



ECS: an interactive tool for data quality assurance

Christian Sieberichs¹ · Simon Geerkens¹ · Alexander Braun¹ · Thomas Waschulzik²

Published online: 8 January 2024

© The Author(s), under exclusive licence to Springer Nature Switzerland AG, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved 2024

Abstract

With the increasing capabilities of machine learning systems and their potential use in safety-critical systems, ensuring high-quality data is becoming increasingly important. In this paper, we present a novel approach for the assurance of data quality. For this purpose, the mathematical basics are first discussed and the approach is presented using multiple examples. This results in the detection of data points with potentially harmful properties for the use in safety-critical systems.

Keywords Equivalence class sets · Data quality assurance · Distance based · Data visualization

1 Introduction

The development of machine learning (ML)-based systems has led to a widespread use in research, industry as well as in the everyday life. Even though ML systems show great performance in solving complex tasks, their use is mostly limited to domains, where wrong decisions only have minor consequences. The application of ML systems in high-risk domains currently is problematic due to the needed quality, lack of trustworthiness and the expected legal basis. To give a legal framework for the application of ML systems, the European AI act [4] is at the moment under development. Simultaneously, multiple projects from research and industry are dealing with the topic of ML systems in high risk areas, such like "KI-Absicherung" [17] and "safetrAIIn" [25]. All of those projects highlight the high requirements that are needed to protect humans from errors made by ML systems. High-risk ML systems have to fulfil the requirements according

to European Commission [4] Chapter 2 "REQUIREMENTS FOR HIGH-RISK AI SYSTEMS" Article 10 "Data and data governance" Point 3: "Training, validation and testing data sets shall be relevant, representative, free of errors and complete". In this paper we introduce a new approach that will contribute to the future fulfilment of this requirement. It is showcased how different relevant aspects of the data can be analysed and how relations between the given data can be used for quality assurance aspects.

The presented approach is part of the QUEEN-method (**Q**ualitäts**g**esicherte **e**ffiziente **E**ntwicklung vorwärtsgerichteter künstlicher **N**euronaler Netze, *quality-assured efficient development of neural networks*) [28] which is a comprehensive approach for the development of quality assured neural networks. In the scope of the QUEEN-method, two data quality assurance methods were developed, namely Q^2 (integrated quality indicator) [7] and ECS (equivalent classes sets) [24]. These methods were developed simultaneous in close cooperation. In this paper, we want to show the mathematical basis and the use of ECS on the topic of quality assurance. The abilities and usage of the Q^2 is covered in another submission [8].

The ECS is particularly used to analyse the local and global composition of data sets. Based on this, a wide variate of data quality properties is addressed. Be it the identification of single data points like outliers, false annotations or isolated data or the identification of groups of data points like decision boundaries and local data point groups of identical output. The ECS makes it possible to identify all data points which do not match specifiable conditions. The method itself is thereby created in such a way that

✉ Christian Sieberichs
christian.sieberichs@hs-duesseldorf.de

Simon Geerkens
simon.geerkens@hs-duesseldorf.de

Alexander Braun
alexander.braun@hs-duesseldorf.de

Thomas Waschulzik
thomas.waschulzik@siemens.com

¹ University of Applied Sciences Düsseldorf,
Düsseldorf 40476, Germany

² Siemens Mobility GmbH, Erlangen 91058, Germany

interactions between the user and the data are supported to simplify and speed up the quality assurance process.

2 Related work/state of the art

Despite the fact that data quality and quality assurance are widely necessary and researched, there exists no single general accepted definition. Instead, there are several attempts to define data quality based on current developments. One example is given by Wang and Strong [27] who define data quality with respect to the intended use of the data. It is argued that data quality has to be a context-dependent term to be appropriately used in the context of a given tasks. In addition, the term of "data quality" is split into multiple properties like accuracy, consistency, completeness, safety and more. In Sidi et al. [23], many of these properties are listed and defined separately. In Pipino et al. [18], data quality is additionally split into subjective and objective assessments of data quality.

A general definition on data quality can thereby not be given. Instead is high data quality considered to be data which is fit for its intended purpose [5]. If data quality is used in standards, it is typically split into the different properties which have to be analysed separately like in Gualo et al. [9].

