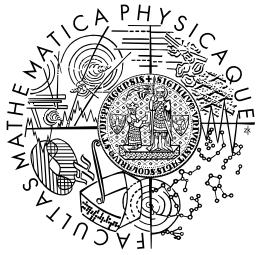

FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

MASTER THESIS

Bc. Adam Harmanec

Registration of images of nuclear fuel assembly

Department of Software and Computer Science Education

Supervisor of the master thesis: RNDr. Jan Blažek, Ph.D.
Study programme: Computer Science
Study branch: Artificial Intelligence

Prague 2021

I declare that I carried out this master thesis independently, and only with the cited sources, literature and other professional sources. It has not been used to obtain another or the same degree.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In ..... date ....

Author's signature

First and foremost, I have to thank my supervisor Jan Blažek for his guidance, patience and for getting me back on track, whenever I delved into irrelevant details.

I also want to thank Marcin Kopeć and Ondřej Pasta, as well as the Centrum Výzkumu Řež for their help with the experiments and for the opportunity to publish our data set and findings.

My acknowledgements also go to the authors of software that was used in the making of this thesis as well as to the authors of various literature cited in this thesis. It was a big help and inspiration.

Last but not least, I would like to thank my girlfriend and my family for their unquenchable love and support.

Title: Registration of images of nuclear fuel assembly

Author: Bc. Adam Harmanec

Department: Department of Software and Computer Science Education

Supervisor: RNDr. Jan Blažek, Ph.D., Department of Software and Computer Science Education

Abstract: Nuclear fuel is visually inspected during regular shutdowns in order to monitor defects and long-term changes. To enable automatic comparison of images of fuel assemblies, it is crucial to perform their registration, the implementation of which has not yet been published in the scientific literature. In this work we present an analysis of image registration techniques and similarity metrics inspired by the focus operators used in autofocus and shape-from-focus. Their performance has been evaluated using a series of experiments that tested their various properties on a novel data set obtained in cooperation with the research organization Centrum výzkumu Řež. Finally, we present and discuss the results and make recommendations on which to use in which scenario.

Keywords: nuclear fuel, image registration, image processing

Contents

Introduction **3**

1 Problem Analysis **4**
1.1 Fuel Assembly . . . 4
1.2 Image Registration . . . 4
1.3 Fuel Assembly Registration . . . 6
1.4 Current method . . . 7
1.5 Registration as an Optimization Problem . . . 7

2 Similarity metrics **8**
2.1 Criteria . . . 8
2.2 Selection of metrics . . . 9
2.2.1 Gradient-based . . . 9
2.2.2 Laplacian-based . . . 10
2.2.3 Statistics-based . . . 10
2.2.4 Wavelet-based . . . 11
2.2.5 DCT-based . . . 12
2.2.6 Miscellaneous . . . 12
2.3 Implementation . . . 13

3 Experiments **14**
3.1 Data set . . . 14
3.2 Setup . . . 14

4 Results **17**
4.1 Baseline . . . 17
4.2 Noise . . . 20
4.2.1 Additive white Gaussian noise . . . 21
4.2.2 Radial distortion . . . 23
4.2.3 Rotation . . . 26

5 Discussion **29**
5.1 Vertical shift . . . 29
5.2 Horizontal shift . . . 29
5.3 Rotation . . . 30
5.4 Findings . . . 31
5.5 Summary . . . 32

Conclusion **34**

Bibliography **37**

List of Figures **38**

List of Tables **39**

List of Abbreviations **40**

A Attachments **41**

Introduction

Nuclear power plants go through regular temporary shutdowns, during which, among other things, a visual inspection of fuel takes place. This is necessary to meet the design criteria of the fuel, which ensures their safe operation. Failure to detect a serious defect can lead to an unplanned shutdown of the reactor and could cause great financial damage as well as the release of radioactivity. During the inspection, mainly nuclear fuel assemblies are examined in order to assess their condition. In a light-water reactor, the fuel assembly (FA) is a cluster of long cylindrical fuel rods held together in a specific shape by spacer grids. They are scanned underwater using a camera remotely controlled by an operator. Previously, due to the complexity of the tests, these videos were mostly hand-checked by experts in a short time window. This is obviously very inefficient and error-prone.

The first step towards automation of the process has been done by Knotek[1] in 2020, who developed a procedure for the stitching of video frames into a one-image-overview (OIO) of a fuel assembly. This speeded up and streamlined the work of the inspectors. This approach, however, is lacking in image registration. In this work, we will focus on the possibility of automating the accurate comparison of these images in time, a task that would have previously been extremely time consuming to do by hand.

The implementation of the registration of images of nuclear fuel assemblies would make it possible to quickly and automatically compare OIOs obtained at different times with an accuracy down to the level of a few pixels. That would make it possible to examine which parts of the FA or positions in the reactor are more prone to defects and to track their evolution over time. In addition, the ability to register two images of the same FA from the same video would allow them to be more accurately stitched together and unlock the possibility to measure other useful properties.

Our goal is to solve the OIO registration and as the first steps (solved in this thesis) we identify:

1. Analyze the properties of images of nuclear fuel assemblies and how that can affect their registration.

2. Design and validate a framework for the evaluation of the quality of their registration.

1. Problem Analysis

1.1 Fuel Assembly

In a light-water reactor, the nuclear fuel assembly (FA) is a bundle of a few hundred fuel rods held together in a predetermined shape by multiple spacer grids. A fuel rod is a long thin tube containing the actual nuclear fuel. A spacer grid is a divider that holds the fuel rods at a specific distance from each other and prevents them from moving and bending. The FA is approximately 4 metres tall with the spacer grids placed dozens of centimeters apart. The rods are about 1 centimeter in diameter and the grids are less than 5 centimeters tall. The FA can be made of different metals and alloys but it generally is visually similar to stainless steel when new, but when used in a reactor it is subject to oxidation and surface defects, becoming darker, less reflective and getting covered with a matte map. In addition, during their performance in the reactor, debris may get stuck on the spacer grids and the rods may start to bend.

The FA is recorded underwater using a camera controlled by an operator, while moving along it with a point light. The acquired images are then blended into one large image using the method developed by Knotek[1], which is discussed in more detail in Section 1.4. This includes some pre-processing steps that make sure the image is verticalized, that is, that the rods are mostly perpendicular to the ground. This OIO is very useful for nuclear experts, who can quickly analyze it and obtain various useful information about the state of the FA.

We carried out our experiments in cooperation with the Research Centre ŘEŽ (CVR). The power plant environment was emulated as closely as possible with some notable differences. We used a shortened VVER-1000 FA mock-up that is only 125 centimeters tall and has the fuel rods arranged in a hexagonal shape. For some experiments we used a tube light instead of a point light to increase the general brightness and detail of the images. The camera moved independently along a preset path.

In the final image of one side of the assembly, viewed perpendicularly from the front, there are 7 to 11 visible vertical rods and several horizontal grids. A second row of rods is visible in between the rods in the front row. It is generally darker but can still create reflections. The spacer grid is mostly flat with some markings and inscriptions and slopes towards the back at the top and bottom edges forming teeth-like shapes in between the rods. The usual image size is around 1400 pixels in height and 600 pixels in width. An example image can be seen in Figure 1.1.

1.2 Image Registration

Image registration is a problem of aligning two or more images of the same object or scene. The images are usually taken at a different time, from a different angle, or by a different sensor, which causes them to differ. The usual approach is to choose a common coordinate space for both images and to search for a transformation in that space that aligns their points as closely as possible. Commonly considered transformations include rotation and shift, the so-called rigid

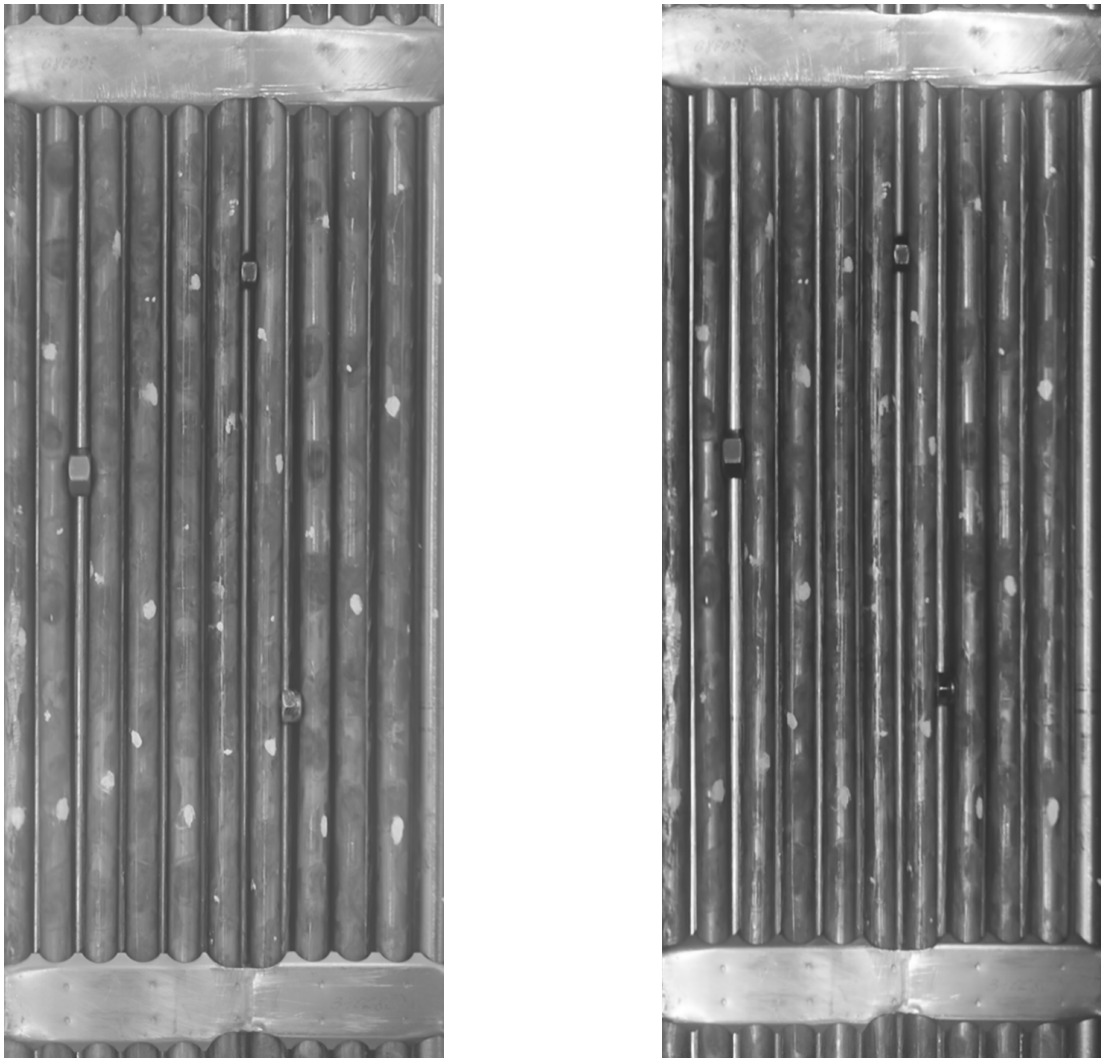

(a) Tube lighting (b) Point lighting

Figure 1.1: Images of FAs with lighting conditions

transformation, but also affine, projective or even more complex transformations. The classic dilemma with choosing the proper transformation comes from balancing between minimizing the number of unknown parameters while preserving accuracy and staying as close to reality as possible.

