Similarity Measures for Visual Comparison and Retrieval of Test Data in Aluminum Production

Nikolina Jekic[1][a], Belgin Mutlu[1,4], Manuela Schreyer[2], Steffen Neubert[2] and Tobias Schreck[3]

[1] *Pro2Future GmbH, Inffeldgasse 25f, 8010 Graz, Austria*

[2] *AMAG Austria Metall AG, Lamprechtshausener Strasse 61, 5282 Ranshofen, Austria*

[3] *Graz University of Technology, Institut of Computer Graphics and Knowledge Visualisation, Inffeldgasse 16c , 8010 Graz, Austria*

[4] *Graz University of Technology, Institut of Interactive Systems and Data Science, Inffeldgasse 16c, 8010 Graz, Austria*

Keywords: Similarity Measures, Visual Analysis, Aluminum Casting.

Abstract: Monitoring, analyzing and determining the production quality in a complex and long-running process such as in the aluminum production is a challenging task. The domain experts are often overwhelmed by the flood of data being generated and collected and have difficulties to analyze and interpret the results. Likewise, experts find it difficult to identify patterns in their data that may indicate deviations and anomalies that lead to unstable processes and lower product quality. We aim to support domain experts in the production data exploration and identifying meaningful patterns. The existing research covers a broad spectrum of pattern recognition methodologies that can be potentially applied to elicit patterns in data collected from industrial production. Hence, in this paper, we further analyze the applicability of different similarity measures to effectively recognize specific ultrasonic patterns which may indicate critical process deviations in aluminum production.

1 INTRODUCTION

The goal of an optimal manufacturing process is to increase productivity and customer satisfaction while minimizing cost, time, and waste. Achieving a high quality of products while remaining competitive, requires companies to continuously improve the performance of their production process. Process data may contain important information such as meaningful relationships and patterns, which could help to improve the quality of the production process (Yin and Kaynak, 2015) (Thalmann et al., 2018). Yet, human beings are overwhelmed by the amount of data being generated in such complex production processes. Visual data analysis has proven to be one of the effective ways to tackle this problem (Soban et al., 2016). The existing research does not only support the exploration of the data and detect hidden patterns/correlations, but they also pave the way to define new methods for improving the production process and increasing production number (Suschnigg et al., 2020) (Sun et al., 2019).

The production process in the aluminum industry is complex and time-consuming. A simplified aluminum production process from the recycled raw material up to the final products includes melting, alloying and further treatment, casting, homogenization, rolling and quality control. In a nutshell, during the cast, each batch results in several ingots. Ingots are molds cast from molten aluminum and are suitable for production processing using methods such as rolling, extrusion, and forging (Vasudevan and Doherty, 2012). These ingots are further rolled to plates and sheets. Finally, to assess quality, experts perform ultrasonic tests (UT) on rolled aluminum plates. Each part of the production process is done in accordance with very high-quality standards. However, due to the complex process dynamics, the final product might not meet them and show certain degrees of indications.

Non-metallic indications that are already contained in the input material or are formed in the foundry production process, can lead to rejects in materials after ultrasonic inspection of the final plates. This leads to reduced capacity, longer delivery times and higher costs. Ultrasonic testing is the last step in the complete process chain and can not be done directly after casting. It can therefore happen that a product goes through the entire process, but finally

[a] https://orcid.org/0000-0001-9884-0929

Jekic, N., Mutlu, B., Schreyer, M., Neubert, S. and Schreck, T.
Similarity Measures for Visual Comparison and Retrieval of Test Data in Aluminum Production.
DOI: 10.5220/0010309302100218
In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2021) - Volume 3: IVAPP*, pages 210-218
ISBN: 978-989-758-488-6

does not meet the required standard. The causes and the influencing factors that lead to UT indications in the casting process are not yet fully understood.

