# SUB-RIEMANNIAN STRUCTURE FOR ORIENTED POINT CLOUDS

**Anonymous authors**
Paper under double-blind review

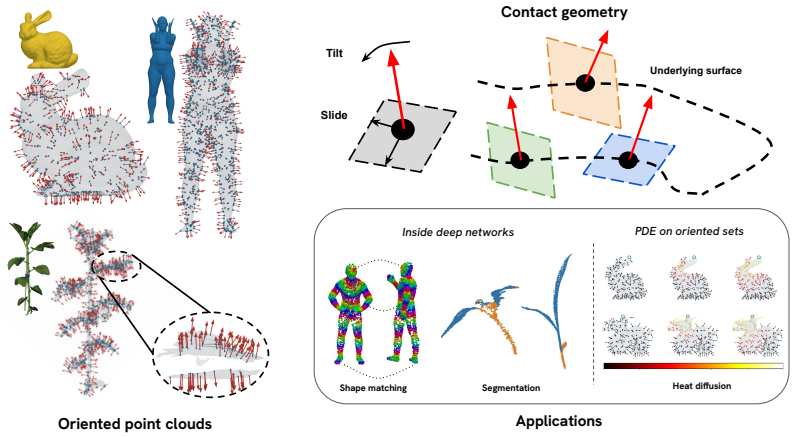



Figure 1: **We introduce a metric structure to define locality within oriented point clouds.** These are lifted to position–orientation space in which tangential translation and normal reorientation define a contact sub-Riemannian structure. An explicit pairwise score enables to compute local neighborhoods that can be used downstream for learning tasks and other shape processing applications such as heat diffusion on oriented point clouds.

## ABSTRACT

Many frameworks dealing with 3D (oriented) point clouds define locality in the ambient Euclidean space: $k$-nearest-neighbor graphs and radius balls in $\mathbb{R}^3$ are critical to define which points are "close to one another" and should exchange information. This can unintentionally merge sampled surfaces that are close in space but geometrically distinct such as the two faces of a thin leaf. Then, features are more prone to leak across boundaries which hinder performance. To address this, we propose to define a sub-Riemannian contact metric on position-orientation space $\mathbb{M} = \mathbb{R}^3 \times \mathbb{S}^2$: points are lifted to $(x, n)$ with $n$ a surface normal, and "motion" is constrained to a non-holonomic horizontal distribution in which spatial velocity must stay in the plane orthogonal to the normal which can reorient at a cost. From this, we derive an explicit score, which locally approximate the distance, to compare oriented points for downstream tasks. We report improved segmentation and correspondence accuracy near contact when using the score as a drop-in change inside PointNet++ aggregation mechanism as well as reduced leakage when using it to simulate heat diffusion on oriented points.

## 1 INTRODUCTION

To deal with unstructured 3D data, such as (oriented) point clouds, triangle soups and Gaussian splats (Kerbl et al., 2023), many networks (Qi et al., 2017; Wang et al., 2019), equivariant architectures (Thomas et al., 2018; Fuchs et al., 2020), and classical geometry processing algorithms (Taubin, 1995; Fleishman et al., 2003), define locality as Euclidean. Neighborhoods are $k$-Nearest Neighbor ($k$NN) sets or radius balls in $\mathbb{R}^3$, and convolutions, message passing, or diffusions act over them.



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

One way or another, this encodes the assumption that points close in $\mathbb{R}^3$ are more likely to exchange information. This assumption is challenged as soon as distinct pieces of surface pass near one another: two stacked leaves in a plant, a wall, or a sheet of metal are separated by a relatively small distance and angle (opposite or co-oriented) normals. A point near an acute angle has a $k$NN ball encompassing both incident faces, folded and layered geometry stacks surfaces within any fixed radius. In all these cases Euclidean neighborhoods mix samples that should belong to distinct geometric semantic (such as for two parallel sheets). This can make denoisers merge distinct folds, segmentation labels may leak through thin structures. The standard workaround such as visibility heuristics, clustering the $k$NN by normal, hand-tuned anisotropic neighborhoods, and others treat symptoms case by case depending on the task. To address this, some approaches starts from the Euclidean neighborhood and re-weight it by normal similarity (Fleishman et al., 2003; Guerrero et al., 2018), or add a normal term to the distance (Rusinkiewicz & Levoy, 2001) but these are limited for describing strong orientation anisotropy. In this paper, we propose a single geometric principle to replace the Euclidean default. We lift each sample point $x \in \mathbb{R}^3$ to position-orientation space $\mathbb{M} = \mathbb{R}^3 \times \mathbb{S}^2$, $g = (x, n)$ with $n$ the (estimated) unit normal, and we define the sub-bundle $\Delta \subset T\mathbb{M}$ as,

$$\Delta_{(x,n)} \;=\; \big\{(\dot{x}, \dot{n}) \;:\; \langle \dot{x}, n \rangle = 0\big\}, \tag{1}$$

A metric is then built as the cost of "motion from a source to a target": we constrain these motions in $\mathbb{M}$ to belong in $\Delta$. This means that spatial motion must stay perpendicular to the normal of the considered point, while the latter may reorient at a cost set by a weight $w > 0$. The forbidden direction is translation along the normal which is intuitively the direction in which Euclidean neighborhoods short-circuit and leak. By tilting (reorienting the normal), sliding, and tilting back, normal displacement can still be achieved, but at an additional cost intrinsic to the metric. This is drawn from the following intuitive hypothesis: **a surface element can slide and tilt but not hop off itself** as illustrated in Figure 2.

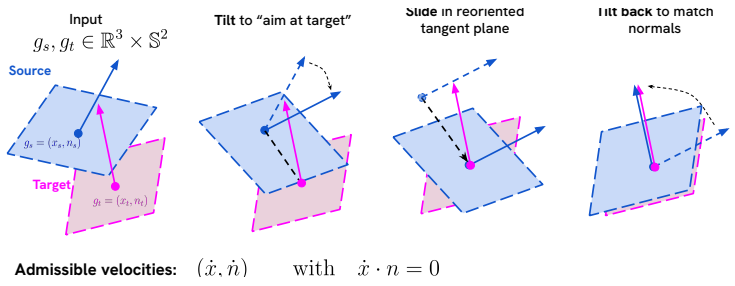



Figure 2: **Illustration of the mechanism behind the metric for matching two oriented points.** A source oriented point $g_s$ has to match a target one $g_t$. To do so, $g_s$ has to tilt first to be able to align the point position and then reorient back to match orientation. The motion has to make a detour to reach the target configuration which is exactly what a sub-Riemannian structure induces.

