal{T}_m$'s axis.

5.1.1. Registering the axes

Let $\overrightarrow{V_m} = \dfrac{\overrightarrow{Q_m R_m}}{\|\overrightarrow{Q_m R_m}\|}$ and $\overrightarrow{V_r} = \dfrac{\overrightarrow{Q_r R_r}}{\|\overrightarrow{Q_r R_r}\|}$.
The first transformation to calculate is the rotation matrix $\mathfrak{R}_p$ such that $\mathfrak{R}_p(\overrightarrow{V_m}) = \overrightarrow{V_r}$[1]. If we define the unit vector $\overrightarrow{r}$ as the direction of the rotation axis, and θ the angle, then:

$$\overrightarrow{r} = \begin{pmatrix} r_x \\ r_y \\ r_z \end{pmatrix} = \frac{\overrightarrow{V_m} \wedge \overrightarrow{V_r}}{\|\overrightarrow{V_m} \wedge \overrightarrow{V_r}\|} \qquad \theta = \cos^{-1}\left(\frac{\overrightarrow{V_m}.\overrightarrow{V_r}}{\|\overrightarrow{V_m}\|.\|\overrightarrow{V_r}\|}\right)$$

$$\implies \mathfrak{R}_p = \begin{pmatrix} r_x^2 + \cos(\theta)(1 - r_x^2) & r_x r_y(1 - \cos(\theta)) + r_z \sin(\theta) & r_x r_z(1 - \cos(\theta)) - r_y \sin(\theta) \\ r_x r_y(1 - \cos(\theta)) - r_z \sin(\theta) & r_y^2 + \cos(\theta)(1 - r_y^2) & r_y r_z(1 - \cos(\theta)) + r_x \sin(\theta) \\ r_x r_z(1 - \cos(\theta)) + r_y \sin(\theta) & r_y r_z(1 - \cos(\theta)) - r_x \sin(\theta) & r_z^2 + \cos(\theta)(1 - r_z^2) \end{pmatrix} \quad (5.1)$$

Note that this transformation does not register the reference cylinders' axes, but makes them parallel (the rotation axis was only defined by its direction, not its position, so it was supposed that it passes through the origin of the coordinate system). To register A_m with A_r, the model is first translated by $-Q_m$ (noted $\mathfrak{T}_{-Q_m}$) so as to have $\mathcal{T}_m$'s axis pass through the origin, then rotated by $\mathfrak{R}_p$, and finally translated by Q_r (noted $\mathfrak{T}_{Q_r}$).

5.1.2. Axis-dependent transformations

The two following transformations, i.e. translation and rotation along and around the model's axis, can be applied in any order (as long as A_m passes through the system's origin, which means before $\mathfrak{T}_{Q_r}$). The angle of rotation θ_c around $\mathfrak{T}_{-Q_m} \circ \mathfrak{R}_p(\overrightarrow{V_m})$ is simply $2\pi\frac{T_x}{w}$ (w being the width in pixels of the projections). The subsequent rotation matrix $\mathfrak{R}_A$ can be developed

[1] For an overview of rotation matrices, consult [21]

in the same way as in 5.1. The translation $\mathfrak{T}_A$ along $\mathfrak{T}_{-Q_m} \circ \mathfrak{R}_p(\overrightarrow{V_m})$, in the same way, equals $\|\overrightarrow{Q_m R_m}\| \frac{T_y}{h_m}$.

5.1.3. Complete transformation

To summarise, the complete transformation is the following sequence of five operations: $\mathfrak{T}_{-Q_m} \circ \mathfrak{R}_p \circ \mathfrak{R}_A \circ \mathfrak{T}_A \circ \mathfrak{T}_{Q_r}$. The displacement of the model's axis during this transformation is illustrated in figure 5.1.

