
HyperSPNs: Compact and Expressive Probabilistic Circuits

Andy Shih Stanford University andyshih@cs.stanford.edu	Dorsa Sadigh Stanford University dorsa@cs.stanford.edu	Stefano Ermon Stanford University ermon@cs.stanford.edu
---	---	--

Abstract

Probabilistic circuits (PCs) are a family of generative models which allows for the computation of exact likelihoods and marginals of its probability distributions. PCs are both expressive and tractable, and serve as popular choices for discrete density estimation tasks. However, large PCs are susceptible to overfitting, and only a few regularization strategies (e.g., dropout, weight-decay) have been explored. We propose HyperSPNs: a new paradigm of generating the mixture weights of large PCs using a small-scale neural network. Our framework can be viewed as a soft weight-sharing strategy, which combines the greater expressiveness of large models with the better generalization and memory-footprint properties of small models. We show the merits of our regularization strategy on two state-of-the-art PC families introduced in recent literature – RAT-SPNs and EiNETs – and demonstrate generalization improvements in both models on a suite of density estimation benchmarks in both discrete and continuous domains.

1 Introduction

One of the core motivations for building models of probability distributions from data is to *reason* about the data distribution. For example, given a model of the distribution over driving routes, we may like to reason about the probability that a driver will take a certain route given that there is congestion on a certain street. Or, given a distribution over weather conditions, we may like to reason about the probability of it raining k times in the next year. Traditional probabilistic models such as Categorical/Gaussian Mixture Models or Hidden Markov Models (HMMs) are well-equipped to answer these types of queries with exact probabilities that are consistent with the model distribution [1]. However, their learned model distribution might not accurately approximate the data distribution due to their simplicity.

On the other end of the spectrum, the modern wave of deep generative models has largely focused on learning accurate approximations of the true data distribution, at the cost of tractability. High capacity models such as autoregressive or flow models can mimic the true data distribution with great fidelity [33, 13, 22]. However, they are not designed for reasoning about the probability distribution beyond computation of likelihoods. More illustrative of this trend are GANs and EBMs, which are expressive enough to sample high-quality images, but give up even the ability to compute exact likelihoods [9, 14].

To maintain the ability to reason about the distribution induced by our model, we need to explore within tractable probabilistic models. Of these model families, probabilistic circuits (e.g. Sum Product Networks) are one of the most expressive and general, subsuming shallow mixture models, and HMMs while maintaining their tractable properties [6, 26, 34]. Their expressive power comes from their hierarchical / deep architecture, allowing them to express a large number of modes in their distribution. Their tractability comes from global constraints imposed in their network structure, enabling efficient and exact computation of likelihoods, marginals, and more. Probabilistic circuits

(PC) are a popular choice for density estimation [35, 18] and approximate inference in discrete settings [30].

Given the status of probabilistic circuits as one of the most expressive models among tractable probabilistic model families, recent works have looked into pushing the limits of expressivity of probabilistic circuits [15, 25, 24]. Naturally, the bigger and deeper the probabilistic circuit, the greater the expressiveness of the model. This has led to a trend of building large probabilistic circuits via design choices such as tensorizing model parameters, forgoing structure learning, and more. For example, [24] report the training of probabilistic circuits with 9.4M parameters.

However, larger probabilistic circuits are more susceptible to overfitting. In addition, the so-called double-descent phenomenon [19] (wherein highly overparameterized models exhibit low generalization gap) has yet to be observed in probabilistic circuits. This puts even more importance on effective regularization strategies for training large probabilistic circuits. Recent works have used dropout [25] and weight decay [37], but there has not been much exploration of other regularization strategies.

In this work, we propose HyperSPNs, which regularize by aggressively limiting the degrees of freedom of the model. Drawing inspiration from HyperNetworks [11], we generate the weights of the probabilistic circuit via a small-scale external neural network. More precisely, we partition the parameters (mixture weights) of the PC into several sectors. For each sector, we learn a low-dimensional embedding, then map that to the parameters of the sector using the neural network. The generated parameters still structurally form a PC, so we retain the same ability to reason about the probability distribution induced by the HyperSPN.

HyperSPNs combine the greater expressiveness of large probabilistic circuits with the better generalization of models with fewer degrees of freedom. The external neural network has much fewer parameters than the original PC, effectively regularizing the PC through a soft weight-sharing strategy. As a bonus, the memory requirement for storing the model is much smaller, and the memory requirement for evaluating the model can also be drastically reduced by materializing sectors of the PC “on the fly”. We verify the empirical performance of HyperSPNs on density estimation tasks for the Twenty Datasets benchmark [12], the Amazon Baby Registries benchmark [8], and the Street View House Numbers (SVHN) [20] dataset, showing generalization improvements over competing regularization strategies in all three settings. Our results suggest that HyperSPNs are a promising framework for training probabilistic circuits.

2 Background


Probabilistic circuits (PCs) are a family of generative models known for their tractability properties [6, 26, 34]. They support efficient and exact inference routines for computing likelihoods and marginals/conditionals of their probability distribution. The internals of a PC consist of a network (edge-weighted Directed Acyclic Graph) of sum and product nodes, stacked in a deep/hierarchical manner and evaluated from leaves to root. PCs gain their tractability by enforcing various structural constraints on their network, such as decomposability, which restricts children of the same product node to have disjoint scopes [5, 6]. Different combination of constraints leading to different types of PCs, including Arithmetic Circuits [6], and most notably Sum Product Networks (SPNs) [26].