Registration methods depend on the nature of the considered images but are usually classified as either area-based or feature-based[2]. One of the area-based methods often used in medical applications is based on selecting an ad-hoc similarity metric and using an optimization algorithm to estimate the optimal transformation by minimizing that metric[3]. Multiple great overviews and surveys have been published describing the contemporary methods and approaches in greater detail[2, 4, 5].

Due to the sensitivity of the data, not many articles have been published on applied nuclear fuel image processing, albeit with a few exceptions[6, 7]. Moreover, to our knowledge, there is no published research in the field of registration of images of nuclear fuel.

1.3 Fuel Assembly Registration

In our case, the registration task has it's specific constraints and difficulties. The goal is to register two images of FA, which can be the final OIO, parts of it, or just two raw frames from a video obtained as described in section 1.1.

Such images don't differ in rotation significantly - for OIOs, the difference isn't greater than a few degrees and in the video frames, the camera is sometimes tilted and may rotate slightly during the recording but, according to the data we gathered, the difference doesn't exceed $\pm 7.8°$ from the ideal position. Additionally, the scale of the most images is very similar because the screening configuration does not change very much. The difference is mostly less than one percent.

The main disparity is caused by shift, with the vertical component being more significant than the horizontal one. The variance can amount to a few tens of pixels in both directions.

Images taken with modern cameras are affected by radial distortion. That means that the sides of the video frames are deformed and don't align perfectly with the previous and subsequent frames making the sides of the final OIOs blurrier than the center.

More variations that need to be taken into account but not eliminated during registration are different lighting conditions, rod bending and often different surface defects, which can develop over time. The latter can be especially significant when registering an image of a new shiny FA and of an older more oxidized one.

Images of nuclear FAs have some unique properties that a registration algorithm can take advantage of. The spacer grids form a large flat area with very uniform intensity. It's bottom and especially top segment is well defined by a noticeable change in brightness. Their teeth along the upper and lower edges are visually separated from the fuel rods underneath, but the transition is quite blurry. The fuel rods also have a similar brightness over their whole surface, but this is significantly thrown off by surface defects and, in addition, due to their round shape, their brightness decreases towards the edges where it fades to dark. The second row of rods is much darker but produces a thin and bright reflection stretching vertically along their entire length. Excessive local oxidation, a type of surface defect, and debris trapped at the edges of the spacer grids also form areas of uniform intensity and very prominent edges.

Typically, the horizontal cross-section of intensities is almost constant in the grid areas, while at the rod areas it resembles a periodic function with wider round peaks corresponding to the rods and is some cases with tall narrow peaks in between where a reflection of the second row of rods is. The vertical cross-section is then mostly constant with high-intensity plateaus in the area of the grids.

These properties depend significantly on the type of lighting that was used in their acquisition. Images with tube lighting are less contrasting, more detailed, and the reflections are less pronounced. On the contrary, images with point lighting have a higher contrast and more intensive reflections. This is well demonstrated in Figure 1.1 on the example of the spacer grid teeth, which are clearly visible under tube lighting, but which are almost imperceptible under a point light.

1.4 Current method

The method developed by Knotek[1] uses a form of registration to properly align frames before blending them into a single OIO. This is done in two steps. First, the rotation is determined by finding lines in the central diamond cutout of each frame using the Canny edge detector and a subsequent Hough transform. A median of the line angles is used. Second, after the images are correctly verticalized, their horizontal shift is selected as the maximum of convolution of the overlapping bands. The vertical shift is constant and is calculated from a manually inserted vertical camera speed parameter.

This approach works well for creating OIOs that are visually examined, but has a limited accuracy, results in blurry edges and cannot be directly used to register two different OIOs. In addition, it's quite slow and relies on some assumptions that might not always be valid. Most notably, the same angle and vertical shift is used for all frames. Instead, we are looking for a more general solution with up to sub-pixel accuracy that could be applied to more use cases and that can handle more complex transformations.

1.5 Registration as an Optimization Problem

Because there aren't any obvious distinctive features that would be stable in time and invariant to the aforementioned transformations and that could be used in feature-based methods, we instead opted for using area-based methods. These generally work with a similarity (or a dissimilarity) metric that is maximized (or minimized). This is a multidimensional optimization problem, where the variables are the parameters of the expected transformation.

Selecting a similarity metric is not trivial as its properties dictate the optimization algorithms that can be used. We want to be able to find the maximum of the metric efficiently and therefore its values must increase monotonically towards its maximum, which in turn should be as close as possible to the real optimum. A more thorough analysis is presented in the next chapter.

2. Similarity metrics

We are looking for a function that evaluates the transformation parameters and quantifies the similarity of the registered images. By doing that, it implicitly defines the optimum transformation, which, if the function has suitable properties, can be found iteratively and effectively[8]. Moreover, because there is no ground truth in a real world application, if there is a similarity metric that can faithfully measure the quality of registration, then any registration algorithm that does not minimize that metric will always produce worse results.

Such a function can be designed in a number of ways, but we decided to investigate functions that can produce a global estimate of image quality, in our case of a blend of two transformed images. Given a 2-D gray-scale blended image, they should return a single real value assessing the accuracy of the fit of the two blended images. An extra benefit of such functions is that they can be used to evaluate the original images directly and to quantify the relative decrease in quality compared to the blended image.

To our knowledge, no one has yet attempted to apply such functions to the evaluation of the quality of blended images of nuclear fuel or similar objects. Because of that, we chose to look to other fields for inspiration. Intuitively, a sub-optimal registration fit will introduce blur and reduce the sharpness of the result. This is somewhat similar to finding the right focus distance in a digital camera or a microscope where an incorrect focus distance results in a blurry image. In addition, for both problems, solutions that are close to the optimum produce a lower amount of blur compared to solutions that are further away.

The so-called focus measures or operators have been studied thoroughly for use in shape-from-focus and autofocus problems. We can therefore build on the rich research in the aforementioned fields and try to apply these metrics to our use case.

2.1 Criteria

For each metric we study the following criteria:

Accuracy measures the *true* error, i.e. how far the maximum of the metric is from the real optimum. In practice it is very difficult to measure as it requires a ground truth or a gold standard as a baseline.

Precision shows the granularity at which the metric is able to operate, both in terms of the range of values around the maximum but also the step size. In the context of shift, this would be quantified in pixels or their fractions and for rotation it would be degrees or minutes.

Robustness is a property of metrics that behave similarly for similar inputs. In other words, if we alter the input images slightly, the optimum should still be roughly the same.

Reliability means that a metric should behave predictably for all reasonable inputs. There shouldn't be a credible input for which the metric completely fails.

Algorithm complexity and derived **computation time** should be reasonably low so that it can be used in real time or on a large number of samples.

2.2 Selection of metrics

Well over a few dozen focus measures or operators have been mentioned in the scientific literature. To select a smaller number of relevant ones, we performed several preliminary experiments and excluded metrics that worked poorly or that were very computationally expensive. Some functions formed natural groups that exhibited very similar behavior. In such cases, we picked the most general and/or the most promising representative from each.

There are several more sophisticated measures based on the way humans perceive blur[9, 10, 11] or methods based on deep convolutional neural networks[12] but we decided against using them because they combine other simpler measures, use non-generalizable heuristics or are uninterpretable. We prefer the ones that use the specific properties of the images that we study.

In the following list we use the 6 categories proposed by Petruz et al.[13]. The gray level intensity value of a pixel with coordinates $[i, j]$ is denoted as $I(i, j)$. The width and height of the image are marked as $W \times H$.

2.2.1 Gradient-based

Metrics in this group use the first derivative of the image to measure the amount and sharpness of edges. The assumption is that blurrier images have less prominent and a lower number of intensity peaks.

1. **Absolute gradient**[14]

$$\phi = \sum_i \sum_j \max(|I_x(i,j)|, |I_y(i,j)|) \tag{2.1}$$

   where I_x and I_y are the first horizontal and vertical derivatives of the image respectively.

   In the literature, a threshold is sometimes used to only select some values that are higher. We decided to omit it, because its correct selection is non-trivial and poorly generalizable. In addition, the original implementation by Jarvis[14] took into account only horizontal derivatives but we took inspiration from the later work of Santos et al.[15] to select the maximum of derivatives from both directions.

2. **Tenengrad**[16, 17]

$$\phi = \sum_i \sum_j (I * S_x)^2(i,j) + (I * S_y)^2(i,j), \tag{2.2}$$
$$S_x = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix}, \quad S_y = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

where $*$ denotes a 2-D convolution with the Sobel operators.

As in the previous case, a threshold was originally used, but for the same reasons we did not use it here either.

2.2.2 Laplacian-based

Similarly to the previous group, metrics in this one measure the amount of edges in an image. Laplacian, the second derivative operator, works as a high spatial frequency pass filter.

3. **Variance of Laplacian**[18]

$$\phi = \frac{1}{W \times H} \sum_i \sum_j (L(i,j) - \overline{L})^2, \tag{2.3}$$
$$L = I * \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

where L is the Laplacian of the image obtained by convolution with the Laplacian mask and $\overline{L}$ is it's mean value.

4. **Modified Laplacian**[19]

$$\phi = \sum_i \sum_j L_m(i,j), \tag{2.4}$$
$$L_m = \left| I * \begin{bmatrix} -1 & 2 & -1 \end{bmatrix} \right| + \left| I * \begin{bmatrix} -1 & 2 & -1 \end{bmatrix}^T \right|$$

where $*$ denotes a 2-D convolution. An alternative definition of the Laplacian L_m is used to ensure that the horizontal and vertical gradients don't cancel each other out if they have opposite signs.

The first implementation

2.2.3 Statistics-based

Global image statistics such as contrast or the image histogram can also be used as texture descriptors to measure the sharpness of the image.