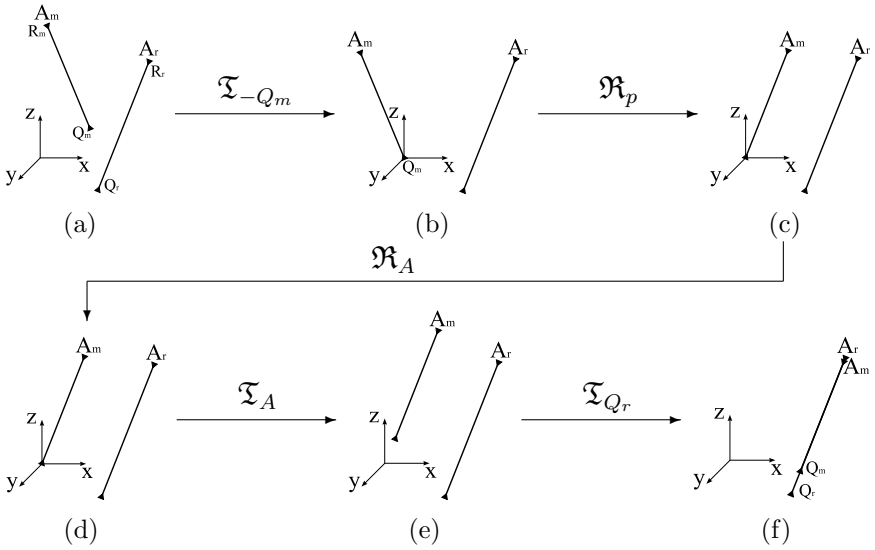

Figure 5.1.: Sequence of transformations of the registration: (a) shows the original positions of the axes A_m and A_r; the first operation centres the model's cylinder by placing Q_m on the system's origin to obtain (b), the second makes the axes parallel as in (c), then the rotation around A_m is applied, but this does not affect A_m, which is why (d) is identical to (c); afterwards, translation along A_m is applied, as shown in (e); the last translation completes the registration, shown in (f), as both axes overlap.

5.2. Results

5.2.1. Application

Once the tomographies have been registered, operations such as substractions between the images is possible, so as to view the portions that have evolved in the sample, as in figure 5.2. This also allows quantitative characterisation of the reaction, such as evolution of permeability, reactive surfaces, or pore network.

5.2.2. Implementation and timing

The different steps of the registration were developed using the C++ language, in separate modules (outer edge of an object, biggest connected component, alignment of 2D images, ...)