The probability of an input \mathbf{x} on the PC is defined to be the value $g_{\text{root}}(\mathbf{x})$ obtained at the root of the PC after evaluating the network topologically bottom-up. Internally, edges leading into a sum node are weighted with mixture probabilities α (that sum to 1), and edges leading into a product node are unweighted. Leaf nodes at the base of the network contain leaf distributions l . The nodes of the PC are evaluated as follows, with $ch(i)$ denoting the children of node i :

$$g_i(\mathbf{x}) = \begin{cases} l_i(\mathbf{x}) & i \text{ is leaf node} \\ \sum_{j \in ch(i)} \alpha_{ij} g_j(\mathbf{x}) & i \text{ is sum node} \\ \prod_{j \in ch(i)} g_j(\mathbf{x}) & i \text{ is product node} \end{cases}$$

As long as the leaf distributions are tractable and the network structure satisfies the constraints of decomposability and smoothness, the output at g_{root} is guaranteed to be normalized, and corresponding marginals/conditionals can be evaluated with linear-time exact inference algorithms [5, 6].

The network structure of a PC refers to the set of nodes and edge connections in the DAG, excluding the edge weights. There are various ways of constructing the network structure of PCs, e.g., through



(a) We show the structure of a RAT-SPN over $n = 4$ variables using layer size $k = 3$. Solid edges are weighted with trainable parameters, and dashed edges are unweighted. Each fully connected layer of solid edges can be treated as a sector.

(b) Counting the # of parameters, given the number of variables n and layer size k .

Figure 1: Given a RAT-SPN, we can partition its set of parameters into sectors. The trainable parameters (solid edges) in Figure 1a can be grouped into fully-connected layers of size $1 \times k$, $k \times k$, and $k \times 2$, oriented in a binary-tree structure. In Figure 1b, we visualize the sectors as blocks, counting the number of blocks in each layer of the binary tree to arrive at the total parameter count. The sector abstraction is a key ingredient in implementing the HyperSPN’s soft weight-sharing strategy.

the use of domain knowledge or through search-based methods [6, 17, 7, 28]. Recently, there has been a trend of abstracting away individual sum and product nodes, and instead dealing directly with layers that satisfy the structural constraints necessary for tractability [25, 24, 30]. Prescribing a network structure through a stacking of these layers leads to memory and hardware efficiencies that allow us to scale up the training of PCs to millions of edges.

Given a PC with a fixed network structure, we can learn its edge weights to best fit the probability distribution of data. We will refer to the full set of mixture probabilities leading into sum nodes in the PC as its parameters. These parameters can be learned using methods such as Expectation-Maximization (EM) [23], signomial program formulations [38], or simply via (projected) gradient descent. EM has been a well-studied choice for optimizing PCs [23, 24], although it requires hand-derived updates that can be tedious and sometimes even incorrect [26]. In this work, we rely on gradient descent, which is straightforward and works out-of-the-box on new families of models [37].

Learning a PC with a large number of parameters can be prone to overfitting. One regularization strategy that has been proposed is “probabilistic dropout” [25], which takes inspiration from the standard dropout technique in neural networks [31] and randomly ignores a subset of edges leading into sum nodes. Weight decay has also been commonly used as another form of regularization [37]. Unfortunately, there has been little exploration done around regularization strategies, and some current approaches only work in limited settings (e.g., dropout has been limited to settings that use PCs as a discriminative classifier [25]).

2.1 RAT-SPNs

Two particular PC structures have been recently proposed, namely RAT-SPNs [25] and EiNETs [24], both of which scale to millions of parameters. They are conceptually similar, the main difference being that EiNETs are designed for better GPU performance. To ground our discussion around regularizing large PCs, we will describe our methods just on RAT-SPNs, although we also apply our methods on EiNETs in later experiments too. In RAT-SPNs, the input variables are organized into a binary tree, with one leaf in the binary tree corresponding to each variable. We start from the base




Figure 2: HyperSPN: we introduce a small-scale external neural network that generates the parameters of the underlying SPN. Using the sector abstraction from Figure 1, we learn an embedding of dimension h for each of the $2n - 1$ sectors. The neural network maps each embedding to the parameters of its corresponding sector, which materializes the underlying SPN.

of the SPN and construct its full structure by stacking sum layers and product layers. For each leaf node in the binary tree, we create one fully connected sum layer on top of the SPN leaves. For each internal node in the binary tree, we create one product layer that merges its two children layers in the SPN, and then stack one fully connected sum layer on top. We illustrate this construction in Figure 1. Finally, RAT-SPNs create so-called replicas of this structure, by repeating the construction but with permuted orderings of input variables in the binary tree leaves. The replicas are then combined under one top-level mixture distribution (for readability, Figure 1 does not show any replicas).

The size of each sum layer in the RAT-SPN is determined by a single constant k . For convenience, we also let the size of each product layer be k . This fully specifies the structure of our SPN. For the rest of the paper, it is helpful to abstractly group the parameters of the RAT-SPN into disjoint sectors. Corresponding to the root of the binary tree, we have a sector with $1 \times k$ parameters. Corresponding to each internal node in the binary tree, we have a sector with $k \times k$ parameters. Finally, corresponding to each leaf in the binary tree, we have a sector with $k \times l$ parameters, where l is the number of SPN leaf distributions per variable (we use $l = 2$ unless otherwise stated).

As shown in Figure 1, for n variables (where n is a power of 2) and r replicas, this corresponds to $2n - 1$ sectors totaling $O(rnk^2)$ parameters. In particular, for each replica there are $(n - 2)k^2 + (2n + 1)k$ parameters. For example, in a problem domain with 1000 variable, using layer size $k = 10$ and replica $r = 10$ leads to a RAT-SPN of around 1M parameters.

3 HyperSPNs

We propose HyperSPNs, an approach of regularizing probabilistic circuits with a large number of parameters by limiting the degrees of freedom of the model using an external neural network. We describe our method concretely on RAT-SPNs, which we will refer to as simply SPNs.