5. **Normalized gray-level variance**[15]

$$\phi = \frac{1}{H \times W \times \overline{I}} \sum_i \sum_j (I(i,j) - \overline{I})^2 \tag{2.5}$$

where $\overline{I}$ is the mean gray level intensity. The normalization in addition to the original gray-level variance[20] is done to compensate for the difference in average intensity among different images.

6. **Mason and Green's histogram method**[21, 20]

$$\phi = \sum_{k>T} H_k(k - T) \tag{2.6}$$

where H_k denotes the number of pixels with intensity k in the image histogram and T is a threshold calculated as

$$\delta(i,j) = 2\Big(I(i,j-1) - I(i,j+1)\Big)^2 + 2\Big(I(i-1,j) - I(i+1,j)\Big)^2 + \Big(I(i-1,j-1) - I(i+1,j+1)\Big)^2 + \Big(I(i-1,j+1) - I(i+1,j-1)\Big)^2,$$

$$T = \frac{\sum_i \sum_j \delta(i,j) I(i,j)}{\sum_i \sum_j \delta(i,j)}$$

2.2.4 Wavelet-based

Coefficients of the discrete wavelet transform describe the frequency and spatial content of an image. Their statistical properties in the high frequency subbands can be leveraged to measure sharpness. It should be noted that metrics in this group are significantly more computationally intensive.

7. **Variance of wavelet coefficients**[22]

$$\phi = \frac{1}{W \times H}\Big[\sum_i \sum_j (W_{LH1}(i,j) - \overline{W_{LH1}})^2 + \sum_i \sum_j (W_{HL1}(i,j) - \overline{W_{HL1}})^2 + \sum_i \sum_j (W_{HH1}(i,j) - \overline{W_{HH1}})^2\Big] \tag{2.7}$$

where W_{LH1}, W_{HL1} and W_{HH1} denote the three detail level-1 sub-bands in the wavelet transform images and $\overline{W_{LH1}}, \overline{W_{HL1}}$ and $\overline{W_{HH1}}$ denote their respective mean values. Following Yang and Nelson[22], we use single level discrete wavelet transform with Daubechies 6 filters.

2.2.5 DCT-based

The discrete cosine transformation gives a representation of an image in the spatial frequency domain. It's coefficients strongly correspond to the image contents. The sum of the AC (non-constant) components of the DCT is equal to the variance of the image intensity[23]. An advantage of this approach is that DCT is very often used in image and video compression and specialized hardware and algorithms can be utilized for it's computation.

8. **DCT reduced energy ratio**[24]

$$\phi = \sum \frac{F(0,1)^2 + F(1,0)^2 + F(2,0)^2 + F(1,1)^2 + F(0,2)^2}{F(0,0)^2} \tag{2.8}$$

where $F(u,v)$ is the discrete cosine transform coefficient of an 8×8 sub-block of the image calculated as

$$F(i,j) = \frac{1}{4} C_i C_j \sum_{x=0}^{7} \sum_{y=0}^{7} I(x,y) \cos\left(\frac{(2x+1)i\pi}{16}\right) \cos\left(\frac{(2y+1)j\pi}{16}\right),$$

$$C_i = \begin{cases} \frac{1}{\sqrt{2}}, & \text{if } i = 0 \\ 1, & \text{otherwise} \end{cases}, \quad C_j = \begin{cases} \frac{1}{\sqrt{2}}, & \text{if } j = 0 \\ 1, & \text{otherwise} \end{cases}$$

This is a faster and improved version of the original DCT energy ratio[25] which uses all of the 63 AC coefficients.

2.2.6 Miscellaneous

Finally, some metrics do not fit into any of the previous categories and are therefore merged into this last group.

9. **Brenner's method**[26, 20]

$$\phi = \sum_i \sum_j \max(|I(i,j) - I(i+2,j))|, |I(i,j) - I(i,j+2)|)^2 \tag{2.9}$$

Just like in the case of **Absolute gradient**, the original definition by Brenner[26] took into account only horizontal differences two pixels away but as Santos et al.[15] later noted, it improves the results significantly if the maximum from both the vertical and horizontal directions is used.

10. **Vollath's autocorrelation**[27]

$$\phi = \sum_i \sum_j I(i,j) I(i+1,j) - \sum_i \sum_j I(i,j) I(i+2,j) \tag{2.10}$$

2.3 Implementation

In preliminary experiments, we found that some implementation details can have a significant effect on the metric performance. To our knowledge, this isn't commonly discussed in literature but we consider it important.

All our code was written in Python. Most implementations are based on `numpy`[28] and `opencv`[29]. The discrete wavelet transform implementation comes from the `pywavelets`[30] library. The discrete cosine transform implementation and some utility methods are from `scipy`[31].

We found that for some metrics, the handling of image borders has a meaningful influence. The image borders can be perceived as edges and create artificial gradients. To avoid this problem, in the case of the **Lapacian-based** and **DCT-based** metrics as well as the **Brenner's method**, we simply discarded the pixels on the edges that might not contain valid data. This has proven to be the most stable solution.

In addition, some metrics, as defined by their authors, transform the two-dimensional array of image intensities into another two-dimensional array of values. But then it may not be clear how to aggregate them. We decided to simply sum the values, which is slightly faster than computing their average with similar results.

As we are interested in precision on the level of fractions of pixels and rotated images usually don't align to the pixel grid, it is necessary to consider some form of interpolation. We therefore tried different interpolation methods on different scenarios. The considered methods were linear, cubic and the Lanczos interpolation on a 8x8 pixel neighbourhood as they are commonly available and relatively fast.

Finally, it should be mentioned that the speed of computation is highly dependent on the specific software implementation, the hardware it is executed on and the image data itself. Even though there are considerable differences between them, all calculations run in less than 100 milliseconds on our test data and common hardware. We believe that virtually all of the tested metrics can be re-implemented to execute faster, so further analysis of their computation time seems unnecessary.

3. Experiments

3.1 Data set

Because real data from nuclear power plants are subject to strict safety rules and cannot be published, we had to create a new artificial data set. All our measurements were performed in the CVR laboratory using a similar setup described by Knotek[1], using a water tank and a custom-made camera-stand. We tried to mimic the real data as much as possible, but the clarity of the images is likely a bit higher, especially when using tube lighting.

We prepared six different configurations of the FA mock-up, which simulate its gradual aging. In the first configuration, the FA consists of new shiny rods that are perfectly straight. In the subsequent configurations, the rods are replaced with older, matte rods with increasing number of artificially simulated defects. These are gradually larger oxides imitated by a white marker and minor rod bending imitated by the insertion of small metal objects between adjacent rods. In all cases, there are two visible spacer grids placed half a meter apart.

For each configuration we captured two underwater videos using a Olympus IM005 camera, one with tube and one with point lighting. In each video, the camera captures the FA first as it moves from top to bottom and then from bottom to top. The procedure for creating OIOs can work with either direction and this allows us to select the one with better result. This way we ended up with twelve videos covering the whole FA twice from which we constructed twelve OIOs, pairs of which are of the same FA in different lighting conditions. The resulting data set can be seen in Figure 3.1.

The amount of samples is relatively small because they are very time consuming to produce and we had access to only a limited number of different components that make up the nuclear fuel assembly. Even so, we believe the new data set is sufficient for our purposes.

3.2 Setup

The purpose of our experiments was to evaluate the metric criteria outlined in section 2.1 in order to compare them. But that turned out not to be straightforward. Just like in real data, it is not possible to obtain ground truth, nor is it possible to do manually with our artificial data, because by blending consecutive frames, the edges are blurred and it is not possible to visually determine where the edge starts and where it ends. This is also true for other parts of the image and can be observed after zooming in in Figure 3.2.

To work around this, we utilized auto-registration, which means that the image is registered with itself. In this way, the ground truth is obvious and we can observe the metric behavior by adding artificial transformations with full control of its parameters. The considered transformations were the following as these are the most prominent in real data.

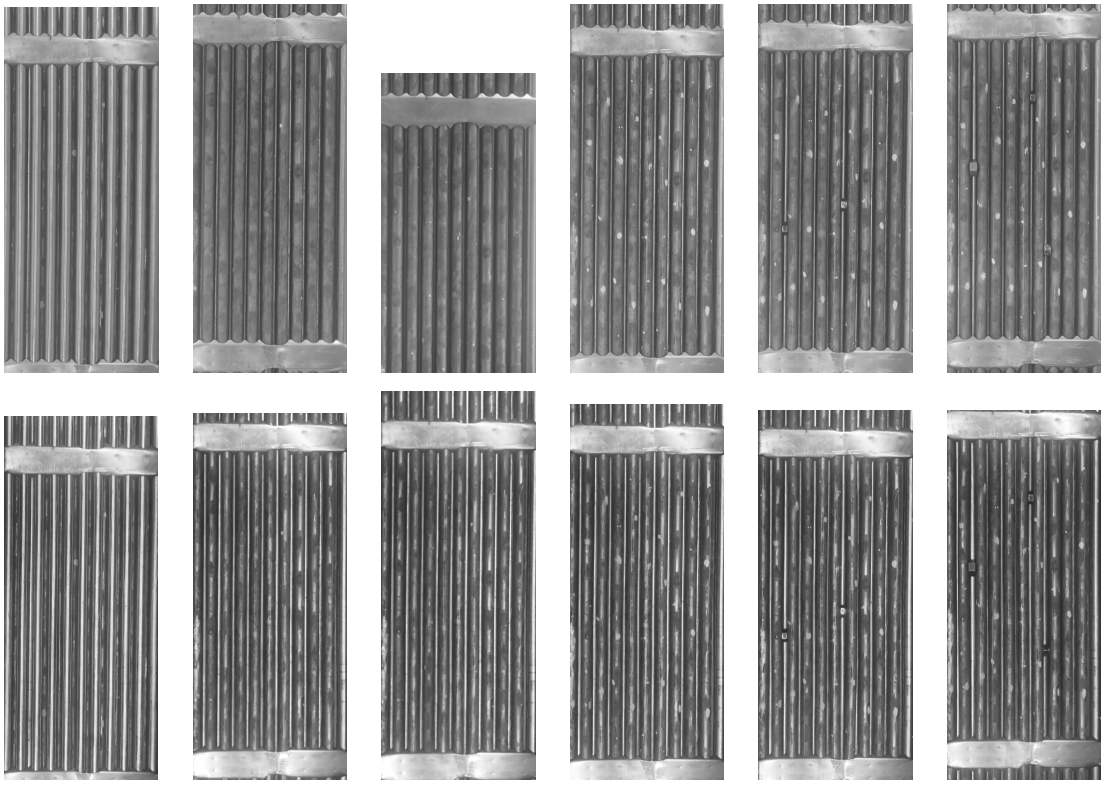

Figure 3.1: The final data set

Shift in both vertical and horizontal directions. We experimented with a small range of $[-2, 2]$ pixels with interpolated values for each $\frac{1}{10}$ of a pixel in between, a medium range of $[-10, 10]$ pixels with extra interpolated values in the middle between each two pixels and a large $[-100, 100]$ range with no interpolation.

