Learning similarity measures from data

Bjørn Magnus Mathisen^{1,2} · Agnar Aamodt^{1,2} · Kerstin Bach^{1,2} · Helge Langseth^{1,2}

Received: 29 April 2019 / Accepted: 15 October 2019 / Published online: 30 October 2019
© The Author(s) 2019

Abstract

Defining similarity measures is a requirement for some machine learning methods. One such method is case-based reasoning (CBR) where the similarity measure is used to retrieve the stored case or a set of cases most similar to the query case. Describing a similarity measure analytically is challenging, even for domain experts working with CBR experts. However, datasets are typically gathered as part of constructing a CBR or machine learning system. These datasets are assumed to contain the features that correctly identify the solution from the problem features; thus, they may also contain the knowledge to construct or learn such a similarity measure. The main motivation for this work is to automate the construction of similarity measures using machine learning. Additionally, we would like to do this while keeping training time as low as possible. Working toward this, our objective is to investigate how to apply machine learning to effectively learn a similarity measure. Such a learned similarity measure could be used for CBR systems, but also for clustering data in semi-supervised learning, or one-shot learning tasks. Recent work has advanced toward this goal which relies on either very long training times or manually modeling parts of the similarity measure. We created a framework to help us analyze the current methods for learning similarity measures. This analysis resulted in two novel similarity measure designs: The first design uses a pre-trained classifier as basis for a similarity measure, and the second design uses as little modeling as possible while learning the similarity measure from data and keeping training time low. Both similarity measures were evaluated on 14 different datasets. The evaluation shows that using a classifier as basis for a similarity measure gives state-of-the-art performance. Finally, the evaluation shows that our fully data-driven similarity measure design outperforms state-of-the-art methods while keeping training time low.

Keywords Similarity measure · Data science · Neural networks · Data analytics · Case-based reasoning · Similarity function · Siamese networks · Similarity metrics · Distance metrics

1 Introduction

Many artificial intelligence and machine learning (ML) methods, such as k-nearest neighbors (k-NN), rely on a similarity (or distance) measure [21] between data points. In case-based reasoning (CBR), a simple k-NN or a more complex similarity function is used to retrieve the stored cases that are most similar to the current query case. The similarity measure used in CBR systems for this purpose is typically built as a weighted Euclidean similarity measure (or as a weight matrix for discrete and symbolic variables). Such a similarity measure is designed with assistance of domain experts by adjusting the weights for each attribute of the cases to represent how important they are (one example can be seen in [32] or generally described in chapter 4 of [3])

In many situations, the design of such a function is non-trivial. Domain experts with an understanding of CBR or machine learning are not easily available. However, before

This work was supported by the Research Council of Norway through the grant number 237790 and the Norwegian Open AI Lab.

✉ Bjørn Magnus Mathisen
bjornmm@ntnu.no
<http://www.idi.ntnu.no>

Agnar Aamodt
agnar@ntnu.no

Kerstin Bach
kerstin.bach@ntnu.no

Helge Langseth
helgel@ntnu.no

¹ Department of Computer Science, Norwegian University of Science and Technology, Trondheim, Norway

² Department of Computer Science, EXPOSED Aquaculture Research Centre, Trondheim, Norway

or during most CBR projects, data are gathered that relate to the problem being solved by the CBR system. These data are used to construct cases for populating the case base. If the data are labeled according to the solution/class, it can be used to learn a similarity measure that is relevant to the task being solved by the system. Exploring efficient methods of learning similarity measures and improving on them are the main motivation of this work.

In the CBR literature, similarity measurement is often described in terms of problem and solution spaces. Problem space is where the features of a problem describe the problem; this is often called feature space in the non-CBR ML literature. Solution space, also referred to as target space, is populated by points describing solutions to points in the problem space. The function that maps a point from the problem space to its corresponding point in the solution space is typically the goal of supervised machine learning. This is illustrated in Fig. 1.

A similarity measure in the problem space represents an approximation of the similarity between two cases or data points in the solution space (i.e., whether these two cases have similar or dissimilar solutions). Such a similarity measure would be of great help in situations where lots of labeled data are available, but domain knowledge is not available, or when the modeling of such a similarity measure is too complex.

Learned similarity measures can also be used in other settings, such as clustering. Another relevant method type is semi-supervised learning in which the labeled part of a dataset is used to cluster or label the unlabeled part.

How to automatically learn similarity measures has been an active area of research in CBR. For instance, Gabel et al. [10] trained a similarity measure by creating a dataset of

collated pairs of data points and their respective similarities. This dataset is then used to train a neural network to represent the similarity measure. In this method, the network needs to extract the most important features in terms of similarity for both data points and then combine these features to output a similarity measure. Recent work (e.g., Martin et al. [22]) has used siamese neural networks (SNN) [5] to learn a similarity measure in CBR. SNNs have the advantage of sharing weights between two parts of the network, in this case the two parts that extract the useful information from the two data points being compared. All of these methods for learning similarity measures have in common that they are trained to compare two data points and return a similarity measurement. Our work of automatically learning similarity measures is also related to the work done by Hüllermeier et al. on preference-based CBR [14,15]. In this work, the authors learn a preference of similarity between cases/data points, which represents a more continuous space between solutions than a typical similarity measure in CBR. This type of approach to similarity measures is similar to learning similarity measures by using machine learning models, in that both can always return a list of possible solutions sorted by their similarity.

In this work, we have developed a framework to show the main differences between various types of similarity measures. Using this framework, we highlight the differences between existing approaches in Sect. 3. This analysis also reveals areas that have not received much attention in the research community so far. Based on this, we developed two novel designs for using machine learning to learn similarity measures from data. Both of the two designs are continuous in their representation of the estimated solution space.

The novelty of our work is threefold: First, we show that using a classifier as a basis for a similarity measure gives adequate performance. Then, we demonstrate similarity measure designed to use as little modeling as possible, while keeping training time low, outperforms state-of-the-art methods. Finally, to analyze the state of the art and compare it to our new similarity measure design we introduce a simple mathematical framework. We show how this is a useful tool for analyzing and categorizing similarity measures.

The remainder of this paper describes our method in more detail. Section 2 describes the novel framework for similarity measurement learning, and Sect. 3 then summarizes previous relevant work in relation to this framework. In Sect. 4, we describe suggestions of new similarity measures and how we design the experimental evaluation. Subsequently, in Sect. 5 we show the results of this evaluation. Finally, in Sect. 6 we interpret and discuss the evaluation results and give some conclusions. We present some of the limitations of our work as well as possible future paths of research.

Fig. 1 Illustration of problem and solution spaces [19]. p_y and p_z are two problem descriptions with features describing a problem each of which has a corresponding (s_y and s_z) solution in solution space. δ_p illustrates the distance between a new problem p_x and a stored problem p_y . Correspondingly, δ_s is the distance between the solution s_y and the solution s_x which is the (unknown) ideal solution to p_x . A fundamental assumption in CBR is that if the similarity between p_x and p_y is high, then the similarity between the unknown solutions s_x and p_y is high ($\delta_p \approx \delta_s$): similar problems have similar solutions

2 A framework for similarity measures

We suggest a framework for analyzing different functions for similarity with $\mathbb{S}$ as a similarity measure applied to pairs of data points $(\mathbf{x}, \mathbf{y})$:

$$\mathbb{S}(\mathbf{x}, \mathbf{y}) = C(G(\mathbf{x}), G(\mathbf{y})), \quad (1)$$

where $G(\mathbf{x}) = \hat{\mathbf{x}}$ and $G(\mathbf{y}) = \hat{\mathbf{y}}$ represent embedding or information extraction from data points $\mathbf{x}$ and $\mathbf{y}$, i.e., $G(\cdot)$ highlights the parts of the data points most useful to calculate the similarity between them. An illustration of this process can be seen in Fig. 2.

