

# SEMI-DISCRETE NORMALIZING FLOWS THROUGH DIFFERENTIABLE VORONOI TESSELLATION

Ricky T. Q. Chen, Brandon Amos, Maximilian Nickel

Facebook AI Research

rtqichen, bda, maxn@fb.com

## ABSTRACT

Mapping between discrete and continuous distributions is a difficult task and many have had to resort to approximate or heuristical approaches. We propose a tessellation-based approach that directly learns quantization boundaries on a continuous space, complete with exact likelihood evaluations. This is done through constructing normalizing flows on convex polytopes defined via a differentiable tessellation. Using a simple homeomorphism with an efficient log determinant Jacobian, we can then cheaply parameterize distributions on bounded domains.

We explore this approach in two application settings, mapping from discrete to continuous and vice versa. Firstly, a *Voronoi dequantization* allows automatically learning quantization boundaries in a multidimensional space. The location of boundaries and distances between regions can encode useful structural relations between the quantized discrete values. Secondly, a *Voronoi mixture model* has constant computation cost for likelihood evaluation regardless of the number of mixture components. Empirically, we show improvements over existing methods across a range of structured data modalities, and find that we can achieve a significant gain from just adding Voronoi mixtures to a baseline model.

## 1 INTRODUCTION

Likelihood-based models have seen increasing usage across multiple data modalities. In particular, the family of normalizing flows stands out as a large amount of structure can be incorporated into the model, aiding its usage in modeling a wide variety of domains such as images [7, 22], graphs [28], invariant distributions [24, 2] and molecular structures [40]. However, the majority of works focus on only continuous functions and continuous random variables. This restriction can make it difficult to apply such models to distributions with implicit discrete structures. In this work, we incorporate discrete structure into standard normalizing flows, while being entirely composable with any other normalizing flow. Specifically, we propose a homeomorphism between an unbounded domain and a convex polytope—defined through a learnable tessellation of the domain. This homeomorphism is cheap to compute, has a cheap inverse, and we additionally provide an efficient and exact formulation of the resulting change in density. In other words, this transformation is highly scalable.




Figure 1: We propose an invertible mapping  $f$  between  $\mathbb{R}^D$  and a convex polytope, which is parameterized based on a differentiable Voronoi tessellation of  $\mathbb{R}^D$ . This mapping adds discrete structure into normalizing flows, and its inverse  $f^{-1}$  and log determinant Jacobian can both be efficiently computed for high dimensions.

## 2 PRELIMINARIES

**Normalizing Flows.** This family of generative models [34, 23] typically includes any model that makes use of invertible transformations  $f$  to map samples between distributions. The relationship between the original and transformed density functions have a closed form expression,

$$p_z(f(x)) = p_x(x) \left| \det \frac{\partial f(x)}{\partial x} \right|^{-1}. \quad (1)$$

If the domain and codomain of  $f$  are different, then  $p_x$  and  $p_z$  can have different supports, *i.e.* regions where probability is non-zero. However, existing designs of  $f$  that have this property only act in a single dimension. The use of general support-modifying invertible transformations have not been discussed extensively in the literature.

**Dequantization.** When modeling discrete data with density models, standard approaches typically rely on dequantization methods. These methods provide a correspondence between discrete values and continuous sub-spaces, where a discrete variable  $\mathbf{y} \in \mathcal{Y}$  is randomly placed within a disjoint subset of  $\mathbb{R}^D$  [42].

Let  $A_{\mathbf{y}}$  denote the subspace corresponding to the value of  $\mathbf{y}$  and let  $q(\mathbf{x}|\mathbf{y})$  be the dequantization model which has a bounded support in  $A_{\mathbf{y}}$ . Then a density model  $p(\mathbf{x})$  with support over  $\mathbb{R}^D$  satisfies

$$\begin{aligned} \log p(\mathbf{y}) &\geq \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x}|\mathbf{y})} [\log(\mathbb{1}_{[\mathbf{x} \in A_{\mathbf{y}}]} p(\mathbf{x})) - \log q(\mathbf{x}|\mathbf{y})] \\ &= \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x}|\mathbf{y})} [\log p(\mathbf{x}) - \log q(\mathbf{x}|\mathbf{y})] \end{aligned} \quad (2)$$

Thus with an appropriate choice of dequantization, maximizing the likelihood under the density model  $p(\mathbf{x})$  is equivalent to maximizing the likelihood under the discrete model  $p(\mathbf{y})$ .

To the best of our knowledge, in all existing dequantization methods, the value of  $D$  and the choice of subsets  $A_{\mathbf{y}}$  are entirely dependent on the type of discrete data (ordinal vs non-ordinal) and the number of discrete values  $|\mathcal{Y}|$  in the non-ordinal case. Furthermore, the subsets  $A_{\mathbf{y}}$  do not interact with one another and are fixed during training. In contrast, we conjecture that important properties such as (dis)similarities between discrete values should be modeled as part of the parameterization of  $A_{\mathbf{y}}$ , and that it'd be useful to be able to automatically adapt these subsets  $A_{\mathbf{y}}$  based on gradients from  $p(\mathbf{x})$ .

**Disjoint mixture models.** Building mixture models is one of the simplest methods for creating more flexible distributions from simpler one. However, in practice it is computationally expensive for large number of mixture components, because evaluating the likelihood under a mixture model requires evaluating the likelihood of each component. To alleviate this issue, Dinh et al. [8] recently proposed to use components with disjoint support. In particular, let  $\{A_k\}_{k=1}^K$  be disjoint subsets of  $\mathbb{R}^D$ , such that each mixture component is defined on one subset and has support restricted to that particular subset. The likelihood of the mixture model then simplifies to

$$p(\mathbf{x}) = \sum_{k=1}^K p(\mathbf{x}|k)p(k) = \sum_{k=1}^K \mathbb{1}_{[\mathbf{x} \in A_k]} p(\mathbf{x}|k)p(k) = p(\mathbf{x}|k = g(\mathbf{x}))p(k = g(\mathbf{x})) \quad (3)$$

where  $g: \mathbb{R}^D \rightarrow \{1, \dots, K\}$  is a set identification function that satisfies  $\mathbf{x} \in A_{g(\mathbf{x})}$ . This framework allows building large mixture models while no longer having a compute cost that scales with  $K$ . In contrast to variational approaches with discrete latent variables, the use of disjoint subsets provides an exact log-likelihood and not a lower bound.

## 3 VORONOI TESSELLATION FOR NORMALIZING FLOWS

We separate the domain into subsets through a Voronoi tessellation [46]. This induces a correspondence between each subset and a corresponding *anchor point*, which provides a parameterization of the tessellation for gradient-based optimization.

Let  $X = \{x_1, \dots, x_K\}$  be a set of anchor points in  $\mathbb{R}^D$ . The Voronoi cell for each anchor point is given by

$$V_k \triangleq \{x \in \mathbb{R}^D : \|x_k - x\| < \|x_j - x\|, j = 1, \dots, K\}, \quad (4)$$




Figure 2: Existing dequantization methods can be seen as special cases of Voronoi dequantization with fixed anchor points. In addition to decoupling the dimension of the embedding space from the number of discrete values, the Voronoi dequantization can learn to model the similarities of discrete values through the positions and boundaries of their Voronoi cells.

*i.e.* it defines a subspace containing all points which have the anchor points as its nearest neighbor. This subspace can equivalently be expressed in the form of a convex polytope,

$$V_k = \{x \in \mathbb{R}^D : Ax < b\}, \quad (5)$$

where  $A_i = 2(x_i - x_k)^\top, b_i = \|x_i\|^2 - \|x_k\|^2$ .

For simplicity, we also include box constraints so that all Voronoi cells are bounded in all directions.

$$V_k = \{x \in \mathbb{R}^D : Ax < b, c_l < x < c_r\} \quad (6)$$

Thus, the learnable parameters of this Voronoi tessellation are the anchor points  $\{x_1, \dots, x_K\}$ , and the box constraints  $c_l, c_r \in \mathbb{R}^D$ .

### 3.1 INVERTIBLE MAPPING ONTO THE VORONOI CELL

Given a cell  $V_k$  for some  $k \in \{1, \dots, K\}$ , we construct an invertible mapping  $f_k : \mathbb{R}^D \rightarrow V_k$  by the following 2-step procedure.

