






HPSCAN: Human Perception-Based Scattered Data Clustering

S. Hartwig^{1,*} , C. v. Onzenoodt¹ , D. Engel¹ , P. Hermosilla²  and T. Ropinski¹ 

*corresponding author: sebastian.hartwig@uni-ulm.de

¹Visual Computing Group, Ulm University, Germany

²Computer Vision Lab, Technical University of Wien, Austria

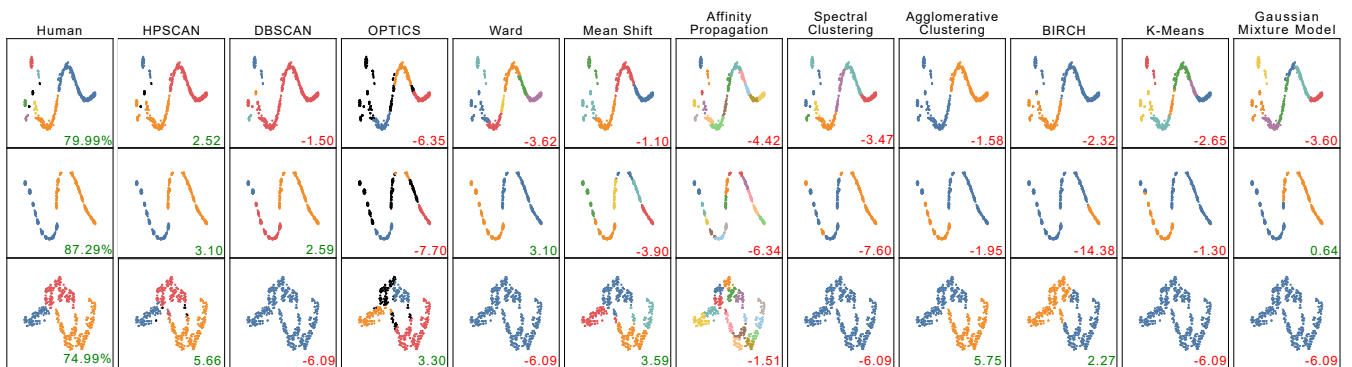


Figure 1: HPSCAN clusters scatterplots in accordance with human cluster perception. Since HPSCAN is trained on point-based clustering data rather than images, it infers clustering directly from the points. Here, we present a comparison of HPSCAN’s perception-aware results against those of several state-of-the-art clustering techniques, which are optimized on our perceptual scatterplot clustering dataset (PSC). In the first column, we display the human annotation with the highest agreement rate from a group of human raters for such stimulus. We report the agreement index κ_α , which measures whether a rating increases or decreases agreement within a group of raters. Higher values indicate better agreement.

Abstract

Cluster separation is a task typically tackled by widely used clustering techniques, such as *k*-means or DBSCAN. However, these algorithms are based on non-perceptual metrics, and our experiments demonstrate that their output does not reflect human cluster perception. To bridge the gap between human cluster perception and machine-computed clusters, we propose HPSCAN, a learning strategy that operates directly on scattered data. To learn perceptual cluster separation on such data, we crowdsourced the labeling of 7,320 bivariate (scatterplot) datasets to 384 human participants. We train our HPSCAN model on these human-annotated data. Instead of rendering these data as scatterplot images, we used their *x* and *y* point coordinates as input to a modified PointNet++ architecture, enabling direct inference on point clouds. In this work, we provide details on how we collected our dataset, report statistics of the resulting annotations, and investigate the perceptual agreement of cluster separation for real-world data. We also report the training and evaluation protocol for HPSCAN and introduce a novel metric, that measures the accuracy between a clustering technique and a group of human annotators. We explore predicting point-wise human agreement to detect ambiguities. Finally, we compare our approach to ten established clustering techniques and demonstrate that HPSCAN is capable of generalizing to unseen and out-of-scope data.

CCS Concepts

• **Human-centered computing** → Visualization theory, concepts and paradigms; Empirical studies in visualization; • **Computing methodologies** → Supervised learning;

1. Introduction

Clustering is often applied to bivariate projections of multidimensional data visualized as scatterplots, which aids humans in identifying cluster patterns within a large set of data points [CD19]. The insights gained from this process can guide further analysis or decision-making. The effectiveness of clustering depends on various factors, such as the choice of the clustering algorithm, its parameters, and the number of clusters to be identified, the development of cluster algorithms has been researched for decades [EKSX96; ZRL97; ABKS99]. Although these algorithms achieve meaningful separation of clusters, the resulting clusters often do not correlate with the perceptual cluster separation of the human visual system, as illustrated in Figure 1. They also require careful hand-tuning of parameters [ASA*19] or reliance on external validation measures [LLX*10]. Recent works even regard human visual perception as the gold standard for evaluating clustering [XLJ*21].

The discrepancy between clustering algorithms and the perceptual cluster separation of the human visual system can hinder humans' ability to interpret clustering results [ASA*19]. When algorithms identify clusters that do not align with the way humans would naturally group data, it becomes more challenging to understand and interpret the patterns and relationships within the data [EdMds*14] or trust the computational clustering model [SSK*15]. Aligning clustering outcomes with human perception significantly enhances the interpretability and the applicability of results, ensuring that decisions are based on insights that are both computationally valid and intuitively understandable. In some instances, intuition may outperform rational or analytical reasoning [SS04; DRP12]. This alignment is especially beneficial in applications where decision-makers heavily rely on visual data interpretation to make informed choices [Hua18; LRG19]. Moreover, existing clustering algorithms lack the ability to indicate the confidence of their output. Therefore, we propose HPSCAN, a learning-based clustering model, that through training on empirical data, integrates the way humans naturally group and detect data. It can automatically quantify grouping patterns in 2D scattered data, separate outliers from clusters, and provide human agreement scores for such estimations.

Our goal is for HPSCAN to be applied directly to point-based data, rather than to scatterplot images. This enables researchers to comfortably interchange existing clustering techniques with HPSCAN. Unlike image-based approaches, such as ScatterNet [MTW*18], our approach focuses solely on the human-aligned clustering of point-based data and does not require the generation of scatterplot images. Learning to predict point-wise cluster assignments from human-annotated point data renders standard classification losses inapplicable due to the arbitrary orderings of cluster IDs in the annotations. Therefore, we adapt meta-classification learning for segmentation tasks, a technique proposed by Hsu *et al.* [HK15; HLS*19], HPSCAN exhibits this property. Similarly, the order of point data can be arbitrary, requiring the ability of input order invariance. By adapting the PointNet++ [QYSG17] as base architecture HPSCAN exhibits such property. To evaluate our model's predictions, we need to measure its agreement with human raters. The standard Van Belle Kappa Index *et al.* [VA09], used to

measure agreement between a single rater and a group of raters, does not work in certain atypical cases that are still frequent in our datasets. This has led us to develop an alternative metric. The collected dataset not only enables us to investigate human rater agreement but also to incorporate this knowledge into the model predictions. Within this paper, we make the following contributions:

- We propose HPSCAN, a point-based neural clustering algorithm that clusters point-based data following human cluster perception.
- We present a novel loss function that accounts for point-to-cluster matching invariant to cluster IDs.
- We release a large-scale point-based clustering dataset, containing 7,320 human annotated scatterplots from real-world data[†].
- We predict human uncertainty in different usage scenarios.
- We introduce an outlier-aware, generalizable metric for measuring agreement between an isolated rater and a group of raters.
- We explore the quality of HPSCAN compared to competitors, sensitivity to various aspects, and generalization to new datasets.

2. Related Work

Many approaches for clustering bi-variate data have been developed in the last few decades. In this section, we provide an overview of existing approaches, first focusing on algorithms based on hand-crafted features and then discussing learning-based algorithms. Then, we will discuss approaches which take into account human perception. Finally, we provide an overview of existing scatterplot datasets for which human judgments were collected.

Conventional clustering. Over the years, a vast array of algorithms for clustering data has been developed. These include density-based clustering [EKSX96], hierarchical clustering [MW17], subspace clustering [EV13], fuzzy clustering [BEF84], and co-clustering [DMM03], with new approaches emerging annually. While some methods have a predetermined maximum number of clusters, others require tuning to determine appropriate values for neighborhood size, point distance threshold, or bandwidth. One of the most widely used methods, DBSCAN [EKSX96], assesses distances between the nearest points and identifies outliers. Zhang *et al.* [ZRL97] propose BIRCH, a tree-based cluster algorithm, that uses an existing agglomerative hierarchical clustering algorithm to cluster leaf nodes. Ankerst *et al.* [ABKS99] presented a density-based clustering algorithm that generates a cluster order, reflecting the cluster structure of a given set of points.

Optimization-based Clustering. One of the earliest clustering techniques is K-Means [HW79], which aims to find a partition that minimizes global within-cluster variance. This was followed by the introduction of learnable clustering algorithms, such as Neural Gas [MBS93]. Du *et al.* [Du10] and Schnellback *et al.* [SK20] provide comprehensive overviews of neural clustering approaches. Self-organized maps (k-SOM), developed by T. Kohonen [Koh91; GK11; AAS04], represents a neural clustering algorithm that operates on a grid of neurons. In this algorithm, the network learns to assign clusters to proximal data points that are close to each other.

[†] The data that support the findings of this study are openly available at <https://github.com/kopetri/HPSCAN>

There are several surveys on clustering using deep learning approaches [AGSC18; KBZ*20; RPY*22; ZXZ*22], which describe techniques for utilizing feature encoding to learn representations conducive to clustering. However, these works primarily focus on optimizing clusters in an unsupervised manner, where clusters are defined by computed quantities and are enforced through cluster loss functions, rather than by human judgments. This work focuses on the development of a deep learning method that is supervised by humans, with the goal of identifying clusters in the underlying data.

