Case-Centred Multidimensional Scaling for Classification Visualisation in Medical Diagnosis

Frank Klawonn[1,2], Werner Lechner[3], and Lorenz Grigull[3,4]

[1] Department of Computer Science
Ostfalia University of Applied Sciences
Salzdahlumer Str. 46/48, D-38302 Wolfenbuettel, Germany
[2] Bioinformatics & Statistics
Helmholtz Centre for Infection Research
Inhoffenstr. 7, D-38124 Braunschweig, Germany
[3] Improved Medical Diagnostics Pte Ltd
190 Middle Road, #19-05 Fortune Centre, Singapore, 188979
[4] Department of Pediatric Haematology and Oncology
Medical University
Carl-Neubergstr. 1, Hannover, Germany

Abstract. Computer-based decision support can assist a medical doctor to find the right diagnosis. The knowledge and experience of the medical doctor is enhanced by a much larger data set of patients than the doctor will ever see in her or his life. The decision support system can derive possible diagnoses for a new patient based on a suitable classifier built on the patients in the patient database. However, since such a system cannot replace a medical doctor and should only support her or him, it should also provide information about the certainty of its recommendation. In this paper, we propose to visualise how close or similar the new patient is to others in the database by a modified multidimensional scaling technique that focuses on the correct positioning of the new patient in the visualisation. In this way, the medical doctor can easily see whether the diagnosis recommended by the system is reliable when all patients close to the new patient have the same diagnosis or whether it is quite uncertain when the new patient is surrounded by patients with different diagnoses.

1 Introduction

To arrive at a medical diagnosis is a complex process. Doctors usually include clinical and laboratory findings to generate a hypothesis and then systematically rule out differential-diagnoses by experience and/or logical thinking. Unfortunately, this process is prone to mistakes. As a consequence, tired doctors or doctors in the emergency department or unexperienced doctors are at risk to pose wrong diagnoses.

We therefore aimed at developing a tool to support the medical diagnostic decision process in the paediatric emergency department. Data sets of about 700 patients with 18 different medical diagnoses frequently encountered at a tertiary childrens hospital were included. Using an ensemble of three different classifiers, a new patient data record could then be allocated to the correct diagnosis with good reliability thus resulting in a good diagnostic support for doctors.

G. Huang et al. (Eds.): HIS 2013, LNCS 7798, pp. 137–148, 2013.

The decision support system for the medical doctor can be incorporated at any phase of the examination of the patient. In a very early state, when only a few attributes or measurements of the patient are available, the decision support system might not be able to propose a diagnosis or only one with great uncertainty. With more and more measurements taken from the patient, the decision support system will be able to provide one or two possible diagnoses with high certainty.

It is important for the medical doctor to obtain information about the status of the decision support system, how certain it is about the proposed diagnosis. The system proposed in [1] which we use in this paper as a case study, will provide an A- and a B-diagnosis. The A-diagnosis is the one considered to be most probable, the B-diagnosis is the one with the second highest probability. The system will not provide any diagnosis if the A-diagnosis has too little certainty. Nevertheless, it is crucial to know whether the system is relatively sure about the proposed A-diagnosis or whether the B-diagnosis is almost as likely as the A-diagnosis.

This could be indicated by probabilities or scored. But a visualisation can include more information than a single number. Therefore, we have developed a visualisation technique that shows how close the new patient is to those ones in the database for which the diagnosis is already known.

Our method is based on principles of multidimensional scaling, a visualisation technique briefly reviewed in Section 2. However, the focus of our visualisation should be put on the new patient, not on data visualisation in general. Therefore, we connect multidimensional scaling as a general purpose data visualisation technique with classifiers and classifier ensembles in Section 3. Since the new patient to be classified should be in the centre of the visualisation, we propose a modified version of multidimensional scaling in Section 4 which we call case-centred multidimensional scaling. The algorithm to compute the visualisation for case-centred multidimensional scaling is explained in Section 5. Section 6 illustrates how our visualisation technique works for artificial and real data. The final conclusions summarise our results and emphasize that – although we have focused on medical diagnosis – our visualisation technique can also be applied in other areas.

