
Differentiable Random Partition Models

Anonymous Author(s)

Affiliation

Address

email

340 A Preliminaries

341 A.1 Hypergeometric Distribution

342 This part is largely based on Sutter et al. [2023].

343 Suppose we have an urn with marbles in different colors. Let $K \in \mathbb{N}$ be the number of different classes
 344 or groups (e.g. marble colors in the urn), $\mathbf{m} = [m_1, \dots, m_K] \in \mathbb{N}^K$ describe the number of elements
 345 per class (e.g. marbles per color), $N = \sum_{k=1}^K m_k$ be the total number of elements (e.g. all marbles in
 346 the urn) and $n \in \{0, \dots, N\}$ be the number of elements (e.g. marbles) to draw. Then, the multivariate
 347 hypergeometric distribution describes the probability of drawing $\mathbf{n} = [n_1, \dots, n_K] \in \mathbb{N}^K$ marbles
 348 by sampling without replacement such that $\sum_{k=1}^K n_k = n$, where n_k is the number of drawn marbles
 349 of class k .

350 In the literature, two different versions of the noncentral hypergeometric distribution exist, Fisher's
 351 [Fisher, 1935] and Wallenius' [Wallenius, 1963, Chesson, 1976] distribution. Sutter et al. [2023]
 352 restrict themselves to Fisher's noncentral hypergeometric distribution due to limitations of the latter
 353 [Fog, 2008]. Hence, we will also talk solely about Fisher's noncentral hypergeometric distribution.

354 **Definition A.1** (Multivariate Fisher's Noncentral Hypergeometric Distribution [Fisher, 1935]). *A*
 355 *random vector \mathbf{X} follows Fisher's noncentral multivariate distribution, if its joint probability mass*
 356 *function is given by*

$$P(\mathbf{N} = \mathbf{n}; \boldsymbol{\omega}) = p(\mathbf{n}; \boldsymbol{\omega}) = \frac{1}{P_0} \prod_{k=1}^K \binom{m_k}{n_k} \omega_k^{n_k} \quad (12)$$

$$\text{where } P_0 = \sum_{(\eta_1, \dots, \eta_K) \in \mathcal{S}} \prod_{k=1}^K \binom{m_k}{\eta_k} \omega_k^{\eta_k} \quad (13)$$

357 *The support S of the PMF is given by $S = \{\mathbf{n} \in \mathbb{N}^K : \forall k \ n_k \leq m_k, \sum_{k=1}^K n_k = n\}$ and*
 358 $\binom{n}{k} = \frac{n!}{k!(n-k)!}$.

359 The class importance $\boldsymbol{\omega}$ is a crucial modeling parameter in applying the noncentral hypergeometric
 360 distribution (see [Chesson, 1976]).

361 A.1.1 Differentiable MVHG

362 Their reparameterizable sampling for the differentiable MVHG consists of three parts:

- 363 1. Reformulate the multivariate distribution as a sequence of interdependent and conditional
 364 univariate hypergeometric distributions.
- 365 2. Calculate the probability mass function of the respective univariate distributions.
- 366 3. Sample from the conditional distributions utilizing the Gumbel-Softmax trick.

367 Following the chain rule of probability, the MVHG distribution allows for sequential sampling over
 368 classes k . Every step includes a merging operation, which leads to biased samples compared to
 369 groundtruth non-differentiable sampling with equal class weights $\boldsymbol{\omega}$. Given that we intend to use
 370 the differentiable MVHG in settings where we want to learn the unknown class weights, we do not
 371 expect a negative effect from this sampling procedure. For details on how to merge the MVHG into a
 372 sequence of unimodal distributions, we refer to Sutter et al. [2023].

373 The probability mass function calculation is based on unnormalized log-weights, which are interpreted
 374 as unnormalized log-weights of a categorical distribution. The interpretation of the class-conditional
 375 unimodal hypergeometric distributions as categorical distributions allows applying the Gumbel-
 376 Softmax trick [Jang et al., 2016, Maddison et al., 2017]. Following the use of the Gumbel-Softmax
 377 trick, the class-conditional version of the hypergeometric distribution is differentiable and reparam-
 378 eterizable. Hence, the MVHG has been made differentiable and reparameterizable as well. Again, for
 379 details we refer to the original paper [Sutter et al., 2023].

380 A.2 Distribution over Random Orderings

381 Yellott [1977] show that the distribution over permutation matrices $p(\pi; \mathbf{s})$ follows a Plackett-
 382 Luce (PL) distribution [Plackett, 1975, Luce, 1959], if and only if the perturbed scores $\tilde{\mathbf{s}}$ are
 383 sampled independently from Gumbel distributions with identical scales. For each item i , sample
 384 $g_i \sim \text{Gumbel}(0, \beta)$ independently with zero mean and fixed scale β . Let $\tilde{\mathbf{s}}$ be the vector of
 385 Gumbel perturbed log-weights such that $\tilde{s}_i = \beta \log s_i + g_i$. Hence,

$$q(\tilde{s}_1 \geq \dots \geq \tilde{s}_n) = \frac{s_1}{Z} \cdot \frac{s_2}{Z - s_1} \cdot \dots \cdot \frac{s_n}{Z - \sum_{i=1}^{n-1} s_i} \quad (14)$$

386 We refer to Yellott [1977] or Grover et al. [2019] for the proof. However, Grover et al. [2019] provide
 387 only an adapted proof sketch from Yellott [1977]. The probability of sampling element i first is given
 388 by its score s_i divided by the sum of all weights in the set

$$q(\tilde{s}_i) = \frac{s_i}{Z} \quad (15)$$

389 For $z_i = \log s_i$, the right hand side of Equation (15) is equal to the softmax distribution $\text{softmax}(z_i) =$
 390 $\exp(z_i) / \sum_j \exp(z_j)$ as already described in [Xie and Ermon, 2019]. Hence, Equation (15) directly
 391 leads to the Gumbel-Softmax trick [Jang et al., 2016, Maddison et al., 2017].

392 A.2.1 Differentiable Sorting

393 In the main text of the paper we rely on a differentiable function $f_\pi(\tilde{\mathbf{s}})$, which sorts the resampled
 394 version of the scores \mathbf{s}

$$\pi = f_\pi(\tilde{\mathbf{s}}) = \text{sort}(\tilde{\mathbf{s}}) \quad (16)$$

395 Here, we summarise the findings from Grover et al. [2019] on how to construct such a differentiable
 396 sorting operator. As already mentioned in Section 2, there are multiple works on the topic [Prillo and
 397 Eisenschlos, 2020, Petersen et al., 2021, Mena et al., 2018], but we restrict ourselves to the work of
 398 Grover et al. [2019] as we see the differentiable generation of permutation matrices as a tool in our
 399 pipeline.

400 **Corollary A.2** (Permutation Matrix [Grover et al., 2019]). *Let $\mathbf{s} = [s_1, \dots, s_n]^T$ be a real-valued*
 401 *vector of length n . Let A_s denote the matrix of absolute pairwise differences of the elements of \mathbf{s}*
 402 *such that $A_s[i, j] = |s_i - s_j|$. The permutation matrix π corresponding to $\text{sort}(\mathbf{s})$ is given by:*

$$\pi = \begin{cases} 1 & \text{if } j = \arg \max[(n+1-2i)\mathbf{s} - A_s \mathbb{1}] \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

403 where $\mathbb{1}$ denotes the column vector of all ones.

404 As we know, the $\arg \max$ operator is non-differentiable which prohibits the direct use of Corollary A.2
 405 for gradient computation. Hence, Grover et al. [2019] propose to replace the $\arg \max$ operator with
 406 softmax to obtain a continuous relaxation $\pi(\tau)$ similar to the GS trick [Jang et al., 2016, Maddison
 407 et al., 2017]. In particular, the i th row of $\pi(\tau)$ is given by:

$$\pi(\tau)[i, :] = \text{softmax}[(n+1-2i)\mathbf{s} - A_s \mathbb{1} / \tau] \quad (18)$$

408 where $\tau > 0$ is a temperature parameter. We adapted this section from Grover et al. [2019] and
 409 we also refer to their original work for more details on how to generate differentiable permutation
 410 matrices.

411 In this, work we remove the temperature parameter τ to reduce clutter in the notation. Hence, we
 412 only write π instead of $\pi(\tau)$, although it is still needed for the generation of the matrix π . For details
 413 on how we select the temperature parameter τ in our experiments, we refer to Appendix C.

414 B Detailed Derivation of the Differentiable Two-Stage Random Partition 415 Model

416 B.1 Two-Stage Partition Model

417 We want to partition n elements $[n] = \{1, \dots, n\}$ into K subsets $\{\mathcal{S}_1, \dots, \mathcal{S}_K\}$ where K is *a priori*
 418 unknown.

419 **Definition B.1** (Partition). A partition ρ of a set of elements $[n] = \{1, \dots, n\}$ is a collection of
 420 subsets $(\mathcal{S}_1, \dots, \mathcal{S}_K)$ such that

$$\mathcal{S}_1 \cup \dots \cup \mathcal{S}_K = [n] \quad \text{and} \quad \forall i \neq j : \mathcal{S}_i \cap \mathcal{S}_j = \emptyset \quad (19)$$

421 Put differently, every element i has to be assigned to precisely one subset \mathcal{S}_k . We denote the size
 422 of the k -th subset \mathcal{S}_k as $n_k = |\mathcal{S}_k|$. Alternatively, we describe a partition ρ as an assignment matrix
 423 $Y = [\mathbf{y}_1, \dots, \mathbf{y}_K]^T \in \{0, 1\}^{K \times n}$. Every row $\mathbf{y}_k \in \{0, 1\}^{1 \times n}$ is a multi-hot vector, where $\mathbf{y}_{ki} = 1$
 424 assigns element i to subset \mathcal{S}_k .

425 In this work, we propose a new two-stage procedure to learn partitions. The proposed formulation
 426 separately infers the number of elements per subset n_k and the assignment of elements to subsets \mathcal{S}_k
 427 by inducing an order on the n elements and filling $\mathcal{S}_1, \dots, \mathcal{S}_K$ sequentially in this order. See Figure 1
 428 for an example.

429 **Definition B.2** (Two-stage partition model). Let $\mathbf{n} = [n_1, \dots, n_K] \in \mathbb{N}_0^K$ be the subset sizes in
 430 ρ , with \mathbb{N}_0 the set of natural numbers including 0 and $\sum_{k=1}^K n_k = n$, where n is the total number
 431 of elements. Let $\pi \in \{0, 1\}^{n \times n}$ be a permutation matrix that defines an order over the n elements.
 432 We define the two-stage partition model of n elements into K subsets as an assignment matrix
 433 $Y = [\mathbf{y}_1, \dots, \mathbf{y}_K]^T \in \{0, 1\}^{K \times n}$ with

$$\mathbf{y}_k = \sum_{i=\nu_k+1}^{\nu_k+n_k} \pi_i, \quad \text{where} \quad \nu_k = \sum_{\ell=1}^{k-1} n_\ell \quad (20)$$

434 such that $Y = [\{\mathbf{y}_k \mid n_k > 0\}_{k=1}^K]^T$.

435 Note that in contrast to previous work on partition models [Mansour and Schork, 2016], we allow \mathcal{S}_k
 436 to be the empty set \emptyset . Hence, K defines the maximum number of possible subsets, not the effective
 437 number of non-empty subsets.

438 To model the order of the elements, we use a permutation matrix $\pi = [\pi_1, \dots, \pi_n]^T \in \{0, 1\}^{n \times n}$
 439 which is a square matrix where every row and column sums to 1. This doubly-stochastic property of
 440 all permutation matrices π [Marcus, 1960] thus ensures that the columns of Y remain one-hot vectors.
 441 At the same time, its rows correspond to n_k -hot vectors \mathbf{y}_k in Definition B.2 and therefore serve as
 442 subset assignment vectors.

443 **Corollary B.3.** A two-stage partition model Y , which follows Definition B.2, is a valid partition
 444 satisfying Definition B.1.

445 *Proof.* By definition, every row π_i and column π_j of π is a one-hot vector, hence every $\sum_{i=\nu_k+1}^{\nu_k+n_k} \pi_i$
 446 results in different, non-overlapping n_k -hot encodings, ensuring $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset \quad \forall i, j$ and $i \neq j$.
 447 Further, since n_k -hot encodings have exactly n_k entries with 1, we have $\sum_{i=\nu_k+1}^{\nu_k+n_k} \sum_{j=1}^n \pi_{ij} = n_k$.
 448 Hence, since $\sum_{k=1}^K n_k = n$, every element i is assigned to a \mathbf{y}_k , ensuring $\mathcal{S}_1 \cup \dots \cup \mathcal{S}_K = [n]$. \square

449 B.2 Two-Stage Random Partition Models

450 An RPM $p(Y)$ defines a probability distribution over partitions Y . In this section, we derive how to
 451 extend the two-stage procedure from Definition B.2 to the probabilistic setting to create a two-stage
 452 RPM. To derive the two-stage RPM's probability distribution $p(Y)$, we need to model distributions
 453 over \mathbf{n} and π . We choose the MVHG distribution $p(\mathbf{n}; \boldsymbol{\omega})$ and the PL distribution $p(\pi; \mathbf{s})$ (see
 454 Section 3).

455 We calculate the probability $p(Y; \boldsymbol{\omega}, \mathbf{s})$ sequentially over the probabilities of subsets $p_{\mathbf{y}_k} := p(\mathbf{y}_k \mid$
 456 $\mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s})$. $p_{\mathbf{y}_k}$ itself depends on the probability over subset permutations $p_{\bar{\pi}_k} := p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \mathbf{s})$,
 457 where a subset permutation matrix $\bar{\pi}$ represents an ordering over n_k out of n elements.