Let  $x \in \mathbb{R}^D$ . If  $x = x_k$ , the anchor point of  $V_k$ , we simply set  $f_k(x) = x$ . Otherwise, the first step is as follows:

1. Find where the ray that starts from  $x_k$  in the direction of  $x$  intersects with the boundary of  $V_k$ .

Since  $V_k$  is a convex polytope, we can frame this as a linear programming problem. First define the direction  $\delta_k(x) \triangleq \frac{x - x_k}{\|x - x_k\|}$  and the ray  $x(\lambda) \triangleq x_k + \lambda \delta_k(x)$ , with  $\lambda > 0$ . Then this problem can be solved through the following linear program:

$$\max \lambda \quad \text{s.t.} \quad x(\lambda) \in \overline{V_k}, \lambda \geq 0 \quad (7)$$

Let  $\lambda^*$  be the solution. This first step solves for the farthest point in the Voronoi cell in the direction of  $x$ . Using this knowledge, we can now map all points that lie on this ray onto the Voronoi cell.

1. Apply an invertible transformation such that any point on the ray  $\{x(\lambda) : \lambda > 0\}$  is mapped onto the line segment  $\{x_k + \alpha(x(\lambda^*) - x_k) : \alpha \in (0, 1)\}$ .

There are many possible choices for designing this transformation. An intuitive choice is to use a monotonic transformation of the distance from  $x$  to the anchor point  $x_k$ .

$$f_k(x) \triangleq x_k + \alpha_k \left( \frac{\|x - x_k\|}{\|x(\lambda^*) - x_k\|} \right) (x(\lambda^*) - x_k) \quad (8)$$

where  $\alpha_k$  is an appropriate *invertible* squashing function from  $\mathbb{R}^+$  to  $(0, 1)$ . In our experiments, we use  $\alpha_k(h) = \text{softsign}(\gamma_k h)$  where  $\gamma_k$  is a learned cell-dependent scale.

**Remarks** The problem in Equation (7) can be solved exactly and without the use of an iterative convex optimization solver, which also provides analytical derivatives of the solution. The inverse computation shares the same first step of this procedure. Finally, we provide a method of computing the exact log determinant Jacobian that only involves dot products, thus efficiently scaling with the number of dimensions. See detailed explanations of these remarks and propositions in Appendix A.




Figure 3: **2D quantized toy data.** Voronoi Flows can model complex relations between discrete values. Each of the two discrete variables have 91 possible values, and no knowledge of ordering is given to the models.




Figure 4: Tessellation can be done in a transformed space; nonlinear boundaries are shown.

**Applications** We apply this approach to dequantization—which maps from discrete variables to continuous variables—and disjoint mixture modeling—which maps from continuous to discrete. Importantly, all existing dequantization boundaries are special cases of Voronoi tessellations (see Figure 2). Our approach allows us to directly parameterize and learn these boundaries through just gradient-based optimization. Detailed explanations are in Appendix B.

## 4 EXPERIMENTS

Additional experiments regarding permutation-invariant itemset modeling and language modeling are in Appendix D.

**Voronoi dequantization** We start by considering synthetic distribution over two discrete variables, to compare against Discrete Flows [43]. Figure 3 shows the target probability mass functions, which are created by a 2D distribution where each dimension is quantized into 91 discrete values. Though the data is ordinal, no knowledge of ordering is given to the models. We see that Discrete Flows has trouble fitting these, presumably because it is difficult to rearrange the target distribution into a factorized distribution. We also show results on UCI discrete data sets in Table 2.

**Voronoi mixture models** We trained disjoint mixture models using the inverse of our transformation, where data is mapped from each bounded Voronoi cell onto the unbounded space, which is then composed with a conditional normalizing flow model. Results are shown in Section 3.1. The boundaries are nonlinear as it is done after a nonlinear transformation.

Method	Connect4	Forests	Mushroom	Nursery	PokerHands	USCensus90
Voronoi Deq.	12.92±0.07	14.20±0.05	9.06±0.05	9.27±0.04	19.86±0.04	24.19±0.12
Simplex Deq.	13.46±0.01	16.58±0.01	9.26±0.01	9.50±0.00	19.90±0.00	28.09±0.08
BinaryArgmax Deq.	13.71±0.04	16.73±0.17	9.53±0.01	9.49±0.00	19.90±0.01	27.23±0.02
Discrete Flow	19.80±0.01	21.91±0.01	22.06±0.01	9.53±0.01	19.82±0.03	55.62±0.35

Table 1: **Discrete UCI data sets.** Negative log-likelihood results on the test sets in nats.

## REFERENCES

- [1] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and Zico Kolter. Differentiable convex optimization layers. *Advances in Neural Information Processing Systems*, 2019.
- [2] Marin Biloš and Stephan Günnemann. Scalable normalizing flows for permutation invariant densities. In *International Conference on Machine Learning*, pp. 957–967. PMLR, 2021.
- [3] Tom Brijs, Gilbert Swinnen, Koen Vanhoof, and Geert Wets. Using association rules for product assortment decisions: A case study. In *Knowledge Discovery and Data Mining*, pp. 254–260, 1999.
- [4] Ricky T. Q. Chen. torchdiffeq, 2018. URL <https://github.com/rtqichen/torchdiffeq>.
- [5] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 2018.
- [6] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [7] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. *International Conference on Learning Representations*, 2017.
- [8] Laurent Dinh, Jascha Sohl-Dickstein, Hugo Larochelle, and Razvan Pascanu. A rad approach to deep mixture models. *arXiv preprint arXiv:1903.07714*, 2019.
- [9] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- [10] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in Neural Information Processing Systems*, 32:7511–7522, 2019.
- [11] Karolien Geurts, Geert Wets, Tom Brijs, and Koen Vanhoof. Profiling high frequency accident locations using association rules. In *Proceedings of the 82nd Annual Transportation Research Board, Washington DC. (USA), January 12-16*, pp. 18pp, 2003.
- [12] Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- [13] Peter M Gruber. Optimum quantization and its applications. *Advances in Mathematics*, 186(2): 456–497, 2004.
- [14] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [15] Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pp. 2722–2730. PMLR, 2019.
- [16] Emiel Hoogeboom, Jorn WT Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. *Advances in Neural Information Processing Systems*, 2019.
- [17] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Towards non-autoregressive language models. *Advances in Neural Information Processing Systems*, 2021.
- [18] Chin-Wei Huang, Laurent Dinh, and Aaron Courville. Augmented normalizing flows: Bridging the gap between generative flows and latent variable models. *arXiv preprint arXiv:2002.07101*, 2020.
- [19] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *International Conference on Learning Representations*, 2017.

- [20] Hyunjik Kim, George Papamakarios, and Andriy Mnih. The lipschitz constant of self-attention. In *International Conference on Machine Learning*, pp. 5562–5571. PMLR, 2021.
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- [22] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in Neural Information Processing Systems*, 2018.
- [23] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [24] Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: exact likelihood generative learning for symmetric densities. In *International Conference on Machine Learning*, pp. 5361–5370. PMLR, 2020.
- [25] Alex Kulesza and Ben Taskar. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286, 2012. ISSN 1935-8237. doi: 10.1561/22000000044.
- [26] Bruno Lévy. A numerical algorithm for 12 semi-discrete optimal transport in 3d. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1693–1715, 2015.
- [27] Phillip Lippe and Efstratios Gavves. Categorical normalizing flows via continuous transformations. *International Conference on Learning Representations*, 2021.
- [28] Jenny Liu, Aviral Kumar, Jimmy Ba, Jamie Kiros, and Kevin Swersky. Graph normalizing flows. *Advances in Neural Information Processing Systems*, 2019.
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *International Conference on Learning Representations*, 2019.
- [30] Xuezhe Ma, Chunting Zhou, Xian Li, Graham Neubig, and Eduard Hovy. Flowseq: Non-autoregressive conditional sequence generation with generative flow. *arXiv preprint arXiv:1909.02480*, 2019.
- [31] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *International Conference on Learning Representations*, 2017.
- [32] Didrik Nielsen, Priyank Jaini, Emiel Hoogeboom, Ole Winther, and Max Welling. Survae flows: Surjections to bridge the gap between vaes and flows. *Advances in Neural Information Processing Systems*, 33, 2020.
- [33] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *arXiv preprint arXiv:1711.00937*, 2017.
- [34] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- [35] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Center for Research in Economics and Statistics Working Papers*, (2017-86), 2017.
- [36] Andres Potapczynski, Gabriel Loaiza-Ganem, and John P Cunningham. Invertible gaussian reparameterization: Revisiting the gumbel-softmax. *arXiv preprint arXiv:1912.09588*, 2019.
- [37] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [38] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, pp. 14866–14876, 2019.