$C(G(\mathbf{x}), G(\mathbf{y})) = C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ models the distance between the two data points based on the embeddings $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$. The functions C and G can be either manually modeled or learned from data configurations of Eq. 1 in Table 1. We then describe their main properties and how they have been implemented in state-of-the-art research. Note that we will use $\mathbb{S}(\cdot)$ to annotate the similarity measurement and $C(\cdot)$ for the sub-part of the similarity measurement that calculates the distance between the two outputs of $G(\cdot)$. $\mathbb{S}(\cdot)$ is distinct from $C(\cdot)$ unless $G(\mathbf{x}) = I(\mathbf{x}) = \mathbf{x}$.

In the following, we characterize the different types of similarity measures:

Type 1 A typical similarity measure in CBR systems would model $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $G(\cdot)$ from domain knowledge. Such a similarity measure is typically modeled by experts

Fig. 2 Illustrating how $G(\cdot)$ from Eq. 1 adds another space, the embedding space, between the problem and the solution space [19] (see Fig. 1). $C(\cdot)$ then combines the two embeddings of p_y and p_x (e_y and e_x , respectively) and calculates the similarity δ_e between them. The main assumption is that distance in embedding space (δ_e) is close to the distance in solution space (δ_s); if the embedded points e_x and e_y are similar, then the true (unknown) solution s_x is similar to solution s_y . The main contribution of $G(\cdot)$ is to create an embedding space such that the distance in embeddings space (δ_e) is a better estimate of the distance in solution space (δ_s) than the distance in problem space (δ_p)

Table 1 Different types of similarity measures in our proposed framework

	$C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$	
	Modeled	Learned
$G(\cdot)$		
Modeled	Type 1	Type 2
Learned	Type 3	Type 4

with the relevant domain knowledge together with CBR experts, who know how to encode this domain knowledge into the similarity measures.

For example when modeling the similarity measure of cars for sale, where the goal is to model the similarity of cars in terms of their final selling price. In this example, domain experts may model the embedding function $G(\cdot)$ so that the amount of miles driven has a greater importance than the color of the car. $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ could be modeled such that difference in miles driven is less important than difference in the number of repairs done on the car. More details and examples can be found in [7].

Type 2 This type represents similarity measures that model $G(\cdot)$ and learn the function $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$. In this context, $G(\cdot)$ can be viewed as an embedding function. Since $G(\cdot)$ is not learned from the data, it is not interesting to analyze it as part of learning the similarity measure, as processing the data through $G(\cdot)$ could be done in batch before applying the data to $\mathbb{S}(\mathbf{x}, \mathbf{y})$. Learning $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ needs to be done with a dataset consisting of triplets of the data points $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, and s being the true similarity between $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$.

A special case of Type 2 is when $G(\cdot)$ is set to be the identity function $I(\mathbf{x}) = G(\mathbf{x}) = \mathbf{x}$, while $C(\mathbf{x}, \mathbf{y})$ is learned from data. Examples of this type are presented, for example, in Gabel et al. [10] where the similarity measure always looks at the two inputs together, never separately.

Type 3 In this type of similarity measure, the embedding/feature extraction $G(\cdot)$ is learned and $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is modeled. Typically, the embedding function learned by $G(\cdot)$ resembles the function that is the goal during supervised machine learning. Within the similarity measurement, $\hat{\mathbf{x}} = G(\mathbf{x})$ is used as an embedding vector for calculating similarity, when in classification $\hat{\mathbf{x}}$ would be the softmax vector output. Using a pre-trained classification model as a starting point for $G(\mathbf{x}) = \hat{\mathbf{x}}$ as input to, e.g., $C(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \|\hat{\mathbf{x}} - \hat{\mathbf{y}}\|_1$ should give good results for similarity measurements if that model had high precision for classification within the same dataset.

However, it is not given that the best embedding vector for calculating similarity is the same as the embedding vector produced by a $G(\mathbf{x})$ trained to do classification.

Type 4 This measure is designed so that both $G(\cdot)$ and $C(\hat{x}, \hat{y})$ are learned.

We will design, implement and evaluate similarity measures based on Type 1, Type 3, Type 2 and Type 4 in Sect. 4. These results are shown in Sect. 5.

To allow $\mathbb{S}$ as a similarity measurement for clustering, e.g., k-nearest neighbors, a similarity measure should fulfill the following requirements:

Symmetry $\mathbb{S}(x, y) = \mathbb{S}(y, x)$ The similarity between x and y should be the same as the similarity between x and y .

Nonnegative $\mathbb{S}(x, y) \geq 0 \forall x, y$ The similarity between two data points cannot be negative.

Identity $\mathbb{S}(x, y) = 1 \iff x = y$ The similarity between two data points should be 1 iff x is equal to y .

Some of these requirements are not satisfied by all types of similarity measures, i.e., symmetry is not a direct design consequence of Type 2 but of Type 3 if $C(\hat{x}, \hat{y})$ is symmetric. Even if symmetry is not present in all similarity measures [30], it is important for reducing training time, as the training set size goes from $N(N-1)$ to $N(\frac{N}{2}-1)$. Symmetry also enables the similarity measure to be used for clustering.

In the next section, we will relate current state of the art to the framework in the context of the different types.

3 Related work

To exemplify the framework presented in the previous section, we will relate previous work to the framework and the types of similarity measurements that derive from the framework. This will also enable us to see possibilities for improvement and further research.

As stated in Sect. 1, our motivation is to automate the construction of similarity measures. Additionally, we would like to do this while keeping training time as low as possible. Thus, we will not focus on Type 1 similarity measures as this type uses no learning. Both Type 2 and Type 4 require a different type of training dataset than a typical supervised machine learning dataset, as $C(x, y)$ is typically dependent on the order of the data points (see Sect. 4). Thus, given our initial motivation, Type 3 similarity measures seem to be the most promising type of similarity measure to focus on. However, it is worth investigating similarity measures of Type 4, to see if the added benefit of learning $C(x, y)$ outweighs the added training time or if it is possible to make it symmetric (as defined in the previous section) so that the training time could become similar to Type 3.

Thus, we will focus on summarizing related work from Type 3 similarity measures, but also add relevant work from Type 1, Type 2 and Type 4 for reference.

Type 1 is a type of similarity measure which is manually constructed. A general overview and examples of this type of similarity measure can be found in [7]. Nikpour et al. [23] presented an alternative method which includes enrichment of the cases/data points via Bayesian networks.

Type 2

In Type 2 similarity measures, only the binary $C(x, y)$ operator of the similarity measure $\mathbb{S}(x, y)$ is learned, while $G(\cdot)$ is either modeled or left as the identity function ($G(x) = I(x) = x$). Stahl et al. have done a lot of work on learning Type 2 similarity measures from data or user feedback. In all of their work, they formulate $C(x, y) = \sum w_i * sim_i(x_i, y_i)$ where for each feature i , sim_i is the local similarity measure and w_i is the weight of that feature. In [27], Stahl et al. described a method for learning the feature weights.