To regularize the SPN, we introduce a smaller-scale external neural network to generate the parameters of the SPN (see Figure 2). We associate each of the sectors of the SPN with a sector embedding. For each sector, we map the embedding to the parameters of the sector via the neural network. Then, we use these mapped parameters directly as the parameters of the SPN, re-normalize such that parameters leading into a sum node add up to 1, and evaluate the SPN as usual. We visualize the computation graph of this process in Figure 3. We materialize the SPN in the forward pass, and propagate the gradients through the SPN in the backward pass. We optimize the neural network and embeddings jointly using gradient descent.




Figure 3: We depict the computation graph when training a HyperSPN. The red modules have trainable parameters, whereas the gray module does not. Red arrows show the flow of gradients during backpropagation.




Figure 4: We visualize the soft weight-sharing effect in HyperSPNs, with training data in blue and testing data in orange. HyperSPNs use a large underlying SPN structure, allowing for a large number of clusters (3 in each quadrant) to fit the data well. Since the SPN parameters correspond to directly mixture weights, the weight-sharing strategy in HyperSPNs prevent overfitting by limiting the degrees of freedom across clusters. We depict this as fixing the shape of the cluster-triples across quadrants.

HyperSPNs can aggressively reduce the degrees of freedom of the original SPN. Using embeddings of size h , the number of learnable parameters is reduced to $O((rn + k^2)h)$. For the same above example with $n = 1000, k = 10, r = 10$, choosing embeddings of dimension $h = 5$ reduces the number of learnable parameters 20 fold, from around 1M down to around 50k parameters.

The lower degrees of freedom mitigates overfitting of the model. At the same time, the materialized SPN has a large number of mixture weights, and thus can express distribution families with a greater number of clusters compared to a small SPN with the same degrees of freedom. We visualize this with a motivating example in Figure 4, where we represent a dataset using a grid that is split recursively for each variable (we only draw the splits for the first 2 variables; the rest is implied). Intuitively, the number of parameters of the SPN is tied closely to the number of clusters in the distribution. As a result, training a large SPN can produce many clusters that overfit to the training points in blue (Figure 4a), and training a small SPN can lead to coarse factorizations and inflexible clusters (Figure 4b). Both cases can lead to poor generalization on the testing points in orange.

On the other hand, the HyperSPN uses a soft weight-sharing strategy by storing its parameters as low-dimensional embeddings and decoding them through a small-scale shared neural network. This allows the model to learn a high number of clusters without suffering from too high of a degree of freedom. In Figure 4c, we visualize the soft weight-sharing as fixing the shape of the cluster-triples within each quadrant, while still giving the model control over the center position of each cluster-triple. In this way, HyperSPNs can improve generalization compared to a large SPN with the same number of mixture weights and a small SPN with the same degrees of freedom. We verify this phenomenon on a hand-crafted dataset in the Appendix, and further observe similar improvements on real world datasets in our experiments.

3.1 Memory-Efficient Evaluation

Another merit of HyperSPNs is the ability for low-memory storage of the model, since we only need to store the embeddings and the neural network, which together have much fewer parameters than the SPN. Importantly, our abstraction of the SPN structure into sectors allows us to forgo storage of the nodes and edges in the SPN network. Moreover, not only can we be memory-efficient during storage, but we can also be memory-efficient during evaluation. To do so, we materialize sectors of the SPN on the fly, and erase them from memory when they become unnecessary.

The sectors of the SPN are structured as a binary tree, so to evaluate each sector we need all the input values leading into the sector. These input values correspond to the output values (each a vector of size k) of its children sectors. Evaluating the binary tree bottom-up depth-wise is possible, but is also costly since the leaf layer has n sectors leading to $O(nk)$ cost for storing the output values of children sectors. Rather, we can proceed greedily, evaluating leaf sectors left-to-right and propagating values up a sectors as soon as all of its children sectors have been evaluated. A simple recursive argument shows that this procedure requires storing only $O(k \log n)$ intermediate values at a time, and is provably optimal.

Proposition 1. *Evaluating a complete binary-tree computation graph requires $\Theta(\log n)$ memory, where n is the number of leaves in the complete binary-tree, and the memory cost of storing each intermediate value and performing each computation is assumed to be a constant.*

Proof. Perform the sequence of computations based on a post-order traversal of the binary tree, erasing children values from memory when their parent value has been computed. This process requires storing at most $\log n + 1$ intermediate values at any time. To show optimality, let $C(i)$ denote the optimal memory cost needed to compute the value for a node i . For non-leaf nodes i , Let j_1 be the first child of i that we begin to process, and j_2 be the second. Since we processed j_1 first (which requires at least 1 unit of memory throughout), we have $C(i) \geq 1 + C(j_2)$. As the base case, for leaf nodes i , we have $C(i) = 1$. Thus, the optimal memory cost for the whole computation is at least $C(\text{root}) = \log n + 1$. \square

Corollary 1. *Evaluating a HyperSPN with an underlying RAT-SPN structure requires $O(k \log n + h(k^2 + rn))$ memory.*

The cost in Corollary 1 can be broken down into two parts, where $O(k \log n)$ memory is required for storing intermediate values in the binary-tree computation graph, and $O(h(rn + k^2))$ memory is required for storing the HyperSPN parameters and materializing the SPN one sector at a time.

The takeaway is that both storing and evaluating a HyperSPN require less memory than storing and evaluating the underlying SPN, which require $O(rnk^2)$ memory. This memory efficiency arises from our key insight of using sector abstractions to enable 1) encoding the SPN parameters as embeddings, 2) forgoing storage of the nodes and edges in the SPN network, and 3) mapping the SPN to a binary-tree computation graph that allows for a low-memory evaluation order.