458 **Definition B.4** (Subset permutation matrix $\bar{\pi}$). A subset permutation matrix $\bar{\pi} \in \{0, 1\}^{n_k \times n}$, where
 459 $n_k \leq n$, must fulfill

$$\forall i \leq n_k : \sum_{j=1}^n \bar{\pi}_{ij} = 1 \quad \text{and} \quad \forall j \leq n : \sum_{i=1}^{n_k} \bar{\pi}_{ij} \leq 1.$$

460 We describe the probability distribution over subset permutation matrices $p_{\bar{\pi}_k}$ using Definition B.4
 461 and Equation (3).

462 **Lemma B.5** (Probability over subset permutations $p_{\bar{\pi}_k}$). *The probability $p_{\bar{\pi}_k}$ of any subset permuta-*
 463 *tion matrix $\bar{\pi} = [\bar{\pi}_1, \dots, \bar{\pi}_{n_k}]^T \in \{0, 1\}^{n_k \times n}$ is given by*

$$p_{\bar{\pi}_k} := p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \mathbf{s}) = \prod_{i=1}^{n_k} \frac{(\bar{\pi} \mathbf{s})_i}{Z_k - \sum_{j=1}^{i-1} (\bar{\pi} \mathbf{s})_j} \quad (21)$$

464 where $\mathbf{y}_{<k} = \{\mathbf{y}_1, \dots, \mathbf{y}_{k-1}\}$, $Z_k = Z - \sum_{j \in \mathcal{S}_{<k}} \mathbf{s}_j$ and $\mathcal{S}_{<k} = \bigcup_{j=1}^{k-1} \mathcal{S}_j$.

465 *Proof.* We provide the proof for $p_{\bar{\pi}_1}$, but it is equivalent for all other subsets. Without loss of
 466 generality, we assume that there are n_1 elements in \mathcal{S}_1 . Following Equation (3), the probability of a
 467 permutation matrix $p(\pi; \mathbf{s})$ is given by

$$p(\pi; \mathbf{s}) = \frac{(\pi \mathbf{s})_1}{Z} \frac{(\pi \mathbf{s})_2}{Z - (\pi \mathbf{s})_1} \dots \frac{(\pi \mathbf{s})_n}{Z - \sum_{j=1}^{n-1} (\pi \mathbf{s})_j} \quad (22)$$

468 At the moment, we are only interested in the ordering of the first n_1 elements. The probability of the
 469 first n_1 is given by marginalizing over the remaining $n - n_1$ elements:

$$p(\bar{\pi} \mid n_1; \boldsymbol{\omega}) = \sum_{\pi \in \Pi_1} p(\pi \mid \mathbf{s}) \quad (23)$$

470 where Π_1 is the set of permutation matrices such that the top n_1 rows select the elements in a specific
 471 ordering $\bar{\pi} \in \{0, 1\}^{n_1 \times n}$, i.e. $\Pi_1 = \{\pi : [\pi_1, \dots, \pi_{n_1}]^T = \bar{\pi}\}$. It follows

$$p(\bar{\pi} \mid n_1; \boldsymbol{\omega}) = \sum_{\pi \in \Pi_1} p(\pi \mid \mathbf{s}) \quad (24)$$

$$= \sum_{\pi \in \Pi_1} \prod_{i=1}^n \frac{(\pi \mathbf{s})_i}{Z - \sum_{j=1}^{i-1} (\pi \mathbf{s})_j} \quad (25)$$

$$= \prod_{i=1}^{n_1} \frac{(\bar{\pi} \mathbf{s})_i}{Z - \sum_{j=1}^{i-1} (\bar{\pi} \mathbf{s})_j} \sum_{\pi \in \Pi_1} \prod_{i=1}^{n-n_1} \frac{(\pi \mathbf{s})_{n_1+i}}{Z - \sum_{j=1}^{n_1} (\bar{\pi} \mathbf{s})_j - \sum_{j=1}^{i-1} (\pi \mathbf{s})_j} \quad (26)$$

$$= \prod_{i=1}^{n_1} \frac{(\bar{\pi} \mathbf{s})_i}{Z - \sum_{j=1}^{i-1} (\bar{\pi} \mathbf{s})_j} \sum_{\pi \in \Pi_1} \prod_{i=1}^{n-n_1} \frac{(\pi \mathbf{s})_{n_1+i}}{Z_1 - \sum_{j=1}^{i-1} (\pi \mathbf{s})_j} \quad (27)$$

472 where $Z_1 = Z - \sum_{j=1}^{n_1} (\bar{\pi} \mathbf{s})_j$. It follows

$$p(\bar{\pi} \mid n_1; \boldsymbol{\omega}) = \prod_{i=1}^{n_1} \frac{(\bar{\pi} \mathbf{s})_i}{Z - \sum_{j=1}^{i-1} (\bar{\pi} \mathbf{s})_j} \quad (28)$$

473 □

474 Lemma B.5 describes the probability of drawing the elements $i \in \mathcal{S}_k$ in the order described by the
 475 subset permutation matrix $\bar{\pi}$ given that the elements in $\mathcal{S}_{<k}$ are already determined. Note that in a
 476 slight abuse of notation, we use $p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s})$ as the probability of a subset permutation $\bar{\pi}$
 477 given that there are n_k elements in \mathcal{S}_k and thus $\bar{\pi} \in \{0, 1\}^{n_k \times n}$. Additionally, we condition on the
 478 subsets $\mathbf{y}_{<k}$ and n_k , the size of subset \mathcal{S}_k . In contrast to the distribution over permutations matrices
 479 $p(\pi; \mathbf{s})$ in Equation (3), we take the product over n_k terms and have a different normalization constant
 480 Z_k . Although we induce an ordering over all elements i in Definition B.2, the probability $p_{\mathbf{y}_k}$ is
 481 invariant to intra-subset orderings of elements $i \in \mathcal{S}_k$.

482 **Lemma B.6** (Probability distribution $p_{\mathbf{y}_k}$). *The probability distribution over subset assignments $p_{\mathbf{y}_k}$*
 483 *is given by*

$$p_{\mathbf{y}_k} := p(\mathbf{y}_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s}) = p(n_k \mid n_{<k}; \boldsymbol{\omega}) \sum_{\bar{\pi} \in \Pi_{\mathbf{y}_k}} p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \mathbf{s})$$

484 where $\Pi_{\mathbf{y}_k} = \{\bar{\pi} \in \{0, 1\}^{n_k \times n} : \mathbf{y}_k = \sum_{i=1}^{n_k} \bar{\pi}_i\}$ and $p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \mathbf{s})$ as in Lemma B.5.

485 *Proof.* We can prove the statement of Lemma B.6 as follows:

$$p_{\mathbf{y}_k} = p(\mathbf{y}_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s})$$

$$= \sum_{n'_k} p(\mathbf{y}_k, n'_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s}) \quad (29)$$

$$= \sum_{n'_k} p(n'_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s}) p(\mathbf{y}_k \mid n'_k, \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s}) \quad (30)$$

$$= \sum_{n'_k} p(n'_k \mid n_{<k}; \boldsymbol{\omega}, \mathbf{s}) p(\mathbf{y}_k \mid n'_k, \mathbf{y}_{<k}; \mathbf{s}) \quad (31)$$

$$= p(n_k \mid n_{<k}; \boldsymbol{\omega}, \mathbf{s}) p(\mathbf{y}_k \mid n_k, \mathbf{y}_{<k}; \mathbf{s}) \quad (32)$$

$$= p(n_k \mid n_{<k}; \boldsymbol{\omega}) \sum_{\bar{\pi} \in \Pi_{\mathbf{y}_k}} p(\bar{\pi} \mid n_k, \mathbf{y}_{<k}; \mathbf{s}) \quad (33)$$

486 Equation (29) holds by marginalization, where n'_k denotes the random variable that stands for the
 487 size of subset \mathcal{S}_k . By Bayes' rule, we can then derive Equation (30). The next derivations stem
 488 from the fact that we can compute $n_{<k}$ if $\mathbf{y}_{<k}$ is given, as the assignments $\mathbf{y}_{<k}$ hold information
 489 on the size of subsets $\mathcal{S}_{<k}$. More explicitly, $n_i = \sum_{j=1}^n y_{ij}$. Further, \mathbf{y}_k is independent of $\boldsymbol{\omega}$ if the
 490 size n'_k of subset \mathcal{S}_k is given, leading to Equation (31). We further observe that $p(\mathbf{y}_k \mid n'_k, \mathbf{y}_{<k}; \mathbf{s})$
 491 is only non-zero, if $n'_k = \sum_{i=1}^n y_{ki} = n_k$. Dropping all zero terms from the sum in Equation (31)
 492 thus results in Equation (32). Finally, by Definition B.2, we know that $\mathbf{y}_k = \sum_{i=\nu_k+1}^{\nu_k+n_k} \boldsymbol{\pi}_i$, where
 493 $\nu_k = \sum_{i=1}^{k-1} n_i$ and $\boldsymbol{\pi} \in \{0, 1\}^{n \times n}$ a permutation matrix. Hence, in order to get \mathbf{y}_k given $\mathbf{y}_{<k}$, we
 494 need to marginalize over all permutations of the elements of \mathbf{y}_k given that the elements in $\mathbf{y}_{<k}$ are
 495 already ordered, which corresponds exactly to marginalizing over all subset permutation matrices $\bar{\boldsymbol{\pi}}$
 496 such that $\mathbf{y}_k = \sum_{i=1}^{n_k} \bar{\boldsymbol{\pi}}_i$, resulting in Equation (33). \square

497 In Lemma B.6, we describe the set of all subset permutations $\bar{\boldsymbol{\pi}}$ of elements $i \in \mathcal{S}_k$ by $\Pi_{\mathbf{y}_k}$. Put
 498 differently, we make $p(\mathbf{y}_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s})$ invariant to the ordering of elements $i \in \mathcal{S}_k$ by marginalizing
 499 over the probabilities of subset permutations $p_{\bar{\boldsymbol{\pi}}_k}$ [Xie and Ermon, 2019].

500 Using Lemmas B.5 and B.6, we propose the two-stage random partition $p(Y; \boldsymbol{\omega}, \mathbf{s})$. Since
 501 $Y = [\mathbf{y}_1, \dots, \mathbf{y}_K]^T$, we calculate $p(Y; \boldsymbol{\omega}, \mathbf{s})$, the PMF of the two-stage RPM, sequentially us-
 502 ing Lemmas B.5 and B.6, where we leverage the PL distribution for permutation matrices $p(\boldsymbol{\pi}; \mathbf{s})$ to
 503 describe the probability distribution over subsets $p(\mathbf{y}_k \mid \mathbf{y}_{<k}; \boldsymbol{\omega}, \mathbf{s})$.

504 **Proposition 4.1** (Two-Stage Random Partition Model). Given a probability distribution over
 505 subset sizes $p(\mathbf{n}; \boldsymbol{\omega})$ with $\mathbf{n} \in \mathbb{N}_0^K$ and distribution parameters $\boldsymbol{\omega} \in \mathbb{R}_+^K$ and a PL probability
 506 distribution over random orderings $p(\boldsymbol{\pi}; \mathbf{s})$ with $\boldsymbol{\pi} \in \{0, 1\}^{n \times n}$ and distribution parameters $\mathbf{s} \in \mathbb{R}_+^n$,
 507 the probability mass function $p(Y; \boldsymbol{\omega}, \mathbf{s})$ of the two-stage RPM is given by

$$p(Y; \boldsymbol{\omega}, \mathbf{s}) = p(\mathbf{y}_1, \dots, \mathbf{y}_K; \boldsymbol{\omega}, \mathbf{s}) = p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\boldsymbol{\pi} \in \Pi_Y} p(\boldsymbol{\pi}; \mathbf{s}) \quad (34)$$

508 where $\Pi_Y = \{\boldsymbol{\pi} : \mathbf{y}_k = \sum_{i=\nu_k+1}^{\nu_k+n_k} \boldsymbol{\pi}_i, k = 1, \dots, K\}$, and \mathbf{y}_k and ν_k as in Definition B.2.