Interactive Clustering. Numerous approaches have been developed for interactive cluster analysis [INZ*07; LSP*10; MV15; BPBO15; Dem17], offering users control over the dimension reduction techniques, visual encoding, cluster parameters, and more. Xia *et al.* [XHL*22] propose an interactive cluster analysis method using contrastive dimensionality reduction. Initially, a neural network generates an embedding to reduce the dimensionality of a given high-dimensional dataset. Subsequently, users interactively select data points to establish 'must-link' and 'cannot-link' connections between clusters. The neural network is then re-trained using contrastive learning to update the embedding. Additionally, Cavallo *et al.* [CD19] introduce *Clustering Tour*, an interactive tool designed to assist users in selecting clustering parameters and evaluating the effectiveness of various clustering outcomes in relation to their analysis goals and expectations.

Image-based Machine Learning Methods for Clustering Image-based deep learning has been successfully applied to bivariate data by learning from scatterplot images. In the work of Pham *et al.* [PND20], a CNN is trained to estimate visual characterizations of scatterplots. Xia *et al.* [XLJ*21] propose a CNN-based approach to analyze the factors of class and cluster perception based on human judgement data. Fan *et al.* [FH18] proposes a CNN trained on density histogram images and human-generated brushing of underlying scatterplots to automatically brush areas in scatterplots, without selecting individual points, which is a selection-targeted clustering technique. These image-based approaches rely on the projection of bivariate data onto a 2D canvas. While in many scenarios scatterplots might be readily available or straightforward to generate, problems such as overdraw can arise when visual encoding is not selected properly. HPSCAN infers clusters from point data directly, without the need to generate images.

Point-based Machine Learning Methods for Clustering In the past years, several point-based learning architectures have been proposed that directly learn from point cloud data [QSMG17; QYSG17; HRV*18; JYC*20; DHN19]. Chenet *et al.* [CLX*19] proposes a deep hierarchical cluster network called ClusterNet, which better adapts to rotations of 3D objects. It is argued that existing data augmentation strategies or rotation equivariant networks cannot guarantee to satisfaction of all rotation-equivariant constraints at each layer. Therefore, ClusterNet introduces a rigorous rotation-invariant representation by employing hierarchical clustering to explore and exploit the geometric structure of the point cloud. In the work of Wöhler *et al.* [WZM*19], PointNet++ was demonstrated to be effective for 2D point cloud learning, estimating the correlation of data dimensions visualized by scatterplots from human-annotated data.

Table 1: Overview of existing scatterplot datasets featuring subjective human judgments. Judgments collected for our dataset consists of richly annotated scatterplots, rather than binary decisions.

Dataset	participants	responses	stimuli	modality	human judgment
ScatterNet [MTW*18]	22	5,135	50,677	real	similarity perception
ClustMe [AASB19]	34	34,000	1,000	synthetic	cluster count (binary)
ASD [QNR22]	70	1,259	5,376	real	visual encoding
VDCP [QR20]	26	1,139	7,500	synthetic	cluster count
HSP [PKF*16]	18	4,446	247	real	similarity perception
WINES [TBB*10]	18	90	18	real	class separability
SDR [SMT13]	2	1,632	816	real+synth	class separability
VCF [XLJ*21]	5	152K	50,864	real+synth	binary cluster separation
CLAMS [JQL*24]	18	1,080	60	synthetic	point-wise cluster separation
PSC	384	7,320	1,464	real	point-wise cluster separation

Perception-based clustering. Most of perception-based approaches have been proposed for visual quality measures. Quadri *et al.* [QR20] crowdsource cluster counts from human observers for synthetic scatterplots. They use distance and density-based algorithms to compute cluster merge trees. Furthermore, they utilize the merge trees with a linear regression model to predict the number of clusters a human would perceive in a scatterplot, without identifying the actual point-to-cluster assignment. Abbas *et al.* [AASB19] propose ClustMe, a visual quality measure to rank scatterplots based on the complexity of their visual patterns. It encodes scatterplots using Gaussian Mixture Models (GMM) before optimizing a model of component-merging based on human judgments. ClustML [HUB*24] enhances ClustMe by utilizing an automatic classifier trained on human perceptual judgments, replacing heuristic-based merging decisions with more accurate, data-driven models that align more closely with human visual assessment. While these approaches are GMM-based quality measures ranking scatterplots based on their grouping patterns, we propose a clustering algorithm that assigns cluster labels to individual points. A recent approach called CLAMS, introduced by Jeon *et al.* [JQL*24], measures cluster ambiguity in visual clustering. They trained a regression model using handcrafted features based on perceptual data to estimate a score representing cluster ambiguity of an input scatterplot. In contrast, HPSCAN learns to represent features of a scatterplot implicitly based on its point encoder. Furthermore, Sedlmair and Aupetit [SMT13; SA15] evaluated 15 visual quality measures for class separability of labeled data based on human judgments. These measures only apply to already labeled data points as a whole and do not allow for the alignment of individual points to clusters. ScatterNet, proposed by Ma *et al.* [MTW*18], is a learned similarity measure that captures perceptual similarities between scatterplots to reflect human judgments. Aupetit *et al.* [ASA*19] provide an evaluation of 6 state-of-the-art clustering techniques on a perception-based benchmark [AASB19]. In their evaluation, they assess a cluster counting task, where the benchmark provides human decisions for a scatterplot on whether one or more than one cluster were perceived. They can show that agglomerative clustering techniques are in substantial agreement with human raters. However, in this work, we asked human raters for point-wise cluster decisions, rather than a binary decision on the number of clusters.

Scatterplot datasets. Our provided dataset is not the first scatterplot dataset available to the research community. Existing

work has already investigated subjective human judgments in the context of scatterplots [MTW*18; AASB19; QNWR22; QR20; QR21; PKF*16; TBB*10; SMT13]. As listed in Table 1, datasets have been collected featuring diverse human judgments, such as similarity perception, class separability, or cluster counts. Sources for real scatterplots include popular datasets like MNIST, Rdatasets [Are23], or scatterplots synthetically generated based on Gaussian Mixture Models. While these datasets provide valuable human judgments for scatterplots, they lack complexity, as point-wise judgments are crucial for investigating human cluster perception. Therefore, we have collected a point-annotated cluster dataset, for which we describe the crowdsourcing process and provide statistics in the following subsections.

Prerequisite. HPSCAN is based on PointNet++, which we detail further for the reader in this paragraph. PointNet++ is a hierarchical neural network that effectively captures local geometric structures in point data. It introduces set abstraction layers composed of three fundamental operations: sampling, grouping, and feature extraction. In the sampling step, Farthest Point Sampling (FPS) is employed to select a subset of representative points from the input, preserving crucial geometrical features while reducing computational complexity. The grouping step organizes neighboring points into local regions centered around each sampled point. For feature extraction within these local neighborhoods, PointNet++ utilizes 1D convolutions to compute point-wise features. These features are then aggregated using symmetric functions such as max pooling, ensuring that the network remains invariant to the order of points, as the aggregation does not depend on the sequence of points. This hierarchical and permutation-invariant design enables PointNet++ to effectively learn both local and global features, making it a suitable point encoder for tasks such as classification and segmentation in point data.

3. HPSCAN

To learn and imitate human cluster perception based on point-based clustered data, we propose HPSCAN, a model that clusters point-based data using human annotations. In this section, we will first discuss the design choices for our approach (see Subsection 3.1). Then, we will describe the architecture of HPSCAN (see Subsection 3.2), before discussing the loss function used for training HPSCAN (see Subsection 3.3).

3.1. Design Considerations

Learning to cluster bivariate data is a classification task where a deep learning model takes as input a vector of N 2D points and produces as output another vector of real numbers that assign each point to one of K classes. We must ensure invariant to input permutation, as different point orderings yield the same scatterplot, and invariance to class assignment, as the identity of the class does not matter.

Clustering Technique. HPSCAN is developed as a clustering technique with the goal of mirroring human perceived clustering on monochrome scatterplots. Therefore, the input to our model is a set of 2D points. HPSCAN assigns each point a cluster

ID indicating cluster affiliations. Human judgments for clustering scatterplots are not identical between individual subjects for the same scatterplot, as shown in Section 5.3. Consequently, HPSCAN estimates for each point a human agreement score indicating the degree of agreement between human raters. In detail, our approach estimates N pairs (C, A) with cluster ID C and agreement score A for a given set of N points.

Point-wise clustering. While clustering by human raters is performed in the image domain through visualization of scattered data, one could argue for developing an image-based clustering technique to imitate human-perceived clustering. We assume that data are already available as pointsets before being encoded as scatterplot. Thus, we design HPSCAN to operate on a point-level basis assigning cluster IDs to each point directly.

Applicable to unstructured data. HPSCAN must be able to learn from unstructured data. Unfortunately, learning from such unstructured data poses several challenges. First, the ordering of our scattered point data, as stored on disk, might vary without actually affecting the visualization. Such effects are sometimes overlooked in the visualization community but have recently also been investigated for line graphs by other researchers [TB21]. For training HPSCAN, we require invariance over point ordering. Another challenge is that our visual stimuli are often sparse, meaning they contain a large number of empty regions, which makes more challenging to extract useful information from the data.

Uncertainty. Since crowdsourcing annotations involve human raters, our collected annotations per scatterplot may contain arbitrary cluster assignments, especially in cases where no clear patterns or shapes are present. Thus, we design HPSCAN to output an agreement score that indicates uncertainty.

Order-invariance. When predicting clusters, we want HPSCAN to be invariant to the cluster IDs, since cluster IDs can be permuted without affecting the correctness of the result. Therefore, we need to implement a model architecture and a loss function that are invariant to these orderings. We utilize meta-classification learning, which is invariant to cluster ID order.