4 Experiments

We experiment with HyperSPN on three sets of density estimation tasks: the Twenty Datasets benchmark, the Amazon Baby Registries benchmark, and the Street View House Numbers (SVHN) dataset. All runs were done using a single GPU. We primarily compare with weight decay as the competing regularization method. We also include comparisons with smaller SPNs that have the same degrees of freedom as our HyperSPN, and with results reported from other works in literature. We omit results on dropout because in our experimentation (and as suggested by [25]), dropout on SPNs worked poorly for these generative density estimation tasks. In all three sets of tasks, we find that HyperSPNs give better generalization performance than weight decay, when comparing using the same underlying SPN structure and optimizer. Furthermore, HyperSPNs outperforms results reported in literature on EiNETs trained using the same optimization method (Adam), and is competitive against results reported on stochastic EM (sEM).

4.1 Twenty Datasets

The Twenty Datasets benchmark [12] is a standard benchmark used in comparing PC models. Our experiment aims to compare the use of HyperSPN against the baseline of regularization via weight decay, using the same underlying SPN structure for both. We use the RAT-SPN structure described in Section 2.1, choosing layer size parameter $k = 5$ and replicas $r = 50$, randomizing the variable orders for each replica. The number of learnable parameters in this SPN structure is shown in the # **Params** column under Weight Decay in Table 1. For Weight Decay, we vary the weight decay value between $1e-3$, $1e-4$, and $1e-5$.

The HyperSPN uses the exact same underlying SPN structure, along with an external neural network (a 2 layer MLP of width 20) and embeddings of dimension ranging between $h = 5, 10, 20$. The number of learnable parameters in the HyperSPN is around 5 times smaller, as shown in the # **Params** column under HyperSPN. We describe more details of our training setup in the Appendix.

In Table 1, we show in bold the better of the results between Weight Decay and HyperSPN. We see that HyperSPN outperforms Weight Decay on all but two of the datasets. On top of that, storing and evaluating the HyperSPN can be much more memory-efficient, as shown by the reduced number of trainable parameters and as described by the techniques in Section 3.1.

Table 1: Twenty Datasets. We plot the test log-likelihood and the number of trainable parameters. Bold values indicate the best results between the two regularization strategies of Weight Decay and HyperSPN, using the same SPN structure and optimizer. Underlined values indicate the best results when compared also to reported values in literature using EiNETs/Adam* and RAT-SPN/sEM[†].

Name	Variables	Adam		Adam		Adam	sEM
		Weight Decay	# Params	HyperSPN	# Params	*	†
		Log-LH	# Params	Log-LH	# Params	Log-LH	Log-LH
NLTCS	16	-6.02	40050	-6.01	9115	-6.04	<u>-6.01</u>
MSNBC	17	-6.04	42550	-6.05	9615	<u>-6.03</u>	-6.04
KDDCup2k	64	-2.14	160050	-2.13	33115	-2.15	<u>-2.13</u>
Plants	69	-13.41	172550	-13.27	35615	-13.74	<u>-13.44</u>
Audio	100	-40.14	250050	-39.74	51115	-40.22	-39.96
Jester	100	-52.99	250050	-52.74	51115	-53.10	-52.97
Netflix	100	-57.18	250050	-56.62	51115	-57.10	-56.85
Accidents	111	-35.55	277550	-35.40	56615	-37.45	-35.49
Retail	135	-10.90	337550	-10.89	68615	-10.97	-10.91
Pumsb-star	163	-31.08	407550	-31.07	82615	-39.23	-32.53
DNA	180	-98.42	450050	-98.79	91115	-97.68	<u>-97.23</u>
Kosarek	190	-10.88	475050	-10.90	96115	-10.92	-10.89
MSWeb	294	-10.14	735050	-9.90	148115	-10.26	-10.12
Book	500	-34.84	1250050	-34.86	251115	-35.15	<u>-34.68</u>
EachMovie	500	-52.85	1250050	-51.62	251115	-55.49	-53.63
WebKB	839	-159.68	2097550	-157.69	420615	-160.51	<u>-157.53</u>
Reuters-52	889	-90.15	2222550	-86.12	445615	-92.76	-87.37
20Newsgrp	910	-154.36	2275050	-152.49	456115	-154.41	<u>-152.06</u>
BBC	1058	-262.77	2645050	-254.44	530115	-267.86	<u>-252.14</u>
Ad	1556	-54.82	3890050	-28.58	779115	-63.82	-48.47

On the two right-most columns, we directly copy over results reported in literature. We compare with training EiNETs on Adam (the same optimizer as our setup) [24], and training RAT-SPNs on sEM (a similar network structure as our setup) [25]. We see that HyperSPN also compares favorably to these reported results (the best values underlined).

To illustrate the regularization effects of HyperSPN, we plot training curves on the training and validation data for the Plants and for Pumsb-star datasets. We use the same HyperSPN model described above. For clarity, let M denote the size of its underlying SPN and m denote its true number of learnable parameters (those in its external neural network and embeddings). We then construct two SPNs – one with size M (SPN-Large) and one with size m (SPN-Small) – and overlay their training curves in Figure 5. For both plots in Figure 5a and 5b, SPN-Large suffers from overfitting on the training data (presumably due to the large degrees of freedom), and SPN-Small fits the data poorly (presumably limited by the size of the SPN). Instead, HyperSPN strikes the balance between the expressiveness of a large underlying SPN, and regularization properties from the compact neural network with low degrees of freedom. This translates to better generalization on the validation data shown in the plots of Figure 5, and on the testing data shown in Table 1.

4.2 Amazon Baby Registries

We conduct similar experiments on the Amazon Baby Registries dataset [8]. We repeat the same experimental setup of maximizing the log-likelihood of data, and use the same model structure for HyperSPN and Weight Decay as their respective structures described in the previous section.

