

STRUDEL: Learning Structured-Decomposable Probabilistic Circuits

Meihua Dang

MHDANG@CS.UCLA.EDU

Antonio Vergari

AVER@CS.UCLA.EDU

Guy Van den Broeck

GUYVDB@CS.UCLA.EDU

Computer Science Department, University of California, Los Angeles

Abstract

Probabilistic circuits (PCs) represent a probability distribution as a computational graph. Enforcing structural properties on these graphs guarantees that several inference scenarios become tractable. Among these properties, structured decomposability is a particularly appealing one: it enables the efficient and exact computations of the probability of complex logical formulas, and can be used to reason about the expected output of certain predictive models under missing data. This paper proposes STRUDEL, a simple, fast and accurate learning algorithm for structured-decomposable PCs. Compared to prior work for learning structured-decomposable PCs, STRUDEL delivers more accurate single PC models in fewer iterations, and dramatically scales learning when building ensembles of PCs. It achieves this scalability by exploiting another structural property of PCs, called determinism, and by sharing the same computational graph across mixture components. We show these advantages on standard density estimation benchmarks and challenging inference scenarios.

Keywords: Probabilistic circuits; structure learning; structured decomposability.

1. Introduction

In several real-world scenarios, decision making requires *advanced* probabilistic reasoning, i.e., the ability to answer *complex probabilistic queries* (Vergari et al., 2020). Consider, for instance, querying a generative model for the probability of events described as logical constraints (Bekker et al., 2015), e.g., rankings of user preferences (Choi et al., 2015; Shen et al., 2017); or the probability of Bayesian classifiers agreeing on their prediction (Choi and Van den Broeck, 2018); or to conform to human expectations (Khosravi et al., 2019b,a). Answering these queries goes beyond the capabilities of intractable probabilistic models like classical Bayesian networks (BNs) and more recent neural estimators such as variational autoencoders and normalizing flows (Papamakarios et al., 2019).

Moreover, in many sensitive domains like healthcare and finance, the result of these queries is required (i) to be *exact*, as approximations without guarantees would make the decision making process brittle, and (ii) to be provided in a *limited amount of time*. This explains the recently growing interest around tractable probabilistic models (TPMs) which guarantee both (i) and (ii) by design.

Probabilistic circuits (PCs) (Choi et al., 2020; Vergari et al., 2020) propose a unifying framework to abstract from the myriad of different TPM representations. Among these, arithmetic circuits (Darwiche, 2003), probabilistic sentential decision diagrams (Kisa et al., 2014), sum-product networks (Poon and Domingos, 2011), and cutset networks (Rahman et al., 2014) naturally fit under the umbrella of PCs. Classical bounded-treewidth graphical models (Koller and Friedman, 2009) and their mixtures (Meilă and Jordan, 2000) are easily cast into a PC. Within the framework of PCs, one can reason about the tractable inference capabilities of a model via the structural properties

of its computational graph. In turn, this enables learning routines that, by enforcing such specific structural properties, deliver PCs guaranteeing tractable inference for the desired classes of queries.

The *structured decomposability* (Kisa et al., 2014) property of PCs enables the largest class of tractable inference scenarios. Indeed, all the advanced probabilistic queries we mentioned in the introductory paragraph can be exactly and efficiently answered using structured-decomposable PCs. In a nutshell, a structured-decomposable PC encodes a probability distribution in a computational graph by recursively decomposing it into smaller distributions according to a hierarchical partitioning of the random variables, also called *vtree*. However, while inference on structured-decomposable PCs has been extensively studied (Bekker et al., 2015; Choi et al., 2015; Shen et al., 2016; Liang and Van den Broeck, 2017), relatively little attention has gone to *learning* these circuits from data. The only learning algorithm fully tailored towards structured-decomposable PCs is LEARNPSDD (Liang et al., 2017). LEARNPSDD had the merit of introducing the task of learning a vtree. However, this vtree learning step is detached from the actual PC structure learning step as local search, which starts from a fully-factorized circuit that ignores the discovered dependencies in the previous step. Furthermore, LEARNPSDD performs a costly local search: it evaluates many candidate PC structures and computes for each of them their penalized likelihood scores, which makes it hard to scale to larger real-world datasets.

In this paper we introduce STRUDEL, a simpler and faster way to learn structured-decomposable PCs. Specifically, STRUDEL drastically simplifies the search by not computing a likelihood score for each candidate structure but greedily growing the circuit. Moreover, we do not perform vtree learning and instead initialize the PC structure in STRUDEL using the best TPM that can be learned with guarantees. It considerably speeds up learning while still delivering accurate PCs. This is even more relevant when learning large mixtures of PCs; here we propose to scale even further by mixing components that share the same structure. We demonstrate these performance gains on 20 standard benchmarks and on the more challenging task of computing the expected predictions of regression models in the context of missing data (Khosravi et al., 2019a). The rest of paper is organized as follows. Sections 2 and 3 introduce the framework of PCs for tractable probabilistic inference. Sections 4 and 5 describe STRUDEL for learning single PCs and mixtures thereof. Lastly, Section 6 discusses our experimental results.

2. Probabilistic Circuits

Recently, the great interest in tractable probabilistic modeling propelled the introduction of a multitude of representations. Many of these representations can be understood under a unifying computational framework, which we refer to as *probabilistic circuits* (PCs) (Choi et al., 2020; Vergari et al., 2020). PCs reconcile and abstract from the different graphical and syntactic representations of recently introduced formalisms such as arithmetic circuits (Darwiche, 2003), probabilistic sentential decision diagrams (PSDDs) (Kisa et al., 2014), sum-product networks (SPNs) (Poon and Domingos, 2011), cutset networks (Rahman et al., 2014), and bounded-treewidth graphical models (Koller and Friedman, 2009; Meilă and Jordan, 2000). Moreover, PCs enable reasoning about the tractable inference scenarios they support in terms of a unified set of structural properties.

Notation. We use upper-case letters for random variables (RVs), e.g., X, Y , and lowercase ones for their assignments e.g., x, y . Analogously, sets of RVs are denoted by upper-case bold letters, e.g., \mathbf{X}, \mathbf{Y} , and their joint values by the corresponding lower-case ones, e.g., \mathbf{x}, \mathbf{y} . Here we consider discrete RVs, specifically represented as Boolean variables, i.e., having values in $\{0, 1\}$.




Figure 1: A *structural decomposable PC* and its corresponding *vtree* and *equivalent BN*. The CLT over RVs $\mathbf{X} = \{X_1, X_2, X_3, X_4\}$ in (1b) with CPTs in (1a) is turned into the structured-decomposable PC in (1d) whose extracted vtree is shown in (1c). In red, all the edges that are “active” in the circuit flow for the input configuration $\{X_1 = 1, X_2 = 0, X_3 = 1, X_4 = 0\}$. Root node is on the right.

Representation. A probabilistic circuit C over RVs \mathbf{X} is a directed acyclic graph (DAG) representing a computational graph that encodes a probability distribution $\mathbf{p}_C(\mathbf{X})$ in a recursive manner. This DAG has three kinds of nodes: *input distributions* (leaves), *product* nodes and *sum* nodes. Figure 1d shows an example of a PC. Each node n encodes a distribution \mathbf{p}_n , defined as follows. Let C_n denote the sub-circuit rooted at node n and $\text{ch}(n)$ its child nodes. An input distribution n encodes a tractable probability distribution \mathbf{p}_n over some RVs $\phi(n) \subseteq \mathbf{X}$, also called its *scope*. A product node n defines the factorized distribution $\mathbf{p}_n(\mathbf{X}) = \prod_{c \in \text{ch}(n)} \mathbf{p}_c(\mathbf{X})$. W.l.o.g. we will consider product nodes to have only two children. A sum node n defines a mixture model $\mathbf{p}_n(\mathbf{X}) = \sum_{c \in \text{ch}(n)} \theta_{n,c} \mathbf{p}_c(\mathbf{X})$ parameterized by edge weights $\theta_{n,c}$. The scope of a product or sum node is the union of the scopes of its children: $\phi(n) = \cup_{c \in \text{ch}(n)} \phi(c)$. In this work targeting Boolean RVs, we consider univariate input distributions, specifically leafs for RV $X_i \in \mathbf{X}$ will be indicator functions of the form $\mathbf{p}(X_i = 1) = \llbracket X_i = 1 \rrbracket$. Let θ denote all the parameters in a circuit C (only attached to sum node edges, since input distributions are indicators).

Even if they are *probabilistic* models expressed via a *graphical* formalism, probabilistic circuits are *not* classical PGMs. The clear *operational* semantics we described above makes PCs peculiar neural networks (Vergari et al., 2019; Peharz et al., 2020) whose inner units are either products (acting as non-linearities) or convex combinations of their inputs. Overall, a PC C defines a multilinear polynomial (Darwiche, 2003) whose indeterminates are the distributions equipping the leaves of C .

Inference. Any PC C over RVs \mathbf{X} that represents a normalized distribution supports computing the likelihood $\mathbf{p}_C(\mathbf{x})$ given a *complete configuration* \mathbf{x} (a complete evidence query) by evaluating the circuit bottom up: starting from the input distributions and computing the output of children before parents. Additional structural properties of the PC, such as *decomposability*, *smoothness* and *determinism*, can extend the set of probabilistic queries that are guaranteed to be answered exactly and in time linear in the *size* of the PC, that is, its number of edges. Then, query answering reduces to traversing the PC bottom-up (forward) and top-down (backward) given values for the leaf nodes.