509 *Proof.* Using Lemmas B.5 and B.6, we write

$$\begin{aligned} p(Y) &= p(\mathbf{y}_1, \dots, \mathbf{y}_K; \boldsymbol{\omega}, \mathbf{s}) = p(\mathbf{y}_1; \boldsymbol{\omega}, \mathbf{s}) \cdots p(\mathbf{y}_K \mid \{\mathbf{y}_j\}_{j < K}; \boldsymbol{\omega}, \mathbf{s}) \\ &= \left(p(n_1; \boldsymbol{\omega}) \sum_{\bar{\pi}_1 \in \Pi_{\mathbf{y}_1}} p(\bar{\pi}_1 \mid n_1; \mathbf{s}) \right) \\ &\quad \cdots \left(p(n_K \mid \{n_j\}_{j < K}; \boldsymbol{\omega}) \sum_{\bar{\pi}_K \in \Pi_{\mathbf{y}_K}} p(\bar{\pi}_K \mid \{n_j\}_{j \leq K}; \mathbf{s}) \right) \end{aligned} \quad (35)$$

$$\begin{aligned} &= p(n_1; \boldsymbol{\omega}) \cdots p(n_K \mid \{n_j\}_{j < K}; \boldsymbol{\omega}) \\ &\quad \cdot \left(\sum_{\bar{\pi}_1 \in \Pi_{\mathbf{y}_1}} p(\bar{\pi}_1 \mid n_1; \mathbf{s}) \cdots \sum_{\bar{\pi}_K \in \Pi_{\mathbf{y}_K}} p(\bar{\pi}_K \mid \{n_j\}_{j \leq K}; \mathbf{s}) \right) \end{aligned} \quad (36)$$

$$= p(\mathbf{n}; \boldsymbol{\omega}) \left(\sum_{\bar{\pi}_1 \in \Pi_{\mathbf{y}_1}} \cdots \sum_{\bar{\pi}_K \in \Pi_{\mathbf{y}_K}} p(\bar{\pi}_1 \mid n_1; \mathbf{s}) \cdots p(\bar{\pi}_K \mid \{n_j\}_{j \leq K}; \mathbf{s}) \right) \quad (37)$$

$$= p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\pi \in \Pi_Y} p(\pi \mid \mathbf{n}; \mathbf{s}) \quad (38)$$

$$= p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\pi \in \Pi_Y} p(\pi; \mathbf{s}) \quad (39)$$

510

□

511 **B.3 Approximating the Probability Mass Function**

512 **Lemma 4.2.** $p(Y; \boldsymbol{\omega}, \mathbf{s})$ can be upper and lower bounded as follows

$$\forall \pi \in \Pi_Y : p(\mathbf{n}; \boldsymbol{\omega}) p(\pi; \mathbf{s}) \leq p(Y; \boldsymbol{\omega}, \mathbf{s}) \leq |\Pi_Y| p(\mathbf{n}; \boldsymbol{\omega}) \max_{\tilde{\pi}} p(\tilde{\pi}; \mathbf{s}) \quad (40)$$

513 *Proof.* Since $p(\pi; \mathbf{s})$ is a probability we know that $\forall \pi \in \{0, 1\}^{n \times n}$ $p(\pi; \mathbf{s}) \geq 0$. Thus, it follows
514 directly that:

$$\forall \pi \in \Pi_Y : p(Y; \boldsymbol{\omega}, \mathbf{s}) = p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\pi' \in \Pi_Y} p(\pi'; \mathbf{s}) \geq p(\mathbf{n}; \boldsymbol{\omega}) p(\pi; \mathbf{s}),$$

515 proving the lower bound of Lemma 4.2.

516 On the other hand, can prove the upper bound in Lemma 4.2 by:

$$\begin{aligned} p(Y; \boldsymbol{\omega}, \mathbf{s}) &= p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\pi' \in \Pi_Y} p(\pi'; \mathbf{s}) \\ &\leq p(\mathbf{n}; \boldsymbol{\omega}) \sum_{\pi' \in \Pi_Y} \max_{\pi \in \Pi_Y} p(\pi; \mathbf{s}) \\ &= p(\mathbf{n}; \boldsymbol{\omega}) \max_{\pi \in \Pi_Y} p(\pi; \mathbf{s}) \sum_{\pi' \in \Pi_Y} 1 \\ &= |\Pi_Y| \cdot p(\mathbf{n}; \boldsymbol{\omega}) \max_{\pi \in \Pi_Y} p(\pi; \mathbf{s}) \\ &\leq |\Pi_Y| \cdot p(\mathbf{n}; \boldsymbol{\omega}) \max_{\pi} p(\pi; \mathbf{s}) \end{aligned}$$

517 We can compute the maximum probability $\max_{\pi} p(\pi; \mathbf{s})$ with the probability of the permutation
518 matrix $f_{\pi}(\mathbf{s})$, which sorts the unperturbed scores in decreasing order. □

519 **B.4 The Differentiable Random Partition Model**

520 We propose the DRPM $p(Y; \boldsymbol{\omega}, \mathbf{s})$, a differentiable and reparameterizable two-stage RPM.

521 **Lemma 4.3** (DRPM). A two-stage RPM is differentiable and reparameterizable if the distribu-
 522 tion over subset sizes $p(\mathbf{n}; \boldsymbol{\omega})$ and the distribution over orderings $p(\pi; \mathbf{s})$ are differentiable and
 523 reparameterizable.

524 *Proof.* To prove that our two-stage RPM is differentiable we need to prove that we can compute
 525 gradients for the bounds in Lemma 4.2 and to provide a reparameterization scheme for the two-stage
 526 approach in Definition B.2.

527 **Gradients for the bounds:** Since we assume that $p(\mathbf{n}; \boldsymbol{\omega})$ and $p(\pi; \mathbf{s})$ are differentiable and repara-
 528 rameterizable, we only need to show that we can compute $|\Pi_Y|$ and $\max_{\tilde{\pi}} p(\tilde{\pi}; \mathbf{s})$ in a differentiable
 529 manner to prove that the bounds in Lemma 4.2 are differentiable. By definition (see Section 4.1),

$$|\Pi_Y| = \prod_{k=1}^K |\Pi_{\mathbf{y}_k}| = \prod_{k=1}^K n_k!.$$

530 Hence, $|\Pi_Y|$ can be computed given a reparametrized version n_k , which is provided by the
 531 reparametrization trick for the MVHG $p(\mathbf{n}; \boldsymbol{\omega})$. Further, from Equation (14) we immediately see that
 532 the most probable permutation is given by the order induced by sorting the original, unperturbed
 533 scores \mathbf{s} from highest to lowest. This implies that $\max_{\tilde{\pi}} p(\tilde{\pi}; \mathbf{s}) = p(\pi_{\mathbf{s}}; \mathbf{s})$, which we can compute
 534 due to $p(\pi_{\mathbf{s}}; \mathbf{s})$ being differentiable according to our assumptions.

535 **Reparameterization of the two-stage approach:** Given reparametrized versions of \mathbf{n} and π , we
 536 compute a partition as follows:

$$\mathbf{y}_k = \sum_{i=\nu_k+1}^{\nu_k+n_k} \boldsymbol{\pi}_i, \quad \text{where } \nu_k = \sum_{l=1}^{k-1} n_l \quad (41)$$

537 The challenge here is that we need to be able to backpropagate through n_k , which appears as an index
 538 in the sum. Let $\boldsymbol{\alpha}_k = \{0, 1\}^n$, such that

$$(\boldsymbol{\alpha}_k)_i = \begin{cases} 1 & \text{if } \nu_k < i \leq \nu_{k+1} \\ 0 & \text{otherwise} \end{cases}$$

539 Given such $\boldsymbol{\alpha}_k$, we can rewrite Equation (41) with

$$\mathbf{y}_k = \sum_{i=1}^n (\boldsymbol{\alpha}_k)_i \boldsymbol{\pi}_i. \quad (42)$$

540 While this solves the problem of propagating through sum indices, it is not clear how to compute
 541 $\boldsymbol{\alpha}_k$ in a differentiable manner. Similar to other works on continuous relaxations [Jang et al., 2016,
 542 Maddison et al., 2017], we can compute a relaxation of $\boldsymbol{\alpha}_k$ by introducing a temperature τ . Let us
 543 introduce auxiliary function $f : \mathbb{N} \rightarrow [0, 1]^n$, that maps an integer x to a vector with entries

$$f_i(x; \tau) = \sigma\left(\frac{x - i + \epsilon}{\tau}\right),$$

544 such that $f_i(x; \tau) \approx 0$ if $\frac{x-i}{\tau} < 0$ and $f_i(x; \tau) \approx 1$ if $\frac{x-i}{\tau} \geq 0$. Note that $\sigma(\cdot)$ is the standard
 545 sigmoid function and $\epsilon \ll 1$ is a small positive constant to break the tie at $\sigma(0)$. We then compute
 546 an approximation of $\boldsymbol{\alpha}_k$ with

$$\tilde{\boldsymbol{\alpha}}_k(\tau) = f(\nu_k; \tau) - f(\nu_{k-1}; \tau),$$

547 $\tilde{\boldsymbol{\alpha}}_k(\tau) \in [0, 1]^n$. Then, for $\tau \rightarrow 0$ we have $\tilde{\boldsymbol{\alpha}}_k(\tau) \rightarrow \boldsymbol{\alpha}_k$. In practice, we cannot set $\tau = 0$ since
 548 this would amount to a division by 0. Instead, we can apply the straight-through estimator [Bengio
 549 et al., 2013] to the auxiliary function $f(x; \tau)$ in order to get $\tilde{\boldsymbol{\alpha}}_k \in [0, 1]^n$ and use it to compute
 550 Equation (42). \square

551 Note that in our experiments, we use the MVHG relaxation of Sutter et al. [2023] and can thus
 552 leverage that they return one-hot encodings for n_k . This allows a different path for computing $\boldsymbol{\alpha}_k$
 553 which circumvents introducing yet another temperature parameter altogether. We refer to our code in
 554 the supplement for more details.

Table 3: Total GPU hours per experiment. We report the cumulative training and testing hours to generate the results shown in the main part of this manuscript. We relied on our internal cluster infrastructure equipped with RTX2080Ti GPUs. Hence, we report the number of compute hours for this GPU-type.

Experiment	Computation Time (h)
Clustering (Section 5.1)	100
Partitioning of Generative Factors (Section 5.2)	480
MTL (Section 5.3)	100

555 **C Experiments**

556 In the following, we describe each of our experiments in more detail and provide additional ablations.
 557 All our experiments were run on RTX2080Ti GPUs. Each run took 6h-8h (Variational Clustering),
 558 4h-6h (Generative Factor Partitioning), or ~ 1 h (Multitask Learning) respectively. We report the
 559 training and test time per model. Please note that we can only report the numbers to generate the final
 560 results but not the development time.

561 **C.1 Variational Clustering with Random Partition Models**

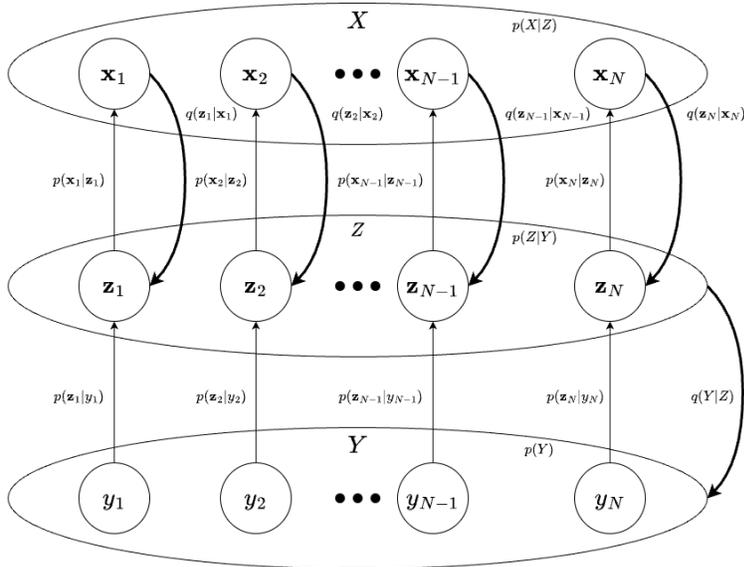


Figure 5: Generative model of the DRPM clustering model. Generative paths are marked with thin arrows, whereas inference is in bold.

562 **C.1.1 Loss Function**

563 As mentioned in Section 5.1, for a given dataset X with N samples, let Z and Y contain the
 564 respective latent vectors and cluster assignments for each sample in X . The generative process
 565 can then be summarized as follows: First, we sample the cluster assignments Y from an RPM,
 566 i.e., $Y \sim P(Y; \omega, \mathbf{s})$. Given Y , we can sample the latent variables Z , where for each \mathbf{y} we have
 567 $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{y}}, \boldsymbol{\sigma}_{\mathbf{y}}^T \mathbb{I}_l)$, $\mathbf{z} \in \mathbb{R}^l$. Finally, we sample X by passing each \mathbf{z} through a decoder like in
 568 vanilla VAEs. Using Bayes rule and Jensen’s inequality, we can then derive the following evidence

569 lower bound (ELBO):

$$\begin{aligned} \log(p(X)) &= \log\left(\int \sum_Y p(X, Y, Z) dZ\right) \\ &\geq \mathbb{E}_{q(Z, Y|X)} \left[\log\left(\frac{p(X|Z)p(Z|Y)p(Y)}{q(Z, Y|X)}\right) \right] \\ &:= \mathcal{L}_{ELBO}(X) \end{aligned}$$

570 We then assume that we can factorize the approximate posterior as follows:

$$q(Z, Y|X) = q(Y|X) \prod_{\mathbf{x} \in X} q(\mathbf{z}|\mathbf{x})$$

571 Note that while we do assume conditional independence between \mathbf{z} given its corresponding \mathbf{x} ,
572 we model $q(Y|X)$ with the DRPM and do not have to assume conditional independence between
573 different cluster assignments. This allows us to leverage dependencies between samples from the
574 dataset. Hence, we can rewrite the ELBO as follows:

$$\begin{aligned} \mathcal{L}_{ELBO}(X) &= \mathbb{E}_{q(Z|X)} [\log(p(X|Z))] \\ &\quad - \mathbb{E}_{q(Y|X)} [KL[q(Z|X)||p(Z|Y)]] \\ &\quad - KL[q(Y|X)||p(Y)] \\ &= \sum_{\mathbf{x} \in X} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \\ &\quad - \sum_{\mathbf{x} \in X} \mathbb{E}_{q(Y|X)} [KL[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|Y)]] \\ &\quad - KL[q(Y|X)||p(Y)] \end{aligned}$$

575 See Figure 5 for an illustration of the generative process and the assumed inference model. Since
576 computing $P(Y)$ and $q(Y|X)$ is intractable, we further apply Lemma 4.2 to approximate the KL-
577 Divergence term in \mathcal{L}_{ELBO} , leading to the following lower bound:

$$\mathcal{L}_{ELBO} \geq \sum_{\mathbf{x} \in X} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \tag{43}$$

$$- \sum_{\mathbf{x} \in X} \mathbb{E}_{q(Y|X)} [KL[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|Y)]] \tag{44}$$

$$- \mathbb{E}_{q(Y|X)} \left[\log \frac{|\Pi_Y| \cdot q(\mathbf{n}; \boldsymbol{\omega}(X))}{p(\mathbf{n}; \boldsymbol{\omega})p(\pi_Y; \mathbf{s})} \right] \tag{45}$$

$$- \log \left(\max_{\tilde{\pi}} q(\tilde{\pi}; \mathbf{s}(X)) \right), \tag{46}$$