Sub-Riemannian (SR) geometry offers a rigorous framework for this construction including in particular a strongly anisotropic metric to encode forbidden short-circuits mentioned earlier. We strongly advise the reader to read through Appendix A for a simplified background on sub-riemannian geometry and contact metric. This type of structure on roto-translation groups has demonstrated its interest for contour perception and crossing-preserving image enhancement (Petitot, 2003; Citti & Sarti, 2006; Franken & Duits, 2009; Duits et al., 2016; Boscain et al., 2014), and have recently powered equivariant PDE networks on images (Smets et al., 2023; Bellaard et al., 2023). Those constructions lift curves in images, aligning spatial motion along the orientation. Here, we propose to adapt this idea to surfaces described as unstructured (oriented) point sets rather than grids.

**Contributions.** The metric structure itself is not new but its use and study as a locality prior within discrete 3D sets. In particular, this paper gather the following

1. **A framework** defining a sub-Riemannian contact metric for sampled surfaces seen as position-orientation sets.
2. **A computable contact score** leveraging the Heisenberg tangent model with a key property: a linear growth in the normal offset leading to a sharper separation power.
3. **Downstream applications** demonstrating the interest of the proposed geometric prior for learning on oriented point clouds and for heat diffusion on oriented geometric measures

## 2 RELATED WORK

**Euclidean locality and its patches.** Point networks build $k$NN/radius graphs in $\mathbb{R}^3$ (Qi et al., 2017; Wang et al., 2019; Monti et al., 2017; Thomas et al., 2019; Wu et al., 2024). Failures at thin structures are well known in geometry processing. To address it, bilateral weighting by normal similarity (Fleishman et al., 2003), learned local shape descriptors (Guerrero et al., 2018) or normal surface fitting Guennebaud & Gross (2007); Öztireli et al. (2009) can be used to mitigate them. These approaches keep the Euclidean candidate set and re-weight it. Here, we change the locality formalism that defines the candidate set, the ranking, and the weights.

**Position-orientation and group-equivariant networks.** Group convolutions (Cohen & Welling, 2016; Bekkers et al., 2018b; Bekkers, 2020; Weiler et al., 2018; Thomas et al., 2018; Fuchs et al., 2020) lift features to (quotients of) $\mathrm{SE}(3)$. PONITA (Bekkers et al., 2024) operates on $\mathbb{R}^3 \times \mathbb{S}^2$ and conditions messages on roto-translation invariants of point pairs: orientation is a feature. Bellaard & Smets (2025) characterize these pair invariants exhaustively and we borrow some of these. Locality, however, remains a Euclidean graph, nothing in the message function is constrained to suppress flow between incompatible orientations at first order. The same goes with Petrov et al. (2026) which lifts points to the product of ambient space with the Grassmannian of tangent planes, iterates the construction so that samples differing in curvature or higher-order contact also separate. However, their lift is unoriented while ours is co-oriented (with $\mathbb{M} \cong S(T^*\mathbb{R}^3)$ and their lifted metric is a product metric with the same linear separation.

**Sub-Riemannian image analysis.** The roto-translation lift of 2D/3D images is used in (Petitot, 2003; Citti & Sarti, 2006), and can further be developed for crossing-preserving diffusion (Franken & Duits, 2009), hypoelliptic inpainting (Boscain et al., 2014), data-driven SR geodesics (Bekkers et al., 2015; Duits et al., 2016; 2018; Mirebeau, 2018), and PDE-based equivariant networks whose sharp anisotropy requires the sub-Riemannian kernel (Smets et al., 2023; Bellaard et al., 2023). There, admissible spatial motion is along the lifted orientation (curves), here it is perpendicular to it. In particular, the distribution has rank 4 instead of 3. Our use of the nilpotentization to compute a score follows from (Bekkers et al., 2018a), who applied the same device to perceptual grouping of blood vessels. (Portegies et al., 2015) give a related kernel approximation on SE(3). Finally, the data itself is not the same: we act on unstructured, oriented point sets instead of grids, and the target is a graph construction rather than a PDE solver.

## 3 A SUB-RIEMANNIAN STRUCTURE ON ORIENTED POINT CLOUDS

To formalize the illustration from Figure 2, we define the state space, admissible motions and properties of contact geometry Geiges (2008) in our context of sampled surfaces.

**State space.** An oriented sample is $g = (x, n) \in \mathbb{M} := \mathbb{R}^3 \times \mathbb{S}^2$. The group $\mathrm{SE}(3)$ acts by $(R, t) \cdot (x, n) = (Rx + t,\, Rn)$. The action is transitive with stabilizer $\mathrm{SO}(2)$ (rotations about the normal), so $\mathbb{M} \cong \mathrm{SE}(3)/\mathrm{SO}(2)$ is a 5-dimensional homogeneous space. A $C^2$ oriented surface $\Sigma \subset \mathbb{R}^3$ lifts to the 2-dimensional sub-manifold $\tilde{\Sigma} = \{(x, n(x)) : x \in \Sigma\} \subset \mathbb{M}$.

**Admissible motions.** An element of the state space can slide in its tangent plane and tilt its normal vector, but cannot slide along it. Reorienting by $\alpha$ costs the same as sliding $w\alpha$ along the surface.