In [28], Stahl et al. introduced learning local similarity measures through an evolutionary algorithm (EA). First, they learn attribute weights (w_i) by adopting the method previously described in [27]. Then, they use an EA to learn the local similarity measures for each feature ($sim_i(x, y)$). In [29], Stahl and Gabel presented work where they learn weights of a modeled similarity measure and the local similarity for each attribute through an ANN. Reategui et al. [24] learned and represented parts of the similarity functions ($C(\hat{x}, \hat{y})$) through ANN. Langseth et al. [18] learned similarity knowledge ($C(\hat{x}, \hat{y})$) from data using Bayesian networks, which still partially relies on modeling the Bayesian networks with domain knowledge.

Abdel-Aziz et al. [1] used the distribution of case attribute values to inform a polynomial local similarity function, which is better than guessing when domain knowledge is missing. So this method extracts statistical properties from the dataset to parametrize $C(\hat{x}, \hat{y})$.

Gabel and Godehardt [10] used a neural network to learn a similarity measure. Their work is done in the context of case-based reasoning (CBR) which uses the measure to retrieve similar cases. They concatenate the two data points into one input vector. Thus, in the context of our framework $G(\cdot)$ is modeled as a identify function $I(x) = x$ and $C(x, y)$ is learned.

Maggini et al. [21] used SIMNNs which they also see as a special case of the Symmetry Networks [26] (SNs). In SIMNNs, $C(\hat{x}, \hat{y})$ and $G(\cdot)$ are both a function of both x and y data points and there is thus no distinct $G(\cdot)$. They also have a specialized structure imposed on their network to make sure the learned function is symmetric. SIMNN is in essence an extended version of a siamese neural network, but without a distinct distance layer usually present in SNN architectures. They focus on the specific properties of the

network architecture and the application of such networks in semi-supervised settings such as k-means clustering. The pair of data points ($\mathbf{x}$ and $\mathbf{y}$) are being compared two times, the first time at the first hidden layer and then at the output layer. Since there are no learnable parameters before this comparison, all the learning is done in $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $G(\mathbf{x})$ is the activation function of the input layer.

Type 3 One way of looking at a similarity measure is as an inverse distance measure, as similarity is the semantic opposite of distance. There has been much work on learning distance measures. Most of this work can be categorized as a Type 3 similarity measure as the learning tasks only aim to learn the embedding function $G(\cdot)$ and then combine the output of this function with a static $C(\cdot)$ (e.g., a $L2$ norm function). The most well-known instance of a Type 3 learned distance measure is siamese neural networks (SNNs); it is highly related to the Type 2 similarity measure by Maggini et al.'s Similarity neural networks (SIMNN) [21].

The main characteristic of SNNs is sharing the weights between the two identical neural networks. The data points we want to measure the similarity for are then input to these networks. This frees the learning algorithm of learning two sets of weights for the same task. This was first used in [5] (using $C(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \cos(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $G(\cdot)$ being learned from data) to measure similarity between signatures. Similar architectures are also discussed in [26].

Chopra et al. [6] used a SNN for face verification and pose the problem as an energy-based model. The output of the SNN is combined through a $L1$ norm (absolute value norm $C(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \|\hat{\mathbf{x}} - \hat{\mathbf{y}}\|$) to calculate the similarity. They emphasize that using a $L2$ norm (Euclidean distance) as part of the loss function would make the gradient too small for effective gradient descent (i.e., create plateaus in the loss function). This work is closely related to Hadsell et al. [11], where they explain the contrastive loss function used for training the SNN (also used in [6,22]) by analogy of a spring system.

Related to this, Vinyals et al. [31] used a similar type of setup for matching an input data point to a support set. It is framed as a discriminative task, where they use two neural networks to parametrize an attention mechanism. They use these two networks to embed the two data points into a feature space where the similarity between them is measured. However, in contrast to SNNs and SIMNNs, their two networks for embedding the data points are not identical, as one network is tailored to embed a data point from the support set, but also given the rest of the support set. Thus, the embedding of the support set data point is also a function of the rest of the support set. With $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ being modeled as a cosine softmax, this is similar to the examples of Type 3 similarity measures mentioned previously (e.g., [4,5]). However, a major difference is that signal extraction functions are not equal: $\mathbb{S}(\mathbf{x}, \mathbf{y}) = C(f(\mathbf{x}), g(\mathbf{y}))$ with $f(\mathbf{x}) \neq g(\mathbf{x})$ (only

stating that $f(\cdot)$ may potentially equal $g(\cdot)$). Since $f(\cdot)$ and $g(\cdot)$ are not sharing weights between them, the architecture is variant (or asymmetric) to the ordering of input pairs. Thus, the architecture needs up to twice as much training to achieve the same performance as a SNN.

In much of the same fashion as Chopra et al. did in [6], Berlemont et al. [4] used SNNs combined with an energy-based model to build a similarity measure between different gestures made with smart phones. However, they adapt the error estimation from using only separate positive and negative pairs to a training subset including a reference sample, a positive sample and a negative sample for every other class. They train $G(\cdot)$ while keeping a static $C(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \cos(\hat{\mathbf{x}}, \hat{\mathbf{y}})$. This training method of using triplets for training SNNs was also described by Lefebvre et al. [20]. A similar approach can be seen in Hoffer et al. [13]; however, they do not use a set of negative examples per reference point for each class as Berlemont et al. did. Instead, they use triples of $(\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-)$, $\mathbf{x}$ being the reference point, $\mathbf{x}^+$ being the same class and $\mathbf{x}^-$ being a different class.

Koch et al. [16] used a convolutional siamese network (CSN), with $G(\cdot)$ implemented as a CNN and $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ implemented as $L1(\hat{\mathbf{x}}, \hat{\mathbf{y}})$. This is done in a semi-supervised fashion for one-shot learning within image recognition. They learn this CSN for determining if two pictures from the Omniglot [17] dataset are within the same class. The model can then be used to classify a data point representing an unseen class by comparing it to a repository of class representatives (support set).

CSNs are also used in the context of CBR by Martin et al. [22] to represent a similarity measure in a CBR system. The CSN is trained with pairs of cases, and the output is their similarity. During training, they have to label pairs of cases as 'genuine' (both cases belong to the same class) or 'impostor' (the cases belong to different classes). This requires that the user has a clear boundary for the classes. In relation to our framework, this similarity measure learns $G(\cdot)$, while $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is static, with $G(\cdot)$ implemented as a convolutional neural network, and $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ implemented as Euclidean distance ($L2$ norm).

In general, using SNNs for constructing similarity measures has a major advantage as you can easily adopt pre-trained models for $G(\cdot)$ to embedding/preprocess the data points. For example, to train a model for comparing two images one could use ResNet [12] for $G(\cdot)$ and then use the $L1$ norm as $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$. This would be a very similar approach to the similarity measure used by Koch et al. [16] with $\mathbb{S}(\mathbf{x}, \mathbf{y}) = \|(G(\mathbf{x}), G(\mathbf{y}))\|_1$, the main difference being that $G(\cdot)$ is designed for bigger pictures.

