

VariGrad: A Novel Feature Vector Architecture for Geometric Deep Learning on Unregistered Data

Emmanuel Hartman¹ (elh18e@fsu.edu) Emery Pierson² (emery.pierson@univ-lille.fr)[†]

Department of Mathematics, Florida State University, Tallahassee, USA¹
Institute of Mathematics, University of Vienna, Vienna, Austria²

Abstract

We present a novel geometric deep learning layer that leverages the varifold gradient (VariGrad) to compute feature vector representations of 3D geometric data. These feature vectors can be used in a variety of downstream learning tasks such as classification, registration, and shape reconstruction. Our model's use of parameterization independent varifold representations of geometric data allows our model to be both trained and tested on data independent of the given sampling or parameterization. We demonstrate the efficiency, generalizability, and robustness to resampling demonstrated by the proposed VariGrad layer.

CCS Concepts

• **Computing methodologies** → **Parametric curve and surface models; Shape analysis;**

1. Introduction

Geometric deep learning is an exciting field that has been rapidly evolving in recent years. It involves applying deep learning techniques to problems that involve geometric data such as point clouds, meshes, graphs, and curves. The ability to effectively learn from geometric data has numerous applications in fields such as computer graphics, computer vision, and robotics. In this project, we aim to contribute to the development of geometric deep learning by proposing a novel layer for deep neural networks. Our layer leverages the gradient of the varifold norm to define a vector field on a template shape which is then seen as a feature vector defined by the input. This method has the potential to improve the accuracy of deep learning models on geometric data, which could have a significant impact on a wide range of applications. In the following article, we will describe our method in detail and present experimental results that demonstrate its effectiveness for 3D curves and shape graphs.

1.1. Background & Related Work

The ability to compute similarity metrics between geometric objects is crucial in many applications. However, traditional methods that rely on aligning parameterizations can be computationally expensive and may not be suitable for large or complex datasets.

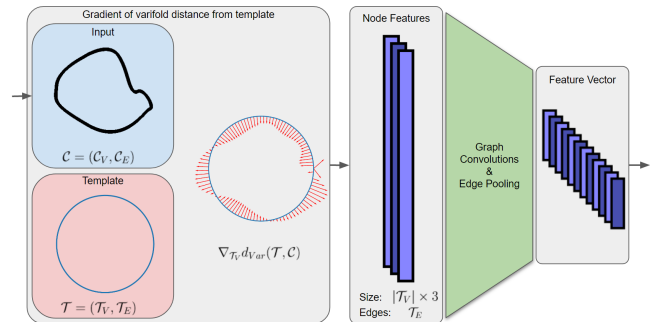


Figure 1: Illustration of the proposed methodology. The input geometric data and template shape are transformed into varifold representations and the gradient of the varifold norm is taken with respect to the vertices of the template. This gradient is passed through graph convolutional and edge pooling layers with respect to the template edges. Resulting feature vectors can be fed into traditional deep learning frameworks for reconstruction and classification tasks.

Varifold representations have emerged as a powerful tool for computing parameterization invariant similarity terms for geometric objects. Unlike traditional approaches, varifold representations are independent of the parameterization of the data, making them well-suited for cases where objects do not share a point-to-point correspondence. The space of varifolds has been equipped with a reproducing kernel Hilbert space (RKHS) norm. The benefits of us-

[†] E. Hartman was supported by NSF grant DMS-1953244. E. Pierson has been partially supported by the Austrian Science Fund (grant no P 35813-N).

ing this norm are that it is robust to noisy data and can be efficiently computed. The RKHS norm is also differentiable, making it suitable as a loss function for optimization-based shape matching [SBC19, HSK*23], or deep learning models for shape reconstruction [BPAD23].