Rotation around the center of the image. As this is a very invasive transformation, the valid extent of the image is reduced by tens or hundreds of pixels, we only considered a smaller and a larger range of $[-1, 1]$ and $[-7, 7]$ degrees with steps of size 0.01 and 0.1 respectively.

For each of the transformations above, we first observed the behavior of the metric when changing a single parameter. It can also be worthwhile to study them jointly, which is quite straightforward for horizontal and vertical shift, but also for rotation and horizontal shift as it is done in the current method by Knotek[1] as described in section 1.2.

Because auto-registration itself is not very representative of reality, in further experiments we added different variants of noise that are present in the real data. The most obvious choice was additive Gaussian noise with varying levels of intensity as it is common in modern cameras. Another type of noise that we experimented with adding that commonly occurs in real data is the presence of another transformation, rotation, or radial distortion that is not currently being solved, to the effect of which the metrics are to be invariant. Last but not least, it is desirable that the metrics work both on the whole image and its parts, and therefore we also applied the experiments to a top, bottom and a middle cutout of the data set images.

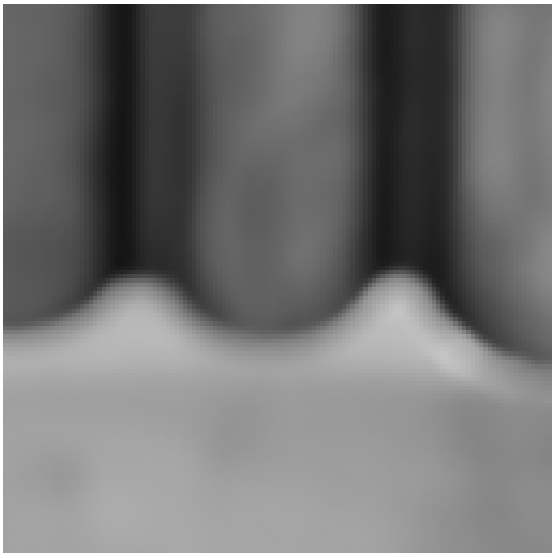

Figure 3.2: A close up of the blurry edge

The actual experiments were carried out by applying the artificial transformation to the image, blending it with the original unchanged image, optionally adding some noise and calculating the metric value. Some transformations reduce the extent of the image that contains valid data, so both blended images must be cropped first.

This way we ended up with a baseline and three additional noise configurations. In each configuration we tested vertical and horizontal shift and rotation on several different value ranges. Each experiment was run for all three different interpolation methods. The testing sample was 48 images from the average of which we obtained the final values.

4. Results

For each of the experiments, we have measured the criteria that were outlined in Section 2.1.

Accuracy was measured as the distance of the metric global maximum to the ground truth, which in the case of auto-registration is zero. More specifically, we're taking the average distance measured over all samples. It follows that the value, denoted in the table as "acc", can only be positive and that it is desirable that it is as low as possible and ideally zero.

Precision is given by the size of the range of values around the metric global maximum where it is strictly monotonically decreasing. The minimum of both directions, towards $-\infty$ and towards ∞, is used and then averaged over all samples. Unlike accuracy, for precision, denoted in the table as "prec", a value as high as possible is desired, ideally spanning the whole range of values. The value can again be only positive.

Robustness is reflected by how much the metrics performance deteriorates with the introduction of noise. It is not quantified in the tables by a specific number, but rather expressed by the change in accuracy and precision between sections 4.1 Baseline and 4.2 Noise.

Reliability is shown by on how many of the data set samples the metric behaves expectantly. Again, it is not quantified in the following tables, but it directly affects accuracy and precision.

Because any interpolation method wasn't decisively better in all scenarios, we've ran each experiment separately with each method to find the right combination of metric and interpolation method for each use case.

The complete results are part of Appendix A. Here we only present the overall results that are relevant to the following discussion.

4.1 Baseline

To get an idea of how the metrics behave in an ideal setting, we tested their performance on transformations without further difficulties. While this is very far from the real application, it gives us a baseline to compare the subsequent experiments to.

Vertical shift						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	**0.0**	**2.0**	0.0	5.14	1.79	63.12
2	0.0	1.77	0.0	7.91	0.0	89.6
3	0.0	1.79	0.0	2.67	0.0	45.1
4	0.0	1.5	0.0	1.5	0.0	2.06
5	0.04	1.33	0.06	7.82	4.02	71.48
6	0.75	0.09	3.25	0.44	46.46	14.06
7	0.0	1.65	0.0	2.58	0.0	41.31
8	0.17	1.02	0.33	7.4	10.56	56.15
9	0.0	1.79	0.0	8.06	**0.0**	**91.46**
10	0.0	1.73	0.02	4.56	23.83	7.65
Cubic						
1	0.2	0.2	0.01	7.95	1.79	63.12
2	0.21	1.22	0.02	9.6	0.0	89.6
3	0.0	1.83	0.0	4.1	0.0	45.1
4	0.1	0.1	0.0	1.54	0.0	2.06
5	0.32	0.95	0.21	9.67	4.02	71.48
6	0.75	0.08	3.25	0.44	46.46	14.06
7	0.0	1.56	0.0	4.36	0.0	41.31
8	0.43	1.19	0.68	9.24	10.56	56.15
9	0.2	1.01	0.03	9.66	**0.0**	**91.46**
10	0.29	0.53	0.17	5.19	23.83	7.65
Lancosz						
1	0.21	0.21	0.05	7.82	1.79	63.12
2	0.01	1.99	0.0	9.98	0.0	89.6
3	0.0	1.78	0.0	4.24	0.0	45.1
4	0.16	0.16	0.0	1.5	0.0	2.06
5	0.13	1.83	0.09	9.89	4.02	71.48
6	0.75	0.06	3.21	0.51	46.46	14.06
7	0.37	1.03	0.59	4.32	0.0	41.31
8	0.25	1.75	0.39	9.58	10.56	56.15
9	0.03	1.97	**0.0**	**10.0**	**0.0**	**91.46**
10	0.29	1.22	0.05	5.94	23.83	7.65

Table 4.1: Baseline vertical shift results

Horizontal shift						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	**0.0**	**2.0**	0.0	4.01	0.0	4.71
2	**0.0**	**2.0**	0.0	4.31	0.0	5.48
3	0.0	1.98	0.0	2.75	0.0	3.0
4	0.0	1.98	0.0	2.65	0.0	2.96
5	0.0	1.87	0.0	9.43	0.0	19.1
6	1.52	0.08	5.11	0.69	48.5	4.71
7	0.0	1.86	0.0	2.32	0.0	2.02
8	0.29	0.94	0.04	6.26	12.94	17.81
9	**0.0**	**2.0**	0.0	4.27	0.0	5.33
10	0.0	1.91	0.0	7.17	**0.0**	**26.75**
Cubic						
1	0.18	0.25	0.05	4.25	0.0	4.71
2	0.1	1.17	0.0	4.74	0.0	5.48
3	**0.0**	**2.0**	0.0	2.78	0.0	3.0
4	0.11	0.11	0.0	2.72	0.0	2.96
5	0.24	0.55	0.11	9.84	0.0	19.1
6	1.52	0.03	5.28	1.41	48.5	4.71
7	0.01	1.93	0.0	2.17	0.0	2.02
8	0.49	1.26	0.35	8.9	12.94	17.81
9	0.1	0.97	0.0	4.74	0.0	5.33
10	0.18	1.27	0.01	9.89	**0.0**	**26.75**
Lancosz						
1	0.1	0.79	0.0	4.47	0.0	4.71
2	**0.0**	**2.0**	0.0	4.96	0.0	5.48
3	0.0	1.94	0.0	2.78	0.0	3.0
4	0.12	0.12	0.0	2.75	0.0	2.96
5	0.01	1.99	0.01	9.99	0.0	19.1
6	1.51	0.06	5.25	0.92	48.5	4.71
7	0.17	1.83	0.8	1.91	0.0	2.02
8	0.33	1.67	0.14	9.21	12.94	17.81
9	**0.0**	**2.0**	0.0	4.95	0.0	5.33
10	**0.0**	**2.0**	**0.0**	**10.0**	**0.0**	**26.75**

Table 4.2: Baseline horizontal shift results

Rotation				
	[-1,1]		**[-7,7]**	
	acc	prec	acc	prec
Linear				
1	**0.0**	0.92	0.0	3.74
2	0.01	0.91	0.0	3.17
3	0.01	0.49	0.01	0.86
4	**0.0**	0.68	0.0	0.9
5	**0.0**	0.43	0.0	5.16
6	0.18	0.01	0.61	0.17
7	0.1	0.05	0.05	0.29
8	0.39	0.07	0.43	0.3
9	0.01	0.91	0.0	3.19
10	0.01	0.6	0.39	3.13
Cubic				
1	0.05	0.3	0.05	2.37
2	0.04	0.37	0.01	2.7
3	0.04	0.47	0.02	0.68
4	0.02	0.24	0.01	0.7
5	0.06	0.1	0.03	4.05
6	0.18	0.01	0.65	0.18
7	0.13	0.09	0.09	0.22
8	0.41	0.05	0.45	0.31
9	0.04	0.49	0.01	2.85
10	0.11	0.1	0.3	0.83
Lancosz				
1	0.03	0.42	0.02	2.79
2	0.01	**0.93**	0.0	3.2
3	0.05	0.34	0.04	0.58
4	0.03	0.26	0.02	0.66
5	0.01	0.48	**0.0**	**5.2**
6	0.18	0.01	0.65	0.19
7	0.2	0.08	0.26	0.17
8	0.4	0.06	0.45	0.3
9	0.01	**0.93**	0.0	3.2
10	0.08	0.41	0.36	1.87

Table 4.3: Baseline rotation results

4.2 Noise

In order to bring the experiments closer to reality, we have supplemented the baseline auto-registration with various other difficulties. Although they are largely artificial, their influence corresponds to reality.

4.2.1 Additive white Gaussian noise

In the case of white noise, the performance of metrics depends very much on its intensity. Therefore, the most telling thing is to choose an amount that is similar to the one that normally occurs in images taken with modern cameras. In the following experiments, we fixed the noise parameters to the mean of 0 and the standard deviation to 15. The noise was sampled independently for both the original and the transformed image and, after it was added, the intensity values were clamped to the original $[0, 255]$ range.