There are only very few examples of Type 4 similarity measures in the literature. In Zagoruyko and Komodakis's work [33], they investigate different types of architectures for learning image comparison using convolutional neural

networks. In all of the architectures, they evaluate $C(\hat{x}, \hat{y})$ is learned, but in some of these architectures $G(\cdot)$ is not symmetric, i.e., $\mathbb{S}(x, y) = C(G(x), H(y))$ where $G(x) \neq H(x)$. Arandjelović and Zisserman's work [2] used a very similar method to many Type 3 similarity measures for calculating similarity. However, their input data are always pairs of two different data types and are as such different from most of the other relevant work leaving $G(\cdot)$ unsymmetrical just as in Zgoruyko et al. [33] and Vinyals et al. [31]. In contrast to the Type 3 similarity measures including [31], Arandjelović et al. also learned $C(\hat{x}, \hat{y})$, which they call a fusion layer.

All similarity measures of Type 3 we found in the literature use a loss function that includes feedback from the binary operator part of $\mathbb{S}(C(\hat{x}, \hat{y}))$. In the case of SNNs, even if $C(x, y)$ is non-symmetric ($C(x, y) \neq C(y, x)$), the loss for each network would be equal as they are equal and share weights. That means that ordering of the two data points being compared during training has no effect, i.e., the training effect of (x, y) is equal to that of (y, x) . This means a lot of saved time during training, as the training dataset could be halved without any negative effect on performance.

However, the implementation of $C(\hat{x}, \hat{y})$ would then decide how much training one would need to adapt a pre-trained model from classifying single data points to measuring similarity between them. One could view the process of starting with a pre-trained model for the dataset and then training the model with loss coming from $C(\hat{x}, \hat{y})$ as adapting the model from classification to similarity measurement.

One way of creating a Type 3 similarity measure using a minimal amount of training would be to pre-train a network on classifying individual data points and then apply that network as $G(\cdot)$ that feeds into a $C(\hat{x}, \hat{y}) = \|\hat{x} - \hat{y}\|$ in a similarity measurement. Evaluation of such a similarity measurement has not been reported in the literature, and such a similarity will be explored in the next section.

4 Method

The framework presented in Sect. 2 and the subsequent analysis of previous relevant work presented in Sect. 3 shows that there are unexplored opportunities within research on similarity measurements.

Given the initial motivation, we seek methods that work well in domains where domain knowledge is very resource demanding. This requires that as much as possible of the similarity measure $\mathbb{S}(x, y) = C(G(\hat{x}), G(\hat{y}))$ is learned from data rather than modeled from domain knowledge. There are some exceptions to this, such as applying general binary operations, such as norms (e.g., $L1$ or $L2$ norm), on the two data points ($\hat{x}$ and $\hat{y}$) preprocessed by $G(\cdot)$. In these cases, there is little domain expertise involved in designing $C(\hat{x}, \hat{y})$

other than intuition that the similarity of two data points is closely related to the norm between $\hat{x}$ and $\hat{y}$.

The most promising type of similarity measures from this perspective is Type 3 and Type 4 where $G(\cdot)$ is learned in Type 3 and both $C(x, y)$ and $G(\cdot)$ are learned in Type 4. However, to test any new design we need to have reference methods to compare against. For reference, we chose to implement one Type 1 similarity measure, two similarity measures of Type 2 (including Gabel et. al's similarity measures and Chopra et. al's Type 3 similarity measure. The Type 1 similarity measure uses a similarity measure that weights each feature uniformly. The Type 2 is identical to the Type 1 similarity measure, except that it uses a local similarity function for each feature which is parametrized by statistical properties of the values of that feature in the dataset.

One unexplored direction of creating similarity measures is creating a SNN similarity measure (Type 3) through training $G(\cdot)$ as a classifier on the dataset later being used for measuring similarity and then using that trained $G(\cdot)$ to construct a SNN similarity measure. This is in contrast to the usual way of training SNNs (as seen in, e.g., [5,6]) where the loss function is a function of pairs of data points, not single data points. The motivation for exploring this type of design is that it shows the similarity measuring performance of using networks pre-trained on classifying data points directly as part of a SNN similarity measure. This is detailed in Sect. 4.2.

Finally, we will explore Type 4 similarity measures which have seen little attention in research so far. To make our design as symmetric as possible, we will use the same design as SNNs for $G(\cdot)$ and introduce novel design to also make $C(\hat{x}, \hat{y})$ symmetric. That way our design is fully symmetric (invariant to ordering of the input pair), and thus, training becomes much more efficient. All of the details of this design are shown in Sect. 4.3. Both of our proposed similarity methods implement $G(\cdot)$ as neural networks. The Type 4 measurement design implements $C(\hat{x}, \hat{y})$ as a combination of a static binary function and neural network.

4.1 Reference similarity measures

As a reference for our own similarity measure, we implemented several reference similarity measures of Type 1, Type 2 and Type 3. First, we implemented a standard uniformly weighted global similarity ($t_{1,1}$) measure which can be defined as:

$$t_{1,1}(x, y) = \mathbb{S}(x, y) = C(x, y) = \sum_i^M w_i \cdot sim_i(x_i, y_i), \quad (2)$$

where $sim_i(x_i, y_i)$ denotes the local similarity of the i th of M attributes. In $t_{1,1}$, all weights and local similarity measures are uniformly distributed and not parametrized by the data.

Fig. 3 Architecture of a ANN similarity measure as used in Gabel [10] (Type 2), where $G(\cdot)$ is set to be the identity function $G(x) = I(x) = x$

We extended this with a Type 2 similarity measure $t_{2,1}$, which is based on the work from Abdel-Aziz et al. [1], where the local similarity measures are parametrized by the data from the corresponding features.

Furthermore, we implemented a Type 2 similarity measure *gabel* as described by Gabel et al. [10]. The architecture of *gabel* is shown in Fig. 3.

Lastly, we implemented the Type 3 similarity measure *chopra* described by Chopra et al. We did not implement the extension done to the contrastive loss function as seen in [4,20] as the change in the training dataset would be too big. This change would make comparisons between the methods harder to justify. Also none of these works showed any comparisons with previous SNNs in terms of any increased performance in relation to regular contrastive loss.

4.2 Type 3 similarity measure

In this subsection, we will detail how we model the Type 3 similarity measure $t_{3,1}$ which uses an embedding function $G(\cdot)$ trained as a classifier. This embedding function maps the input point, x , to an embedding space (see Fig. 2) in which dimensions represent the probabilities of x belonging to a class. We then model the similarity measure between two points as the a static function ($C(\cdot)$) between their two respective embeddings.

For this, we choose the $L2$ norm. So replacing $C(\cdot)$ for $L2$ in Eq. 1: $C(\hat{x}, \hat{y}) = \|\hat{x} - \hat{y}\|_2$, we can redefine Eq. 1 to be:

$$\mathbb{S}(x, y) = t_{3,1}(x, y) = 1.0 - \|G(x) - G(y)\|_2 \quad (3)$$

where $G(\cdot)$ outputs the modeled solution as a n dimensional vector (the feature vector output from the network to the softmax function for n classes) for the case based on the problem attributes of data point x . This means that if the $G(x)$ evaluates the two cases as very similar in terms of classification $G(x) \approx G(y)$ and $\|G(x) - G(y)\| \approx 0$, then $\mathbb{S}(x, y) \approx 1.0$. This architecture is also illustrated in Fig. 4

Fig. 4 Architecture of the $t_{3,1}$ similarity measure where $G(\cdot)$ is trained to output softmax vectors for classification and the similarity is calculated as a modeled $L2$ norm between these two vectors (Type 3)

Following the model for the $t_{3,1}$ similarity measure, we define the loss estimate as log-loss between $G(x) = \hat{x}$ and t , where t is the true classification softmax vector; $\hat{x}$ is the class probability vector output from $G(x)$. Notice that the error estimate of $t_{3,1}$ does not depend on the output of $C(\hat{x}, \hat{y})$.