578 where π_Y is the permutation that lead to Y during the two-stage resampling process. Further, we
579 want to control the regularization strength of the KL divergences similar to the β -VAE [Higgins
580 et al., 2016]. Since the different terms have different regularizing effects, we rewrite Equations (45)
581 and (46) and weight the individual terms as follows, leading to our final loss:

$$\mathcal{L} := - \sum_{\mathbf{x} \in X} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \tag{47}$$

$$+ \beta \cdot \sum_{\mathbf{x} \in X} \mathbb{E}_{q(Y|X)} [KL[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}|Y)]] \tag{48}$$

$$+ \gamma \cdot \mathbb{E}_{q(Y|X)} \left[\log \left(\frac{|\Pi_Y| \cdot q(\mathbf{n}; \boldsymbol{\omega}(X))}{p(\mathbf{n}; \boldsymbol{\omega})} \right) \right] \tag{49}$$

$$+ \delta \cdot \mathbb{E}_{q(Y|X)} \left[\log \left(\frac{\max_{\tilde{\pi}} q(\tilde{\pi}; \mathbf{s}(X))}{p(\pi_Y; \mathbf{s})} \right) \right] \tag{50}$$

582 C.1.2 Architecture

583 The model for our clustering experiments is a relatively simple, fully-connected autoencoder with a
584 structure as seen in Figure 6. We have a fully connected encoder E with three layers mapping the input

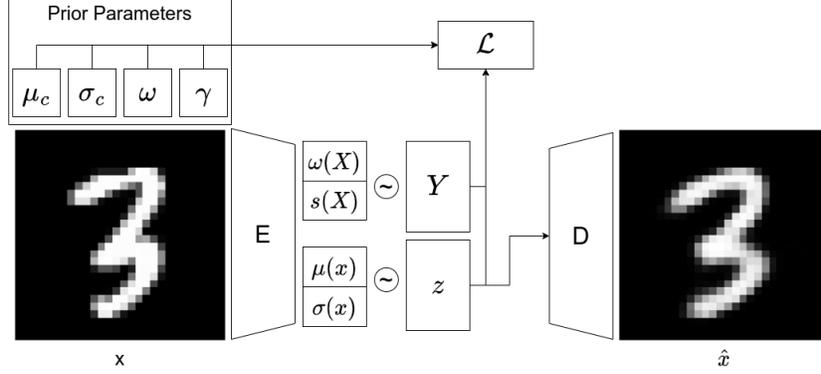


Figure 6: Autoencoder architecture of the DRPM-VC model.

585 to 500, 500, and 2000 neurons, respectively. We then compute each parameter by passing the encoder
 586 output through a linear layer and mapping to the respective parameter dimension in the last layer.
 587 In our experiments, we use a latent dimension size of $l = 10$ for MNIST and $l = 20$ for FMNIST,
 588 such that $\boldsymbol{\mu}(\boldsymbol{x}), \boldsymbol{\sigma}(\boldsymbol{x}) \in \mathbb{R}^l$. To understand the architecture choice for the DRPM parameters, let us
 589 first take a closer look at Equation (48). For each sample \boldsymbol{x} , this term minimizes the expected KL
 590 divergence between its approximate posterior $q(\boldsymbol{z}|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{\mu}(\boldsymbol{x}), \text{diag}(\boldsymbol{\sigma}(\boldsymbol{x})))$ and the prior at index
 591 \boldsymbol{y} given by the partition Y sampled from the DRPM $q(Y|X; \boldsymbol{s}, \boldsymbol{\omega})$, i.e., $\mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{y}}, \text{diag}(\boldsymbol{\sigma}_{\boldsymbol{y}}))$. Ideally,
 592 the most likely partition should assign the approximate posterior to the prior that minimizes this KL
 593 divergence. We can compute such $\boldsymbol{s}(X)$ and $\boldsymbol{\omega}(X)$ given the parameters of the approximate posterior
 594 and priors as follows:

$$\forall \boldsymbol{x}_i \in X : s_i(\boldsymbol{x}_i) = u \cdot (K - \arg \min_k (KL[\mathcal{N}(\boldsymbol{\mu}(\boldsymbol{x}_i), \text{diag}(\boldsymbol{\sigma}(\boldsymbol{x}_i)) || \mathcal{N}(\boldsymbol{\mu}_k, \text{diag}(\boldsymbol{\sigma}_k))]))$$

$$\boldsymbol{\omega}(X) = \frac{1}{|X|} \sum_{\boldsymbol{x} \in X} \left\{ \frac{\mathcal{N}(\boldsymbol{x} | \boldsymbol{\mu}_k, \text{diag}(\boldsymbol{\sigma}_k))}{\sum_{k'=1}^K \mathcal{N}(\boldsymbol{x} | \boldsymbol{\mu}_{k'}, \text{diag}(\boldsymbol{\sigma}_{k'}))} \right\}_{k=1}^K,$$

595 where u is a scaling constant that controls the probability of sampling the most likely partition. Note
 596 that $\boldsymbol{\omega}$ and \boldsymbol{s} minimize Equation (48) if defined this way when given the distribution parameters of the
 597 approximate posterior and the priors. The only thing that is left unclear is how much u should scale
 598 the scores \boldsymbol{s} . Ultimately, we leave u as a learnable parameter but detach the rest of the computation of
 599 \boldsymbol{s} and $\boldsymbol{\omega}$ from the computational graph to improve stability during training. Finally, once we resample
 600 $\boldsymbol{z} \sim \mathcal{N}(\boldsymbol{\mu}(\boldsymbol{x}), \boldsymbol{\sigma}(\boldsymbol{x}))$, we pass it through a fully connected decoder D with four layers mapping \boldsymbol{z} to
 601 2000, 500, and 500 neurons in the first three layers and then finally back to the input dimension in
 602 the last layer to end up with the reconstructed sample $\hat{\boldsymbol{x}}$.

603 C.1.3 Training

604 As in vanilla VAEs, we can estimate the reconstruction term in Equation (47) with MCMC by
 605 applying the reparametrization trick [Kingma and Welling, 2014] to $q(\boldsymbol{z}|\boldsymbol{x})$ to sample M samples
 606 $\boldsymbol{z}^{(i)} \sim q(\boldsymbol{z}|\boldsymbol{x})$ and compute their reconstruction error to estimate Equation (47). Similarly, we can

607 sample from $q(Y|X)$ L times to estimate the terms in Equations (48) to (50), such that we minimize

$$\begin{aligned} \tilde{\mathcal{L}} := & - \sum_{\mathbf{x} \in X} \frac{1}{M} \sum_{i=1}^M \log p(\mathbf{x} | \mathbf{z}^{(i)}) \\ & + \frac{\beta}{L} \cdot \sum_{\mathbf{x} \in X} \sum_{i=1}^L KL[q(\mathbf{z} | \mathbf{x}) || p(\mathbf{z} | Y^{(i)})] \\ & + \frac{\gamma}{L} \cdot \sum_{i=1}^L \log \left(\frac{|\Pi_{Y^{(i)}}| \cdot q(\mathbf{n}^{(i)}; \boldsymbol{\omega}(X))}{p(\mathbf{n}^{(i)}; \boldsymbol{\omega})} \right) \\ & + \frac{\delta}{L} \cdot \sum_{i=1}^L \log \left(\frac{\max_{\tilde{\pi}} q(\tilde{\pi}; \mathbf{s}(X))}{p(\pi_{Y^{(i)}}; \mathbf{s})} \right) \end{aligned}$$

608 In our experiments, we set $M = 1$ and $L = 100$ since the MVHG and PL distributions are not
 609 concentrated around their mean very well, and more Monte Carlo samples thus lead to better
 610 approximations of the expectation terms. We further set $\beta = 1$ for MNIST and $\beta = 0.1$ for FMNIST,
 611 and otherwise $\gamma = 1$, and $\delta = 0.01$ for all experiments.

612 To resample \mathbf{n} and π we need to apply temperature annealing [Grover et al., 2019, Sutter et al.,
 613 2023]. To do this, we applied the exponential schedule that was originally proposed together with the
 614 Gumbel-Softmax trick [Jang et al., 2016, Maddison et al., 2017], i.e., $\tau = \max(\tau_{final}, \exp(-rt))$,
 615 where t is the current training step and r is the annealing rate. For our experiments, we choose
 616 $r = \frac{\log(\tau_{final}) - \log(\tau_{init})}{100000}$ in order to anneal over 100000 training step. Like Jang et al. [2016], we
 617 set $\tau_{init} = 1$ and $\tau_{final} = 0.5$.

618 Similar to Jiang et al. [2016], we quickly realized that proper initialization of the cluster parameters
 619 and network weights is crucial for variational clustering. In our experiments, we pretrained the
 620 autoencoder structure by adapting the contrastive loss of [Li et al., 2022], as they demonstrated that
 621 their representations manage to retain clusters in low-dimensional space. Further, we also added
 622 a reconstruction loss to initialize the decoder properly. To initialize the prior parameters, we fit a
 623 GMM to the pretrained embeddings of the training set and took the resulting Gaussian parameters to
 624 initialize our priors. Note that we used the same initialization across all baselines. See Appendix C.1.4
 625 for an ablation where we pretrain with only a reconstruction loss similar to what was proposed with
 626 the VADE baseline.

627 To optimize the DRPM-VC in our experiments, we used the AdamW [Loshchilov and Hutter, 2019]
 628 optimizer with a learning rate of 0.0001 with a batch size of 256 for 1024 epochs. During initial
 629 experiments with the DRPM-VC, we realized that the pretrained weights of the encoder would often
 630 lose the learned structure in the first couple of training epochs. We suspect this to be an artifact of
 631 instabilities induced by temperature annealing. To deal with these problems, we decided to freeze
 632 the first three layers of the encoder when training the DRPM-VC, giving us much better results. See
 633 Appendix C.1.5 for an ablation where we applied the same optimization procedure to VADE.

634 Finally, when training the VADE baseline and the DRPM-VC on FMNIST, we often observe a local
 635 optimum where the prior distributions collapse and become identical. We can solve this problem
 636 by refitting the GMM in the latent space every 10 epochs and by using the resulting parameters to
 637 reinitialize the prior distributions.

638 C.1.4 Reconstruction Pretraining

639 While the results of our variational clustering method depend a lot on the specific pretraining, we
 640 want to demonstrate that improvements over the baselines do not depend on the chosen pretraining
 641 method. To that end, we repeat our experiments but initialize the weights of our model with an
 642 autoencoder that has been trained to minimize the mean squared error between the input and the
 643 reconstruction. This initialization procedure was originally proposed in [Jiang et al., 2016]. We
 644 present the results of this ablation in Table 4. Simply minimizing the reconstruction error does not
 645 necessarily retain cluster structures in the latent space. Thus, it does not come as a surprise that overall
 646 results get about 10% to 20% worse across most metrics, especially for MNIST, while results on
 647 FMNIST only slightly decrease. However, we still beat the baselines across most metrics, suggesting

Table 4: We compare the clustering performance of the DRPM-VC on test sets of MNIST and FMNIST between GMM in latent space (Latent GMM) and Variational Deep Embedding (VADE) initializing weights using an autoencoder trained on a reconstruction objective. We measure performance in terms of the Normalized Mutual Information (NMI), Adjusted Rand Index (ARI), and cluster accuracy (ACC) over five seeds and put the best model in bold.

	MNIST			FMNIST		
	NMI	ARI	ACC	NMI	ARI	ACC
LATENT GMM	0.75±0.00	0.66±0.01	0.75 ±0.01	0.56±0.02	0.41±0.03	0.57±0.02
VADE	0.77 ±0.02	0.62±0.04	0.69±0.04	0.53±0.07	0.35±0.08	0.47±0.09
DRPM-VC	0.74±0.00	0.67 ±0.01	0.75 ±0.02	0.59 ±0.01	0.47 ±0.02	0.62 ±0.01

Table 5: We compare the clustering performance of the DRPM-VC on test sets of MNIST and FMNIST between GMM in latent space (Latent GMM), and Variational Deep Embedding (VADE) when freezing the encoder. We measure performance in terms of the Normalized Mutual Information (NMI), Adjusted Rand Index (ARI), and cluster accuracy (ACC) over five seeds and put the best model in bold.

	MNIST			FMNIST		
	NMI	ARI	ACC	NMI	ARI	ACC
LATENT GMM	0.86±0.02	0.83±0.06	0.88±0.07	0.60±0.00	0.47±0.01	0.62±0.01
VADE	0.90 ±0.02	0.88 ±0.06	0.92±0.06	0.64 ±0.01	0.47±0.01	0.59±0.03
DRPM-VC	0.89±0.01	0.88 ±0.03	0.94 ±0.02	0.64 ±0.00	0.51 ±0.01	0.65 ±0.00

648 that modeling the implicit dependencies between cluster assignments helps to improve variational
649 clustering performance.