- [39] Walter Rudin et al. *Principles of mathematical analysis*, volume 3. McGraw-hill New York, 1964.
- [40] Victor Garcia Satorras, Emiel Hoogeboom, Fabian Bernd Fuchs, Ingmar Posner, and Max Welling. E(n) equivariant normalizing flows. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021.
- [41] Yichuan Tang, Ruslan Salakhutdinov, and Geoffrey Hinton. Deep mixtures of factor analysers. *International Conference on Machine Learning*, 2012.
- [42] Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *International Conference on Learning Representations*, 2016.
- [43] Dustin Tran, Keyon Vafa, Kumar Agrawal, Laurent Dinh, and Ben Poole. Discrete flows: Invertible generative models of discrete data. *Advances in Neural Information Processing Systems*, 32:14719–14728, 2019.
- [44] Rianne van den Berg, Alexey A Gritsenko, Mostafa Dehghani, Casper Kaae Sønderby, and Tim Salimans. Idf++: Analyzing and improving integer discrete flows for lossless compression. In *International Conference on Learning Representations*, 2020.
- [45] Aäron Van Den Oord and Benjamin Schrauwen. Factoring variations in natural images with deep gaussian mixture models. In *Neural Information Processing Systems*, 2014.
- [46] Georges Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. premier mémoire. sur quelques propriétés des formes quadratiques positives parfaites. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1908(133): 97–102, 1908.
- [47] Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. *International Conference on Learning Representations*, 2019.
- [48] Zachary Ziegler and Alexander Rush. Latent normalizing flows for discrete sequences. In *International Conference on Machine Learning*, pp. 7673–7682. PMLR, 2019.

## A ADDITIONAL REMARKS AND PROPOSITIONS

**Box constraints.** There can be continuity problems if a Voronoi cell is unbounded, as the solution to Equation (8) does not exist if  $x(\lambda^*)$  diverges. Furthermore, when solving Equation (7) it can be difficult to numerically differentiate between an unbounded cell and one whose boundaries are very far away. It is for these reasons that we introduce box constraints (Equation 6) in the formulation of Voronoi cells which allows us to sidestep these issues for now.

**Solving for  $\lambda^*$ .** While Equation (7) can be solved using a convex optimization library, this approach is prone to numerical errors and requires differentiating through convex optimization solutions for gradient-based learning [1]. We instead note that the solution of Equation (7) can be expressed in closed form, since it is always going to be the intersection of the ray  $x(\lambda)$  and a linear constraint.

Let  $\mathbf{a}_i^\top \mathbf{x} = b_i$  be the plane that represents one of the linear constraints in Equation (6). Let  $\lambda_i$  be the intersection of this plane with the ray, *i.e.* it is the solution to  $\mathbf{a}_i^\top \mathbf{x}(\lambda_i) = b_i$ . This can be solved exactly as

$$\lambda_i = \frac{b_i - \mathbf{a}_i^\top \mathbf{x}_k}{\mathbf{a}_i^\top \delta_k(\mathbf{x})}. \quad (9)$$

Then the solution is simply the smallest positive  $\lambda_i$ , which satisfies all the linear constraints,  $\lambda^* = \min\{\lambda_i : \lambda_i > 0\}$ . There are a total of  $K + 2D - 1$  linear constraints, including the Voronoi cell boundaries and box constraints, which can be computed fully in parallel.

**Automatic differentiation.** Note that the above formulation of the solution  $\lambda^*$  also provides gradients  $\frac{d\lambda^*}{d\delta_k(\mathbf{x})}$  through automatic differentiation, since it only involves primitive operations. This allows end-to-end differentiation of the mapping  $f_k$ , and will allow us to learn the parameters of  $f_k$ , *i.e.* parameters of  $\alpha_k$  and the Voronoi tessellation.

**Proposition 1.** *The mapping  $f_k : \mathbb{R}^D \rightarrow V_k$  as defined in the 2-step procedure is a homeomorphism.*

That is to say,  $f_k$  is a bijection, and both  $f_k$  and  $f_k^{-1}$  are continuous. This allows us to use  $f_k$  within the normalizing flows framework, as a mapping between a distribution  $p_x$  defined on  $\mathbb{R}^D$  and the transformed distribution  $p_z$  on  $V_k$ , with change in density given by

$$p_z(f_k(x)) = p_x(x) \left| \det \frac{df_k(x)}{dx} \right|^{-1}. \quad (10)$$

**Proposition 2.** *If  $p_x(x)$  is continuous, then the transformed density  $p_z(f_k(x))$  is continuous almost everywhere.*

This comes from the Jacobian being continuous a.e.

### A.1 COMPUTING THE INVERSE MAPPING

We will also be using the inverse mapping  $f_k^{-1} : V_k \rightarrow \mathbb{R}^D$ , so we next describe how to compute this. Let  $\mathbf{z} = f_k(\mathbf{x})$ .

Conveniently, since both  $\mathbf{x}$  and  $\mathbf{z}$  lie on the ray  $\{\mathbf{x}(\lambda) : \lambda > 0\}$ , we know  $\delta_k(\mathbf{x}) = \delta_k(\mathbf{z})$ . So the first step is the same as the forward procedure: we solve for  $\lambda^*$  and  $\mathbf{x}(\lambda^*)$ . Following this, we then recover  $\mathbf{x}$  by inverting Step 2 of the forward procedure.

This inverse transformation is given by

$$\tilde{\alpha} = \frac{\mathbf{z} - \mathbf{x}_k}{\mathbf{x}(\lambda^*) - \mathbf{x}_k} \quad (11)$$

$$\tilde{\Delta} = \alpha_k^{-1}(\tilde{\alpha}_1) \quad (12)$$

$$\Delta = \tilde{\Delta} \|\mathbf{x}(\lambda^*) - \mathbf{x}_k\| \quad (13)$$

$$\mathbf{x} = \Delta \delta_k(\mathbf{z}) + \mathbf{x}_k \quad (14)$$



Equation (11) is an element-wise division. Since  $\tilde{\alpha}$  will be the same in all dimensions, we can simply pick a dimension in Equation (12). In our experiments, the inverse  $\alpha_k^{-1}$  can be computed analytically, though since it is just a scalar function, simple methods like bisection can also work when the inverse is not known analytically. Lastly, Equation (14) follows from the observation that  $\delta_k(x) = \delta_k(z)$ .

## A.2 EFFICIENT COMPUTATION OF THE LOG DET JACOBIAN

As  $f_k$  is a mapping in  $D$  dimensions, computing the log determinant Jacobian for likelihood computation can be costly and will scale poorly with  $D$  if computed naively. Instead, we note that the Jacobian of  $f_k$  is highly structured.

Intuitively, because  $f_k$  depends only on the direction  $\delta_k(x)$  and the distance away from  $x_k$ , it only has two degrees of freedom regardless of  $D$ . In fact, the Jacobian of  $f_k$  can be represented as a rank-2 update on a scaled identity matrix. This allows us to use the matrix determinant lemma to reformulate the log determinant Jacobian in a compute- and memory-efficient form. We summarize this in a proposition.

**Proposition 3.** *Let the transformation  $f_k(x)$  and all intermediate quantities be as defined in Section 3.1 for some given input  $x$ . Then the Jacobian factorizes as*

$$\frac{\partial f_k(x)}{\partial x} = cI + u_1 v_1^\top + u_2 v_2^\top \quad (15)$$

for some  $c \in \mathbb{R}$ ,  $u_i \in \mathbb{R}^D$ ,  $v_i \in \mathbb{R}^D$ , and its log determinant has the form

$$\begin{aligned} & \log \left| \det \frac{\partial f_k(x)}{\partial x} \right| \\ &= \log |1 + w_{11}| + \log \left| 1 + w_{22} - \frac{w_{12}w_{21}}{1 + w_{11}} \right| + D \log c \end{aligned} \quad (16)$$

where  $w_{ij} = c^{-1} v_i^\top u_j$ . This expression for the log determinant only requires dot products between vectors in  $\mathbb{R}^D$ . To reduce notational clutter, the exact formulas for  $c, u_1, u_2, v_1, v_2$  can be found in Appendix F.