3.2. Model Architecture

To address the challenges outlined above, we decided to adapt the PointNet++ architecture [QYSG17] to learn from our point-based scatterplot dataset. PointNet++ has been demonstrated to perform well in classification and segmentation tasks, and it has been applied to 2D point clouds sampled from MNIST [LBBH98] images. Following their example, we fix the input size to $N = 512$ points in Euclidean space and set the z -axis of all points to 0 which transforms the model into a 2D architecture as all values of computations and weight matrices along the third dimensions are removed. The feature extractor consists of 4 hierarchical layers, for both down-sampling and up-sampling stages, where we use point cloud sizes 256, 128, 64, 8 for farthest point sampling (FPS). Note that we use a randomly initialized FPS during training and fix it for inference. Additionally, we change the implementation of Point-




Figure 2: Overview of the input and output parameters of HPSCAN: A point set of N points is processed by a PointNet++, which outputs for each point three outputs: agreement score, noise and cluster affiliations. The latter two outputs are combined to produce the final point-wise clustering estimation.

Net++, that for each point, the four feature propagation stages of PointNet++ have output feature sizes 256, 256, 128, 128. Finally, the point encoder projects feature vectors of size 128 using a final 1D convolution layer, outputting an agreement score α , cluster probability P_c , and noise probability P_{noise} for each point. The output of our network has the shape $N \times (C + 2)$, where C is the maximum number of clusters, see Figure 2.

3.3. Training Loss

Commonly used loss functions, such as negative log-likelihood, would require a fixed cluster order and are therefore not applicable. This is because they would imply that the position of a cluster is related to the number of clusters present in the target. Specifically, we need a loss function that computes identical gradients for predictions $A = [c_0, c_1, c_2]$ and $B = [c_2, c_0, c_1]$ given that there are three clusters in the ground truth. Without this, the model would chase arbitrary gradients, resulting in a lack of convergence, which underscores the importance of a specialized loss function. Therefore, we draw from the field of meta-classification learning to use a contrastive loss, as proposed by Hsu *et al.* [HK15; HLK17; HLS*19]. Meta-classification learning addresses a multi-class problem by reformulating it as a binary-class problem. It optimizes a binary classifier for pairwise similarity prediction and through this process learns a multi-class classifier as a submodule. Consequently, we represent our cluster targets as a similarity matrix \mathbb{S} , which encodes point-wise cluster affiliation. It has the form $\mathbb{S} \in N \times N$, where N is the number of points in the scatterplot, and it is defined as:

$$\mathbb{S}_{i,j}(C) = \begin{cases} 1, & C_i, C_j \text{ in the same cluster (positive)} \\ 0, & \text{in different clusters (negative)} \end{cases} \quad (1)$$

where C represents the point-wise cluster IDs for N points. This matrix contains positive values for similar points that belong to the same cluster and zero values for points that are dissimilar, i.e., not in the same cluster. For clarification, we demonstrate the representation of our targets in Figure 3. It is evident that permuting the order of points results in the same matrix with equal permutation applied.

We define the loss term based on this adjacency matrix representation, which does not require the identification of clusters, thereby making it invariant to cluster identity, as follows:



Figure 3: Visualization of the encoding of ground truth and predicted clusters, along with the corresponding contrastive loss. Scattered data with color coded clustering (left), and corresponding similarity matrix \mathbb{S} , see Equation (1) (right).

$$\mathcal{L}_{\text{mcl}} = - \sum_{ij} \omega_D^{ij} (\mathbb{S}_{ij} \log \hat{\mathbb{S}}_{ij} + (1 - \mathbb{S}_{ij}) \log(1 - \hat{\mathbb{S}}_{ij})) \quad (2)$$

where ω_D is a weight matrix used to rescale the influence of negative samples in our contrastive loss, and $\hat{\mathbb{S}}$ is the estimated similarity matrix, which can be derived from the output. Furthermore, we use it to correct the imbalanced error contribution from small clusters by applying an increasing weight factor to the loss function. We define this weight matrix as follows:

$$\omega_D^{ij}(\mathbb{S}) = \begin{cases} \frac{1}{w_c^2}, & \mathbb{S}_{ij} = 1 \\ D, & \mathbb{S}_{ij} = 0 \end{cases} \quad (3)$$

where D is the momentum of negative samples and w_c is the cluster-specific contribution, defined as $w_c = N_c/N$. Here, N_c is the number of points for a specific cluster c , and N is the total number of points.

Separating clusters in scattered data is not always well-defined. While it might be obvious in some cases, in others, points may not cluster at all, or outliers might be present in the scatterplot. In the following, we treat non-clustering points and outliers the same, referring to both as NOISE. To better capture the aspect of NOISE, we introduce another loss term $\mathcal{L}_{\text{noise}}$. It is computed using a weighted binary cross entropy loss to counter the strong class imbalances between positive (CLUSTER) and negative (NOISE) cases. In our experiments, we found that scaling the weighting of negative cases by 9 is beneficial. We reserve cluster ID 0 for points annotated as NOISE. To convert a multi-cluster target, which contains multiple cluster IDs, into a binary target (0 = noise and 1 = not noise), we replace all cluster IDs greater than 0 with 1. The loss term is then defined as follows:

$$\mathcal{L}_{\text{noise}} = -\frac{1}{N} \sum_i^N 0.9 v_i \log(\hat{v}_i) + 0.1 (1 - v_i) \log(1 - \hat{v}_i) \quad (4)$$

where v_i is the binary ground truth label, and \hat{v}_i is the predicted probability of point i being classified as noise.

Finally, we aim to estimate an agreement score using HPSCAN. To compute the agreement between two clusterings, R_a and R_b , obtained from raters a and b for the same point set, we propose the following measure, referring to it as the agreement score \mathcal{A} . We define it for a single point i as follows:

$$\mathcal{A}_i(R_a, R_b) = \frac{1}{N} \sum_j^N 1 - |\mathbb{S}_{ij}(R_a) - \mathbb{S}_{ij}(R_b)| \quad (5)$$

where R_a and R_b are the annotations of two participants for the same scatterplot R . The group G of raters from our study is a sample of the population of all possible human raters we aim to model. We can estimate what would be the agreement among any pairs of raters in the general population by averaging the $K = \binom{M}{2}$ pairwise rater agreements taken among the M raters we have in G :

$$\alpha_i(G) = \frac{1}{K} \sum_k^K \mathcal{A}_i(R_{a_k}, R_{b_k}) \quad (6)$$

Finally, training HPSCAN to predict human agreement is done by adding the agreement loss term to Equation (8), which is defined by:

$$\mathcal{L}_{\text{agree}} = \frac{1}{N} \sum_i^N |\alpha_i - \hat{\alpha}_i| \quad (7)$$

where α_i is the aggregated agreement score between all human raters, see Equation (6) and $\hat{\alpha}_i$ is the prediction by HPSCAN.

Now, we can combine all loss terms using a sum, whereby our loss term $\mathcal{L}_{\text{total}}$ is composed by rescaling \mathcal{L}_{mcl} by a factor of 0.1 adjusting it to similar loss contribution. This results in the following overall loss:

$$\mathcal{L}_{\text{total}} = 0.1 \mathcal{L}_{\text{mcl}} + \mathcal{L}_{\text{noise}} + \mathcal{L}_{\text{agree}} \quad (8)$$

3.4. Training Protocol

To train HPSCAN, we use the Adam optimizer [KB14], with betas = (0.9, 0.999) and a learning rate of 1e-5, reducing the learning rate by a factor of 10 when the validation loss stagnates for 50 epochs, with a batch size of 32. During training, we apply random data augmentation to the input point clouds, including horizontal and vertical flips, and random rotations between -180 and 180 degrees. We then normalize all points clouds to be in the range of -1.0 to 1.0 , centered around the origin, identical to the procedure described in Section 4.1. Additionally, we adopt a random

crop transformation, which performs random horizontal and vertical cuts through the points. Points from the left (top) side of the cutting line are moved to the right (bottom) side, and vice versa. We keep track of the clusters that get cut, thereby increasing the number of cluster annotations accordingly. For details on how we implement this data augmentation strategy, we refer the reader to Section 6 in the supplemental material.

4. Perceptual Scatterplot Clustering Dataset

This section provides details on the collection of our perceptual scatterplot clustering dataset (PSC), which is used to train HPSCAN. To collect a large number of annotations suitable for training HPSCAN, we crowdsourced annotations for our dataset online using *Prolific*.

4.1. Stimuli Selection

The correct selection of stimuli is crucial for collecting high-quality annotations. It was essential to choose real-world stimuli, rather than more simplistic ones generated with Gaussian mixture models, to test our approach in real-world scenarios. Therefore, we downloaded data from <https://data.gov>, the United States government's open data website, which offers access to datasets covering a broad range of topics, including agriculture, climate, crime, education, finance, health, energy, and more. At the time of collecting the dataset, the site provided access to more than 240,000 datasets. We chose to collect only data available as CSV files for easier processing. In the first preprocessing step, we filtered out datasets with fewer than 512 or more than 10,000 rows to generate comparable stimuli. Additionally, we used only CSV files with more than two columns, as we require at least two data dimensions to create two-dimensional plots. For datasets with more than 512 rows, we randomly sampled a fixed number of 512 rows. Following the approach of Sedlmair *et al.* [SMT13], we applied dimensionality reduction techniques to these datasets, yielding 1,464 scatterplots. For dimensionality reduction, we used t-SNE [vdMH08] and PCA [Jol02] from the scikit-learn framework [PVG*11], using default parameters. Finally, we normalized all datasets by centering the points around (0,0) and applying minimum/maximum normalization, with positions lying within the range of $[-1, 1]$. The resulting dataset has been used to generate stimuli for our crowdsourcing process.

4.2. Crowdsourcing Process

In the past, crowdsourcing experiments have proven useful for collecting large amounts of annotated data [HSvO*22; HB10; BKG11; vOVR22; YMH*23]. To crowdsource our training data from a large group of raters, we built a web-based framework supporting mouse and keyboard interactions. Crowd workers were exposed to this framework and tasked with segmenting clusters in scatterplots generated from the stimuli selected as discussed above. For intuitive interaction, clusters had to be segmented by brushing points, with the brush color varied per cluster.