A first step to assure the data quality is the use of descriptive statistics [11]. Herein statistical methods are used to gain greater insights into the given data. Common methods are the visualization via scatter plots and histograms, often combined with the measurement of central tendencies, dispersion and location parameters. Our proposed method extends the descriptive statistical methods, enables the visualization of multiple quality assurance aspects in one plot and enables a direct interaction between the visualization of the quality indicator visualization and the data.

When trying to assure the quality of data, another possible approach is the representation of given data points in lower dimensional space using methods of dimensionality reduction. Commonly used methods are PCA [13], tSNE [15] or UMAP [16]. These methods often produce representations interpretable by humans if the output dimensionality is chosen to be low enough. However, such methods often result in considerable loss of information. The ECS on the other hand is computed on the original values and takes all the given information into account.

Some approaches try to cover as many dimensions of the data quality as possible. One way to do this is by testing the data against predefined rules and assumptions. An example of such an approach is the pointblank R package [12] which is created for an agent-based data quality assurance. In this package, specific elements of the data

are tested against predefined functions. As part of this, it can be tested if the data is greater, equal, lower and so on. Another method is given by DEEQU published in Schelter et al. [21] and Schelter et al. [22]. This package allows for assumption based unit tests which can be defined by the user. Tests on specific parameters of the data, similar to those already mentioned with regard to the pointblank package, are possible as well. A last method that should be mentioned here is shown in Heinrich et al. [10]. This approach showcases a probability-based method which calculates a value representing the probability that a data set is free of internal errors with respect to entered rules. The entered rules are based on the presence of data of certain values, comparable to the package pointblank. The main problem in using the mentioned approaches is the large amount of required knowledge about the data to create accurate assumptions. On top of this, the efficient creation of assumptions is only given if the user is aware that the data quality is influenced in some regards. Due to the reliance on the relationship between data points, our methods do not need any assumptions or rules that are to be specified by an user. Instead our approach can be used without any knowledge about the data.

A different approach is the focus on just a single dimension of the data quality. On the topic of outlier detection, these are for example density-based algorithms like Breunig et al. [2]. In this approach, the amount of local neighbouring data points is calculated and the thereby generated local density is compared with the nearest neighbours. Another approach is using the DBSCAN algorithm [3] to cluster the given data. Based on this clustering the method proposed by Thang and Kim [26] calculates values to identify clusters of minimal sizes. These clusters are then regarded as possible anomalies. Another data quality property is the detection of possible outliers, which can also be solved by density based clustering. One example of such an algorithm is given by Fawzy et al. [6]. This algorithm uses a fixed clustering to identify clusters followed by the computation of cluster distances. The clusters are classified as anomalous based on the inter-cluster distances and the deviation from the mean inter-cluster distance. Two quite similar approaches are Samara et al. [19] and Samara et al. [20]. Both approaches use a clustering of the given data in a first step. The first one uses the previously mentioned DBSCAN, the second one uses a cluster algorithm named OPTICS [1]. In a second step, anomalous clusters are identified, once based on inverse distance weighting (IDW) and once using the kringing method. The main advantage of all of the mentioned methods is the reliable calculation of their data quality property. However, due to the methods focus on one specific data quality property, they are only useful if the assumption exists that this property could contain errors. The advantage of our

proposed method is that multiple data quality properties may be analysed with one approach.

3 Method

The ECS is based on the idea that a data set can be split into input data and output data. The input data defines the dimensions of the data, henceforth are called features, which can be used to predict the output features. The amount of all possible inputs creates the input space I . Accordingly is the output space O created by all possible outputs. To use the ECS properly, all feature values must be numbers. Features which are not created by numbers have to be represented in some way as a number or a combination of numbers.

To start the calculation of the ECS two metrics are needed. These metrics should be chosen in such a way that "similar" data points according to the semantics of the task that has to be solved have a relatively small distance to each other. At the same time, "dissimilar" data points should have a relatively large distance. The distances between two data points can be calculated in the input space and in the output space independent from each other. By doing so, it is possible to use different metrics for the distances in I and in O . Which metric is best suited for the data set depends on the given type of data and the task to be solved. In the following, the difference between data points in the input space is named input distance d_{RI} . Accordingly, the difference between data points in the output space is called output distance d_{RO} .