650 C.1.5 Baselines with fixed Encoder

651 For the experiments in the main text, we wanted to implement the VADE baseline similar to the
652 original method proposed in Jiang et al. [2016]. This means, in contrast to our method, we used
653 their optimization procedure, i.e., Adam with a learning rate of 0.002 with a decay of 0.95 every
654 10 steps, and did not freeze the encoder as we do for the DRPM-VC. To ensure our results do not
655 stem from this minor discrepancy, we perform an ablation experiment on VADE using the same
656 optimizer and learning rate as with the DRPM-VC and freeze the encoder backbone. The results of
657 this additional experiment can be found in Appendix C.1.5. As can be seen, VADE results do improve
658 when adjusting the optimization procedure in this way. However, we still match or improve upon the
659 results of VADE in most metrics, especially in ARI and ACC, suggesting purer clusters compared
660 to VADE. We suspect this is because we assign samples to fixed clusters when sampling from the
661 DRPM, whereas VADE performs soft assignments by marginalizing over a categorical distribution.

662 C.1.6 Additional Partition Samples

663 In Section 5.1, we have seen a sample of a partition of the DRPM-VC trained on FMNIST. We
664 provide additional samples for both MNIST and FMNIST at the end of the appendix in Figures 12
665 and 13. We can see that for both datasets, the DRPM-VC learns coherent representations of each
666 cluster that easily allow us to generate new samples from each class.

667 C.1.7 Samples per cluster

668 In addition to sampling partitions and then generating samples according to the sampled cluster
669 assignments, we can also directly sample from each of the learned priors. We show some examples
670 of this for both MNIST and FMNIST at the end of the appendix in Figures 14 and 15. We can again
671 see that the DRPM-VC learns accurate cluster representations since each of the samples seems to

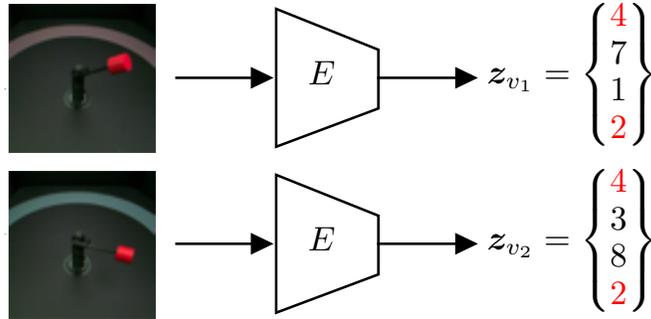


Figure 7: Motivation for the Partitioning of Generative Factors under weak supervision. The knowledge about the data collection process provides a weak supervision signal. We have access to a dataset of pairs of images of the same robot arm with a subset of shared generative factors (in red). We want to learn the shared and independent generative factors in addition to learning from the data. The images of the robot arms are taken from Locatello et al. [2020] but originate from the mpi3d toy dataset (see https://github.com/rr-learning/disentangle_dataset). The image is from Sutter et al. [2023] and their ICLR 2023 presentation video (see <https://iclr.cc/virtual/2023/poster/10707>).

672 correspond to one of the classes in the datasets. Further, the clusters also seem to capture the diversity
 673 in each cluster, as we see a lot of variety across the generated samples.

674 C.2 Variational Partitioning of Generative Factors

675 We assume that we have access to multiple instances or views of the same event, where only a
 676 subset of generative factors changes between views. The knowledge about the data collection process
 677 provides a form of weak supervision. For example, we have two images of a robot arm as depicted
 678 here on the left side (see [Gondal et al., 2019]), which we would describe using high-level concepts
 679 such as color, position or rotation degree. From the data collection process, we know that a subset
 680 of these generative factors is shared between the two views. We do not know how many generative
 681 factors there are in total nor how many of them are shared. More precisely, looking at the robot arm,
 682 we do not know that the views share two latent factors, depicted in red, out of a total of four factors.
 683 Please note that we chose four generative in Figure 7 only for illustrative reason as there are seven
 684 generative factors in the *mpi3d* toy dataset. Hence, the goal of learning under weak supervision is not
 685 only to infer good representations, but also inferring the number of shared and independent generative
 686 factors. Learning what is shared and what is independent lets us reason about the group structure
 687 without requiring explicit knowledge in the form of expensive labeling. Additionally, leveraging weak
 688 supervision and, hence, the underlying group structure holds promise for learning more generalizable
 689 and disentangled representations (see [e.g., Locatello et al., 2020]).

690 C.2.1 Generative Model

691 We assume the following generative model for DRPM-VAE

$$p(\mathbf{X}) = \int_{\mathbf{z}} p(\mathbf{X}, \mathbf{z}) d\mathbf{z} \quad (51)$$

$$= \int_{\mathbf{z}} p(\mathbf{X} | \mathbf{z}) p(\mathbf{z}) d\mathbf{z} \quad (52)$$

692 where $\mathbf{z} = \{z_s, z_1, z_2\}$. The two frames share an unknown number n_s of generative latent factors
 693 z_s , and an unknown number, n_1 and n_2 , of independent factors z_1 and z_2 . The RPM infers n_k and

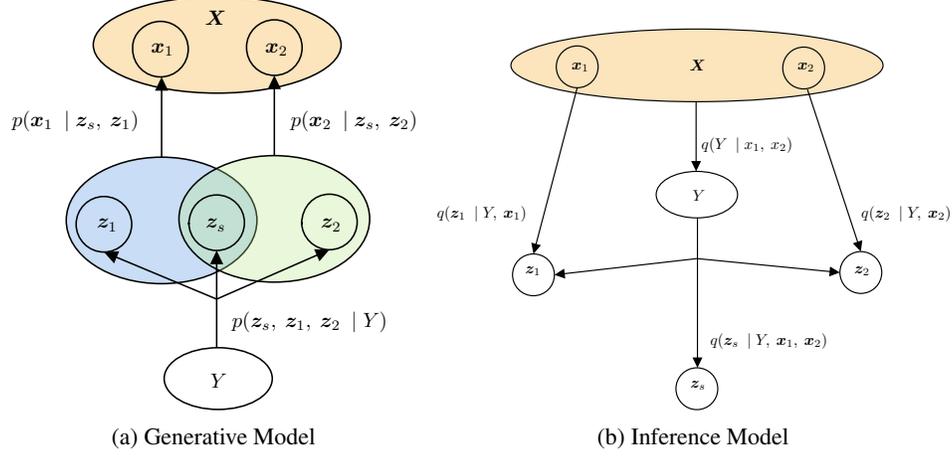


Figure 8: Graphical Models for DRPM-VAE models in the weakly-supervised experiment.

694 z_k using Y . Hence, the generative model extends to

$$\begin{aligned}
p(\mathbf{X}) &= \int_{\mathbf{z}} p(\mathbf{X} | \mathbf{z}) \sum_Y p(\mathbf{z} | Y) p(Y) d\mathbf{z} \\
&= \int_{\mathbf{z}} p(\mathbf{x}_1, \mathbf{x}_2 | \mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2) \sum_Y p(\mathbf{z} | Y) p(Y) d\mathbf{z} \\
&= \int_{\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2} p(\mathbf{x}_1 | \mathbf{z}_s, \mathbf{z}_1) p(\mathbf{x}_2 | \mathbf{z}_s, \mathbf{z}_2) \sum_Y p(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2 | Y) p(Y) d\mathbf{z}_s d\mathbf{z}_1 d\mathbf{z}_2 \quad (53)
\end{aligned}$$

695 Figure 8 shows the generative and inference models assumptions in a graphical model.

696 C.2.2 DRPM ELBO

697 We derive the following ELBO using the posterior approximation $q(\mathbf{z}, Y | \mathbf{X})$

$$\mathcal{L}_{ELBO}(\mathbf{X}) = \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log p(\mathbf{X} | \mathbf{z}, Y) - \log \frac{q(\mathbf{z}, Y | \mathbf{X})}{p(\mathbf{z}, Y)} \right] \quad (54)$$

$$= \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log p(\mathbf{X} | \mathbf{z}) - \log \frac{q(\mathbf{z} | Y, \mathbf{X}) q(Y | \mathbf{X})}{p(\mathbf{z}) p(Y)} \right] \quad (55)$$

$$= \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log p(\mathbf{x}_1, \mathbf{x}_2 | \mathbf{z}) - \log \frac{q(\mathbf{z} | Y, \mathbf{X})}{p(\mathbf{z})} - \log \frac{q(Y | \mathbf{X})}{p(Y)} \right] \quad (56)$$

$$\begin{aligned}
&= \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_1 | \mathbf{z}_s, \mathbf{z}_1)] - \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_2 | \mathbf{z}_s, \mathbf{z}_2)] \\
&\quad - \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log \frac{q(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2 | Y, \mathbf{X})}{p(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2)} \right] - \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log \frac{q(Y | \mathbf{X})}{p(Y)} \right] \quad (57)
\end{aligned}$$

698 Following Lemma 4.2, we are able to optimize DRPM-VAE using the following ELBO $\mathcal{L}_{ELBO}(\mathbf{X})$:

$$\mathcal{L}_{ELBO} \geq \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_1 | \mathbf{z}_s, \mathbf{z}_1)] - \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_2 | \mathbf{z}_s, \mathbf{z}_2)] \quad (58)$$

$$- \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log \frac{q(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2 | Y, \mathbf{X})}{p(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2)} \right] \quad (59)$$

$$- \mathbb{E}_{q(Y | \mathbf{X})} \left[\log \left(\frac{|\Pi_Y| \cdot q(\mathbf{n} | \mathbf{X}; \boldsymbol{\omega})}{p(\mathbf{n}; \boldsymbol{\omega}_p) p(\pi_Y; \mathbf{s}_p)} \right) \right] \quad (60)$$

$$- \log \left(\max_{\tilde{\pi}} q(\tilde{\pi} | \mathbf{X}; \mathbf{s}) \right), \quad (61)$$

699 where π_Y is the permutation that lead to Y during the two-stage resampling process. Further, we
700 want to control the regularization strength of the KL divergences similar to the β -VAE [Higgins et al.,

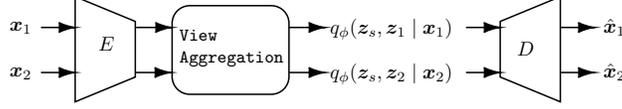


Figure 9: Setup for the weakly-supervised experiment. The three methods differ only in the View Aggregation module.

2016]. The ELBO $\mathcal{L}(\mathbf{X})$ to be optimized can be written as

$$\mathcal{L}_{ELBO} = \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_1 | \mathbf{z}_s, \mathbf{z}_1)] + \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} [\log p(\mathbf{x}_2 | \mathbf{z}_s, \mathbf{z}_2)] \quad (62)$$

$$- \beta \cdot \mathbb{E}_{q(\mathbf{z}, Y | \mathbf{X})} \left[\log \frac{q(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2 | Y, \mathbf{X})}{p(\mathbf{z}_s, \mathbf{z}_1, \mathbf{z}_2)} \right] \quad (63)$$

$$- \gamma \cdot \mathbb{E}_{q(Y | \mathbf{X})} \left[\log \left(\frac{|\Pi_Y| \cdot q(\mathbf{n}; \boldsymbol{\omega}(\mathbf{X}))}{p(\mathbf{n}; \boldsymbol{\omega}_p)} \right) \right] \quad (64)$$

$$- \delta \cdot \mathbb{E}_{q(Y | \mathbf{X})} \left[\log \left(\frac{\max_{\tilde{\pi}} q(\tilde{\pi}; \mathbf{s}(\mathbf{X}))}{p(\pi_Y; \mathbf{s}_p)} \right) \right] \quad (65)$$

where $\mathbf{s}(\mathbf{X})$ and $\boldsymbol{\omega}(\mathbf{X})$ denote distribution parameters, which are inferred from \mathbf{X} (similar to the Gaussian parameters in the vanilla VAE).

As in vanilla VAEs, we can estimate the reconstruction term in Equation (58) with MCMC by applying the reparametrization trick [Kingma and Welling, 2014] to $q(\mathbf{z} | Y, \mathbf{X})$ to sample L samples $\mathbf{z}^{(l)} \sim q(\mathbf{z} | Y, \mathbf{X})$ and compute their reconstruction error to estimate Equation (58). Similarly, we can sample from $q(Y | \mathbf{X})$ L times. We use $L = 1$ to estimate all expectations in \mathcal{L}_{ELBO} .

C.2.3 Implementation and Hyperparameters

In this experiment, we use the `disentanglement_lib` from Locatello et al. [2020]. We use the same architectures proposed in the original paper for all methods we compare to. The baseline algorithms, LabelVAE [Bouchacourt et al., 2018, Hosoya, 2018] and AdaVAE [Locatello et al., 2020] are already implemented in `disentanglement_lib`. For details on the implementation of these methods we refer to the original paper from Locatello et al. [2020]. HGVAE is implemented in Sutter et al. [2023]. We did not change any hyperparameters or network details. All experiments were performed using $\beta = 1$ as this is the best performing β (according to Locatello et al. [2020]). For DRPMVAE we chose $\gamma = 0.25$ for all runs. All models are trained on 5 different random seeds and the reported results are averaged over the 5 seeds. We report mean performance with standard deviations.

We adapted Figure 9 from Sutter et al. [2023]. It shows the baseline architecture, which is used for all methods. As already stated in the main part of the paper, the methods only differ in the View Aggregation module, which determines the shared and independent latent factors. Given a subset S of shared latent factors, we have

$$q_\phi(\mathbf{z}_i | \mathbf{x}_j) = \text{avg}(q_\phi(\mathbf{z}_i | \mathbf{x}_1), q_\phi(\mathbf{z}_i | \mathbf{x}_2)) \quad \forall i \in S \quad (66)$$

$$q_\phi(\mathbf{z}_i | \mathbf{x}_j) = q_\phi(\mathbf{z}_i | \mathbf{x}_j) \quad \text{else} \quad (67)$$

where avg is the averaging function of choice [Locatello et al., 2020, Sutter et al., 2023] and $j \in \{1, 2\}$. The methods used (i. e. Label-VAE, Ada-VAE, HG-VAE, DRPM-VAE) differ in how to select the subset S .