In Table 2, we see that HyperSPN clearly outperforms Weight Decay as a regularization technique, giving better test log-likelihoods (in bold) while using around 5 times fewer trainable parameters. On the right-most column we include results as reported in literature [37] on training EiNETs using sEM. Although the comparison with this right-most column is hard to interpret due to the differences in model architecture, optimization procedure, and regularization, we still note that HyperSPNs compare favorably on the majority of the datasets (best values underlined).




Figure 5: We plot the training curves on Plants and Pumsb-star from the Twenty Datasets benchmark. The HyperSPN has the same underlying SPN as SPN-Large, but limits its degrees of freedom to the same as those of SPN-Small. In both plots, HyperSPN fits the data better than SPN-Small, and is better regularized than SPN-Large, leading to the best results on the validation set (solid orange line).

The experimental results on the Amazon Baby Registries benchmark confirm our findings in the previous subsection. On both tasks, HyperSPNs generalizes much better than Weight Decay when compared on the same underlying SPN structure, while enabling more memory-efficient storage and evaluation of the model.

Table 2: Amazon Baby Registries. We plot the test log-likelihood and the number of trainable parameters. Bold values indicate the best results between the two regularization strategies of Weight Decay and HyperSPN, using the same SPN structure and optimizer. Underlined values indicate the best results when compared also to reported values in literature using EiNETs/sEM*.

Name	Variables	Adam Weight Decay		Adam HyperSPN		sEM *
		Log-LH	# Params	Log-LH	# Params	Log-LH
Apparel	100	-9.32	250050	-9.28	50655	<u>-9.24</u>
Bath	100	-8.54	250050	-8.52	50655	<u>-8.49</u>
Bedding	100	-8.59	250050	-8.57	50655	<u>-8.55</u>
Carseats	34	-4.72	85050	-4.65	17655	<u>-4.72</u>
Diaper	100	-9.99	250050	-9.91	50655	<u>-9.86</u>
Feeding	100	-11.35	250050	-11.30	50655	<u>-11.27</u>
Furniture	32	-4.54	80050	-4.32	16655	<u>-4.38</u>
Gear	100	-9.21	250050	-9.21	50655	<u>-9.18</u>
Gifts	16	-3.43	40050	-3.40	8655	<u>-3.42</u>
Health	62	-7.49	155050	-7.41	31655	<u>-7.47</u>
Media	58	-7.86	145050	-7.83	29655	<u>-7.82</u>
Moms	16	-3.48	40050	-3.47	8655	<u>-3.48</u>
Safety	36	-4.48	90050	-4.34	18655	<u>-4.39</u>
Strollers	40	-5.25	100050	-5.00	20655	<u>-5.07</u>
Toys	62	-7.83	155050	-7.79	31655	<u>-7.84</u>

Sample Quality In addition to generalization performance on log-likelihood, we can examine the sample quality of HyperSPNs in comparison to standard SPNs. To estimate sample quality, we take the approach of applying kernel density estimation / Parzen windows [3] using the test dataset. Out of the 35 datasets from the above Twenty Datasets and Amazon Baby Registries, we observe that HyperSPNs report better sample quality for 29 of them, compared to SPNs trained with weight decay. We report the full results of this experiment in the Appendix.

4.3 Street View House Numbers (SVHN)

Lastly, we experiment on the SVHN dataset [20]. We use the infrastructure provided by [24], which groups the train/valid/test data into 100 clusters, then learns a separate PC for each cluster. The final density estimator is a combination of all the individual PCs under one mixture distribution. Here, we




Figure 6: We plot the improvement in test log-likelihood on the Street View House Numbers (SVHN) dataset when using HyperSPN instead of weight decay. Following the setup from [24], the dataset is split into 100 clusters and a separate HyperSPN is learned for each cluster. We sort the cluster index by improvement for readability. The dashed blue line corresponds to the average (weighted by cluster size) log-likelihood improvement of 9.

study the performance of each individual PC on each cluster separately, comparing between weight decay and HyperSPN when trained using Adam.

The underlying PC structure is an EiNET with leaves that hold continuous distributions [24]. To introduce soft weight-sharing of the HyperSPN on the EiNET, we partition the parameters into sectors of size 40, which is the number of sum nodes used in each layer in the setup of [24]. We use an embedding size of $h = 8$, which reduces the number of learnable edge-parameters in the EiNET by about 5 fold. In Figure 6, we plot the difference between HyperSPN and weight decay in log-likelihood on the test set, sorted in increasing order for readability. The higher the value, the more improvement we observe when using HyperSPN as the regularization method. In around 70% of the clusters, we see an improvement in generalization, and on average (weighted by cluster size) the HyperSPN achieves an improvement in log-likelihood of 9 ± 4 .

Our experiments suggest that HyperSPNs give a superior regularization strategy for PCs compared to the primary alternative of weight decay. HyperSPNs also lead to better performance than other reported approaches in literature that use the same optimizer (Adam), as well as other reported methods that use sEM. Our experimental results are general, showing that HyperSPNs are applicable across three standard benchmarks, across PCs with discrete and continuous leaf distributions, and across different state-of-the-art structures such as RAT-SPNs and EiNETs.

5 Related Work

Many existing works focus on structure learning for probabilistic circuits, covering a wide range of PC types [7, 28, 27, 15]. These approaches often regularize by limiting the complexity of the structure [7, 35]. However, structure learning in general can be tedious and difficult to tune [25].