Different groups of production parameters influence the quality of batches and ingots, and this may be an implication that only some batches and ingots would have specific indication patterns. Influencing parameters can be possibly found, for example, in the input material, in process parameters of the furnace and casting plant, where several hundred different parameters are continuously recorded during the casting, or in the chemical composition of the different batches. By far, most of the ingots are of good quality, with no or very few indications, which do not lead to rejects after ultrasonic inspection. Some are of mixed quality, which means that some parts of the ingot are of good quality, with very few indications, and some parts are of poor quality with many indications. In this case, some of the plates made from the entire ingot are rejects after the ultrasonic testing. In certain situations, the whole ingot is of poor quality, which means that all the plates made of the entire ingot are rejects, whereas neighboring ingots from the same batch are of perfect quality. The number of indications in an ingot depends also on some product-specific parameters, for example, the material of the ingot, the casting format and the format of the final product, which makes the analysis even more difficult. Additionally, the total number of test batches is limited, especially if one wants to concentrate on a specific alloy and/or format. The overall small number of indications and rejects complicates the analysis and research even further. Thus, in the following, we restricted the analysis to one specific material, casting- and final-product-format in order to lay out the methodology to the inclined reader.

Grouping and recommending batches and ingots with similar indication patterns, is highly desired to support the end-users, i.e., material engineers in casting and rolling, in production data exploration and inspection of possible influencing parameters on product quality. Distance or similarity measures are essential to solving many pattern recognition problems such as classification, clustering, and retrieval problems (Cha, 2007). There are many measures of similarity and selecting the right one is one of the challenges encountered by researchers. Depending on the application, some of the similarity measures do not always have optimal behavior (Shirkhorshidi et al., 2015). In this paper, we studied the capability of different similarity measures to effectively recognize specific indication patterns in production data. Furthermore, we introduce a concept for visual analysis and interactive pattern search in ultrasonic images of aluminum ingots. To do so, we aim to help domain experts to identify specific patterns in production data which may indicate critical process deviations. Lastly, we evaluate the benefit of interactive pattern search in ultrasonic images of aluminum ingots.

2 RELATED WORK

In this section, we analyze and discuss relevant work conducted in the research areas of visual analysis for industrial application and similarity measures.

2.1 Visual Analysis for Industrial Application

The trend of digitalization in the industry (so-called Industry 4.0, or also, smart production) generates large amounts of production data. In this scenario, domain experts are often overwhelmed by the amount of data and unable to obtain useful information that could help them to analyze their production processes. Visual data analysis, in which users interact with data to explore and analyze it, using visual displays, has been proven to be an effective approach for gaining insight from production data (Lee et al., 2014) (Wu et al., 2018).

A growing number of visualization solutions targeting production scenarios have been presented in recent years (Matkovic et al., 2002), (Jo et al., 2014), (Xu et al., 2016). Recently, a survey on visualization and visual analysis applications for smart manufacturing has been published (Zhou et al., 2019). The survey provides an overview of several studies conducted for industrial applications, with a few examples available for smart manufacturing applications in the iron and steel industry. In an early study (Wu, 2001), the problem of metal ingot casting and production planning is presented. This work reports that visualization of the production schedule provides the basis for interactive decision support. Zhou et al. (Zhou et al., 2016) proposed the integration of advanced simulation and visualization for the manufacturing process addressing issues on energy, environment, productivity, safety, and quality in the steel industry.

The increasing amount of data becoming available has to date triggered the use of visualization and development of visual data analysis tools in a variety of industrial domains (Liu et al., 2014). However, still many manufacturing systems are not ready to allow production specialists to efficiently and effectively analyze growing amounts of production data, also due to a lack of analytics tools and interfaces (Lee et al., 2013).

2.2 Similarity Measures

Manufacturing and various areas are becoming increasingly more data-driven, which increases the necessity of identifying the similarity between datasets. Datasets have several representations such as scalar values, vectors, or matrices. Mathematically, there are many measures of similarity or dissimilarity between the different forms of these datasets (Deza and Deza, 2006). This work is partly inspired by similarity techniques used for the comparison of distributions in image processing and computer vision. Several works have supported the visual retrieval and exploration of large numbers of scatter plot images. In (Scherer et al., 2011), feature vectors based on correlation coefficients are proposed to rank and cluster scatter plots for comparison. In (Behrisch et al., 2014), a relevance-feedback approach was proposed to learn to distinguish scatter plots of interest to specific users and tasks. In addition, in (Shao et al., 2016) an approach to describe scatter plots by the set of local patterns occurring was introduced and applied to filter interesting scatter plots from a larger number of plots. Recently, Bazan et al. (Bazan et al., 2019) present research work on a qualitative analysis of the similarity measures most used in the literature and the Earth Mover's Distance. In (Hernández-Rivera et al., 2017) it is demonstrated how similarity metrics can be used to quantify differences between sets of diffraction patterns. Although there exist many well-known similarity metrics, still a selection of metrics to measure the similarity between two distributions is crucial, because depending on the application, they do not always have optimal behavior. In the following, the discussion about qualitative analysis of different similarity measures on the problem of grouping similar batches and ingots will be presented.