**Remark on computational cost.** All vectors used in computing Equation (16) are either readily available as intermediate quantities after computing  $f_k(x)$ , or are gradients of scalar functions and can be efficiently computed through reverse-mode automatic differentiation. The only operations on these vectors are dot products, and no large  $D$ -by- $D$  matrices are ever constructed as part of any intermediate steps. Compared to explicitly constructing the Jacobian matrix, this is more efficient in both compute and memory and can readily scale to large values of  $D$ .

## B APPLICATION SETTINGS

We discuss two applications of our method to likelihood-based modeling of discrete and continuous data.

The first is a novel dequantization method that allows training a model of discrete data using density models that normally only act on continuous variables such as normalizing flows. Compared to existing dequantization methods [7, 36, 17], the Voronoi dequantization is not restricted to any fixed-dimensional continuous space and can benefit from learning similarities between discrete values. In fact, previous approaches can be seen as special cases of a Voronoi dequantization with fixed equidistant anchor points.

The second is a novel formulation of disjoint mixture models, where each Voronoi cell represents a single component in a large mixture model. To the best of our knowledge, disjoint mixture models have not been explored significantly in the literature and have not been successfully applied in more than a couple dimensions. Compared to an existing method [8], the Voronoi mixture model is not restricted to acting individually for each dimension and we empirically show can scale to thousands of dimensions.

## B.1 VORONOI DEQUANTIZATION

Let  $\mathbf{y}$  be a discrete random variable with finite support  $\mathcal{Y}$ . Then we can use Voronoi tessellation to define subsets for dequantizing  $\mathbf{y}$  as long as there are at least as many Voronoi cells—equivalently, anchors points—as the number of discrete values, *i.e.*  $K \geq |\mathcal{Y}|$ . By assigning each discrete value to a Voronoi cell, we can then define a dequantization model  $q(\mathbf{x}|\mathbf{y})$  by first sampling from a base (*e.g.* Gaussian) distribution  $\mathbf{z} \sim q(\mathbf{z}|\mathbf{y})$  in  $\mathbb{R}^D$  and then applying the mapping  $\mathbf{x} = f_k(\mathbf{z})$  from Section 3.1 to construct a distribution over  $V_k$ . We can then obtain probabilities  $q(\mathbf{x}|\mathbf{y})$  efficiently using Proposition 3 and train the dequantization alongside the density model  $p(\mathbf{x})$  on  $\mathbb{R}^D$ .

Automatic differentiation through  $p(\mathbf{x})$  and  $f_k$  (as discussed in Appendix A) provides gradients for training  $q(\mathbf{x}|\mathbf{y})$ , which include gradient with respect to parameters of the Voronoi tessellation. We are also free to choose the number of dimensions  $D$ , where a smaller  $D$  induces dequantized distributions that may be easier to fit while a larger  $D$  allows the anchor points and Voronoi cells more room to change over the course of training.

Sampling from the model  $p(\mathbf{y}|\mathbf{x})$  is straightforward and deterministic after sampling  $\mathbf{x}$ . We can write  $\mathbf{y} = g(\mathbf{x})$  where  $g$  is the set identification function satisfying  $\mathbf{x} \in V_{g(\mathbf{x})}$ . From the definition of Voronoi tessellation Equation (4), it is easy to see that  $g(\mathbf{x})$  is the nearest neighbor operation,

$$g(\mathbf{x}) = \arg \min_k \|\mathbf{x} - \mathbf{x}_k\|. \quad (17)$$

Depending on the position of anchor points, we can recover the disjoint subsets used by prior methods as special cases. We illustrate this in Figure 2.

## B.2 VORONOI MIXTURE MODELS

Let  $\{V_k\}$  be a Voronoi tessellation of  $\mathbb{R}^D$ . Then a disjoint mixture model can be constructed by defining distributions on each Voronoi cell. Here we make use of the inverse mapping  $f_k^{-1} : V_k \rightarrow \mathbb{R}^D$  so that we only need to parameterize distributions over  $\mathbb{R}^D$ . Let  $\mathbf{x}$  be a point in  $\mathbb{R}^D$ , our Voronoi mixture model assigns the density

$$\begin{aligned} \log p_{\text{mix}}(\mathbf{x}) &= \log p_{\text{comp}}(f_{k=g(\mathbf{x})}^{-1}(\mathbf{x})|k = g(\mathbf{x})) \\ &\quad + \log \left| \det \frac{\partial f_k^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right| \\ &\quad + \log p(k = g(\mathbf{x})) \end{aligned} \quad (18)$$

where  $p_{\text{comp}}$  can be any distribution over  $\mathbb{R}^D$ .

A Voronoi mixture model can also be viewed as a type of normalizing flow where, in addition to the change of variable due to  $f_k^{-1}$ , we also need to apply the change in density resulting from choosing one out of  $K$  components,  $p(k)$ . Per the observation of Dinh et al. [8], we have transformed  $\mathbf{x}$  into a tuple of a continuous variable  $\mathbf{z} = f_k^{-1}(\mathbf{x})$  and a discrete auxiliary variable  $k$ . A natural extension is to repeatedly apply this transformation, effectively forming a deep mixture model [41, 45], and resulting in a sequence of discrete auxiliary variables. This lets us define large mixture models hierarchically without the need to create a tessellation with a large number of Voronoi cells. Crucially, even though the number of mixture components increases exponentially with respect to depth, the computational cost does not scale with the number of components.

We also provide a simple proposition that the resulting density function is amenable to gradient descent.

**Proposition 4.** *For any distribution  $p(k)$  with support over  $\{1, \dots, K\}$ , define the mixture distribution,*

$$p(\mathbf{x}) = \sum_{k=1}^K p(\mathbf{x}|k)p(k), \quad (19)$$

where  $p(\cdot|k)$  is the distribution mapped onto the Voronoi cell  $V_k$  (*i.e.* Equation 10) from a density that is continuous. Then the density function of the mixture is continuous a.e.

Unlike Dinh et al. [8], we don't need to worry about the density being discontinuous for certain parameterizations.

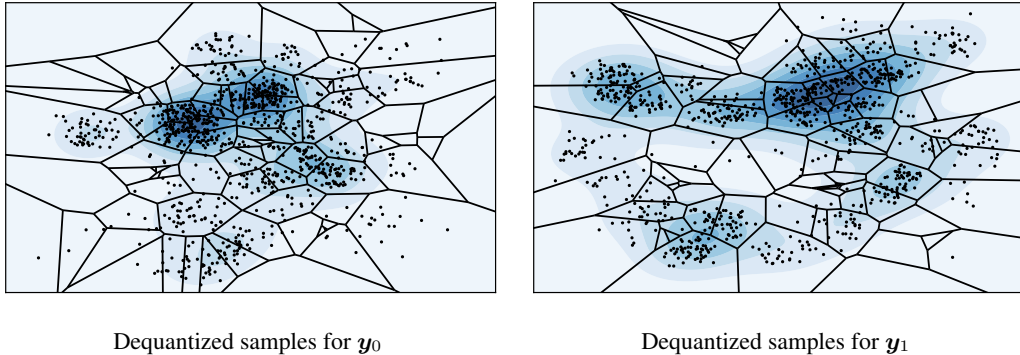


Figure 5: Dequantized samples from model trained on quantized `8gaussians` data set. Each plot shows samples from a single discrete variable, mapped onto a Voronoi tessellation of  $\mathbb{R}^2$ , along with a kernel density estimator. The model learns to cluster discrete values with similar probability values.

## C RELATED WORK

**Normalizing flows for discrete data.** Invertible mappings have been proposed for discrete data, where discrete values are effectively rearranged from a factorized distribution. In order to parameterize the transformation, Hoogeboom et al. [16], van den Berg et al. [44] use quantized ordinal transformations, while Tran et al. [43] takes a more general approach of doing modulo addition on one-hot vectors. These approaches suffer from gradient bias due to the need to use discontinuous operations and do not have universality guarantees since it’s unclear whether simple rearrangement is sufficient to transform any joint discrete distribution into a factorized distribution. In contrast, the dequantization approach provides a universality guarantee since the lower bound in Equation (2) is tight when  $p(\mathbf{x})\mathbb{1}_{[\mathbf{x}\in A_{\mathbf{y}}]} \propto q(\mathbf{x}|\mathbf{y})$ , with a proportionality equal to  $p(\mathbf{y})$ .