To differentiate between "similar" and "dissimilar" data points, the distances can be separated into different groups. The minimal approach is to create two groups. One group for relatively small distances and another one for relatively large distances. Doing so requires a threshold, which is called δ_{in} for distances in the input space and δ_{out} for distances in the output space. These δ can be absolute distance values or a percentage of the maximum known distance between data points. They are set based on the data quality properties that should be identified and the used data type. By comparing two data points with each other, four possible scenarios can be distinguished:

- Small input distance–small output distance
- Small input distance–large output distance
- Large input distance–small output distance
- Large input distance–large output distance

Each of these scenarios shows a relation between the data points. If for example the distances are both small, than the data points may showcase a common use case with a typical output. A small input distance in combination with

a large output distance on the other hand could showcase complex areas in the input space or an outlier. Either way, the identification of data properties based on two data points is not enough. Due to this, the following four ECS-sets are calculated. In these sets, the compared data points are saved, which are part of one of the above scenarios (Table 1).

$$ECS_{EE}(D) := \{d_c | d_c \in D^2 \wedge d_{RI}(B) \leq \delta_{in} \wedge d_{RA}(B) \leq \delta_{out}\} \tag{1}$$

$$ECS_{EU}(D) := \{d_c | d_c \in D^2 \wedge d_{RI}(B) \leq \delta_{in} \wedge d_{RA}(B) > \delta_{out}\} \tag{2}$$

$$ECS_{UE}(D) := \{d_c | d_c \in D^2 \wedge d_{RI}(B) > \delta_{in} \wedge d_{RA}(B) \leq \delta_{out}\} \tag{3}$$

$$ECS_{UU}(D) := \{d_c | d_c \in D^2 \wedge d_{RI}(B) > \delta_{in} \wedge d_{RA}(B) > \delta_{out}\} \tag{4}$$

Each of the four ECS sets represents all comparisons between data points which result in one of the four scenarios. Thereby, an E showcases a small distance whereas a U showcases a large one. The first of the two letters of the ECS-sets represents the input distance and the second represents the output distance. Following this, ECS_{EU} contains all data point comparisons which result in a small input and a large output distance. The ECS_{UE} on the other hand contains comparisons which result in a large input and a small output distance.

The information of the ECS-sets can be used to analyse the data points for each of the four scenarios. This way, it is possible to identify data points with specific properties. It would, for example, be possible to identify all data points which have the many dissimilar data points in close proximity. It would also be possible to identify data points which showcase small distances in the input and the output space. By doing so, certain areas of the input space can be identified which correlate with a certain outputs of the output space. It could also be possible to identify features which differentiate certain data points from each other.

The ECS-sets contain all information of the data set which could be used for quality assurance. However, the

Table 1 parameters which are needed to calculate the ECS-sets

d_{RI}	Distance in I
d_{RA}	Distance in O
δ_{in}	Inputdelta
δ_{out}	Outputdelta
D	Dataset
D^2	All possible comparisons of data points out of D
d_1, d_2	Two data points out of D
d_c	Combination of two data points (d_1, d_2)

formatting of the sets is difficult for humans to read. This is especially the case when entire data sets should be analysed and not just a small subset of the data. The solution is the comprehensive representation of the ECS sets in such a way that interesting data points can easily be identified. Before this can be done, it has to be determined which combinations of data points are the most interesting ones. The expectation hereby would be that similar input data would create output data that is related in some way. Based on this, it can be assumed that a combination of data points with a small input distance also has a small output distance. On the other hand, it would not be expected that data points with large input distances to each other would showcase similar output data. The most interesting combinations of data points would thereby be combinations which result in a small input distance. This comparisons can be display particularly by sorting the data point comparisons based on the input distance. An example of the sorted representation of the *ECS_EE* is shown on the right side of the Fig. 1. Listed on the x-axis is the comparison between data points. This comparison is data point based and showcases the comparisons of any data point with the k th smallest distance in the input space. On the y-axis, it is displayed how many of these comparisons are part of the current ECS set, which is in this case the *ECS_EE*. In this process, functions are created representing every data point. To showcase the entire data set these functions are superimposed over each other. Every function visually displays if and which of the nearest data points are part of the *ECS_EE*. The data set which was used to create the displayed *ECS_EE* is shown on the left side of Fig. 1. It is a simple data set created by two input features (a, b) and one output feature (colour and shape). Increasing functions display that most of the data points with the k th smallest distance are part of the current ECS-set. Functions which do not increase display that the comparisons are part of another




Fig. 1 Left: data set created by two input features and one output feature. The data points are located in clusters which can be separated from each other. The cluster in the bottom left and the top right corner are showcasing the same output. Right: *ECS_EE* of the displayed data set. Noticeable are steep and early increases which shows the existence of nearby data point with small input and output distance

ECS-set. It should be emphasized here that the function in the k th position only increases for one of the four ECS-sets.