3 INTERACTIVE PATTERN SEARCH

Our concept supports an interactive pattern search in ultrasonic images of aluminum ingots. The concept contains several steps which will be explained further in the following sections.

3.1 Data Preprocessing

To assess the quality, experts perform ultrasonic tests on rolled aluminum plates to meet the high-quality standards in the final products. Ultrasonic testing (UT) is used to locate the position and size of indications on rolled aluminum plates. The explanation and eventual reduction of these indications is a key priority in production process analysis. The dataset was obtained from ultrasonic tests conducted on cut aluminum plates with different length, width, and thickness from the cast and rolled ingots. One of the biggest challenges is to match the indications, detected on the final plates, to the ingot length. The data preparation was done with Pandas, one of the main tools used by data analysts in the programming language Python (McKinney, 2012). In the first steps, the tasks of data reduction, cleaning and transformation were performed, i.e., selection of relevant data; handling incomplete data, missing values and outliers; removing duplicates and recalculating values from the final plates back to the original ingot. Regarding the measuring unit, we note that these are the tons of tested material as an example. A single cast ingot with dimension $450 \times 1400 \times 7000$ mm weighs around 12.3 tons.

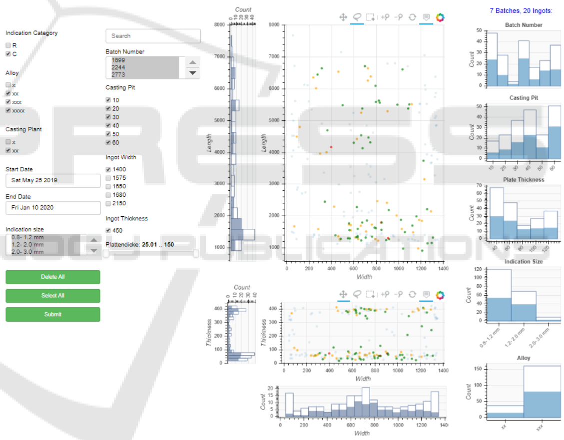

Figure 1: Interactive visual data analysis solution ADAM. Ingot visualization is composed of scatter plots and frequency histograms, showing the front and the top view of the ingot. The figures showing the front view of the ingot were used as UT images in the user study.

3.2 Visual Data Analysis Design ADAM

Our interactive visual data analysis solution ADAM, an acronym for Aluminum production Data Analysis and Monitoring, is based on the ideas presented in (Jekic et al., 2019). For more details, we refer the reader to this publication.

For the analysis tool ADAM and for the procedure of an interactive pattern search, which will be explained in the following section, the data preparation was a major challenge. Data from a wide variety of data sources had to be extracted and combined, such as the process data from various melting and casting furnaces, quality data such as the chemical composition of the batches, the input material and the UT test results, together with the material due to technical requirement that is generated in various process steps along the production process. The UT indications, that were detected in the individual final plates, together with the amount of the technical requirement, then had to be calculated back to the exact ingot position to make a comparison with, among other things, the process data that were recorded and/or calculated back to the casting length. The analysis tool ADAM enables the user to view the exact position of the indications in the ingot, that were detected in the final plates. In the next development steps, an interactive pattern search should enable the user to clearly display a part of the large amount of data recorded during the casting process, where several hundred parameters are recorded in high resolution, for a large number of ingots, which are produced. In the end, a smaller, clearly visible group of similar ingots should be proposed to the user, who can then compare the associated casting process data and thus could identify possible influencing parameters. This aspect will be dealt with in more detail in the following subsection.

The design of ADAM is shown in Figure 1. A set of tightly linked views of production parameters with cross-filtering capability supports the inspection of factors possibly influencing the product quality. Our approach was designed in an iterative development cycle guided by domain requirements obtained from a team of production experts. ADAM was developed using Bokeh, a Python software library (Bokeh Development Team, 2018). Two scatter plots for visualization were selected, showing the front and the top view of the ingot. Further, the figures showing the front view of the ingot were used in the user study. Color-coded circles (yellow, orange, and red) in the scatter plots represent the values of indications with specific diameters. ADAM is successfully integrated into the aluminum producer's system landscape and used by domain experts several times a week for data exploration and internal technical reporting. Our domain experts determined, using the visual analysis tool ADAM for several months, target use cases shown in Figure 2.