A dataset of size N would then be defined as:

$$T = \left[(x^1, t^1) \dots (x^N, t^N) \right], \quad (4)$$

where x^N is the problem part of the N -th data point and t^N is the solution/target part of the N -th data point.

If the relation between the problem part of the data point (x) and the solution part of the data point (t) is complex, the network architecture needs to be able to represent the complexity and any generalizations of patterns in that complexity.

4.3 Type 4 similarity measure

As previously explained, Type 4 similarity measures are currently the most unexplored type of similarity measure. It is also the type of similarity measure that requires the least amount of modeling. In principle, Type 4 similarity measures learn two things: $G(\cdot)$ learns a useful embedding, where the most useful parts of x and y are encoded into $\hat{x}$ and $\hat{y}$. $C(\hat{x}, \hat{y})$ learns how to combine those embeddings to calculate the similarity of the original x and y .

In Type 4 similarity measures, both $C(\hat{x}, \hat{y})$ and $G(\cdot)$ are learned. In our Type 4 similarity method, we will use an ANN to represent both $G(\cdot)$ and $C(\hat{x}, \hat{y})$. This has the advantage that the learning on $\mathbb{S}(x, y)$ is an end-to-end process. The loss computed after $C(\hat{x}, \hat{y})$ can be used to compute gradients for both $C(\hat{x}, \hat{y})$ and $G(\cdot)$. $C(\hat{x}, \hat{y})$ will learn the binary combination best suited to calculate the similarity of the two embeddings, while $G(\cdot)$ will learn to embed the two data points optimally for calculating their similarity in $C(\hat{x}, \hat{y})$. In principle, any ML method could be used to learn $G(\cdot)$ and $C(\hat{x}, \hat{y})$, but not all ML methods lend themselves naturally

Fig. 5 Architecture of a *eSNN* where we combine the symmetry of SNNs with the ability to learn $C(\hat{x}, \hat{y})$. $C(\hat{x}, \hat{y})$ is expanded in this picture to highlight the $ABS(\hat{x} - \hat{y})$ operation done as the first operation of $C(\hat{x}, \hat{y})$ to keep C invariant to the ordering of inputs. It also illustrates the two additional loss signals to $G(\cdot)$ which helps train the similarity measure

to back-propagating the error signal from $C(\hat{x}, \hat{y})$ through $G(\cdot)$ and back to the input.

We define our Type 4 similarity method, extended siamese neural network (*eSNN*) as shown in Fig. 5.

Given that this similarity method outputs similarity and the loss function is a function of the input, we get a new general loss function for similarity, defined per data point as follows:

$$L_s(\mathbf{x}, \mathbf{y}, s) = |s - C(G(\mathbf{x}), G(\mathbf{y}))|, \quad (5)$$

where s is the true similarity of cases $\mathbf{x}$ and $\mathbf{y}$. Since this loss function is dependent on pairs of data points and the true similarity between them, we need to create a new dataset based on the original dataset. This new dataset consists of triplets of two different data points from the original dataset and the true similarity of these two data points:

$$\mathbf{T} = \left[(\mathbf{x}^1, \mathbf{y}^1, s^1) \dots (\mathbf{x}^N, \mathbf{y}^N, s^N) \right], \quad (6)$$

where s^N is 1 if $\mathbf{x}^N$ and $\mathbf{y}^N$ belong to the same class and 0 otherwise.

It is worth mentioning that this dataset is of size $N(N-1)$ for the similarity measure to train on all possible combinations of the N data points. Certain similarity measure architectures (e.g., *gabel* from Gabel et al. [10] or Zagoruyko et al.'s similarity measures [33]) need to train on a dataset containing all possible combinations of data points (of size $N(N-1)$) as training on the triplet $(\mathbf{x}, \mathbf{y}, s)$ does not guarantee that the model learns that $\mathbb{S}(\mathbf{y}, \mathbf{x}) = s$. Thus, the training dataset must also include the triplet $(\mathbf{y}, \mathbf{x}, s)$. However, this may be largely avoided by using architectures (such as those seen in SNNs and SNs) that exploit symmetry and weight sharing. To achieve this, we modeled $C(\mathbf{x}, \mathbf{y})$ as a ANN where the first layer is an absolute difference operator on

two vectors: $\mathbf{z} = ABS(\hat{\mathbf{x}} - \hat{\mathbf{y}})$. where $\mathbf{z}$ is the element-wise absolute difference between $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$. The rest of $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is hidden layers of ANN that operate on $\mathbf{z}$. This way $C(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ becomes invariant to the ordering of inputs to $S(\mathbf{x}, \mathbf{y})$. Consequently, the model only needs to train on order-invariant unique pairs of data points, reducing training set size from $N(N-1)$ to $N(\frac{N}{2}-1)$. The resulting architecture of *eSNN* is shown in Fig. 5.

In Sect. 4.2, we argue why $G(\cdot)$ trained to correctly classify its input is a good embedding function for calculating similarity. As a result, we added two loss signals to $G(\cdot)$ during training. These loss signals are calculated from the difference between the embedding of the data point produced by $G(\cdot)$ and the correct softmax classification vector.

This also introduced an opportunity for exploring the relative importance of the embedding function $G(\cdot)$ and the binary similarity function $C(\cdot)$ in terms of the performance of the similarity measure. This could be done by weighting the three different loss signals ($\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ and similarity as shown in Fig. 5) during training and measuring the effect of that weighting on the performance. We define our weighted loss function as such:

$$L(\alpha, \mathbf{x}, \mathbf{y}, s) = \frac{(1-\alpha)}{2} \cdot (L_c(\mathbf{x}, \mathbf{t}_x) + L_c(\mathbf{y}, \mathbf{t}_y)) + \alpha \cdot L_s(\mathbf{x}, \mathbf{y}, s), \quad (7)$$

where $L_s(\cdot)$ is defined in Eq. 5, $\mathbf{t}_x$ is the true label of data point $\mathbf{x}$, $\mathbf{t}_y$ is the true label of data point $\mathbf{y}$ and $L_c(\mathbf{v}_1, \mathbf{v}_2)$ is the categorical cross-entropy loss between two softmax vectors. We use this formula and tested with different 100 different values of α in the range $[0, 1]$ to find the weighting scheme best for performance. The results are shown in Fig. 6.

Fig. 6 Showing results from weighting the three different outputs in terms of signal strength to loss measured on the UCI dataset balance scale [8] (5-fold cross-validation and repeated 5 times). This measurement was done using training data of size $N(\frac{N}{2}-1)$. The effect of α is much less impactful on the validation result after 200 or more epochs of training when training on $N(N-1)$ datasets. However, choosing the correct α using $N(\frac{N}{2}-1)$ datasets does impact the speed of training for *eSNN* when training on $N(N-1)$ datasets

Fig. 7 Testing how the RProp algorithm performs in comparison with ADAM and RMSProp. Our proposed architecture performs best using the RProp algorithm (fivefold cross-validation and repeated 5 times)

Figure 6 indicates that $\alpha = 0.15$ is ideal for this dataset. We have used $\alpha = 0.15$ throughout the experiments reported in Sect. 5.

