
Point Cloud Synthesis Using Inner Product Transforms

Ernst Röell^{1,2,3} Bastian Rieck^{1,2}

¹AIDOS Lab, University of Fribourg, Switzerland

²Institute of AI for Health, Helmholtz Munich, Germany

³Technical University of Munich, Germany

Abstract

Point cloud synthesis, i.e. the generation of novel point clouds from an input distribution, remains a challenging task, for which numerous complex machine learning models have been devised. We develop a novel method that encodes geometrical-topological characteristics of point clouds using inner products, leading to a highly-efficient point cloud representation with provable expressivity properties. Integrated into deep learning models, our encoding exhibits high quality in typical tasks like reconstruction, generation, and interpolation, with inference times orders of magnitude faster than existing methods.

1 Introduction

Point clouds are a data modality of crucial relevance for numerous domains. While computer graphics is the predominant application area, where point clouds are often used as precursor to more structured representations like meshes, they also occur in higher dimensions in the form of sensor data, for instance. However, the synthesis of hitherto-unseen point clouds from a given distribution still proves to be a challenging task, with numerous models aiming to address it [39]. The complexity arises because of the sparsity and ‘set-like’ structure of point clouds, making it hard to generalise existing machine learning architectures directly. State-of-the-art methods thus typically require large amounts of compute, exhibiting long training and inference times. With recent work [33] demonstrating that a change of perspective—like exchanging raw coordinates for Fourier-based features—can make comparatively simple deep-learning architectures competitive in computer-vision tasks, our paper explores the question to what extent novel representations of point clouds can lead to gains in computational performance *without* sacrificing too much quality.

To obtain such representations, we build on a multi-scale geometrical-topological descriptor [35] based on *inner products* of coordinates representing a high-dimensional shape. Using suitable approximations, this descriptor permits us to represent a point cloud as a *single* 2D image. Unlike other single-view representations, however, the coordinates of the image represent a different ‘domain,’ namely geometrical-topological aspects. In fact, this mapping is computationally efficient and *injective*, making it theoretically possible to reconstruct a point cloud from its descriptor (we will demonstrate that this property also holds empirically when working with approximations). We refer to this mapping as the *Inner Product Transform* (IPT) and use it for the generation of 3D point clouds.¹ Our work primarily focuses on building highly-efficient models that enable real-time inference in settings of limited computational resources while also being trainable on commodity hardware.

¹For readers familiar with computational topology, the IPT is a special case of the *Euler Characteristic Transform* (ECT). Our focus on point clouds permits us to formulate the IPT without any background knowledge in computational topology, while also resulting in substantially simplified proofs and more powerful statements about its theoretical properties.

Our paper is built on *two core notions*, the first one being that we treat point cloud generation as a two-step task, with the first step being an image generation task, yielding a two-dimensional descriptor, followed by a (multi-modal) image-to-point-cloud *reconstruction* task. We realise the two steps using separate machine learning models, thus substantially decreasing architectural complexity—in effect, we represent a point cloud as a special image that can be *generated*, providing a bridge to the point cloud domain. The second core notion is that the generated image is a faithful, i.e. *injective*, representation of the point cloud, indicating that it is possible to perfectly reconstruct a point cloud from its descriptor. Our paper makes the following **contributions**:

1. We present a novel generative image-to-point-cloud pipeline based on inner products that allows training and inference times to be *orders of magnitude faster* while retaining high quality generation.
2. We show that our representation yields a stable latent space, which permits (i) high-quality interpolation tasks, and (ii) solving different out-of-distribution tasks *without* the need for retraining while still maintaining high quality.
3. We demonstrate injectivity and other advantageous properties of our descriptor, allowing the generalisation of our method to point clouds of arbitrary dimensions.

Our code is available at <https://github.com/aidos-lab/inner-product-transforms> and is released under a BSD-3-Clause license.

2 Related Work

Point clouds being a nigh-ubiquitous data modality, numerous models already exist to tackle classification or generation tasks [39]. The lack of structure, as well as the requirement of permutation invariance, imposes constraints on the underlying computational architecture, typically substantially increasing model complexity [29, 41]. To solve generation tasks, many methods opt for *jointly* learning the generation of the shape, i.e. the surface or object the points are sampled from, as well as the mapping of points onto that object. This core idea drives several recent state-of-the-art models, including Point-Voxel CNN [21], PointFlow [40], SoftFlow [17], Point Voxel Diffusion [44, PVD] and LION [42]. While these models exhibit high-quality results, their architectures require long training and inference times.

By contrast, our approach only needs to model a distribution of geometrical-topological descriptors, represented as 2D images, from which we subsequently *reconstruct* a point cloud again. While reconstructing point clouds from images is an active field of research (see Fahim et al. [11] for a survey on single-view reconstruction), such images are typically depth images or snapshots taken from a specific position around the object. As such, they are not necessarily yielding a faithful, unique representation of an object. Another close analogue to our method is given by ‘structure-from-motion’ approaches [27], which reconstruct complex geometries based on *sets* of images, taken from different spatial viewpoints. For our method, however, the viewpoints are represented as unit vectors on a sphere, which are used to ‘probe’ the point cloud from a specific direction.

Hence, a crucial property of our method, the *Inner Product Transform* (IPT) is that it theoretically yields a faithful representation, hence permitting an injective mapping between the image domain and the shapes we aim to reconstruct (or generate). We observe this property to hold empirically when working with discrete approximations. This is due to the fact that the IPT is a special case of a general geometrical-topological descriptor, the *Euler Characteristic Transform* (ECT), which studies shapes at multiple scales and from multiple directions, providing a unique characterisation [35]. Being a stable [10, 13] and efficient descriptor, the ECT is often used to solve questions in data science, mostly in the form of ‘hand-crafted’ features for classification and regression tasks [4, 8, 22, 24, 25, 34, 37]. Recent work addressed this shortcoming and enabled the use of the ECT in machine-learning applications in the form of a generic differentiable *computational layer* [31, 32] or a *positional encoding* [3]. Notably, the ECT remains a unique characterisation even when using a finite number of directions [9, 14], meaning that, theoretically, it can be inverted to reconstruct the input data. Practically, however, inversion is presently possible only for select data modalities like planar graphs [12]. **While restricted to point clouds, our method is thus the first to enable the inversion of such a descriptor for input data of arbitrary dimensionality, making it possible to use it in the context of generative models.**

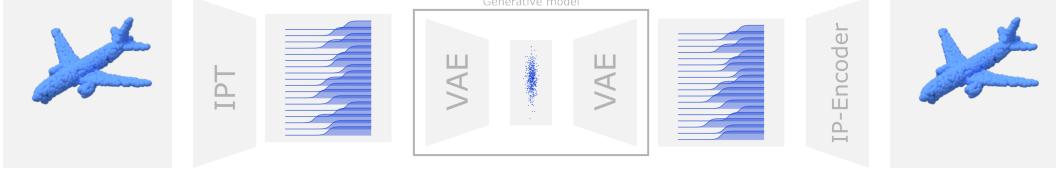


Figure 1: Given a point cloud on the left, we compute its *Inner Product Transform* (IPT). For generative tasks, we train a generative model (middle) to reconstruct and generate the distribution of IPTs. The (possibly-generated) IPT is then passed through the encoder model to obtain the reconstructed (or novel) point cloud. Our pipeline is decoupled, permitting *any* generative image model to be used to generate point clouds.

3 Methods

Our method, the *Inner Product Transform* (IPT), can be intuitively understood as a filtering process of point clouds, employing different sets of hyperplanes, created by a set of *directions* (i.e. normal vectors). The inner products of point cloud coordinates with a direction vector are then used to parametrise a *curve* that counts the number of points below the hyperplane. By stacking curves, we obtain a 2D image representation of the point cloud. The IPT is a special case of the *Euler Characteristic Transform* [35, ECT], but our subsequent description is self-contained, does not require any knowledge of topology, and presents simplified proofs of all properties.

3.1 Inner Product Transforms

Let S^{n-1} denote the unit sphere in \mathbb{R}^n . Given a point cloud $X \subset \mathbb{R}^n$, a fixed direction vector $\xi \in S^{n-1}$, and a *height* $h \in \mathbb{R}$, we define the set $X_{\xi, h} := \{x \in X \mid \langle x, \xi \rangle \leq h\}$, where $\langle x, \xi \rangle$ is the Euclidean dot product. The set $X_{\xi, h}$ contains all points below the hyperplane spanned by $\langle x, \xi \rangle = h$, and we denote its cardinality by $\chi(X_{\xi, h})$.² We then define the *Inner Product Transform* as

$$\begin{aligned} \text{IPT}(X) : S^{n-1} \times \mathbb{R} &\rightarrow \mathbb{N} \\ (\xi, h) &\mapsto \chi(X_{\xi, h}). \end{aligned} \tag{1}$$

A point $x \in X$ is included in $X_{\xi, h}$, thus affecting $\chi(X_{\xi, h})$, if and only if its height $h_x := \langle x, \xi \rangle$ along ξ is less than h . We can thus formulate the contribution of a point x to the IPT along each direction in terms of an *indicator function*:

$$\mathbb{1}_x(\xi, h) := \begin{cases} 1 & \text{if } \langle \xi, x \rangle \leq h \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

This enables us to rewrite Eq. (1) as

$$\begin{aligned} \text{IPT}(X) : S^{n-1} \times \mathbb{R} &\rightarrow \mathbb{N} \\ (\xi, h) &\mapsto \sum_{x \in X} \mathbb{1}_x(\xi, h). \end{aligned} \tag{3}$$

Following ideas from Röell & Rieck [32], we can replace all indicator functions with sigmoid functions, i.e.

$$\begin{aligned} \widehat{\text{IPT}}(X) : S^{n-1} \times \mathbb{R} &\rightarrow \mathbb{R} \\ (\xi, h) &\mapsto \sum_{x \in X} S(\lambda(\langle \xi, x \rangle - h)), \end{aligned} \tag{4}$$

where λ denotes a scale parameter, which controls how closely the sigmoid function approximates the indicator function (see Appendix D for additional ablations demonstrating the stability of this approach). In practice, we sample n_d directions and discretise all heights with n_h steps, thus representing $\widehat{\text{IPT}}$ as an *image* of resolution $n_h \times n_d$.