**Dequantization methods.** Within the normalizing flows literature, the act of adding noise was originally used for ordinal data as a way to combat numerical issues [7]. Later on, appropriate dequantization approaches have been shown to lower bound the log-likelihood of a discrete model [42, 15]. For non-ordinal data, many works have proposed simplex-based approaches. Early works on relaxations [19, 31] proposed continuous distribution on the simplex that mimic the behaviors of discrete random variables; however, these were only designed for the use with a Gumbel base distribution. Potapczynski et al. [36] extend this to a Gaussian distribution—although it is not hard to see this can work with any base distribution—by designing invertible transformations between  $\mathbb{R}^D$  and the probability simplex with  $K$  vertices, with  $D = K - 1$ , where  $K$  is the number of classes of a discrete random variable.

Intuitively, after fixing one of the logits, the softmax operation is an invertible transformation from  $\mathbb{R}^{K-1}$  to  $\{\mathbf{x} \in \mathbb{R}^K : \sum_{i=1}^K x_i = 1, x_i > 0\}$ . The  $(K - 1)$ -simplex can then be broken into  $K$  subsets, each corresponding to a particular discrete value.

$$A_k = \{\mathbf{x} \in \mathbb{R}^K : \sum_{i=1}^K x_i = 1, x_k > x_i \forall i \neq k\}. \quad (20)$$

More recently, Hoogeboom et al. [17] proposed ignoring the simplex constraint and simply use

$$A_k = \{\mathbf{x} \in \mathbb{R}^K : x_k > x_i \forall i \neq k\}, \quad (21)$$

which effectively increases the number of dimensions by one compared to the simplex approach. However, both simplex-based approaches force the dimension of the continuous space  $D$  to scale with  $K$ . In order to make simplex-based dequantization work when  $K$  is large, they propose reducing all discrete variables to a set of binary variables before applying dequantization. In contrast, Voronoi dequantization has full flexibility is choosing  $D$  regardless of  $K$ .

A number of works [30, 27] also proposed removing the constraint that subsets are disjoint, and instead work with general mixture models with unbounded support, relying on the conditional model  $p(\mathbf{y}|\mathbf{x})$  being sufficiently weak so that the task of modeling is forced onto a flow-based prior. Similar to general approaches that combine normalizing flows with variational inference approaches [48, 18, 32],

Method	Connect4	Forests	Mushroom	Nursery	PokerHands	USCensus90
Voronoi Deq.	12.92±0.07	14.20±0.05	9.06±0.05	9.27±0.04	19.86±0.04	24.19±0.12
Simplex Deq.	13.46±0.01	16.58±0.01	9.26±0.01	9.50±0.00	19.90±0.00	28.09±0.08
BinaryArgmax Deq.	13.71±0.04	16.73±0.17	9.53±0.01	9.49±0.00	19.90±0.01	27.23±0.02
Discrete Flow	19.80±0.01	21.91±0.01	22.06±0.01	9.53±0.01	19.82±0.03	55.62±0.35

Table 2: **Discrete UCI data sets.** Negative log-likelihood results on the test sets in nats.

they have achieved good performance on a number of tasks. However, they lose the computational savings and the deterministic decoder  $p(\mathbf{y}|\mathbf{x})$  gained from using disjoint subsets. On the other hand, quantization based on nearest neighbor have been for learning discrete latent variables [33, 38], but no likelihoods are constructed, the boundaries are not explicitly differentiated through, and the model relies on training with modified gradients based on heuristics.

**Disjoint mixture models.** The computational savings from using disjoint subsets was pointed out by Dinh et al. [8]. However, their method only works in each dimension individually. They transform samples using a linear spline, which is equivalent to creating subsets based on the knots of the spline and apply a linear transformation within each subset. Furthermore, certain parameterizations of the spline can lead to discontinuous density functions, whereas our disjoint mixture has a density function that is continuous almost everywhere. Overall, splines are interesting because each segment is independent, and the use of monotonic splines have been combined with normalizing flows [10]. However, the splines in Dinh et al. [8] are not enforced to be monotonic, so the full transformation is not bijective and acts like a disjoint mixture model. Ultimately, their experiments were restricted to two-dimensional synthetic data sets, and it remained an interesting research question whether disjoint mixture models can be successfully applied in high dimensions.

## D EXPERIMENTS

We experimentally validate our semi-discrete approach of combining Voronoi tessellation within likelihood-based modeling on a variety of data domains: discrete-valued UCI data sets, itemset modeling, language modeling, and disjoint mixture modeling. The goal of these experiments is not to show state-of-the-art results in these domains but to showcase the relative merit of our use of Voronoi tessellation compared to existing methods. For this reason, we just use basic coupling blocks with affine transformations [6, 7] as a base flow model for our experiments, unless stated otherwise. When comparing to Discrete Flows [43], we use their bipartite layers. Details regarding preprocessing for data sets can be found in Appendix G, and detailed experimental setup is in Appendix H.

### D.1 2D SYNTHETIC DATA

We start by considering synthetic distribution over two discrete variables, to compare against Discrete Flows [43]. Figure 3 shows the target probability mass functions, which are created by a 2D distribution where each dimension is quantized into 91 discrete values. Though the data is ordinal, no knowledge of ordering is given to the models. We see that Discrete Flows has trouble fitting these, presumably because it is difficult to rearrange the target distribution into a factorized distribution.

For our model, we dequantize each discrete variable with a Voronoi tessellation in  $\mathbb{R}^2$ . We then learn a flow model on the combined  $\mathbb{R}^4$ , parameterized by multilayer perceptrons (MLPs). In Figure 5 we visualize the learned Voronoi tessellation and samples from our model. The learned tessellation seems to group some of discrete values that occur frequently together, so the resulting model can have less modes.

### D.2 DISCRETE-VALUED UCI DATA

We experiment with complex data sets where each discrete variable can have a varying number of classes. Furthermore, these discrete variables may have hidden semantics. To this end, we use

Model	Dequantization	Retail	Accidents
Equivariant CNF	Voronoi	9.44 $\pm$ 2.34	7.81 $\pm$ 2.84
Equivariant CNF	Simplex	24.16 $\pm$ 0.21	19.19 $\pm$ 0.01
Equivariant CNF	BinaryArgmax	10.47 $\pm$ 0.42	6.72 $\pm$ 0.23
Determinantal Point Process		20.35 $\pm$ 0.05	15.78 $\pm$ 0.04

Table 3: **Permutation-invariant itemset modeling.** Negative log-likelihood results on the test sets in nats.

Dequantization	text8	enwik8
Voronoi ( $D=2$ )	1.39 $\pm$ 0.01	1.46 $\pm$ 0.01
Voronoi ( $D=4$ )	1.37 $\pm$ 0.00	1.41 $\pm$ 0.00
Voronoi ( $D=6$ )	1.37 $\pm$ 0.00	1.40 $\pm$ 0.00
Voronoi ( $D=8$ )	1.36 $\pm$ 0.00	1.39 $\pm$ 0.01
Argmax	1.38	1.42

Table 4: **Character-level language modeling.** Negative log-likelihood results on the test sets in bits per character.

unprocessed data sets from the UCI database [9]. The only pre-processing we perform is to remove variables that only have one discrete value. We then take 80% as train, 10% as validation, and 10% as the test set. Most of these datasets have a combination of both ordinal and non-ordinal variables, and we expect the non-ordinal variables to exhibit relations that are unknown (*e.g.* spatial correlations).

We see that Discrete Flows can be competitive with dequantization approaches, but can also fall short on difficult data sets such as the USCensus90, the largest data set we consider with 2.7 million examples and 68 different discrete variables of varying types. For dequantization methods, simplex [36] and binary argmax [17] approaches are mostly on par. We do see a non-negligible gap in performance between these baselines and Voronoi dequantization for most of the data sets, likely due to the ability to learn semantically useful relations between the values of each discrete variable. For instance, the Connect4 dataset contains positions of a two-player board game with the same name, which exhibit complex spatial dependencies between the board pieces, and the USCensus90 dataset contains highly correlated variables and is often used to test clustering algorithms.