A PC is *decomposable* if for every product node, the children have disjoint scopes. That is, product nodes encode well-defined factorized probability distributions. A PC is *smooth* if for every sum node, the children have the same scope. That is, sum nodes encode mixtures of distributions that are well-defined over identical sets of RVs. Smooth and decomposable PCs enable the tractable computation of any marginal query (Darwiche, 2003). SPNs (Poon and Domingos, 2011) are

examples of smooth and decomposable PCs. A circuit is deterministic if for every sum node n and assignment \mathbf{x} , at most one of the children of n have a non-zero output. That is, a deterministic sum defines a mixture model whose components have disjoint support, enabling tractable MAP queries (Chan and Darwiche, 2006). Examples of smooth, decomposable and deterministic PCs are cutset networks (Rahman et al., 2014; Di Mauro et al., 2015b) and selective SPNs (Peharz et al., 2014).

Structured-Decomposable PCs. More recently, the stronger property of *structured decomposability* has been introduced to enable a larger class of tractable inference scenarios (Pipatsrisawat and Darwiche, 2008; Kisa et al., 2014). Briefly, the product nodes in a structured-decomposable PC cannot decompose in arbitrary ways, but must agree on a “contract.” A PC is structured-decomposable if it is *normalized* for a *vtree*, a binary tree encoding a hierarchical decomposition of RVs. Each leaf in a vtree denotes a RV, while an internal node indicates how to decompose a set of RVs in two subsets mapping to its left and right branch. A PC is normalized for a vtree if the scope of every product node decomposes over its children as its corresponding node in the vtree. An example of a vtree and a structured-decomposable PC normalized for it are shown in Figure 1c and Figure 1d.

By enforcing structured decomposability, several classes of advanced probabilistic queries become computable exactly and efficiently. For instance, structured-decomposable PCs allow to compute symmetric and group queries (Bekker et al., 2015) and, given certain constrained vtrees, same-decision probabilities (Oztok and Darwiche, 2015), their expected version (Choi et al., 2017) and classifier agreement (Choi and Van den Broeck, 2018). Moreover, if two PCs conform to the same vtree, it is possible to efficiently compute the KL divergence between them (Liang and Van den Broeck, 2017); to multiply them (Shen et al., 2016); or to take the expectation of one according to the distribution encoded by the other (Choi et al., 2015; Khosravi et al., 2019a). AND-OR graphs (Dechter and Mateescu, 2007) and PSDDs (Kisa et al., 2014) are examples of smooth, deterministic and structured-decomposable PCs.¹ Despite all these advanced inference scenarios that structured-decomposable PCs enable, relatively little attention has been dedicated to learning these circuits from data, and the only attempt so far is difficult to scale (Liang et al., 2017).

3. Circuit Flows: Fast Inference and Parameter Learning

Before explaining STRUDEL, we briefly introduce *circuit flows* – a computational tool that allows us to scale up our learner. Determinism not only makes MAP inference tractable, but also enables closed-form parameter estimation in PCs (Kisa et al., 2014) and dramatically speeds up inference by leveraging circuit flows. We refer to Liang and Van den Broeck (2019) for a formal definition. Figure 1d shows an example of a circuit flow. Intuitively, a circuit flow in a deterministic PC C encodes which parameters are activated by different input configurations.

Fast inference. For a given sample \mathbf{x} , a flow acts as a transformation $f_C : \mathbf{X} \mapsto \{0, 1\}^{|\theta|}$ mapping \mathbf{x} to a binary encoding, called *flow embedding*, with as many entries as there are parameters in C . The k -th entry in $f_C(\mathbf{x})$ is 1 if sample \mathbf{x} flows through the edge associated with the k -th parameter in θ , reaching the output, and 0 otherwise. As such, the log-likelihood $\mathcal{L}\mathcal{L}_C(\theta; \mathbf{x})$ of a deterministic circuit C parameterized by θ , given a single input configuration \mathbf{x} , is efficiently computed as

$$\mathcal{L}\mathcal{L}_C(\theta; \mathbf{x}) = \log(\mathbf{p}_C(\mathbf{x})) = f_C(\mathbf{x})^T \cdot \log(\theta).$$

1. In their original formulation (Kisa et al., 2014), PSDDs required a stronger notion of determinism, which has no practical implication for tractable probabilistic inference.

Similarly, the circuit flow of a batch of samples \mathcal{D} can be represented as a binary matrix $\mathbf{F}_C(\mathcal{D}) \in \{0, 1\}^{|\mathcal{D}| \times |\theta|}$. Then, the log-likelihood given the entire batch at once is efficiently vectorized as

$$\mathcal{L}\mathcal{L}_C(\boldsymbol{\theta}; \mathcal{D}) = \mathbf{F}_C(\mathcal{D}) \cdot \log(\boldsymbol{\theta}).$$

This formulation of the likelihood has the clear benefit that a vectorized computation can yield significant speedups. Furthermore, the matrix $\mathbf{F}_C(\mathcal{D})$ can be computed by propagating bit-vectors up and down the circuit C . Recall that flow embeddings are binary vectors by virtue of the determinism property. Bit-vector arithmetic is extremely efficient, much more so than using floating-point vectors to compute the same likelihoods. Lastly, circuit flows will greatly benefit inference in our large ensembles sharing the same structure (cf. Section 5). Since a flow only depends on the circuit structure, for a mixture $\mathcal{M} = \{C_i\}_{i=1}^k$ of k PCs sharing the same structure, we need to evaluate a single flow $f_{\mathcal{M}}$ once. As such, the log-likelihood of \mathcal{M} given \mathbf{x} can be efficiently computed as:

$$\mathcal{L}\mathcal{L}_{\mathcal{M}}(\boldsymbol{\Theta}; \mathbf{x}) = \text{logsumexp}(f_{\mathcal{M}}(\mathbf{x})^T \cdot \log(\boldsymbol{\Theta}) + \log(\mathbf{w})), \quad (1)$$

where $\mathbf{w} = \{w_i\}_{i=1}^k$ are the mixture weights, and $\boldsymbol{\Theta}$ is the matrix whose columns are the parameters $\boldsymbol{\theta}_i$ of the i th PC in the mixture. We empirically show these speedups in Section 6 and Appendix C.

Parameter learning. More instrumental to our purpose, flow embeddings can be used for parameter learning of a PC C . We define the *aggregate flow* $a_C(i, j; \mathcal{D})$ of one edge $e_{i,j}$ associated with the k th weight $\theta_{i,j}$, as the total number of configurations in the dataset \mathcal{D} that flow through edge $e_{i,j}$, and whose likelihood therefore contains the k th weight $\theta_{i,j}$ as a factor. That is, $a_C(i, j; \mathcal{D}) = \sum_{h=1}^{|\mathcal{D}|} \mathbf{F}_C(\mathcal{D})[h, k]$. Now, the maximum likelihood estimator (MLE) of weight $\theta_{i,j}$ can be computed in closed form as the ratio

$$\theta_{i,j}^{\text{MLE}} = \frac{a_C(i, j; \mathcal{D})}{\sum_* a_C(i, *, \mathcal{D})}. \quad (2)$$

In other words, $\theta_{i,j}^{\text{MLE}}$ can be computed as the ratio of the number of samples in \mathcal{D} flowing through edge $e_{i,j}$ over the total number of samples flowing through node i .

In summary, the circuit flow formulation has the benefit of (1) vectorizing the computation, (2) allowing for a single computation of the flow embeddings to be reused across PCs with the same structure but different parameters, for example in large ensembles, and (3) yielding simple closed-form parameter estimates. Moreover, we will show in Section 4 that flows can act as one of the building blocks in structure learning. Thus, both our learner for expressive single models in Section 4 and for scalable ensembles in Section 5 will make use of this efficient formulation.

4. STRUDEL: Learning Structured-Decomposable Probabilistic Circuits

The objective of structure learning for PCs is to find a circuit structure and parameters that approximate well the data distribution. If the learned PC has to guarantee tractable inference for certain classes of queries, its structure has to enforce the corresponding properties discussed in Section 2. For the advanced inference scenarios we are interested in, and to retain efficient parameter learning (cf. Section 3), we require structured decomposability and determinism. So far, the only alternative learner to deliver such PCs is LEARNPSDD (Liang et al., 2017), which however performs a search procedure that can be computationally inefficient. We review LEARNPSDD next, before proposing STRUDEL and explaining how it addresses some deficiencies of the prior work.

LEARNPSDD and its limitations. LEARNPSDD performs a local search over the space of possible structured-decomposable PCs, given a vtree as input. To learn a vtree from data, a hierarchical

clustering step is performed over the RVs discovering some independence relationships: they are recursively grouped bottom up (or split top down) so as to maximize their pairwise mutual information. Next, local search starts from a fully-factorized PC, that is, one where all RVs are considered to be independent, reshaped to conform to the learned vtree. Each search iteration locally changes the circuit while preserving its semantics and structural properties of smoothness, determinism and structured decomposability. To propose candidates, LEARNPSDD consistently applies two structural transformations – *split* and *clone* – to all possible nodes in the circuits. These candidates are then ranked by their log-likelihood score, penalized by their circuit size.