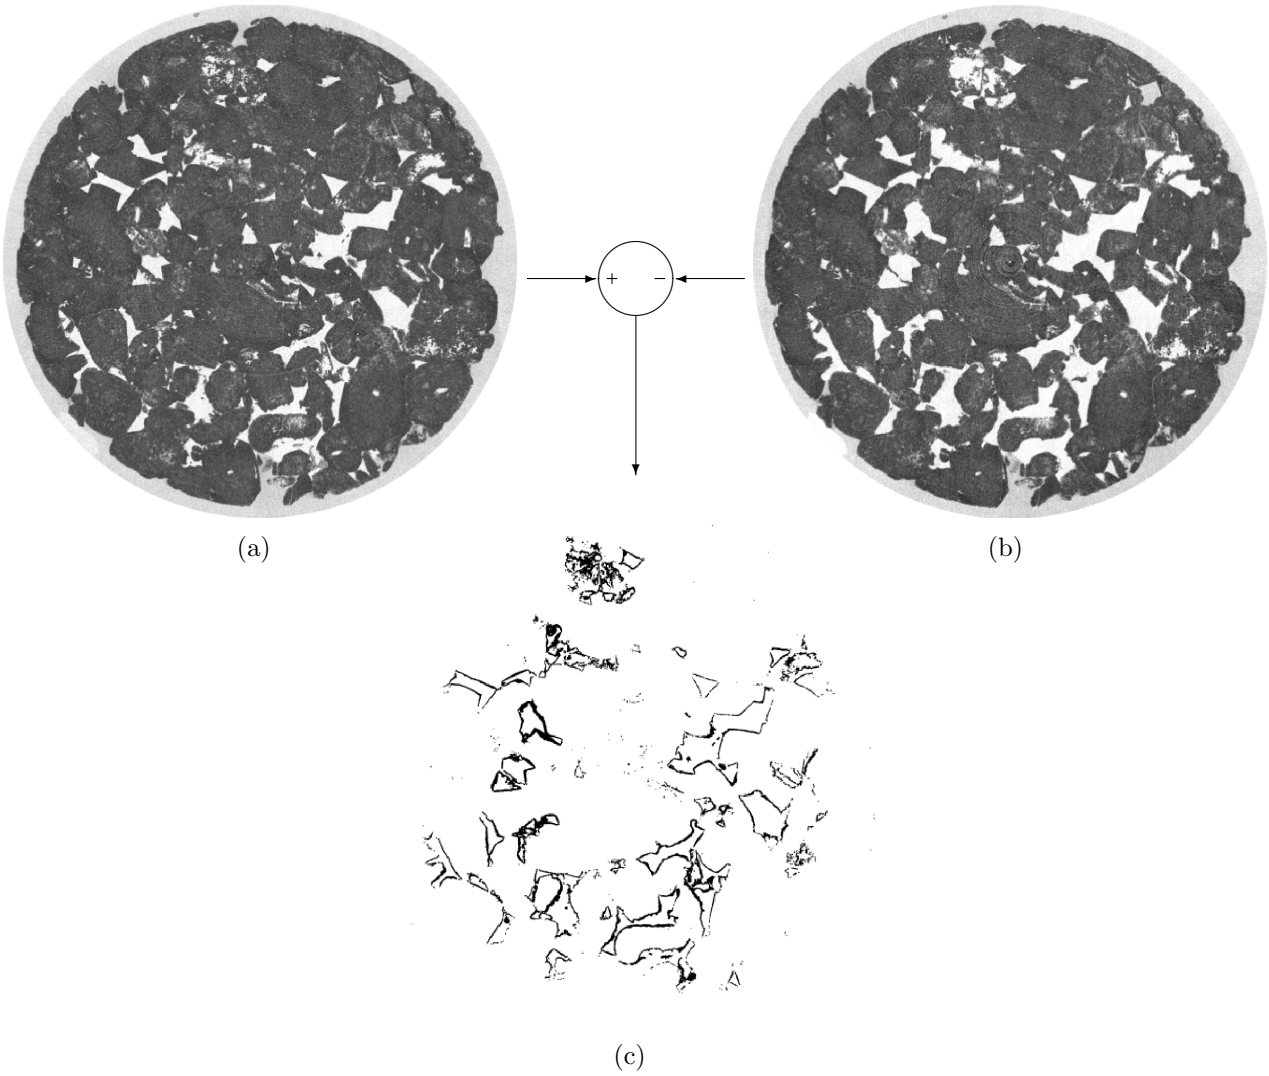

Figure 5.2.: Example of a registered sample: (b) is the same sample as (a) but after further percolation, and (c) is the difference between the two, highlighting the modified portions.

integrated in the visualisation program Amira from Mercury Computer Systems. The entire process was launched through an Amira script, which uses Tool Command Language. Tests were run on a PC with a Pentium 4 processor at 3.4 gigahertz, 2 gigabytes of DDR-RAM, and an Ultra ATA hard drive (on which the tomographies were stored). In average, registration of two tomographies takes about 45 minutes, counting between 17 and 20 minutes to extract $\mathcal{T}$ and $\mathcal{D}(\mathcal{P}, \mathcal{T})$ for each tomography. In those 20 minutes, each axial cross-section is processed in about 0.7 seconds, 22 % for mineral segmentation and finding its outer edge, 9 % for finding the biggest connected component and 27 % for finding the smallest surrounding circle. The remaining 58 % is taken by the interface with Amira and the Tcl interpreter.