2 Multidimensional Scaling and Distance-Based Visualisation

Multidimensional scaling (MDS) is a dimension reduction technique which is mainly used for visualisation of high-dimensional data. MDS assumes that a p-dimensional data set $X = \{x_1, x_2, ..., x_n\} \subseteq \mathbb{R}^p$ is given. By $d(v, w)$ we denote the Euclidean distance $\|v - w\|$ between points $v \in \mathbb{R}^p$ and $w \in \mathbb{R}^p$. In MDS each of the high-dimensional data points x_i has to be mapped to a low-dimensional representative y_i. The projection of X is denoted as $Y = \{y_1, y_2, ..., y_n\} \subseteq \mathbb{R}^q$ where $1 \leq q < p$ (typically $q \in \{2, 3\}$). A perfect distance-preserving projection of X to Y would keep the distances $d^x_{ij} = d(x_i, x_j)$ of the high-dimensional space identical to the distances $d^y_{ij} = d(y_i, y_j)$ of the projected data objects, that is, $d^x_{ij} = d^y_{ij}$ holds. A perfect projection is, however, impossible except for a few trivial cases. Therefore, MDS seeks to minimise the error introduced by the projection ($|d^x_{ij} - d^y_{ij}|$ for all i, j). Common objective functions for MDS are [2]:

$$E_1 = \frac{1}{\sum_{i=1}^{n}\sum_{j=i+1}^{n}\left(d_{ij}^x\right)^2}\sum_{i=1}^{n}\sum_{j=i+1}^{n}\left(d_{ij}^y - d_{ij}^x\right)^2, \tag{1}$$

$$E_2 = \sum_{i=1}^{n}\sum_{j=i+1}^{n}\left(\frac{d_{ij}^y - d_{ij}^x}{d_{ij}^x}\right)^2, \tag{2}$$

$$E_3 = \frac{1}{\sum_{i=1}^{n}\sum_{j=i+1}^{n} d_{ij}^x}\sum_{i=1}^{n}\sum_{j=i+1}^{n}\frac{\left(d_{ij}^y - d_{ij}^x\right)^2}{d_{ij}^x}. \tag{3}$$

E_1 is based on the absolute errors for the distances, E_3 on the relative errors and E_2 is a compromise between the absolute and the relative error. E_2 is also called stress and MDS based on E_2 is called Sammon mapping. The minimisation of any of the objective functions for MDS poses a non-linear optimisation problem, so that the selected objective function is usually minimised by a numerical optimisation technique.

It should be noted that MDS only needs the distances d_{ij}^x between the data points, not the coordinates of the points themselves. MDS does not have to be based on the Euclidean distance. Any other distance measure between the data points can also be used to define the values d_{ij}^x. This fact will also be exploited in the following section when we explain the connection MDS and the visualisation of a classification result.

3 Visualisation of Classification Decisions

In this paper, a modified version of MDS is introduced, not for the primary purpose of visualising the patients in a database for whom the diagnosis is already known, but for illustrating how similar a new patient is to other patients in the database.

Figure 1 uses an artificial data set to illustrate how such a visualisation could look. The red cone represents the new patient. Patients in the database with diagnosis A are represented by green spheres, whereas patients in the database with diagnosis B are marked by blue cubes. In this case, the new patient fits very well to the patients with diagnosis A and is not very close to the patients with diagnosis B.

But how would one generate such a visualisation based on MDS? MDS requires the pairwise distances between the patients in the database as well as distances of the new patient to the patients in the database. One could simply use the Euclidean distance as a distance measure. This would require that all attributes are numerical, so that categorical attributes would have to be converted into numerical attributes. Furthermore, in order to be independent of the influence of the measurement unit, each attribute should be normalised in a suitable way, for instance by scaling its range to the unit interval. Conversion of categorical attributes to numerical ones and normalisation is not the topic of this paper and we refer to a more detailed discussion on these topics to [3].