We highlight the following shortcomings of LEARNPSDD: (i) vtree learning as a separate pre-processing step has a limited effect on structure learning, which starts from a fully-factorized distribution, discarding the dependencies discovered in vtree learning. Moreover, while circuit flows speed up likelihood computation in deterministic circuits, (ii) using likelihood to score candidate structures drastically slows down learning, especially in large data regimes. As a result, when employed in mixture models LEARNPSDD has not been able to scale beyond tens of components.

To overcome these shortcomings, we propose to: (i) extract a vtree structure from the best graphical model that can be learned in tractable time, and then compile it into a structured-decomposable PC, which provides a more informative starting point, (ii) dramatically reduce learning time by employing a greedier local search employing a single transformation, split, and (iii) effectively use circuit flows to speed up parameter learning and likelihood computation. The resulting algorithm is a *simpler, faster* structure learning scheme, yet yielding competitively *accurate* PCs and enabling fast learning of large mixtures. We name it STRUDEL: a STRUctured-DEcomposable Learner.

4.1 From Chow-Liu Trees to Structured-Decomposable PCs

Instead of coming up with a fixed vtree and distilling an initial PC structure out of it, we propose to do the inverse: *learn the best initial PC from data, and distill a vtree from it.*

Chow-Liu trees. Here we use *Chow-Liu trees* (CLTs) as the “best” initial PC structure. CLTs are tree-shaped BNs that satisfy our desiderata: they guarantee to encode the best tree model in terms of KL divergence with the data distribution; they support linear time marginals and MAP inference; and can be learned in time $O(|\mathbf{X}|^2|\mathcal{D}|)$. Furthermore, as we will show next, we can quickly compile CLTs into smooth, deterministic and structured-decomposable PCs and extract a vtree from them.

More formally, a CLT \mathcal{T} over RVs \mathbf{X} is a tree-shaped BN equipped with parameters $\theta_{i|\text{Pa}_i}$ defining the conditional probability table (CPT) of node i associated to RV X_i with parent node Pa_i . An example of a CLT is shown in Figure 1b. The classic Chow-Liu (1968) algorithm learns a CLT \mathcal{T} from data \mathcal{D} by running a maximum spanning tree algorithm over a complete graph induced by the pairwise mutual information (MI) matrix over variables \mathbf{X} as estimated from data \mathcal{D} . These MI estimates are used to compute the $\theta_{i|\text{Pa}_i}$ parameters, and can be smoothed by adding a Laplace correction factor α . See Appendix A for a detailed algorithm. Computation-wise, the cost of learning a CLT is the same as learning a vtree in LEARNPSDD, i.e., that of computing the pairwise MI. Learning-wise, we are not discarding this information but embed it in the initial PC structure.

Compiling CLTs. We now turn our attention to compiling a CLT into a structured-decomposable PC. Compiling generic BNs into smooth, deterministic and decomposable PCs,² has been extensively researched in the literature (Darwiche, 2003), and compilation of a BN into a structure decomposable circuit has also been explored (Choi et al., 2013; Shen et al., 2016). However, compiling a BN (even

2. Specifically to Arithmetic Circuits (Darwiche, 2003) represented as DAGs having parameters attached to leaf nodes.

a CLT) for an arbitrary vtree can lead to an exponentially larger PC. Therefore, we adopt a simple strategy tailored for CLTs that extracts a vtree guaranteeing a linear-size PC in the number of RVs.

We start from the observation that a *rooted* CLT provides a natural variable decomposition. While rooting the CLT can be done arbitrarily, we root it at its *Jordan center* as a heuristics to minimize the resulting vtree depth and thus yielding smaller PCs. Then we traverse the CLT top-down to build the vtree, for each node $X_i \in \mathcal{T}$, if X_i is a leaf node, compile it to a vtree leaf node v_i containing variable X_i ; otherwise build an inner node with v_i as its one branch, and the vtree for its children $\text{ch}(X_i)$ as its other branch. After a vtree is fully grown, we proceed compiling the CLT bottom-up, caching the previously compiled sub-circuits, which guarantees that we obtain a PC of linear size (Darwiche, 2003). For every node $X_i \in \mathcal{T}$ that we visit, and for every parent configuration, we introduce sum nodes selecting a value of X_i with edge weights $\mathbf{p}(X_i | X_{\text{Pa}_i} = x_{\text{Pa}_i})$, i.e., its distribution conditioned on the parent configuration. A corresponding indicator leaf, following the vtree structure, is introduced in the product. This yields a smooth deterministic sum node branching over the possible values for the considered RV. Figure 1d illustrates the compiled PC for the example CLT in Figure 1b.

4.2 How And What to Split

From this initial circuit, STRUDEL applies the *split* operation at each step, performing a *greedy search*. It waives the necessity for computing the penalized log-likelihoods of possible candidate structures.

Our split operation is based on the one proposed in Liang et al. (2017), which builds on Lowd and Domingos (2008). Given a PC structure C^t at iteration t , we create C^{t+1} in a *two-step procedure*. First, we select one *edge* $e_{n,c}$ to split, from one sum node n to one of its child product nodes c , and one *variable* X_i to split in the scope of c . Second, we make two *partial copies* of sub-circuits rooted at c conditioned on $\llbracket X_i = 0 \rrbracket$ and $\llbracket X_i = 1 \rrbracket$ respectively. These partial copies are carried out by copying nodes up to a certain depth bound. We thus create C^{t+1} by removing child c from node n in C^t , while adding both of the new copies. This splitting operation preserves smoothness and determinism since sums are conditioned on the RV X_i , while copying preserves structured decomposability.

Selecting edges to split. To select an edge and variable to split, the simplest but uninformed strategy would be to pick one edge and one variable randomly. We name these two strategies ERAND and vRAND . Clearly, a more informed heuristics, but less expensive than the computation of the likelihood, would benefit search. We introduce two novel heuristics, eFLOW and vMI for selecting an edge and a variable respectively. Specifically, eFLOW selects edge $e_{i,j}$ which maximizes the aggregate circuit flow (cf. Section 3):

$$\text{score}_{\text{eFLOW}}(e_{i,j}; C^t, \mathcal{D}) := a_{C^t}(i, j; \mathcal{D}) \quad (3)$$

That is, the eFLOW picks the edges where more samples in \mathcal{D} flow through, indicating that introducing a new sub-circuit there could potentially better model the distribution over those samples.

Selecting RVs to split. Once we pick edge $e_{i,j}$, we then select the RV among those in the scope of node j . Specifically our vMI heuristics selects the RV X_k sharing more dependencies with the others in the scope. That is, we maximize the score:

$$\text{score}_{\text{vMI}}(X_k; C^t, \mathcal{D}) := \sum_{X_h \neq X_k} \overline{\text{MI}}(X_h; X_k) \quad (4)$$

where $\overline{\text{MI}}$ is the pairwise mutual information estimated on the samples of \mathcal{D} “flowing” through edge $e_{i,j}$. By introducing new parameters for highly dependent RVs we can learn more accurate PCs. The entire local search loop performed by STRUDEL is summarized in Algorithm 2 in Appendix B.

5. Fast Mixtures with STRUDEL

Learning mixtures of PCs greatly improves their performance as density estimators (Vergari et al., 2015; Di Mauro et al., 2015a, 2017; Rahman and Gogate, 2016).

Issues in learning structured-decomposable mixtures. Building a mixture of several PCs results in a joint non-deterministic PC, as it introduces a sum node marginalizing a latent variable. While such a PC would not allow exact MAP inference, it could still be used for queries requiring structured decomposability. However, to answer complex queries like the expectation of predictive models (Khosravi et al., 2019a), one would require a *structured-decomposable mixture of PCs*, i.e., an ensemble whose components are structured-decomposable *and* share the same vtree. *To force such a constraint, while preserving the mixture expressiveness or containing its circuit size, is a non-trivial research question.* Consider learning several PCs with STRUDEL while requiring them to share the same vtree. If we learn each component from a different CLT, compiling them to PCs while enforcing a unique vtree might lead to an exponential blow-up in the size of some PCs. Alternatively, enforcing Algorithm 1 to output a CLT that can be compactly compiled according to a vtree, would result in losing the algorithm’s optimality guarantee.

Shared-structure mixtures. We propose a simpler and faster ensembling strategy which proves to be very effective in practice. *We build ensembles of PCs sharing the same structure*, concretely the structure learned by STRUDEL for single models on that data, while letting each mixture component have different parameters. This has a number of advantages: we (i) need to perform structure learning only once, (ii) materialize a single flow f_M once (as identical structures will generate the same flow), and (iii) can evaluate the likelihood of the whole mixture efficiently, as shown in Eq. 1. This strategy is compatible with classical ensembling schemes such as EM, bagging and boosting. All these scenarios, e.g., each M step in EM, reduce to learning each mixture component parameters on a weighted version of the original data. Determinism allows to do this in closed form (cf. Eq. 2).

6. Experiments

In this section, we rigorously evaluate STRUDEL empirically.³ The natural competitor for STRUDEL is LEARNPSDD, as they both aim to learn PCs with the same structural properties (cf. Section 4). We evaluate both learners as density estimators on a series of 20 standard benchmark datasets. Specifically, we aim to answer the following research questions: **(Q1)** What is the effect of initializing structure learning with a CLT? **(Q2)** How is the splitting heuristic in STRUDEL affecting structure learning? **(Q3)** How do single PCs learned by STRUDEL compare to those learned by LEARNPSDD? **(Q4)** Are ensembles of PCs learned by STRUDEL competitive with LEARNPSDD? **(Q5)** Is our inference approach based on circuit flows speeding up likelihood computations on ensembles of PCs? **(Q6)** Are ensembles of PCs learned by STRUDEL helpful for advanced probabilistic queries?