5.2.3. Conclusion

The experiments of the PICOREF programme on the percolation of porous sedimentary minerals by CO_2-enriched water result in a series of micro-tomographies showing the sample at different stages of percolation, for the purpose of acquiring experimental measurements for fluid-rock interaction model validation. The first step, once the tomographies have been reconstructed, is to register these 3D images. The registration is simple in theory, as the sample is rigid (i.e. there is no scaling or shearing involved). In practice, due to the number of pixels in the tomography (several billion), serious attention must be paid to the computational complexity of the process.

Since the percolation alters the microstructure of the sample, feature-based registration is favoured. The most pertinent feature or information which remains invariant during the percolation and visible on the tomographies is the interior edge of the coat of resin which holds the mineral together. The registration is based on this feature and on the cylindrical shape of the sample, and can be summarised as follows:

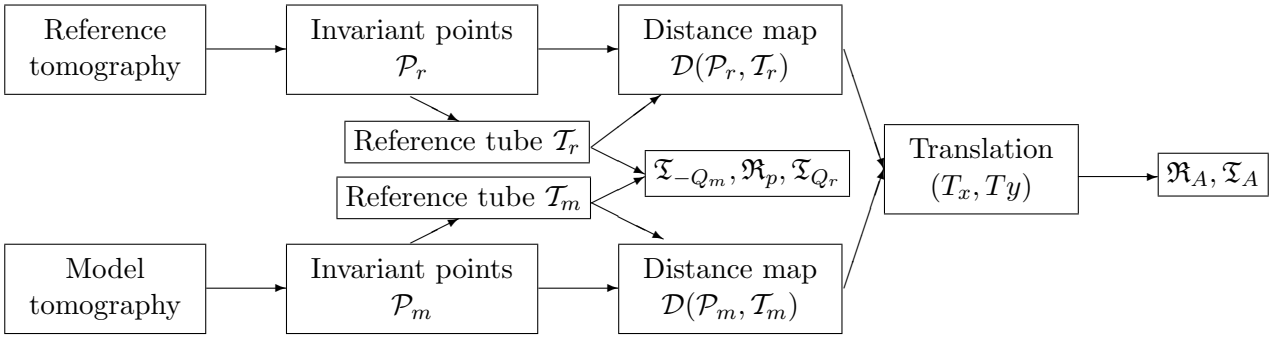

The final transformation is the sequence of rotations and translations $\mathfrak{T}_{-Q_m} \circ \mathfrak{R}_p \circ \mathfrak{R}_A \circ \mathfrak{T}_A \circ \mathfrak{T}_{Q_r}$.

Although this method seems to yield correct results, appreciation of its accuracy is, for the time being, only subjective. For a more impartial appreciation, an idea would be to calculate the covariance between the reference tomography and the registered model, vary slightly the rotational and translational parameters of the registration transformation individually, and verify that the new covariances are greater than the original value.

Finally, it is difficult to say which changes in the method can improve the quality of the result, but improvement of the computation's speed is definitely more tangible. Although algorithmic changes, such as using the fast Fourier transform with the cross-correlation theorem for the alignment of the distance maps (section 4.1.3), might help save a few minutes, the most significant acceleration would be achieved by avoiding the integration in Amira and the use of a scripting language, as they take up more than half the total computation time.

Bibliography

[1] Svante Arrhenius. On the influence of carbonic acid in the air upon the temperature of the ground. *Philosophical Magazine*, 41(5):237–276, 1896.

[2] Tuncel Yegulalp, Klaus Lackner, and Hans Ziock. A review of emerging technologies for sustainable use of coal for power generation. *International journal of surface mining reclamation and environment*, 15(1):52–68, 2001.

[3] Zaihua Liu and Daoxian Yuan. Contribution of carbonate rock weathering to the atmospheric CO_2 sink. *Environmental Geology*, 39:1053–1058, 2000.

[4] Hans Rogner. An assessment of world hydrocarbon resources. *Annual Review Energy Environment*, 22:217–262, 1997.

[5] Klaus Lackner. A guide to CO_2 sequestration. *Science*, 300:1677–1678, 2003.