4.4 Network parameters

For all similarity measures tested using ANN and all datasets except MNIST, $G(\cdot)$ and $C(\cdot)$ were implemented with two hidden layers of 13 nodes. This was done to replicate the network parameters used by Gabel et al. to ensure we had comparable results. For the MNIST dataset test, both *chopra* and *eSNN* used three hidden layers of 128 nodes for $G(\cdot)$ and the same for $C(\cdot)$.

Other than the network architecture, we also wanted to choose which optimizer to use for learning the ANN model. We wanted to choose the RProp [25] to be more comparable with the results from Gabel et al. which also used RProp optimizer. Our tests seen in Fig. 7 show that RProp outperforms all other optimizers tested (ADAM and RMSProp). This is consistent with the results reported by Florescu and Igel [9]. This should hold true until the run-time performance of RProp degrades with dataset size, as RProp uses full batch sizes.

4.5 Evaluation protocol and implementation

The different similarity measures presented earlier in this section require different training datasets. The reference Type 1 similarity measures ($t_{1,1}$) require no training, while $t_{2,1}$ and $t_{3,1}$ do not require a similarity training consisting of triplets as described in Eq. 6. All other similarity measures evaluated were trained using identical training datasets. As a result, all similarity measures were trained on a dataset consisting of all possible combinations of data points (as explained in Sect. 4.3) as this is required by the *gabel* similarity measure. However, results highlighting the differences in training

performance when using the different training datasets are shown in Fig. 13.

The results reported in the next section are all fivefold stratified cross-validation repeated 5 times for robustness. The performance reported is an evaluation of each similarity measurement using the part of the dataset (validation partition) that was not used for training. Using the similarity measure being evaluated, we computed the similarity between every data point in the validation partition (V) and every data point in the training partition (T). For each validation data point ($x_v \in V$), we find the data point in the training set T with the highest similarity ($x_t = \arg \max_{x_i \in T} (\mathbb{S}(x_v, x_i))$). If x_t has the same class as x_v from the validation partition, we scored it as 1.0; if not, we scored it as 0.0.

The implementation was done in Keras¹ with Tensorflow as backend. The methods were measured on 14 different datasets available from the UCI machine learning repository [8]. Results were recorded after 200 epochs and 2000 epochs (the latter number to be consistent with Gabel et al. [10]) to reveal how fast the different methods were achieving their performance.

5 Experimental evaluation

To calculate the performance of our similarity measure, we chose to use the same method of evaluation as Gabel et al. [10] to make the similarity metrics more easily comparable. In addition, this evaluation method of using publicly available datasets from the UCI machine learning repository [8] makes the results easy to reproduce. We selected a subset of the original 19 datasets, choosing not to use regression datasets, resulting in a set of 14 classification datasets. The datasets' numerical features were all normalized; categorical features were replaced by a one-hot vector.

The validation losses from evaluating the similarity measures on the 14 datasets are shown in Figs. 8 and 9. Figure 8 shows the results after training for 200 epochs, while Fig. 9 shows the results after 2000 epochs. This has been done to illustrate how the differences between the similarity measures develop during training. In addition, the 200 and 2000 epoch runs are independent runs (i.e., Fig. 9 shows not the same models as seen in Fig. 8 1800 epochs later)

The numbers that are the basis of these figures are also reported in Table 2 for 200 epochs and Table 3 for 2000 epochs. The tables are highlighted to show the best result per dataset. In some cases, the differences between two methods for one dataset were smaller than the standard deviation, thus highlighting more than one result.

¹ Code available at NTNU OpenAI lab github page: <https://github.com/ntnu-ai-lab>.

Fig. 8 Performance of $eSNN$ in comparison with reference similarity measures and state-of-the-art similarity methods over all test datasets trained over 200 epochs

Fig. 9 Performance of $eSNN$ in comparison with reference similarity measures and state-of-the-art similarity methods over all test datasets trained over 2000 epochs

Finally, to illustrate that $eSNN$ scales to larger datasets we report results from the MNIST dataset in Fig. 10. The MNIST results are not validation results, as calculating the similarity between all the data points in the test set and the training set (as per the evaluation protocol described in Sect. 4.5) was too resource demanding.

Table 2 shows the validation losses of the different similarity measures on the different datasets. Our proposed Type 4 similarity measure $eSNN$ has 11% less validation loss than the second best (Type 3) similarity measure $chopra$ (Chopra et al. [6]). The other Type 3 similarity measures follow with $t_{3,1}$ having 51% higher loss and $gabel$ (Gabel et al. [10]) with 52% more loss. The Type 1 similarity measure had 61%

Table 2 Validation retrieval training, in comparison with state-of-the-art methods

	<i>eSNN</i>	<i>chopra</i>	<i>gabel</i>	$t_{3,1}$	$t_{1,1}$	$t_{2,1}$
bal	0.01	0.00	0.14	0.10	0.42	0.81
car	0.04	0.02	0.19	0.16	0.25	0.25
cmc	0.52	0.53	0.54	0.55	0.54	0.58
eco	0.22	0.20	0.46	0.35	0.21	0.22
glass	0.08	0.08	0.12	0.10	0.06	0.07
hay	0.19	0.21	0.26	0.17	0.33	0.37
heart	0.21	0.24	0.28	0.24	0.24	0.23
iris	0.04	0.03	0.18	0.07	0.05	0.04
mam	0.21	0.25	0.26	0.27	0.28	0.29
mon	0.28	0.33	0.39	0.45	0.29	0.29
pim	0.28	0.30	0.35	0.35	0.31	0.32
ttt	0.03	0.03	0.17	0.07	0.32	0.07
use	0.07	0.08	0.08	0.39	0.21	0.18
who	0.29	0.45	0.33	0.45	0.46	0.45
Sum	2.47	2.75	3.75	3.72	3.97	4.17
Average	0.18	0.20	0.27	0.27	0.28	0.30

The best result for each dataset is highlighted in bold
eSNN has the smallest loss in 8 of 14 datasets

more loss but managed to be the best similarity measure for the glass dataset. At last, Type 2 similarity measure had 69% higher loss than *eSNN* on average.

The results when training for 2000 epochs are quite different from those at 200 epochs, as seen by how much closer the other similarity measures are in Fig. 9 than in Fig. 8. *eSNN* still outperforms all other similarity measures on average, but the second best similarity measure $t_{3,1}$ is much closer with just 6.9% higher loss. *gabel* is 11.8% worse, *chopra* is 14.7% worse, $t_{1,1}$ is 61.2% worse, and finally, $t_{2,1}$ is 69% worse than *eSNN*.

The gap between *eSNN* and the state of the art is considerable at 200 epochs. This gap shrinks from 11% at 200 epochs to 6.9% at 2000 epochs, which is still a considerable difference.