Most of the pioneering works use well-defined shape descriptors to obtain a suitable representation of curves or shape graphs. Indeed, it is theoretically possible to retrieve a curve from its geometric moments, or its affine integral signature [FKK10], that have been applied in the context of curve classification. However, retrieving curves from such features implies solving costly optimization problems [KS21]. The Square Root Velocity Field framework has been used in elastic matching [SKJJ10] on 3D curves. It has been used on face-extracted curves for face identification [VvJD*11] and on extremal human curves [SWD14] for human pose retrieval. However, this elastic matching approach requires solving the problem of parameterization for each pair of curves. Moreover, their extension to shape graphs remains an active area of research [GSS21, GBNS22, SBC22].

In the recent years, a few works have shown the superiority of deep learning-based methods over the previous approaches. However, some approaches require strong hypotheses on the data, such as a common parameterization [HSC*21], or an ordering of the curve or shape graph [HE18, XJB21]. Meanwhile, a few approaches using convolutional neural networks on binary images have been proposed for sketch representation [QYC*22, MKDM22]. Not only those approaches don't take profit from the geometric information of curves, but they also do not extend to 3D data. Moreover, most of those approaches require training on large datasets [SBHH16, HE18] to be efficient.

Geometric deep learning [BBL*17] tries to circumvent those problems compared to the previous approaches. In particular, a few methods address the problem of dealing with data that lacks point-to-point correspondences or fixed shapes. Those methods interpret point clouds with an arbitrary number of inputs as a feature vector of a single-shaped object. PointNet [CSKG17] is a notable example of such methods, designed specifically for processing and classifying point clouds. PointNet employs a shared multi-layer perceptron (MLP) network to transform individual points into a high-dimensional feature space. The feature vectors are then aggregated using a symmetric function, such as max pooling, to generate a fixed-length representation of the entire point cloud. This representation can be fed into a fully connected neural network for classification or other downstream tasks. To overcome the limitation of solely relying on point cloud information without considering connectivity and local relationships between points, the Dynamic Graph Convolutional Neural Network (DGCNN) [WSL*19] method was introduced. DGCNN utilizes k-nearest neighbors (kNN) to establish edges between points, enabling the construction of a graph representation. By employing graph convolutional and edge pooling layers, DGCNN generates high-dimensional feature vectors for each point and aggregates them using a symmetric function. This results in a fixed-length representation of the entire point cloud that depends on the local relationships between points, which can be fed into a fully connected neural network for classification or other tasks. While some other architec-

tures [MKK21, SACO22, WNEH22] have shown more expressivity on 3D point clouds and surfaces, they rely on surface-specific operators such as the Laplace-Beltrami operator [SACO22], which cannot be used for 3D curves or shape graphs.

1.2. Contributions

In contrast, our proposed method directly operates on curves and shape graphs rather than their vertices as point clouds. We introduce a novel approach that leverages varifold representations to construct consistent-sized feature vectors, enabling the representation of geometric data irrespective of its sampling or registration. By utilizing varifold representations, our method captures the intrinsic properties of curves and shape graphs, leading to enhanced representations for geometric data. We propose a novel solution to this problem that leverages the differentiability of the varifold similarity metric. We compute the gradient of the distance between a template shape and the input shapes with respect to the points of the template shape. The resulting gradient is a vector field on the template shape that describes the optimal direction of deformation of the vertices of the template to minimize the distance between the shapes. This vector field can be interpreted as a feature vector for each input shape whose dimension is fixed regardless of the discretization or parameterization of the input shape. Traditional deep learning layers can process this feature vector, which is represented as a vector field lying on our template shape. Additionally, the vector field can be convolved with respect to the geometry of our template shape. For open curves, this is a 1D convolution, and for closed curves, this convolution is performed with cyclic padding. For surfaces and shape graphs, a graph convolution can be performed with respect to the underlying connectivity of the template shape. This approach provides a robust and flexible method for converting unregistered geometric objects into a fixed-dimensional feature vector that can be utilized in traditional and geometric deep-learning models.

2. Feature Vectors for 3D Curves and Shape Graphs

We focus on discrete 3D curves and shape graphs which are composed of vertices $V \subseteq \mathbb{R}^3$ and linear edges E connecting these vertices. However, a major challenge in working with such geometric data is the reliance on specific parameterizations, which can make downstream learning tasks and analysis difficult. To address this challenge, we adopt varifold representations as a powerful tool for characterizing and studying geometric objects without being dependent on their parameterizations.