[6] Stephen Pacala. *Could carbon sequestration in leaky reservoirs solve the greenhouse problem ?* Carbon Mitigation Initiative, Princeton University, 2001.

[7] Elizabeth Wilson and David Keith. Geological carbon storage: understanding the rules of the underground. In J. Gale and Y. Kaya, editors, *Proceedings of the 6th Greenhouse Gas Control Conference, Kyoto Japan*, pages 229–234, Oxford, UK, 2003. Pergamon.

[8] Étienne Brosse and. Modelling fluid-rock interaction induced by the percolation of co_2 enriched solutions in core samples: the role of reactive surface. *Oil & Gas science and technology*, 60(2):287–305, 2005.

[9] Catherine Noiriel, Dominique Bernard, Philippe Gouze, and Xavier Thibault. Hydraulic properties and microgeometry evolution accompanying limestone dissolution by acidic water. *Oil & Gas science and technology*, 60(1):177–192, 2005.

[10] Lisa Brown. A survey of image registration techniques. *ACM Computing Surveys*, 24(4): 325–376, 1992.

[11] Twan Maintz and Max Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, April 1998.

[12] Donald Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison-Wesley, Reading, Massachusetts, second edition, 10 January 1973.

[13] Tat Kong and Azriel Rosenfeld. Digital topology: introduction and survey. *Comput. Vision Graph. Image Process.*, 48(3):357–393, 1989.

[14] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[15] Gilles Bertrand, Jean-Christophe Everat, and Michel Couprie. Image segmentation through operators based upon topology. *Journal of Electronic Imaging*, 6(4):395–405, 1997.

[16] James Sylvester. A question in the geometry of situation. *Quarterly Journal of Mathematics*, 1:79, 1857.

[17] Jack Elzinga and Donald Hearn. Geometrical solutions for some minimax location problems. *Transportation Science*, 6:379–394, 1972.

[18] Zvi Drezner and Saharon Shelah. On the complexity of the Elzinga-Hearn algorithm for the 1-center problem. *Math. Oper. Res.*, 12(2):255–261, 1987.

[19] Nimrod Megiddo. Linear-time algorithms for linear programming in $\mathbb{R}^3$ and related problems. *SIAM Journal on Computing*, 12(4):759–776, November 1983.

[20] Charles Corbett. The unweighted center problem on the Euclidean plane: Two algorithms. Working paper. Econometrisch Instituut, Erasmus Universiteit Rotterdam, Netherlands, 1989.

[21] Eric Weisstein. Rotation matrix. `http://mathworld.wolfram.com/RotationMatrix.html`.

[22] Johann Radon. Über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten. *Mathematische-Physische Klasse 69*, pages 262–267, 1917. Berichte Sächsische Akadamie der Wissenschaften.

[23] Avinah Kak and Malcom Slaney. *Principles of computerized tomographic imaging*. IEEE Press, New York, N.Y., 1988.

[24] Stanley Deans. *The Radon transform and some of its applications*. Wiley and sons, New York, N.Y., 1983.

[25] José Baruchel, Jean-Yves Buffière, Eric Maire, Paul Merle, and Gilles Peix. *X-Ray Tomography in Material Science*. Hermes Sciences Publications, 2000. ISBN 2-7462-0115-1.

A. Principles of parallel-beam X-ray tomography

Tomography is the non-invasive process of generating images of cross-sections from a series of transmission or reflection data acquired by illuminating the object from different directions. In the case of X-ray tomography, for each direction of illumination, the process is similar to a radiography: the acquired transmission data gives an attenuation map of the rays. Figure A.1 illustrates this: the centre object, consisting of an ellipse and a circle, is illuminated from two directions, the acquired images are called projections, which are used to reconstruct cross-sections of the object. In reality, many more projections are needed for an accurate reconstruction (it depends of course on the number of pixels in the reconstructed image).

The problem of tomographic imaging is to find a way to reconstruct the cross-sections from the projections, problem which has been solved by Radon in 1917 [22], giving his name to the projection transformation as well as demonstrating the reversibility of the process.