Missing values must also be taken into account. Missing values in patient data are quite normal, since not all possible measurements will be taken from each patient. The Euclidean distance between two patients could be computed based only on those attributes without missing values. But this would introduce a bias to smaller distances for

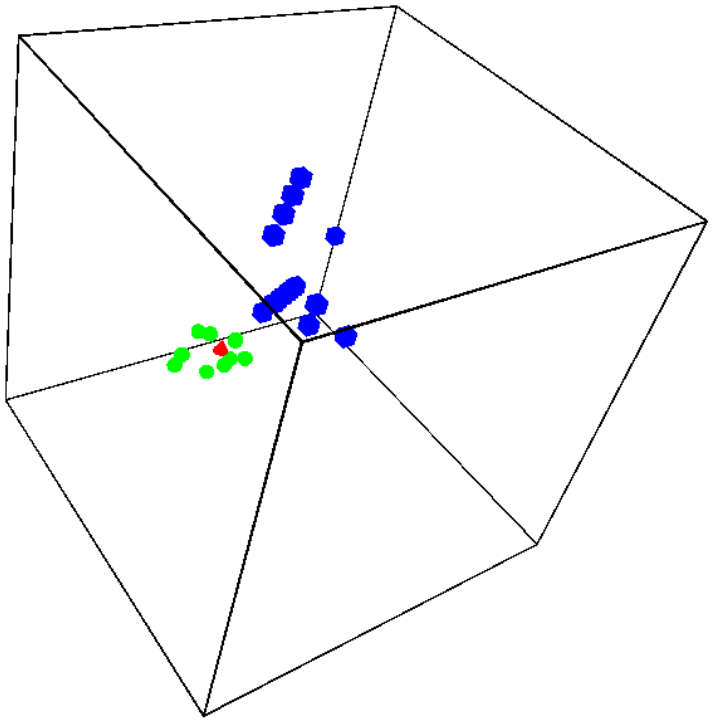

Fig. 1. Visualisation for an artificial data set where the classes or diagnoses are well separated and the new patient is in the centre of patients from one diagnosis. The cube is drawn to provide a better 3D perception.

patients with a larger number of missing values. To avoid this bias, this distance could be divided by the number of attributes which where available for the distance calculation. Instead of the Euclidean distance, suitable kernel-based distances [4] could also be considered. However, a visualisation based on the Euclidean or a kernel-based distance would essentially illustrate how well a nearest neighbour classifier would perform.

But the system for diagnostic support might not be based on a nearest neighbour classifier, but on other classification techniques. In our case [1], none of the single classifiers we tried showed a fully satisfactory performance. Therefore, we rely on a classifier ensemble [5] of three classifiers: a support vector machine, a neural network and a fuzzy rule-based system.

The visualisation of the performance of a classifier ensemble could be based on the distances between the outputs of the different classifiers instead of the Euclidean distances based on the attributes of the patients. This idea is illustrated by Table 3. For the new patient and also for each patient in the database, we have k different possible diagnoses and each of the c classifiers of the classifier ensemble provides a score for each of the diagnoses. This means, each patient is represented by a $(k \cdot c)$-dimensional vector. The Euclidean distance between these vectors could be used as basis for the MDS visualisation. In our specific application, we have 18 diagnoses and 3 classifiers, so that the

Table 1. Structure of the data table for the distance calculation based on the results of the single classifiers

	Classifier 1			...	Classifier c		
Patient	Diagnosis 1	...	Diagnosis k	...	Diagnosis 1	...	Diagnosis k
New patient	•	•	•	•	•	•	•
Patient 1	•	•	•	•	•	•	•
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Euclidean distance between 54-dimensional vectors would have to be computed. However, this could lead to certain problems. Two patients might be considered very similar mainly because they do definitely belong to a larger subset of diagnoses and have very similar values for these diagnoses. But they might still be different in the two or three still possible diagnoses. Therefore, we display the new patient together only with those patients in the database from the two most probable diagnoses that the classifier indicates for the new patient. Therefore, the distance between patients is computed based on 6- instead of 18-dimensional vectors. We also limit the number of patients to be displayed and show only those cases from the database being closest to the new patient.

4 Case-Centred Multidimensional Scaling

In the previous section, we have explained how the distances between the new patient and patients in the database and also distances between patients in the database can be computed. Based on these distances, MDS can be applied to obtain a visualisation that indicates how well the new patient can be classified by the computer-based decision support system. In this way, the medical doctor obtains information about where the decision support system stands and whether further information about the patient is needed to be able to distinguish well between the two closest diagnoses.