The created ECS histograms consist of a large number of functions. Areas in the ECS histogram, in which large amounts of functions showcase the same behaviour, are displayed darker. Accordingly smaller amounts of functions are display brighter. For the representation of the amount of functions, gamma correction is used. This way, even singular functions should stay visible.

As stated before, it would be expected that a small input distance influences the output distance. An ideal data set would have a strong correlation between the position in the input and the output space. The resulting data point combinations would just have small input and output combinations for all the nearest neighbouring data points. This would result in a steep increase of all functions in the *ECS_EE* until all possible similar data points are combined with each other. From this point on the functions in the *ECS_EE* do not increase any further. The *ECS_EE* function created by a single data point in such an ideal data set is shown schematically in the Fig. 2. The main diagonal is thereby displaying the maximum speed at which a function is able to increase. Data sets or individual data points which are not ideal do create different functions. One extreme example would be that a function would not increase at all in the *ECS_EE*. This would be the case because there are no possible combinations with small input distance, small output distance or both. Which of these possibilities is actually the case, can be tested by using the other ECS histograms.

The benefit of the representation of the information of the data set in the form of the ECS histogram is the data point based presented information. The neighbouring data points to each data point and there relation to each other are shown and can be compared to set expectations. This basis helps at identifying functions which do not live up to the expectations fairly easy. The reason why the functions are not behaving the way they should, is intrinsically given by the combination of the behaviour and the ECS set. Additionally,




Fig. 2 Schematic *ECS_EE* created by a data point in an ideal data set




Fig. 3 Created data set with two input features, one output feature and 1000 data points. The data points are located in clusters of different density and with different amounts of data points

it would be possible to define limits in the ECS histogram which can be tested autonomously.

4 Application

We want to showcase the abilities of the ECS by the application on two different data sets. Simultaneously we want to show how the ECS can be used explicitly to detect certain data quality properties. With this in mind we created a data set which is used as an example. The data set is created in such a way that properties detected by the ECS can be verified by displaying the critical data points. In addition is the ECS used on the MNIST data [14] set to display the detection of data quality properties on a commonly known example. The application on the data sets is focused on the data quality properties created by outliers, isolated data points and local groups of data points with identical output values.

4.1 ECS on point cloud

To demonstrate the usage of the ECS, an artificial data set is created which is displayed at the left side of Fig. 3. This data set is similar to the one displayed in Fig. 1. The most important difference is, that the clustered data points can not clearly be separated from each due to the clusters overlapping partially. The data set contains 1000 data points which are grouped in four clusters. As in Fig. 1, each cluster has a different amount and a different density of data points.

The data set was created this way, because it demonstrates a simple classification task. At the same time, properties like outlier and local groups with identical output are present and can be visualized. In the following, it is shown how these properties can be identified by using the ECS. All the ECS histograms are created using a δ_{in} of 0.3 times the maximum distance in the input space and a δ_{out} of 0 to differentiate between all differing outputs.

4.1.1 Outliers

An outlier is a data point which has an unexpected output for the given input. This output is typically very different from an output that would be expected. Here, an outlier is just be considered to vary in the output space. Unwanted variations in the input space are treated in the following section. The reasons for an outlier can be different. The output may for example be wrong or the data point showcases a rare but correct input.

Due to their character, outlier appear in areas which are dominated by data points with a different output. Given this information, it can be stated that an outlier has close neighbours with large output distances. The *ECS_EU* is used to identify these cases. Functions in the *ECS_EU* are increasing if there are data point combinations with small input and large output distances. Such functions that already increase for the nearest neighbours can thus be regarded as outliers. By targeting these functions, the corresponding outliers can be identified. How many combinations for how many nearest neighbours should be part of the *ECS_EU* is dependent on the given data set. In the given point cloud example, 100 nearest neighbours were chosen to be enough to represent the local data points. If out of these 100 combinations more than 70 have a large output distance, then the data point is regarded as an outlier.

The *ECS_EU* is shown in Fig. 4. The area of importance in which the functions of outliers appear is highlighted by a rectangle. It can be noticed that some data points in the point cloud are highlighted which means that they are considered to be outliers. It can also be noticed that most of the functions are not increasing by much.