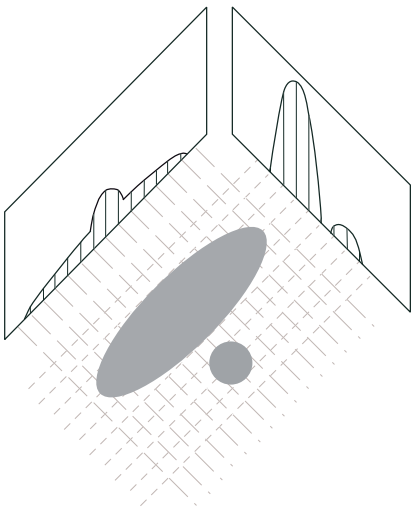

Figure A.1.: Illustration of the principle for parallel beam X-ray tomography

The theory of X-ray tomographic imaging was first applied for clinical purposes with the invention of the X-ray computed tomographic scanner by Godfrey Hounsfield, for which he received a Nobel Prize in 1972, shared with Allan Cormack who found some of the reconstruction algorithms. Hounsfield used algebraic techniques for the reconstruction, but another method, called convolution-backprojection, developed by Ramachandran and Lakshminarayanan in 1971 and later by Shepp and Logan in 1974, has a better accuracy and shorter computation time [23]. For more details on this subject, a complete bibliography is compiled in [24].

Tomography has numerous uses in the medical field, but it also has promising industrial applications (defect detection, local characterisation). Progress in this field now allows for enhanced spatial resolution (under the micron, thus the prefixe micro in micro-tomography) with better image quality, especially with the use of a modern synchrotron radiation source with a monochromatic beam [25]. It is therefore possible to study the microstructure of objects, and furthermore, since tomography is a non-destructive technique, their evolution.

B. Useful geometric properties

These properties are helpful when implementing the Elzinga-Hearn algorithm presented in 3.2.2, they present certain ways of obtaining the necessary geometric characterisations, and can prove crucial for two reasons: these properties can be verified very rapidly, which is important when large amounts of data need to be processed, and especially when this specific part of the registration (the 1-centre problem) is one of the most time-consuming; and secondly, they avoid operations that can be inaccurate in certain situations, e.g. using the standard inverse sine function for angles very close to $\frac{\pi}{2}$.

Let ABC be a triangle as shown in the following diagram:

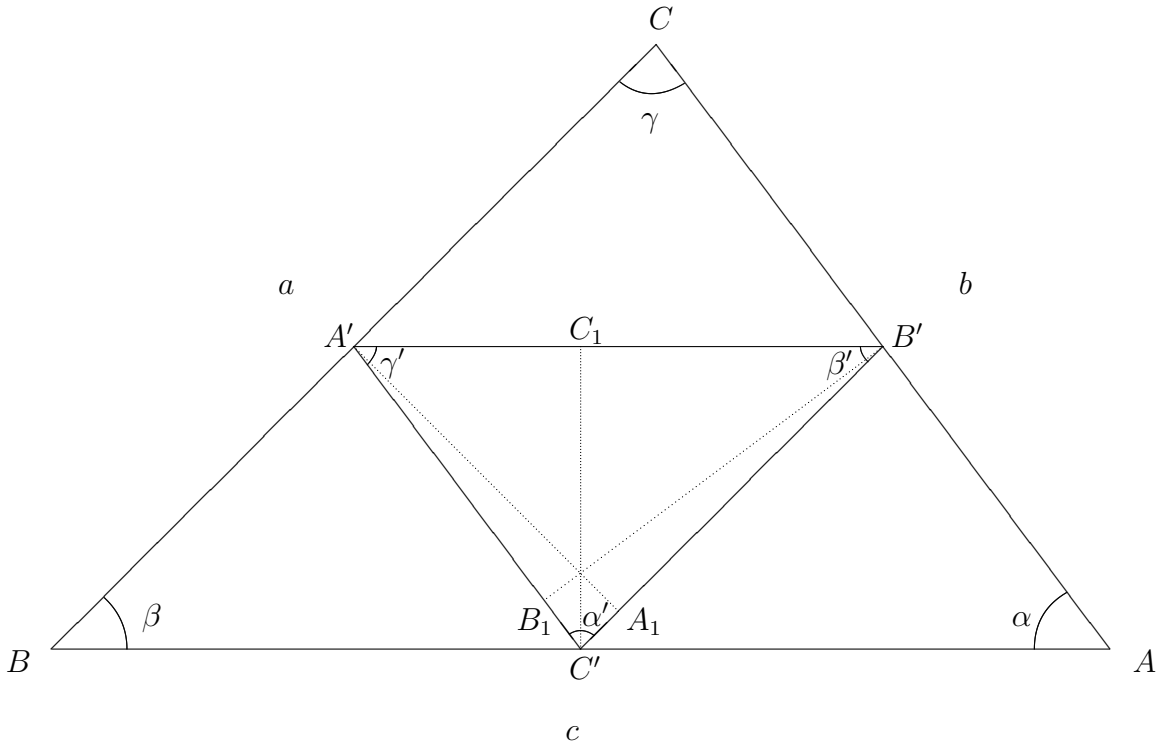

where a, b and c are the lengths of the edges, α, β and γ are the angles, and the vertices A, B and C are respectively defined in $\mathbb{R}^2$ as $\binom{A_x}{A_y}$, $\binom{B_x}{B_y}$ and $\binom{C_x}{C_y}$. We note A', B' and C' the centres of the triangle's edges, α', β' and γ' the angles of the triangle $A'B'C'$, and A_1, B_1 and C_1 the altitudes of this triangle. We also suppose $c \geqslant a$ and $c \geqslant b$.

B.1. Acuteness

$$\text{Law of cosines} \iff c^2 = a^2 + b^2 - 2ab\cos\gamma \tag{B.1a}$$

$$(B.1a) \iff \left(0 < \gamma < \frac{\Pi}{2} \Leftrightarrow 2ab\cos\gamma > 0\right) \tag{B.1b}$$

$$(B.1b) \iff \left(c^2 < a^2 + b^2 \Leftrightarrow \text{the triangle } ABC \text{ is acute}\right)\blacksquare$$

This allows, in step 3, to very effectively test whether a triangle is acute or not, since a^2, b^2 and c^2 are easily computed ($a^2 = A_x^2 + A_y^2$).

B.2. Circumcentre

Property 1
The perpendicular bisectors of the triangle ABC are the altitudes of the triangle $A'B'C'$.

Demonstration
Let Δ_1 be the perpendicular bisector of $[BC]$.

$$B' \text{ and } C' \text{ centres of } [AC] \text{ and}[AB] \implies \frac{AB'}{AC} = \frac{AC'}{AB} \tag{B.2a}$$

$$(B.2a) \iff (B'C') /\!\!/ (BC) \tag{B.2b}$$

$$(B.2b) \implies \Delta_1 \perp (B'C') \tag{B.2c}$$

$$(B.2c) \wedge (A' \in \Delta_1) \iff \Delta_1 \text{ perpendicular bisector of } A'B'C'\blacksquare$$

The same applies to the two others perpendicular bisectors of ABC.

Property 2
The orthocentre of the triangle $A'B'C'$ is the barycentre of the system
$\{(A', \tan\alpha'), (B', \tan\beta'), (C', \tan\gamma')\}$

Demonstration
Let H be the barycentre of the system
$\{(A', \tan\alpha'), (B', \tan\beta'), (C', \tan\gamma')\}$.