(Q1) Effect of using CLTs. To evaluate the impact of different initializations in STRUDEL, we compare the CLT initialization scheme we proposed in Section 4.1 against the scheme of LEARNPSDD, where the initial PC is a fully-factorized distribution, normalized for a vtree learned in advance. For both, we learn PCs with up to 1000 splits on every dataset, then report the mean test bits-per-dimension (bpd)⁴ averaged across all datasets as a function of iterations in Figure 2 (left). On

3. Code and experiments are available at <https://github.com/UCLA-StarAI/Strudel>.

4. $\text{bpd}(\theta; \mathcal{D}) = -\sum_{i=1}^{|\mathcal{D}|} \mathcal{L}\mathcal{L}(\theta; x_i) / (\log(2) \cdot |\mathcal{D}| \cdot m)$ where m is the number of features in dataset \mathcal{D} .




Figure 2: *Effect of different initializations and heuristics in STRUDEL.* We report the mean test bits per dimensions (bpd) averaged across all datasets (y-axis) for each iteration (x-axis) as scored by the different initialization schemes (left) and with four possible splitting heuristics (cf. Section 4) in STRUDEL when using the CLT initialization (center). On the right, the per-dataset bpd scores (y-axis) versus total learning time (in secs, x-axis) and learned PC sizes (proportional to circle radiuses times 1000, cf. legend) of the heuristics eFLOW+vMI (blue) and eFLOW+vRAND (red).

average, employing CLTs in STRUDEL not only delivers more accurate initial PCs as expected, but better bpd in the long run. Detailed per-dataset curves can be found in Appendix D.5.

(Q2) Effect of splitting heuristics. We adopt the same setting of Q1 and we mix and match all possible combinations of splitting heuristics: 1) eFLOW-vMI, 2) eFLOW-vRAND, 3) eRAND-vMI, 4) eRAND-vRAND (cf. Section 4.2). Figure 2 (center) reports the mean test bpd per iteration averaged across all datasets. Detailed per-datasets plots can be found in Appendix D.6. It is apparent how selecting edges at random with eRAND delivers suboptimal PCs when compared to eFLOW, regardless of the heuristic to select the RV. On the other hand, eFLOW-vRAND delivers slightly more accurate PCs, on average, than eFLOW-vMI. However, this comes at a high price which is highlighted in Figure 2 (right): circuits learned by eFLOW-vRAND are one or two orders of magnitude larger and each splitting iteration on them is much slower than for eFLOW-vMI. This can be explained as follows: vRAND can greatly increase the PC size by arbitrarily picking a RV far from the root of the sub-circuit selected by eFLOW hence duplicating a larger PC, while vMI picks more informative RVs which generally are closer to the root of such sub-circuit. Given all the above, we employ eFLOW-vMI and the CLT initialization scheme in our remaining experiments.

(Q3) Single models. We perform early stopping after 100 iterations with no improvement for STRUDEL, and LEARNPSDD which we re-run on all datasets by using the hyperparameters reported in Liang et al. (2017). Table 2 in Appendix D.1 reports the mean test log-likelihood for each dataset. For single models, STRUDEL consistently learns more (or equally) accurate PCs than LEARNPSDD on 10 datasets out of 20.⁵ More strikingly, *STRUDEL delivers equivalently accurate PCs sooner than LEARNPSDD.* This is clearly shown in Figure 3 (left & center) for all datasets: it takes fewer iterations for STRUDEL to achieve comparable bpd and each of its iterations takes at least one order of magnitude less time than LEARNPSDD. Lastly, PCs learned with STRUDEL are still comparable on many datasets to other PC learners performing local search like selective SPNs (SELSPN) (Peharz et al., 2014) and normal SPNs (SEASPN) (Dennis and Ventura, 2015). Note however, that these competitors learn PCs with less structural requirements (SELSPNs are not structured-decomposable and SEASPNs are not deterministic), and hence support less tractable inference scenarios (cf. Section 2). Detailed learning times, circuit sizes and statistical tests are reported in Appendix D.

5. The likelihoods we are reporting for LEARNPSDD are significantly better than those originally reported in Liang et al. (2017). Compared to those original results, STRUDEL is more accurate than LEARNPSDD 16 times out of 20.

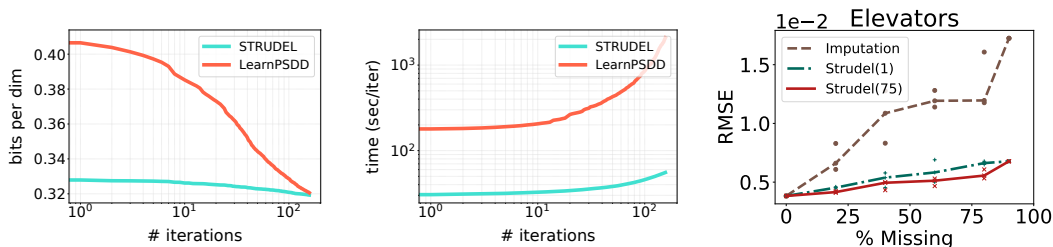


Figure 3: *Comparing STRUDEL and LEARNPSDD accuracy and learning times.* We report the mean bpd averaged across all datasets (y-axis) (left), and seconds per iteration during learning (y-axis) (center) for each iteration (x-axis) as scored by single models learned by STRUDEL (blue) and LEARNPSDD (red). *Expected prediction benchmarks.* On the right, the Root Mean Square Error (RMSE) for median imputation, STRUDEL and STRUDEL-Ensemble models (number in parenthesis shows the number of components in ensemble models).

(Q4) Mixtures with STRUDEL. We evaluate the shared-structure mixtures of PCs learned by STRUDEL by reusing for each dataset the best single PC learned for setting **Q1**, and performing parameter learning with EM (STRUDEL-EM) or a combination of EM and bagging (STRUDEL-BEM) to alleviate overfitting as in Liang et al. (2017). For STRUDEL-EM, we choose the number of mixture components k_{EM} by a grid search on $\{2, 5, 10, 15, 20, 25, 30\}$. STRUDEL-BEM instead trains mixtures of PCs using EM on 10 different bagged datasets. Table 9 in Appendix E reports the mean test log-likelihoods of our mixtures and the corresponding ones learned by LEARNPSDD. On 15 datasets STRUDEL-EM outperforms LEARNPSDD-EM, while STRUDEL-BEM is more accurate than its counterpart 11 times. This is remarkable if one notes that PCs in LEARNPSDD-(B)EM are allowed to take arbitrary structures and update the structures for each component during learning. As expected, STRUDEL drastically reduces the learning times of large mixtures, as single PCs can be learned much faster, and mixtures can be learned in a fraction of the time by virtue of shared flows (cf. Section 3).

(Q5) Effectiveness of flows. Appendix C reports detailed timings comparing our circuit flow approach to classical circuit evaluation for computing mixture likelihoods. The circuit flow approach is orders of magnitude faster (around 10^2 – 10^3), and up to 4591 times faster on the ‘msnbc’ dataset.

(Q6) Advanced probabilistic queries. Finally, we evaluate how PCs learned with STRUDEL can be exploited for advanced inference scenarios requiring structured decomposability. We adopt the experimental setting of Khosravi et al. (2019a) aiming to compute the expected predictions of a regressor r w.r.t. a generative model represented as a structured-decomposable PC sharing the same vtree of r . We compute the expected predictions of regressors learned on 4 different benchmarks and employ a single PC or a mixture learned by STRUDEL-BEM with 5 bags and a number of EM components cross-validated in $\{5, 10, 15, 20\}$. Figure 3 (right) and Figure 4 in Appendix F show the RMSE of our models for different percentages of missing values, when compared to common imputation schemes like median imputation. Not only do single PC perform better than the baseline, but the cheap mixtures of PCs with shared structure help further reduce the error.

7. Conclusions

In this paper we introduced STRUDEL to learn structured-decomposable PCs in a fast and simple way. STRUDEL delivers accurate single PCs in a fraction of the time of its competitor and effectively

scales up to learning mixtures of PCs sharing the same structure. We consider STRUDEL as an initial stepping stone to learn PCs for several application scenarios where advanced probabilistic inference is required and out-of-the-scope of the current landscape of tractable probabilistic models.

Acknowledgments This work is partially supported by NSF grants #IIS-1943641, #IIS-1633857, #CCF-1837129, DARPA grant #N66001-17-2-4032, a Sloan Fellowship, Intel, and Facebook.