Recent works have embraced a more deep-learning-like framework by prescribing the structure of the PC and learning only the model parameters [25, 24]. This framework leads to the training of large PCs with millions of parameters. However, little exploration has been done around regularization in this new parameter-centric framework. For generative settings, dropout was suggested as unsuitable [25], leaving only weight decay as the primary parameter regularization method. Concurrently, other works [16] look into data softening and entropy regularization techniques. In our work, we propose the new regularization strategy of HyperSPNs which show promising improvements in generalization.

Finally, the idea of generating parameters with an external neural network comes from HyperNetworks [11], which were introduced as an effective architecture for recurrent models. Soft weight-sharing strategies in general [21] have also shown good empirical results [2] and proven useful for neural network compression [32]. There have also been recent work on using similar techniques to parameterize conditional SPNs [29, 36].

6 Conclusion

We propose HyperSPNs, a new paradigm of training large probabilistic circuits using a small-scale external neural network to generate the parameters of the probabilistic circuit. HyperSPNs provide regularization by aggressively reducing the degrees of freedom through a soft weight-sharing strategy. We show empirical improvements in the generalization performance of PCs compared to competing regularization methods, across both discrete and continuous density estimation tasks, and across different state-of-the-art PC structures. As a bonus, HyperSPNs enable memory-efficient storage and evaluation, combining the expressiveness of large PCs with the compactness of small models.

Limitations and Future Work Our analysis of the memory efficiency of HyperSPNs is limited their storage and evaluation, since training HyperSPNs involves manipulating of the gradients through the underlying SPN. Future work can look into combining HyperSPNs with techniques from memory-efficient back propagation [4, 10] to train very large SPNs that do not fit in memory.

7 Acknowledgments

This research was supported by NSF (#1651565, #1522054, #1733686), ONR (N00014-19-1-2145), AFOSR (FA9550-19-1-0024), ARO (W911NF-21-1-0125) and Sloan Fellowship.

References

- [1] Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- [2] Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc V. Le. Can weight sharing outperform random architecture search? an investigation with tunas. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [3] Yoshua Bengio, Grégoire Mesnil, Yann N. Dauphin, and Salah Rifai. Better mixing via deep representations. In *Proceedings of the 30th International Conference on Machine Learning, ICML*, 2013.
- [4] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. *CoRR*, abs/1604.06174, 2016.
- [5] Adnan Darwiche. Decomposable negation normal form. *J. ACM*, 48(4):608–647, 2001.
- [6] Adnan Darwiche. A differential approach to inference in bayesian networks. *J. ACM*, 50(3):280–305, 2003.
- [7] Robert Gens and Pedro M. Domingos. Learning the structure of sum-product networks. In *Proceedings of the 30th International Conference on Machine Learning, ICML*, 2013.
- [8] Jennifer Gillenwater, Alex Kulesza, Emily B. Fox, and Benjamin Taskar. Expectation-maximization for learning determinantal point processes. In *Advances in Neural Information Processing Systems 27, NIPS*, 2014.
- [9] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27 NIPS*, 2014.
- [10] Audrunas Gruslys, Rémi Munos, Ivo Danihelka, Marc Lanctot, and Alex Graves. Memory-efficient backpropagation through time. In *Advances in Neural Information Processing Systems 29, NIPS*, 2016.
- [11] David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *5th International Conference on Learning Representations, ICLR*, 2017.

- [12] Jan Van Haaren and Jesse Davis. Markov network structure learning: A randomized feature generation approach. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, 2012.
- [13] Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron C. Courville. Neural autoregressive flows. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, 2018.
- [14] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- [15] Yitao Liang and Guy Van den Broeck. Learning logistic circuits. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI*, 2019.
- [16] Anji Liu and Guy Van den Broeck. Tractable regularization of probabilistic circuits. In *UAI Workshop on Tractable Probabilistic Modeling (TPM)*, 2021.
- [17] Daniel Lowd and Pedro M. Domingos. Learning arithmetic circuits. In *Proceedings of the 24th Conference in Uncertainty in Artificial Intelligence, UAI*, 2008.
- [18] Nicola Di Mauro, Antonio Vergari, Teresa Maria Altomare Basile, and Floriana Esposito. Fast and accurate density estimation with extremely randomized cutset networks. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD*, 2017.
- [19] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. In *8th International Conference on Learning Representations, ICLR*, 2020.
- [20] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [21] Steven J. Nowlan and Geoffrey E. Hinton. Simplifying neural networks by soft weight-sharing. *Neural Comput.*, 4(4):473–493, 1992.
- [22] George Papamakarios, Eric T. Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *CoRR*, abs/1912.02762, 2019.
- [23] Robert Peharz. *Foundations of sum-product networks for probabilistic modeling*. PhD thesis, PhD thesis, Medical University of Graz, 2015.
- [24] Robert Peharz, Steven Lang, Antonio Vergari, Karl Stelzner, Alejandro Molina, Martin Trapp, Guy Van den Broeck, Kristian Kersting, and Zoubin Ghahramani. Einsum networks: Fast and scalable learning of tractable probabilistic circuits. In *Proceedings of the 37th International Conference on Machine Learning, ICML*, 2020.
- [25] Robert Peharz, Antonio Vergari, Karl Stelzner, Alejandro Molina, Martin Trapp, Xiaoting Shao, Kristian Kersting, and Zoubin Ghahramani. Random sum-product networks: A simple and effective approach to probabilistic deep learning. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI*, 2019.
- [26] Hoifung Poon and Pedro M. Domingos. Sum-product networks: A new deep architecture. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, 2011.
- [27] Tahrima Rahman, Prasanna Kothalkar, and Vibhav Gogate. Cutset networks: A simple, tractable, and scalable approach for improving the accuracy of chow-liu trees. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD*, 2014.
- [28] Amirmohammad Rooshenas and Daniel Lowd. Learning sum-product networks with direct and indirect variable interactions. In *Proceedings of the 31th International Conference on Machine Learning, ICML*, 2014.