$$A_1 \in [B'C'] \implies \frac{1}{A_1B'}\overrightarrow{A_1B'} + \frac{1}{A_1B'}\overrightarrow{A_1B'} + \frac{1}{A_1B'}\overrightarrow{A_1B'} = \overrightarrow{0} \tag{B.3a}$$

$$\left.\begin{array}{r}(B.3a)\\(\tan\beta' = \frac{A'A_1}{A_1B'})\\(\tan\gamma' = \frac{A'A_1}{A_1C'})\end{array}\right\} \implies \tan\beta'\overrightarrow{A_1B'} + \tan\gamma'\overrightarrow{A_1B'} = \overrightarrow{0} \tag{B.3b}$$

$$(B.3b) \implies A_1 \text{ barycentre of } \{(B', \tan\beta'), (C', \tan\gamma')\} \tag{B.3c}$$

$$\textit{idem for } B_1 \implies B_1 \text{ barycentre of } \{(A', \tan\alpha'), (C', \tan\gamma')\} \tag{B.3d}$$

$$\left.\begin{array}{r}(\text{barycentre associativity})\\(B.3c)\end{array}\right\} \implies H \text{ barycentre of } \{(A', \tan\alpha'), (A_1, \tan\beta' + \tan\gamma')\} \tag{B.3e}$$

$$\left.\begin{array}{r}(\text{barycentre associativity})\\(B.3d)\end{array}\right\} \implies H \text{ barycentre of } \{(B', \tan\beta'), (B_1, \tan\alpha' + \tan\gamma')\} \tag{B.3f}$$

$$(B.3e) \wedge (B.3f) \implies (H \in (A'A_1)) \wedge (H \in (B'B_1)) \tag{B.3g}$$

$$(B.3g) \implies H \text{ orthocentre of } A'B'C' \blacksquare$$

Property 3
$(\alpha' = \alpha) \wedge (\beta' = \beta) \wedge (\gamma' = \gamma)$
Demonstration

$$(B.2b) \wedge (\frac{AB'}{AC} = \frac{1}{2}) \wedge (\frac{AB'}{AC} = \frac{B'C'}{BC}) \implies B'C' = \frac{BC}{2} \tag{B.4}$$

$$\textit{idem for } A'C' \text{ and } A'B' \implies (A'C' = \frac{AC}{2}) \wedge (A'B' = \frac{AB}{2}) \tag{B.5}$$

$$(B.4) \wedge (B.5) \implies A'B'C' \text{ is a homothety of } ABC \tag{B.6}$$

$$(B.6) \wedge (\text{scaling conserves angles}) \implies (\alpha' = \alpha) \wedge (\beta' = \beta) \wedge (\gamma' = \gamma) \blacksquare$$

Properties 1,2 and 3 imply that the circumcentre of ABC is the barycentre of the system $\{(A', \tan\alpha), (B', \tan\beta), (C', \tan\gamma)\}$

B.3. Tangent

For the determination of the circumcentre as described previously, the tangent of angles are needed. The straightforward method would be first to find the angle, probably using the identity $\widehat{ABC} = \cos^{-1}\frac{\overrightarrow{BA}.\overrightarrow{BC}}{\|\overrightarrow{BA}\|.\|\overrightarrow{BC}\|}$, and then to use the tangent function provided by the system. For angles very close to $\frac{\pi}{2}$, its cosine is very close to 0, a slight imprecision in the inverse cosine function can corrupt the resulting angle by a few degrees. The sensibility of the tangent function is also important with angles approaching $\frac{\pi}{2}$, bringing further imprecision. In the case of the Elzinga-Hearn algorithm, the circumcentre's inaccurate positionning can lead to considering points theroetically inside the circumscribing circle to be outside, therefore selected for the next steps, possibly leading to an infinite loop. Instead, the formula $\tan = \frac{\sin}{\cos}$ cleverly avoids calculating the value of the angle, which gives $\tan(\widehat{ABC}) = \frac{\sqrt{1-\left(\frac{\overrightarrow{BA}.\overrightarrow{BC}}{\|\overrightarrow{BA}\|.\|\overrightarrow{BC}\|}\right)^2}}{\frac{\overrightarrow{BA}.\overrightarrow{BC}}{\|\overrightarrow{BA}\|.\|\overrightarrow{BC}\|}}$.