Future extension of ADAM will support automated data exploration tasks by automatically suggesting to the user similar batches and ingots of interest, and batches and ingots with atypical distributions. Grouping batches and ingots with similar patterns are important to investigate parameters that are possibly influencing the quality.

3.3 Concept for Interactive Pattern Search

The procedure for an interactive pattern search can be divided into two main steps and some possible additional steps.

The first step includes a selection of reference ingots and batches. Ingots and batches with interesting/atypical distribution of indications should be selected automatically and suggested to the user. Therefore, a standard distribution should be defined and ingots/batches selected for which the corresponding distribution of indications differs greatly from the standard distribution.

The second step is the selection of similar ingots and batches. Ingot and batches should be automatically selected, where the distribution of the indications is similar compared to the selected reference in the first step. Figure 3 shows the reference ingot and the automatically selected ingots with a similar distribution. Furthermore, it should be possible to automatically find ingots and batches with some specific predefined patterns as in case of accumulation of indications at a specific location in the ingot (see Figure 2 the reference 431630). Different methods, which can be used for similarity search, are described and evaluated in the next sections.

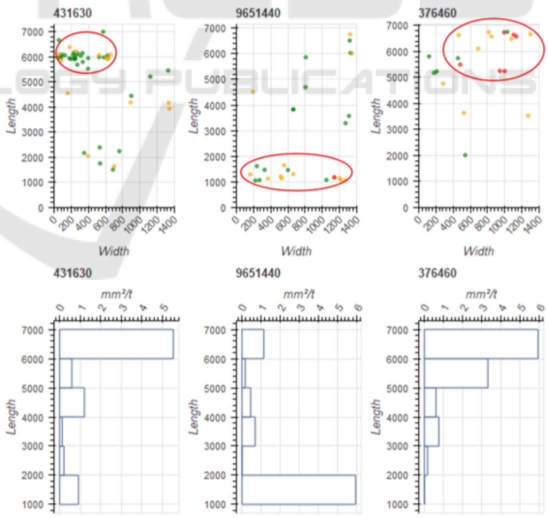

Figure 2: Three visual prominent patterns(from left to right): 1) group of indications at specific location in the ingot, 2) group of indications at the beginning of the casting, 3) group of indications at the end of the casting.

The third step includes the selection of conspicuous signals in the process data. In this part, we consider the process data corresponding to ingots/batches selected in the second step and compare it to the pro-

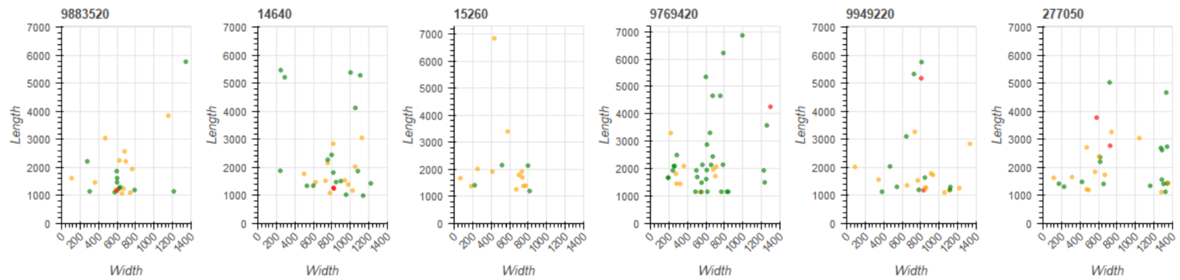

Figure 3: The first image is the reference image of aluminum ingot and the other five are the most similar images ranked by distance measure.

cess data corresponding to other ingots/batches, to detect patterns that are influencing the product quality.

Additionally, a further step in which the user can label patterns would be possible. The user should be able to add new patterns of interest to a list of predefined patterns (for example accumulation of indications) and export it in a report. To implement the three main steps, it is necessary to find good methods for the similarity search and to construct a suitable target value for the similarity search.