4.1.2 Isolated data points

Isolated data points are data points which have a large input distance to many or all of there nearest neighbours. This means that the data point showcases an input which is rare




Fig. 4 Left: Data set with marked outlier. Right: *ECS_EU* of the displayed data set. The area is marked in which functions of outlier appear

or possibly wrong. In the literature, these type of data points are often referred to by "Out-of-distribution-data".

The ECS_{UE} and the ECS_{UU} are used to identify data points which have large distances to the nearest neighbours. In both ECS sets are combinations saved which have large input distances to each other. The difference between these two ECS histograms is the differently sized output distance, which is not considered for isolated data points. The corresponding functions of isolated data points increase very early in the ECS histograms. Most of the time, the functions increase in the ECS_{UE} as well as in the ECS_{UU} . This is the case, because the nearest neighbours themselves may have large input distances to each other and thereby showcase very different outputs. The sooner a function increases, the fewer data points are given in the local area of an isolated data point. The amount of neighbouring data points that should exist is on the given task and data set. Typically, this means that every data point should at least have a few neighbouring data points with small input distance. If many data points have no near neighbours, an adjustment of the parameter δ_{out} can be considered.

By using the ECS_{UE} and the ECS_{UU} , it can be stated that, there are no isolated data points in the current data set with less than 50 close neighbours. This can be confirmed by the fact that the sample data set used here was created with clusters of data points

4.1.3 Local groups of identical output

A local group with identical output is a structure created by multiple data points. All data points in such a group have small distances to each other regarding the input and the output distances. There are no greater amounts of data points which showcase a large output distance, besides possible outliers or false data points. The identification of these groups showcase the ability of the used metric to differentiate between different outputs on the basis of the corresponding input. This means that the input data of the groups share similar features which in turn leads to the differentiation. It would be possible to solve the given task at least for these groups based on these similar features.

The combination of small input distances and small output distances can be identified using the ECS_{EE} . The functions correlating with data points as part of a local group with identical output increase strongly. The functions will increase as long as there exist data points with small distances in the input and output space in the data set. These strongly increasing functions showcase every data point which is part of such a group of data points. Using the ECS_{EE} , there is the possibility to also identify groups of different amounts of data points. This can be done by choosing the function increasing the strongest for different amounts of neighbours. If the functions have increase up to

the chosen amount, it means that there is a minimum of this amount of data points in the group.

In the given case in Fig. 5, groups with 100 data points and identical output should be identified. The area of importance in the ECS_{EE} is marked by a rectangle in the upper right corner. It should be noticed that not just the function increasing the strongest were marked but also some functions which increase a little bit slower. This has been done to make the identified groups more robust against false data points and outliers. In the given case, this means that also functions with 95 out of 100 data points are regarded as local groups with identical output. In addition, it can be noticed that most of the functions in Fig. 5 are increasing. This indicates that there are many data points arranged in local groups. This is logical, because the point cloud was created this way as four groups of clustered data points. The detected data points which are part of a local group are marked on the left side of the Fig. 5.

4.2 ECS on MNIST

In contrast to the previous example, MNIST is a data set which was created to represent the specific task of classifying handwritten numbers. The data set consists of 60000 images of size 28×28 serving as the input and as many numbers between zero and nine for classifying the input.

The most important difference between the previous used point cloud and the MNIST data set is the amount of data and input features. The much greater amount leads to much more functions in the ECS histograms. The ECS histograms thereby get more complicated. This can be counteracted by applying more specific metrics to the data. Here, the pixel-wise euclidean distance is chosen as a metric. The euclidean distance is typically not used on images due to its bad performance. But in the case of MNIST, this metric is applicable, as the image pixels are given as centred grayscale values. It can be shown that the abilities of the ECS are still given using the euclidean metric. Another problem which appears




Fig. 5 Left: Data set in which local groups of identical output are highlighted. Right: ECS_{EE} of the data set with a marked area in the top right for the functions which represent data points in local groups with identical output




Fig. 6 Left: Representation of the MNIST data set with numbers which showcase the output value of the clusters. Right: *ECS_EU* which is used to identify outlier

by using data with many features is the curse of dimensionality through which all distances are getting closer to each other. As a result a larger δ_{in} of 0.75 times, the maximum distance in the input space is used in the following. The δ_{out} is still 0 to differentiate between all differing outputs.