The entire crowdsourcing process is divided into three parts. First, each crowd worker received an introduction, where they

watched a 3-minute video describing the task, example stimuli with corresponding clustering, and examples of good and bad clustering. We refer the reader to the supplementary material for the tutorial video. To segment the scatterplots on a point-based level, we presented the plots vector-based, rendered on a 500×500 pixel-sized canvas using a marker size of 5 pixels. We allowed participants to brush the points using a circular brush with variable size, enabling them to make both coarse and fine-grained annotations. Initially, all points within the scatterplot are colored black; during the segmentation process, participants use the brush to colorize the points within a cluster using a selected color. Initially, the interface provides only a single color for brushing points. However, participants can add more colors (up to 20) by clicking a + button. If participants were unable to delineate any clusters, we enabled them to explicitly state this using a checkbox labeled *I can't see any cluster*. Participants were further instructed that points without cluster affiliation should remain black, signaling that they are considered outliers or noise. Once confident with a segmentation, participants pressed a button below the plot to continue to the next stimulus.

In the second part of the crowdsourcing process, each participant practiced the task during a short training phase and received performance feedback. While this part used stimuli not included in the main study, these stimuli formed obvious clusters, chosen to align with the expected agreement between clustering algorithm results and capable human observers. Accordingly, we used this clustering as ground truth for the participant feedback. We could also verify participant performance through these stimuli, allowing participants to continue to the study's final part only if they completed all training stimuli successfully.

In the third part of the crowdsourcing process, participants annotated clusters in 20 stimuli each during the main study. To detect bots or *click-through* behavior, we added three additional sanity checks. Such a stimulus displays multiple spatially separated Gaussian blobs forming visually distinct clusters. We present such a stimulus used as a sanity check, for which we predefined ground truth cluster separation. Participant segmentations that diverged more than 30% from the target failed this sanity check, and we discarded the data from participants who failed more than one sanity check. In Section 1 of the supplemental material, we show the user interface of the web application used, along with an example of a sanity check.

Using this procedure, we collected 7,320 clusterings for 1,464 scatterplots from 384 participants, with each stimulus annotated by at least 5 individuals. On average, a participant took 15.1 minutes to complete our study, with 257 male, 125 female, and 2 participants who preferred not to provide their sex. The average age of the participants was 32.7 years. We had to reject a single participant for failing the sanity checks and an additional 17 participants for discontinuing the study. In the following, we report statistics regarding the collected dataset. Compared to existing datasets, see Table 1, our study obtained annotations from 5 times more subjects than previous dataset studies.

4.3. Annotation Analysis

To investigate the agreement between human judgments, we computed an agreement score α , which is defined in Section 5.1. The




Figure 4: Our dataset, PSC, consists of 1,464 stimuli. We show the distribution of stimuli that were annotated with a certain number of clusters. Rater agreement is visualized in red for each number of clusters. The white numbers correspond to the number of stimuli for each cluster count, while the black numbers indicate the corresponding amounts. The red line represents the average agreement for the entire dataset.

average agreement score for PSC is 81.9%. In Figure 4, we display the average agreement scores stratified by the number of clusters. The number of clusters is determined by checking if more than half of the participants agree on the same number of clusters. If this agreement cannot be found, we compute the average number of clusters from all five annotations. The results show that the participants more often annotated fewer clusters, as indicated by the decreasing number of annotations with an increasing number of clusters.

To illustrate our agreement score, Figure 5 shows two stimuli with corresponding annotations and different agreement scores, whereby both stimuli having been annotated by five participants. In the top row, the five annotations demonstrate strong agreement in clustering, resulting in an average agreement score of 99.79%. In contrast, the bottom row shows annotations with an agreement score below 50%, where participants not only disagree on which points belong together, but one participant (most right) even indicated the absence of any clusters. Additionally, Figure 4 reveals that the occurrence of data samples decreases with an increasing number of clusters. For cluster numbers greater than six, the amount of data samples falls below 5%, which is too small a proportion to make general assumptions. Therefore, we exclude such stimuli from our training set and include only those highlighted by orange bars in Figure 4.

5. Evaluation

For the evaluation of HPSCAN, we report performance for the test split of PSC with the following dataset sizes: train (1,171 stimuli), validation (87 stimuli), and test (206 stimuli).

Similar to existing clustering algorithms like DBSCAN [EKSX96], Birch [ZRL97], and Optics [ABKS99], HPSCAN can predict on bivariate point sets directly, rather than rendered scatterplot images. As described in Section 4, the scatterplots used for crowdsourcing human annotations originate

from real bivariate data. For the evaluation of our experiments, we compute the three metrics proposed in Section 5.1 and report results averaged per group of cluster numbers from Figure 4. To compensate for the imbalanced numbers of data samples per group, we compute the weighted average. Note that κ_n is computed only for data samples that contain points labeled as noise. Additionally, we provide a comprehensive evaluation of this experiment in Section 7 and Section 8 of the supplemental material, where we elaborate on further metrics like the Silhouette Index [Rou87] and investigate the effect of dimensionality reduction on model performance.

5.1. Outlier-Aware Rater Agreement

Many measures exist that can quantify performance regarding different aspects. Therefore, selecting a metric that can capture the intended solution to our problem is crucial. To evaluate HPSCAN, we need to measure the agreement between an isolated rater and a group of raters, which is computed by the Vanbelle Kappa Index [VA09], defined as:

$$\kappa_v = \frac{\Pi_T - \Pi_E}{\Pi_M - \Pi_E} \quad (9)$$

with Π_T being the theoretical agreement, Π_M the maximum attainable agreement and Π_E the agreement expected by chance. Thus, the Vanbelle Kappa Index computes values in the range of -1 (minimal agreement) to 1 (perfect agreement). A value of 0 corresponds to the agreement expected by chance. Note that Vanbelle and Albert indicate in their work that if there is no variability in the rating from the isolated rater or from the group of raters, their index reduces to $\kappa_v = 1$ for perfect agreement, and $\kappa_v = 0$ otherwise when $\Pi_M = \Pi_E$. In cases where raters agree on NO CLUSTER and SINGLE CLUSTER, all ratings are identical, leading to $\kappa_v = 0$, if the group of raters does not perfectly agree, resulting in unreasonable scores computed by κ_v . In Section 13 of the supplementary material, we provide toy examples demonstrating such shortcomings of the Vanbelle Kappa Index.

Therefore, we develop a new measure that additionally accounts for outliers within the group of raters, and that can even measure if an isolated rating can improve agreement within the group. To

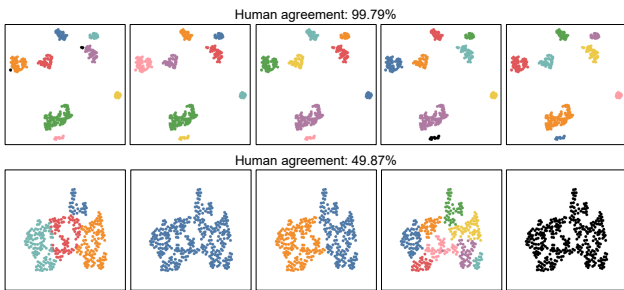


Figure 5: Comparison of two scatterplots used in our study. The top row shows a strong average agreement score of 99.79%, while the bottom row has a low average agreement score below 50%.

do so, we define α by averaging Equation (6) over all points and compute an averaged m-fold score κ_α for a given prediction R and the group of raters G . We pick an annotation within the group and replace it with R , yielding a modified group G^R . Then we compute $\alpha(G^R)$ and repeat this M times, where M is equal to the number of annotations in the group. Thus, we compute an agreement index κ_α , which indicates whether a given prediction improves agreement for G when $\kappa_\alpha > 0$ or decreases it when $\kappa_\alpha < 0$. We define κ_α as follows:

$$\kappa_\alpha(R, G) = \frac{1}{M} \sum \alpha(G_m^R) - \alpha(G) \quad (10)$$

where G is the group of annotations, and G_m^R is the modified version of G with the m^{th} annotation replaced by R .

κ_v and κ_α implicitly capture outliers; however, to explicitly measure noise detection, we further adopt a noise index κ_n , which computes the average of an m-fold binary Jaccard Index \mathfrak{J} over the group of raters and the prediction. This measure is defined as:

$$\kappa_n(\hat{R}, \hat{G}) = \frac{1}{M} \sum \mathfrak{J}(\hat{R}, \hat{R}_m) \quad (11)$$

where \hat{R} is the binary prediction for outliers, encoding points associated with noise as 0 and points affiliating with a cluster as 1 . \hat{R}_m is the binary annotation of the m^{th} rater of group G .

Having defined these metrics, we will use them in the following sections to evaluate the results of HPSCAN.

5.2. Contrastive Loss Weighting Analysis

In Section 3.3, we proposed our contrastive loss function \mathcal{L}_{mcl} , along with weight matrix ω_D , which is used to rescale the momentum of cases when two points are in dissimilar clusters. Depending on the scaling factor D , the weight matrix emphasizes gradients for dissimilar points, separating them into different clusters. In this experiment, we demonstrate how this can be used to adjust the behavior of our model using different values for $D \in [0.1, 1.0, 10.0, 50.0, 100.0]$. To do so, we train five models using the protocol from Section 3.4, and we report performance in Table 2a. Each model was trained for 37K steps, and looking at the results, the model trained with a weighting factor $D = 50.0$ shows slightly better results for κ_α and κ_v , while noise prediction accuracy is behind the others. From the results presented in Section 2 of the supplementary material, we observe that the model separates clusters more distinctly for higher values of D and merges clusters for smaller values. With medium values of $D = 10$ and $D = 50$, the clustering results are more variable, increasing clustering performance but also making it harder for the model to differentiate between noise and cluster decisions. However, in subsequent experiments, we show that this effect diminishes with a larger number of training steps, and we therefore use $D = 50.0$ in the remaining experiments unless stated otherwise.

Table 2: Results of the contrastive loss weighting analysis (a) in Section 5.2, the human agreement analysis (b) in Section 5.3, and the fine-tuning analysis (c) in Section 5.4. We report performance results computing three metrics: κ_α , κ_v , κ_n based on our test set.