3.4 Definition of the Target Value: Calculation of Different Quality Criteria per Bin

In the following, only ingots with a thickness of 450 mm are taken into account, which represents the largest quantity of all ingots produced. For the calculation, we only consider indications between 1000 mm and 7000 mm. After recalculation of the position of the indications in the ingot, the proportion of the not tested area in the ingot depends on the plate thickness of the final product. For the greater number of ingots, the area between 0 and 1000 mm is not tested, or only partly tested.

Non-weighted Bins. During several workshops and with feedback from domain experts, we defined for each bin width 1000 mm. According to that ingot was divided into 6 parts (bin 1 to 6) and for each bin a quality criterion was calculated, given as:

$$quality\ criteria = \frac{indication\ area[mm^2]\ per\ bin}{tested\ material[t]\ per\ bin} \qquad (1)$$

The calculated area considered different plate thickness groups. The result is a quality criterion for each bin and ingot and also for each thickness group.

Weighted Bins. In many cases, it is hard to identify to which bin the indication in the ingot belongs, e.g., if there is an accumulation of indications at the boundary of two bins. Therefore we use a smoothing procedure, which also takes into account the indications in the neighboring bin. Similar to the previous discussion, we again consider 6 bins. The position of the center of bin $B_i, i = 1,\ldots,6$ is at $y_i = 1000 + i \cdot \frac{6000}{7}$ and the bin-width for each bin is $\frac{6000}{7}$. For each bin $B_i, i = 1,\ldots,6$, the indications at casting-length $l \in (y_i - \frac{6000}{7}, y_i + \frac{6000}{7})$ are weighted according to the weight-function $g_i(l) = 1 - 6 \cdot (\frac{7 \cdot |l - y_i|}{6000})^2 + 8 * (\frac{7 \cdot |l - y_i|}{6000})^3 - 3 * (\frac{7 \cdot |l - y_i|}{6000})^4$ and the quality criterion per bin is calculated similar to the previous calculation.

3.5 Analysis of Similarity Measures of Distributions

There are many measures of similarity that, depending on the application, do not always have optimal behavior. In many different application domains, there are several ways to define the nearness between distributions. A distance is defined as a quantitative measurement of how far apart two entities are. The similarity and the dissimilarity represent, respectively, how alike or how different two distributions are. If distributions are close, they will have high similarity and if distributions are far, they have a low similarity. To consider the similarity between ingots, in our case represented as scatter plots, we calculated 1-d histograms of the indications the length of the ingot and then compared histograms based on their distance. The smaller the distance between the histograms, the higher the similarity of the scatterplots. There are quite a few ways to apply distance metrics to compare histograms. We tested six different and popular distance measures: Euclidean distance, Manhattan distance, Chebyshev distance, Cosine similarity, Correlation distance and Bray-Curtis distance (Cha, 2007) (Deza and Deza, 2006). To assess the applicability of these measures in detecting similar patterns in ultrasonic images of aluminum ingots (in the further text: UT images) we have performed a user study where we have asked the domain experts to evaluate the results of our method. We used the results of this study to measure the accuracy of our approach.

4 USER STUDY AND RESULTS

In this section, we present a discussion of the first results achieved with our method, including examples to demonstrate the value of our findings. To demonstrate an evaluation of our approach, we conducted a study with four domain experts who represent the target user group. The reference result set and target use cases (see Figure 2) for our assessment were obtained by our domain experts using the visual analytics tool ADAM. Hence, to create the ground truth against which to evaluate, we use a set of queries capturing typical analysis tasks. We are specifically interested to evaluate the applicability of similarity measures with a specific focus on domain tasks to detect similar patterns in UT images of aluminum ingots.

4.1 Initial Results

For the pattern search in UT images of aluminum ingots, we consider data with the restrictions type of material x, casting plant y, thickness 450 mm, and width 1400 mm. Figure 2 shows visually prominent patterns regarding the casting length. Ingots should be automatically selected, where the distribution of the indications is very similar compared to a selected reference pattern. If the material is not free of indications, the standard behavior is that most of the indications appear at the beginning of the casting length and the number of indication decreases towards the end. During visual data exploration using ADAM, domain experts noticed that atypical distributions occur in some batches. Consequently, they wanted to find similar batches that appeared over time to analyze and link these cases with production parameters from the casting process. In the next section, we provide an evaluation based on capturing ground truth from domain users.