### D.3 PERMUTATION-INVARIANT ITEMSET MODELING

A major appeal of using normalizing flows in continuous space is the ability to incorporate invariance to specific group symmetries into the density model. For instance, this can be done by ensuring the ordinary differential equation is equivariant [24, 2, 40] in a continuous normalizing flow [5, 12] framework. Here we focus on invariance with respect to permutations, *i.e.* sets of discrete variables. This invariance cannot be explicitly modeled by Discrete Flows as they require an ordering or a bipartition of discrete variables.

We preprocessed a data set of retail market basket data from an anonymous Belgian retail store [3] and a data set of anonymized traffic accident data [11], which contain sets of discrete variables each with 765 and 269 values, respectively for each data set. Given the large number of discrete values, we expected Voronoi dequantization to perform better. However, we see that in Table 4, binary argmax is quite competitive, likely because it decomposes the full distribution using multiple binary variables. We did not do this binarization trick for simplex dequantization, and we see that it suffers from having to use large embedding dimensions, performing worse than a determinantal point process baseline [25].

#### D.4 LANGUAGE MODELING

As a widely used benchmark for discrete models [43, 27, 17], we also experiment with our method on language modeling. Here we used the open source code provided by Hooeboom et al. [17] with the exact same autoregressive flow model and optimizer setups. The only difference is replacing their binary argmax with our Voronoi dequantization. Results are shown in Table 4, where we tried out multiple embedding dimensions  $D$ . Generally, we find  $D = 2$  to be too low and can stagnate training since the Voronoi cells are more constrained. The flexibility in choosing  $D$  gives us a slight edge in performance.

## E CONCLUSION AND DISCUSSION

We combine Voronoi tessellation with normalizing flows to construct a new invertible transformation that has learnable discrete structure. This acts as a learnable mapping between discrete and continuous distributions. We propose two applications of this method: a Voronoi dequantization that maps discrete values into a learnable convex polytope, and a Voronoi mixture modeling approach that has the same compute cost as a single component. We showcased the relative merit of our approach across a range of data modalities with varying structure.

### E.1 LIMITATIONS AND FUTURE DIRECTIONS

**Diminishing density on boundaries.** The distributions within each Voronoi cell necessarily go to zero on the boundary between cells due to the use of a homeomorphism from an unbounded domain. This is a different and undesirable property as opposed to the partitioning method used by [8]. However, alternatives could require more compute cost in the form of solving normalization constants. Balancing expressiveness and compute cost is a delicate problem and sits at the core of tractable probabilistic modeling.

**Design of homeomorphisms and tessellations.** Our proposed homeomorphism is simple and scalable, but this comes at the cost of smoothness and expressiveness. As depicted in Figure 1, the transformed density has a non-smooth probability density function. This non-smoothness exists when the ray intersections with multiple boundaries. This may cause optimization to be difficult as the gradients of the log-likelihood objective can be discontinuous. Additionally, our use of Euclidean distance can become problematic in high dimensions, as this metric can cause almost all points to be nearly equidistant, resulting in a behavior where all points lie seemingly very close to the boundary. Improvements on the design of homeomorphisms to bounded domains could help alleviate these problems. Additionally, more flexible tessellations—such as the Laguerre tessellation—and additional concepts from semi-discrete optimal transport [35, 26, 13] may be adapted to improve semi-discrete normalizing flows.

## F PROOFS

**Proposition 3.** *Let the transformation  $f_k(x)$  and all intermediate quantities be as defined in Section 3.1 for some given input  $x$ . Then the Jacobian factorizes as*

$$\frac{\partial f_k(x)}{\partial x} = cI + u_1 v_1^\top + u_2 v_2^\top \quad (15)$$

for some  $c \in \mathbb{R}$ ,  $u_i \in \mathbb{R}^D$ ,  $v_i \in \mathbb{R}^D$ , and its log determinant has the form

$$\begin{aligned} & \log \left| \det \frac{\partial f_k(x)}{\partial x} \right| \\ &= \log |1 + w_{11}| + \log \left| 1 + w_{22} - \frac{w_{12} w_{21}}{1 + w_{11}} \right| + D \log c \end{aligned} \quad (16)$$

where

$$\Delta \triangleq \|x - x_k\| \quad (22)$$

$$c \triangleq \alpha_k(\tilde{\Delta})\lambda^* \Delta^{-1} \quad (23)$$

$$u_1 \triangleq \left[ \alpha_k(\tilde{\Delta})\Delta^{-1} - \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right] \delta_k(x) \quad (24)$$

$$v_1 \triangleq \frac{\partial \lambda^*}{\partial \delta_k(x)} \quad (25)$$

$$u_2 \triangleq \left[ \left( \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right) - \alpha_k(\tilde{\Delta})\lambda^* \Delta^{-1} - \alpha_k(\tilde{\Delta})\Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \delta_k(x)} \right)^\top \delta_k(x) \right) \right] \delta_k(x) \quad (26)$$

$$+ \left( \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right) \left( \left( \frac{\partial \lambda^*}{\partial \delta_k(x)} \right)^\top \delta_k(x) \right) \delta_k(x) \quad (27)$$

$$v_2 \triangleq \delta_k(x) \quad (28)$$

$$w_{ij} \triangleq c^{-1} v_i^\top u_j, \text{ for } i, j \in \{1, 2\} \quad (29)$$

*Proof.* We first write the Jacobian of  $f_k$  in the form of  $cI + u_1 v_1^\top + u_2 v_2^\top$  where  $c$  is a scalar, and  $u_1, u_2, v_1, v_2$  are vectors of size  $D$ . Define the shorthand  $\Delta \triangleq \|x - x_k\|$ ,  $\Delta^* \triangleq \|x(\lambda^*) - x_k\|$ , and  $\tilde{\Delta} \triangleq \Delta/\Delta^*$ . To simplify notation, we use the short-hands  $\alpha_k = \alpha_k(\tilde{\Delta})$ ,  $\delta_k = \delta_k(x)$ . Then the Jacobian follows

$$\frac{\partial f_k(x)}{\partial x} = \frac{\partial}{\partial x} (x_k + \alpha_k(x(\lambda^*) - x_k)) \quad (30)$$

$$= (x(\lambda^*) - x_k) \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \frac{\partial \tilde{\Delta}}{\partial x} \right)^\top + \alpha_k \left( \frac{\partial x(\lambda^*)}{\partial x} \right) \quad (31)$$

$$= (x(\lambda^*) - x_k) \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \frac{1}{\Delta^*} \delta_k^\top - \frac{\Delta}{(\Delta^*)^2} (x(\lambda^*) - x_k)^\top \frac{\partial x(\lambda^*)}{\partial x} \right) + \alpha_k \left( \frac{\partial x(\lambda^*)}{\partial x} \right) \quad (32)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \Delta \delta_k \delta_k^\top \frac{\partial x(\lambda^*)}{\partial x} + \alpha_k \left( \frac{\partial x(\lambda^*)}{\partial x} \right) \quad (33)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top + \left( \alpha_k - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \Delta \delta_k \delta_k^\top \right) \left( \frac{\partial x(\lambda^*)}{\partial x} \right) \quad (34)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top + \left( \alpha_k - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \Delta \delta_k \delta_k^\top \right) \left( \lambda^* \Delta^{-1} I - \lambda^* \Delta^{-1} \delta_k \delta_k^\top + \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \left( \frac{\partial \delta_k}{\partial x} \right) \right) \quad (35)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top + \left( \alpha_k - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \Delta \delta_k \delta_k^\top \right) \left( \lambda^* \Delta^{-1} I - \lambda^* \Delta^{-1} \delta_k \delta_k^\top + \Delta^{-1} \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \left( I - \delta_k \delta_k^\top \right) \right) \quad (36)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top + \left( \alpha_k - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \Delta \delta_k \delta_k^\top \right) \left( \lambda^* \Delta^{-1} I - \lambda^* \Delta^{-1} \delta_k \delta_k^\top + \Delta^{-1} \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top - \Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) \delta_k \delta_k^\top \right) \quad (37)$$