In principle, one could apply MDS directly to the corresponding distances or precalculate an MDS visualisation for the patients in the database and use the technique described in [6] to add the new patient to the visualisation. This would, however, imply that all distances are equally important: the distances between the new patient and the patients in the database in the same way as distances between patients in the database. But the focus of our visualisation is the new patient and where he should be positioned in comparison to the patients in the database.

Therefore, we introduce a case-centred version of MDS for a 3D-visualisation. The new patient is placed in the centre of the visualisation, i.e. at the origin of the coordinate system and we preserve the distance of the new patient to each of the patients in the database exactly. This is achieved by placing each patient in the database on the surface of a sphere around the origin of the coordinate system. The distance to the new patient is used as the radius of the sphere. This can be considered as a constrained MDS problem. The patients in the database cannot be positioned arbitrarily, but are constraint to the surface of their corresponding sphere.

The formalisation of this constraint MDS problem and an algorithm to solve it is presented in the following section. We use polar coordinates to simplify the problem.

Polar coordinates where already used for MDS in [7] for two-dimensional and in [8] for three-dimensional representations. However, in contrast to our proposed method, these approaches – like ordinary MDS – do not focus on a specific patient or data object in the visualisation. All patients or data objects are treated equally. But for our purposes, it is essential to primarily focus on the new patient in relation to the patients in the database and only then to consider the relations between the patients in the database.

5 Algorithm for Case-Centred Multidimensional Scaling

For each patient in the database, we have the freedom to choose the position on the surface of the sphere with the radius corresponding to the distance to the new patient. The objective function for the Sammon mapping (3) becomes for a 3D-visualisation

$$E \;=\; \frac{1}{4}\sum_{i=1}^{n-1}\sum_{j=i+1}^{n}\frac{\big((x_i-x_j)^2+(y_i-y_j)^2+(z_i-z_j)^2-d_{ij}\big)^2}{d_{ij}} \tag{4}$$

where we assume that we have n patients in the database. The constant factor $\frac{1}{4}$ is introduced for convenience reasons as we will see later on. d_{ij} is the distance between patients i and j in the database. (x_i, y_i, z_i) are the coordinates for positioning patient i in the visualisation.

The objective function (4) should be minimised under the constraints

$$x_i^2 + y_i^2 + z_i^2 \;=\; r_i^2 \tag{5}$$

where r_i^2 is the distance of the new patient to patient i in the database. In this way it is guaranteed that the distances between the new patient and the patients in the database are represented without error in the visualisation.

The minimisation of the objective function (4) under the constraints (5) is based on a gradient descent method. In order to take the constraints (5) into account, we rewrite Eq. (4) in spherical coordinates. For the computation of the gradient, we need to calculate the partial derivatives with respect to all parameters. For the sake of simplicity, we only use spherical coordinates for the record i for which we want to calculate the corresponding derivatives. Then the addend in Eq. (4) becomes

$$\frac{\big((r_i\sin(\theta_i)\cos(\varphi_i)-x_j)^2+(r_i\sin(\theta_i)\sin(\varphi_i)-y_j)^2+(r_i\cos(\theta_i)-z_j)^2-d_{ij}\big)^2}{d_{ij}} \tag{6}$$

where $0 \le \theta_i < \pi$ and $0 \le \varphi_i < 2\pi$.

The partial derivatives w.r.t. θ_i and φ_i are

$$\frac{\partial E}{\partial \theta_i} \;=\; r_i \cdot \sum_{\substack{j=i\\ j\neq i}}^{n} \frac{1}{d_{ij}}.$$

$$\Big((r_i \sin(\theta_i)\cos(\varphi_i) - x_j)^2 + (r_i \sin(\theta_i)\sin(\varphi_i) - y_j)^2 + (r_i \cos(\theta_i) - z_j)^2 - d_{ij}\Big)\cdot$$

$$\begin{aligned}\Big(\quad & (r_i \sin(\theta_i)\cos(\varphi_i) - x_j)\cos(\theta_i)\cos(\varphi_i) \\ & + (r_i \sin(\theta_i)\sin(\varphi_i) - y_j)\cos(\theta_i)\sin(\varphi_i) \\ & - (r_i \cos(\theta_i) - z_j)\sin(\theta_i) \quad \Big)\end{aligned} \tag{7}$$