Vertical shift with white noise						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	0.44	0.14	0.64	0.82	2.38	2.19
2	0.24	0.16	**0.17**	2.44	**0.21**	6.54
3	0.71	0.13	3.27	0.69	16.5	1.44
4	0.73	0.13	2.91	0.66	9.56	1.4
5	0.38	0.14	0.52	1.86	4.25	**20.75**
6	1.59	0.09	4.95	0.68	53.06	1.73
7	0.66	0.13	3.17	0.69	9.12	1.27
8	0.55	0.13	0.99	0.77	11.19	3.6
9	0.23	0.15	0.27	1.78	0.4	5.46
10	**0.16**	0.15	0.27	1.49	23.15	2.81
Cubic						
1	0.52	0.14	0.62	0.8	2.38	2.19
2	0.32	0.14	0.36	2.79	**0.21**	6.54
3	0.68	0.13	3.2	0.71	16.5	1.44
4	0.7	0.13	2.97	0.65	9.56	1.4
5	0.4	0.13	0.55	1.98	4.25	**20.75**
6	1.44	0.05	4.91	0.73	53.06	1.73
7	0.71	0.14	3.19	0.7	9.12	1.27
8	0.59	0.13	0.97	0.77	11.19	3.6
9	0.3	0.12	0.34	2.22	0.4	5.46
10	0.27	0.12	0.4	2.3	23.15	2.81
Lancosz						
1	0.54	0.14	0.66	0.79	2.38	2.19
2	0.3	0.14	0.29	**3.44**	**0.21**	6.54
3	0.68	0.13	3.2	0.69	16.5	1.44
4	0.71	0.13	2.97	0.65	9.56	1.4
5	0.38	0.13	0.52	2.56	4.25	**20.75**
6	1.2	0.07	5.1	0.7	53.06	1.73
7	0.76	**0.15**	3.3	0.7	9.12	1.27
8	0.58	0.13	0.97	0.8	11.19	3.6
9	0.29	0.13	0.32	2.3	0.4	5.46
10	0.29	0.13	0.34	2.41	23.15	2.81

Table 4.4: White noise vertical shift results

Horizontal shift with white noise						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	0.07	0.45	0.03	3.34	0.02	4.15
2	**0.0**	1.06	**0.0**	4.23	0.0	4.83
3	0.19	0.16	0.58	1.42	7.5	2.38
4	0.14	0.22	0.41	1.8	7.38	2.56
5	0.03	0.72	**0.0**	9.41	**0.0**	**18.6**
6	1.95	0.03	4.01	1.41	35.08	5.67
7	0.12	0.31	0.49	1.69	7.75	1.96
8	0.32	0.2	0.08	4.95	12.94	15.83
9	**0.0**	0.94	0.01	4.11	0.0	4.73
10	0.34	0.13	0.34	1.1	1.67	3.46
Cubic						
1	0.19	0.17	0.08	3.56	0.02	4.15
2	0.1	0.26	0.01	4.48	0.0	4.83
3	0.32	0.12	0.59	1.69	7.5	2.38
4	0.28	0.14	0.52	2.09	7.38	2.56
5	0.2	0.14	0.09	9.66	**0.0**	**18.6**
6	1.93	0.01	3.85	2.19	35.08	5.67
7	0.16	0.17	0.58	1.8	7.75	1.96
8	0.41	0.14	0.34	7.72	12.94	15.83
9	0.13	0.21	0.02	4.27	0.0	4.73
10	0.38	0.13	0.48	1.34	1.67	3.46
Lancosz						
1	0.16	0.22	0.07	3.49	0.02	4.15
2	0.06	**1.2**	0.01	4.61	0.0	4.83
3	0.24	0.12	0.66	1.62	7.5	2.38
4	0.24	0.14	0.49	1.91	7.38	2.56
5	0.1	0.5	0.02	**9.98**	**0.0**	**18.6**
6	1.88	0.01	3.97	2.19	35.08	5.67
7	0.23	0.17	0.57	1.7	7.75	1.96
8	0.41	0.13	0.24	8.42	12.94	15.83
9	0.08	0.92	0.02	4.33	0.0	4.73
10	0.36	0.12	0.41	1.48	1.67	3.46

Table 4.5: White noise horizontal shift results

Rotation with white noise				
	[-1,1]		**[-7,7]**	
	acc	prec	acc	prec
Linear				
1	0.01	0.03	0.01	2.76
2	**0.01**	**0.14**	0.0	2.35
3	0.1	0.01	0.99	0.25
4	0.04	0.01	0.02	1.86
5	0.01	0.05	**0.0**	**4.23**
6	0.24	0.01	0.29	0.16
7	0.1	0.02	0.7	0.25
8	0.4	0.01	0.43	0.27
9	0.01	0.13	0.01	2.42
10	0.03	0.02	0.35	0.52
Cubic				
1	0.04	0.02	0.04	2.57
2	0.04	0.01	0.01	2.05
3	0.11	0.01	0.8	0.25
4	0.06	0.01	0.05	1.92
5	0.05	0.01	0.04	3.94
6	0.26	0.01	0.38	0.19
7	0.12	0.01	0.43	0.21
8	0.41	0.01	0.46	0.3
9	0.04	0.01	0.01	2.09
10	0.11	0.01	0.28	0.36
Lancosz				
1	0.03	0.01	0.03	2.58
2	0.02	0.05	0.0	2.38
3	0.1	0.01	0.63	0.26
4	0.07	0.02	0.05	1.8
5	0.02	0.01	0.01	4.22
6	0.27	0.01	0.31	0.18
7	0.15	0.01	0.58	0.16
8	0.41	0.01	0.44	0.2
9	0.02	0.04	0.01	2.4
10	0.08	0.02	0.39	0.46

Table 4.6: White noise rotation results

4.2.2 Radial distortion

Radial distortion can be parameterized in various ways, but we decided to use the common Brown's distortion model[32]. It works with an infinite number of coefficients and also takes into account tangential distortion, but for our purposes it was sufficient to use only the first coefficient of radial distortion k_1, which was set to 0.25. For our hypothetical camera, we set the optical center to the center of the image and the focal length to twice the size of the image.

Vertical shift with distortion						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	**0.0**	1.84	0.0	5.8	1.75	63.92
2	**0.0**	1.74	0.0	8.21	0.0	89.19
3	**0.0**	1.97	0.0	3.89	0.0	53.04
4	**0.0**	1.73	0.0	1.96	0.0	3.81
5	0.06	1.27	0.12	7.78	4.02	71.0
6	1.71	0.06	4.36	1.97	51.1	18.17
7	0.04	1.46	0.42	3.11	4.38	20.96
8	0.19	0.9	0.38	6.01	10.71	55.19
9	**0.0**	1.68	0.0	8.33	**0.0**	**90.98**
10	0.02	1.53	0.02	5.35	24.52	7.88
Cubic						
1	0.12	1.65	0.01	8.93	1.75	63.92
2	0.22	1.13	0.03	9.7	0.0	89.19
3	**0.0**	1.93	0.0	6.33	0.0	53.04
4	**0.0**	1.76	0.0	2.79	0.0	3.81
5	0.33	0.96	0.31	9.69	4.02	71.0
6	1.71	0.02	4.32	2.94	51.1	18.17
7	0.6	1.04	0.8	4.43	4.38	20.96
8	0.46	1.26	0.76	9.09	10.71	55.19
9	0.22	1.41	0.05	9.78	**0.0**	**90.98**
10	0.27	0.89	0.07	5.84	24.52	7.88
Lancosz						
1	0.02	**1.98**	0.0	9.55	1.75	63.92
2	0.02	**1.98**	**0.0**	**9.97**	0.0	89.19
3	0.03	1.94	0.0	6.27	0.0	53.04
4	**0.0**	1.78	0.0	2.65	0.0	3.81
5	0.12	1.88	0.15	9.85	4.02	71.0
6	1.71	0.02	4.33	2.6	51.1	18.17
7	0.64	1.1	1.0	4.46	4.38	20.96
8	0.29	1.71	0.51	9.41	10.71	55.19
9	0.02	**1.98**	**0.0**	**9.97**	**0.0**	**90.98**
10	0.08	1.9	0.02	6.31	24.52	7.88

Table 4.7: Distortion vertical shift results

Horizontal shift with distortion						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	0.0	1.91	0.0	4.29	0.0	4.88
2	**0.0**	**2.0**	0.0	4.52	0.0	5.54
3	0.0	1.98	0.0	2.83	0.0	3.19
4	**0.0**	**2.0**	0.0	2.83	0.0	3.08
5	0.0	1.84	0.0	9.86	0.0	19.67
6	1.66	0.19	5.48	1.2	47.65	5.06
7	0.04	1.51	0.23	2.21	14.08	2.27
8	0.23	1.04	0.0	7.15	16.04	20.48
9	**0.0**	**2.0**	0.0	4.48	0.0	5.5
10	0.0	1.86	0.0	7.83	**0.0**	**27.54**
Cubic						
1	0.16	1.04	0.01	4.47	0.0	4.88
2	0.12	0.68	0.0	4.94	0.0	5.54
3	**0.0**	**2.0**	0.0	2.98	0.0	3.19
4	**0.0**	**2.0**	0.0	2.89	0.0	3.08
5	0.26	0.58	0.11	9.89	0.0	19.67
6	1.64	0.14	5.8	2.02	47.65	5.06
7	0.17	1.83	0.42	2.17	14.08	2.27
8	0.45	1.31	0.38	8.83	16.04	20.48
9	0.12	0.74	0.0	4.93	0.0	5.5
10	0.2	1.26	0.01	9.91	**0.0**	**27.54**
Lancosz						
1	0.01	1.99	0.0	4.7	0.0	4.88
2	**0.0**	**2.0**	0.0	5.24	0.0	5.54
3	0.01	1.99	0.0	2.98	0.0	3.19
4	**0.0**	**2.0**	0.0	2.86	0.0	3.08
5	0.03	1.97	**0.0**	**10.0**	0.0	19.67
6	1.67	0.25	5.49	2.31	47.65	5.06
7	0.18	1.82	0.42	2.09	14.08	2.27
8	0.3	1.7	0.05	9.34	16.04	20.48
9	**0.0**	**2.0**	0.0	5.21	0.0	5.5
10	0.01	1.99	**0.0**	**10.0**	**0.0**	**27.54**

Table 4.8: Distortion horizontal shift results

Rotation with distortion				
	[-1,1]		**[-7,7]**	
	acc	prec	acc	prec
Linear				
1	**0.0**	**0.99**	0.0	4.0
2	0.0	0.94	0.0	3.33
3	0.0	0.28	0.0	0.75
4	0.0	0.75	0.0	1.02
5	0.0	0.54	**0.0**	**5.37**
6	0.35	0.03	1.09	0.29
7	0.13	0.12	1.36	0.21
8	0.28	0.13	0.29	0.36
9	0.0	0.94	0.0	3.42
10	0.04	0.45	0.24	2.95
Cubic				
1	0.04	0.28	0.01	3.49
2	0.04	0.25	0.01	2.87
3	0.01	0.39	0.0	0.72
4	0.01	0.65	0.0	0.95
5	0.06	0.08	0.03	4.74
6	0.35	0.03	1.11	0.29
7	0.14	0.11	1.28	0.28
8	0.31	0.07	0.31	0.36
9	0.04	0.25	0.01	2.88
10	0.1	0.1	0.3	2.16
Lancosz				
1	0.0	0.95	0.0	3.84
2	0.01	0.94	0.0	3.12
3	0.01	0.37	0.0	0.72
4	0.0	0.62	0.0	1.02
5	0.0	0.64	0.0	5.35
6	0.35	0.03	1.11	0.27
7	0.15	0.11	1.29	0.3
8	0.29	0.11	0.29	0.39
9	0.01	0.94	0.0	3.14
10	0.06	0.46	0.24	2.97

Table 4.9: Distortion rotation results

4.2.3 Rotation

During the image acquisition process, the camera can be and often is tilted and so the FA in the resulting frames is skewed. To simulate this we have added artificial rotation to both blended images. Its amount does not have much influence, so we chose an angle of 1 degree for the following experiments.