**Definition 3.1** (Horizontal distribution and SR metric)**.** *For $w > 0$ define*

$$\Delta_{(x,n)} = \big\{(p, u) \in \mathbb{R}^3 \times T_n\mathbb{S}^2 : \langle p, n\rangle = 0\big\}, \qquad G_{(x,n)}((p, u), (q, v)) = \langle p, q\rangle + w^2\langle u, v\rangle \quad (2)$$

*A path $t \mapsto \gamma(t) = g(t) = (x(t), n(t))$ is* horizontal *if $\dot\gamma(t) \in \Delta_{\gamma(t)}$ a.e., its length is $l(\gamma) = \int \sqrt{G(\dot\gamma)}\, dt$, and the Carnot-Carathéodory distance $d_{\mathrm{CC}}(g(0), g(1))$ is the minimal length of horizontal paths joining $g(0)$ to $g(1)$.*

$\Delta$ has rank 4 in a 5-manifold: two tangential translations and two tilts. The missing direction is translation along $n$.

**Proposition 3.2** (Contact structure)**.** *Let $\theta_{(x,n)}(\dot x, \dot n) = \langle \dot x, n\rangle$, so $\Delta = \ker\theta$. Then, $\theta$ is a contact form: $\theta \wedge (d\theta)^2$ is nowhere zero. Consequently $(\mathbb{M}, \Delta)$ is the unit cotangent bundle $S(T^*\mathbb{R}^3)$ with its canonical contact structure.*

**Corollary 3.3** (Lifted surfaces are Legendrian)**.** *$\tilde\Sigma$ is an integral submanifold of $\Delta$ of the maximal dimension* 2*.*

**Proposition 3.4** (Step-2 bracket-generating). *$\Delta$ is bracket-generating of step 2: $\Delta + [\Delta, \Delta] = T\mathbb{M}$. Consequently, because of Chow-Rashevskii theorem, any two points of $\mathbb{M}$ can be joined by a horizontal path and $d_{\mathrm{CC}}$ is a finite metric inducing the manifold topology.*

This means that the (hard) constraint cannot be integrated but the forbidden direction can be reached, at a higher cost that we will now quantify and explain in more details.

## 4 A COMPUTABLE LOCAL CONTACT SCORE

In the following, $P, Q$ are oriented points with $x_P, x_Q$ their respective position in $\mathbb{R}^3$ and $n_P, n_Q$ their attached normal vectors. The true Carnot-Carathéodory distance $d_{CC}$ is

$$d_{CC}(P,Q) = \inf_{\gamma} \int_0^1 \sqrt{\|\dot{x}\|^2 + w^2 \|\dot{n}\|^2} \mathrm{d}t, \qquad \langle \dot{x}, n \rangle = 0 \tag{3}$$

This is a true metric but obtaining it require solving an optimal control problem: it has no global closed-form solution but it can be locally approximated.In this section, we (i) define a pair $(P, Q)$ through rigid invariants, (ii) replace $d_{CC}$ by the exact distance of its tangent model.

### 4.1 PAIR INVARIANTS

Let's first consider the four rigid-motion invariants

$$\|r\| = \|x_Q - x_P\|, \quad \alpha = \angle(n_P, n_Q) \in [0, \pi], \quad \zeta_P = r \cdot n_P, \quad \zeta_Q = r \cdot n_Q, \tag{4}$$

and set

$$\zeta = \tfrac{1}{2}(\zeta_P + \zeta_Q) = \left\langle r, \frac{n_P + n_Q}{2} \right\rangle \qquad \rho^2 = \|r\|^2 - \zeta^2 \tag{5}$$

Then, $\zeta$ is the symmetric weight-two contact coordinate (the normal offset) and $\rho$ is the tangential spatial separation. Note that $\zeta = 0$ means the vector $r$ makes equal and opposite angles with the tangent planes.

### 4.2 THE HEISENBERG TANGENT MODEL

Mitchell's theorem (Mitchell, 1985) identifies the metric tangent cone of a sub-Riemannian manifold at g with the Carnot group given by the nilpotentization of $(\Delta, G)$ at g. Since $\Delta$ is a contact distribution, this group is a Heisenberg group which in our case is five-dimensional $\mathbb{H}^5$. In this space, horizontal displacements of norm $a$ and vertical coordinate $\tau$ projects to a circular arc of radius $R$ and angle $\phi \in (0, 2\pi)$, $\tau$ being the area between the arc and its chord: $a = 2R\sin(\phi/2), |\tau| = \frac{1}{2}R^2(\phi - \sin\phi)$ and length $R\phi$ which gives the squared Heisenberg distance $H$ as

$$H(a^2, \tau) = R^2\phi^2 = a^2 \left( \frac{\phi}{2\sin(\phi/2)} \right)^2 = \frac{2|\tau|\phi^2}{\phi - \sin\phi} \tag{6}$$

where $\phi$ is the unique solution of

$$\frac{\phi - \sin\phi}{1 - \cos\phi} = \frac{4|\tau|}{a^2} = s \tag{7}$$

which is an isoperimetric problem. From this, we can define a "**detour cost**" $q$ such that

$$q(s) := \frac{\phi(s)^2}{4\sin^2(\phi(s)/2)} = 1 + \frac{3}{4}s^2 - \frac{9}{80}s^4 + O(s^6) \tag{8}$$

To specialize the tangent model to our metric, set $a^2 := \rho^2 + w^2\alpha^2, \tau := w\zeta$. Both terms in $a$ have length units, while $\tau$ has length-squared units, as required by the weight-two Heisenberg scaling.

**Proposition 4.1** (The Heisenberg distance approximates $d_{CC}$). *Fix $P$ and let $Q \to P$ and $\epsilon = \|r_T\| + w\alpha + \sqrt{w|\zeta|}$ with $r_T$ the projection of $r$ onto $n_P^{\perp}$ such that horizontal coordinates are $O(\epsilon)$ and the missing normal-translation coordinate is $O(\epsilon^2)$. Up to the sign convention for the central Heisenberg coordinate, $(a, \tau)$ defined above are the horizontal norm and weight-two coordinate of the $\mathbb{H}^5$ nilpotentization to leading order. Consequently,*

$$d_{CC}(P,Q)^2 = H(a^2, w\zeta) + o(\epsilon^2)$$

We then define the proposed **local contact score** $\ell_w$ such that

$$\ell_w(P,Q)^2 = H(a^2, w\zeta) = a^2 q\left(\frac{4|\tau|}{a^2}\right) \tag{9}$$