References

- J. Bekker, J. Davis, A. Choi, A. Darwiche, and G. Van den Broeck. Tractable learning for complex probability queries. In *NeurIPS*, 2015.
- H. Chan and A. Darwiche. On the robustness of most probable explanations. In *UAI*, 2006.
- A. Choi, D. Kisa, and A. Darwiche. Compiling probabilistic graphical models using sentential decision diagrams. In *ECSQARU*, 2013.
- A. Choi, G. Van den Broeck, and A. Darwiche. Tractable learning for structured probability spaces: A case study in learning preference distributions. In *IJCAI*, 2015.
- Y. Choi and G. Van den Broeck. On robust trimming of bayesian network classifiers. In *IJCAI*, 2018.
- Y. Choi, A. Darwiche, and G. Van den Broeck. Optimal feature selection for decision robustness in bayesian networks. In *IJCAI*, 2017.
- Y. Choi, A. Vergari, and G. Van den Broeck. Probabilistic circuits: A unifying framework for tractable probabilistic modeling. 2020.
- C. K. Chow and C. N. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Trans. Inf. Theory*, 1968.
- A. Darwiche. A differential approach to inference in Bayesian networks. *JACM*, 2003.
- R. Dechter and R. Mateescu. And/or search spaces for graphical models. *AI*, 2007.
- A. Dennis and D. Ventura. Greedy Structure Search for Sum-product Networks. In *IJCAI*, 2015.
- N. Di Mauro, A. Vergari, and T. M. Basile. Learning bayesian random cutset forests. In *ISMIS*. 2015a.
- N. Di Mauro, A. Vergari, and F. Esposito. Learning accurate cutset networks by exploiting decomposability. In *AI*IA 2015, Advances in Artificial Intelligence*. 2015b.
- N. Di Mauro, A. Vergari, T. M. Basile, and F. Esposito. Fast and accurate density estimation with extremely randomized cutset networks. In *ECML-PKDD*. Springer, 2017.
- P. Khosravi, Y. Choi, Y. Liang, A. Vergari, and G. Van den Broeck. Tractable computation of the moments of predictive models. In *NeurIPS*, 2019a.
- P. Khosravi, Y. Liang, Y. Choi, and G. Van den Broeck. What to expect of classifiers? Reasoning about logistic regression with missing features. In *IJCAI*, 2019b.

- D. Kisa, G. Van den Broeck, A. Choi, and A. Darwiche. Probabilistic sentential decision diagrams. In *KR*, 2014.
- D. Koller and N. Friedman. Probabilistic graphical models: principles and techniques. 2009.
- Y. Liang and G. Van den Broeck. Towards compact interpretable models: Shrinking of learned probabilistic sentential decision diagrams. In *XAI workshop*, 2017.
- Y. Liang and G. Van den Broeck. Learning logistic circuits. In *AAAI*, 2019.
- Y. Liang, J. Bekker, and G. Van den Broeck. Learning the structure of probabilistic sentential decision diagrams. In *UAI*, 2017.
- D. Lowd and P. Domingos. Learning arithmetic circuits. In *UAI*, 2008.
- M. Meilă and M. I. Jordan. Learning with mixtures of trees. *JMLR*, 2000.
- U. Oztok and A. Darwiche. A top-down compiler for sentential decision diagrams. In *IJCAI*, 2015.
- G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- R. Peharz, R. Gens, and P. Domingos. Learning selective sum-product networks. In *Workshop on LTPM*. 2014.
- R. Peharz, S. Lang, A. Vergari, K. Stelzner, A. Molina, M. Trapp, G. V. d. Broeck, K. Kersting, and Z. Ghahramani. Einsum networks: Fast and scalable learning of tractable probabilistic circuits. In *ICML*, 2020.
- K. Pipatsrisawat and A. Darwiche. New compilation languages based on structured decomposability. In *AAAI*, 2008.
- H. Poon and P. Domingos. Sum-Product Networks: a New Deep Architecture. *UAI*, 2011.
- T. Rahman and V. Gogate. Learning ensembles of cutset networks. In *AAAI*, 2016.
- T. Rahman, P. Kothalkar, and V. Gogate. Cutset networks: A simple, tractable, and scalable approach for improving the accuracy of chow-liu trees. In *ECML-PKDD*, 2014.
- Y. Shen, A. Choi, and A. Darwiche. Tractable operations for arithmetic circuits of probabilistic models. In *NeurIPS*, 2016.
- Y. Shen, A. Choi, and A. Darwiche. A tractable probabilistic model for subset selection. In *UAI*, 2017.
- A. Vergari, N. Di Mauro, and F. Esposito. Simplifying, Regularizing and Strengthening Sum-Product Network Structure Learning. In *ECML-PKDD*, 2015.
- A. Vergari, N. Di Mauro, and F. Esposito. Visualizing and understanding sum-product networks. *Machine Learning Journal*, 2019.
- A. Vergari, Y. Choi, R. Peharz, and G. Van den Broeck. Probabilistic circuits: Representations, inference, learning and applications. *AAAI Tutorial*, 2020.

Appendix A. Learning a Chow-Liu Tree

We list the algorithm for learning Chow-Liu Trees (Chow and Liu, 1968) here.

Algorithm 1: LearnCLT($\mathcal{D}, \mathbf{X}, \alpha$)

Input : a dataset \mathcal{D} over RVs, $\mathbf{X} = \{X_i\}_{i=1}^n$, Laplace smoothing factor α
Output : a Chow-Liu tree model $\langle \mathcal{T}, \theta = \{\theta_{i|\text{Pa}_i}\}_{i=1}^n \rangle$ estimating $p(\mathbf{X})$

- 1 $\text{MI} \leftarrow \mathbf{0}_{n \times n}$
- 2 **for each** $X_i, X_j \in \mathbf{X}$ **do**
- 3 $\text{MI}_{ij} \leftarrow \text{estimateMI}(\mathcal{D}, X_i, X_j, \alpha)$
- 4 **end**
- 5 $T \leftarrow \text{maximumSpanningTree}(\text{MI})$
- 6 $\mathcal{T} \leftarrow \text{traverseTree}(T)$
- 7 $\theta \leftarrow \{\theta_{i,\text{Pa}_i} \leftarrow \text{estimateCPT}(\mathcal{D}, X_i, X_{\text{Pa}_i}, \alpha)\}$
- 8 **return** $\langle \mathcal{T}, \theta \rangle$

Appendix B. STRUDEL Algorithm

The complete pseudocode for STRUDEL is listed in Algorithm 2.

Algorithm 2: STRUDEL(\mathcal{D}, \mathbf{X})

Input : a dataset \mathcal{D} over RVs \mathbf{X}
Output : a structured-decomposable PC C

- 1 $\mathcal{T} \leftarrow \text{LearnCLT}(\mathcal{D}, \mathbf{X})$
- 2 $\mathcal{V} \leftarrow \text{getVTree}(\mathcal{T})$
- 3 $C \leftarrow \text{compile}(\mathcal{T}, \mathcal{V})$
- 4 **while** C is not overfitting **do**
- 5 $e_{i,j}^* \leftarrow \text{argmax}_{e_{i,j} \in \text{edges}(C)} \text{score}_{\text{eFLOW}}(e_{i,j}; C, \mathcal{D})$
- 6 $X^* \leftarrow \text{argmax}_{X_k \in \phi(C_i)} \text{score}_{\text{vMI}}(X_k; C, \mathcal{D})$
- 7 $C \leftarrow \text{SplitOperation}(C, e_{i,j}^*, X^*; \mathcal{D})$
- 8 **return** C

Appendix C. Circuit Flows for Fast Inference

To empirically show that our shared circuit flows implementation (Section 3) benefits the efficient evaluation of ensembles, compared a vectorized version of the classical algorithm that evaluates the circuit bottom-up (Darwiche, 2003). We report the time taken to compute likelihoods for an ensemble as a function of the number of components in Table 1.