To illustrate the difference in terms of training efficiency between different types similarity measure, we show the validation loss for *gabel*, *chopra* and *eSNN* during training. Specifically, for each epoch we test the loss of each similarity measure by the same method as described in Sect. 4.5. Figures 11 and 12 shows validation loss during training of *eSNN*, *chopra* and *gabel* on the UCI Iris and Mammographic mass datasets [8] respectively. This exemplifies the training performance of these methods in relation to the Iris and Mammographic mass dataset results reported in Table 1 and Table 2. One can also note that in training for the Mammographic mass dataset as seen in Fig. 11 *chopra* never achieves the same performance as *eSNN*. In contrast, while training on the Iris dataset (as seen in Fig. 12), which is a less com-

Table 3 Validation retrieval loss after 2000 epochs of training, in comparison with state-of-the-art methods

	<i>eSNN</i>	<i>chopra</i>	<i>gabel</i>	$t_{3,1}$	$t_{1,1}$	$t_{2,1}$
bal	0.02	0.00	0.08	0.01	0.43	0.83
car	0.01	0.01	0.06	0.02	0.24	0.24
cmc	0.52	0.53	0.54	0.53	0.54	0.58
eco	0.22	0.20	0.22	0.18	0.19	0.21
glass	0.06	0.07	0.08	0.09	0.05	0.06
hay	0.18	0.21	0.20	0.15	0.32	0.34
heart	0.21	0.27	0.23	0.22	0.24	0.23
iris	0.08	0.05	0.07	0.04	0.06	0.05
mam	0.21	0.27	0.25	0.27	0.29	0.28
mon	0.26	0.30	0.33	0.27	0.32	0.32
pim	0.27	0.31	0.25	0.30	0.30	0.31
ttt	0.03	0.03	0.07	0.03	0.32	0.08
use	0.08	0.10	0.07	0.08	0.18	0.16
who	0.30	0.46	0.29	0.43	0.47	0.45
Sum	2.45	2.81	2.74	2.62	3.95	4.14
Average	0.18	0.20	0.20	0.19	0.28	0.30

The best result for each dataset is highlighted in bold
eSNN has the smallest validation retrieval loss in 6 of 14 datasets in addition to the lowest average loss

Fig. 10 Training loss (not validation retrieval loss) during training on the MNIST dataset for *chopra* and *eSNN*. *gabel* could not be evaluated as training on a $N(N - 1)$ -sized dataset for MNIST is too resource demanding

plex dataset than the Mammographic mass dataset, *chopra* achieves the same performance as *eSNN*.

Figure 13 shows the validation loss during training when *chopra* and *eSNN* are using a training dataset of size N and *gabel* is using a training dataset of size $N(N - 1)$. This figure illustrates how much fewer evaluations a SNN similarity measure like *chopra* or symmetric Type 4 similarity measure such as *eSNN* needs than a similarity measurement that is not invariant to input ordering, while still having excellent relative performance.

Fig. 11 Validation retrieval loss during training on the Mammograph UCI ML dataset [8]. The figure shows that the mammograph dataset is a dataset that needs learning outside of embedding via $G(\cdot)$. *chopra* starts out good as $C(\hat{x}, \hat{x})$ is already designed as the $L1$ norm. However, *eSNN* and *gabel* catch up when it learns an equivalent and better $C(\hat{x}, \hat{x})$ function

Fig. 12 Validation retrieval loss during training on the Iris UCI ML dataset [8]. Since *chopra* starts out with very low validation loss, it seems probable that the static $L1$ norm $C(\hat{x}, \hat{x})$ used by *chopra* is close to optimal for correctly identifying if the two data points belong to the same class or not. The performance increase done by *chopra* is a slight optimization of $G(\cdot)$. The performance increase done during training by *gabel* and *eSNN* is mainly by learning a $C(\hat{x}, \hat{x})$ equivalent in function to that used by *chopra*, and secondary a slight optimization of $G(\cdot)$. *eSNN* catches up to *chopra* in performance after around 20 epochs; however, *gabel* takes longer (5% validation loss at 2000 epochs) as shown in Table 3

Finally, in Figs. 14 and 15 we show how *eSNN* can be used for semi-supervised clustering. The figures show PCA and T-SNE clustering of embeddings produced untrained and trained *eSNN* networks, respectively, from the MNIST dataset. The embeddings are the vector output of $G(\cdot)$ for each of the data points in the test set. The embeddings shown are computed from a test set that is not used for training. The

Fig. 13 Validation retrieval loss during training on the balance dataset, which illustrates the difference in amount of evaluations needed to achieve acceptable performance. *Chopra* achieves good performance very quickly, but is outperformed by *eSNN* soon. Both have very good performance before having evaluated less (N) data points than used by *gabel* ($N(N-1)$)

(a) PCA clustering on the MNIST dataset before training

(b) PCA clustering on the MNIST dataset after training

Fig. 14 PCA clustering showing the two first principal components ($PCA1$ and $PCA2$) of the embeddings produced by *eSNN* from MNIST input before (a) and after (b) training

(a) T-SNE clustering of the MNIST dataset before training

(b) T-SNE clustering of the MNIST dataset after training

Fig. 15 T-SNE clustering of embeddings produced by *eSNN* from MNIST input before (a) and after (b) training

figures show that *eSNN* learns a way to correctly cluster data points that it has not used for training.

6 Conclusions and future work

Section 5 shows that all of the learned similarity measures outperformed the classical similarity measure $t_{1,1}$ and also $t_{2,1}$ where the local (per feature) similarity measures were adapted to the statistical properties of the features [1]. In practice, one should weight the importance of each feature according to how important it is in terms of similarity measurement. In many situations, the number of possible attributes to include in such a function can be overwhelming, and modeling them in the way we did in $t_{1,1}$ and $t_{3,1}$ also overlooks possible covariations between the attributes. Both of these problems can be addressed using the proposed method to model the similarity using machine learning on a dataset that maps from case problem attributes to case solution attributes.

However, one should be careful to note that all of the learned similarity measures are built on the assumption that similar data points have similar target values ($\delta_s \approx \delta_e \approx \delta_p$ in Fig. 2). If this assumption does not hold, learning the similarity measure might be much more difficult.

We have also presented a framework for how to analyze and group different types of similarity measures. We have used this framework to analyze previous work and highlight different strengths and weaknesses of the different types of similarity measures. This also highlighted unexplored types of similarity measures, such as Type 4 similarity measures.

As a result, we designed and evaluated a Type 3 similarity measure $t_{3,1}$ based on a classifier. The evaluations showed that using a classifier as a basis for a similarity measure achieves comparable results to state-of-the-art methods, while using much less training evaluations to achieve that performance.

We then combined strengths from Type 4 and Type 3 similarity measures into a new Type 4 similarity measure, called extended siamese neural networks (*eSNN*), which

- Learns an embedding of the data points using $G(\cdot)$ in the same way as Type 3 similarity measures, but using shared weights in the same way as SNNs to make the operation symmetrical.
- Learns $C(\hat{x}, \hat{y})$, thus enabling extended performance in relation to SNN and other Type 3 similarity measurements.
- Restricts $C(\hat{x}, \hat{y})$ to make it invariant to input ordering and thus obtaining end-to-end symmetry through the similarity measure.