2.1. Varifold Representations of 3D Curves and Shape Graphs

In recent years, varifold representations of geometric data have seen many applications in the field of shape analysis [SBC19, KCC17, CT13, CCG*20]. The space of varifolds is given by Borel measures on the Cartesian product of \mathbb{R}^3 and the unit sphere S^2 denoted by $\mathcal{M}(\mathbb{R}^3 \times S^2)$. Each discrete shape graph, $\mathcal{C} = (V, E)$, is mapped to a finitely supported measure $\mu \in \mathcal{M}(\mathbb{R}^3 \times S^2)$ where each support corresponds to each edge and is given by the centroid and normalized tangent vector of the edge. The mass of the support associated with an edge in the varifold representation is

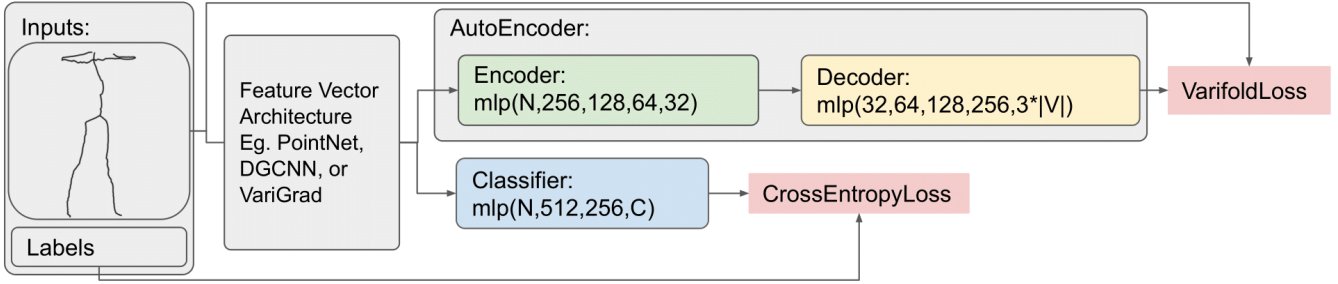


Figure 2: Network Architectures for Downstream Tasks: A diagram of the classifier and auto-encoder networks used for numerical experiments. Unparameterized 3D curves or shape graphs are given as inputs and feature vectors are extracted via a feature vector architecture. These feature vector architectures are then combined with the downstream architectures described in this figure and the combined networks are trained for their respective tasks.

determined by its length. This mapping ensures that regardless of the specific parameterization or ordering of points and edges in a shape graph, it will be identified with the same varifold representation. The space of varifolds is equipped with a Reproducing Kernel Hilbert Space (RKHS) norm, we define a suitable kernel function that measures the similarity between varifolds. The kernel function captures the inner product structure in the space of varifolds, allowing us to compute distances and perform various operations. This kernel function can be defined coordinate-wise by taking the product of a Gaussian kernel on \mathbb{R}^3 and a Binet kernel on S^2 . The inner product defined by this kernel is given by

$$\langle \mu, \nu \rangle_{Var} = \int_{\mathbb{R}^3 \times S^2} \int_{\mathbb{R}^3 \times S^2} e^{-a\|x_1 - x_2\|^2} \langle u_1, u_2 \rangle^2 d\mu d\nu$$

where a is a balancing parameter. From this inner product, a distance on the space of varifolds is given by

$$\|\mu - \nu\|^2 = \langle \mu, \mu \rangle_{Var} + \langle \nu, \nu \rangle_{Var} - 2\langle \mu, \nu \rangle_{Var}.$$