$$= \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \delta_k \delta_k^\top + \alpha_k \lambda^* \Delta^{-1} I - \alpha_k \lambda^* \Delta^{-1} \delta_k \delta_k^\top + \alpha_k \Delta^{-1} \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top - \alpha_k \Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) \delta_k \delta_k^\top \quad (38)$$

$$- \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \lambda^* \delta_k \delta_k^\top + \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \lambda^* \left( \delta_k^\top \delta_k \right) \delta_k \delta_k^\top - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \delta_k^\top \delta_k \right) \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \quad (39)$$

$$+ \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \delta_k^\top \delta_k \right) \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) \delta_k \delta_k^\top \quad (40)$$

$$= \alpha_k \lambda^* \Delta^{-1} I + \left[ \alpha_k \Delta^{-1} - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \delta_k^\top \delta_k \right) \right] \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \quad (41)$$

$$+ \left[ \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) - \alpha_k \lambda^* \Delta^{-1} - \alpha_k \Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) - \lambda^* \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \delta_k^\top \delta_k - 1 \right) + \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \delta_k^\top \delta_k \right) \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) \right] \delta_k \delta_k^\top \quad (42)$$

$$= \alpha_k \lambda^* \Delta^{-1} I + \left[ \alpha_k \Delta^{-1} - \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \right] \delta_k \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \quad (43)$$

$$+ \left[ \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) - \alpha_k \lambda^* \Delta^{-1} - \alpha_k \Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) + \left( \frac{\partial \alpha_k}{\partial \tilde{\Delta}} \right) \left( \left( \frac{\partial \lambda^*}{\partial \delta_k} \right)^\top \delta_k \right) \right] \delta_k \delta_k^\top \quad (44)$$



Now that  $\frac{\partial f_k(x)}{\partial x}$  is in the form of  $cI + u_1v_1^\top + u_2v_2^\top$ , where

$$c \triangleq \alpha_k(\tilde{\Delta})\lambda^*\Delta^{-1} \quad (45)$$

$$u_1 \triangleq \left[ \alpha_k(\tilde{\Delta})\Delta^{-1} - \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right] \boldsymbol{\delta}_k(\mathbf{x}) \quad (46)$$

$$v_1 \triangleq \frac{\partial \lambda^*}{\partial \boldsymbol{\delta}_k(\mathbf{x})} \quad (47)$$

$$u_2 \triangleq \left[ \left( \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right) - \alpha_k(\tilde{\Delta})\lambda^*\Delta^{-1} - \alpha_k(\tilde{\Delta})\Delta^{-1} \left( \left( \frac{\partial \lambda^*}{\partial \boldsymbol{\delta}_k(\mathbf{x})} \right)^\top \boldsymbol{\delta}_k(\mathbf{x}) \right) + \left( \frac{\partial \alpha_k(\tilde{\Delta})}{\partial \tilde{\Delta}} \right) \left( \left( \frac{\partial \lambda^*}{\partial \boldsymbol{\delta}_k(\mathbf{x})} \right)^\top \boldsymbol{\delta}_k(\mathbf{x}) \right) \right] \boldsymbol{\delta}_k(\mathbf{x}) \quad (48)$$

$$v_2 \triangleq \boldsymbol{\delta}_k(\mathbf{x}) \quad (49)$$

Applying the matrix determinant lemma twice, we can show that

$$\det(cI + u_1v_1^\top + u_2v_2^\top) = (1 + v_2^\top(cI + u_1v_1^\top)^{-1}u_2) \det(cI + u_1v_1^\top) \quad (50)$$

$$= \left[ \left( 1 + c^{-1}v_2^\top u_2 - \frac{c^{-2}}{1 + c^{-1}v_1^\top u_1} \right) v_2^\top u_1 v_1^\top u_2 \right] (1 + c^{-1}v_1^\top u_1) c^D \quad (51)$$

We simplify this by defining the scaled dot products,

$$w_{ij} \triangleq c^{-1}v_i^\top u_j, \text{ for } i, j \in \{1, 2\}. \quad (52)$$

Then

$$\log \left| \det \frac{\partial f_k(x)}{\partial x} \right| = \log |1 + w_{11}| + \log \left| 1 + w_{22} - \frac{w_{12}w_{21}}{1 + w_{11}} \right| + D \log c \quad (53)$$

□

Intermediate steps above used the following gradient identities.

$$\begin{aligned} \frac{\partial \boldsymbol{\delta}_k(\mathbf{x})}{\partial x} &= \frac{\partial}{\partial x} (x - x_k) \Delta^{-1} \\ &= \Delta^{-1} I + (x - x_k) \left( \frac{\partial}{\partial x} (\Delta^2)^{-\frac{1}{2}} \right)^\top \\ &= \Delta^{-1} I + (x - x_k) \left( \frac{-1}{2\Delta^3} \right) \left( \frac{\partial \Delta^2}{\partial x} \right)^\top \\ &= \Delta^{-1} I + (x - x_k) \left( \frac{-1}{2\Delta^3} \right) 2(x - x_k)^\top \\ &= \Delta^{-1} (I - \boldsymbol{\delta}_k(\mathbf{x}) \boldsymbol{\delta}_k(\mathbf{x})^\top) \end{aligned} \quad (54)$$

$$\frac{\partial \Delta}{\partial x} = \frac{\partial \Delta}{\partial x} = \frac{\partial}{\partial x} (\Delta^2)^{\frac{1}{2}} = \frac{1}{\Delta} (x - x_k) = \boldsymbol{\delta}_k(\mathbf{x}) \quad (55)$$

**Log determinant of  $\frac{\partial f_k^{-1}(z)}{\partial z}$ .** We can also use Proposition 3 to compute the log determinant of the inverse transform without needing to recompute  $f_k(x)$ . The only difference is a sign:  $\log \left| \det \frac{\partial f_k^{-1}(z)}{\partial z} \right| = -\log \left| \det \frac{\partial f_k(x)}{\partial x} \right|$ . The required quantities,  $\Delta$ ,  $\mathbf{x}(\lambda^*)$ , and  $\boldsymbol{\delta}_k(\mathbf{x})$ , are readily available after computing  $x = f_k^{-1}(z)$ . The gradients with respect to quantities of  $x$  can be expressed using gradients with respect to quantities of  $z$ ,

$$\frac{\partial \alpha_k(\Delta)}{\partial \Delta} = \left( \frac{\partial \alpha_k^{-1}(\hat{\alpha})}{\partial \hat{\alpha}} \right)^{-1} \quad \text{and} \quad \frac{\partial \lambda^*}{\partial \boldsymbol{\delta}_k(\mathbf{x})} = \frac{\partial \lambda^*}{\partial \boldsymbol{\delta}_k(z)}, \quad (56)$$

which are accessible through automatic differentiation.

**Proposition 1.** *The mapping  $f_k : \mathbb{R}^D \rightarrow V_k$  as defined in the 2-step procedure is a homeomorphism.*



*Proof.* Let  $\mathbf{x} \in \mathbb{R}^D$  and  $\mathbf{x}_k$  be a given anchor point corresponding to a Voronoi cell  $V_k$ . If  $\mathbf{x} \neq \mathbf{x}_k$ , then  $\mathbf{x}$  is uniquely represented by the tuple  $(\Delta, \delta)$  where  $\Delta = \|\mathbf{x} - \mathbf{x}_k\|$  and  $\delta = \frac{\mathbf{x} - \mathbf{x}_k}{\|\mathbf{x} - \mathbf{x}_k\|}$ , since  $\mathbf{x} = \mathbf{x}_k + \Delta\delta$ . Since  $\delta$  uniquely defines the ray  $\{\mathbf{x}_k + \lambda\delta; \lambda > 0\}$  and  $\mathbf{x}(\lambda^*)$ , then because  $\alpha_k$  in Equation (8) is a bijection in  $\Delta$ ,  $f_k$  is a bijection. For  $\mathbf{x} \neq \mathbf{x}_k$ , the continuity of  $f_k$  follows from Rudin et al. [39, Theorem 4.7] since  $\Delta$  and  $\mathbf{x}(\lambda^*)$  are continuous in  $\mathbf{x}$  and  $\alpha_k$  is continuous in  $\Delta$ . Then since  $f_k(\mathbf{x}) \rightarrow \mathbf{x}_k$  as  $\mathbf{x} \rightarrow \mathbf{x}_k$  from all directions, this justifies the choice of setting  $f_k(\mathbf{x}) = \mathbf{x}_k$  when  $\mathbf{x} = \mathbf{x}_k$ . Finally, by the invariance of domain theorem, since  $V_k$  is an open set in  $\mathbb{R}^D$ ,  $f_k$  is an open map and the inverse  $f_k^{-1}$  is continuous, and we can conclude  $f$  is a homeomorphism between  $\mathbb{R}^D$  and  $V_k$ .  $\square$