D	κ_α	κ_v	κ_n	T_{agree}	#samples	κ_α	κ_v	κ_n	D	T_{agree}	κ_α	κ_v	κ_n
0.1	-5.64	0.53	49.97%	UNFILTERED	1171	-1.66	0.62	45.91%	0.01	70%	-0.85%	0.66	57.07%
1.0	-6.33	0.43	48.09%	50%	1148	-1.52	0.62	47.51%	0.1	70%	-0.52%	0.69	60.91%
10.0	-2.01	0.61	46.26%	60%	1049	-1.34	0.63	50.24%	1.0	70%	-0.68%	0.67	52.36%
50.0	-1.66	0.62	45.91%	70%	883	-0.55	0.68	58.48%					
100.0	-1.78	0.61	50.76%	80%	672	-1.27	0.57	54.08%					
				90%	448	-1.66	0.58	53.92%					

(a) Analysis of different weights D for negative samples in our contrastive loss

(b) Human agreement analysis by selecting different subsets of our data set

(c) Fine-tuning analysis of our best model by decreasing momentum for negative samples

5.3. Human Agreement Maximization

In Section 4.3, we discussed the agreement rate between raters, and Figure 4 shows that it is similar for different numbers of clusters. In this experiment, we investigate the impact of increasing human agreement during training on our model’s performance. To do so, we use a threshold T_{agree} to filter our training dataset. For each threshold, we discard training samples where the agreement score $\alpha(G) < T_{\text{agree}}$. In this way, we construct five subsets using threshold values: 50%, 60%, 70%, 80%, 90%. In Section 12 of the supplemental material, we provide details on the selection of these threshold values. Then, for each subset, we train HPSCAN using the protocol from Section 3.4 and compare their performance results using PSC. In this experiment, we use a fixed negative momentum $D = 50.0$.

The obtained results are shown in Table 2b. The first row shows results from the model trained on all available training data, UNFILTERED, and is identical to the best model from the first experiment (fourth row of Table 2a). The results suggest, that while increasing the level of agreement in the annotations, the performance increases until reaching a maximum, as seen in row number four. From this point on, filtering for higher agreement rates reduces the number of training samples and decreases model performance. We conclude from this experiment, that a certain degree of variance in the annotations, helps to improve model performance, and that our model achieves robustness through such variability. As a result, we use the best model, trained with a threshold of 70%, for our remaining experiments, as shown in Table 2b.

5.4. Clustering Techniques Comparison

In this experiment, we investigate the gap between existing clustering techniques and HPSCAN. Therefore, ten state-of-the-art clustering algorithms are compared using the implementation from scikit-learn [PVG*11]. Since all approaches require parameterization, we first conduct a parameter search using our training dataset to find the optimized parameters for each technique. Second, we fine-tune HPSCAN using the insights from Section 5.2 and Section 5.3, investigating whether smaller values of D improves single cluster predictions. We train three different models using momentum values $D = [0.01, 0.1, 1]$. In Table 2c, we report the performance results for all three models. Since the model has already

learned to separate clusters, fine-tuning it to reduce its cluster separation ambitions must be done carefully. Therefore, we fine-tune each mode using a reduced learning rate of 10^{-7} for 8K training steps. The results show that fine-tuning our model using a weighting factor $D = 0.1$ helps to improve single cluster predictions and even the overall performance of our model. In Section 8 of the supplemental material, we evaluate the performance of such a model for specific numbers of clusters, separately. We demonstrate that the largest error contribution originates from predictions for single clusters. As a baseline for our evaluation, we adopt an image-based model utilizing a pre-trained image segmentation U-Net [RFB15], which we fine-tune it on our collected scatter plot stimuli. Details on the training protocol can be found in Section 4 of the supplementary material.

We then evaluate each approach using PSC and compare it against HPSCAN using the best model from the fine-tuning experiment. We report performance results in the upper part of Table 3. Note that some techniques do not compute outliers, for which we do not calculate κ_n . The baseline model performs better for the κ_α metric compared to the standard clustering techniques that do not use prior knowledge about the number of clusters. This indicates that predictions are well-aligned with human ratings. The results for κ_v and κ_n indicate average alignment compared to existing techniques. Looking at the results for HPSCAN, we can see that it outperforms all compared state-of-the-art approaches on all metrics. The κ_α measure shows that predictions from HPSCAN almost perfectly agree with the group of raters, indicated by a value close to zero. The κ_v also suggests superior performance of HPSCAN over the other clustering techniques, which we interpret as indicating that HPSCAN predictions agree with human annotators. Examining the closest competitors: Ward, K-means, and Gaussian Mixture Models, which have similar κ_v scores compared to HPSCAN, we see that traditional clustering algorithms can be used to reflect human judgments. This is also supported by κ_α values close to zeros; however, these methods require a prior of known number of clusters, which is inferred automatically by our method. Finally, examining the accuracy for noise or outlier prediction, HPSCAN has a slight advantage over DBSCAN, while both approaches outperform Optics. Additionally, we measured the running time for inference of cluster predictions utilizing the test split of PSC for 206 inference steps resulting in 5.16 sec. That translates to 0.025 sec per step utilizing an Nvidia RTX 3060 and batch size of 64. As

Table 3: Comparison between HPSCAN, the image baseline and ten state-of-the-art clustering techniques using two test datasets: PSC (Section 5.4) and SDR [SMT13] (Section 5.5). We do not compute κ_n for cluster techniques which do not compute outliers. Results highlighted with a * are from clustering techniques that require priors about the number of clusters as per the ground truth annotation.

	HPSCAN	Baseline	DBSCAN	OPTICS	Ward	Mean Shift	Affinity Propagation	Spectral Clustering	Agglomerative Clustering	BIRCH	K-means	Gaussian Mixture Model
			[EKXS96]	[ABKS99]	[ML11]	[CM02]	[FD07]	[NJW01]	[MW17]	[ZRL97]	[HW79]	[BCRR97]
Evaluation on our test split												
$\kappa_\alpha \uparrow$	-0.52	-2.40	-4.13	-5.56	-1.04*	-6.90	-10.34	-2.65*	-7.02	-7.68	-1.49*	-0.90*
$\kappa_V \uparrow$	0.69	0.51	0.62	0.56	0.66*	0.53	0.39	0.53*	0.38	0.38	0.63*	0.67*
$\kappa_n \uparrow$	59.67%	36.13%	52.12%	21.29%	-	-	-	-	-	-	-	-
Evaluation on SDR [SMT13] data set.												
$\kappa_\alpha \uparrow$	-1.34	-5.35	-3.19	-3.15	-7.21*	-5.28	-13.04	-7.21*	-5.85	-6.47	-7.21*	-7.21*
$\kappa_V \uparrow$	0.64	0.43	0.50	0.60	0.62*	0.60	0.31	0.62*	0.43	0.46	0.62*	0.62*
$\kappa_n \uparrow$	46.28%	33.86%	28.36%	36.11%	-	-	-	-	-	-	-	-

*The ground truth number of cluster was given to compute these scores.

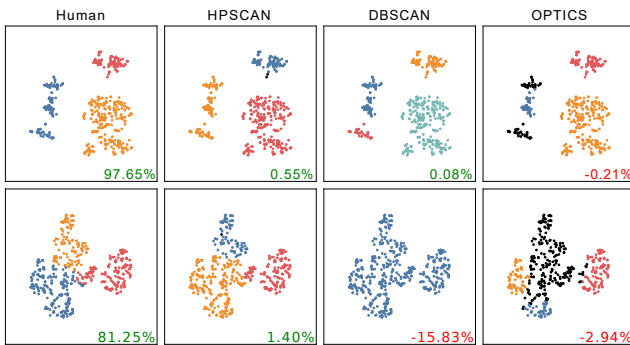


Figure 6: Qualitative evaluation results: Two point sets (rows) sampled from the SDR [SMT13] dataset are shown. Each column presents outcomes for different clustering techniques. We compare HPSCAN to the two best performing standard techniques, along with human ratings that display the highest agreement among raters, shown in the first column. For each technique, the corresponding κ_α index is shown in green, when the clustering improves group agreement, and in red otherwise.

comparison DBSCAN has a running time of 3.48 sec or 0.017 sec per step.

5.5. Generalization Analysis

In this experiment, we aim to investigate the ability of HPSCAN to generalize to unseen data. To this end, we utilize the dataset provided by the work of Sedlmair [SMT13], which consists of 2D scatterplots derived using dimensionality reduction techniques (SDR) and includes both real and synthetic data. We apply the same crowdsourcing annotation procedure as described in Section 4 to collect clustering annotations from 68 participants for 272 stimuli. We then evaluate the best model from Section 5.4 using this dataset and report results for inferring clustering in Table 3 and for regressing agreement predictions in Table 4. We demonstrate that

HPSCAN outperforms existing clustering techniques on all three measures, κ_α , κ_V , κ_n indicating better alignment with human ratings and showcasing the ability of HPSCAN to generalize to unseen data. Qualitative results for HPSCAN generated during inference are shown in the upper row of Figure 6. These estimated clusterings for the SDR [SMT13] dataset underline the well-aligned predictions of HPSCAN with human judgments, and we refer the reader to Section 5 of the supplemental material for further evaluation results.

5.6. Human Agreement Estimation

After investigating the effect of maximizing human agreement during training to learn a robust model in Section 5.3, one question remains: what if human raters do not agree on a clustering? On the left side of Figure 7, we display clusterings proposed by five human annotators for the same stimulus, as well as the corresponding agreement rate amongst the group of raters, shown in the sixth column. Consequently, we extend HPSCAN to also output an agreement score α , as proposed in Equation (6), for each point. The estimation of our network has the shape $N \times 1$ and is computed by applying a Sigmoid activation to the network output. We use the same fine-tuning procedure as described in Section 5.4, with identical hyperparameters, and train HPSCAN on PSC. We then evaluate it on the test split of PSC and SDR [SMT13], reporting regression results measured in Mean Squared Error (MSE) and Mean Absolute Error (MAE) in Table 4. Looking at the results, there appears to be no large gap between the two datasets considering the regression values, and together with qualitative evaluation results, there is an indication that HPSCAN is capable to reflect human agreement with an error of 13 points in MAE for PSC and 15 points in MAE for SDR [SMT13] dataset. In the last column of Figure 7, we display prediction results of HPSCAN for the SDR [SMT13] dataset.