To show the input data of the MNIST data set in a way, a representation is used in the following chapters. This representation is created by using UMAP [16], a dimensionality reduction method. The cluster, which are created this way are marked with a number to showcase the corresponding output. The ECS is used based on the original MNIST input and output data.

4.2.1 Outliers

As shown in the chapter "ECS on point cloud - outliers", the *ECS_EU* used to identify outliers. The *ECS_EU* of the MNIST data set for the nearest 200 neighbours is shown in Fig. 6. As mentioned before, this representation has much more functions. These are too dense for any singular function to be identified without an interaction. But it can be noticed that most functions do not increase by a lot, as indicated by the darker visualization. The amount of functions ($|F|$) which increase to a specific value of fulfilment (v_f) until the 200th neighbour are shown in the following Table 2.

It must be noticed that more than half of the given data points have a maximum of 10 combinations showcasing a small input distance and a large output distance. On the other hand, there are more than 6000 data points for which half of the nearest 200 neighbours have a large output distance. Not all of these are outliers, some may be positioned between classes others may have badly assigned distances. To identify outliers only functions performing worse than a random assignment of distances are used. This means that all data points having more than 180 data points with large output distance among the nearest 200 neighbours are interpreted as outliers. By choosing these functions 804 data points were identified as outliers. In Fig. 7, a random sample of nine of these outliers is displayed. It is noticeable that all of

Table 2 Amount of data point combinations which are part of the *ECS_UE* for the 200 nearest neighbours

v_f	$ F $
101–200	6021
51–100	7337
11–50	14914
0–10	31728
0	16813



Fig. 7 Detected outlier in the MNIST data set. The output of these data points is from left to right: 1,2,5,2,7,7,6,9,3

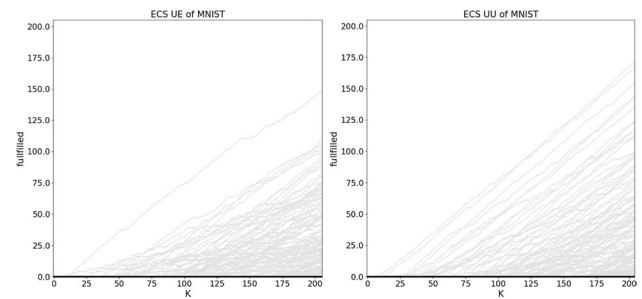


Fig. 8 *ECS_UE* and *ECS_UU* of the MNIST data set for the detection of isolated data points

these data points do look strange. Most may also be mistaken with a different number. It would for example be possible to remove these data points from the MNIST data set to achieve a higher data quality.

4.2.2 Isolated data points

The identification of isolated data points in MNIST is identical with the identification of data points in the point cloud. The *ECS_UE* and *ECS_UU* used are shown in Fig. 8. In the *ECS_UE* are 129 and in the *ECS_UU* 132 data points with less than 200 neighbouring data points. Most of the correlated data points appear in the *ECS_UE* as well as in the *ECS_UU*. Due to the relatively small amount of increasing functions, the histogram is created darker.

The earliest functions start increasing in the *ECS_UE* and *ECS_UU* for less than 10 neighbours. The input data of the earliest increasing functions is shown in Fig. 9. One function in the *ECS_UE* is noticeable do to its steep and early increase. The corresponding input data is shown in Fig. 9 in the third image from left. This data point has a large distance to its closest neighbours. At the same time, most of these neighbours have the same output "4". This lead to the conclusion that the data point has still some of the most important features which

are correlated with the output "4", even if the data point is very isolated. Overall, it is noticeable that most isolated data points shown use many input pixels to display the number. This is not often the case in the MNIST data set. In addition, the pixel-wise euclidean distance used reacts especially on pixel-wise differences by assigning higher distances in the input space.

4.2.3 Local groups of identical output

The *ECS_{EE}*, which is used for the identification of local groups of data points with identical output, is shown on the right side of Fig. 10 for the nearest 500 neighbours. As it is the case for the detection of outlier, the amount of functions is much larger than in the point cloud example. It is also not possible to identify single functions but instead overall trends of the functions. It can be noticed that most of the functions are increasing very steeply. This means that most data points do have small input as well as output distances in combination with their nearest neighbours. This in turn means that most data points are located in local groups with identical output. The amount of data points which should be part of the groups can be changed by using different amounts of neighbours in the *ECS_{EE}*. In Table 3, the amount of data points (*ldpl*) which belong to a local group of different size (*gs*) is shown. These amounts were created by allowing a maximum of 5 data points a different output which could exist due to outliers.