From this, as $a \to 0$ (meaning from $P$ to $Q$, normals are co-oriented and displacement is purely normal), we have $\ell_w^2 \to 4\pi w|\zeta|$ which means a pure normal offset of "size" $\delta$ costs $\sqrt{4\pi w\delta}$. $\ell_w$ is symmetric, $\mathrm{SE}(3)$-invariant and vanishes when $P \to Q$ but the triangle inequality is not verified in general which means $\ell_w$ is a semi-metric. For unoriented normals, we extend the score as

$$\ell_w^{\mathrm{proj}}((x,[n]),(y,[m])) = \min_{s\in\{-1,1\}} \ell_w((x,n),(y,sm)), \tag{10}$$

Beyond the asymptotic property of Proposition 4.1, we derive non-asymptotic properties of the score, ensuring it is motivated in general. In particular, the score grows linearly in the normal offset $|\zeta|$ while quadratic off-plane penalties such as PTP barely affects small offsets. We also study the, although ideal, useful case of stacked sheets with co-oriented and opposing normals. Details are given in Appendix B.

## 5 Experiments

We introduced this contact score as a geometric prior to increase the separation power in "contact-prone" contexts. Supporting this claim with practical applications beyond synthetic evaluation is tricky, but we report a set of results in this section. Complementary details and explorations can be found in Appendix C. In Section 5.1, we show how the strong anisotropy induced by the metric allows to address structural limitations when defining neighborhood in contact-prone structures. Then, in the following Section 5.2, we show that replacing the neighborhood mechanism in an established point-based neural network can improve its capabilities on these structures. Finally, we go further and use the kernel induced by the metric as a "similarity score" in a recent algorithm to define a diffusion process on weighted oriented discrete point sets (also called discrete varifolds (Charon & Trouvé, 2013; Kaltenmark et al., 2017; Buet et al., 2017) in the shape analysis literature), leveraging orientation in a more principled way than using it as feature. We denote the proposed contact metric structure (CONTACT). For an ordered pair $i = (x_i, n_i)$, $j = (x_j, n_j)$, we write $r = x_j - x_i$, $\alpha = \angle(n_i, n_j)$, $\zeta_i = \langle r, n_i\rangle$, $\zeta_j = \langle r, n_j\rangle$, $\zeta = \frac{1}{2}(\zeta_i + \zeta_j)$ and $\rho^2 = \|r\|^2 - \zeta^2$.

**Baseline methods.** To support its interest compared to classical ways to define "locality" in oriented point sets, we compare to the following baselines: euclidean EUC, bilateral (Fleishman et al., 2003) BIL (which is Euclidean neighbors whose features are weighted by $e^{-\alpha^2/2\sigma_\alpha^2}$ before max-pooling), product metric PROD and point-to-plane (Rusinkiewicz & Levoy, 2001) PTP.

**"Ablation" of the score.** To test whether the exact Heisenberg formulation is needed for the learning gains, we also replace the grouping score by two functions of the same invariants formalized in Table 1. LIN preserves linear growth in the normal offset, whereas QUAD uses a quadratic penalty.

### 5.1 Synthetic validations

We start by exhibiting the type of setting in which the contact structure we propose is the most relevant. We consider two coaxial cylinders of radii $R$ and $R + \delta = 1.08R$ (very close) and sample 12 equally spaced points from each with co-oriented normals. The task is, when selecting a point, computing its 2 closest neighbors (on each side). Such a neighbor is correct if it belongs to the cylinder of the queried point. We compare the resulting reconstruction and showcase in Figure 3 how the CONTACT score is the only approach among the mentioned baselines to correctly reconstruct the two cylinders. In particular, we see that the cylinders are separated by the weight-two coordinate $\zeta \approx \delta$. Note that if the cylinders would have been with opposing normals, $\zeta = 0$ and $\ell_w^2 = \rho^2 + w^2\alpha^2 \geq \pi^2 w^2$, independently of thickness. This will be even more visible with the heat diffusion experiment in 5.3. We also mention that this synthetic experiment is pushed further in Appendix C with effect of normal noise and staggered sampling.

### 5.2 Downstream learning applications

We consider two tasks: segmentation and shape correspondence. For both, we employ a PointNet++ (Qi et al., 2017) model, widely used for point cloud processing. It is notably relying on euclidean

Table 1: **Chosen baselines to define different notions of locality in oriented point sets.** The "Off-plane behavior" column states how a score depends on displacement along the normal, which is the direction Euclidean neighborhoods short-circuit. All scores are $SE(3)$-invariant, symmetric and definite. BIL is applied as a weight on a Euclidean computation and the weight depends on the angle $\alpha$ through the scale $\sigma_\alpha = 0.4$.

| Score | Squared pairwise score $D^2(i,j)$ | Off-plane behavior |
|---|---|---|
| EUC | $\|r\|^2$ | — |
| BIL | $\|r\|^2$, weight $\times\, e^{-\alpha^2/2\sigma_\alpha^2}$ | — |
| PROD (SECTION) | $\|r\|^2 + w^2\|n_i - n_j\|^2$ | product penalty |
| PROD (ANGULAR) | $\|r\|^2 + w^2\alpha^2$ | product penalty |
| PTP | $\|r\|^2 + 2\pi(\zeta_i^2 + \zeta_j^2)$ | quadratic off-plane |
| LIN | $\rho^2 + w^2\alpha^2 + 2\pi w\|\zeta\|$ | symmetrized, linear |
| QUAD | $\rho^2 + w^2\alpha^2 + 4\pi\zeta^2$ | symmetrized, quadratic |
| CONTACT | $H(\rho^2 + w^2\alpha^2,\ w\zeta)$ | symmetrized, quadratic to linear |