and

$$\frac{\partial E}{\partial \varphi_i} \;=\; r_i \cdot \sum_{\substack{j=i \\ j\neq i}}^{n} \frac{1}{d_{ij}}\cdot$$

$$\Big((r_i \sin(\theta_i)\cos(\varphi_i) - x_j)^2 + (r_i \sin(\theta_i)\sin(\varphi_i) - y_j)^2 + (r_i \cos(\theta_i) - z_j)^2 - d_{ij}\Big)\cdot$$

$$\begin{aligned}\Big(\quad & -(r_i \sin(\theta_i)\sin(\varphi_i) - x_j)\sin(\theta_i)\sin(\varphi_i) \\ & + (r_i \sin(\theta_i)\sin(\varphi_i) - y_j)\sin(\theta_i)\cos(\varphi_i) \quad \Big),\end{aligned} \tag{8}$$

respectively.

One could first use random values for the parameters x_i and y_i ($i = 1, \ldots, n$), i.e. position each of the patients in the database randomly on the surface of his corresponding sphere. But we are more interested to make smaller mistakes between the patients in the database who are close to the new patients, i.e. those patients whose associated spheres have a radii. We insert the patients stepwise into the visualisation in ascending order with respect to their distances to the new patient (radii of their associated spheres).

By reordering the patients in the database with respect to their distance to the new patient, we can assume without loss of generality that $r_1 \leq r_2 \leq \ldots \leq r_n$ holds. We then proceed in the following way.

1. The first patient in the database, i.e. the one closest to the new patient, will be positioned at point $(x_1, y_1, z_1) = (r_1, 0, 0)$.
2. If $r_1 + \sqrt{d_{12}} \geq r_2$ and $r_1 + r_2 \geq \sqrt{d_{12}}$ hold, then the patient in the database second closest to the new patient can be positioned in such a way that the distance $\sqrt{d_{12}}$ between the first two patients in the datbase can also be represented exactly in the visualisation. In this case, choose $(x_2, y_2, z_2) = (r_2\cos(\alpha), r_2\sin(\alpha), 0)$ for the second patient in the database where

$$\alpha \;=\; \arccos\left(\frac{r_1^2 + r_2^2 - d_{12}}{2r_1 r_2}\right).$$

   If $r_1 + \sqrt{d_{12}} < r_2$ holds, then the second patient in the database will be positioned at $(x_2, y_2, z_2) = (r_2, 0, 0)$ in order to minimise the error for the distance d_{12}.
   If $r_1 + r_2 < \sqrt{d_{12}}$ holds, then the second patient in the database will be positioned at $(x_2, y_2, z_2) = (-r_2, 0, 0)$ in order to minimise the error for the distance d_{12}.

3. For positioning the third point, the problem of three intersecting spheres needs to be solved: The third point should have a distance of r_3 to the origin of the coordinate system – where the new patient is located – and distances d_{13} and d_{23} to the points representing the first and the second patient in the database. In case there is a solution to the problem of intersecting spheres, choose (x_3, y_3, z_3) where

$$x_3 = \frac{r_1^2 + r_3^2 - d_{13}}{2r_1},$$
$$y_3 = \frac{r_3^2 - d_{23} + x_2^2 + y_2^2}{2y_2} - \frac{x_2}{y_2}x_3,$$
$$z_3 = \sqrt{r_3 - x_3^2 - y_3^2}.$$

   If there is no solution for the problem of three intersecting spheres, i.e. if $r_3 - x_3^2 - y_3^2 < 0$ holds, then the position of the third point is determined by a gradient method as described in the following step.
4. When the positions $(x_1, y_1, z_1), \ldots, (x_k, y_k, z_k)$ for the patients with $r_1 \leq \ldots \leq r_k \leq \ldots \leq r_n$ have been determined, the position $(x_{k+1}, y_{k+1}, z_{k+1})$ for the next patient is determined by a gradient descent method based on the gradient in terms of the variables θ_{k+1} and φ_{k+1} given in Equations (7) and (8), respectively, where $i = k + 1$. For the position $(x_{k+1}, y_{k+1}, z_{k+1})$ we only consider the patients $1, \ldots, k$.