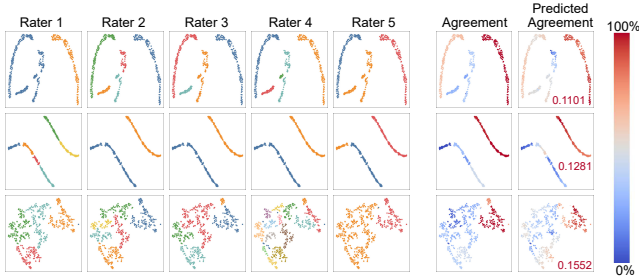


Figure 7: The first five columns display human annotations collected during our online crowdsourcing study. We compare annotations for three stimuli. In the sixth column, the computed agreement score per point is shown. Finally, in the last column, the prediction of HPSCAN is shown, along with the averaged absolute error over all points.

Test dataset	MSE	MAE
HPSCAN	0.0305	0.1393
SDR [SMT13]	0.0317	0.1564

Table 4: We evaluate our model, that is trained on our dataset, for the human agreement estimation task. We compute two regression metrics to measure its performance on two datasets: HPSCAN and the dataset SDR [SMT13], that consists of real and synthetic scatterplots.

5.7. Cluster Counting Experiment

In this experiment, we apply HPSCAN to the ClustMe dataset [AASB19], which consists of 1,000 stimuli with human judgments available. Each stimulus was annotated by 34 human raters to determine whether one or more than one cluster is perceived by the participant. We collect predictions from HPSCAN for each stimulus and convert the clustering result into the estimated number of clusters. We compute the performance results of HPSCAN using the Vanbelle Kappa Index, accuracy, and F1 score, and report the results in Table 5. To compute accuracy and F1 score, the majority cluster count from the ground truth was used as a binary label. The results indicate poor performance results compared to the approaches presented in ClustMe. To contextualize the results of HPSCAN, we show the performance of three dummy models: the random cluster model, which estimates randomly between single and multiple clusters; the single cluster model, which always estimates a single cluster; and the multi-cluster model, which always estimates more than one cluster. Indicated by the accuracy of 80% for the single cluster model and 20% for the multi-cluster model, ClustMe has a class imbalance favoring single clusters. Additionally, ClustMe datasets are generated using arbitrarily configured Gaussian Mixture Models (GMM). Consequently, HPSCAN achieves a Vanbelle Kappa Index of 0.5, an accuracy of 84%, and an F1 score of 55%, suggesting a tendency to estimate mainly single clusters for ClustMe. This experiment reveals a domain gap between dataset distributions originating from projections through dimensionality reduction and GMM-based scattered data.

model	κ_v	Acc	F1
HPSCAN	0.50	84%	55%
Θ_{random}	0.00 ± 0.03	$50 \pm 2\%$	$28 \pm 2\%$
Θ_{single}	0.00	80%	0%
Θ_{multi}	0.00	20%	33%

Table 5: Evaluation results on the ClustMe [AASB19] dataset, which consists of 1,000 stimuli annotated by 34 human raters. Each annotation is a binary decision whether a single cluster or multiple cluster were perceived by the human.

6. Discussion and Future Work

We demonstrate HPSCAN, a clustering model based on human-perceived clustering annotations. We collected a large-scale dataset using Prolific to crowdsource human point-wise clustering annotations. Investigation of the collected data shows agreement above 80% between human subjects. This dataset enables us to train HPSCAN, a learned clustering model that mimics point clustering as performed by the human visual system. In multiple experiments, our protocol for training a point-based model is demonstrated, and we show how we fine-tune our model to adjust it to human annotations. To evaluate HPSCAN, we further proposed a novel metric, that measures agreement improvement while also being sensitive to annotation consistency. Ultimately, we evaluate our model using SDR [SMT13] and Data.gov, which are datasets featuring unseen and out-of-scope data, and compare it against ten state-of-the-art clustering techniques, with HPSCAN outperforming all compared algorithms, demonstrating its ability to generalize to new data. In our experiments, we also present how to deal with multiple correct clusterings, e.g., when multiple human judgments disagree on a particular stimulus. HPSCAN can be used to estimate human agreement, enabling the detection of such ambiguities.

While HPSCAN is the first of its kind and performs well on human-like clustering, it has some limitations. PointNet++ is order invariant with respect to the input point cloud, and our contrastive loss term is invariant to the order of cluster IDs. However, constructing the similarity matrix has a complexity of $O(n^2)$. In our experiments, we used a fixed size for point clouds, $N = 512$, to keep computational costs low. As a result, we had to discard or randomly sample from real-world datasets, to match such requirements. In the field of contrastive learning, approaches like SimCLR [CKNH20] or SwAV [CMM*20] exist, which do not have this limitation and could be used to cluster points based on their latent codes. Note that during training time, HPSCAN is limited by a fixed size of N , but during inference, PointNet++ and thus our approach as well can process arbitrary numbers of points, as demonstrated in the evaluation of SDR [SMT13] dataset in Section 5.5. During training, we set C to 20, which is also the maximum number of clusters that HPSCAN can detect. As a result, our model is currently unable to handle datasets with more than 20 clusters. However, based on our observations, it is unlikely that a human judge would perceive more than this number of distinct clusters in a typical scatterplot, so we consider this limitation reasonable for the intended applications.

HPSCAN is trained on human-annotated data collected during a crowdsourced study. Therefore, we rendered datasets using a fixed

visual encoding (marker size, opacity, color, etc.), which limits the conclusions about the agreement and perceived numbers of clusters to such visual encoding. In cases where visual encoding affects cluster perception, this would not be captured by HPSCAN. However, as shown by Quadri et al. [QNWR22] visual encodings interfere with perceived clustering, and clusterings proposed by HPSCAN might look inaccurate if, for instance, larger markers are used in the scatterplot. Nevertheless, if desired, our crowdsourced study could have also included variable visual encodings, and provided it as point features to our network during training. Thus, this is not strictly a limitation of HPSCAN, as more stimuli data could incorporate visual encodings during cluster prediction.

HPSCAN is developed as an alternative to existing clustering techniques, specifically aligned to human perception. However, this work fixes the visual encoding during dataset construction in order to keep the design space manageable. Therefore, exploring the effect of visual encoding is subject to future research, while this work provides a promising baseline that is able to generalize beyond unseen data. Further, our analysis of human agreement on clustering shows a large variety of agreement degree depending on the structure of the underlying data. It is questionable how meaningful a single clustering prediction is for cases when a group of human raters provides ambiguous annotations. A generative model could be trained to produce many different such human-aligned clusterings, reflecting the same human ambiguous distribution of annotations. In this work, we decided to predict an agreement score for each point, providing neural feedback about the learned human discord. For a user of HPSCAN, this provides some understanding of the reliability of the corresponding estimated clustering.

While we see HPSCAN as a milestone for perception-based clustering, we see several endeavors for future research. First, we could investigate the influence of visual stimuli on cluster perception. Furthermore, we could leverage the presented technology to learn other scatterplot tasks, such as noise detection. Third, there is room for future incorporation of GMM-based data during training to achieve more robustness on synthetic data. Finally, we would like to investigate how HPSCAN can be used to optimize scatterplot visualization parameters along the lines of Wang et al. [WFC*17; WCG*18].

7. Data Availability Statement

The data that support the findings of this study are openly available in github.com at <https://github.com/kopetri/HPSCAN>.