4.2 Data Labeling

Data labeling is important for many practical applications. To evaluate our approach we created three classes, established by the similarity of the figures to the reference image from 1 to 3 (similar, partially similar, not similar). To set the border values which are used for data labeling we performed the first part of the user study. As this is not a generated dataset, but a real production dataset, the labeling borders were verified by our domain experts. The dataset contains

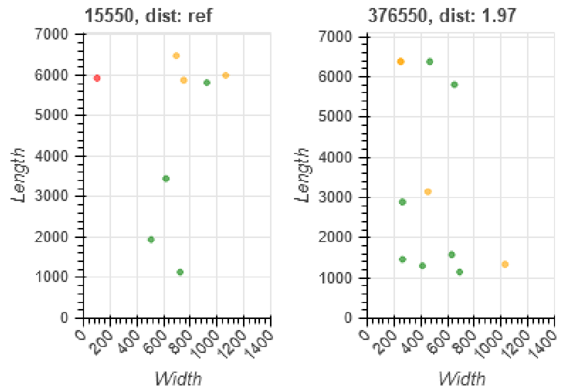

Figure 4: UT ingot images used in the first part of user study- data labeling.

1200 UT images of ingots. We calculated different distances measures for the complete dataset. A subset of 50 selected pairs of ingots were presented to the domain experts. Each of the 50 images contained a reference ingot compared against another ingot (Figure 4). Data samples were carefully selected based on different patterns discovered during the acquired experience in using the ADAM tool. It may happen that for one reference ingot multiple similar ingots exist as well as that only a few similar ingots exist. The results were compared with the different distance measures of the sample ingots. Finally, we managed to define the border values between similar, partially similar and not similar based on the results of the domain experts. The border values depend on the distance measure taken into account. Data labeling is a very difficult task, where domain experts may be uncertain about their answers. The problem with manual labeling is that the labels generated are usually subjective and can easily be biased towards the user's personal opinions. In our case, the accuracy between labeling among users i.e., interrater reliability was approximately 75%. In future work, we will consider other methods to improve labeling.

4.3 User Study

For each task, 30 figures representing UT images of ingots are given to test users. Our domain experts were asked to rank the similarity of the figures to the reference image from 1 to 3 (similar, partially similar, not similar). Results from experts are compared with results from interactive similarity search. The dataset containing 1200 UT images of ingots was considered for the evaluation of our approach. To compare the different methods (distance measures), UT images of ingots were ranked by similarity to one specific reference ingot using different distance measures and la-

Table 1: Accuracy of different distance measures for the reference ingot image 9883520.

Distance measures	Non-weighted	Weighted
Euclidean distance	0.77	0.63
Manhattan distance	0.73	0.63
Chebyshev distance	0.73	0.57
Cosine similarity	0.73	0.73
Correlation distance	0.73	0.73
Bray-Curtis distance	0.5	0.56

beled to the three classes (explained in the previous subsection). Additionally, for every task, we showed to the user a selection of ingots with similar patterns obtained from our method as in Figure 3.

In the first task, domain experts needed to rank UT images of ingots based on similarity regarding the reference UT image. The first reference UT image was 9883520 (Figure 3). In this case, the pattern contains indications at the beginning of casting length. The results of comparison with different binning extraction (non-weighted and weighted bins) are presented in Table 1. The ultimate goal of our method is the ability to predict the target class determined by the user. The highest accuracy in detecting similar patterns is achieved with Euclidean distance using the non-weighted quality criteria. However, we have a certain number of UT images classified not correctly (see Figure5). A confusion matrix was used for the purpose of distinguishing between different classes and see which classes actually cause confusion. What looks promising is that for the similarity measure with the highest accuracy i.e. euclidean distance there is only one UT ingot image belonging to class similar, classified as not similar class. Considering different methods for class labeling in future work could improve accuracy.

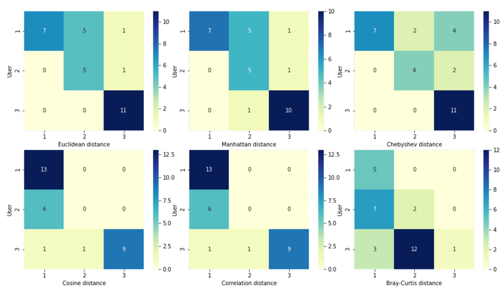

Figure 5: Confusion matrix for non-weighted bins.