We define the dissimilarity between two discrete shape graphs $C_1 = (V_1, E_1)$ and $C_2 = (V_2, E_2)$ as the varifold distance between the associated finitely supported varifolds μ_1, μ_2 respectively. In particular, we write the dissimilarity between two shape graphs as $d_{Var}(C_1, C_2) = \|\mu_1 - \mu_2\|$. As these varifolds are finitely supported the inner product can be approximated by a sum over E_1 and E_2 given by

$$\langle \mu_1, \mu_2 \rangle_{Var} = \sum_{e_i \in E_1} \sum_{e_j \in E_2} e^{-a\|c_i - c_j\|^2} \langle u_i, u_j \rangle^2 l_i l_j.$$

Here c_i, c_j denote the centroids of e_i, e_j , u_i, u_j denote the normalized tangent vectors of e_i, e_j , and l_i, l_j denote the lengths of e_i, e_j .

By utilizing varifold representations equipped with a RKHS norm, we can effectively and efficiently capture the essential geometric information of 3D curves and shape graphs in a parameterization-independent manner. This allows for meaningful and efficient comparisons, analysis, and processing of geometric data, as well as facilitating tasks such as shape matching, recognition, and reconstruction. Moreover, the invariance of varifold rep-

resentations to parameterization changes provides robustness and consistency in geometric deep learning tasks.

2.2. VariGrad Feature Vector Architecture

Our architecture is designed for data-driven applications where the geometric data used for training or applying the model can be matched to a template shape \mathcal{T} . Leveraging this template shape, we compute the gradient of the varifold distance between each input shape and the template shape, specifically with respect to the vertices of the template shape. This computation results in a vector field on the template that can be viewed as a representation of the input. This data is fed into a sequence of convolutional layers where the convolution is performed with respect to the template shape [KW17]. The resulting vector field is concatenated into a fixed-dimension representation of the input data regardless of the sampling of the input data. A visual representation of the VariGrad architecture is given in Figure 1. Moreover, we have made a Pytorch implementation of this framework as well as the trained VariGrad models available on Github[†].

Our framework sees improved results compared to methods that combine pointwise Multi-Layer Perceptron (MLP) and feature point aggregation. This is mainly due to the utilization of the geometric prior given by the template shape on the input data. Moreover, rather than treating the points individually our method is able to efficiently process whole curves and shape graphs even in the unregistered setting, thanks to their varifold representations. Moreover, as they use MLP to embed each point in the point cloud in a high dimensional space, PointNet and DGCNN require more trainable parameters than the graph edge convolutions [KW17] used by VariGrad.

3. Experimental Results

The following experiments aim to compare the performance of the proposed VariGrad layer to the commonly used PointNet and

[†] https://github.com/emmanuel-hartman/Pytorch_VariGrad

DGCNN architectures in geometric deep learning applications where point-to-point correspondences are unknown. Each experiment constructs models that utilize the PointNet, DGCNN, and VariGrad architectures respectively to learn a consistent size feature vector. The feature vectors from these architectures are used in further learning tasks such as classification or shape reconstruction. Other than the differences in resources required by the different feature vector architectures these models use the same trainable parameters, loss functions, optimization parameters, and number of training epochs.

3.1. Datasets

To evaluate the performance of the methods, similar networks will be trained using these layers on two different datasets. We consider one dataset of 3D curves and one dataset of 3D shape graphs that are extracted from 3D surface data. We first extract landmarks from the shapes available in those datasets and compute the surface geodesics on the given shapes, giving curves that represent the pose of the 3D shapes. An illustration of the process is visualized in Figure 3. We use this geometric data as input of the different models. In both datasets, the extraction process produces unregistered curves and shape graphs, meaning that there is no prior point correspondence present in the data. For both of the datasets we consider, we perform a random split between training and testing set, ensuring that the classes (face or body identities) are well-balanced between the two splits. Moreover, we take the template shape to be a random shape from the corresponding dataset. The first dataset comprises boundary curves of human faces extracted from the COMA dataset of human face motions [RBSB18]. The extraction process is performed on all 20465 elements of the COMA dataset, resulting in 20465 closed curves in three dimensions. These curves are unparameterized and have a variety of samplings due to the extraction process with numbers of vertices ranging from 64 to 96. This dataset is then separated into a training set of 18506 curves and a validation set of 1959 curves. Along with each curve, a label indicating which of the eight identities present in the COMA dataset from which the curve was extracted. The second dataset consists of 41220 shape graphs representing human pose, extracted from the Dynamic FAUST dataset of 4D human motion scans [BRPMB17]. Again these shape graphs are unparameterized where the numbers of vertices in the shape graphs range from 111 to 356. Similar to the dataset of curves, the data is split into 35000 training samples and 6220 testing samples and an identity label is included for each shape graph in the dataset, corresponding to one of the twelve identities. By training and comparing the performance of the networks on these datasets, insights can be gained regarding the suitability and effectiveness of the VariGrad layer when compared to the state-of-the-art PointNet and DGCNN methods.