The above described algorithm positions the points for the patients in the visualisation step by step in increasing order of their distances to the new case, in order to obtain a quick solution. Of course, one could also use this as an initialisiation and then optimise the positions again based on a similar gradient descent method, but now always taking all points into account. However, this would also lead to higher computational costs.

6 Examples

As mentioned above, we include only those patients from the database in the visualisation with the two closest diagnoses that are considered possible for the new patient. An ideal result would look like the one for the artificial data set in Figure 1 where the new patient – the red cone – is only surrounded directly by patients represented by green spheres, i.e. all patients close to the new patient have the same diagnosis. Patients with the second most probable diagnosis marked by blue cubes are all quite far away from the new patient.

Figure 2 shows the data of a 17 year old boy suffering from leukemia and pneumonia. He was initially admitted for pneumonia, but further tests revealed an acute leukemia (ALL). Interestingly, both diagnoses are valued by the system. The green spheres indicate other patients with malignant hematological diseases. The red triangle symbolizes the patient under discussion and the blue cubes stand for children with pneumonia. Of note, the doctors first treated only pneumonia in this patient and the leukemia has only

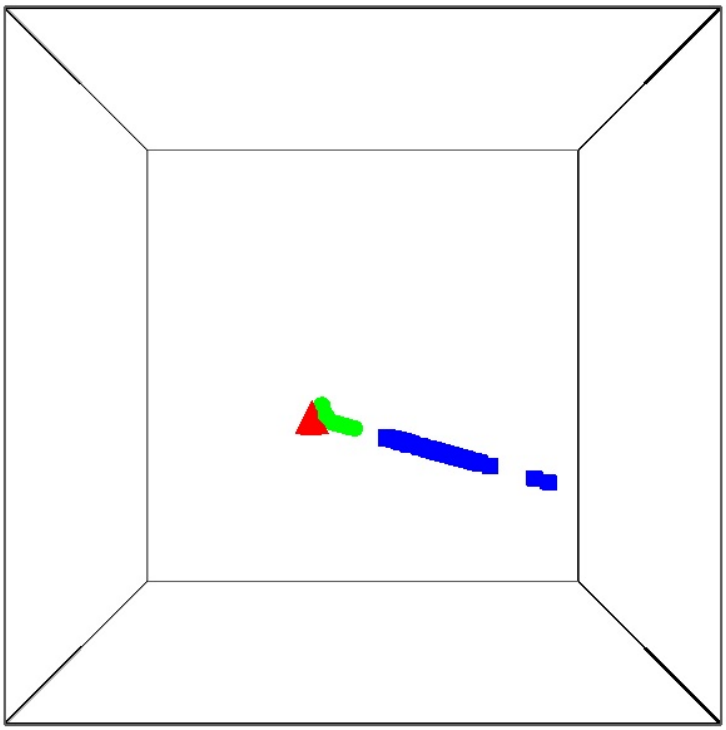

Fig. 2. A difficult patient for the doctors

been detected after transferall to a tertiary hospital. Usage of the diagnostic aid would have been helpful for treatment of this boy.

Figure 3 shows the case of patient where it was difficult to find the right diagnosis. In this 8 year old boy with high fever and headache an inflammation of the brain (meningitis) was suspected. Results of additional diagnostic tests were negative for this differential diagnosis. Two days later, he developped severe abdominal pain. Now, the doctors in charge assumed that "appendicitis" could be the cause of the problems. However, the surgeons remained sceptical and ordered a computer tomography (CT) of the abdomen. Surprisingly, the CT scan showed a large abscess formation in the kidney. To our big surprise, the computer system was already suggesting an inflammation of the urinary tract followed by systemic infection. Both differential diagnoses are apparently correct and better than those posed by the doctor on duty.

The patient in Figure 4 came to the pediatric emergency department with a short history of abdominal pain. The 11 year old boy had fever and nausea. Further tests showed a normal leukocyte number and a moderately elevated C-reactive protein. An appendicitis was suspected, but the surgeon was not convinced. Therefore, sonography was performed. Here, small amounts of fluid in the abdominal cavity were seen. As a consequence, an appendectomy was done revealing a severe inflammation of the appendix. The diagnostic tool gave strong arguments for the differential diagnosis "appendicitis" already in admission. This might help for fast track procedures.