References

- [AAS04] AMARASIRI, RASIKA, ALAHAKOON, DAMMINDA, and SMITH, KATE A. “HDGSOM: A Modified Growing Self-Organizing Map for High Dimensional Data Clustering”. *Fourth International Conference on Hybrid Intelligent Systems (HIS’04)*. IEEE. 2004, 216–221 2.
- [AASB19] ABBAS, MOSTAFA M., AUPETIT, MICHAËL, SEDLMAIR, MICHAEL, and BENSMAIL, HALIMA. “ClustMe: A Visual Quality Measure for Ranking Monochrome Scatterplots Based on Cluster Patterns”. *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library, 2019, 225–236 3, 4, 11.
- [ABKS99] ANKERST, MIHAEL, BREUNIG, MARKUS M., KRIEGEL, HANS-PETER, and SANDER, JÖRG. “OPTICS: Ordering Points to Identify the Clustering Structure”. *ACM SIGMOD Record* 28.2 (1999), 49–60 2, 7, 10.
- [AGSC18] ALJALBOUT, ELIE, GOLKOV, VLADIMIR, SIDDIQUI, YAWAR, and CREMERS, DANIEL. “Clustering with Deep Learning: Taxonomy and New Methods”. *arXiv preprint arXiv:1801.07648* (2018). URL: <https://arxiv.org/abs/1801.07648> 3.
- [Are23] AREL-BUNDOCK, VINCENT. *Rdatasets: A Collection of Datasets Originally Distributed in Various R Packages*. R package version 1.0.0. 2023. URL: <https://vincentarelbundock.github.io/Rdatasets4>.
- [ASA*19] AUPETIT, MICHAËL, SEDLMAIR, MICHAEL, ABBAS, MOSTAFA M., et al. “Toward Perception-Based Evaluation of Clustering Techniques for Visual Analytics”. *2019 IEEE Visualization Conference (VIS)*. IEEE. 2019, 141–145 2, 3.
- [BCRR97] BENSMAIL, HALIMA, CELEUX, GILLES, RAFTERY, ADRIAN E., and ROBERT, CHRISTIAN P. “Inference in Model-Based Cluster Analysis”. *Statistics and Computing* 7 (1997), 1–10 10.
- [BEF84] BEZDEK, JAMES C., EHRLICH, ROBERT, and FULL, WILLIAM. “FCM: The Fuzzy C-Means Clustering Algorithm”. *Computers & Geosciences* 10.2–3 (1984), 191–203 2.
- [BKG11] BUHRMESTER, MICHAEL, KWANG, TRACY, and GOSLING, SAMUEL D. “Amazon’s Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?”. *Perspectives on Psychological Science* 6.1 (2011), 3–5 6.
- [BPBO15] BRUNEAU, PASCAL, PINHEIRO, PATRICIA, BROEKSEMA, BERT, and OTJACQUES, BENOÎT. “Cluster Sculptor: An Interactive Visual Clustering System”. *Neurocomputing* 150 (2015). Special Issue on Information Processing and Machine Learning for Applications of Engineering; Solving Complex Machine Learning Problems with Ensemble Methods; Visual Analytics using Multidimensional Projections, 627–644. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2014.09.062. URL: <https://www.sciencedirect.com/science/article/pii/S0925231214012892> 3.
- [CD19] CAVALLO, MARCO and DEMIRALP, ÇAĞATAY. “Clustrophile 2: Guided Visual Clustering Analysis”. *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2019), 267–276. DOI: 10.1109/TVCG.2018.2864477 2, 3.
- [CKNH20] CHEN, TING, KORNBLITH, SIMON, NOROUZI, MOHAMMAD, and HINTON, GEOFFREY. “A Simple Framework for Contrastive Learning of Visual Representations”. *International Conference on Machine Learning*. PMLR. 2020, 1597–1607 11.
- [CLX*19] CHEN, CHAO, LI, GUANBIN, XU, RUIJIA, et al. “ClusterNet: Deep Hierarchical Cluster Network with Rigorously Rotation-Invariant Representation for Point Cloud Analysis”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019 3.
- [CM02] COMANICIU, DORIN and MEER, PETER. “Mean Shift: A Robust Approach Toward Feature Space Analysis”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24.5 (2002), 603–619 10.
- [CMM*20] CARON, MATHILDE, MISRA, ISHAN, MAIRAL, JULIEN, et al. “Unsupervised Learning of Visual Features by Contrasting Cluster Assignments”. *Advances in Neural Information Processing Systems* 33 (2020), 9912–9924 11.
- [Dem17] DEMIRALP, ÇAĞATAY. “Clustrophile: A Tool for Visual Clustering Analysis”. *arXiv preprint arXiv:1710.02173* (2017). URL: <https://arxiv.org/abs/1710.02173> 3.
- [DHN19] DING, ZHIPENG, HAN, XU, and NIETHAMMER, MARC. “VoteNet: A Deep Learning Label Fusion Method for Multi-Atlas Segmentation”. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2019*. Springer. 2019, 202–210 3.

- [DMM03] DHILLON, Inderjit S., MALLELA, SUBRAMANYAM, and MODHA, DHARMENDRA S. "Information-Theoretic Co-Clustering". *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2003, 89–98 2.
- [DRP12] DANE, ERIK, ROCKMANN, KEVIN W., and PRATT, MICHAEL G. "When Should I Trust My Gut? Linking Domain Expertise to Intuitive Decision-Making Effectiveness". *Organizational Behavior and Human Decision Processes* 119.2 (2012), 187–194 2.
- [Du10] DU, K.-L. "Clustering: A Neural Network Approach". *Neural Networks* 23.1 (2010), 89–107 2.
- [EdMdS*14] ETEMADPOUR, RONAK, da MOTTA, ROBSON CARLOS, de SOUZA PAIVA, JOSE GUSTAVO, et al. "Role of Human Perception in Cluster-Based Visual Analysis of Multidimensional Data Projections". *2014 International Conference on Information Visualization Theory and Applications (IVAPP)*. IEEE. 2014, 276–283 2.
- [EK SX96] ESTER, MARTIN, KRIEGEL, HANS-PETER, SANDER, JÖRG, and XU, XIAOWEI. "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise". *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD '96)*. 1996, 226–231 2, 7, 10.
- [EV13] ELHAMIFAR, EHSAN and VIDAL, RENÉ. "Sparse Subspace Clustering: Algorithm, Theory, and Applications". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.11 (2013), 2765–2781 2.
- [FD07] FREY, BRENDAN J. and DUECK, DELBERT. "Clustering by Passing Messages Between Data Points". *Science* 315.5814 (2007), 972–976 10.
- [FH18] FAN, CHAORAN and HAUSER, HELWIG. "Fast and Accurate CNN-Based Brushing in Scatterplots". *Computer Graphics Forum*. Vol. 37. 3. Wiley Online Library. 2018, 111–120 3.
- [GK11] GHASEMINEZHAD, M. H. and KARAMI, ALI. "A Novel Self-Organizing Map (SOM) Neural Network for Discrete Groups of Data Clustering". *Applied Soft Computing* 11.4 (2011), 3771–3778 2.
- [HB10] HEER, JEFFREY and BOSTOCK, MICHAEL. "Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2010, 203–212 6.
- [HK15] HSU, YEN-CHANG and KIRA, ZSOLT. "Neural Network-Based Clustering Using Pairwise Constraints". *arXiv preprint arXiv:1511.06321* (2015) 2, 5.
- [HLK17] HSU, YEN-CHANG, LV, ZHAOYANG, and KIRA, ZSOLT. "Learning to Cluster in Order to Transfer Across Domains and Tasks". *arXiv preprint arXiv:1711.10125* (2017) 5.
- [HLS*19] HSU, YEN-CHANG, LV, ZHAOYANG, SCHLOSSER, JOEL, et al. "Multi-Class Classification Without Multi-Class Labels". *arXiv preprint arXiv:1901.00544* (2019) 2, 5.
- [HRV*18] HERMOSILLA, PEDRO, RITSCHER, TOBIAS, VÁZQUEZ, PERE-PAU, et al. "Monte Carlo Convolution for Learning on Non-Uniformly Sampled Point Clouds". *ACM Transactions on Graphics (TOG)* 37.6 (2018), 1–12 3.
- [HSvO*22] HARTWIG, SEBASTIAN, SCHELLING, MICHAEL, van ONZENODT, CHRISTIAN, et al. "Learning Human Viewpoint Preferences from Sparsely Annotated Models". *Computer Graphics Forum*. Vol. 41. 6. Wiley Online Library. 2022, 453–466 6.
- [Hua18] HUANG, LAURA. "The Role of Investor Gut Feel in Managing Complexity and Extreme Risk". *Academy of Management Journal* 61.5 (2018), 1821–1847 2.
- [HUB*24] HAMZA, MOSTAFA M., ULLAH, EHSAN, BAGGAG, ABDELKADER, et al. "ClustML: A Measure of Cluster Pattern Complexity in Scatterplots Learnt from Human-Labeled Groupings". *Information Visualization* 23.2 (2024), 105–122. DOI: 10.1177/14738716231220536. URL: <https://doi.org/10.1177/14738716231220536>.
- [HW79] HARTIGAN, JOHN A. and WONG, MANCHEK A. "Algorithm AS 136: A k-Means Clustering Algorithm". *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28.1 (1979), 100–108 2, 10.
- [INZ*07] IMRE, DEME, NAM, EUN HA, ZELENYUK, ALLA, et al. "ClusterSculptor: A Visual Analytics Tool for High-Dimensional Data". *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (VAST)*. Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2007, 75–82. DOI: 10.1109/VAST.2007.4388999. URL: <https://doi.ieeecomputersociety.org/10.1109/VAST.2007.4388999>.
- [Jol02] JOLLIFFE, IAN T. *Principal Component Analysis for Special Types of Data*. Springer, 2002 6.
- [JQL*24] JEON, HYEON, QUADRI, GHULAM JILANI, LEE, HYUNWOOK, et al. "CLAMS: A Cluster Ambiguity Measure for Estimating Perceptual Variability in Visual Clustering". *IEEE Transactions on Visualization and Computer Graphics* 30.1 (2024), 770–780. DOI: 10.1109/TVCG.2023.3327201 3.
- [JYC*20] JIANG, HAIYONG, YAN, FEILONG, CAI, JIANFEI, et al. "End-to-End 3D Point Cloud Instance Segmentation Without Detection". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 12796–12805 3.
- [KB14] KINGMA, DIEDERIK P. and BA, JIMMY. "Adam: A Method for Stochastic Optimization". *arXiv preprint arXiv:1412.6980* (2014) 6.
- [KBZ*20] KARIM, MD REZAUL, BEYAN, OYA, ZAPPA, ACHILLE, et al. "Deep learning-based clustering approaches for bioinformatics". *Briefings in Bioinformatics* 22.1 (Feb. 2020), 393–415. ISSN: 1477-4054. DOI: 10.1093/bib/bbz170. eprint: <https://academic.oup.com/bib/article-pdf/22/1/393/35934885/bbz170.pdf>. URL: <https://doi.org/10.1093/bib/bbz170>.
- [Koh91] KOHONEN, TEUVO. "Self-Organizing Maps: Optimization Approaches". *Artificial Neural Networks*. Elsevier, 1991, 981–990 2.
- [LBBH98] LECUN, YANN, BOTTOU, LÉON, BENGIO, YOSHUA, and HAFNER, PATRICK. "Gradient-Based Learning Applied to Document Recognition". *Proceedings of the IEEE* 86.11 (1998), 2278–2324 4.
- [LLX*10] LIU, YANCHI, LI, ZHONGMOU, XIONG, HUI, et al. "Understanding of Internal Clustering Validation Measures". *2010 IEEE International Conference on Data Mining*. IEEE. 2010, 911–916 2.
- [LRG19] LUAN, SHENGHUA, REB, JOCHEN, and GIGERENZER, GERD. "Ecological Rationality: Fast-and-Frugal Heuristics for Managerial Decision Making under Uncertainty". *Academy of Management Journal* 62.6 (2019), 1735–1759 2.
- [LSP*10] LEX, ALEXANDER, STREIT, MARC, PARTL, CHRISTIAN, et al. "Comparative Analysis of Multidimensional, Quantitative Data". *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), 1027–1035. DOI: 10.1109/TVCG.2010.138 3.
- [MBS93] MARTINETZ, THOMAS M., BERKOVICH, STANISLAV G., and SCHULTEN, KLAUS J. "'Neural-Gas' Network for Vector Quantization and Its Application to Time-Series Prediction". *IEEE Transactions on Neural Networks* 4.4 (1993), 558–569 2.
- [ML11] MURTAGH, FIONN and LEGENDRE, PIERRE. "Ward's Hierarchical Clustering Method: Clustering Criterion and Agglomerative Algorithm". *arXiv preprint arXiv:1111.6285* (2011) 10.
- [MTW*18] MA, YUXIN, TUNG, ANTHONY K. H., WANG, WEI, et al. "ScatterNet: A Deep Subjective Similarity Model for Visual Analysis of Scatterplots". *IEEE Transactions on Visualization and Computer Graphics* 26.3 (2018), 1562–1576 2–4.
- [MV15] METSALU, TAUNO and VILO, JAAK. "ClustVis: A Web Tool for Visualizing Clustering of Multivariate Data Using Principal Component Analysis and Heatmap". *Nucleic Acids Research* 43.W1 (May 2015), W566–W570. ISSN: 0305-1048. DOI: 10.1093/nar/gkv468. eprint: <https://academic.oup.com/nar/article-pdf/43/W1/W566/7476330/gkv468.pdf>. URL: <https://doi.org/10.1093/nar/gkv468>.