3.2. Classification

We evaluate the classification models based on their classification accuracy on the testing data. The results show that the model utilizing the VariGrad architecture outperforms the PointNet and DGCNN architectures. This difference in performance is likely due to the sparser set of points provided by the vertices of the three-dimensional curves considered in this experiment. The lower num-

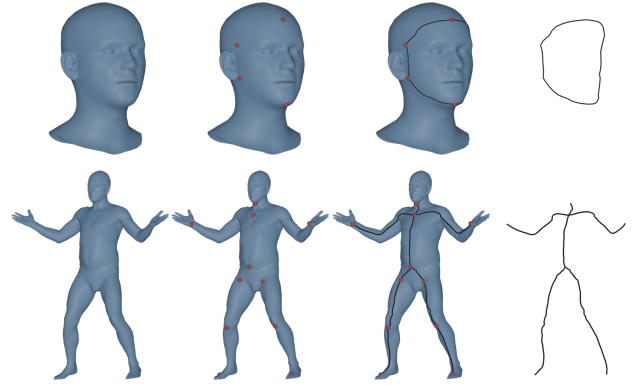


Figure 3: Examples of curve and shape graph extraction used for the creation of the example datasets. Surface meshes are given as inputs, landmarks are identified, geodesics along the surface are found between the landmark points, and finally these geodesic curves are concatenated to create the curves and shape graphs.

ber of points may not be sufficient for the PointNet and DGCNN architectures to learn representative feature vectors. The downstream network is an MLP, for which the architecture is outlined in Figure 2. Here N denotes the dimension of the feature vector recovered by the respective architectures and C denotes the number of classes in the data set. The input and hidden fully connected layers are followed by a ReLU activation function and a batch normalization layer. We train models using each feature vector architecture

	Face Curve Data		Shape Graph Data	
	Classification Accuracy	Train Time (s / batch)	Classification Accuracy	Train Time (s / batch)
PointNet	29.36	0.157	78.08	0.148
DGCNN	62.71	0.099	83.29	0.113
VariGrad	84.37	0.045	83.31	0.046

Table 1: Classifications Results: We present the results of the classification network utilizing three different feature vector architectures. We present the percentage of correct identity classifications by our model on an unseen testing set. Further, we present the mean training time per batch (10 samples) of each network to underscore the efficiency of the VariGrad approach.

on both the COMA curve training set and DFAUST shape graph training set. We report the results of these experiments in Table 1 along with the average training time of each model. In both the cases of curves and shape graphs, the models using the VariGrad feature vector architecture produce higher classification accuracies than the models utilizing the PointNet or DGCNN architectures. Moreover, the reliance on less trainable parameters means that the average training time of the VariGrad architecture is significantly lower than that of PointNet and DGCNN.



Figure 4: Qualitative comparison of shape reconstruction results for three examples from the testing set of the human pose shape graph dataset. We display the target shape graph (orange), with the reconstructions of PointNet (purple), DGCNN (green), and VariGrad (blue).