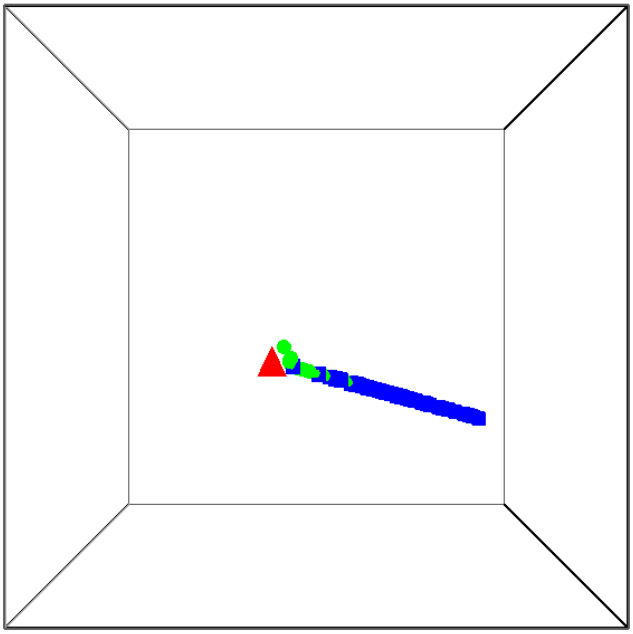

Fig. 3. A case where the system proposes the right diagnosis, but with less certainty

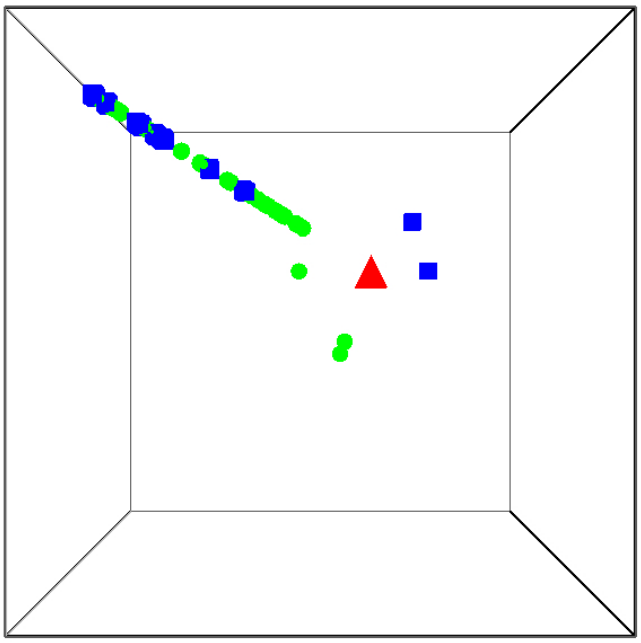

Fig. 4. Another case where the system proposes the right diagnosis, but with less certainty

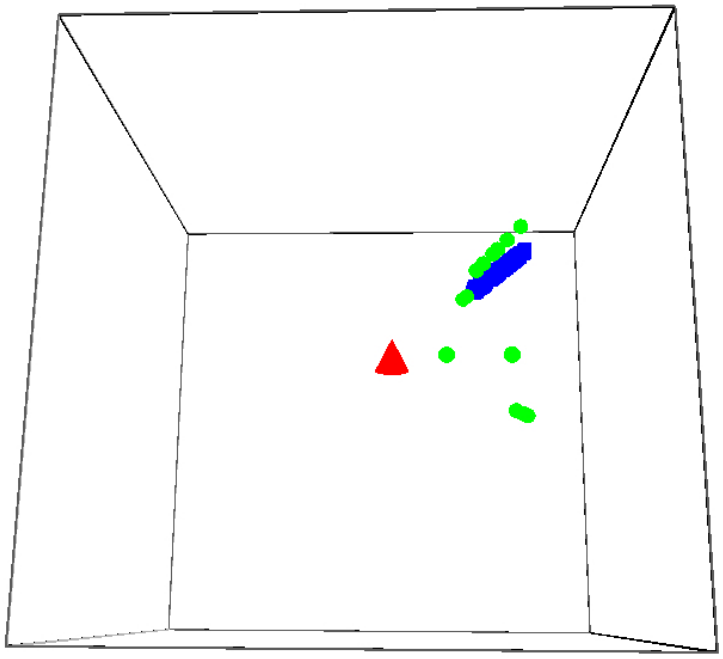

Fig. 5. Visualisation of a data set obtained from a healthy person.