- [MW17] MOSELEY, BENJAMIN and WANG, JOSHUA. “Approximation Bounds for Hierarchical Clustering: Average Linkage, Bisecting k-Means, and Local Search”. *Advances in Neural Information Processing Systems* 30 (2017) 2, 10.
- [NJW01] NG, ANDREW, JORDAN, MICHAEL, and WEISS, YAIR. “On Spectral Clustering: Analysis and an Algorithm”. *Advances in Neural Information Processing Systems* 14 (2001) 10.
- [PKF*16] PANDEY, ANSHUL VIKRAM, KRAUSE, JOSUA, FELIX, CRISTIAN, et al. “Towards Understanding Human Similarity Perception in the Analysis of Large Sets of Scatter Plots”. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016, 3659–3669 3, 4.
- [PND20] PHAM, VUNG, NGUYEN, NGAN V. T., and DANG, TOMMY. “ScagCNN: Estimating Visual Characterizations of 2D Scatterplots via Convolution Neural Network”. *Proceedings of the 11th International Conference on Advances in Information Technology*. 2020, 1–9 3.
- [PVG*11] PEDREGOSA, FABIAN, VAROQUAUX, GAËL, GRAMFORT, ALEXANDRE, et al. “Scikit-learn: Machine Learning in Python”. *Journal of Machine Learning Research* 12 (2011), 2825–2830. URL: <https://scikit-learn.org/stable/modules/classes.html#5Cmodule-sklearn.manifold> 6, 9.
- [QNWR22] QUADRI, GHULAM JILANI, NIEVES, JENNIFER ADORNO, WIERNIK, BRENTON M., and ROSEN, PAUL. “Automatic Scatterplot Design Optimization for Clustering Identification”. *IEEE Transactions on Visualization and Computer Graphics* (2022) 3, 4, 12.
- [QR20] QUADRI, GHULAM JILANI and ROSEN, PAUL. “Modeling the Influence of Visual Density on Cluster Perception in Scatterplots Using Topology”. *IEEE Transactions on Visualization and Computer Graphics* 27.2 (2020), 1829–1839 3, 4.
- [QR21] QUADRI, GHULAM JILANI and ROSEN, PAUL. “A Survey of Perception-Based Visualization Studies by Task”. *IEEE Transactions on Visualization and Computer Graphics* (2021) 4.
- [QSMG17] QI, CHARLES R., SU, HAO, MO, KAICHUN, and GUIBAS, LEONIDAS J. “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, 652–660 3.
- [QYSG17] QI, CHARLES RUIZHONGTAI, YI, LI, SU, HAO, and GUIBAS, LEONIDAS J. “PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space”. *Advances in Neural Information Processing Systems* 30 (2017) 2–4.
- [RFB15] RONNEBERGER, OLAF, FISCHER, PHILIPP, and BROX, THOMAS. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III*. Springer. 2015, 234–241 9.
- [Rou87] ROUSSEEUW, PETER J. “Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis”. *Journal of Computational and Applied Mathematics* 20 (1987), 53–65 8.
- [RPY*22] REN, YAZHOU, PU, JINGYU, YANG, ZHIMENG, et al. “Deep Clustering: A Comprehensive Survey”. *arXiv preprint arXiv:2210.04142* (2022). URL: <https://arxiv.org/abs/2210.04142> 3.
- [SA15] SEDLMAIR, MICHAEL and AUPETIT, MICHAËL. “Data-Driven Evaluation of Visual Quality Measures”. *Computer Graphics Forum*. Vol. 34. 3. Wiley Online Library. 2015, 201–210 3.
- [SK20] SCHNELLBACH, JANIK and KAJO, MARTON. “Clustering with Deep Neural Networks—An Overview of Recent Methods”. *Network* 39 (2020) 2.
- [SMT13] SEDLMAIR, MICHAEL, MUNZNER, TAMARA, and TORY, MELANIE. “Empirical Guidance on Scatterplot and Dimension Reduction Technique Choices”. *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), 2634–2643 3, 4, 6, 10, 11.
- [SS04] SADLER-SMITH, EUGENE and SHEFY, ERELLA. “The Intuitive Executive: Understanding and Applying ‘Gut Feel’ in Decision-Making”. *Academy of Management Perspectives* 18.4 (2004), 76–91 2.
- [SSK*15] SACHA, DOMINIK, SENARATNE, HANSI, KWON, BUM CHUL, et al. “The Role of Uncertainty, Awareness, and Trust in Visual Analytics”. *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2015), 240–249 2.
- [TB21] TRAUTNER, THOMAS and BRUCKNER, STEFAN. “Line Weaver: Importance-Driven Order Enhanced Rendering of Dense Line Charts”. *Computer Graphics Forum*. Vol. 40. 3. Wiley Online Library. 2021, 399–410 4.
- [TBB*10] TATU, ANDRADA, BAK, PETER, BERTINI, ENRICO, et al. “Visual Quality Metrics and Human Perception: An Initial Study on 2D Projections of Large Multidimensional Data”. *Proceedings of the International Conference on Advanced Visual Interfaces*. 2010, 49–56 3, 4.
- [VA09] VANBELLE, SOPHIE and ALBERT, ADELIN. “Agreement Between an Isolated Rater and a Group of Raters”. *Statistica Neerlandica* 63.1 (2009), 82–100 2, 8.
- [vdMH08] Van der MAATEN, LAURENS and HINTON, GEOFFREY. “Visualizing Data Using t-SNE”. *Journal of Machine Learning Research* 9.11 (2008) 6.
- [vOVR22] Van ONZENODT, CHRISTIAN, VÁZQUEZ, PERE-PAU, and ROPINSKI, TIMO. “Out of the Plane: Flower vs. Star Glyphs to Support High-Dimensional Exploration in Two-Dimensional Embeddings”. *IEEE Transactions on Visualization and Computer Graphics* (2022) 6.
- [WCG*18] WANG, YUNHAI, CHEN, XIN, GE, TONG, et al. “Optimizing Color Assignment for Perception of Class Separability in Multiclass Scatterplots”. *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), 820–829 12.
- [WFC*17] WANG, YUNHAI, FENG, KANG, CHU, XIAOWEI, et al. “A Perception-Driven Approach to Supervised Dimensionality Reduction for Visualization”. *IEEE Transactions on Visualization and Computer Graphics* 24.5 (2017), 1828–1840 12.
- [WZM*19] WÖHLER, LESLIE, ZOU, YUXIN, MÜHLHAUSEN, MORITZ, et al. “Learning a Perceptual Quality Metric for Correlation in Scatterplots”. *Vision, Modeling and Visualization*. Ed. by SCHULZ, HANS-JÖRG, TESCHNER, MATTHIAS, and WIMMER, MICHAEL. The Eurographics Association, 2019. ISBN: 978-3-03868-098-7. DOI: [10.2312/vmv.20191318](https://doi.org/10.2312/vmv.20191318) 3.
- [XHL*22] XIA, JIAZHI, HUANG, LINQUAN, LIN, WEIXING, et al. “Interactive Visual Cluster Analysis by Contrastive Dimensionality Reduction”. *IEEE Transactions on Visualization and Computer Graphics* (2022) 3.
- [XLJ*21] XIA, JIAZHI, LIN, WEIXING, JIANG, GUANG, et al. “Visual Clustering Factors in Scatterplots”. *IEEE Computer Graphics and Applications* 41.5 (2021), 79–89 2, 3.
- [YMH*23] YANG, FUMENG, MA, YUXIN, HARRISON, LANE, et al. “How Can Deep Neural Networks Aid Visualization Perception Research? Three Studies on Correlation Judgments in Scatterplots”. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI ’23. Hamburg, Germany: Association for Computing Machinery, 2023. ISBN: 9781450394215. DOI: [10.1145/3544548.3581111](https://doi.org/10.1145/3544548.3581111). URL: <https://doi.org/10.1145/3544548.3581111> 6.
- [ZRL97] ZHANG, TIAN, RAMAKRISHNAN, RAGHU, and LIVNY, MIRON. “BIRCH: A New Data Clustering Algorithm and Its Applications”. *Data Mining and Knowledge Discovery* 1.2 (1997), 141–182 2, 7, 10.
- [ZZX*22] ZHOU, SHENG, XU, HONGJIA, ZHENG, ZHUONAN, et al. “A Comprehensive Survey on Deep Clustering: Taxonomy, Challenges, and Future Directions”. *arXiv preprint arXiv:2206.07579* (2022). URL: <https://arxiv.org/abs/2206.07579> 3.