Table 2 shows the accuracy of the results for the second reference UT image of the ingot. The pattern in reference 431630 (Figure 2) contains an accumulation of indications in a small area of the ingot.

The highest accuracy in detecting similar patterns is achieved with Chebyshev distance using the weighted quality criteria. Euclidean distance, with the highest accuracy in the previous task, and Chebyshev distance come from the same family of distances(Minkowski distance family). We have earlier discussed that if an accumulation of indications is exactly between two bins, patterns may be recognized better using the weighted quality criteria. In this case, the accuracy is lower than the accuracy for the reference 9883520. One explanation for this is that in the case of an accumulation of indications in a small area, the user sometimes can not estimate the actual amount of indications, while the distance measure calculation highly depends on the quality criteria per bin. Another explanation could be that in this specific case, the border of labeled classes should be considered further.

In the end, we showed to our domain experts UT images ranked by distance measure as in Figure 3. Domain experts report the five UT images ranked by the Minkowski distance family, showed good results for both reference ingots. Domain experts agree that the use of such methods could allow them to gather information to identify specific patterns, distribution of indication, and correlation with process data. An important requirement for analyzing ultrasonic test data is an extensive professional and domain-specific knowledge of users. However, we need to be careful as the similarity is subjective and is highly dependent on the domain and application. The user study showed promising results in detecting specific ultrasonic patterns. Regarding the similarity measures, the Minkowski distance family provided the highest recognition rates among the distances used in this work.

Table 2: Accuracy of different distance measures for the reference ingot image 431630.

Distance measures	Non-weighted	Weighted
Euclidean distance	0.6	0.63
Manhattan distance	0.63	0.63
Chebyshev distance	0.6	0.67
Cosine similarity	0.57	0.6
Correlation distance	0.6	0.57
Bray-Curtis distance	0.27	0.23

5 SUMMARY AND FUTURE WORK

Selecting the right distance measure is one of the challenges encountered by professionals and researchers when attempting to apply different methods in real-world applications. The variety of similarity measures can cause confusion and difficulties in choosing a suitable measure. The performance of similarity measures may vary depending on different datasets. In this paper, we studied a quantitative comparison for different similarity measures on UT images of ingots. The aim of this study was to clarify which similarity measures are more appropriate and applicable when searching for specific ultrasonic patterns. Further, we conducted interviews with domain experts in the analysis of UT indications images comparison and used this feedback to define a ground truth for our evaluation. We provided a discussion and demonstrated the possible insights enabled by our approach and its potential to support production data exploration.

Future work includes investigation of process data corresponding to groups of similar ingots and batches, and potentially discovering key influential parameters in the process data. As future work, we also want to include advanced multidimensional data visualizations, to support pattern detection and parameter correlation. Furthermore, automatic classification of certain quality patterns, based on interactively provided expert examples, is considered an interesting future extension of an existing visual analytics solution.

ACKNOWLEDGEMENTS

This research work is done by Pro2Future and AMAG Austria Metall AG. Pro2Future is funded within the Austrian COMET Program-Competence Centers for Excellent Technologies under the auspices of the Austrian Federal Ministry of Transport, Innovation and Technology, the Austrian Federal Ministry for Digital and Economic Affairs and of the Provinces of Upper Austria and Styria. COMET is managed by the Austrian Research Promotion Agency FFG.

REFERENCES

Bazan, E., Dokládal, P., and Dokladalova, E. (2019). Quantitative analysis of similarity measures of distributions.

Behrisch, M., Korkmaz, F., Shao, L., and Schreck, T. (2014). Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 43–52. IEEE.

Bokeh Development Team (2018). *Bokeh: Python library for interactive visualization*.

Cha, S.-H. (2007). Comprehensive survey on distance/similarity measures between probability density functions. *City*, 1(2):1.

Deza, M.-M. and Deza, E. (2006). *Dictionary of distances*. Elsevier.

Hernández-Rivera, E., Coleman, S. P., and Tschopp, M. A. (2017). Using similarity metrics to quantify differences in high-throughput data sets: application to x-ray diffraction patterns. *ACS combinatorial science*, 19(1):25–36.

Jekic, N., Mutlu, B., Faschang, M., Neubert, S., Thalmann, S., and Schreck, T. (2019). Visual analysis of aluminum production data with tightly linked views. In *EuroVis (Posters)*, pages 49–51.