Noticeable are most of the data points located in groups of a few hundred data points. But still more than 4000 data points could be detected which are part of local groups with more than 1500 data points. The position of these data points is highlighted in dark in the UMAP representation in Fig. 10. It can be noticed, that especially data points with an output of 1, but also of 0 and 6 show local groups of identical output. This means, that the used metric has the ability to differentiate these data points from each other. The local groups which can be identified can then be used to solve the given task, based on their location in the input space.

5 Conclusion

In this paper, we presented a novel approach for the data quality assurance based on local similarities. It was shown how the ECS is calculated and can be used on an artificial example. The thereby presented procedure was used to detect data quality properties on the MNIST data set. Besides the possibility to detect outlier, isolated data points and local



Fig. 9 Detected isolated data points in the MNIST data set

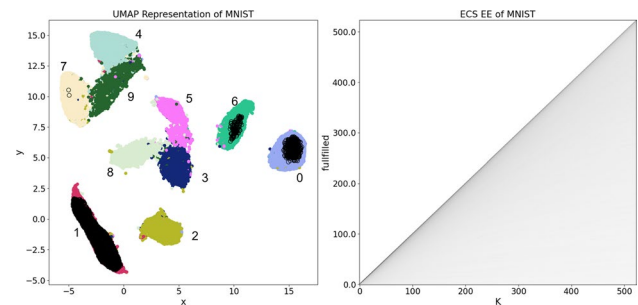


Fig. 10 Left: Representation of the MNIST data set with number which showcase the output value of the clusters. Data points which are part of a local group with 1500 data points or more are highlighted in dark. Right: *ECS_{EE}* for the MNIST data set for the nearest 500 neighbours

Table 3 Amount of data points which are part of the different sized local groups with identical output. The entire data set contains 60000 data points

<i>gs</i>	<i>ldpl</i>
100	38383
200	27351
500	14745
1000	7851
1500	4329

groups of similar output has the versatile applicability of the ECS been shown. ECS could also be used to validate quantitative data set requirements for data quality properties. These can state the minimum amount of elements per group, the amount of outliers or a maximum amount of local groups. Some of these properties, like the amount of accepted outliers in local groups, may be dependent from associated safety requirements and the required safety integrity level

Funding Open Access funding enabled and organized by Projekt DEAL. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 532148125 and supported by the central publication fund of Hochschule Düsseldorf University of Applied Sciences.

Data availability The MNIST dataset and the UMAP visualisation are publicly available. The code used to create the ECS sets and the visualisation of the ECS sets is the property of Siemens AG and may not be published.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will