- [29] Xiaoting Shao, Alejandro Molina, Antonio Vergari, Karl Stelzner, Robert Peharz, Thomas Liebig, and Kristian Kersting. Conditional sum-product networks: Imposing structure on deep probabilistic architectures. In *International Conference on Probabilistic Graphical Models, PGM*, pages 401–412. PMLR, 2020.
- [30] Andy Shih and Stefano Ermon. Probabilistic circuits for variational inference in discrete graphical models. In *Advances in Neural Information Processing Systems 33, NeurIPS*, 2020.
- [31] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 2014.
- [32] Karen Ullrich, Edward Meeds, and Max Welling. Soft weight-sharing for neural network compression. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [33] Benigno Uribe, Marc-Alexandre Côté, Karol Gregor, Iain Murray, and Hugo Larochelle. Neural autoregressive distribution estimation. *J. Mach. Learn. Res.*, 17:205:1–205:37, 2016.
- [34] Antonio Vergari, YooJung Choi, Robert Peharz, and Guy Van den Broeck. Probabilistic circuits: Representations, inference, learning and applications. In *Tutorial at the The 34th AAAI Conference on Artificial Intelligence*, 2020.
- [35] Antonio Vergari, Nicola Di Mauro, and Floriana Esposito. Simplifying, regularizing and strengthening sum-product network structure learning. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD*, 2015.
- [36] Matej Zečević, Devendra Singh Dhami, Athresh Karanam, Sriraam Natarajan, and Kristian Kersting. Interventional sum-product networks: Causal inference with tractable probabilistic models. In *Advances in Neural Information Processing Systems 34, NeurIPS*, 2021.
- [37] Honghua Zhang, Brendan Juba, and Guy Van den Broeck. Probabilistic generating circuits. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, 2021.
- [38] Han Zhao, Pascal Poupart, and Geoffrey J. Gordon. A unified approach for learning the parameters of sum-product networks. In *Advances in Neural Information Processing Systems 29, NIPS*, 2016.

The code for our experiments is available at <https://github.com/AndyShih12/HyperSPN>.

A Hand-Crafted Example

To examine the merits of HyperSPNs as discussed in Section 3, we construct a hand-crafted dataset to test the three types of models described in Figure 4: SPN-Large, SPN-Small, and HyperSPN. The hand-crafted dataset is procedurally generated with 256 binary variables and 10000 instances, broken into train/valid/test splits at 70/10/20%. The generation procedure is designed such that the correlation between variable i and j is dependent on the path length between leaves i and j of a complete binary tree over the 256 variables. The exact details can be found in our code.

SPN-Large has the same number of SPN edges as the HyperSPN, while SPN-Small has roughly the same number of trainable parameters as HyperSPN. Both SPN-Large and SPN-Small are regularized via weight-decay. As we can see in Table 3, HyperSPN gives the best generalization performance on the test split of our hand-crafted dataset, when compared to standard SPNs with either similar number of SPN edges (SPN-Large) and similar number of trainable parameters (SPN-Small).

Table 3: Testing HyperSPNs on a hand-crafted toy dataset with 256 variables.

	Log-Likelihood	# Params
SPN-Large	-166.90 ± 0.03	640050
SPN-Small	-167.00 ± 0.01	102450
HyperSPN	-166.32 ± 0.04	129115

B Experimental Details

Here, we provide more details on our experimental setup. In Table 4, we give the hyperparameters used for training our models on the Twenty Datasets and Amazon Baby Registries benchmarks. For both methods we do early stopping by training until the validation performance plateaus/declines (we train some up to 80k steps). Then we take the version of the model that performed best on the validation set, and use it for evaluation on the test set.

Recall that for the standard SPN, we take gradient descent steps on the mixture weights of the SPN. For the HyperSPN, we take gradient descent steps on the parameters of the external neural network. Empirically, we found that higher learning rates are more suitable for the SPN (e.g. $2e-2$), and lower learning rates are more suitable for the HyperSPN (e.g. $5e-3$).

Table 4: Training hyperparameters for Twenty Datasets and Amazon Baby Registries

	SPN (Weight Decay)	HyperSPN
Learning Rate	$2e-2$	$5e-3$
Weight Decay	{ $1e-3$, $1e-4$, $1e-5$ }	–
Embedding Dim	–	{5, 10, 20}
Batch Size	500	500

For the SVHN experiment, we build on the code provided at <https://github.com/cambridge-mlg/EinsumNetworks>. We train each method for 100 epochs using Adam, and also use early stopping based on the validation set. The weight decay hyperparameter used for the weights of the models are the same as that shown in Table 4, and we scale down the learning rate for both models to $1e-3$ and $3e-4$, respectively. We found the slower learning rate to be more suitable for both models on this benchmark, with HyperSPNs still giving the better performance.

For the Twenty Datasets and Amazon Baby Registries, the leaf nodes are binary indicator random variables, hence there are no trainable parameters. For SVHN, the leaf nodes are factorized Gaussians. The training of the leaf distributions was kept the same for SPNs and HyperSPNs (i.e. for HyperSPNs, we do not generate the parameters of the leaf distributions using the external neural network).

C Additional Experiments and Ablations

We provide additional experiments that examine different hyperparameter choices for the embedding size h and for the weight decay parameter. We also show more detailed results for sample quality as measured using Parzen windows. Finally, we report error bars for the experiments from Tables 1 & 2.

C.1 Weight Decay and Embedding Size h

We try out different hyperparameter values for the weight decay value used for regularizing standard SPNs, and for the embedding size of the HyperSPNs. We report these results in Table 5, where we see that weight decay of $1e-4$ for standard SPNs and embedding size of $h = 10$ for HyperSPNs gave slightly better performance than the alternative settings.