3.3. Shape Reconstruction

For the task of shape reconstruction, three fully connected auto-encoder models are trained to learn a latent code representation of the shapes based on the feature vectors produced by PointNet, DGCNN, and VariGrad, respectively. The auto-encoder architecture is outlined in Figure 2 where N denotes the dimension of the feature vector recovered by the respective architectures and $|V|$ denotes the number of desired vertices for the shape reconstruction. Again, we train models using each feature vector architecture on

	Face Curve Data		Shape Graph Data	
	Mean Error	Train Time (s / batch)	Mean Error	Train Time (s / batch)
PointNet	0.0053	0.149	0.0039	0.174
DGCNN	0.0010	0.113	0.0015	0.133
VariGrad	0.0004	0.073	0.0005	0.072

Table 2: Shape Reconstruction Results: We present the results of the shape auto-encoder network utilizing three different feature vector architectures. We present the mean varifold error for an unseen testing set as well as the mean training time per batch (10 samples).

both the COMA curve training set and DFAUST shape graph training set. To evaluate each network we compute the mean of the squared varifold norm between unseen testing sets of shapes and the reconstructions produced by each network. We report the results of these experiments in Table 2 with the average training times of each model. Again, the VariGrad architecture requires less training time than PointNet or DGCNN. Meanwhile, the VariGrad architecture produces a lower mean error than PointNet or DGCNN for both curves and shape graphs. For the testing set of shape graphs, we show a qualitative comparison of the reconstructions produced by each of the feature vector architectures in Figure 4. This comparison confirms that VariGrad not only produces numerically superior results but also visually more accurate reconstructions of 3D geometric data.

3.4. Generalizability and Reparameterization Invariance

The final set of experiments we present demonstrates the generalizability of our model and the invariance of our model to reparameterizations of the input data. To do this, we consider two additional testing sets of shape graphs. First, we consider resampled versions of the DFAUST shape graph testing set. For this dataset,

we apply random reparameterizations to each curve in the shape graph. This includes identity labels so that we can evaluate both the reconstruction and classification networks on the data. Second, we extract 90 shape graphs from the FAUST dataset of 3D scans of human poses. This dataset includes less extreme poses than the DFAUST shape graph dataset but does include some variability in the rigid alignment of the shape graphs not seen in the training data. In Table 3, we report the evaluation metrics for the trained classi-

	Reparameterized DFAUST		FAUST
	Classification Accuracy	Reconstruction Mean Error	Reconstruction Mean Error
PointNet	49.98	0.0040	0.0061
DGCNN	50.63	0.0049	0.0064
VariGrad	80.91	0.0005	0.0040

Table 3: Generalizability Results: Here we test our pre-trained shape graph reconstruction and classification models on two additional datasets which include variability unseen in the training.



Figure 5: Qualitative comparison of each model's invariance to reparameterization. We display 100 resampled versions of a single shape graph (orange) and the 100 reconstructions produced by PointNet (purple), DGCNN (green), and VariGrad (blue).

fication and reconstruction models applied to the reparameterized DFAUST shape graphs and for the trained reconstruction model applied to the FAUST shape graphs. In these results, we show that the VariGrad architecture produces superior results for both classification and reconstruction tasks. In particular, for the reparameterized DFAUST data, we report similar results to the original testing set in both classification and reconstruction tasks. Meanwhile, the DGCNN architecture sees a significant decrease in classification and reconstruction accuracy for the reparameterized data.

Finally, we propose an experiment to evaluate qualitatively the parameterization invariance of the different models. We created 100 reparameterizations of a single randomly selected shape graph from the DFAUST shape graph data set and applied them as input for the trained autoencoders. We collected the 100 outputs for each method and display them in Figure 5. We observe that our model keeps a very low variation of output shapes, as opposed to both PointNet and DGCNN. This visually confirms the invariance of our model to resampling and reparameterization.