²For readers familiar with topology, this notation is an allusion to the *Euler Characteristic*.

Properties. The IPT and its approximation $\widehat{\text{IPT}}$ from Eq. (4) satisfy several important properties. We first focus on the case of the IPT and note that it is *injective*, i.e. a point cloud can be perfectly reconstructed from its descriptor. **Unlike existing injectivity results [14, 35], which focus on geometric simplicial complexes, we prove that for point clouds in \mathbb{R}^n , it is sufficient to use $n + 1$ affinely-independent directions for the reconstruction.**

Theorem 1. *Given two point clouds X, Y with $X \neq Y$, we have $\text{IPT}(X) \neq \text{IPT}(Y)$.*

An important consequence of injectivity is that it allows us to formulate a metric for the set of point clouds contained in a ball of fixed radius.

Lemma 1. *Let $B_n(R)$ denote a ball of radius R in \mathbb{R}^n . For two point clouds $X, Y \subset B_n(R)$, define their distance as*

$$d(X, Y) := \frac{1}{|\Xi|} \sum_{\xi \in \Xi} \|\text{IPT}(X)|_{\xi} - \text{IPT}(Y)|_{\xi}\|_2, \quad (5)$$

where Ξ is a finite set of directions and $\|\cdot\|_2$ is the L^2 -norm restricted to the interval $[-R, R]$. The function $d(\cdot, \cdot)$ satisfies the definition of a metric.

When approximating the IPT via Eq. (4) and a finite number of directions in practice, the metric defined in Lemma 1 corresponds to the pixel-wise mean squared error between discretised IPTs. Hence, given a sufficient number directions, we can formulate a *loss function* based on this metric, which admits a highly-efficient implementation and will turn out to lead to high-quality results. We also obtain a result that permits us to calculate IPTs from disjoint unions of point clouds.

Lemma 2. *Let $X, Y \subset \mathbb{R}^n$ be disjoint point clouds, then*

$$\text{IPT}(X \cup Y) = \text{IPT}(X) + \text{IPT}(Y). \quad (6)$$

Finally, we can prove that the IPT is *surjective* on convex linear combinations.

Theorem 2. *Given two point clouds $X, Y \subset \mathbb{R}^n$, the IPT is surjective for the rational linear subspace spanned by $\text{IPT}(X)$ and $\text{IPT}(Y)$, up to a rational scaling factor. In particular, for $p, q \in \mathbb{N}_0$ with $0 \leq p \leq q$ and $q > 0$ we have*

$$\frac{p}{q} \text{IPT}(X) + \frac{q-p}{q} \text{IPT}(Y) = \frac{1}{q} \text{IPT}(Z), \quad (7)$$

where $Z = \cup_p X \cup_{q-p} Y$.

A direct consequence of Theorem 2 is that, along the *linear interpolation* between two IPTs, there are only valid IPTs, i.e. each interpolation step affords a perfect reconstruction in theory. We analyse this aspect further in Section 4.5. **We note that most of these properties (except injectivity) only hold for the IPT but, as we will later demonstrate, $\widehat{\text{IPT}}$, the approximation to the IPT, retains these properties in practice.**

3.2 The IP-Encoder

The existence of Theorem 1 unfortunately does not lead to a practical algorithm for ‘inverting’ an IPT. We therefore suggest an approach based on neural networks and describe the IP-Encoder, which encodes an IPT to a point cloud. Subsequently, this will enable us to learn *distributions* of IPTs and reconstruct new point clouds in a generative setting and, as we shall see, even in the setting of out-of-distribution data. Since our representation is permutation-invariant, our IP-Encoder model directly inherits this invariance, which substantially reduces the complexity of its architecture.

Model architecture. Given the structure of the IPT as an image, CNN architectures provide suitable base models. However, since the direction vectors $\xi \in S^{n-1}$ cannot be consistently ordered along one dimension, we need to reframe the input data. Specifically, we consider an IPT, normalised to $[-1, 1]$, as a *multi-channel, one-dimensional signal*. Our IP-Encoder model then consists of multiple 1D convolutional layers followed by fully-connected layers, resulting in a conceptually simple and efficient architecture. To *generate* a novel IPT, we can use any generative model for images. In our experiments, we will use an architecture based on a convolutional variational autoencoder, denoted by IP-VAE (see Appendix F for more architectural details).

The IPT as an optimisation problem. As an alternative to the machine learning model above, we also investigate whether the inversion of an IPT can be turned into an *optimisation problem*. Due to the differentiability of our approximation $\widehat{\text{IPT}}$ and given a known IPT, minimising a loss function, such as the one described in [Lemma 1](#), using backpropagation should result in a suitable approximation of the unknown point cloud one wishes to recover. However, while this works in *theory*, it requires the input and target to be a ‘correct’ $\widehat{\text{IPT}}$ of a point cloud. In practice, this is *not* guaranteed: Generative models may sample from their latent spaces and output samples that are *close* to being an $\widehat{\text{IPT}}$ of a point cloud without actually satisfying all structural constraints. To some extent, we argue that this is even the desired behaviour since a generative model would otherwise just rehash its inputs [\[16\]](#). Hence, any optimisation-based method will not necessarily result in realistic reconstructions in a generative setting. We nevertheless investigate such an optimisation procedure in an ablation study in the appendix ([Figure S.3](#)).

Topological loss functions. The fact that [Lemma 1](#) shows the IPT to be a metric naturally poses the question if its approximation $\widehat{\text{IPT}}$ can be used as an effective and efficient loss term. While in principle the $\widehat{\text{IPT}}$ is sufficient as a loss term for training a point cloud reconstruction model, training time can be reduced through a combination of a low-resolution (64×64) $\widehat{\text{IPT}}$ combined with the Chamfer Distance (CD). The CD is a fast-to-compute (pseudo-)metric often used in point cloud evaluation (see [Appendix E](#) for more details). The $\widehat{\text{IPT}}$ loss ensures that the overall geometry and point cloud density of the object is captured, whereas the CD loss ensures that fine-grained details are accounted for. Building a joint loss, combining the $\widehat{\text{IPT}}$ and CD, thus results in a density-aware loss term that takes the global density of a point cloud into account.

Latent space. Interpolation between samples provides valuable insight into the capacities of our generative model. [Theorem 2](#) provides an intuition for interpolation using $\widehat{\text{IPT}}$ s. To this end, suppose we perform linear interpolation between point clouds X, Y , using a parameter p over the interval $[0, 1]$, which is partitioned in q equidistant steps. Then, at step p/q , the point cloud $Z = \cup_p X \cup_{q-p} Y$ is the union of p copies of X and $q - p$ copies of Y . During the interpolation, the number of copies of X is increased and the number of copies of Y is reduced. We may interpret this intuitively as transporting ‘mass’ from Y to X over the course of q steps. Our IP-Encoder model averages all copies during reconstruction in a natural fashion, resulting in smooth transitions between point clouds.

4 Experiments

We demonstrate the effectiveness, efficiency, and overall utility of our IP-Encoder through a comprehensive suite of experiments. **Subsequently, to simplify the notation, we will only refer to the IPT, with the understanding that we are calculating an approximation of it according to Eq. (4).** Throughout our experiments, the emphasis and motivating questions regard the expressivity of the IPT and its capacity to faithfully and *effectively* represent point clouds in a practical setting. Good results with a minimal architecture underpin the fact the IPT is an effective representation as it shows that the data distribution is easy to learn. Among other things, we show that our IP-Encoder model can (i) effectively reconstruct shapes, (ii) create novel point clouds from *generated* IPTs, and (iii) effectively downsample point clouds.

All our experiments rely on a subset of the ShapeNet dataset, and we adopt the preprocessing and evaluation workflow introduced by Yang et al. [\[40\]](#). Each point cloud in the dataset consists of 2048 points sampled on the surface of three shape classes (airplanes, chairs and cars). We report the *minimum matching distance* [\[2\]](#), MMD based on the Chamfer Distance (MMD-CD) or the *Earth Mover’s Distance* (MMD-EMD) between reconstructed point clouds (see [Appendix E](#) for a brief description of these evaluation metrics). Unless otherwise mentioned, the reported Chamfer Distance is scaled by 1.00×10^4 and the Earth Mover’s Distance is scaled by 1.00×10^3 . Our hardware consists of an NVIDIA GeForce RTX 4070 with 12GB VRAM and an 13th Gen Intel(R) Core(TM) i7-13700K with 32GB RAM. We compare our methods to several state-of-the-art models, namely (i) PointFlow [\[40\]](#), (ii) SoftFlow [\[17\]](#), (iii) ShapeGF [\[6\]](#), (iv) Canonical VAE [\[7\]](#) and (v) LION [\[42\]](#).

Table 1: Reconstruction results on three ShapeNet classes. The best-performing model is highlighted in **bold**, while the second-best is shown in *italics*. Our method consistently ranks second in terms of the EMD, which is recognised as the best metric for reconstruction quality [44].

MODEL	Airplane		Chair		Car	
	CD (\downarrow)	EMD (\downarrow)	CD (\downarrow)	EMD (\downarrow)	CD (\downarrow)	EMD (\downarrow)
PointFlow	1.30 ± 0.00	5.36 ± 0.06	10.43 ± 0.02	17.54 ± 0.16	6.94 ± 0.01	12.93 ± 0.19
SoftFlow	1.19 ± 0.00	4.28 ± 0.06	11.05 ± 0.03	17.68 ± 0.08	6.82 ± 0.01	11.44 ± 0.10
ShapeGF	1.05 ± 0.00	4.42 ± 0.04	5.96 ± 0.07	12.23 ± 0.11	5.68 ± 0.01	9.26 ± 0.18
Canonical VAE	0.98 ± 0.00	3.19 ± 0.03	6.56 ± 0.02	8.60 ± 0.07	5.44 ± 0.01	6.13 ± 0.02
LION	0.30 ± 0.00	0.12 ± 0.00	0.70 ± 0.00	0.14 ± 0.00	0.60 ± 0.00	0.09 ± 0.00
IP-Encoder (ours)	<i>1.05 ± 0.00</i>	<i>1.57 ± 0.00</i>	9.24 ± 0.00	6.19 ± 0.00	5.82 ± 0.00	3.18 ± 0.00

4.1 Reconstructing Point Clouds

Our first set of experiments assesses (i) the reconstruction quality of the IP-Encoder, (ii) the efficacy of the IPT as a loss, and (iii) its computational efficiency.