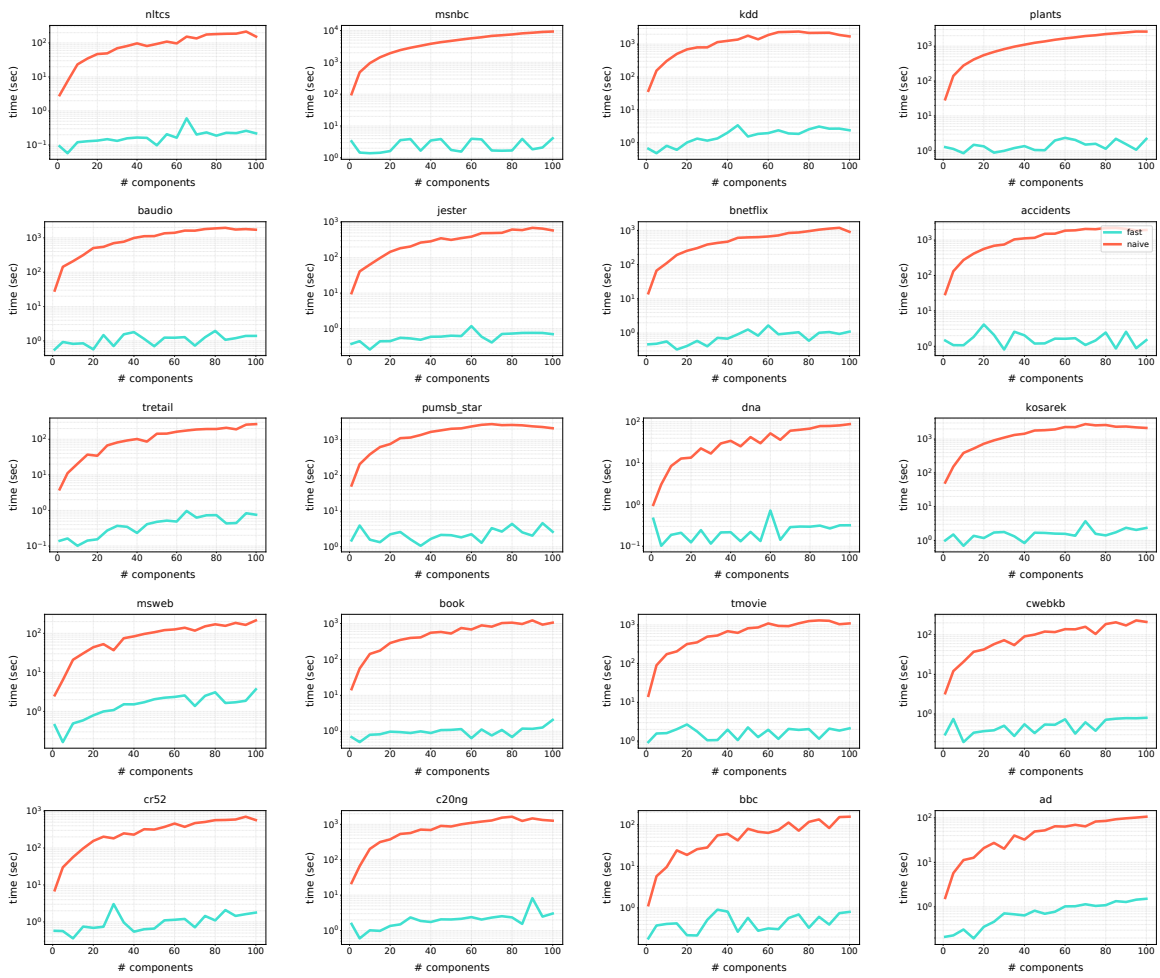


Table 1: *Circuit flows for fast inference of ensembles for each dataset.* We report the time (seconds) used evaluate the ensemble circuits on each dataset (y-axis) for different number of component in the ensemble (x-axis), comparing the circuit flows implementation (blue) and the classical algorithm of naively evaluating the circuits bottom-up (red).

Appendix D. Single Models with STRUDEL

D.1 Density Estimation Performance

We compare STRUDEL and LEARNPSDD as density estimators. The average test-set log-likelihood is reported in Table 2.

DATASET	LEARNPSDD	STRUDEL	SELSPN	SEASPN
NLTCS	-6.03	-6.06	-6.03↓	-6.07↑
MSNBC	-6.04	-6.05	-6.04↓	-6.06↑
KDD	-2.17	-2.17	-2.16↓	-2.16↓
PLANTS	-13.49	-13.72	-12.97↓	-13.12↓
AUDIO	-41.51	-42.26	-41.23↓	-40.13↓
JESTER	-54.63	-55.30	-54.38↓	-53.08↓
NETFLIX	-58.53	-58.68	-57.98↓	-56.91↓
ACCIDENTS	-28.291	-29.46	-26.88↓	-30.02↓
RETAIL	-10.92	-10.90	-10.88↓	-10.97↑
PUMSB-STAR	-25.40	-25.28	-22.66↓	-28.69↓
DNA	-83.02	-87.10	-80.44↓	-81.76↓
KOSAREK	-10.99	-10.98	-10.85↓	-11.00↓
MSWEB	-9.93	-10.19	-9.93↓	-10.25↓
BOOK	-36.06	-35.77	-36.01↓	-34.91↓
EACHMOVIE	-55.41	-59.47	-55.73↓	-53.28↓
WEBKB	-161.42	-160.50	-158.52↓	-157.88↓
ROUTERS-52	-93.30	-92.38	-88.48↓	-86.38↓
20NEWS-GRP	-160.43	-160.77	-158.68↓	-153.63↓
BBC	-260.24	-258.96	-259.35↑	-253.13↓
AD	-20.13	-16.52	-16.94↑	-16.77↑

Table 2: *Density estimation benchmarks: single models*. Average test log-likelihood for STRUDEL and LEARNPSDD models. The bold values indicate STRUDEL is better than or statistically equivalent with (cf. Appendix D.4) LEARNPSDD. On the right other state-of-the-art structure learners, which are *not* targeting structured-decomposable circuits (see text). ↑ (resp. ↓) indicates that STRUDEL is more accurate (resp. less accurate).

D.2 Circuit Sizes

We compare the circuit size of models learned by LEARNPSDD and STRUDEL in Table 3. As expected, STRUDEL delivers larger circuits but almost always on the same order of magnitude, which makes sense since STRUDEL does not penalize on the circuit size.

DATASET	LEARNPSDD	STRUDEL
NLTCS	1304	4373
MSNBC	5465	20795
KDD	2915	6572
PLANTS	11583	119194
AUDIO	18208	55675
JESTER	11322	27713
NETFLIX	10997	27173
ACCIDENTS	8418	75363
RETAIL	2989	3979
PUMSB-STAR	8298	108876
DNA	3068	17507
KOSAREK	7173	37583
MSWEB	6581	2347
BOOK	10978	54839
EACHMOVIE	20648	123996
WEBKB	11033	25406
ROUTERS-52	10410	36343
20NEWS-GRP	15793	58749
BBC	12335	29532
AD	12238	13152

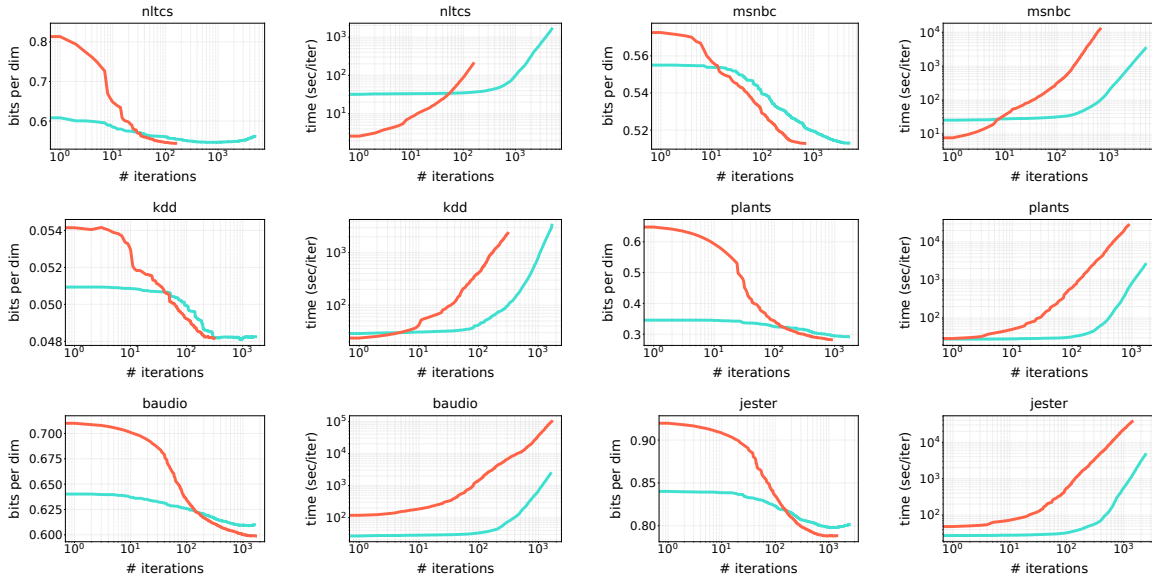
Table 3: *Density estimation benchmarks: single models.* The circuit sizes of best learned single models for LEARNPSDD and STRUDEL. As expected, STRUDEL learns larger PCs than LEARNPSDD since the latter explicitly penalizes larger sizes via its likelihood score. Nevertheless, the increase in size is contained with the vMI heuristics (cf. Figure 2 (right)) and is a negligible price to pay for a much faster learner (cf. Table 4).

D.3 Learning Times

To compare the efficiency of LEARNPSDD and STRUDEL, we reproduce the experimental setting of LEARNPSDD and rerun the single model experiments with the latest version of the LEARNPSDD code. Here we report the learning times – both the seconds per iteration (SEC PER ITER) and total seconds during learning (SEC SUM) – in Table 4, from which it is clear that STRUDEL is more efficient than LEARNPSDD.

DATASET	LEARNPSDD	STRUDEL	LEARNPSDD	STRUDEL
	SEC/ITER	SEC/ITER	SEC TOT	SEC TOT
NLTCs	1.3	0.1	193.0	42.1
MSNBC	18.6	0.5	12573.9	1930.3
KDD	7.5	0.4	2325.3	131.4
PLANTS	31.2	1.3	27438.4	1911.3
AUDIO	59.0	0.8	98457.1	848.4
JESTER	23.0	0.4	23573.1	279.6
NETFLIX	52.3	0.3	51300.7	436.4
ACCIDENTS	54.0	0.8	29664.2	3393.9
RETAIL	66.5	0.1	16018.3	44.2
PUMSB-STAR	83.7	1.1	40508.6	4737.3
DNA	3.0	0.5	595.4	86.6
KOSAREK	45.8	0.4	30196.8	398.9
MSWEB	176.6	4.3	93589.7	12.9
BOOK	28.2	1.1	32567.2	2316.8
EACHMOVIE	31.1	0.7	61792.3	1427.2
WEBKB	5.6	0.4	6333.0	884.3
ROUTERS-52	93.2	0.5	98796.0	1561.0
20NEWS-GRP	34.6	0.7	53829.2	2593.2
BBC	6.2	0.9	8000.8	2136.3
AD	3.7	1.1	3748.6	300.8

Table 4: Total times (in seconds) taken (SEC TOT) and averaged times per iteration (SEC/ ITER) to learn the best single models on each dataset for LEARNPSDD and STRUDEL. STRUDEL requires a fraction of the time of LEARNPSDD greatly speeding up learning.