need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ankerst, M., Breunig, M. M., Kriegel, H.-P., Sander, J.: OPTICS: Ordering Points to Identify the Clustering Structure. In: Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data, SIGMOD '99, 49–60. New York, NY, USA: Association for Computing Machinery. ISBN 1-58113-084-8. Event-place: Philadelphia, Pennsylvania, USA (1999)
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., Sander, J.: LOF: Identifying Density-Based Local Outliers. In: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, SIGMOD '00, 93–104. New York, NY, USA: Association for Computing Machinery. ISBN 1-58113-217-4. Event-place: Dallas, Texas, USA (2000)
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X.: A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96, pp. 226–231. AAAI Press. Event-place: Portland, Oregon (1996)
- European Commission. Laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain union legislative acts (2021)
- Fadahunsi, K.P., Akinlua, J.T., O'Connor, S., Wark, P.A., Gallagher, J., Carroll, C., Majeed, A., O'Donoghue, J.: Protocol for a systematic review and qualitative synthesis of information quality frameworks in eHealth. *BMJ Open* **9**(3), e024722 (2019)
- Fawzy, A., Mokhtar, H.M.O., Hegazy, O.: Outliers detection and classification in wireless sensor networks. *Egypt. Inf. J.* **14**(2), 157–164 (2013)
- Geerkens, S.: Anwendung und Validierung des SHLQI² auf realen Beispieldmengen und neuronale Netzwerke (2021)
- Geerkens, S., Sieberichs, C., Braun, A., Waschulzik, T.: QI2 - an Interactive Tool for Data Quality Assurance (2023)
- Gualo, F., Rodriguez, M., Verdugo, J., Caballero, I., Piattini, M.: Data quality certification using ISO/IEC 25012: Industrial experiences. *J. Syst. Softw.* **176**, 110938 (2021)
- Heinrich, B., Klier, M., Schiller, A., Wagner, G.: Assessing data quality - a probability-based metric for semantic consistency. *Decis. Support Syst.* **110**, 95–106 (2018)
- Holcomb, Z.: Fundamentals of Descriptive Statistics. Routledge, 0 edition. ISBN 978-1-351-97033-4 (2016)
- Iannone, R., Vargas, M.: pointblank: Data Validation and Organization of Metadata for Local and Remote Tables. R package version 0.11.4.9000. <https://rstudio.github.io/pointblank/> (2022)
- Jolliffe, I.T.: Principal component analysis: a beginner's guide - I. Introduction and application. *Weather* **45**(10), 375–382 (1990)
- LeCun, Y., Cortes, C.: MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist/>(2010)
- Maaten, L.v.d., Hinton, G.E.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008)
- McInnes, L., Healy, J., Melville, J.: UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv:1802.03426* [cs, stat] (2020)
- Mock, M., Scholz, S., Blank, F., Hüger, F., Rohatschek, A., Schwarz, L., Stauner, T.: An Integrated Approach to a Safety Argumentation for AI-Based Perception Functions in Automated Driving. In: Habli, I., Suján, M., Gerasimou, S., Schoitsch, E., and Bitsch, F. (eds.) *Computer Safety, Reliability, and Security. SAFECOMP 2021 Workshops*, vol. 12853, pp. 265–271. Springer International Publishing, Cham. ISBN 978-3-030-83905-5 978-3-030-83906-2. Series Title: Lecture Notes in Computer Science (2021)
- Pipino, L.L., Lee, Y.W., Wang, R.Y.: Data quality assessment. *Commun. ACM* **45**(4), 211–218 (2002)
- Samara, M.A., Bennis, I., Abouaissa, A., Lorenz, P.: Enhanced efficient outlier detection and classification approach for WSNs. *Simul. Model. Pract. Theory* **120**, 102618 (2022)
- Samara, M.A., Bennis, I., Abouaissa, A., Lorenz, P.: Complete outlier detection and classification framework for WSNs based on OPTICS. *J. Netw. Comput. Appl.* **211**, 103563 (2023)
- Schelter, S., Lange, D., Schmidt, P., Celikel, M., Biessmann, F., Grafberger, A.: Automating large-scale data quality verification. *Proc. VLDB Endow.* **11**(12), 1781–1794 (2018)
- Schelter, S., Schmidt, P., Rukat, T., Kiessling, M., Taptunov, A., Biessmann, F., Lange, D.: DEEQU - Data quality validation for machine learning pipelines. In: *Machine Learning Systems workshop at the conference on Neural Information Processing Systems (NeurIPS) 2018*. <https://ssc.io/pdf/deequ.pdf> (2018b)
- Sidi, F., Shariat Panahy, P. H., Affendey, L. S., Jabar, M. A., Ibrahim, H., Mustapha, A.: Data quality: A survey of data quality dimensions. In: *2012 International Conference on Information Retrieval & Knowledge Management*, 300–304. Kuala Lumpur: IEEE. ISBN 978-1-4673-1091-8 978-1-4673-1090-1 (2012)
- Sieberichs, C.: Anwendung und Validierung des ECS auf reale Beispieldmengen und neuronale Netzwerke (2021)
- Siemens, A. G.: Digitalisierung und Automatisierung für fahrerlosen Regionalverkehr. <https://press.siemens.com/global/de/pressmitteilung/digitalisierung-und-automatisierung-fuer-fahrerlosen-regionalverkehr> (2022)
- Thang, T.M., Kim, J.: The Anomaly Detection by Using DBSCAN Clustering with Multiple Parameters. In: *2011 International Conference on Information Science and Applications*, 1–5. Jeju Island: IEEE. ISBN 978-1-4244-9222-0 (2011)
- Wang, R.Y., Strong, D.M.: Beyond accuracy: what data quality means to data consumers. *J. Manag. Inf. Syst.* **12**(4), 5–33 (1996)
- Waschulzik, T.: Qualitätsgesicherte effiziente Entwicklung vorwärtsgerichteter künstlicher Neuronaler Netze mit überwachtem Lernen (QUEEN). Ph.D. thesis, Technische Universität München, München (1999)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.