Architecture and experimental setup. We consider the IPT as a 1D signal, with each direction corresponding to a channel, sample 128 directions uniformly from the unit sphere, and discretise each direction into 128 steps, thus obtaining an IPT with a resolution of 128×128 . Sampling the directions *randomly* is motivated by the lack of canonical ordering of unit vectors in three and higher dimensions (notice that in two dimensions, the angle affords a parametrisation that results in a natural ordering). While ‘pseudo-ordered’ directions, for instance via a spiral along the unit sphere, *might* potentially lead to better results, they come at the cost of generality, and we thus refrained from doing so in our experiments. To improve this for future work, one could use a *positional encoding* of the directions; at present, such an ordering is only implicitly present through the fixed ordering of columns in the image we use to represent the IPT. Given such a 2D representation of an IPT, our IP-Encoder consists of four 1D convolutional layers with batch normalisation, max-pooling, and SiLU activation functions (cf. Appendix F). After mapping the IPT into a latent space, we apply a final 3-layer MLP to predict the final point cloud. We use ReLU activation functions for the first two layers and a tanh activation function for the last layer, since tanh is better suitable for bounded outputs. Subsequently, we train the IP-Encoder *separately* for each of the classes for $5k$ epochs, using a CD + IPT-64 loss, denoting a weighted sum of the CD and the IPT with a resolution of 64×64 .

Table 2: Ablation study wrt. the loss function. Our joint loss, combining CD and IPT, yields *balanced* results without overfitting any of the two metrics.

LOSS FUNCTION	Airplane (\downarrow)		Chair (\downarrow)		Car (\downarrow)	
	CD	EMD	CD	EMD	CD	EMD
CD	1.00	8.89	10.44	32.43	5.97	14.96
IPT-64	2.41	1.09	13.06	4.29	7.75	2.47
CD + IPT-64	1.03	1.46	9.52	8.44	6.12	4.16

an ablation study concerning the loss term, showing that the *combination* of CD and IPT is crucial for high reconstruction quality. As Table 2 shows, a joint loss yields the best quality. In line with prior work, CD on its own tends to adversely affect point-cloud density [2], resulting in larger (i.e. *worse*) EMD scores. Conversely, training only with a ECT loss shows that the reconstructed follow the density of the underlying object, but may become blurry. This is partially due to the low resolution of 64, making the combination of CD (for details) and the ECT (for global density and structure) crucial. While it would also possible to use the EMD as a loss term, we find that our IPT-64 loss function is *substantially faster* than the EMD (0.0006 s versus 0.0291 s), making it the preferred loss function.

We observe the IPT to be an effective representation of point clouds, leading to a conceptually simple model with exceptionally fast training times.

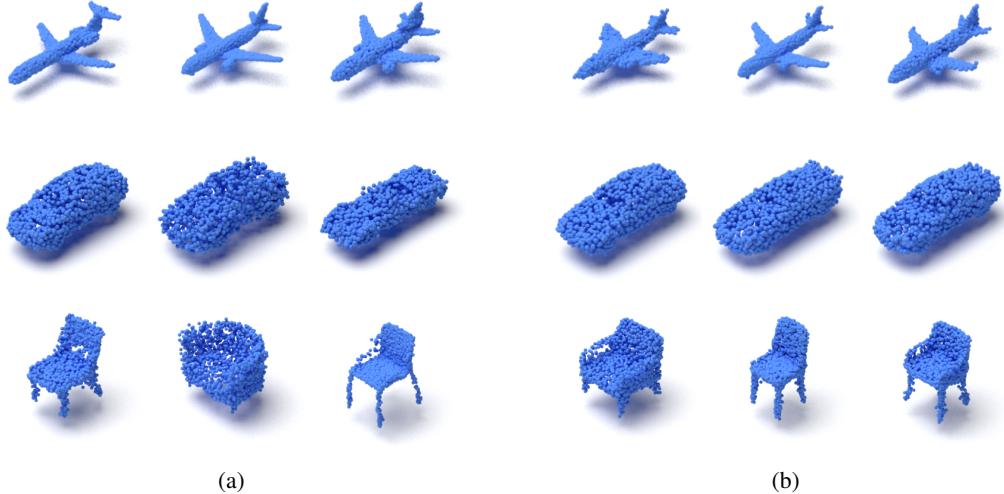


Figure 2: Examples of *reconstructed* (a) and *generated* (b) point clouds using our IP-Encoder model for three classes in the ShapeNet dataset.

4.2 Generating Novel Point Clouds

We now view the IPT as a latent space, from which we can *sample* and subsequently *reconstruct* point clouds, demonstrating that (i) the IPT constitutes a stable latent space, (ii) a *distribution* of shapes can be learned through the IPT, and (iii) the IP-Encoder works well in out-of-sample settings.

Architecture and experimental setup. We use IP-VAE, a CNN-style VAE [15] as our generative model, thus turning point cloud generation into an image generation task. Our encoder uses four convolutional blocks with Leaky ReLU activation functions, batch normalisation, and a final linear layer for the 64-dimensional latent space embedding. We also add additional models, namely (i) Set-VAE [18], (ii) Point Voxel Diffusion [44, PVD], (iii) Point Straight Flows [38], and (iv) XCube [30]. These models lack reconstruction capabilities or the respective code, so we excluded them from the experiment in Section 4.1. For training the IP-VAE, we follow the β -VAE setup, using KL-divergence and MSE loss terms with $\beta = 1.00 \times 10^{-4}$. We sample latent vectors from the IP-VAE and consider them to be IPTs, which we subsequently map to a point cloud using the IP-Encoder model. Our evaluation of generative performance follows the setup of Yang et al. [40].

Table 3: Inference time T in s for all methods, measured on the same hardware. Our model is orders of magnitude faster than all others.

MODEL	DEVICE	T (s)
PointFlow		0.270
SoftFlow		0.120
ShapeGF	GPU	0.340
SetVAE		0.030
PSF		0.040
PVD		29.900
IP-VAE	GPU	0.001
IP-VAE	CPU	0.006

Results. Table 4 reports numerical results, whereas Figure 2b depicts generated samples from each class. We include different quality metrics and note that, despite its conceptual simplicity, our model consistently ranks among the best two models in terms of MMD-EMD, and exhibits performance on a par with more elaborate architectures. This is noteworthy since the IP-VAE is not constrained to produce *exact* IPTs, hence its outputs may contain artefacts. The fact that our IP-Encoder model (remaining fixed throughout this experiment) can reconstruct point clouds from such out-of-sample data *without* additional regularisation or retraining, shows that the IPT describes a stable and expressive latent space. We observe fast *training times*, with the IP-VAE taking approximately 15.00 min, making the full training pipeline run in less than 1.00 h. This is in stark contrast to models like LION (550.00 h) or Soft-Flow (\approx 144.00 h). Finally Table 3 shows that our *inference times* on commodity hardware are orders of magnitude faster than existing models.

The IPT yields a suitable representation for generating high-quality point clouds, while exhibiting inference times that are orders of magnitude faster than existing models. More expressive generative models are likely to improve generative quality (at the expense of longer runtimes).

Table 4: Evaluation metrics for point cloud generation tasks, cited from their respective papers. We highlight the best result in **bold** and the second-best in *italics*. Our implementation of the EMD follows PVD [44]. Whenever available, we report coverage (COV), minimum matching distance (MMD), and 1-NNA (1-nearest neighbour accuracy) for both EMD and CD.

Model	Airplane						Chair						Car					
	MMD (↓)		COV (↑)		1-NNA (↓)		MMD (↓)		COV (↑)		1-NNA (↓)		MMD (↓)		COV (↑)		1-NNA (↓)	
	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD	EMD
PointFlow	0.22	0.39	<i>47.90</i>	46.41	75.68	69.44	2.41	1.60	<i>42.90</i>	50.00	60.88	59.89	0.90	0.81	46.88	50.00	60.65	62.36
SoftFlow	0.23	0.37	46.91	<i>47.90</i>	70.92	69.44	2.53	1.68	41.39	47.43	59.95	63.51	1.19	0.86	42.90	44.60	62.63	64.71
Shape-GF	2.70	0.66	40.74	40.49	80.00	76.17	2.89	1.70	46.67	48.03	68.96	65.48	9.23	0.76	<i>49.43</i>	50.28	63.20	56.53
SetVAE	0.20	0.37	43.70	48.40	75.31	<i>77.65</i>	2.55	1.59	46.83	44.26	58.76	61.48	0.88	0.73	49.15	46.59	59.66	61.48
PVD	0.22	0.38	48.88	52.09	73.82	64.81	2.62	1.56	49.84	<i>50.60</i>	56.26	53.32	1.08	0.79	41.19	50.56	54.55	53.83
LION	0.22	0.37	47.16	49.63	67.41	61.23	2.64	1.55	48.94	<i>52.11</i>	53.70	52.34	0.91	0.75	50.00	56.53	53.41	51.14
PSF					71.11	<i>61.09</i>					58.92	54.45					57.19	56.07
XCube						52.85	49.75				53.99	48.60					57.96	54.43
IP-VAE (Ours)	0.19	0.34	45.67	45.67	77.28	68.40	2.34	1.48	41.24	48.34	61.33	63.07	0.85	0.70	35.23	49.72	58.52	59.23

4.3 Point Cloud Downsampling

Motivated by the promising results in terms of generative performance, we further investigate the capacity of our IP-Encoder model to upsample a downsampled point cloud.

Table 5: Reconstruction performance for consecutive down- and upsampling, with the best result shown in **bold**. For reference purposes, the last row repeats our results without downsampling.