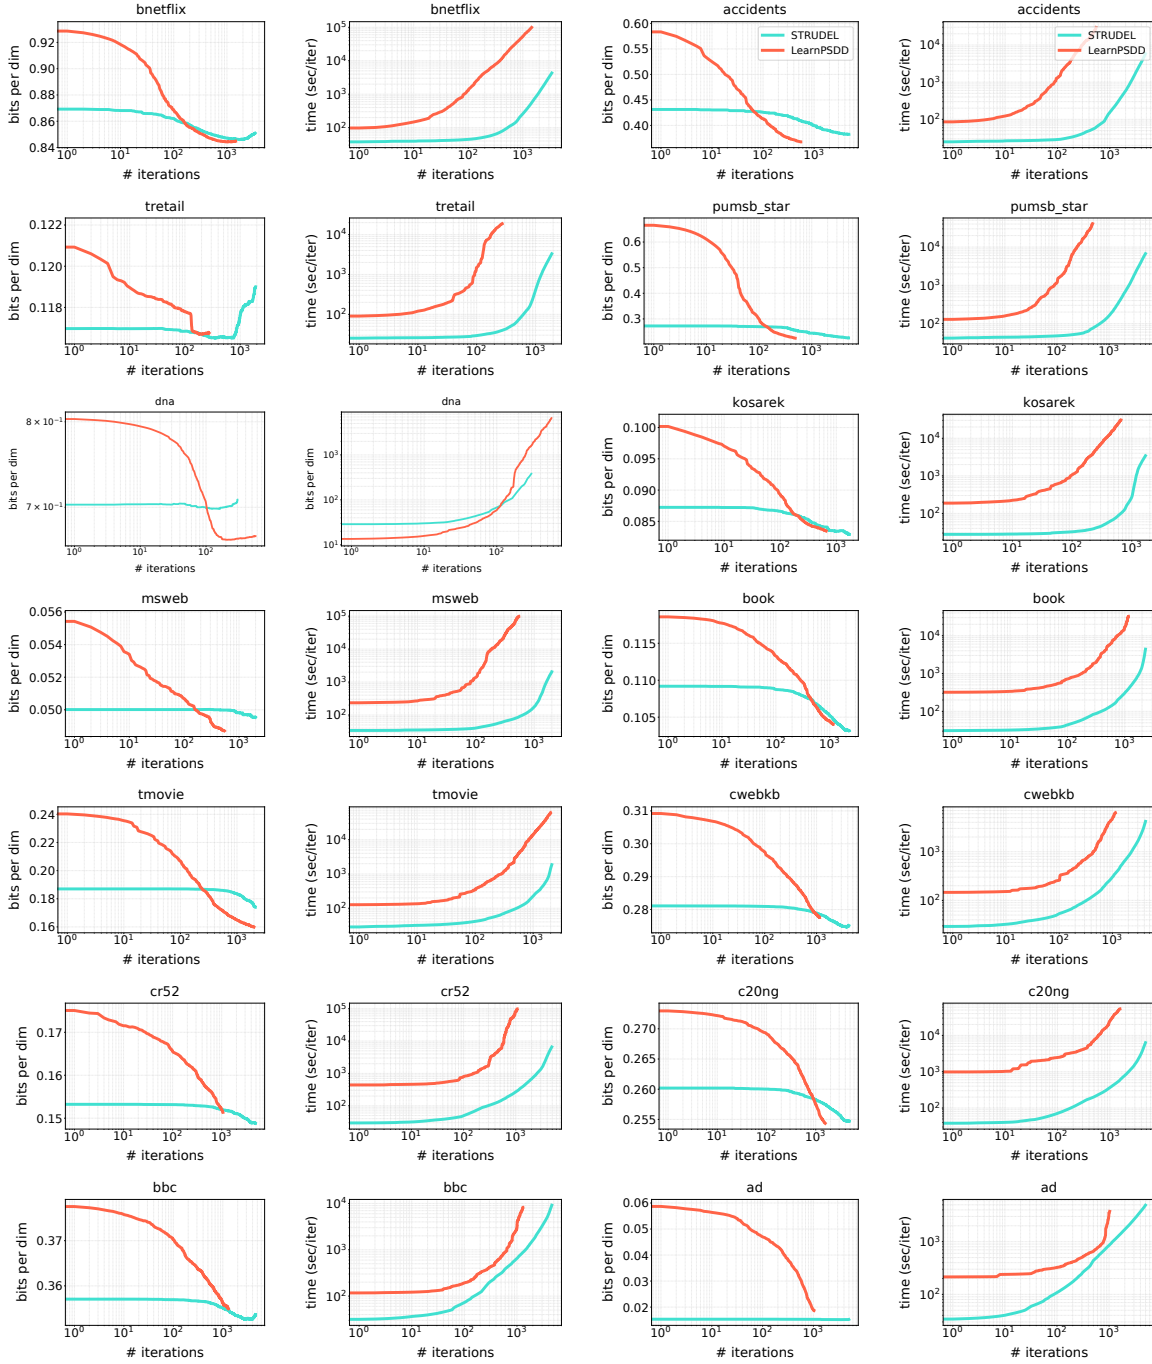


Table 5: *Comparing STRUDEL and LEARNPSDD single models performance and learning times.* For each of the 20 benchmark datasets, we report 1) the test bits per dimensions (bpd) (y-axis) for each iteration (x-axis) on column 2 and column 4, and 2) seconds per iteration during learning (y-axis) for each iteration (x-axis) on column 2 and column 4 as scored by STRUDEL (blue) and LEARNPSDD (red). The plots for the same dataset is aligned next to each other as comparison.

D.4 Statistical Tests

Since we only have the log-likelihood per sample available for SelSPN, we run pairwise Wilcoxon signed rank tests to compare STRUDEL and SelSPN more rigorously. The p-values are reported in Table 6, in 4 out of 20 datasets, the results are statistically equivalent

DATASET	LEARNPSDD	SELSPN
NLTCS	5.34E-02	1.35E-02
MSNBC	0.00E+00	0.00E+00
KDD	0.00E+00	0.00E+00
PLANTS	6.58E-25	3.46E-56
AUDIO	1.67E-30	4.73E-38
JESTER	2.37E-45	2.05E-20
NETFLIX	5.29E-04	1.38E-08
ACCIDENTS	1.46E-102	3.08E-279
RETAIL	2.48E-01	6.54E-05
PUMSB-STAR	1.52E-12	2.83E-269
DNA	8.39E-164	2.95E-189
KOSAREK	5.17E-03	5.93E-21
MSWEB	2.76E-100	7.06E-114
BOOK	6.48E-02	7.60E-01
EACHMOVIE	2.42E-53	2.33E-45
WEBKB	1.08E-01	4.07E-01
ROUTERS-52	3.56E-04	2.51E-87
20NEWS-GRP	1.40E-02	2.37E-07
BBC	1.96E-04	4.33E-01
AD	9.90E-65	1.64E-09

Table 6: Pairwise Wilcoxon signed rank test p-values for the comparisons of test log-likelihoods for STRUDEL against LEARNPSDD SelSPN on all datasets. Bold values indicate the results are not statistically significant when picking confidence value 99%.

D.5 Initializations

Here we report the effect of different initialization methods (CLT and independent) on each dataset in Table 7.