Vertical shift with rotation						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	**0.0**	1.77	0.0	6.06	3.44	64.65
2	**0.0**	1.77	0.0	8.1	0.0	84.98
3	**0.0**	1.98	0.0	4.11	0.0	65.48
4	**0.0**	1.87	0.0	2.69	0.0	3.81
5	0.08	1.28	0.15	7.39	4.62	67.42
6	1.62	0.09	4.57	2.21	51.08	23.27
7	**0.0**	1.49	0.31	3.4	1.06	37.81
8	0.21	0.86	0.54	6.81	10.71	53.21
9	**0.0**	1.67	0.0	8.29	**0.0**	**87.5**
10	0.04	1.51	0.0	5.0	27.02	7.1
Cubic						
1	0.08	1.92	0.0	8.75	3.44	64.65
2	0.22	1.19	0.03	9.74	0.0	84.98
3	**0.0**	1.96	0.0	6.66	0.0	65.48
4	**0.0**	1.83	0.0	2.66	0.0	3.81
5	0.31	0.93	0.31	9.69	4.62	67.42
6	1.51	0.14	3.86	4.94	51.08	23.27
7	0.07	1.55	0.3	5.09	1.06	37.81
8	0.43	1.21	0.86	9.14	10.71	53.21
9	0.21	1.32	0.04	9.75	**0.0**	**87.5**
10	0.25	0.85	0.05	5.48	27.02	7.1
Lancosz						
1	0.01	**1.99**	0.0	8.88	3.44	64.65
2	0.02	1.98	**0.0**	**9.98**	0.0	84.98
3	**0.0**	1.96	0.0	5.71	0.0	65.48
4	**0.0**	1.61	0.0	1.88	0.0	3.81
5	0.11	1.89	0.17	9.83	4.62	67.42
6	1.51	0.27	3.86	5.03	51.08	23.27
7	0.09	1.5	0.5	4.72	1.06	37.81
8	0.28	1.72	0.56	9.44	10.71	53.21
9	0.02	1.98	**0.0**	**9.98**	**0.0**	**87.5**
10	0.08	1.92	0.01	6.08	27.02	7.1

Table 4.10: Rotation vertical shift results

Horizontal shift with rotation						
	[-2,2]		**[-10,10]**		**[-100,100]**	
	acc	prec	acc	prec	acc	prec
Linear						
1	0.0	1.95	0.0	4.43	0.0	5.02
2	**0.0**	**2.0**	0.0	4.73	0.0	5.81
3	**0.0**	**2.0**	0.0	3.09	0.0	3.23
4	**0.0**	**2.0**	0.0	3.12	0.0	3.17
5	0.02	1.61	0.0	9.83	0.0	20.38
6	1.19	0.36	5.32	1.41	45.98	5.69
7	0.0	1.81	0.08	2.46	7.29	2.19
8	0.08	1.01	0.0	7.72	4.77	22.79
9	**0.0**	**2.0**	0.0	4.71	0.0	5.71
10	0.0	1.7	0.0	8.1	**0.0**	**26.83**
Cubic						
1	0.14	1.55	0.0	4.22	0.0	5.02
2	0.12	1.05	0.0	4.61	0.0	5.81
3	**0.0**	**2.0**	0.0	2.8	0.0	3.23
4	**0.0**	**2.0**	0.0	2.76	0.0	3.17
5	0.25	0.76	0.03	9.91	0.0	20.38
6	1.05	0.74	5.45	2.22	45.98	5.69
7	0.06	1.93	0.41	2.28	7.29	2.19
8	0.38	1.26	0.31	8.3	4.77	22.79
9	0.11	1.09	0.0	4.58	0.0	5.71
10	0.18	1.2	0.0	9.79	**0.0**	**26.83**
Lancosz						
1	**0.0**	**2.0**	0.0	4.49	0.0	5.02
2	**0.0**	**2.0**	0.0	5.05	0.0	5.81
3	**0.0**	**2.0**	0.0	2.82	0.0	3.23
4	**0.0**	**2.0**	0.0	2.74	0.0	3.17
5	0.03	1.97	**0.0**	**10.0**	0.0	20.38
6	0.99	0.97	5.51	2.25	45.98	5.69
7	0.05	1.93	0.42	2.09	7.29	2.19
8	0.2	1.8	0.05	9.06	4.77	22.79
9	**0.0**	**2.0**	0.0	5.03	0.0	5.71
10	0.01	1.99	**0.0**	**10.0**	**0.0**	**26.83**

Table 4.11: Rotation horizontal shift results

5. Discussion

5.1 Vertical shift

Vertical shift detection is useful not only for registering different OIOs, but also for determining the speed at which the camera moves when capturing the FA. This can lead to sharper and higher quality OIOs, and moreover, one does not need to enter it as a parameter manually.

In general, registration on the y-axis appears to be more complicated probably because there are relatively few horizontal edges in the images of nuclear fuel that the algorithm could rely on. This is especially the case when the image does not contain a part of a spacer grid or when the subject FA consists of new shiny rods with no oxidization. On the other hand, vertical shift registration should be possible on a relatively large range, precisely because the image edges do not align anywhere else better than at the optimal setting.

Our experiments in section 4.2.1 showed that of all the difficulties, white noise has the greatest negative effect on the detection of vertical shift. Our hypothesis is that due to the white noise, false local edges appear in the image, which confuses most of the tested metrics. The other two artificial transformations, radial distortion and rotation, didn't have quite as significant influence, but even so, the performance of most metrics deteriorated noticeably by their cause.

Most metrics performed very well on the small strongly interpolated $[-2, 2]$ range in the baseline test but failed with the introduction of noise. Overall however, the **Gradient-based** and the **Laplacian-based** metrics, as well as the **Brenner's method** turned out to be the most appropriate. There isn't an obvious best interpolation method but the Lancosz interpolation is recommended with distortion or rotation.

On the medium $[-10, 10]$ range, we can see a significant drop in precision for the **Laplacian-based** and the **Wavelet-based** metrics. Here the best pick is **Tenengrad** with white noise, but also the **Normalized gray-level variance**, the **DCT reduced energy** and the **Brenner's method** are all valid choices with the right interpolation method, which usually is the Lancosz method.

Here the best are the **Gradient-based** metrics and the **Brenner's method** and especially **Tenengrad** when noise is present.

In the case of the largest $[-100, 100]$ range, we see that in the **Miscellaneous** metrics the **Vollath's autocorrelation** reaches the limit of its precision, while the **Brenner's method** looks like the best choice. Even in this case, the **Gradient-based** metrics behave nicely and in addition also the **Normalized gray-level variance** with white noise.

5.2 Horizontal shift

Solving registration in the horizontal direction is useful because it can compensate for camera shake and possible inaccuracies caused by its tilt, which need to be corrected by rotating the affected frames. Of course, it is also necessary to register the stitched OIOs and can improve their quality.

Finding the optimal horizontal shift is apparently a bit easier, probably because the images of nuclear fuel contain a large number of vertical edges, regardless of which part of the image it is. The range of values on which registration can be performed is limited by the fact that the centers of the individual rods are approximately 50 pixels apart on the x-axis and from a certain point the algorithm is derailed and tries to align mismatched rods.

As in the case of the vertical shift, white noise had the most negative effect here, albeit to a smaller extent. Intuitively, the larger number of vertical edges is less likely to be eclipsed by the white noise. Radial distortion and rotation had almost no effect on the metrics performance.

Most metrics again performed well in the small $[-2, 2]$ range with the **Gradient-based**, **Laplacian-based** and **Miscellaneous** metrics all being a sensible choice. Unlike most the other metrics, the **Tenengrad** remained notably robust even with white noise. The selection of an interpolation method depends on the used metric and on whether accuracy or precision is more important.

Again, as in the previous case, here we see a significantly limited precision of the **Laplacian-based** and the **Wavelet-based** metrics in the medium $[-10, 10]$ range. The **Normalized gray-level variance**, the **DCT reduced energy ratio** and the **Vollath's autocorrelation** all seem to be the most appropriate, the first two of which are fairly invariant even to white noise. At this range, the Lancosz interpolation is almost always the best choice, especially for the **DCT reduced energy ratio**.

In the large $[-100, 100]$ range we can see that the precision of most metrics is very limited. This isn't the case for the **Vollath's autocorrelation**. In case noise is present, the **Normalized gray-level variance** or the **DCT reduced energy ratio** is a better choice.

5.3 Rotation

Rotation correction is very often necessary because the camera is usually slightly tilted, not more than a few degrees, when recording, and the resulting OIO requires that the rods be perpendicular to the ground before they can be stitched together. On the other hand, when registering the finished OIOs, rotation can usually be omitted.

Rotation registration is somewhat different from the previous translations. First of all, it does not work in whole units of pixels, which inevitably leads to interpolations, and in addition it is a transformation that reduces the extent of valid pixels in the image, so cropping is almost always necessary. Although it may not always be true, Knotek[1] worked on the assumption that the rotation of the camera remains the same during the video acquisition, so it should be sufficient to rotate all images by the same angle. That implies that the optimal solution can be searched for using multiple frames.

In line with the previous sections, the biggest problem for rotation was white noise. This is especially the case for the smaller range of values, where interpolation took place in small steps. For rotation, only radial distortion could be used as another artificial transformation, which again did not pose a big problem.

For the smaller $[-1, 1]$ range, the **Gradient-based** methods and the **Brenner's method** proved to be the best. In the case of a significant amount of white noise, the **Tenengrad** is particularly suitable. The cubic interpolation is generally the worst choice here.