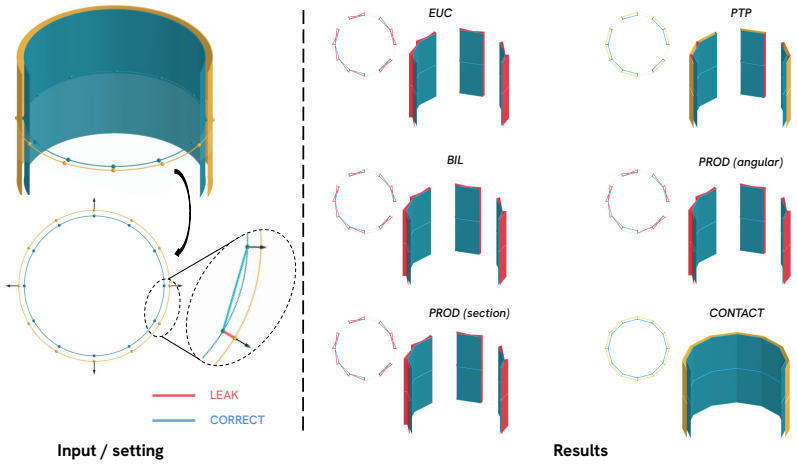



Figure 3: **Setting and results for the cylinder reconstruction experiment.** For each method, we take the undirected union of its two-neighbor graph and extrude every edge into a triangulated strip. Red marks cross-cylinder edges (leaks) and false strips. CONTACT is the only among compared approach to exhibit two separate polygonal cylinders.

neighborhood aggregation to produce a hierarchical pooling of features to better encode structure. We replace this aggregation with the baselines and report the final metrics.

**Plant segmentation.** *Data.* Pheno4D (Schunck et al., 2021) is a set of high-resolution laser scans of growing maize and tomato plants. We use every annotated scan: 126 point clouds over 14 plants (7 maize $\times$ 7 days, 7 tomato $\times$ 11 days), each carrying 1.5–4.5 million points. Both use the convention $0 =$ soil, $1 =$ stem, $2, 3, \ldots =$ leaf instances. We drop the soil, which is $85\%$ of every scan and separable by a height threshold, and keep the plant. Each cloud is subsampled to $2^{14} = 16{,}384$ points, equipped with unoriented normals estimated from local PCA over 24 neighbors, and normalized to the unit sphere. Training draws $4{,}096$ points per cloud. Because the normals are unoriented, we use equation 10.

*Task.* Binary segmentation on the plant: **stem or leaf**. Train/test splits are by plant, since the scans of one plant are the same plant on consecutive days and are not independent. Maize06, Maize07, Tomato06 and Tomato07 are held out for test, Maize05 and Tomato05 for validation, and the remaining 8 plants are used for training.

*Method.* We train a PointNet++ segmentation network (Qi et al., 2017) with four set-abstraction levels at 1024, 256, 64 and 32 centroids, $k = 32$ neighbors per centroid, a feature-propagation decoder, and a two-class head (leaf or stem). The network trained for 100 epochs, replacing only the grouping mechanism. Normals are an input feature in each variant, including EUC, so the baseline sees exactly

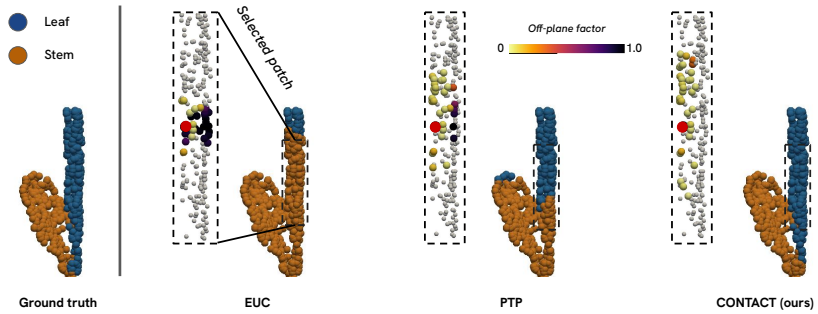



Figure 4: **Qualitative focus on the Pheno4D segmentation task.** For each method, we show a selected patch from the chosen neighborhood formulation and color points depending on an "off-plane" factor defined as $|\zeta|/\|r\|$.

the same information. All variants use $k$-nearest-neighbor grouping rather than a ball query, because a radius is not comparable across scores with different units. Three seeds are run. More training details in Appendix C.

*Results.* In Table 2, we report the test metrics. We highlight how using the proposed contact neighborhood increase the model's ability to segment in this context.

| grouping | mIoU | accuracy | med. $\lvert\zeta\rvert/\lVert r\rVert$ | med. normal dev. |
|---|---|---|---|---|
| EUC | $0.750 \pm 0.027$ | $0.913 \pm 0.014$ | 0.0680 | 12.62° |
| BIL | $0.789 \pm 0.026$ | $0.929 \pm 0.014$ | 0.0665 | 11.93° |
| PROD (Section) | $0.788 \pm 0.030$ | $0.930 \pm 0.013$ | 0.0666 | 11.96° |
| PTP | $0.790 \pm 0.011$ | $0.933 \pm 0.006$ | 0.0574 | 11.44° |
| LIN | $0.827 \pm 0.018$ | $0.944 \pm 0.009$ | 0.0476 | 10.38° |
| QUAD | $0.816 \pm 0.008$ | $0.942 \pm 0.003$ | 0.0551 | 10.08° |
| CONTACT | $\mathbf{0.835 \pm 0.023}$ | $\mathbf{0.950 \pm 0.007}$ | 0.0484 | 10.14° |

Table 2: **Pheno4D stem/leaf segmentation.** We report mean $\pm$ standard deviation over three seeds. Only the grouping geometry differs between rows. In the last two columns, we report how much normals agree on computed patch (from the same queried points) on average over ten point clouds.

*Analysis.* What the grouping changes is the shape of the patch. We measure the weight-two coordinate itself over the neighbors each score selects through $|\zeta|/\|r\|$, which express how far off the query's own tangent plane each grouped point lies. We do this over ten clouds and report the average in the last two columns of Table 2. We observe that the columns are almost monotone together. The contact score selects patches much "flatter" than Euclidean grouping and point-to-plane, the local MLP is handed a cleaner tangential neighborhood. Of course, correlation does not imply causation but we still report the trend.

**Dense shape correspondence on FAUST.** *Data.* FAUST (Bogo et al., 2014) provides 100 registered human body meshes: 10 subjects in 10 poses each, each one with the same 6890-vertex topology. Vertex $i$ is therefore the same semantic point on every shape, which gives dense ground-truth correspondence with no annotation step and no ambiguity. For this one, we either estimate normals by local PCA as before or we take the true oriented normals from the meshes.