METHOD	Airplane (↓)			Chair (↓)			Car (↓)		
	CD	EMD	CD	EMD	CD	EMD	CD	EMD	CD
IP-Downsampler	1.18	2.58	11.70	11.42	6.42	5.81			
Uniform	3.10	4.71	15.38	15.34	9.30	9.09			
IP-Encoder	1.03	1.46	9.53	8.45	6.12	4.17			

Architecture and experimental setup. We train a downsampling model with the same architecture as the IP-Encoder, resulting in a point cloud with 256 points. Here, the purpose of the loss term is to minimise the discrepancy of the original point cloud to its downsampled version, measured using the IPT, which we normalise because of its dependency on point cloud cardinality. We then pass the IPT of the downsampled point cloud to the IP-Encoder model to obtain an upsampled version, which, ideally, should be close to the original point cloud. Reporting both CD and EMD scores, we compare these results to a uniform subsampling method.

Results. Table 5 depicts numerical scores (cf. Figure S.2 in the appendix for sample point clouds). Unlike uniform subsampling, our downsampling model preserves uniform density of each object, showing that it has learned the underlying shape. We observe some loss in quality compared to reconstruction without downsampling, which is to be expected. Despite potential compounding errors, the loss in quality remains low. Due to the similar setup, we can compare results with Table 1 and we observe that even when downsampling, we *still* perform on a par with the best model in the reconstruction task. Notably, the IP-Encoder has not been retrained for this task, so the IPT of the subsampled point cloud is out-of-distribution. This indicates that our model has captured the ‘true’ underlying shape characteristics.

Our IP-Encoder can downsample point clouds with a minimal loss of quality, demonstrating that the model is effective in encoding relevant shape properties.

4.4 Practical Inversion of the IPT

While theoretical invertibility (see Theorem 1) ensures that the IPT yields an *expressive* summary, practical algorithms for inversion of an IPT remain an open question. However, in our setting, making use of an approximation to the IPT, i.e. $\widehat{\text{IPT}}$, we may use backpropagation to invert it *without* requiring an IP-Encoder. The core observation is the differentiability of the $\widehat{\text{IPT}}$ with respect to the point cloud coordinates, and the fact that it constitutes a metric (see Lemma 1). The main idea is to randomly initialise a point cloud, compute its $\widehat{\text{IPT}}$, and subsequently calculate the loss between said $\widehat{\text{IPT}}$ and the $\widehat{\text{IPT}}$ of the target point cloud. Differentiability ensures we can compute the derivative with respect to the point cloud coordinates and update them to reduce this loss. Applying backpropagation thus ultimately results in convergence to the original point cloud. We show that

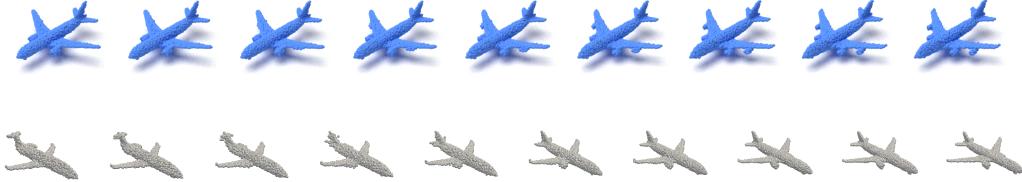


Figure 3: Linear interpolation between IPTs results in a smooth interpolation in the point-cloud domain. Given the IPT of two airplanes, we linearly interpolate between the pixel values and pass each step through the IP-Encoder to obtain a prediction of the intermediary point cloud. Although the IP-Encoder has not been specifically trained on such data, it is able to produce meaningful reconstructions, since it has learned the distribution of shapes. We provide two examples of the interpolation process, one with high pairwise similarity (top), the other with low pairwise similarity (bottom).

(i) IPTs can be inverted in practice, (ii) a resolution of 128 yields high-quality reconstructions, and (iii) the IPT has *stable gradients* (which is not clear a priori, given that we are approximating an indicator function). For this experiment, we apply backpropagation for 2000 epochs with the Adam optimiser using a learning rate of 0.5, which is halved in epochs 50, 100, 200, and 1000, respectively. We repeat all experiments with three resolutions and for three choices of scaling factor λ in Eq. (4), which we set either to be equal, half, or a quarter of selected resolution. As a quality metric, we report the EMD and CD between the optimised and target point cloud.

Table 6: Inversion of the IPT with backpropagation leads to high quality reconstruction.

DATASET	CD (\downarrow)	EMD (\downarrow)
Airplane	0.52 ± 0.00	0.12 ± 0.00
Car	2.04 ± 0.00	0.39 ± 0.00
Chair	1.83 ± 0.00	0.35 ± 0.00

Results. Table 6 and Figure S.3 in the appendix depict the numerical results and show some reconstructed samples, respectively. Comparing the result with Table 1, we observe that backpropagation outperforms the IP-Encoder model. Table S.1 in the appendix provides the full set of results and an ablation with respect to the scale and resolution. We observe that all approximations become gradually coarser as the resolution is decreased. As expected, we observe problems with vanishing gradients for large values of λ , since the sigmoid function approaches an indicator function in this case. This requires balancing λ with respect to the resolution, and we empirically observed that around a quarter of the IPT resolution is sufficient for high-quality results, which is how we pick the parameter in practice. It is important to note that while this method of inverting works for an IPT without any artifacts, it fails to accurately converge if the IPT is only an approximation of a true IPT, as is the case with a *generated* IPT. This also further motivates the use of a *learned* method for reconstruction.

Backpropagation through IPTs directly is possible and outperforms the IP-Encoder in reconstruction quality. However, this optimisation-based scheme is unable deal with noisy or generated IPTs, motivating the need of the IP-Encoder.

4.5 Interpolation in Latent Space

As our final experiment, we consider the linear *interpolation* between the IPT of two point clouds. In the setting of infinite directions, Theorem 2 guarantees that we only encounter valid IPTs along each interpolation. However, in practice, given our use of image-generative models, it is not guaranteed that our input is an IPT in the *strict* sense. Hence, it is crucial to understand the characteristic properties of the IPT in practice since not every sum of sigmoid functions is guaranteed to correspond to a proper shape (in other words, IPTs are not *surjective*). In our discretised representation, interpolating between IPTs can be achieved by interpolating per pixel. To visually assess the quality of the latent space, we reconstruct each intermediary IPT during the interpolation using the IP-Encoder. Figure 3 depicts the resulting point clouds; we observe that the intermediary point clouds still remain plausible reconstructions. Overall, this serves to highlight the utility of the IPT latent space.

The IPT permits linear interpolations, resulting in smooth interpolations between point clouds.

5 Discussion

We propose the *Inner Product Transform* (IPT), which enables us to represent point clouds using 2D images by evaluating inner products. Next to showing that this representation has advantageous theoretical properties, we develop an end-to-end-trainable pipeline for *point-cloud generation*. Notably, we simplify the point-cloud generation process by introducing an intermediate step, the generation of the IPT descriptor, for which one may use *any* image-generation model. Our method is (i) exceptionally fast since it uses only inner products and other simple computational building blocks, (ii) well-grounded in theory (it is a sufficient statistic on the space of point clouds), and, (iii) we observe that its *reconstruction performance* and its *generative capabilities* are, for the most part, on a par with more complex models. This is due to the conceptual simplicity of the model and its provably highly-stable latent space, which facilitates performing out-of-distribution tasks *without* the need for retraining or additional regularisation.

Limitations. An unfortunate limitation arising from the disconnect between theory and practice is that our injectivity results do not guarantee that one can *find* suitable directions, in particular given the complex interplay between *directions* and *resolutions* of the IPT. We believe that additional guidance for selecting such parameters would be helpful. The current formulation of the IPT is also *not* invariant with respect to rotations, which could be a desired property for some applications (or, at the very least, might make the method more robust in the regime of smaller sample sizes). However, with increasing sample size, we observe that *data augmentation* enables IPT-based models to ‘learn’ equivariance despite not being intrinsically invariant (see [Figure S.5](#) in the appendix for preliminary results). Moreover, if point clouds become high-dimensional, our direction-sampling procedure becomes inefficient—this is not a problem for the shape-generation tasks we are tackling in this paper, though. A larger limitation in point-cloud generation involves the question of evaluation metrics: In the absence of ground-truth information, it is hard to choose the ‘best’ model (unlike point-cloud reconstruction, where well-defined metrics exist). In terms of other generative-model evaluation metrics like COV and 1-NNA, our models exhibits only middling performance, ranking slightly better in terms of COV than in terms of 1-NNA. This is in stark contrast to its excellent reconstruction performance and its strong MMD scores. Taken together, this indicates that the $\widehat{\text{IPT}}$ is capable of generating high-quality point clouds, but its generative part, based on simple VAEs, is not sufficiently powerful to provide *diverse* representations.

Future work. Hence, future research concerning more complex—and expressive—model architectures and their capabilities would be intriguing and potentially lead to improved results in terms of diversity metrics, albeit at the cost of computational efficiency. We believe *diffusion models* to be particularly suitable for generating new IPTs, given their proven track record in image generation. Another possible extension of the IP-Encoder could involve a transformer-based architecture [\[36\]](#), due to its high capacity for tasks involving sets. This could be achieved by considering the IPT as a bag of tokens with associated directions. Along those lines, the selected directions for computing the IPT would benefit from a proper *positional encoding*, as opposed to our implicit (but simple) encoding; this would enable a model to better capture dependencies between individual directions. In addition, given the connection of the IPT to more involved geometrical-topological descriptors for graphs, meshes, or higher-order complexes, a natural question is to what extent our inversion results apply to such data. This would further serve to contextualise the IPT within the emerging field of topological deep learning [\[28\]](#) but requires extensions to handle datasets without strong geometrical signals [\[5\]](#), for instance. Another feasible next step would be the extension to *graph generation tasks*, thus building a bridge between *image generation* and *graph generation*. The caveats we raised about generative metrics are all the more relevant in the context of graph learning, though [\[19, 26\]](#). Finally, while we have not yet explored any *single-shot* or even *zero-shot* experiments, we believe that our work may pave the path towards novel, efficient *point cloud foundation models*.

Acknowledgments

The authors are grateful for the productive discussions with the reviewers and their useful feedback, which has significantly strengthened our work. We are indebted to them and the area chair, who, like us, believe in the merits of this paper. The authors are also grateful for the thorough feedback received from Julius von Rohrscheidt. This work has received funding from the Swiss State Secretariat for Education, Research, and Innovation (SERI).