For DRPM-VAE, we infer $\boldsymbol{\omega}$ from the pairwise KL-divergences KL_{pw} between the latent vectors of the two views.

$$KL_{pw}(\mathbf{x}_1, \mathbf{x}_2) = \frac{1}{2} KL[q(\mathbf{z}_1 | \mathbf{x}_1) || q(\mathbf{z}_2 | \mathbf{x}_2)] + \frac{1}{2} KL[q(\mathbf{z}_2 | \mathbf{x}_2) || q(\mathbf{z}_1 | \mathbf{x}_1)] \quad (68)$$

where $q(\mathbf{z}_j | \mathbf{x}_j)$ are the encoder outputs of the respective images. We do not average or sum across dimensions in the computation of $KL_{pw}(\cdot)$ such that the $KL_{pw}(\cdot)$ is d -dimensional, where d is the latent space size. The encoder E in Figure 9 maps to $\boldsymbol{\mu}(\mathbf{x}_j)$ and $\boldsymbol{\sigma}(\mathbf{x}_j)$ of a Gaussian distribution. Hence, we can compute the KL divergences above in closed form. Afterwards, we feed the pairwise KL divergence KL_{pw} to a single fully-connected layer, which maps from d to K values

$$\log \boldsymbol{\omega} = FC(KL_{pw}(\mathbf{x}_1, \mathbf{x}_2)) \quad (69)$$



Figure 10: Samples from the noisyMultiMNIST dataset with increasing noise ratio in the right task.

732 where $d = 10$ and $K = 2$ in this experiment. d is the total number of latent dimensions and K is the
 733 number of groups in the latent space. To infer the scores $s(\mathbf{X})$ we again rely on the pairwise KL
 734 divergence KL_{pw} . Instead of using another fully-connected layer, we directly use the log-values of
 735 the pairwise KL divergence

$$\log s = \log KL_{pw}(\mathbf{x}_1, \mathbf{x}_2) \quad (70)$$

736 Similar to the original works, we also anneal the temperature parameter for $p(\mathbf{n}; \boldsymbol{\omega})$ and $p(\pi; \mathbf{s})$
 737 [Grover et al., 2019, Sutter et al., 2023]. We use the same annealing function as in the clustering
 738 experiment (see Appendix C.1). We anneal the temperature τ from 1.0 to 0.5 over the complete
 739 training time.

740 C.3 Multitask Learning

741 C.3.1 MultiMNIST Dataset

742 The different tasks in multitask learning often vary in difficulty. To measure the effect of discrepancies
 743 in task difficulties on DRPM-MTL, we introduce the noisyMultiMNIST dataset.

744 The noisyMultiMNIST dataset modifies the MultiMNIST dataset [Sabour et al., 2017] as follows.
 745 In the right image, we set each pixel value to zero with probability $\alpha \in [0, 1]$. This is done before
 746 merging the left and right image in order to only affect the difficulty of the right task. Note that for
 747 $\alpha = 0$ noisyMultiMNIST is equivalent to MultiMNIST and for $\alpha = 1$ the right task can no longer be
 748 solved. This allows us to control the difficulty of the right task, without changing the difficulty of the
 749 left. A few examples are shown in Figure 10.

750 C.3.2 Implementation & Architecture

751 The multitask loss function for the *MultiMNIST* dataset is

$$\mathbb{L} = w_L \mathbb{L}_L + w_R \mathbb{L}_R \quad (71)$$

752 where w_L and w_R are the loss weights, and \mathbb{L}_L and \mathbb{L}_R are the individual loss terms for the respective
 753 tasks L and R . In our experiments, we set the task weights to be equal for all dataset versions, i.e.
 754 $w_L = w_R = 0.5$. We use these loss weights for the DRPM-MTL and ULS method. For the ULS
 755 method, it is by definition and to see the influence of a mismatch in loss weights. The DRPM-MTL
 756 method on the other hand does not need additional weighting of loss terms. The task losses are
 757 defined as cross-entropy losses

$$\mathbb{L}_t = - \sum_{c=1}^{C_t} \mathbf{g}t_c \log p_c = -\mathbf{g}t^T \log \mathbf{p} \quad (72)$$

758 where $C_L = C_R = 10$ for MultiMNIST, $\mathbf{g}t$ is a one-hot encoded label vector and \mathbf{p} is a categorical
 759 vector of estimated class assignments probabilities, i.e. $\sum_c p_c = 1$.

760 The predictions for the individual tasks \mathbf{p}_t are given as

$$\mathbf{p}_t = h_{\theta_t}(z), \quad \text{where} \quad (73)$$

$$z = \text{enc}_{\theta}(\mathbf{x}) \quad (74)$$

761 for a sample $\mathbf{x} \in \mathbf{X}$ (see also Figure 11). We use an adaptation of the LeNet-5 architecture LeCun
 762 et al. [1998] to the multitask learning problem [Sener and Koltun, 2018]. Both DRPM-MTL and ULS
 763 use the same network $\text{enc}_{\theta}(\cdot)$ with shared architecture up to some layer for both tasks, after which
 764 the network branches into two task-specific sub-networks that perform the classifications. Different

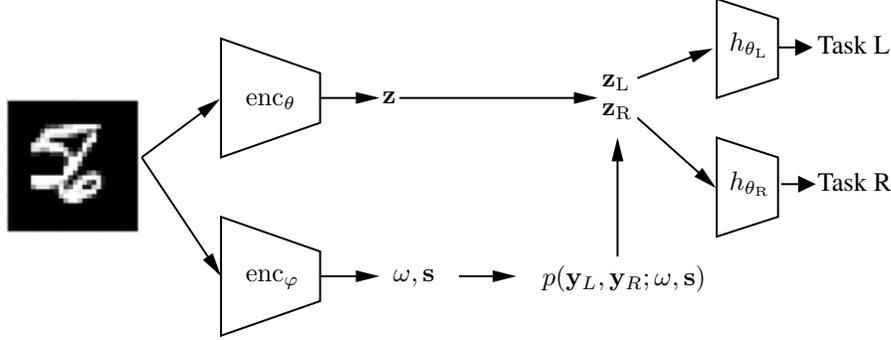


Figure 11: Overview of the multitask learning pipeline of the DRPM-MTL method.

765 to the ULS method, the task-specific networks in the DRPM-MTL pipeline predict the digit using
 766 only a subset of z . DRPM-MTL uses the following prediction scheme

$$767 \quad \mathbf{p}_t = h_{\theta_t}(\mathbf{z}_t), \quad \text{where} \quad (75)$$

$$768 \quad \mathbf{z}_t = \mathbf{z} \odot \mathbf{y}_t \quad (76)$$

$$769 \quad \mathbf{y}_t = \text{DRPM}(\boldsymbol{\omega}, \mathbf{s})_t = \text{DRPM}(\text{enc}_\varphi(\mathbf{x}))_t \quad (77)$$

770 The DRPM-MTL encoder first predicts a latent representation $\mathbf{z} \leftarrow \text{enc}_\theta(\mathbf{x})$, where \mathbf{x} is the input
 771 image. Using the same encoder architecture but different parameters φ , we predict a partitioning
 772 encoding $\mathbf{z}' \leftarrow \text{enc}_\varphi(\mathbf{x})$. With a single linear layer per DRPM log-parameter $\log \boldsymbol{\omega}$ and $\log \mathbf{s}$ are
 773 computed. Next we infer the partition masks $\mathbf{y}_L, \mathbf{y}_R \sim p(\mathbf{y}_L, \mathbf{y}_R; \boldsymbol{\omega}, \mathbf{s})$. We then feed the masked
 774 latent representations $\mathbf{z}_L \leftarrow \mathbf{z} \odot \mathbf{y}_L$ and $\mathbf{z}_R \leftarrow \mathbf{z} \odot \mathbf{y}_R$ into the task specific classification networks
 775 $h_{\theta_L}(\mathbf{z}_L)$ and $h_{\theta_R}(\mathbf{z}_R)$ respectively to obtain the task specific predictions. Since the two tasks in the
 776 MultiMNIST dataset are of similar nature, the task-specific networks h_{θ_L} and h_{θ_R} share the same
 777 architecture, but have different parameters.

775 C.3.3 Training

776 For both the ULS and the DRPM-MTL model, we use the Adam optimizer with learning rate 0.0005
 777 and train them for 200 epochs with a batch size of 256. We again choose an exponential schedule for
 778 the temperature τ and anneal it over the training time, as is explained in Appendix C.1.3.

779 In our ablation we use $\alpha \in \{0, 0.1, 0.2, \dots, 0.9\}$ and train each model with five different seeds.
 780 The reported accuracies and partition sizes are then means over the five seeds with the error bands
 781 indicating the variance and standard deviation respectively. We evaluate each model after the epoch
 782 with the best average test accuracy.

783 C.3.4 CelebA for MTL

784 In addition to the experiment shown in Section 5.3, we show additional results for DRPM-MTL on
 785 the CelebA dataset [Liu et al., 2015]. In MTL, each of the 40 attributes of the CelebA dataset serves
 786 as an individual task. Hence, using CelebA for MTL results is a 40 task learning problem making the
 787 scaling of different task losses more difficult compared to MultiMNIST (see Section 5.3) where we
 788 only need to scale two different tasks.

789 We again use the newly introduced DRPM-MTL method and compare it to the ULS model. We use
 790 the same pipeline as for MultiMNIST dataset but with different encoders and hyperparameters (see
 791 Appendices C.3.2 and C.3.3). We use the pipeline of Sener and Koltun [2018] with a ResNet-based en-
 792 coder to map an image to a representation of $d = 64$ dimensions. For architectural details, we refer to
 793 Sener and Koltun [2018] and <https://github.com/isl-org/MultiObjectiveOptimization>.

794 Again, ULS inputs all $d = 64$ dimensions to the task-specific sub-networks whereas DRPM-MTL
 795 partitions the intermediate representations into n_T different subsets, which are then fed to the
 796 respective task networks. n_T is the number of tasks.

797 Compared to the MultiMNIST experiment (see Appendix C.3.2), we introduce an additional regular-
 798 ization for the DRPM-MTL method. The additional regularization is based on the upper bound in

Table 6: Results for the MTL experiment on the CelebA dataset. We compare the DRPM-MTL again to the ULS method. We assess the performance of both methods on two sub-experiments of the CelebA experiment. In Table 6a, we form a MTL experiment with 10 different tasks. In Table 6b, we form a MTL experiment with 20 different tasks where the first 10 tasks are the same as in the 10 tasks experiment. We train both methods for 50 methods using a learning rate of 0.0001 and a batch size of 128. The temperature annealing schedule remains the same as in the MultiMNIST experiment. We report the per task classification accuracy in percentages (%) as well as the average task accuracy in the bottom row of both subtables.

(a) 10 Tasks			(b) 20 Tasks		
	ULS	DRPM		ULS	DRPM
T0	92.0±0.5	92.4±0.5	T0	92.4±0.7	93.0±0.2
T1	83.8±0.4	83.7±0.2	T1	83.7±0.6	83.9±0.7
T2	80.2±0.5	80.2±0.4	T2	79.9±0.6	80.1±0.4
T3	81.9±0.8	82.2±0.6	T3	82.4±0.5	83.0±0.7
T4	98.5±0.2	98.5±0.1	T4	98.6±0.1	98.6±0.1
T5	95.2±0.2	95.3±0.2	T5	95.2±0.1	95.5±0.0
T6	80.0±1.4	82.4±0.4	T6	82.0±1.3	84.4±0.4
T7	82.0±0.3	82.2±0.2	T7	82.5±0.1	82.8±0.2
T8	89.7±0.7	90.7±0.2	T8	90.1±0.9	91.0±0.4
T9	94.6±0.5	95.0±0.2	T9	94.7±0.2	95.1±0.1
avg(Tasks)	87.8±0.3	88.3±0.1	T10	95.9±0.1	95.9±0.1
			T11	84.9±0.1	84.6±0.3
			T12	91.0±0.4	91.6±0.2
			T13	94.7±0.1	94.9±0.1
			T14	95.4±0.3	96.0±0.1
			T15	99.2±0.0	99.2±0.1
			T16	95.8±0.3	96.0±0.1
			T17	97.3±0.3	97.5±0.2
			T18	91.2±0.3	91.2±0.1
			T19	87.0±0.3	87.3±0.2
			avg(Tasks)	90.7±0.2	91.1±0.1

799 Lemma 4.2 and is penalizing size of $|\Pi_Y|$ for a given n . Hence, the loss function changes to

$$\mathbb{L} = \frac{1}{n_T} \sum_{t=1}^{n_T} \mathbb{L}_t + \lambda \cdot \mathbb{L}_{\text{reg}} \quad (78)$$

$$\text{where } \mathbb{L}_{\text{reg}} = \log \left(\prod_{t=1}^{n_T} n_t! \right) = \sum_{t=1}^{n_T} \log \Gamma(n_t + 1) \quad (79)$$

800 For both versions of the experiment (i.e. $n_T = 10$ and $n_T = 20$), we set $\lambda = 0.015 \approx \frac{1}{64}$, which is
 801 the number of elements we want to partition. The task losses \mathbb{L}_t are simple BCE losses similar to the
 802 MultiMNIST experiments but with two classes per task only.

803 We perform two different experiments based on the CelebA experiment. First, we use form a MTL
 804 experiments using the first 10 attributes out of the 40 attributes. Second, we increase the number of
 805 different tasks to 20. Because we sort the attributes alphabetically in both cases, the first 10 tasks are
 806 shared between the two experiment versions.