By far not all metrics can handle the larger $[-7, 7]$ range, but either the **Normalized gray-level variance**, the **Gradient-based** or the **Miscellaneous** metrics can be used with the first one being the most robust to white noise. The selection of an interpolation method for this range depends on the metric and scenario.

5.4 Findings

First of all, we must mention that, despite the promising results in our preliminary experiments, **Mason and Green's histogram method** proved to be completely unsuitable for this use case. We believe that this is not an error in implementation, but rather the instability and lack of reliability of the threshold selection and possibly the histogram method as a whole. In the same vein, the **DCT-based** methods have proven to be fairly unreliable, especially with regard to accuracy. Further analysis is necessary to explain these phenomenons.

Another interesting finding that emerged from our experiments is that most metrics are extremely prone to white noise. As mentioned earlier, we think this is due to the false edges created by the noise, which overwhelms the original structure of the image. Interestingly, the effect is even more pronounced with interpolation. It is probably worth mentioning that adding white noise to an already stitched OIO does not correspond to reality, because it is the blending of frames that averages their values and removes most of the noise. However, the original raw frames usually contain some noise, so it is desirable to know how it affects the metrics.

Of the metrics tested, two proved to be suitable in situations with an increased amount of noise. The **Normalized gray-level variance** is likely robust and reliable because of its inherent normalization and in the case of **Tenengrad** it is the Sobel filters that are less susceptible to noise than, for example, Laplacian.

The **Gradient-based** metrics and especially **Tenengrad** seem to be a good overall compromise for all use cases and is recommended when a single metric needs to be chosen. Its accuracy and precision is one of the best while also being very robust and reliable.

The **Laplacian-based** metrics have a small precision which is caused by the Mexican hat shape that its values form around the optimum. They do have a very sharp and pronounced peak which makes them useful at small ranges. The **Variance of wavelet coefficients** manifests the same behavior but generally performs worse while also being significantly slower so it cannot be recommended with this implementation. As mentioned in Section 2.2.4, we've used the Daubechies 6 wavelets following the recommendation by Yang and Nelson[22]. Further experiments with other wavelets are needed in order to reject the **Wavelet-based** metrics as a whole.

The **Normalized gray-level variance** has a very smooth function surface spanning large ranges of values very suitable for gradient optimization algorithms. It does somewhat lack in terms of accuracy.

Both **Miscellaneous** metrics are very susceptible to white noise but have shown to be exceptionally viable for vertical and horizontal shift in case of the **Brenner's method** and the **Vollath's autocorrelation** respectively. This would also suggest that including both the horizontal and vertical direction in the computation of the **Brenner's method** as suggested by Santos et al.[15] is a very sensible choice.

The interpolation method results aren't very conclusive. The linear interpolation is often the best and generally the most accurate at the smallest, most interpolated range. The cubic interpolation is rarely the best option and so cannot really be recommended. The Lancosz interpolation is generally better at the larger ranges with perhaps a slightly worse accuracy.

In the presence of a white noise, linear interpolation usually maintains the best accuracy and the Lancosz interpolation has a slightly better precision. The same is generally true for radial distortion and rotation.

Overall, however, the choice of an interpolation method depends most on the metric used. For example, the **DCT reduced energy ratio** works significantly better with one of the more sophisticated methods.

It should also be noted that the speed of computation of each metric is not the same. In our experiments the cubic and Lancosz interpolations were about 20% and 90% slower than the linear interpolation respectively. Nevertheless, we believe that this slowdown may be worth it in some cases.

5.5 Summary

The results of our experiments lead us to believe there isn't a single overall best metric or interpolation method which could be used exclusively. Instead, their combination can be used playing into their respective strengths. The registration can then be solved hierarchically, starting with a large range of values and finding the approximate location of the optimum for all parameters of the transformation and then gradually refining them in a smaller neighborhood, while utilizing the metric and interpolation method most suitable at that point.

In the first rough estimation, metrics with large precision are necessary. For this purpose, the **Normalized gray-level variance** is the most versatile but the **Brenner's method** can also be used for vertical shift or rotation, the **Vollath's autocorrelation** for horizontal shift or the **Gradient-based** metrics for rotation. The Lancosz interpolation generally has the better precision and is the most suitable here.

In the refinement step, a well pronounced peak is desirable. That is a known feature of the **Laplacian-based** but also of the **Gradient-based** metrics which makes them very suitable. The fine-tuning can be done for each parameter individually and in multiple iterations if high accuracy is required. The linear interpolation has the best accuracy and should be used at this stage.

We have put together a proof-of-concept implementation to verify the approach outlined above. In the first version, the value of the metric for both shifts

and rotation is directly maximized using the bounded Powell method[8]. This approach was only viable for the **Tenengrad** and the **Normalized gray-level variance**. In the second version, we've employed a hierarchical scheme. First, rotation is solved using the **Normalized gray-level variance**, then the approximate horizontal shift using the **Vollath's autocorrelation** and the approximate vertical shift using the **Brenner's method**, all using the Lancosz interpolation, and finally the horizontal shift and then the vertical shift is refined using the **Absolute gradient** with linear interpolation. This is done separately for each parameter using the bounded Brent method[33].

Some example results can be seen in Figure 5.1. As we can see, both versions are generally sharper. In the case of the first version, however, the rotation is clearly not entirely correct. In the second version, this problem no longer occurs.

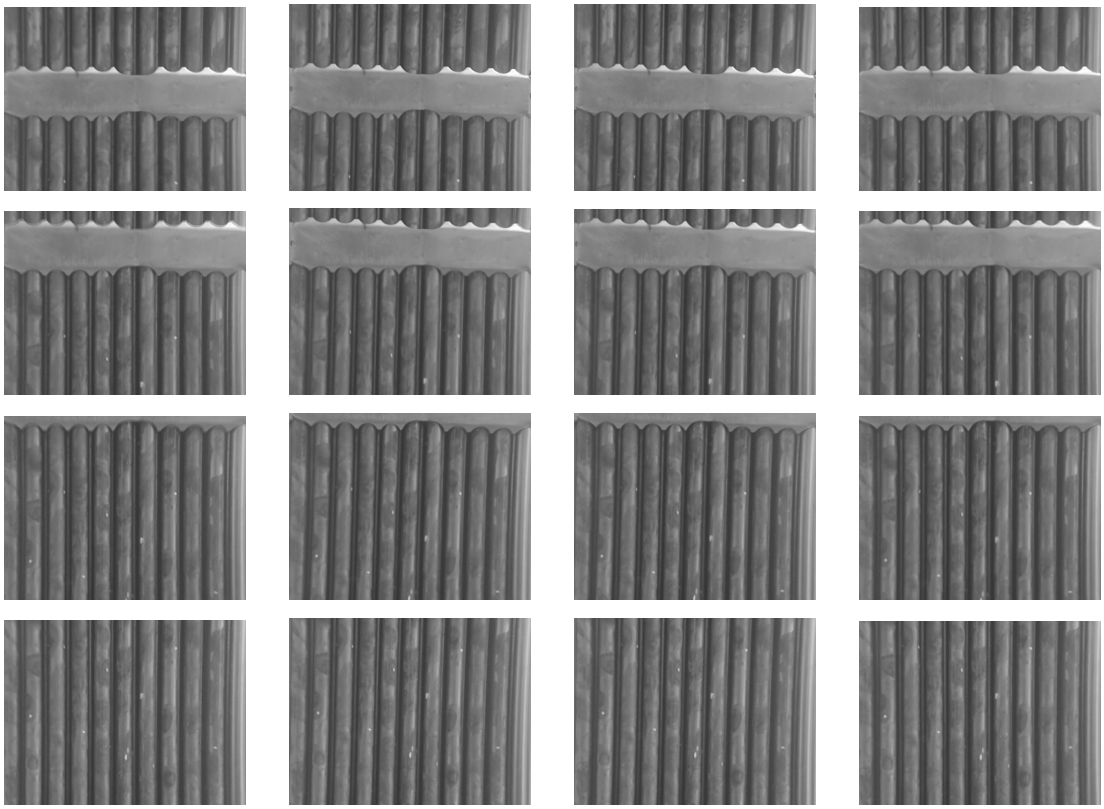

(a) Current Method (b) Tenengrad (c) Normalized gray-level variance (d) Hierarchical

Figure 5.1: A comparison of the current method and the new proof-of-concept methods. The blended frames are from the same video, 50 frames apart.

Conclusion

In this thesis, we performed an analysis of the method of obtaining and the resulting properties of nuclear fuel images. We have summarized what image registration is, current approaches to solving it and how it is possible to apply them to the problem of registration of images of nuclear fuel assembly.

Then, we took inspiration from the field of autofocus and shape-from-focus and made an overview of similarity metrics that can be used to guide an optimization algorithm for registration. We outlined criteria that are important for their practical use, grouped them based on their main principle and, based on preliminary experiments, chose the most promising representatives from each group. We also briefly mentioned several implementation details that aren't usually discussed in literature.

Next, we designed and ran experiments to test how the individual selected metrics meet the specified criteria. This also included the creation of a new data set in cooperation with the Research Centre ŘEŽ (CVR). Apart from artificial experiments in ideal conditions, we tried to approach reality by adding various forms of noise and transformations present in real data.

Finally, we discussed our results and made recommendations on when it is appropriate to use which metrics. We showcased how our findings can be used to improve the current method and outlined the possibilities for future work.