Impact Statement

This paper presents work whose goal is to advance the field of machine learning, specifically the problem of high-quality point cloud (or shape) generation. Given the computational efficiency of our method, we believe it to be beneficial for ‘democratising’ access to machine learning algorithms. There are many other potential societal consequences of point cloud generation methods, none which we feel must be specifically highlighted here.

References

- [1] Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., Bodenstein, S. W., Evans, D. A., Hung, C.-C., O’Neill, M., Reiman, D., Tunyasuvunakool, K., Wu, Z., Žemgulytė, A., Arvaniti, E., Beattie, C., Bertolli, O., Bridgland, A., Cherepanov, A., Congreve, M., Cowen-Rivers, A. I., Cowie, A., Figurnov, M., Fuchs, F. B., Gladman, H., Jain, R., Khan, Y. A., Low, C. M. R., Perlin, K., Potapenko, A., Savy, P., Singh, S., Stecula, A., Thillaisundaram, A., Tong, C., Yakneen, S., Zhong, E. D., Zielinski, M., Žídek, A., Bapst, V., Kohli, P., Jaderberg, M., Hassabis, D., and Jumper, J. M. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630(8016):493–500, 2024.
- [2] Achlioptas, P., Diamanti, O., Mitliagkas, I., and Guibas, L. Learning representations and generative models for 3D point clouds. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 40–49. PMLR, 2018.
- [3] Amboage, J., Röell, E., Schnider, P., and Rieck, B. LEAP: Local ECT-based learnable positional encodings for graphs, 2025. URL <https://arxiv.org/abs/2510.00757>.
- [4] Amézquita, E. J., Quigley, M. Y., Ophelders, T., Landis, J. B., Koenig, D., Munch, E., and Chitwood, D. H. Measuring hidden phenotype: Quantifying the shape of barley seeds using the Euler characteristic transform. *in silico Plants*, 4(1):diab033, 12 2021.
- [5] Ballester, R., Röell, E., Schmid, D. B., Alain, M., Escalera, S., Casacuberta, C., and Rieck, B. MANTRA: The manifold triangulations assemblage. In *International Conference on Learning Representations*, 2025.
- [6] Cai, R., Yang, G., Averbuch-Elor, H., Hao, Z., Belongie, S., Snavely, N., and Hariharan, B. Learning gradient fields for shape generation. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J.-M. (eds.), *Computer Vision – ECCV 2020*, pp. 364–381, Cham, Switzerland, 2020. Springer.
- [7] Cheng, A.-C., Li, X., Liu, S., Sun, M., and Yang, M.-H. Autoregressive 3D shape generation via canonical mapping. In Avidan, S., Brostow, G., Cissé, M., Farinella, G. M., and Hassner, T. (eds.), *Computer Vision – ECCV 2022*, pp. 89–104, Cham, Switzerland, 2022. Springer.
- [8] Crawford, L., Monod, A., Chen, A. X., Mukherjee, S., and Rabadán, R. Predicting clinical outcomes in Glioblastoma: An application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.
- [9] Curry, J., Mukherjee, S., and Turner, K. How many directions determine a shape and other sufficiency results for two topological transforms. *Transactions of the American Mathematical Society, Series B*, 9(32):1006–1043, 2022.

[10] Dłotko, P. On the shape that matters—topology and geometry in data science. *European Mathematical Society Magazine*, pp. 5–13, 2024.

[11] Fahim, G., Amin, K., and Zarif, S. Single-view 3D reconstruction: A survey of deep learning methods. *Computers & Graphics*, 94:164–190, 2021.

[12] Fasy, B. T., Micka, S., Millman, D. L., Schenfisch, A., and Williams, L. Challenges in reconstructing shapes from Euler characteristic curves, 2018. URL <https://arxiv.org/abs/1811.11337>.

[13] George, J., Osborn, O. L., Munch, E., Ridgley II, M., and Wang, E. X. On the stability of the Euler characteristic transform for a perturbed embedding, 2025. URL <https://arxiv.org/abs/2506.19991>.

[14] Ghrist, R., Levanger, R., and Mai, H. Persistent homology and Euler integral transforms. *Journal of Applied and Computational Topology*, 2(1):55–60, 2018.

[15] Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. β -VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.

[16] Kamb, M. and Ganguli, S. An analytic theory of creativity in convolutional diffusion models. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 28795–28831. PMLR, 2025.

[17] Kim, H., Lee, H., Kang, W. H., Lee, J. Y., and Kim, N. S. SoftFlow: Probabilistic framework for normalizing flow on manifolds. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 16388–16397. Curran Associates, Inc., 2020.

[18] Kim, J., Yoo, J., Lee, J., and Hong, S. SetVAE: Learning hierarchical composition for generative modeling of set-structured data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15059–15068, 2021.

[19] Krimmel, M., Hartout, P., Borgwardt, K., and Chen, D. Polygraph discrepancy: A classifier-based metric for graph generation, 2025. URL <https://arxiv.org/abs/2510.06122>.

[20] Lee, J., Lee, Y., Kim, J., Kosiorek, A., Choi, S., and Teh, Y. W. Set transformer: A framework for attention-based permutation-invariant neural networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3744–3753. PMLR, 2019.

[21] Liu, Z., Tang, H., Lin, Y., and Han, S. Point-Voxel CNN for efficient 3D deep learning. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[22] Marsh, L., Zhou, F. Y., Qin, X., Lu, X., Byrne, H. M., and Harrington, H. A. Detecting temporal shape changes with the Euler characteristic transform. *Transactions of Mathematics and its Applications*, 8, 2024.

[23] Mémoli, F. and Sapiro, G. Comparing point clouds. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pp. 32–40, 2004.

[24] Munch, E. An invitation to the Euler characteristic transform. *The American Mathematical Monthly*, 132(1):15–25, 2025.

[25] Nadimpalli, K. V., Chattopadhyay, A., and Rieck, B. Euler characteristic transform based topological loss for reconstructing 3D images from single 2D slices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 571–579, 2023.

[26] O'Bray, L., Horn, M., Rieck, B., and Borgwardt, K. Evaluation metrics for graph generative models: Problems, pitfalls, and practical solutions. In *International Conference on Learning Representations*, 2022.

[27] Özyeşil, O., Voroninski, V., Basri, R., and Singer, A. A survey of structure from motion. *Acta Numerica*, 26:305–364, 2017.

[28] Papamarkou, T., Birdal, T., Bronstein, M., Carlsson, G., Curry, J., Gao, Y., Hajij, M., Kwitt, R., Liò, P., Lorenzo, P. D., Maroulas, V., Miolane, N., Nasrin, F., Ramamurthy, K. N., Rieck, B., Scardapane, S., Schaub, M. T., Veličković, P., Wang, B., Wang, Y., Wei, G.-W., and Zamzmi, G. Position: Topological deep learning is the new frontier for relational learning. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, number 235 in Proceedings of Machine Learning Research, pp. 39529–39555. PMLR, 2024.

[29] Qi, C. R., Yi, L., Su, H., and Guibas, L. J. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[30] Ren, X., Huang, J., Zeng, X., Museth, K., Fidler, S., and Williams, F. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4209–4219, June 2024.

[31] Rieck, B. Topology meets machine learning: An introduction using the Euler characteristic transform. *Notices of the American Mathematical Society*, 72(7):719–727, 2025.

[32] Röell, E. and Rieck, B. Differentiable Euler characteristic transforms for shape classification. In *International Conference on Learning Representations*, 2024.

[33] Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 7537–7547. Curran Associates, Inc., 2020.

[34] Toscano-Duran, V., Rottach, F., and Rieck, B. Molecular machine learning using Euler characteristic transforms. In López Fernández, A., Rodríguez-González, A., Leirós-Rodríguez, R., Mata Miquel, C., and González Suárez, V. M. (eds.), *Artificial Intelligence in Biomedicine*, pp. 391–405, Cham, Switzerland, 2026. Springer.

[35] Turner, K., Mukherjee, S., and Boyer, D. M. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 12 2014.

[36] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[37] von Rohrscheidt, J. and Rieck, B. Diss-l-ECT: Dissecting graph data with local Euler characteristic transforms. In Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., and Zhu, J. (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 61790–61809. PMLR, 2025.

[38] Wu, L., Wang, D., Gong, C., Liu, X., Xiong, Y., Ranjan, R., Krishnamoorthi, R., Chandra, V., and Liu, Q. Fast point cloud generation with straight flows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9445–9454, 2023.

[39] Xu, Q.-C., Mu, T.-J., and Yang, Y.-L. A survey of deep learning-based 3D shape generation. *Computational Visual Media*, 9(3):407–442, 2023.

[40] Yang, G., Huang, X., Hao, Z., Liu, M.-Y., Belongie, S., and Hariharan, B. PointFlow: 3D point cloud generation with continuous normalizing flows. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4540–4549, 2019.

- [41] Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [42] Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., and Kreis, K. LION: Latent point diffusion models for 3D shape generation. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 10021–10039. Curran Associates, Inc., 2022.
- [43] Zhang, X., Zhang, Z., Zhang, C., Tenenbaum, J., Freeman, B., and Wu, J. Learning to reconstruct shapes from unseen classes. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [44] Zhou, L., Du, Y., and Wu, J. 3D shape generation and completion through Point-Voxel Diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5826–5835, 2021.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The core contributions of our work are clearly stated and reflected in the abstract and presented in the introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: An extensive discussion is provided in a separate section as well as potential avenues of future work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide clear statements as to what our assumptions are for all theoretical claims. Full proofs are provided in the appendix for brevity.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide detailed information for the models used in the experiments and an interested reviewer should be able to recreate our experiments based on the manuscript. Moreover, the full implementation is publicly available at <https://github.com/aidos-lab/inner-product-transforms>.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We provide the code with adequate instruction for reproducing all our experiments in the supplementary materials. Our implementation and full experimental setup are publicly available.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We provide full details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: Where applicable we provide error bars, if in the literature it is *not* common to report error bars for certain metrics we follow the standards and not provide them.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Extensive tables are provided and we list both training and inference times for all our models. The hardware specifications of the compute resources are also provided in the main text.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We fully comply with the code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We provide an impact statement.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not provide pretrained models and / or scraped data.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets and models used are properly referenced.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: We will provide an anonymized version of the code with instructions for running the main experiments during the review process and after the review process the code will become publicly available.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: Our data and experiments do not use data obtained from human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [\[NA\]](#)

Justification: We do not use data obtained from crowdsourcing nor data from human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: No LLM has been used for the writing, coding, editing nor the research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Appendix

A Changelog	23
B Overview of the IPT	24
C Additional Downsampling Figures	25
D Additional Experiments	26
E Metrics for Point Clouds	30
F Architectural Details	31
G Properties of the Inner Product Transform	32

A Changelog

We provide a brief list of changes that arose during the review process, commenting shortly on how we implemented the feedback received by the reviewers.