807 Table 6 shows the results of both methods, ULS and DRPM-MTL. We see that the DRPM-MTL
 808 scales better to a larger number of tasks compared to the ULS method, highlighting the importance of
 809 finding new ways of automatic scaling between tasks. Interestingly, the DRPM-MTL outperforms
 810 the ULS method on most tasks for the 20-tasks experiment even though it has only access to
 811 $d/n_T = 64/20 = 3.2$ dimensions on average. On the other hand, the ULS method can access the
 812 full set of 64 dimensions for every single task.

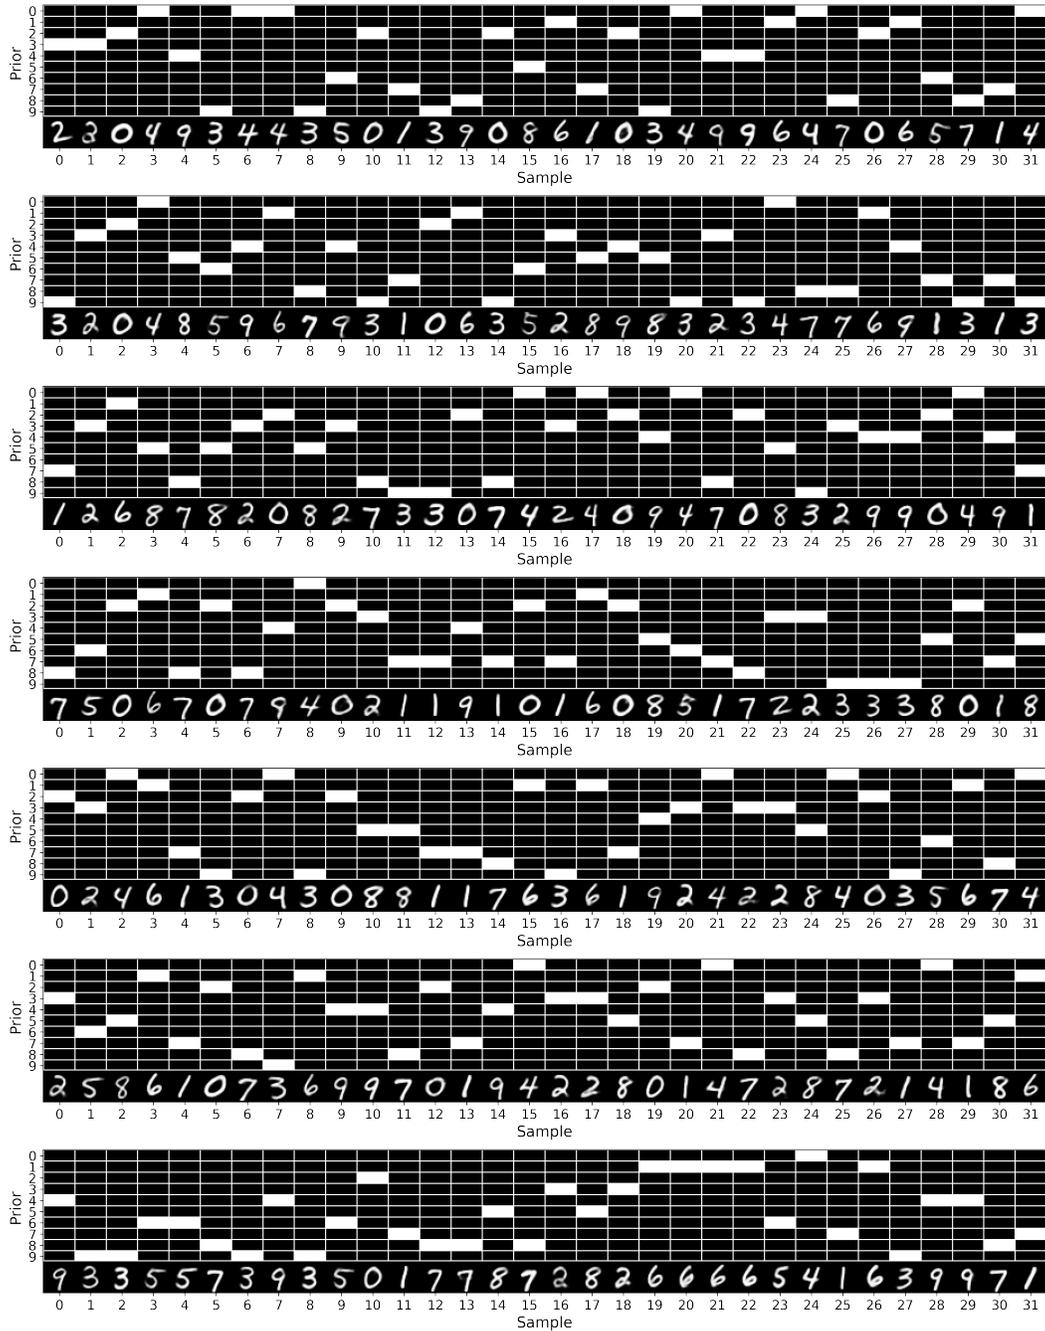


Figure 12: Additional partition samples from the DRPM-VC trained on MNIST. The different sets of each partition match each of the digits very well, even after repeatedly sampling from the model.

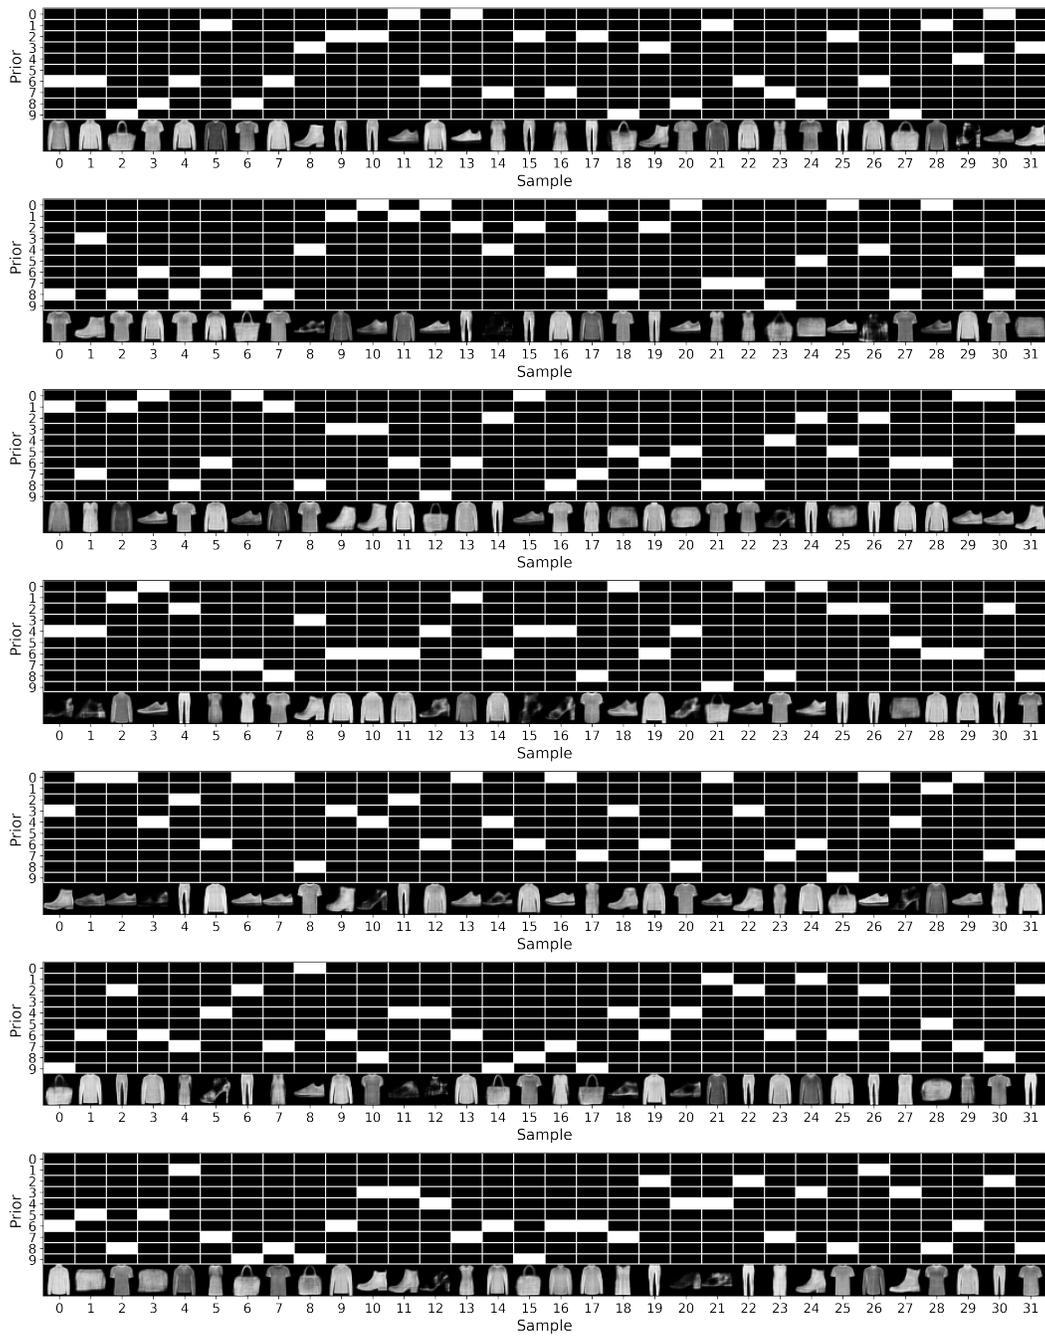


Figure 13: Additional partition samples from the DRPM-VC trained on FMNIST. Most clusters accurately represent one of the clothing categories and generate new samples very well. The only problem is with the handbag class, where the DRPM-VC learns two different clusters for different kinds of handbags (cluster 5 and 6).



Figure 14: Various samples from each of the generative priors. Each prior learns to represent one of the digits. Further, we see a lot of variation between the different samples, suggesting that the clusters of the DRPM-VC manage to capture some of the diversity present in the dataset.

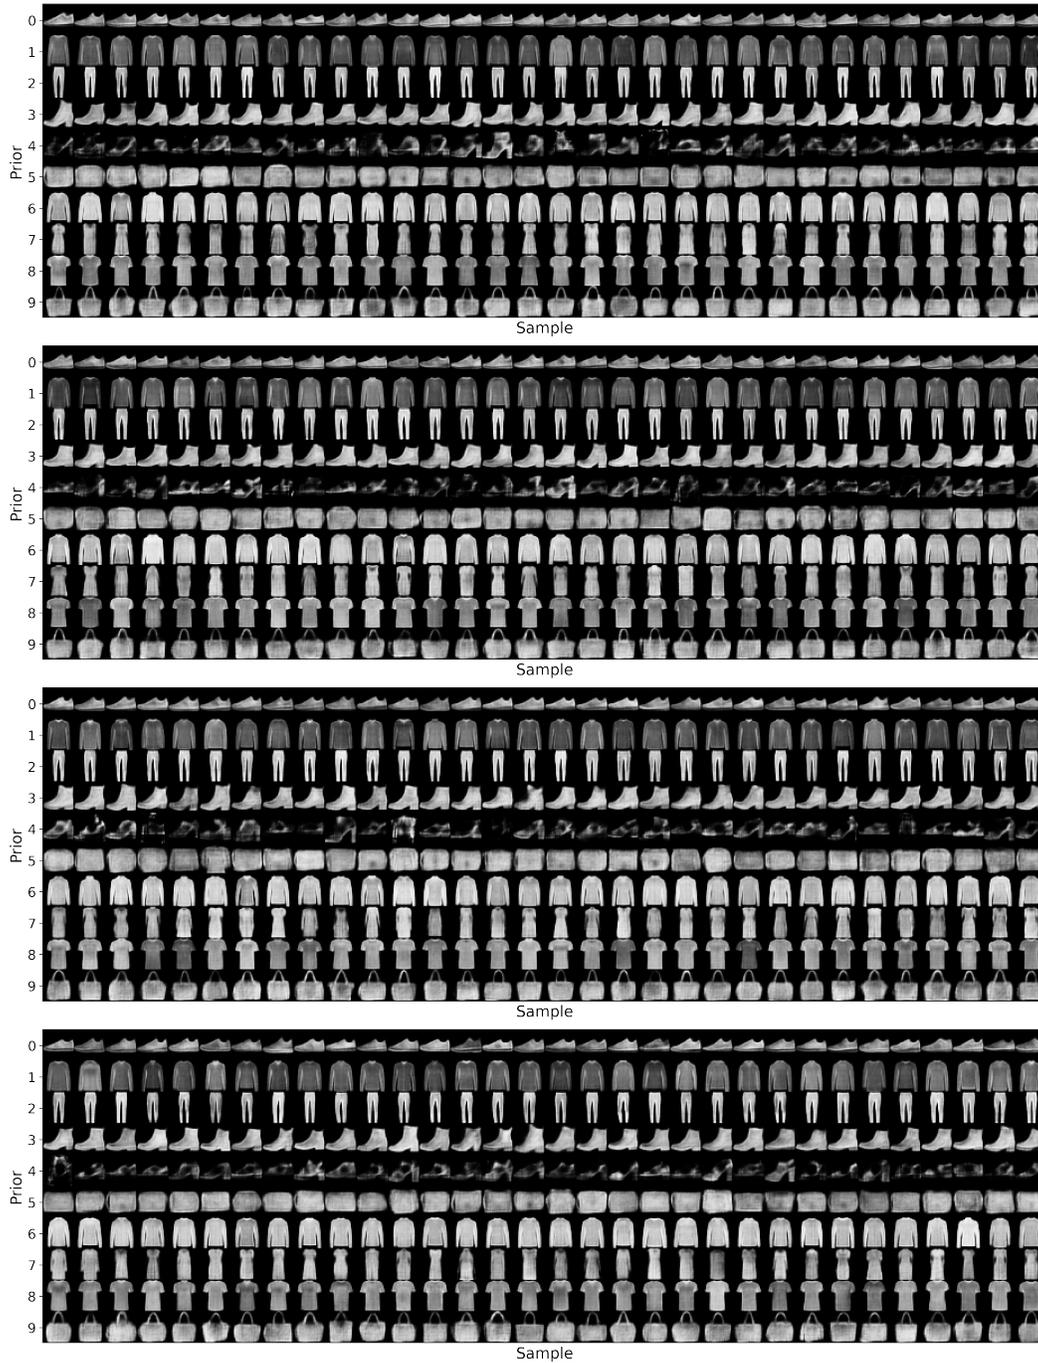


Figure 15: Various samples from each of the generative priors. Each prior learns to represent one of the digits. The DRPM-VC learns nice representations that provide coherent generations of most classes. For high-heels (cluster 4), generating new samples seems difficult due to the heterogeneity within that class.