4. Conclusion

In this paper, we presented VariGrad, a new feature vector architecture for learning tasks on geometric data such as curves and shape graphs. As we demonstrated in the experiments, by leveraging the varifold representation, this architecture produces superior results in downstream learning tasks. Moreover, varifold representations are not dependent on the given parameterization of the data and thus the VariGrad layer is robust to reparameterization and resampling of geometric data. We plan in the short future to extend this work to new applications, such as hand sketches. However, as the topology of hand sketches often vary in a same class of drawings, the choice of VariGrad’s template shape would have to be done more carefully. Moreover, as the varifold representation of curves has a natural extension to surfaces and point clouds, we plan to focus on extending VariGrad to such objects in the future.

References

- [BBL*17] BRONSTEIN M. M., BRUNA J., LECUN Y., SZLAM A., VANDERGHEYNST P.: Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42. 2
- [BPAD23] BESNIER T., PIERSON E., ARGUILLÈRE S., DAUDI M.: Toward mesh-invariant 3d generative deep learning with geometric measures. *Computers & Graphics* (2023). URL: <https://www.sciencedirect.com/science/article/pii/S009784932300122X>, doi:<https://doi.org/10.1016/j.cag.2023.06.027>. 2
- [BRPMB17] BOGO F., ROMERO J., PONS-MOLL G., BLACK M. J.: Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (July 2017). 4
- [CCG*20] CHARON N., CHARLIER B., GLAUNÈS J., GORI P., ROUSSILLON P.: 12 - fidelity metrics between curves and surfaces: currents, varifolds, and normal cycles. In *Riemannian Geometric Statistics in Medical Image Analysis*, Pennec X., Sommer S., Fletcher T., (Eds.). Academic Press, 2020, pp. 441–477. URL: <https://www.sciencedirect.com/science/article/pii/B9780128147252000212>, doi:<https://doi.org/10.1016/B978-0-12-814725-2.00021-2>. 2
- [CSKG17] CHARLES R. Q., SU H., KAICHUN M., GUIBAS L. J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Jul 2017), IEEE, p. 77–85. 2
- [CT13] CHARON N., TROUVÉ A.: The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM Journal on Imaging Sciences* 6, 4 (2013), 2547–2580. URL: <https://doi.org/10.1137/130918885>, doi:[10.1137/130918885](https://doi.org/10.1137/130918885). 2
- [FKK10] FENG S., KOGAN I., KRIM H.: Classification of curves in 2d and 3d via affine integral signatures. *Acta applicandae mathematicae* 109, 3 (2010), 903–937. 2
- [GBNS22] GUO X., BAL A. B., NEEDHAM T., SRIVASTAVA A.: Statistical shape analysis of brain arterial networks (ban), 2022. [arXiv:2007.04793](https://arxiv.org/abs/2007.04793). 2
- [GSS21] GUO X., SRIVASTAVA A., SARKAR S.: A quotient space formulation for generative statistical analysis of graphical data. *Journal of Mathematical Imaging and Vision* 63 (07 2021). doi:[10.1007/s10851-021-01027-1](https://doi.org/10.1007/s10851-021-01027-1). 2
- [HE18] HA D., ECK D.: A neural representation of sketch drawings. In *International Conference on Learning Representations* (2018). URL: <https://openreview.net/forum?id=Hy6GHpkCW>. 2
- [HSC*21] HARTMAN E., SUKURDEEP Y., CHARON N., KLASSEN E., BAUER M.: Supervised deep learning of elastic srv distances on the shape space of curves. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2021), pp. 4425–4433. 2
- [HSK*23] HARTMAN E., SUKURDEEP Y., KLASSEN E., CHARON N., BAUER M.: Elastic shape analysis of surfaces with second-order sobolev metrics: a comprehensive numerical framework. *International Journal of Computer Vision* 131, 5 (2023), 1183–1209. 2