*Task and metric.* Given a human mesh, label every point with the template vertex it corresponds to: a 6890-way per-point classification. The metric is the geodesic distance along the template surface between the predicted and the true vertex, normalized by the geodesic diameter. We report its mean geodesic error overall and focused on a "contact stratum": points at which some vertex within three sampling scales is more than 0.3 geodesic unit away (normalized scale) along the surface. This concerns 385 test points per mesh, and it is where the grouping geometry matter more. Geodesic distances are precomputed once on the template with the heat method (Crane et al., 2013).

*Method.* The network and protocol are similar to the plant segmentation experiment: a single-scale-grouping PointNet++ (Qi et al., 2017) with four set-abstraction levels (1024, 256, 64,32 centroids, $k = 32$), a feature-propagation decoder, and a 6890 head. Splits are by identity, test 8-9, validation 7, training 0-6. Three seeds are run.

*Results.* We report metrics in Table 3 and we observe that overall geodesic error is improved with the CONTACT neighborhood computation with PCA normals. With oriented normals, the differences are less significant and within one standard deviation.

| normals | grouping | geodesic error | **contact stratum** |
|---|---|---|---|
| unoriented (PCA) | EUC | $0.0137 \pm 0.0010$ | $0.0828 \pm 0.0110$ |
| | BIL | $0.0129 \pm 0.0015$ | $0.0675 \pm 0.0122$ |
| | PROD (SECTION) | $0.0132 \pm 0.0009$ | $0.0701 \pm 0.0080$ |
| | PTP | $0.0131 \pm 0.0008$ | $0.0740 \pm 0.0126$ |
| | LIN | $0.0126 \pm 0.0002$ | $0.0668 \pm 0.0090$ |
| | QUAD | $0.0128 \pm 0.0006$ | $0.0670 \pm 0.0076$ |
| | CONTACT | $\mathbf{0.0121 \pm 0.0007}$ | $\mathbf{0.0568 \pm 0.0073}$ |
| oriented (GT) | EUC | $0.0098 \pm 0.0005$ | $0.0434 \pm 0.0085$ |
| | BIL | $0.0102 \pm 0.0004$ | $0.0476 \pm 0.0072$ |
| | PROD (SECTION) | $0.0097 \pm 0.0004$ | $0.0414 \pm 0.0086$ |
| | PTP | $0.0100 \pm 0.0006$ | $0.0438 \pm 0.0056$ |
| | LIN | $0.0097 \pm 0.0005$ | $0.0390 \pm 0.0040$ |
| | QUAD | $0.0098 \pm 0.0003$ | $0.0392 \pm 0.0063$ |
| | CONTACT | $\mathbf{0.0095 \pm 0.0004}$ | $\mathbf{0.0385 \pm 0.0050}$ |

Table 3: **FAUST dense correspondence on test subjects.** We report mean $\pm$ standard deviation over three seeds. Geodesic error is normalized by the geodesic diameter.

In particular, we see that the effect is localized to the contact stratum. With PCA normals, against Euclidean grouping, the contact-stratum error drops from 0.0828 to 0.0568.

Interestingly, for both experiments (segmentation and correspondence), we see a trend supporting the use of the invariants from equation 5 as a strong geometric prior for learning. In particular, the strong performance of LIN supports that the linearity in $|\zeta|$ matters.

### 5.3 HEAT DIFFUSION ON ORIENTED POINT CLOUDS

Beyond neighborhoods, the score $\ell_w$ defines a similarity kernel we can use to simulate diffusions with the algorithm proposed in Kessler et al. (2026): it builds diffusion operator from **similarity kernels**. We represent an oriented point cloud by the discrete measure $\mu_N = \sum_i m_i \delta_{P_i}$, where $P_i = (x_i, n_i)$ and $m_i > 0$ is an area weight. A temperature signal $u_i$ carries total heat $\sum_i m_i u_i$. We then define a **contact similarity kernel** as $K_{ij} = \exp[-\ell_w(P_i, P_j)^2/(2\sigma^2)]$ using the contact score.

**From similarities to diffusion.** The symmetric Sinkhorn algorithm of Kessler et al. (2026) repeatedly adjust one positive scale factor per sample until each row of the averaging operator sums to one. Applying the same factors at both ends of each interaction preserves symmetry with respect to the area weights. With $M = \mathrm{diag}(m_i)$, the updates and resulting operator are

$$\lambda_i \leftarrow \sqrt{\frac{\lambda_i}{\sum_j K_{ij} m_j \lambda_j}}, \qquad Q = \Lambda K \Lambda M, \qquad \Lambda = \mathrm{diag}(\lambda_i). \tag{11}$$

At convergence, $Q\mathbf{1} = \mathbf{1}$ and $MQ = Q^T M$: averaging preserves both a uniform temperature and the integrated heat. The CONTACT kernel does not need to be positive semi-definite, so its balanced operator can have negative eigenvalues. We therefore construct continuous heat flow as

$$A = \frac{Q - \mathrm{diag}(Q\mathbf{1})}{s}, \qquad u(t) = \exp(tA)u(0), \qquad s = \frac{\sigma^2}{2}. \tag{12}$$

Non-negative off-diagonal rates and symmetry ensure positive temperatures and heat conservation: at exact convergence, $A = (Q - I)/s$. This continuous-time extension is distinct from the native Gaussian baseline $u_G(ks) = Q_G^k u(0)$, obtained by replacing $\ell_w$ with Euclidean distance. Here $s$ is the nominal Gaussian time scale, without fitting diffusion speeds. The comparison below evaluates the complete diffusion constructions used for the CONTACT and Gaussian kernels.