Figure 5 shows an interesting example with an "ideal" virtual healthy patient. The red triangle does not represent a specific patient, but is a virtual person obtained by averaging the data of all healthy persons in the database. The green spheres correspond to healthy patients in the database. Although the green spheres, representing the correct diagnosis, are closest to the virtual healthy person, they are as close to the red triangle as in the other examples. This is probably an indication for the curse of dimensionality [9,10]. The original data set contains 26 clinical and laboratory parameters, so that – according to the curse of dimensionality – the density at the centre of this 26-dimensional data set is quite low.

7 Conclusions

Wrong diagnoses carry an enormous risk for patients. Clinical decision support systems try to reduce this risk, but such applications are not widely used in daily practice. Our new diagnostic data mining tool demonstrated good results to compute a diagnosis using 26 clinical and laboratory parameters frequently used in a paediatric emergency department. To even increase the potential benefit for a user of this diagnostic support system, the display of the data is of considerable relevance. Any (tired) doctor will profit from data display immediately indicating "diagnostic clarity" or even "diagnostic vagueness". This feed-back given from the computer back to the doctor might result

in an impulse to order additional diagnostic tests to provide additional security. Consequently, the risk of posing wrong diagnoses should be decreased.

Our method is not restricted to patient data, but can be applied to any type of classification problem to visualise how well a nearest neighbour classifier or a classifier ensemble can classify a specific new object.

In this paper, we have focused on 3D-visualisation. It is straightforward to derive a corresponding algorithm for a simplified 2D-visualisation.

References

1. Grigull, L., Lechner, W.: Supporting diagnostic decisions using hybrid and complementary data mining applications: a pilot study in the pediatric emergency department. Pediatric Research 71, 725–731 (2012)
2. Kruskal, J., Wish, M.: Multidimensional Scaling. SAGE Publications, Beverly Hills (1978)
3. Berthold, M., Borgelt, C., Höppner, F., Klawonn, F.: Guide to Intelligent Data Analysis: How to Intelligently Make Sense of Real Data. Springer, London (2010)
4. Shawe-Taylor, J., Cristianini, N.: Kernel Methods for Pattern Analysis. Cambridge University Press, Cambridge (2004)
5. Kuncheva, L.: Combining Pattern Classifiers: Methods and Algorithms. Wiley, Chichester (2004)
6. Pekalska, E., de Ridder, D., Duin, R., Kraaijveld, M.: A new method of generalizing Sammon mapping with application to algorithm speed-up. In: Boasson, M., Kaandorp, J., Tonino, J., Vosselman, M. (eds.) ASCI 1999: Proc. 5th Annual Conference of the Advanced School for Computing and Imaging, Delft, ASCI, pp. 221–228 (1999)
7. Rehm, F., Klawonn, F., Kruse, R.: MDS_{polar}: A new approach for dimension reduction to visualize high dimensional data. In: Famili, A.F., Kok, J.N., Peña, J.M., Siebes, A., Feelders, A. (eds.) IDA 2005. LNCS, vol. 3646, pp. 316–327. Springer, Heidelberg (2005)
8. Rehm, F., Klawonn, F.: Improving angle based mappings. In: Tang, C., Ling, C.X., Zhou, X., Cercone, N.J., Li, X. (eds.) ADMA 2008. LNCS (LNAI), vol. 5139, pp. 3–14. Springer, Heidelberg (2008)
9. François, D., Wertz, V., Verleysen, M.: The concentration of fractional distances. IEEE Trans. Knowl. Data Eng. 19(7), 873–886 (2007)
10. Jayaram, B., Klawonn, F.: Can unbounded distance measures mitigate the curse of dimensionality? Int. Journ. Data Mining, Modelling and Management 4, 361–383 (2012)