- [KCC17] KALTENMARK I., CHARLIER B., CHARON N.: A general framework for curve and surface comparison and registration with oriented varifolds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (July 2017). 2
- [KS21] KOUSHOLT A., SCHULTE J.: Reconstruction of convex bodies from moments. *Discrete & Computational Geometry* 65, 1 (2021), 1–42. 2
- [KW17] KIPF T. N., WELING M.: Semi-supervised classification with graph convolutional networks, 2017. [arXiv:1609.02907](https://arxiv.org/abs/1609.02907). 3
- [MKDM22] MANDA B., KENDRE P. P., DEY S., MUTHUGANAPATHY R.: Sketchcleanet—a deep learning approach to the enhancement and correction of query sketches for a 3d cad model retrieval system. *Computers & Graphics* 107 (2022), 73–83. 2
- [MKK21] MITCHEL T. W., KIM V. G., KAZHDAN M.: Field convolutions for surface cnns. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 10001–10011. 2
- [QYC*22] QIN J., YUAN S., CHEN J., AMOR B. B., FANG Y., HOANG-XUAN N., CHU C.-B., NGUYEN-NGOC K.-N., CAO T.-T., NGO N.-K., ET AL.: Shrec’22 track: Sketch-based 3d shape retrieval in the wild. *Computers & Graphics* 107 (2022), 104–115. 2
- [RBSB18] RANJAN A., BOLKART T., SANYAL S., BLACK M. J.: Generating 3D faces using convolutional mesh autoencoders. In *European Conference on Computer Vision (ECCV)* (2018), pp. 725–741. URL: <http://coma.is.tue.mpg.de/>. 4
- [SACO22] SHARP N., ATTAIKI S., CRANE K., OVSJANIKOV M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)* 41, 3 (2022), 1–16. 2
- [SBC19] SUKURDEEP Y., BAUER M., CHARON N.: An inexact matching approach for the comparison of plane curves with general elastic metrics. In *2019 53rd Asilomar Conference on Signals, Systems, and Computers* (2019), pp. 512–516. doi:[10.1109/IEEECONF44664.2019.9049031](https://doi.org/10.1109/IEEECONF44664.2019.9049031). 2
- [SBC22] SUKURDEEP Y., BAUER M., CHARON N.: A new variational model for shape graph registration with partial matching constraints. *SIAM Journal on Imaging Sciences* 15, 1 (2022), 261–292. 2
- [SBHH16] SANGKLOY P., BURNELL N., HAM C., HAYS J.: The sketchy database: Learning to retrieve badly drawn bunnies. *ACM Trans. Graph.* 35, 4 (jul 2016). URL: <https://doi.org/10.1145/2897824.2925954>, doi:[10.1145/2897824.2925954](https://doi.org/10.1145/2897824.2925954). 2
- [SKJJ10] SRIVASTAVA A., KLASSEN E., JOSHI S. H., JERMYN I. H.: Shape analysis of elastic curves in euclidean spaces. *IEEE transactions on pattern analysis and machine intelligence* 33, 7 (2010), 1415–1428. 2
- [SWD14] SLAMA R., WANNOUS H., DAUDI M.: 3D human motion analysis framework for shape similarity and retrieval. *Image and Vision Computing* 32, 2 (Feb. 2014), 131–154. 2

- [VvJD*11] VELTKAMP R. C., VAN JOLE S., DRIRA H., AMOR B. B., DAOUDI M., LI H., CHEN L., CLAES P., SMEETS D., HERMANS J., ET AL.: Shrec'11 track: 3d face models retrieval. In *3DOR@ Eurographics* (2011), pp. 89–95. [2](#)
- [WNEH22] WIERSMA R., NASIKUN A., EISEMANN E., HILDEBRANDT K.: Deltaconv: Anisotropic operators for geometric deep learning on point clouds. *Transactions on Graphics* 41, 4 (jul 2022). [doi:10.1145/3528223.3530166](https://doi.org/10.1145/3528223.3530166). [2](#)
- [WSL*19] WANG Y., SUN Y., LIU Z., SARMA S. E., BRONSTEIN M. M., SOLOMON J. M.: Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)* 38, 5 (2019), 1–12. [2](#)
- [XJB21] XU P., JOSHI C. K., BRESSON X.: Multigraph transformer for free-hand sketch recognition. *IEEE Transactions on Neural Networks and Learning Systems* 33, 10 (2021), 5150–5161. [2](#)