813 **References**

- 814 R. P. Adams and R. S. Zemel. Ranking via sinkhorn propagation. *arXiv preprint arXiv:1106.1925*,
815 2011.
- 816 Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic
817 neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- 818 C. M. Bishop and M. Svensen. Robust Bayesian Mixture Modelling. *To appear in the Proceedings of*
819 *ESANN*, page 1, 2004. Publisher: Citeseer.
- 820 D. Bouchacourt, R. Tomioka, and S. Nowozin. Multi-level variational autoencoder: Learning
821 disentangled representations from grouped observations. In *Thirty-Second AAAI Conference on*
822 *Artificial Intelligence*, 2018.
- 823 R. Caruana. Multitask learning: A knowledge-based source of inductive bias. In *Machine learning:*
824 *Proceedings of the tenth international conference*, pages 41–48, 1993.
- 825 R. Caruana and V. R. de Sa. Promoting poor features to supervisors: Some inputs work better as
826 outputs. *Advances in Neural Information Processing Systems*, 9, 1996.
- 827 J. Chesson. A non-central multivariate hypergeometric distribution arising from biased sampling with
828 application to selective predation. *Journal of Applied Probability*, 13(4):795–797, 1976. Publisher:
829 Cambridge University Press.
- 830 R. De la Cruz-Mesía, F. A. Quintana, and P. Müller. Semiparametric Bayesian Classification with
831 longitudinal Markers. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 56
832 (2):119–137, Mar. 2007. ISSN 0035-9254. doi: 10.1111/j.1467-9876.2007.00569.x. URL
833 <https://doi.org/10.1111/j.1467-9876.2007.00569.x>.
- 834 N. Dilokthanakul, P. A. M. Mediano, M. Garnelo, M. C. H. Lee, H. Salimbeni, K. Arulkumaran,
835 and M. Shanahan. Deep unsupervised clustering with Gaussian mixture variational autoencoders.
836 *arXiv preprint arXiv:1611.02648*, 2016.
- 837 R. A. Fisher. The logic of inductive inference. *Journal of the royal statistical society*, 98(1):39–82,
838 1935. Publisher: JSTOR.
- 839 A. Fog. Calculation methods for Wallenius’ noncentral hypergeometric distribution. *Communications*
840 *in Statistics—Simulation and Computation*®, 37(2):258–273, 2008. Publisher: Taylor & Francis.
- 841 M. W. Gondal, M. Wuthrich, D. Miladinovic, F. Locatello, M. Breidt, V. Volchkov, J. Akpo,
842 O. Bachem, B. Schölkopf, and S. Bauer. On the Transfer of Inductive Bias from Simulation
843 to the Real World: a New Disentanglement Dataset. In *Advances in Neural Information Processing*
844 *Systems*, volume 32. Curran Associates, Inc., 2019.
- 845 R. L. Graham, D. E. Knuth, O. Patashnik, S. L.-C. i. Physics, and undefined 1989. Concrete
846 mathematics: a foundation for computer science. *aip.scitation.org*, 3(5):165, 1989. doi: 10.1063/
847 1.4822863. URL <https://aip.scitation.org/doi/pdf/10.1063/1.4822863>. Publisher:
848 AIP Publishing.
- 849 A. Grover, E. Wang, A. Zweig, and S. Ermon. Stochastic Optimization of Sorting Networks via
850 Continuous Relaxations. In *International Conference on Learning Representations*, 2019. URL
851 <https://openreview.net/forum?id=H1eSS3CcKX>.
- 852 J. A. Hartigan. Partition models. *Communications in Statistics - Theory and Methods*, 19(8):
853 2745–2756, 1990. doi: 10.1080/03610929008830345. Publisher: Taylor & Francis.
- 854 I. Higgins, L. Matthey, A. Pal, C. Burgess, and X. Glorot. beta-vae: Learning basic visual concepts
855 with a constrained variational framework. 2016. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=Sy2fzU9gl)
856 [Sy2fzU9gl](https://openreview.net/forum?id=Sy2fzU9gl).
- 857 H. Hosoya. A simple probabilistic deep generative model for learning generalizable disentangled
858 representations from grouped data. *CoRR*, abs/1809.0, 2018.

- 859 E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint*
860 *arXiv:1611.01144*, 2016.
- 861 Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou. Variational deep embedding: An unsupervised and
862 generative approach to clustering. *arXiv preprint arXiv:1611.05148*, 2016.
- 863 M. E. Johnson. *Multivariate statistical simulation: A guide to selecting and generating continuous*
864 *multivariate distributions*, volume 192. John Wiley & Sons, 1987.
- 865 D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on*
866 *Learning Representations*, 2014.
- 867 V. Kurin, A. D. Palma, I. Kostrikov, S. Whiteson, and M. P. Kumar. In Defense of the Unitary
868 Scalarization for Deep Multi-Task Learning. *CoRR*, abs/2201.04122, 2022. URL [https://](https://arxiv.org/abs/2201.04122)
869 arxiv.org/abs/2201.04122.
- 870 Y. LeCun, C. Cortes, and C. J. Burges. Gradient-based learning applied to document recognition. In
871 *Proceedings of the IEEE*, volume 86, pages 2278–2324, 1998. Issue: 11.
- 872 C. J. Lee and H. Sang. Why the Rich Get Richer? On the Balancedness of Random Partition Models.
873 *arXiv preprint arXiv:2201.12697*, 2022.
- 874 Y. Li, M. Yang, D. Peng, T. Li, J. Huang, and X. Peng. Twin Contrastive Learning for Online
875 Clustering. *International Journal of Computer Vision*, 130(9):2205–2221, Sept. 2022. ISSN
876 0920-5691, 1573-1405. doi: 10.1007/s11263-022-01639-z. URL [http://arxiv.org/abs/](http://arxiv.org/abs/2210.11680)
877 [2210.11680](http://arxiv.org/abs/2210.11680). arXiv:2210.11680 [cs].
- 878 S. W. Linderman, G. E. Mena, H. J. Cooper, L. Paninski, and J. P. Cunningham. Reparameterizing
879 the Birkhoff Polytope for Variational Permutation Inference. In A. J. Storkey and F. Pérez-Cruz,
880 editors, *International Conference on Artificial Intelligence and Statistics, {AISTATS} 2018, 9-11*
881 *April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, volume 84 of *Proceedings of Machine*
882 *Learning Research*, pages 1618–1627. PMLR, 2018. URL [http://proceedings.mlr.press/](http://proceedings.mlr.press/v84/linderman18a.html)
883 [v84/linderman18a.html](http://proceedings.mlr.press/v84/linderman18a.html).
- 884 Z. Liu, P. Luo, X. Wang, and X. Tang. Deep Learning Face Attributes in the Wild, Sept. 2015. URL
885 <http://arxiv.org/abs/1411.7766>. arXiv:1411.7766 [cs].
- 886 F. Locatello, B. Poole, G. Rätsch, B. Schölkopf, O. Bachem, and M. Tschannen. Weakly-supervised
887 disentanglement without compromises. In *International Conference on Machine Learning*, pages
888 6348–6359. PMLR, 2020.
- 889 I. Loshchilov and F. Hutter. Decoupled Weight Decay Regularization, Jan. 2019. URL [http:](http://arxiv.org/abs/1711.05101)
890 [//arxiv.org/abs/1711.05101](http://arxiv.org/abs/1711.05101). arXiv:1711.05101 [cs, math].
- 891 R. D. Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 1959.
- 892 J. MacQueen. Classification and analysis of multivariate observations. In *5th Berkeley Symp. Math.*
893 *Statist. Probability*, pages 281–297, 1967.
- 894 C. Maddison, A. Mnih, and Y. Teh. The concrete distribution: A continuous relaxation of discrete
895 random variables. In *International Conference on Learning Representations*, 2017.
- 896 L. Manduchi, K. Chin-Cheong, H. Michel, S. Wellmann, and J. E. Vogt. Deep Conditional Gaussian
897 Mixture Model for Constrained Clustering. *ArXiv*, abs/2106.0, 2021.
- 898 T. Mansour and M. Schork. *Commutation relations, normal ordering, and Stirling numbers*. CRC
899 Press Boca Raton, 2016.
- 900 M. Marcus. Some Properties and Applications of Doubly Stochastic Matrices. *The American*
901 *Mathematical Monthly*, 67(3):215–221, 1960. ISSN 00029890, 19300972. URL [http://www.](http://www.jstor.org/stable/2309679)
902 [jstor.org/stable/2309679](http://www.jstor.org/stable/2309679). Publisher: Mathematical Association of America.

- 903 G. E. Mena, D. Belanger, S. W. Linderman, and J. Snoek. Learning Latent Permutations with Gumbel-
904 Sinkhorn Networks. In *6th International Conference on Learning Representations, {ICLR} 2018,*
905 *Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings.* OpenReview.net,
906 2018. URL <https://openreview.net/forum?id=Byt3oJ-0W>.
- 907 G. Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*,
908 pages 666–704, 1781.
- 909 F. Petersen, C. Borgelt, H. Kuehne, and O. Deussen. Monotonic Differentiable Sorting Networks. In
910 *International Conference on Learning Representations*, 2021.
- 911 J. Pitman. Some Developments of the Blackwell-Macqueen URN Scheme. *Lecture Notes-Monograph*
912 *Series*, 30:245–267, 1996. ISSN 07492170. URL <http://www.jstor.org/stable/4355949>.
913 Publisher: Institute of Mathematical Statistics.
- 914 R. L. Plackett. The analysis of permutations. *Journal of the Royal Statistical Society: Series C*
915 *(Applied Statistics)*, 24(2):193–202, 1975. Publisher: Wiley Online Library.
- 916 S. Prillo and J. Eisenschlos. SoftSort: A Continuous Relaxation for the argsort Operator. In
917 *Proceedings of the 37th International Conference on Machine Learning, {ICML} 2020, 13-18 July*
918 *2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 7793–7802.
919 PMLR, 2020. URL <http://proceedings.mlr.press/v119/prillo20a.html>.
- 920 G.-C. Rota. The Number of Partitions of a Set. *The American Mathematical Monthly*, 71(5):498,
921 May 1964. ISSN 00029890. doi: 10.2307/2312585. Publisher: JSTOR.
- 922 Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval.
923 *International journal of computer vision*, 40(2):99, 2000. Publisher: Springer Nature BV.
- 924 S. Sabour, N. Frosst, and G. E. Hinton. Dynamic Routing Between Capsules. In
925 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and
926 R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Cur-
927 ran Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper/2017/file/](https://proceedings.neurips.cc/paper/2017/file/2cad8fa47bbef282badbb8de5374b894-Paper.pdf)
928 [2cad8fa47bbef282badbb8de5374b894-Paper.pdf](https://proceedings.neurips.cc/paper/2017/file/2cad8fa47bbef282badbb8de5374b894-Paper.pdf).
- 929 R. Santa Cruz, B. Fernando, A. Cherian, and S. Gould. Deeppermnet: Visual permutation learning.
930 In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages
931 3949–3957, 2017.
- 932 O. Sener and V. Koltun. Multi-task learning as multi-objective optimization. *Advances in neural*
933 *information processing systems*, 31, 2018.
- 934 R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The*
935 *annals of mathematical statistics*, 35(2):876–879, 1964. Publisher: JSTOR.
- 936 T. M. Sutter, L. Manduchi, A. Ryser, and J. E. Vogt. Learning Group Importance using the Differ-
937 entiable Hypergeometric Distribution. In *International Conference on Learning Representations*,
938 2023.
- 939 L. L. Thurstone. A law of comparative judgment. In *Scaling*, pages 81–92. Routledge, 1927.
- 940 K. T. Wallenius. Biased sampling; the noncentral hypergeometric probability distribution. Technical
941 report, Stanford Univ Ca Applied Mathematics And Statistics Labs, 1963.
- 942 H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmark-
943 ing Machine Learning Algorithms, Sept. 2017. URL <http://arxiv.org/abs/1708.07747>.
944 arXiv:1708.07747 [cs, stat].
- 945 S. M. Xie and S. Ermon. Reparameterizable subset sampling via continuous relaxations. *arXiv*
946 *preprint arXiv:1901.10517*, 2019.
- 947 D. Xin, B. Ghorbani, A. Garg, O. Firat, and J. Gilmer. Do Current Multi-Task Optimization Methods
948 in Deep Learning Even Help? *arXiv preprint arXiv:2209.11379*, 2022.
- 949 J. I. Yellott. The relationship between Luce’s choice axiom, Thurstone’s theory of comparative
950 judgment, and the double exponential distribution. *Journal of Mathematical Psychology*, 15(2):
951 109–144, 1977. Publisher: Elsevier.