Bibliography

1. KNOTEK, Jaroslav. *Automation of nuclear fuel visual inspection*. Prague, 2020. MA thesis. Charles University, Faculty of Mathematics, Physics, Department of Software, and Computer Science Education.
2. ZITOVÁ, Barbara; FLUSSER, Jan. Image registration methods: a survey. *Image and Vision Computing*. 2003, vol. 21, no. 11, pp. 977–1000. ISSN 0262-8856. Available from DOI: 10.1016/S0262-8856(03)00137-9.
3. LIU, Zhe; DENG, Xiang; WANG, Guang-zhi. Accuracy Validation for Medical Image Registration Algorithms: a Review. *Chinese Medical Sciences Journal*. 2012, vol. 27, no. 3, pp. 176–181. ISSN 1001-9294. Available from DOI: 10.1016/S1001-9294(14)60052-3.
4. WYAWAHARE, Medha; PATIL, P.M.; ABHYANKAR, Hemant. Image Registration Techniques: An overview. *International Journal of Signal Processing, Image Processing and Pattern Recognition*. 2009, vol. 2.
5. JAIN, Shweta; KANWAL, Navdeep. Overview on image registration. In: *2014 International Conference on Medical Imaging, m-Health and Emerging Communication Systems (MedCom)*. 2014, pp. 376–381. Available from DOI: 10.1109/MedCom.2014.7006036.
6. CHO, Jai; CHOI, Young-Soo; JEONG, Kyung-Min; SHIN, Jung-Cheol. Measurement of nuclear fuel rod deformation using an image processing technique. *Nuclear Engineering and Technology*. 2011, vol. 43. Available from DOI: 10.5516/NET.2011.43.2.133.
7. RAJPUT, Hitesh; SOM, Tanmoy; KAR, Soumitra. Fuel Assembly Verification for Safe Operation of Nuclear Power Plants. *Nuclear Technology*. 2015, vol. 192, no. 2, pp. 125–132. Available from DOI: 10.13182/NT14-154.
8. POWELL, M. J. D. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*. 1964, vol. 7, no. 2, pp. 155–162. ISSN 0010-4620. Available from DOI: 10.1093/comjnl/7.2.155.
9. VU, Cuong T.; CHANDLER, Damon M. S3: A Spectral and Spatial Sharpness Measure. In: *2009 First International Conference on Advances in Multimedia*. 2009, pp. 37–43. Available from DOI: 10.1109/MMEDIA.2009.15.
10. FERZLI, Rony; KARAM, Lina J. A No-Reference Objective Image Sharpness Metric Based on the Notion of Just Noticeable Blur (JNB). *IEEE Transactions on Image Processing*. 2009, vol. 18, no. 4, pp. 717–728. Available from DOI: 10.1109/TIP.2008.2011760.
11. HASSEN, Rania; WANG, Zhou; SALAMA, Magdy M. A. Image Sharpness Assessment Based on Local Phase Coherence. *IEEE Transactions on Image Processing*. 2013, vol. 22, no. 7, pp. 2798–2810. Available from DOI: 10.1109/TIP.2013.2251643.

12. PEI, Zhao; HUANG, Li; ZHANG, Yanning; MA, Miao; PENG, Yali; YANG, Yee-Hong. Focus Measure for Synthetic Aperture Imaging Using a Deep Convolutional Network. *IEEE Access*. 2019, vol. 7, pp. 19762–19774. Available from DOI: 10.1109/ACCESS.2019.2896655.

13. PERTUZ, Said; PUIG, Domenec; GARCIA, Miguel Angel. Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*. 2013, vol. 46, no. 5, pp. 1415–1432. ISSN 0031-3203. Available from DOI: 0.1016/j.patcog.2012.11.011.

14. JARVIS, RA. Focus optimization criteria for computer image processing. *Microscope*. 1976, vol. 24, no. 2, pp. 163–180.

15. SANTOS, A; ORTIZ-DE-SOLORZANO, Carlos; VAQUERO, Juan Jose; PEÑA, J; MALPICA, Norberto; DEL POZO GUERRERO, Francisco. Evaluation of autofocus functions in molecular cytogenetic analysis. *Journal of Microscopy*. 1997, vol. 188, no. 3, pp. 264–272. Available from DOI: 10.1046/j.1365-2818.1997.2630819.x.

16. TENENBAUM, J. M. *Accommodation in computer vision*. 1970. PhD thesis. Stanford University, Department of Electrical Engineering.

17. KROTKOV, E. Focusing. *International Journal of Computer Vision*. 1988, vol. 1, pp. 223–237.

18. PECH-PACHECO, J.L.; CRISTOBAL, G.; CHAMORRO-MARTINEZ, J.; FERNANDEZ-VALDIVIA, J. Diatom autofocusing in brightfield microscopy: A comparative study. In: *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*. 2000, vol. 3, 314–317 vol.3. ISBN 0-7695-0750-6. Available from DOI: 10.1109/ICPR.2000.903548.

19. NAYAR, S. K.; NAKAGAWA, Y. Shape from focus. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1994, vol. 16, no. 8, pp. 824–831. Available from DOI: 10.1109/34.308479.

20. FIRESTONE, Lawrence; COOK, Kitty; CULP, Kevin; TALSANIA, Neil; PRESTON JR., Kendall. Comparison of autofocus methods for automated microscopy. *Cytometry*. 1991, vol. 12, no. 3, pp. 195–206. Available from DOI: 10.1002/cyto.990120302.

21. MASON, DC; GREEN, DK. Automatic focusing of a computer-controlled microscope. *IEEE transactions on biomedical engineering*. 1975, vol. 22, no. 4, pp. 312–317.

22. YANG, Ge; NELSON, B.J. Wavelet-based autofocusing and unsupervised segmentation of microscopic images. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2003, vol. 3, 2143–2148 vol.3. Available from DOI: 10.1109/IROS.2003.1249188.

23. BAINA, J.; DUBLET, J. Automatic focus and iris control for video cameras. In: *Fifth International Conference on Image Processing and its Applications, 1995*. 1995, pp. 232–235. Available from DOI: 10.1049/cp:19950655.

24. LEE, Sang-Yong; YOO, Jae-Tack; KUMAR, Yogendera; KIM, Soo-Won. Reduced Energy-Ratio Measure for Robust Autofocusing in Digital Camera. *IEEE Signal Processing Letters*. 2009, vol. 16, no. 2, pp. 133–136. Available from DOI: 10.1109/LSP.2008.2008938.

25. SHEN, Chun-Hung; CHEN, H.H. Robust focus measure for low-contrast images. In: *2006 Digest of Technical Papers International Conference on Consumer Electronics*. 2006, pp. 69–70. Available from DOI: 10.1109/ICCE.2006.1598314.

26. BRENNER, J F; DEW, B S; HORTON, J B; KING, T; NEURATH, P W; SELLES, W D. An automated microscope for cytologic research a preliminary evaluation. *Journal of Histochemistry & Cytochemistry*. 1976, vol. 24, no. 1, pp. 100–111. Available from DOI: 10.1177/24.1.1254907. PMID: 1254907.

27. VOLLATH, D. The influence of the scene parameters and of noise on the behaviour of automatic focusing algorithms. *Journal of Microscopy*. 1988, vol. 151, no. 2, pp. 133–146. Available from DOI: 10.1111/j.1365-2818.1988.tb04620.x.

28. HARRIS, Charles R.; MILLMAN, K. Jarrod; WALT, Stéfan J. van der; GOMMERS, Ralf; VIRTANEN, Pauli; COURNAPEAU, David; WIESER, Eric; TAYLOR, Julian; BERG, Sebastian; SMITH, Nathaniel J.; KERN, Robert; PICUS, Matti; HOYER, Stephan; KERKWIJK, Marten H. van; BRETT, Matthew; HALDANE, Allan; RÍO, Jaime Fernández del; WIEBE, Mark; PETERSON, Pearu; GÉRARD-MARCHANT, Pierre; SHEPPARD, Kevin; REDDY, Tyler; WECKESSER, Warren; ABBASI, Hameer; GOHLKE, Christoph; OLIPHANT, Travis E. Array programming with NumPy. *Nature*. 2020, vol. 585, no. 7825, pp. 357–362. Available from DOI: 10.1038/s41586-020-2649-2.

29. BRADSKI, G. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*. 2000.

30. LEE, Gregory R.; GOMMERS, Ralf; WASELEWSKI, Filip; WOHLFAHRT, Kai; O'LEARY, Aaron. PyWavelets: A Python package for wavelet analysis. *Journal of Open Source Software*. 2019, vol. 4, no. 36, p. 1237. Available from DOI: 10.21105/joss.01237.

31. VIRTANEN, Pauli; GOMMERS, Ralf; OLIPHANT, Travis E.; HABERLAND, Matt; REDDY, Tyler; COURNAPEAU, David; BUROVSKI, Evgeni; PETERSON, Pearu; WECKESSER, Warren; BRIGHT, Jonathan; VAN DER WALT, Stéfan J.; BRETT, Matthew; WILSON, Joshua; MILLMAN, K. Jarrod; MAYOROV, Nikolay; NELSON, Andrew R. J.; JONES, Eric; KERN, Robert; LARSON, Eric; CAREY, C J; POLAT, İlhan; FENG, Yu; MOORE, Eric W.; VANDERPLAS, Jake; LAXALDE, Denis; PERKTOLD, Josef; CIMRMAN, Robert; HENRIKSEN, Ian; QUINTERO, E. A.; HARRIS, Charles R.; ARCHIBALD, Anne M.; RIBEIRO, Antônio H.; PEDREGOSA, Fabian; VAN MULBREGT, Paul; SCIPY 1.0 CONTRIBUTORS. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. 2020, vol. 17, pp. 261–272. Available from DOI: 10.1038/s41592-019-0686-2.

32. BROWN, Duane C. Decentering distortion of lenses. In: 1966.

33. BRENT, Richard P. An algorithm with guaranteed convergence for finding a zero of a function. *The Computer Journal*. 1971, vol. 14, no. 4, pp. 422–425.

List of Figures

1.1 Images of FAs with lighting conditions 5

3.1 The final data set 15
3.2 A close up of the blurry edge 16

5.1 A comparison of the current method and the new proof-of-concept methods. The blended frames are from the same video, 50 frames apart. 33

List of Tables

4.1 Baseline vertical shift results 18
4.2 Baseline horizontal shift results 19
4.3 Baseline rotation results 20
4.4 White noise vertical shift results 21
4.5 White noise horizontal shift results 22
4.6 White noise rotation results 23
4.7 Distortion vertical shift results 24
4.8 Distortion horizontal shift results 25
4.9 Distortion rotation results 26
4.10 Rotation vertical shift results 27
4.11 Rotation horizontal shift results 28

List of Abbreviations

CVR Research Centre ŘEŽ (The research organisation Centrum výzkumu Řež (CVŘ) was founded on 9th October 2002 as 100% daughter company of ÚJV Řež, a. s. (NRI – Nuclear Research Institute). The main aim of the research organisation is research, development and innovations in the field of power generation (especially nuclear).)

FA fuel assembly (A bundle of a few hundred fuel rods held together in a predetermined hexagonal shape by spacer grids.)

OIO one-image-overview (An image of one whole side of a nuclear fuel assembly stitched together from consecutive video frames.)

A. Attachments

In the attached file, there are two folders:

- `data` - containing the data set and other images used for experiments
  - `preliminary.png` - the OIO used in the preliminary experiments
  - `<1-6>-point.png` and `<1-6>-tube.png` - the data set of 16 OIOs
  - `frame-<1-5>.png` - the raw frames taken from the video that was used to obtain the `3-tube.png` image that were used for visualisation and comparison with the current method
- `notebooks` - Jupyter notebooks with the metric implementations, replicable experiments and visualisations
  - `preliminary-experiments.ipynb` - for the selection of metrics
  - `final-experiments.ipynb` - as outlined in Chapter 3
  - `visualisation.ipynb` - interesting and useful graphs and plots