**Validation.** We compare with the point-cloud Laplacian of Sharp & Crane (2020) and ordinary cotangent mesh heat as reference, each evolved as $\exp(-tM^{-1}L)u(0)$. We evaluate on 3 shapes: the Stanford bunny and dragon and a set of 3 stacked sheet in which two are co-oriented and the last one has opposite normals. We measure the normalized discrepancy as follows: let $u_i(t)$ be the heat

field computed by a given method and $u_i^{mesh}(t)$ the cotangent mesh reference, $m_i$ is still the lumped mesh vertex mass. The discrepancy $D$ is then

$$D(t) = \frac{1}{2} \sum_i \left| \frac{m_i u_i(t)}{\sum_j m_j u_j(t)} - \frac{m_i u_i^{mesh}(t)}{\sum_j m_j u_j^{mesh}(t)} \right| \tag{13}$$

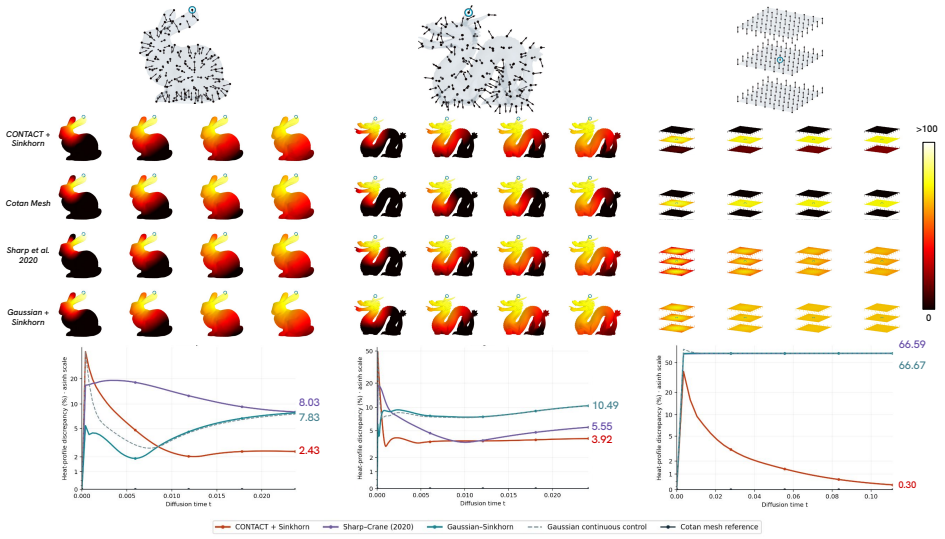



Figure 5: **Heat diffusion comparison.** We compare the computed heat diffusion process against a reference mesh diffusion (from the cotangent Laplacian). On the graphs is reported the normalized heat-profile discrepancy $D(t)$ (in %) which is the total variation measured with common mesh areas. Reported numbers are the final discrepancy. Lower means closer to the mesh diffusion.

Results are reported in Figure 5. In these three examples, the CONTACT construction substantially reduces leakage and diffusion shortcuts across nearby disconnected surfaces and achieves the best final normalized heat-profile agreement, while exhibiting a pronounced short-time mismatch with cotangent mesh diffusion. An additional example and a wave propagation simulation with a similar approach are given in the appendix.

## 6 Discussion, Limitations, and Conclusion

**Limitations.** While we demonstrated several advantage of using a sub-Riemannian contact geometry to describe locality in oriented point clouds, we still mention limitations. First, normals are an input and the approach is sensible to noise as additional results with synthetic data reported in Appendix C and its effect should be studied in depth in future work. Similarly, the weight $w$ is fixed but a locally adapted or learned computation is a natural extension. Also if the sampling is not dense enough, we drift from the "local approximation" regime where the Heisenberg distance can be used to approximate the true Carnot-Carathéodory distance.

**Conclusion and future work.** We introduced a metric structure on oriented point clouds to reduce shortcuts and leakage across contact when connecting neighbors. We showcased its use as a drop-in replacement of local aggregation operations in both deep learning framework and PDE computation (heat diffusion). This resulted in better handling of contact-prone structures without relying on a clean connectivity matrix. From this, many directions are worth pursuing, either on the theoretical side with potential extensions of the metric to textured shapes, curvature driven contact similar to Petrov et al. (2026) but with stronger anisotropy. Or, it could be used in applications involving normals such as refining mesh reconstruction algorithms from Gaussian splatting using normal maps (Gomez et al., 2026) or defining hypoelliptic heat-kernel signatures as pose-invariant descriptors generalizing Sun et al. (2009) as a purely geometric feature descriptor on oriented point clouds. Either way, we hope this work will open ways to improve geometry processing and learning algorithms on oriented sets in general.

## REFERENCES


Erik J Bekkers. B-spline CNNs on Lie groups. In *International Conference on Learning Representations (ICLR)*, 2020.

Erik J Bekkers, Remco Duits, Alexey Mashtakov, and Gonzalo R Sanguinetti. A PDE approach to data-driven sub-Riemannian geodesics in $SE(2)$. *SIAM Journal on Imaging Sciences*, 8(4): 2740–2770, 2015.

Erik J. Bekkers, Da Chen, and Jorg M. Portegies. Nilpotent approximations of sub-riemannian distances for fast perceptual grouping of blood vessels in 2d and 3d. *J. Math. Imaging Vis.*, 60 (6):882–899, July 2018a. ISSN 0924-9907. doi: 10.1007/s10851-018-0787-z. URL https://doi.org/10.1007/s10851-018-0787-z.

Erik J Bekkers, Maxime W Lafarge, Mitko Veta, Koen AJ Eppenhof, Josien PW Pluim, and Remco Duits. Roto-translation covariant convolutional networks for medical image analysis. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2018b.

Erik J Bekkers, Sharvaree Vadgama, Rob Hesselink, Putri A van der Linden, and David W Romero. Fast, expressive SE$(n)$ equivariant networks through weight-sharing in position-orientation space. In *International Conference on Learning Representations (ICLR)*, 2024.

Gijs Bellaard and Bart M. N. Smets. Roto-translation invariant metrics on position-orientation space. In *International Conference on Geometric Science of Information*, 2025. URL https://api.semanticscholar.org/CorpusID:277596215.

Gijs Bellaard, Daan LJ Bon, Gautam Pai, Bart MN Smets, and Remco Duits. Analysis of (sub-)Riemannian PDE-G-CNNs. *Journal of Mathematical Imaging and Vision*, 65(6):819–843, 2023.

Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, June 2014. IEEE.

Ugo Boscain, Roman A Chertovskih, Jean-Paul Gauthier, and Alexey O Remizov. Hypoelliptic diffusion and human vision: A semidiscrete new twist. *SIAM Journal on Imaging Sciences*, 7(2): 669–695, 2014.

Blanche Buet, Gian Leonardi, and Simon Masnou. A varifold approach to surface approximation. *Archive for Rational Mechanics and Analysis*, 226, 11 2017. doi: 10.1007/s00205-017-1141-0.

Nicolas Charon and Alain Trouvé. The varifold representation of nonoriented shapes for diffeomorphic registration. *SIAM Journal on Imaging Sciences*, 6(4):2547–2580, 2013. doi: 10.1137/130918885. URL https://doi.org/10.1137/130918885.

Giovanna Citti and Alessandro Sarti. A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision*, 24(3):307–326, 2006.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International Conference on Machine Learning (ICML)*, 2016.

Keenan Crane, Clarisse Weischedel, and Max Wardetzky. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics*, 32(5):1–11, 2013.

Remco Duits, Arpan Ghosh, Tom Dela Haije, and Alexey Mashtakov. Sub-Riemannian geodesics in $SE(3)$ whose spatial projections do not have cusps. *Journal of Dynamical and Control Systems*, 22(4):771–805, 2016.

Remco Duits, Stephan PL Meesters, Jean-Marie Mirebeau, and Jorg M Portegies. Optimal paths for variants of the 2D and 3D Reeds-Shepp car with applications in image analysis. *Journal of Mathematical Imaging and Vision*, 60(6):816–848, 2018.

Shachar Fleishman, Iddo Drori, and Daniel Cohen-Or. Bilateral mesh denoising. In *ACM SIGGRAPH*, pp. 950–953, 2003.

Erik Franken and Remco Duits. Crossing-preserving coherence-enhancing diffusion on invertible orientation scores. *International Journal of Computer Vision*, 85(3):253–278, 2009.

Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. SE(3)-transformers: 3D roto-translation equivariant attention networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Hansjörg Geiges. *An Introduction to Contact Topology*, volume 109 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2008. ISBN 9780521865852.

Diego Gomez, Antoine Guédon, Nissim Maruani, Bingchen Gong, and Maks Ovsjanikov. From blobs to spokes: High-fidelity surface reconstruction via oriented gaussians, 2026.

Gaël Guennebaud and Markus Gross. Algebraic point set surfaces. *ACM Trans. Graph.*, 26(3):23–es, July 2007. ISSN 0730-0301. doi: 10.1145/1276377.1276406. URL https://doi.org/10.1145/1276377.1276406.

Paul Guerrero, Yanir Kleiman, Maks Ovsjanikov, and Niloy J Mitra. PCPNet: Learning local shape properties from raw point clouds. In *Computer Graphics Forum (Proc. Eurographics)*, volume 37, pp. 75–85, 2018.

Irène Kaltenmark, Benjamin Charlier, and Nicolas Charon. A general framework for curve and surface comparison and registration with oriented varifolds. 07 2017. doi: 10.1109/CVPR.2017.487.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023. URL https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/.

Nathan Kessler, Robin Magnet, and Jean Feydy. Sinkhorn normalization of diffusion kernels. In *International Conference on Machine Learning (ICML)*, 2026.

Jean-Marie Mirebeau. Fast-marching methods for curvature penalized shortest paths. *Journal of Mathematical Imaging and Vision*, 60(6):784–815, 2018.

John Leslie Mitchell. On carnot-carathéodory metrics. *Journal of Differential Geometry*, 21:35–45, 1985. URL https://api.semanticscholar.org/CorpusID:120608865.

Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

Jean Petitot. The neurogeometry of pinwheels as a sub-Riemannian contact structure. *Journal of Physiology-Paris*, 97(2-3):265–309, 2003.

Alice Petrov, Mohammad Sina Nabizadeh, Ana Dodik, and Justin Solomon. Tangent blow-ups for processing non-manifold geometry. *Computer Graphics Forum*, 2026. doi: https://doi.org/10.1111/cgf.70527. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70527.

Jorg M. Portegies, Gonzalo Sanguinetti, Stephan P. L. Meesters, and Remco Duits. New approximation of a scale space kernel on se(3) and applications in neuroimaging. In *Scale Space and Variational Methods in Computer Vision*, 2015. URL https://api.semanticscholar.org/CorpusID:7968693.

Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the ICP algorithm. In *Third International Conference on 3D Digital Imaging and Modeling (3DIM)*, pp. 145–152. IEEE, 2001.

David Schunck, Federico Magistri, Radu Alexandru Rosu, André Cornelißen, Nived Chebrolu, Stefan Paulus, Jens Léon, Sven Behnke, Cyrill Stachniss, Heiner Kuhlmann, and Lasse Klingbeil. Pheno4d: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis. *PLOS ONE*, 16(8):e0256340, 2021. doi: 10.1371/journal.pone.0256340.

Nicholas Sharp and Keenan Crane. A laplacian for nonmanifold triangle meshes. *Computer Graphics Forum*, 39, 2020. URL https://api.semanticscholar.org/CorpusID:221607415.

Bart MN Smets, Jim Portegies, Erik J Bekkers, and Remco Duits. PDE-based group equivariant convolutional neural networks. *Journal of Mathematical Imaging and Vision*, 65(1):209–239, 2023.

Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer Graphics Forum (Proc. SGP)*, volume 28, pp. 1383–1392, 2009.

Gabriel Taubin. A signal processing approach to fair surface design. *ACM SIGGRAPH*, pp. 351–358, 1995.

Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6410–6419, 2019.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics*, 38(5):1–12, 2019.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler, faster, stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4840–4851, 2024.

A. C. Öztireli, G. Guennebaud, and M. Gross. Feature preserving point set surfaces based on non-linear kernel regression. *Computer Graphics Forum*, 28(2):493–501, 2009. doi: https://doi.org/10.1111/j.1467-8659.2009.01388.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2009.01388.x.