**Proposition 2.** *If  $p_x(x)$  is continuous, then the transformed density  $p_z(f_k(x))$  is continuous almost everywhere.*

*Proof.* See proof of Proposition 3 below for the form of the Jacobian. For the case where  $\mathbf{x} \neq \mathbf{x}_k$ , all quantities in Equations (22) to (29) are continuous with respect to  $\mathbf{x}$ . Hence the Jacobian  $\frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}}$  is continuous, and since it is always full-rank, then the composition  $\left| \det \frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}} \right|$  is continuous [39, Theorem 4.7] and so is the product  $p_x(x) \left| \det \frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}} \right|$  [39, Theorem 4.9].  $\square$

**Proposition 4.** *For any distribution  $p(k)$  with support over  $\{1, \dots, K\}$ , define the mixture distribution,*

$$p(\mathbf{x}) = \sum_{k=1}^K p(\mathbf{x}|k)p(k), \quad (19)$$

where  $p(\cdot|k)$  is the distribution mapped onto the Voronoi cell  $V_k$  (i.e. Equation 10) from a density that is continuous. Then the density function of the mixture is continuous a.e.

*Proof.* Proposition 2 ensures the distribution is continuous a.e. within each Voronoi cell. It’s straightforward to see that the density of  $p_z(f_k(x))$  approaches zero as  $f_k(x)$  approaches the boundaries of  $V_k$  for any properly normalized distribution  $p_x$ .  $\square$

## G DATA SETS

### G.1 UCI DATA SETS

The main preprocessing we did was to (i) remove the “label” attribute from each data set, and (ii) remove attributes that only ever take on one value. Apart from this, the USCensus90 dataset contains a unique identifier for each row, which was removed. Descriptions for all dataset are below.

**Connect4** [webpage] This dataset contains all legal 8-ply positions in the game of Connect Four in which neither player has won yet, and in which the next move is not forced. The original task was to predict which player would win, which has been removed during preprocessing. There are a total **42** discrete variables (one for each location on the board), each with **3** possible discrete values (taken by player 1, taken by player 2, blank). This data set was randomly split into 54045 training examples, 6755 validation examples, and 6757 test examples.

**Forests** [webpage] This dataset contains cartographic variables regarding forests including four wilderness areas located in the Roosevelt National Forest of northern Colorado. These areas represent forests with minimal human-caused disturbances. The original task was to predict the forest cover type, which has been removed during preprocessing. There are a total of **54** discrete variables, with **10** being the highest number of discrete values. This data set was randomly split into 464809 training examples, 58101 validation examples, and 58102 test examples.

**Mushroom** [webpage] This data set includes descriptions of hypothetical samples corresponding to 23 species of gilled mushrooms in the Agaricus and Lepiota Family. The original task was to predict whether each species is edible, which has been removed during preprocessing. There are a total of **21** discrete variables, with **12** being the highest number of discrete values. This data set was randomly split into 6499 training examples, 812 validation examples, and 813 test examples.

**Nursery** [webpage] This data set contains attributes of applicants to nursery schools, during a period when there was excessive enrollment to these schools in Ljubljana, Slovenia, and the rejected applications frequently needed an objective explanation. All data have been completely anonymized. The original task was to predict whether an applicant would be recommended for acceptance by hierarchical decision model, which has been removed during preprocessing. There are a total of **8** discrete variables, with **5** being the highest number of discrete values. This data set was randomly split into 10367 training examples, 1296 validation examples, and 1297 test examples.

**PokerHands** [webpage] This data set contains poker hands consisting of five playing cards drawn from a standard deck of 52. Each card is described using two attributes (suit and rank), for a total of 10 predictive attributes. There is one Class attribute that describes the “poker hand”. The original task was to predict the poker hand class (pairs, full house, royal flush, etc.), which has been removed during preprocessing. There are a total of **10** discrete variables, with **13** being the highest number of discrete values. This data set was randomly split into 820008 training examples, 102501 validation examples, and 102501 test examples.

**USCensus90** [webpage] This data set contains a portion of the data collected as part of the 1990 census in the United States, with the data completely anonymized. There are a total of **68** discrete variables, with **18** being the highest number of discrete values. This data set was randomly split into 2212456 training examples, 122914 validation examples, and 122915 test examples.

## G.2 ITEMSET DATA SETS

These data sets were taken from the Frequent Itemset Mining Dataset Repository [webpage]. Each row is interpreted as a set of items with no emphasis on the ordering of items.

**Retail** [3] This data set contains anonymized retail market basket data from an anonymous Belgian retail store. We first removed rows with less than 4 items, then randomly sampled a subset of 4 items for every row. Items that appear in less 300 rows were dropped from the data set. The final preprocessed data set contains **765** distinct items. This data set was randomly split into 24280 training examples, 3035 validation examples, and 3036 test examples.

**Accidents** [11] This data set contains contains anonymized traffic accident data. Data on traffic accidents are obtained from the National Institute of Statistics (NIS) for the region of Flanders (Belgium) for the period 1991-2000. We first removed rows with less than 4 items, then randomly sampled a subset of 4 items for every row. This subsampling occurred 10 times if a row has 16 or more items, 5 times if the row has 8 to 15 items, and once if the row has 4-7 items. Items that appear in less 300 rows were dropped from the data set. The final preprocessed data set contains **213** distinct items. This data set was randomly split into 270129 training examples, 33766 validation examples, and 33767 test examples.

## H EXPERIMENTAL DETAILS

**2D synthetic data sets** Continuous data sets were quantized into 91 bins for each coordinate. For Voronoi dequantization, we dequantized each coordinate into an embedding space of 2 dimensions, with 91 Voronoi cells. The dequantization model is parameterized by 4 layers of coupling blocks, each with a 2 hidden layer MLP with 256 hidden units each, where the Swish activation function was used [37]. The flow model is similarly parameterized but with 16 layers of coupling blocks. Each block alternated between 4 different partitioning schemes: masking out the first half, masking out the second half, masking out the odd indices, and masking out the even indices. We trained with the Adam optimizer [21] with a learning rate of  $1e-3$ .

**UCI data sets** For Voronoi dequantization, we dequantized each coordinate into an embedding space of 4 or 6 dimensions, with the number of Voronoi cells set to the highest number of discrete values over all discrete variables. The dequantization model is parameterized by 4 layers of coupling blocks, each with a 2 hidden layer MLP with 256, 512, or 1024 hidden units each, where the Swish activation function was used [37]. The flow model is similarly parameterized but with 16 or 32 layers of coupling blocks. Each block alternated between 4 different partitioning schemes: masking out the first half, masking out the second half, masking out the odd indices, and masking out the even indices. We trained with the Adam optimizer [21] with a learning rate sweep over  $\{1e-3, 5e-4, 1e-4\}$ .

**Itemset data sets** For Voronoi dequantization, we dequantized each coordinate into an embedding space of 6 dimensions, with the number of Voronoi cells set to the number of items in the data set. We used a continuous normalizing flow (CNF) with the ordinary differential equation (ODE) defined using a Transformer architecture and a  $L_2$ -distance based multihead attention layer [20] and the GeLU activation function [14]. No positional embeddings were provided to the model to ensure the model is equivariant to permutations. We composed 12 CNF layers, each defined using a Transformer model that has 2 or 3 layers of alternating multihead attention and fully connected residual connections. To solve the ODE and train our model, we used the `dopri5` solver from the `torchdiffeq` library [4] with `atol=rtol=1e-5`. We trained with the AdamW optimizer [29, 47] with a learning rate of  $1e-3$  and weight decay of  $1e-6$ . For the Voronoi dequantization, we set  $D=6$  for both data sets, though it may be possible to improve performance by tuning  $D$ .

**Character-level language modeling** We used the provided hyperparameters from the open source repository [URL]. Some parts of code had to be adapted for our usage, but model architecture and optimizer remained the same.