Jo, J., Huh, J., Park, J., Kim, B., and Seo, J. (2014). Livegantt: Interactively visualizing a large manufacturing schedule. *IEEE transactions on visualization and computer graphics*, 20(12):2329–2338.

Lee, J., Kao, H.-A., Yang, S., et al. (2014). Service innovation and smart analytics for industry 4.0 and big data environment. *Procedia Cirp*, 16(1):3–8.

Lee, J., Lapira, E., Bagheri, B., and Kao, H.-a. (2013). Recent advances and trends in predictive manufacturing systems in big data environment. *Manufacturing letters*, 1(1):38–41.

Liu, S., Cui, W., Wu, Y., and Liu, M. (2014). A survey on information visualization: recent advances and challenges. *The Visual Computer*, 30(12):1373–1393.

Matkovic, K., Hauser, H., Sainitzer, R., and Groller, M. E. (2002). Process visualization with levels of detail. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002.*, pages 67–70. IEEE.

McKinney, W. (2012). *Python for data analysis: Data wrangling with Pandas, NumPy, and IPython.* " O'Reilly Media, Inc.".

Scherer, M., Bernard, J., and Schreck, T. (2011). Retrieval and exploratory search in multivariate research data repositories using regressional features. In *Proceedings of the 11th annual international ACM/IEEE joint conference on Digital libraries*, pages 363–372.

Shao, L., Schleicher, T., Behrisch, M., Schreck, T., Sipiran, I., and Keim, D. A. (2016). Guiding the exploration of scatter plot data using motif-based interest measures. *Journal of Visual Languages & Computing*, 36:1–12.

Shirkhorshidi, A. S., Aghabozorgi, S., and Wah, T. Y. (2015). A comparison study on similarity and dissimilarity measures in clustering continuous data. *PloS one*, 10(12):e0144059.

Soban, D., Thornhill, D., Salunkhe, S., and Long, A. (2016). Visual analytics as an enabler for manufacturing process decision-making. *Procedia Cirp*, 56:209–214.

Sun, D., Huang, R., Chen, Y., Wang, Y., Zeng, J., Yuan, M., Pong, T.-C., and Qu, H. (2019). Planningvis: A visual analytics approach to production planning in smart factories. *IEEE transactions on visualization and computer graphics*, 26(1):579–589.

Suschnigg, J., Ziessler, F., Brillinger, M., Vukovic, M., Mangler, J., Schreck, T., and Thalmann, S. (2020). Industrial production process improvement by a process engine visual analytics dashboard. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*, pages 1320–1329.

Thalmann, S., Mangler, J., Schreck, T., Huemer, C., Streit, M., Pauker, F., Weichhart, G., Schulte, S., Kittl, C., Pollak, C., et al. (2018). Data analytics for industrial process improvement a vision paper. In *2018 IEEE 20th Conference on Business Informatics (CBI)*, volume 2, pages 92–96. IEEE.

Vasudevan, A. K. and Doherty, R. D. (2012). *Aluminum Alloys–Contemporary Research and Applications: Contemporary Research and Applications*. Elsevier.

Wu, P. Y. (2001). Visualizing capacity and load in production planning. In *Proceedings Fifth International Conference on Information Visualisation*, pages 357–360. IEEE.

Wu, W., Zheng, Y., Chen, K., Wang, X., and Cao, N. (2018). A visual analytics approach for equipment condition monitoring in smart factories of process industry. In *2018 IEEE Pacific Visualization Symposium (PacificVis)*, pages 140–149. IEEE.

Xu, P., Mei, H., Ren, L., and Chen, W. (2016). Vidx: Visual diagnostics of assembly line performance in smart factories. *IEEE transactions on visualization and computer graphics*, 23(1):291–300.

Yin, S. and Kaynak, O. (2015). Big data for modern industry: challenges and trends [point of view]. *Proceedings of the IEEE*, 103(2):143–146.

Zhou, C., Wang, J., Tang, G., Moreland, J., Fu, D., and Wu, B. (2016). Integration of advanced simulation and visualization for manufacturing process optimization. *Jom*, 68(5):1363–1369.

Zhou, F., Lin, X., Liu, C., Zhao, Y., Xu, P., Ren, L., Xue, T., and Ren, L. (2019). A survey of visualization for smart manufacturing. *Journal of Visualization*, 22(2):419–435.