- Changes in the *main text*:
 - We fixed the typo concerning 1D and 2D convolutions in [Section 3](#) and [Section 4](#).
 - We updated the constraints [Theorem 2](#) to enforce q to be non-zero.
 - We fixed the typo concerning the training times in [Section 4](#).
 - We clarified all architectures in [Section 3](#) and now discuss transformer-based models in [Section 5](#).
 - We clarified the ordering of directions in [Section 3](#) and comment on improvements (via *positional encodings*) in [Section 5](#).
 - We added details about our hardware in [Section 4](#).
 - We updated [Section 4.4](#) to better reflect the fact that we provide the first *practical* method for inverting an IPT.
- Changes in the *appendix*:
 - We added a new experiment on reconstructing point cloud from *partial views* in [Table S.2](#).
 - We demonstrate *out-of-distribution experiments* in [Table S.3](#).
 - We add a new ablation study with respect to the number of directions in [Table S.4](#).
 - We provided new results on downsampling in [Table S.5](#).
 - We provide new results on reconstructing large point clouds in [Table S.5](#) (bottom row).
 - We added new results for *multi-class generation* on the MNIST and ShapeNet13 datasets in [Figure S.4](#) and [Table S.6](#).
 - We provide a new experiment on *learning equivariance* via *data augmentation* in [Figure S.5](#).
 - We wrote preliminary code for using latent diffusion models and added a first visualisation of the results in [Figure S.6](#).
- We prefer to address the following aspects in *future work*, rather than adding an inadequate discussion to our work:
 - A discussion on how to use spherical harmonics for the IPT (we believe that this could be essential for a completely new approach for calculating IPT, beyond the scope of this paper).
 - A full comparison with models like DeepSDF (we lacked compute to finish this).
 - A rate-distortion analysis (we rewrote the compression aspects at the suggestion of the reviewers and we believe that such an analysis should be better contextualised in a follow-up work).
 - A report of FLOPs (we were unable to collect all this data for the larger baseline models in time and found some reporting inconsistencies).

We very much appreciate the feedback by reviewers, which helped us substantially improve our work.

B Overview of the IPT

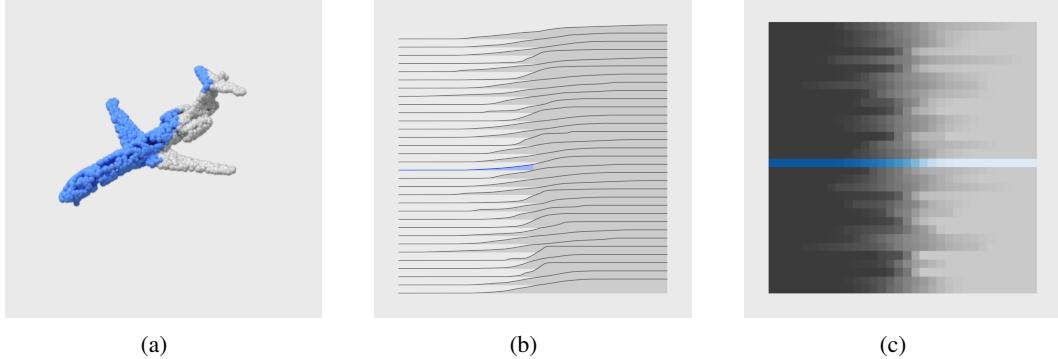
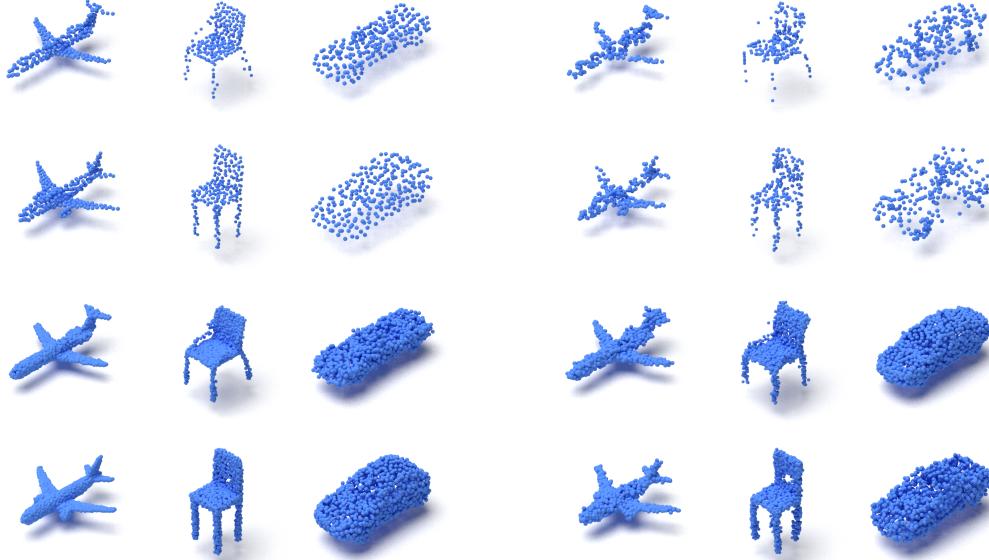


Figure S.1: An overview of the IPT calculation for 32 directions and resolution of 32. For a given direction vector ξ , we filter the point cloud with hyperplanes. The partially filtered point cloud is shown in (a) and points included in the filtration are coloured blue. For each of the 32 directions, sampled uniformly from the sphere, the respective curves along each direction are shown in (b) and the partially-completed curve from (a) is highlighted. Note that neighbouring curves are *not* necessarily related, requiring us to treat each curve as its own signal. Each of the 32 curves is discretised in 32 steps and stacked to form an *image* representation of the point cloud (c) of size 32×32 . The row corresponding to the full curve from (a) and each of the rows corresponds to the curve in the same index in (b). For complex geometries such as the ShapeNet data, we empirically observe that *at least* 128 directions are required, using a resolution of 128.

C Additional Downsampling Figures

We provide a visualisation of [Table 5](#) for the comparison of downsampling with the IP-Downsampler (left) and uniform subsampling (right) for all three classes.



(a) Downsampled point-clouds with the IP-Downsampler (top) and upsampled with the IP-Encoder (bottom).

(b) Downsampled point clouds with uniform subsampling (top) and upsampled with the IP-Encoder (bottom).

Figure S.2: A visual comparison between down- and upsampling using our IP-Downsampler **(a)** and uniform subsampling **(b)**. The downsampled point clouds using the IP-Downsampler are equidistantly spread, in contrast to uniform subsampling. Compare to uniform subsampling, we achieve higher upsampling quality. Our model has a bias towards equidistant point clouds and leads to a better representation of the underlying shape.

D Additional Experiments

This section provides supplementary figures and tables for the experiments in the main text. Please note that some results should be considered *preliminary work*, pointing towards potential future applications, which we included in the interest of full transparency regarding our work.

Table S.1: Reconstruction results for rendering of the three ShapeNet classes. A point cloud is randomly initialised and optimised with the IPT as a loss function. We repeat this experiment for different resolutions and different scales. The ideal scale is 1/4 of the resolution (bottom row).

RESOLUTION	SCALE FACTOR	Airplane		Car		Chair	
		CD (\downarrow)	EMD (\downarrow)	CD (\downarrow)	EMD (\downarrow)	CD (\downarrow)	EMD (\downarrow)
128	128	1.27 \pm 0.00	0.29 \pm 0.00	2.97 \pm 0.00	0.55 \pm 0.00	3.40 \pm 0.01	0.65 \pm 0.00
	64	0.75 \pm 0.00	0.18 \pm 0.00	2.15 \pm 0.00	0.42 \pm 0.00	2.10 \pm 0.00	0.41 \pm 0.00
	32	0.52 \pm 0.00	0.12 \pm 0.00	2.04 \pm 0.00	0.39 \pm 0.00	1.83 \pm 0.00	0.35 \pm 0.00
64	64	2.65 \pm 0.01	0.57 \pm 0.00	4.76 \pm 0.01	0.85 \pm 0.00	5.83 \pm 0.01	1.10 \pm 0.00
	32	1.88 \pm 0.01	0.40 \pm 0.00	4.05 \pm 0.01	0.70 \pm 0.00	4.29 \pm 0.00	0.78 \pm 0.00
	16	1.22 \pm 0.00	0.26 \pm 0.00	3.80 \pm 0.00	0.64 \pm 0.00	3.62 \pm 0.00	0.62 \pm 0.00
32	32	6.21 \pm 0.01	1.32 \pm 0.00	9.85 \pm 0.01	1.99 \pm 0.00	16.92 \pm 0.02	3.54 \pm 0.01
	16	4.11 \pm 0.01	0.90 \pm 0.00	8.56 \pm 0.02	1.63 \pm 0.00	12.55 \pm 0.02	2.53 \pm 0.00
	8	2.77 \pm 0.00	0.61 \pm 0.00	7.24 \pm 0.01	1.35 \pm 0.00	9.05 \pm 0.00	1.73 \pm 0.00

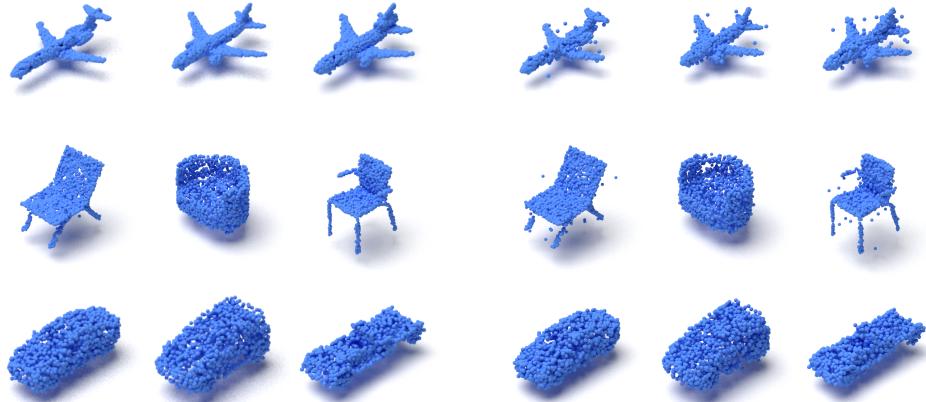


Figure S.3: Examples of *rendered* point clouds using our optimisation scheme for the three classes of ShapeNet. On the left, the point clouds are rendered at a resolution of 128 with a scale of 32, leading to high-quality reconstructions. The right shows the *same* rendering with the scale set to 128, where we observe unstable gradients. The resulting point clouds have only partially converged, with some clear outliers.