Keeping *eSNN* symmetrical end to end enables the user of this similarity measure to train on much smaller datasets than required by other types of similarity measures. Type 3 measures based on SNNs also have this advantage, but our results show that the ability to learn $C(\hat{x}, \hat{y})$ is important for performance in many of the 14 datasets we tested. Our results showed that *eSNN* outperformed state-of-the-art methods on average over the 14 datasets by a large margin. We also demonstrated that *eSNN* achieved this performance much faster given the same dataset than current state of the art. In addition, the symmetry of *eSNN* enables it to train on datasets that are orders of magnitude smaller. Our case study of clustering embeddings produced from *eSNN* shows that the *eSNN* model can be used for semi-supervised clustering.

Finally, we demonstrated that the training of this similarity measure scales to large datasets like MNIST. Our main motivation for this work was to automate the construction of similarity measures while keeping training time as low as possible. We have shown that *eSNN* is a step toward this. Our evaluation shows that it can learn similarity measures across a wide variety of datasets. We also show that it scales well

in comparison with similar methods and scales to datasets of some size such as MNIST.

The applications for *eSNN* as a similarity measure are not only as a similarity measure in a CBR system. It can also be used for semi-supervised clustering: training *eSNN* on labeled data, then use the trained *eSNN* for clustering unlabeled data. In much the same fashion, it could be used for semi-supervised clustering, using *eSNN* as a matching network in the same fashion as the distance measure is applied in Vinyals et al. [31].

In continuation of this work, we would like to explore what is actually encoded by learned similarity measures. This could be done by varying the different features of a query data point q in $\mathbb{S}(x, q)$ and discovering when that data point would change from one class to another (when the class of the closest other data point changes)—this would form a multi-dimensional boundary for each class. This boundary could be explored to determine what the similarity measure actually encoded during the learning phase.

Another interesting avenue of research would be to apply recurrent neural networks to embed time series into embedding space (see Fig. 2) to enable the similarity measure to calculate similarity between time series which is currently a non-trivial problem.

The architecture of similarity measures still requires more investigation, e.g., is the optimal embedding from $G(\cdot)$ different from the softmax classification vector used in normal supervised learning? If so, it is worth investigating why it is different.

Acknowledgements We would like to thank the EXPOSED project (grant number 237790) and NTNU Open AI Lab for the support to do this work. Thanks also to Gunnar Senneset and Hans Vanhauwaert Bjelland for their great support during our work.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Abdel-Aziz, A., Strickert, M., Hüllermeier, E.: Learning solution similarity in preference-based CBR. In: International Conference on Case-Based Reasoning, pp. 17–31. Springer, Berlin (2014)
2. Arandjelovic, R., Zisserman, A.: Look, listen and learn. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 609–617. IEEE (2017)
3. Bergmann, R.: Experience Management: Foundations, Development Methodology, and Internet-Based Applications. Springer, Berlin (2002)
4. Berlemont, S., Lefebvre, G., Duffner, S., Garcia, C.: Siamese neural network based similarity metric for inertial gesture classification and rejection. In: 2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG), vol. 1, pp. 1–6. IEEE (2015)
5. Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a “siamese” time delay neural network. In: Advances in neural information processing systems, pp. 737–744 (1994)
6. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005. CVPR 2005, vol. 1, pp. 539–546. IEEE (2005)
7. Cunningham, P.: A taxonomy of similarity mechanisms for case-based reasoning. IEEE Trans. Knowl. Data Eng. **21**(11), 1532–1543 (2009)
8. Dheeru, D., Karra Taniskidou, E.: UCI machine learning repository. <http://archive.ics.uci.edu/ml> (2017). Accessed 1 June 2019
9. Florescu, C., Igel, C.: Resilient backpropagation (RPROP) for batch-learning in tensorflow. In: ICLR 2018 Workshop Permission Proceedings (to appear) (2018)
10. Gabel, T., Godehardt, E.: Top-down induction of similarity measures using similarity clouds. In: Hüllermeier, E., Minor, M. (eds.) Case-Based Reasoning Research and Development, pp. 149–164. Springer International Publishing, Cham (2015)
11. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06) IEEE, pp. 1735–1742 (2006)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
13. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: International Workshop on Similarity-Based Pattern Recognition, pp. 84–92. Springer, Berlin (2015)
14. Hüllermeier, E., Cheng, W.: Preference-based CBR: general ideas and basic principles. In: IJCAI, pp. 3012–3016 (2013)
15. Hüllermeier, E., Schlegel, P.: Preference-based CBR: first steps toward a methodological framework. In: International Conference on Case-Based Reasoning, pp. 77–91. Springer, Berlin (2011)
16. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: ICML Deep Learning Workshop, vol. 2 (2015)
17. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. Science **350**(6266), 1332–1338 (2015)
18. Langseth, H., Aamodt, A., Winnem, O.M.: Learning retrieval knowledge from data. In: Sixteenth International Joint Conference on Artificial Intelligence, Workshop ML-5: Automating the Construction of Case-Based Reasoners, pp. 77–82. Stockholm, Citeseer (1999)
19. Leake, D.B.: Case-Based Reasoning: Experiences, Lessons and Future Directions. MIT Press, Cambridge (1996)
20. Lefebvre, G., Garcia, C.: Learning a bag of features based nonlinear metric for facial similarity. In: 2013 10th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp. 238–243. IEEE (2013)
21. Maggini, M., Melacci, S., Sarti, L.: Learning from pairwise constraints by similarity neural networks. Neural Netw. **26**, 141–158 (2012)
22. Martin, K., Wiratunga, N., Sani, S., Massie, S., Clos, J.: A convolutional siamese network for developing similarity knowledge in the selfback dataset. In: Proceedings of the International Conference on Case-Based Reasoning Workshops, CEUR Workshop Proceedings, ICCBR (Organisers), pp. 85–94 (2017)
23. Nikpour, H., Aamodt, A., Bach, K.: Bayesian-supported retrieval in BNCreek: A knowledge-intensive case-based reasoning system. In: International Conference on Case-Based Reasoning, pp. 323–338. Springer, Berlin (2018)

24. Reategui, E.B., Campbell, J.A., Leao, B.F.: Combining a neural network with case-based reasoning in a diagnostic system. *Artif. Intell. Med.* **9**(1), 5–27 (1997)
25. Riedmiller, M., Braun, H.: A direct adaptive method for faster backpropagation learning: The RPROP algorithm. In: *IEEE International Conference on Neural Networks*, pp. 586–591. IEEE (1993)
26. Shawe-Taylor, J.: Symmetries and discriminability in feedforward network architectures. *IEEE Trans. Neural Netw.* **4**(5), 816–826 (1993)
27. Stahl, A.: Learning feature weights from case order feedback. In: *International Conference on Case-Based Reasoning*, pp. 502–516. Springer, Berlin (2001)
28. Stahl, A., Gabel, T.: Using evolution programs to learn local similarity measures. In: *International Conference on Case-Based Reasoning*, pp. 537–551 (2003)
29. Stahl, A., Gabel, T.: Optimizing similarity assessment in case-based reasoning. In: *Proceedings of the National Conference on Artificial Intelligence*, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, vol. 21, p. 1667 (2006)
30. Tversky, A.: Features of similarity. *Psychol. Rev.* **84**(4), 327 (1977)
31. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems*, pp. 3630–3638 (2016)
32. Wienhofen, L.W.M., Mathisen, B.M.: Defining the Initial Case-Base for a CBR Operator Support System in Digital Finishing, pp. 430–444. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-47096-2_29
33. Zagoruyko, S., Komodakis, N.: Learning to compare image patches via convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4353–4361 (2015)

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