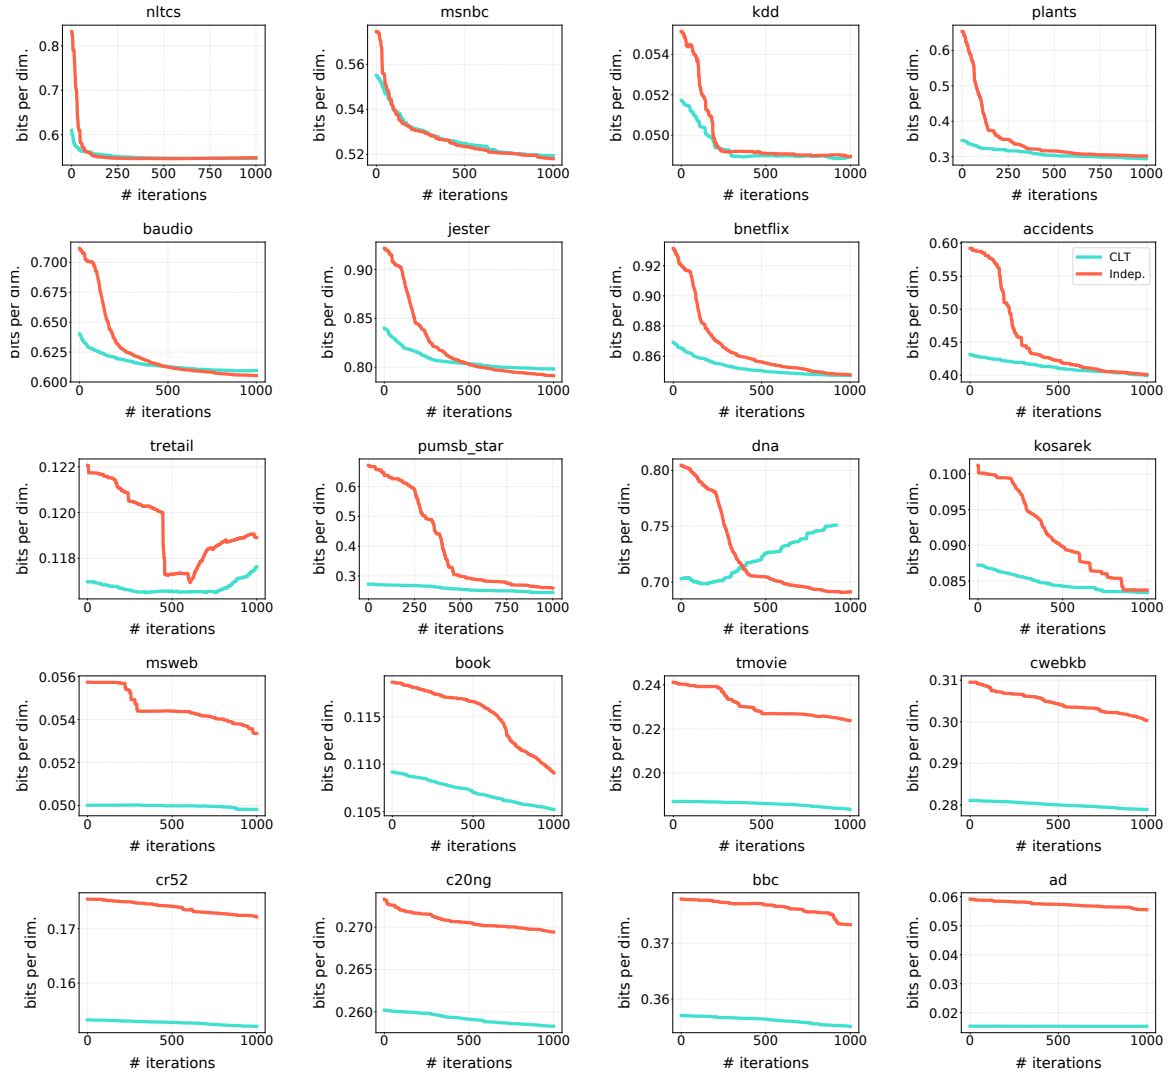


Table 7: *Effect of different initializations in STRUDEL for each dataset.* We report the mean test bits per dimensions (bpd) on each dataset (y-axis) for each iteration (x-axis) as scored by the different initializations: CLT (blue) and independent (red).

D.6 Heuristics

Here we report the effect of different heuristics on each dataset in Table 8.

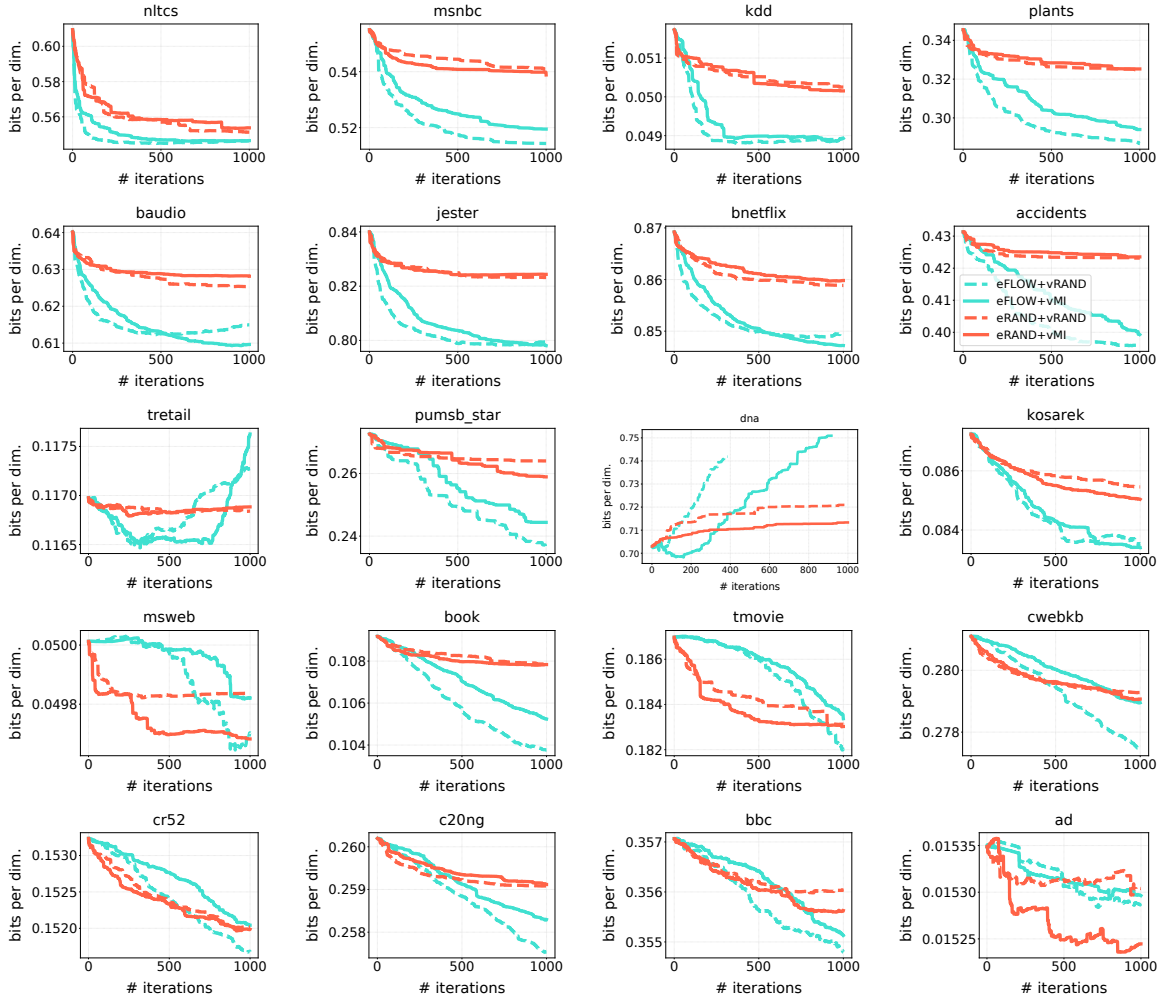


Table 8: *Effect of different heuristics in STRUDEL for each dataset.* We report the mean test bits per dimensions (bpd) on each dataset (y-axis) for each iteration (x-axis) as scored by the 4 combinations of different heuristics.

Appendix E. Mixtures with STRUDEL

We report the test set log-likelihood result in Table 9 of mixtures with STRUDEL and mixtures with LEARNPSDD.

DATASET	LEARNPSDD	STRUDEL	LEARNPSDD	STRUDEL
	EM	EM	BEM	BEM
NLTCS	-6.03↓	-6.07	-5.99 ↓	-6.06
MSNBC	-6.04 ↑	-6.04	-6.04 ↑	-6.04
KDD	-2.12↓	-2.14	-2.11 ↓	-2.13
PLANTS	-13.79↑	-13.22	-13.02↑	-12.98
AUDIO	-41.98↑	-41.20	-39.94 ↓	-41.50
JESTER	-53.47↓	-54.24	-51.29 ↓	-55.03
NETFLIX	-58.41↑	-57.93	-55.71 ↓	-58.69
ACCIDENTS	-33.64↑	-29.05	-30.16↑	-28.73
RETAIL	-10.81↓	-10.83	-10.72 ↓	-10.81
PUMSB-STAR	-33.67↑	-24.39	-26.12↑	-24.12
DNA	-92.67↑	-87.15	-88.01↑	-86.22
KOSAREK	-10.81↑	-10.70	-10.52 ↓	-10.68
MSWEB	-9.97↑	-9.74	-9.89↑	-9.71
BOOK	-34.97↑	-34.49	-34.97↓	-34.99
EACHMOVIE	-58.01↑	-53.72	-58.01↑	-53.67
WEBKB	-161.09↑	-154.83	-161.09↑	-155.33
ROUTERS-52	-89.61↑	-86.35	-89.61↑	-86.22
20NEWS-GRP	-160.09↑	-153.87	-155.97↑	-154.47
BBC	-253.19↓	-256.53	-253.19 ↓	-254.41
AD	-31.78↑	-16.52	-31.78↑	-16.38

Table 9: *Density estimation benchmarks: ensembles*. Average test log-likelihood for STRUDEL and LEARNPSDD models. Two versions of ensembles (EM vs. BEM) are compared separately, ↑ (resp. ↓) indicates that STRUDEL is more accurate (resp. less accurate). Bold values indicates the best on certain dataset over 4 methods.

Appendix F. Expected predictions

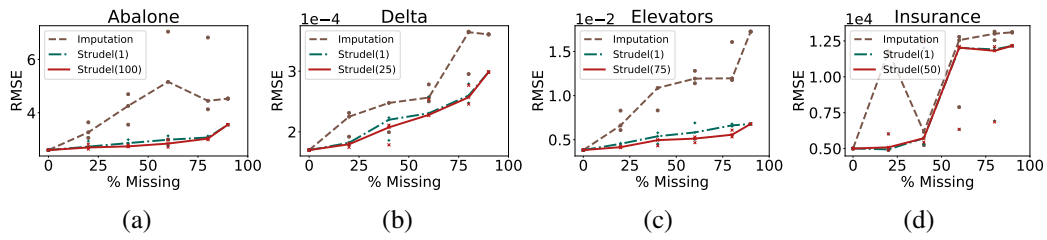


Figure 4: *Expected prediction benchmarks*. Root Mean Square Error (RMSE) for median imputation, STRUDEL and STRUDEL-Ensemble models (number in parenthesis shows the number of components in the ensemble models).