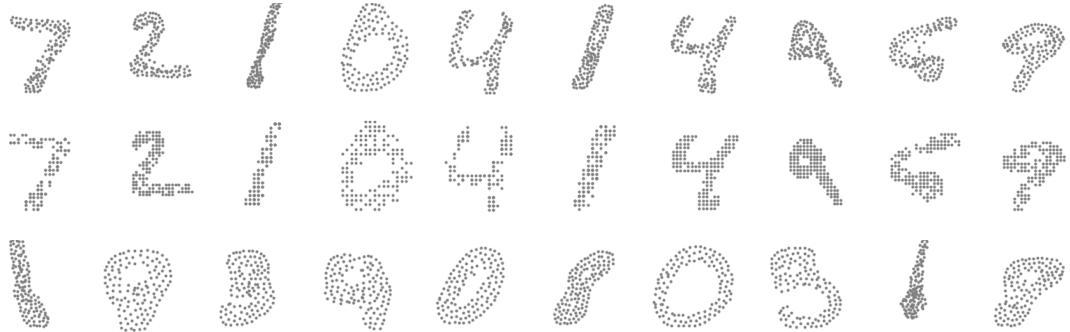


Figure S.4: To ensure our pipeline can capture *multi-class distributions*, we train the IP-Encoder and IP-VAE on the MNIST dataset. The top row depicts the reconstructed point clouds with the ground truth in the middle row. The last row depicts point clouds generated with the IP-VAE. A latent vector is sampled from the latent distribution and decoded into an IPT. The novel IPT is subsequently converted to a point cloud with the IP-Encoder. Using our evaluation we find the 1-NNA CD and EMD scores to be 70.31 and 62.50, respectively. This indicates that our pipeline can handle multi-class data distributions.



Figure S.5: Recent work showed that *equivariance* with respect to certain operations like rotations can also be achieved through data augmentation, thus obviating the need for more complex architectures [1]. To assess the capabilities of the IP-Encoder in this context, we follow Qi et al. [29] and use a point cloud version the MNIST dataset of handwritten digits. During training, we apply a random rotation to each point cloud and then compute the IPT, thus permitting the model to learn an equivariant representation of the data. As the figure shows, data augmentation is *sufficient* to encode rotations, resulting in an equivariant model without having to specifically add equivariance as a separate inductive bias.

Table S.2: Reconstruction of point clouds based on *partial views* using the GenRe [43] dataset. The point clouds in the dataset are sampled from depth image in a particular direction and the task consists of reconstructing the full point from the partial view. We asses the capability of the IP-Encoder to reconstruct the full point cloud from the partial view *without* any additional training. Thus, the model is only trained on the set of *complete* point clouds. We use the trained IP-Encoder from Section 4.1; CD and EMD have to be multiplied by 1.00×10^3 and 1.00×10^2 , respectively.

MODEL	<i>Chair</i> (\downarrow)		<i>Car</i> (\downarrow)	
	CD	EMD	CD	EMD
SoftFlow	2.79	3.30	1.85	2.79
PointFlow	2.71	3.65	1.80	2.85
PVD	3.21	2.94	1.77	2.15
IP-Encoder	1.10	1.24	0.71	0.91

Table S.3: Out-of-distribution results for the IP-Encoder model. The IP-Encoder trained on the class of airplanes is used to reconstruct the category of cars and chairs to provide an insight into true out-of-distribution reconstruction and capacity to generalise. Results indicate that our model is not *fully general* and thus incapable of reconstructing arbitrary point cloud configurations. Due to the simplicity of our model architecture, this is partially to be expected. We hope to improve the multi-modal capabilities of our model in future work.

AIRPLANE TO	CD	EMD
Car	22.68	9.36
Chair	203.69	52.29

Table S.4: Ablation with respect to the number of directions used in the IPT to investigate the impact on reconstruction quality. The spatial resolution is kept at 128 (the default) and we progressively reduce the number of directions. Although the reconstruction quality remains *consistently* high, the number of directions has, in general, a positive impact on reconstruction quality.

Airplane (↓)		
NUMBER OF DIRECTIONS	CD	EMD
4	1.29	2.05
8	1.18	1.88
16	1.09	1.79
32	1.07	1.69
64	1.03	1.56
128	1.05	1.57

Table S.5: Additional downsampling experiments. We follow the same setup as in [Section 4.3](#) and train a separate downsampling model for each cardinality. With each model, we predict a downsampled version of the point cloud, which is subsequently upsampled with the IP-Encoder, to obtain a reconstruction of the original point cloud. We observe high quality over all cardinalities, even when the original point cloud is summarized with only 32 points (!). As a comparison, we add the original reconstruction result (cf. [Table 1](#)) without downsampling in the second to last row. The bottom part of the table demonstrates the capacity of the IP-Encoder to reconstruct larger point clouds (without downsampling). Although overall performance is still high, our model has the *fundamental limitation* that it predicts the point clouds in a *single* $N \times 3$ vector. To generate large point clouds, we thus believe that different architectures, such as set transformers [20], will provide better scalability with respect to the cardinality.

NUMBER OF POINTS	Airplane (↓)		Chair (↓)		Car (↓)	
	CD	EMD	CD	EMD	CD	EMD
32	1.40	3.38	18.25	11.61	9.99	6.05
64	1.20	2.68	14.34	9.79	7.57	4.97
128	1.15	2.31	12.07	8.88	6.64	4.51
256	1.13	2.17	11.59	8.53	6.29	4.23
512	1.12	2.09	11.55	8.34	6.32	4.36
2048	1.05	1.57	9.24	6.19	5.82	3.18
4096	0.87	1.61	8.17	6.33	5.13	3.27

Table S.6: Preliminary results for *multi-class reconstruction* on the ShapeNet13 dataset. The dataset contains 13 classes (airplanes, cars, chairs, lamps, tables, sofas, cabinets, benches, telephones, speakers, monitors, vessels, and rifles) from the full ShapeNet dataset. We follow the experimental setup of Zhou et al. [44] and cite their results as comparison partners. Our model performs well despite its conceptual simplicity; we envision more expressive architectures to perform even better.

MODEL	CD (\downarrow)	EMD (\downarrow)
PVD	58.65	57.85
PointFlow	63.25	66.05
LION	51.85	48.95
IP-Encoder (Ours)	63.92	50.28

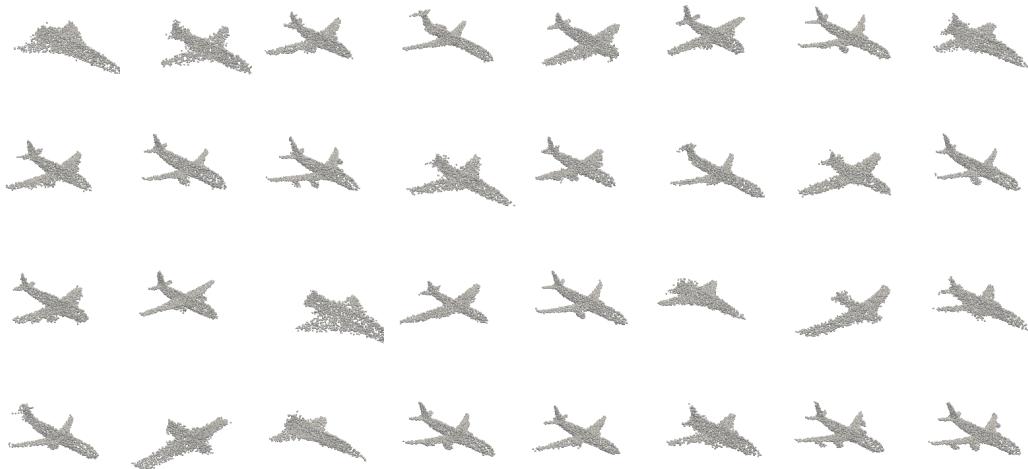


Figure S.6: Preliminary results for generating samples of the *airplane* class using a latent diffusion model. We train a VAE on the set of IPTs of airplanes, followed by training a diffusion model on the latent embeddings of the IPTs to generate new IPTs. Again, we use the IP-Encoder to map a generated IPT into a point cloud. This demonstrates the flexibility and generality of our model, which we look forward to exploiting better in future work.

E Metrics for Point Clouds

A good metric for point cloud generation balances computational speed and theoretical guarantees. Finding such metrics is a challenging task, since often computations require the consideration of all pairs of points between the two point clouds. An example is the Gromov–Hausdorff distance [23], which has advantageous theoretical properties, but is hard to evaluate. Two metrics are commonly used to describe the distance between point clouds, the Chamfer Distance (CD) and Earth Mover’s Distance (EMD), for which we present a self-contained summary here. Although not a metric in the mathematical sense, CD poses a good balance between computational speed and quality and is defined for point clouds X and Y as

$$\text{CD}(X, Y) = \sum_{x \in X} \min_{y \in Y} \|x - y\| + \sum_{y \in Y} \min_{x \in X} \|x - y\|. \quad (8)$$

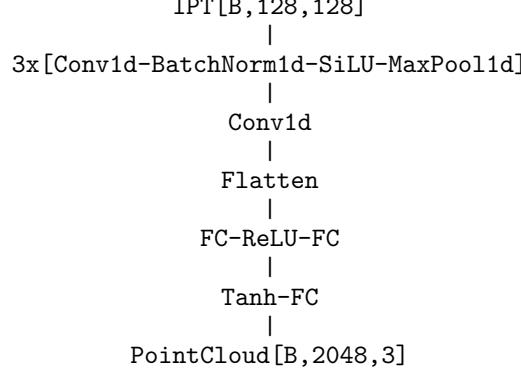
For the EMD, by contrast, the distance between point clouds is viewed as the cost required to transport one point cloud into the other, i.e.

$$\text{EMD}(X, Y) = \min_{\phi: X \rightarrow Y} \sum_{x \in X} \|x - \phi(x)\|, \quad (9)$$

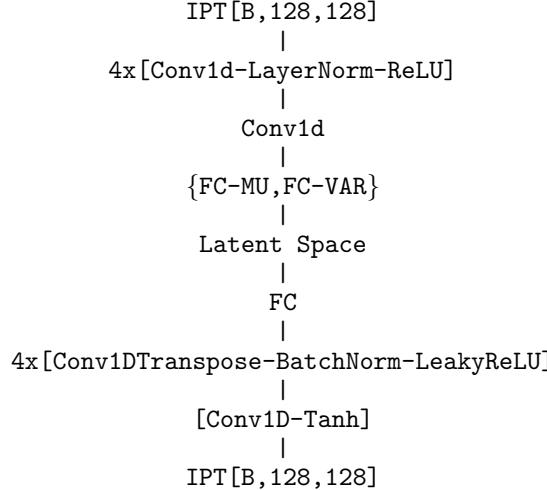
where ϕ solves an optimal transport problem. Solving the optimal transport problem is a computationally intensive task that becomes prohibitive for medium to large point clouds. The properties of using the CD as loss term were investigated in Achlioptas et al. [2], revealing that reconstructions had non-uniform surface density, compared to the uniformly sampled ground truth points. A strong advantage of the CD is its computational efficiency for medium to large point cloud cardinalities.

F Architectural Details

We provide the details of both the IP-Encoder and IP-VAE architecture. The IPT as defined in Eq. (4) is viewed as an image of shape $n_h \times n_d$, where n_h is the discretisation of the heights along each direction and n_d is the number of directions. In the architecture, n_d denotes channels in the CNN; to accommodate for this, we expect the input to be of the form $n_b \times n_d \times n_h$, where n_b is the batch size. In all our experiments we set $n_d = n_h = 128$ and denote a batch of IPTs in the diagram with $\text{IPT}[\mathbf{B}, 128, 128]$. The internal architecture of the IP-Encoder is:



In the architecture above, $\text{PointCloud}[\mathbf{B}, 2048, 3]$ denotes the *final* predicted batch of point clouds. Our IP-VAE has the following architecture:



In the diagram above, latent embeddings are denoted by FC-MU and FC-VAR, respectively.

G Properties of the Inner Product Transform

This section discusses all proofs for statements from the main paper. While the first result about the injectivity is already known, we provide a novel, highly-accessible proof, which only requires basic concepts from linear algebra. In the following, we will consider two point clouds $X, Y \subset \mathbb{R}^n$ to be equal if they are equal in the sense of sets; in particular, this implies that they need to have the same cardinality. We will also assume that all points are in *general position*.

Theorem 1. *Given two point clouds X, Y with $X \neq Y$, we have $\text{IPT}(X) \neq \text{IPT}(Y)$.*

Proof. Let $X, Y \subset \mathbb{R}^n$ with $X \neq Y$ and let $\Xi \subset S^{n-1}$ be a set of $n + 1$ affinely-independent directions. Given $\xi \in \Xi$, we write $X_\xi := \{\langle x, \xi \rangle \mid x \in X\}$ for the projection of X onto the one-dimensional subspace along ξ , with $\langle x, \xi \rangle$ denoting the standard Euclidean inner product. For $t \in \mathbb{R}$, we define $X_{\xi, t} := \{x \in X \mid \langle x, \xi \rangle \leq t\}$. If for any direction $\xi \in \Xi$ we have $X_\xi \neq Y_\xi$, we are done because this means that we can find thresholds $t_1, t_2 \in \mathbb{R}$ with $t_1 < t_2$ such that for $t \in [t_1, t_2]$, the cardinality of $X_{\xi, t}$ changes but the cardinality of $Y_{\xi, t}$ does not change or vice versa. Thus, let us assume that $X_\xi = Y_\xi$ for all directions ξ . For each ξ , we may sort the values by magnitude and calculate differences, i.e. expressions of the form $\langle x, \xi \rangle - \langle y, \xi \rangle$. We have $\langle x, \xi \rangle - \langle y, \xi \rangle = 0$ for all directions $\xi \in \Xi$ by assumption, which we may rewrite as $\langle x - y, \xi \rangle = 0$. Treating this as a system of $n + 1$ linear equations, this is equivalent to stating that the kernel of the corresponding linear map is the whole domain. However, since the $n + 1$ directions are affinely independent, all coefficients must be zero, implying that the point clouds are the same. This is a contradiction, so our initial assumption must have been wrong. Thus, there is $\xi \in \Xi$ such that $X_\xi \neq Y_\xi$, so we have $\text{IPT}(X) \neq \text{IPT}(Y)$. \square

Lemma 1. *Let $B_n(R)$ denote a ball of radius R in \mathbb{R}^n . For two point clouds $X, Y \subset B_n(R)$, define their distance as*

$$d(X, Y) := \frac{1}{|\Xi|} \sum_{\xi \in \Xi} \|\text{IPT}(X)|_\xi - \text{IPT}(Y)|_\xi\|_2, \quad (5)$$

where Ξ is a finite set of directions and $\|\cdot\|_2$ is the L^2 -norm restricted to the interval $[-R, R]$. The function $d(\cdot, \cdot)$ satisfies the definition of a metric.

Proof. If $d(X, Y) = 0$ then $\|\text{IPT}(X) - \text{IPT}(Y)\|_2 = 0$ and since $\|\cdot\|_2$ is a metric, it follows that $\text{IPT}(X) = \text{IPT}(Y)$. Since the IPT is injective, we conclude that $X = Y$. Equality in this case is seen as equality the of sets, that is up to permutation. Both the reverse implication and symmetry follow from the definitions. For the triangle inequality, we verify

$$\begin{aligned} d(X, Y) &= \|\text{IPT}(X) - \text{IPT}(Y)\|_2 \\ &= \|(\text{IPT}(X) - \text{IPT}(Z)) - (\text{IPT}(Y) - \text{IPT}(Z))\|_2 \\ &\leq \|\text{IPT}(X) - \text{IPT}(Z)\| + \|\text{IPT}(Y) - \text{IPT}(Z)\|_2 \\ &= d(X, Z) + d(Z, Y). \end{aligned} \quad (10)$$

\square

Lemma 2. *Let $X, Y \subset \mathbb{R}^n$ be disjoint point clouds, then*

$$\text{IPT}(X \cup Y) = \text{IPT}(X) + \text{IPT}(Y). \quad (6)$$

Proof. It follows from the definition of the IPT that for all $\xi \in S^{n-1}$ and $h \in \mathbb{R}$ we have

$$\begin{aligned} \text{IPT}(X \cup Y) &= \sum_{x \in X \cup Y} \mathbb{1}_x(\xi, h) \\ &= \sum_{x \in X} \mathbb{1}_x(\xi, h) + \sum_{y \in Y} \mathbb{1}_y(\xi, h) \\ &= \text{IPT}(X) + \text{IPT}(Y). \end{aligned} \quad (11)$$

The second equality uses that the intersection of X and Y is empty, implying that a point is either in X or in Y . \square

Theorem 2. *Given two point clouds $X, Y \subset \mathbb{R}^n$, the IPT is surjective for the rational linear subspace spanned by $\text{IPT}(X)$ and $\text{IPT}(Y)$, up to a rational scaling factor. In particular, for $p, q \in \mathbb{N}_0$ with $0 \leq p \leq q$ and $q > 0$ we have*

$$\frac{p}{q}\text{IPT}(X) + \frac{q-p}{q}\text{IPT}(Y) = \frac{1}{q}\text{IPT}(Z), \quad (7)$$

where $Z = \cup_p X \cup_{q-p} Y$.

Proof. To show surjectivity, it suffices to show that for any rational linear combination of IPTs there exists a point cloud $Z \subset \mathbb{R}^n$ that has an IPT equal to that linear combination, up to a rational coefficient. This is to say that $\forall \alpha, \beta \in \mathbb{Q} \exists \gamma \in \mathbb{Q}$ such that

$$\alpha\text{IPT}(X) + \beta\text{IPT}(Y) = \gamma\text{IPT}(Z). \quad (12)$$

Let $\alpha = p/q$ and $\beta = r/s$, then

$$\begin{aligned} qs \left[\frac{p}{q}\text{IPT}(X) + \frac{r}{s}\text{IPT}(Y) \right] &= sp\text{IPT}(X) + qr\text{IPT}(Y) \\ &= \text{IPT}(\cup_{sp} X \cup_{qr} Y). \end{aligned} \quad (13)$$

Thus, setting $\gamma = 1/qs$ and $Z = \cup_{sp} X \cup_{qr} Y$ does the trick. We conclude that the IPT is surjective on rational linear combinations. Let $0 \leq p \leq q$, statement follows from the equalities

$$\begin{aligned} \frac{p}{q}\text{IPT}(X) + \frac{q-p}{q}\text{IPT}(Y) \\ &= \frac{1}{q}[p\text{IPT}(X) + (q-p)\text{IPT}(Y)] \\ &= \frac{1}{q}[\text{IPT}(\cup_p X) + \text{IPT}(\cup_{q-p} Y)] \\ &= \frac{1}{q}\text{IPT}(\cup_p X \cup_{q-p} Y). \end{aligned} \quad (14)$$

□