Table 5: For standard SPNs (left), we vary the weight decay value between $1e-3$, $1e-4$, and $1e-5$. For HyperSPNs (right), we vary the embedding size h between 5, 10, and 20.

Name	Variables	SPN			HyperSPN		
		$w = 1e-3$	$w = 1e-4$	$w = 1e-5$	$h = 5$	$h = 10$	$h = 20$
NLTCS	16	-6.02	-6.02	-6.02	-6.03	-6.01	-6.02
MSNBC	17	-6.06	-6.04	-6.04	-6.05	-6.06	-6.07
KDDCup2k	64	-2.14	-2.14	-2.14	-2.14	-2.13	-2.13
Plants	69	-13.44	-13.41	-13.45	-13.31	-13.27	-13.27
Audio	100	-40.16	-40.14	-40.17	-39.86	-39.74	-39.75
Jester	100	-53.01	-52.99	-53.05	-52.92	-52.74	-52.83
Netflix	100	-57.18	-57.20	-57.21	-56.73	-56.66	-56.62
Accidents	111	-35.94	-35.55	-35.65	-36.02	-35.52	-35.40
Retail	135	-10.92	-10.92	-10.90	-10.89	-10.92	-10.92
Pumsb-star	163	-31.89	-31.08	-31.48	-31.46	-31.39	-31.07
DNA	180	-98.59	-98.42	-98.45	-98.79	-99.05	-98.88
Kosarek	190	-10.91	-10.89	-10.88	-10.90	-10.92	-10.92
MSWeb	294	-10.28	-10.14	-10.14	-9.93	-9.92	-9.90
Book	500	-34.90	-34.84	-35.02	-34.95	-35.02	-34.86
EachMovie	500	-52.85	-53.21	-54.01	-51.62	-52.10	-52.00
WebKB	839	-159.68	-160.10	-160.06	-157.69	-158.35	-158.24
Reuters-52	889	-92.82	-90.15	-90.99	-86.93	-86.12	-86.76
20Newsgrp	910	-154.36	-154.70	-155.01	-152.57	-152.82	-152.49
BBC	1058	-267.47	-262.77	-268.99	-256.07	-254.44	-255.81
Ad	1556	-56.34	-54.90	-54.82	-31.50	-28.58	-29.84
Apparel	100	-9.32	-9.33	-9.33	-9.30	-9.28	-9.28
Bath	100	-8.54	-8.59	-8.60	-8.52	-8.52	-8.52
Bedding	100	-8.64	-8.59	-8.63	-8.59	-8.58	-8.57
Carseats	34	-4.72	-4.82	-4.76	-4.77	-4.65	-4.66
Diaper	100	-9.99	-10.03	-10.00	-9.91	-9.94	-9.93
Feeding	100	-11.37	-11.35	-11.41	-11.31	-11.30	-11.31
Furniture	32	-4.56	-4.54	-4.60	-4.48	-4.46	-4.32
Gear	100	-9.24	-9.21	-9.21	-9.41	-9.21	-9.21
Gifts	16	-3.48	-3.43	-3.48	-3.44	-3.40	-3.42
Health	62	-7.49	-7.50	-7.49	-7.75	-7.41	-7.43
Media	58	-7.86	-7.91	-7.90	-7.87	-7.83	-7.87
Moms	16	-3.49	-3.48	-3.49	-3.48	-3.50	-3.47
Safety	36	-4.48	-4.48	-4.57	-4.44	-4.37	-4.34
Strollers	40	-5.27	-5.29	-5.25	-5.23	-5.02	-5.00
Toys	62	-7.83	-7.83	-7.83	-7.83	-7.81	-7.79
Average	-	-37.18	-36.91	-37.17	-35.79	-35.63	-35.68

C.2 Sample Quality

Next, we report estimates of sample quality of the trained SPN / HyperSPN using Parzen windows on the test data. We treat the binary data as real vectors, and use a Gaussian kernel with a fixed variance on each of the test data points. Then, we sample 500 data points from the trained SPN / HyperSPN, and compute the log-likelihood of the samples.

Table 6: We examine an estimate of sample quality using Parzen windows on the test data, taking 500 samples per dataset. We see that HyperSPNs give better sample quality under this metric (higher is better).

Name	Variables	SPN	HyperSPN
NLTCS	16	-2.346	-2.348
MSNBC	17	-2.247	-2.247
KDDCup2k	64	-2.066	-2.066
Plants	69	-2.927	-2.926
Audio	100	-3.606	-3.594
Jester	100	-4.329	-4.319
Netflix	100	-4.470	-4.467
Accidents	111	-3.557	-3.554
Retail	135	-2.290	-2.290
Pumsb-star	163	-4.247	-4.238
DNA	180	-5.693	-5.676
Kosarek	190	-2.310	-2.318
MSWeb	294	-2.299	-2.299
Book	500	-2.758	-2.740
EachMovie	500	-4.053	-3.972
WebKB	839	-6.307	-6.111
Reuters-52	889	-4.632	-4.558
20Newsgrp	910	-5.683	-5.705
BBC	1058	-9.401	-8.942
Ad	1556	-3.187	-3.147
Apparel	100	-2.257	-2.251
Bath	100	-2.219	-2.209
Bedding	100	-2.233	-2.224
Carseats	34	-2.151	-2.152
Diaper	100	-2.273	-2.262
Feeding	100	-2.308	-2.301
Furniture	32	-2.146	-2.144
Gear	100	-2.238	-2.228
Gifts	16	-2.134	-2.134
Health	62	-2.205	-2.205
Media	58	-2.229	-2.227
Moms	16	-2.135	-2.135
Safety	36	-2.141	-2.142
Strollers	40	-2.163	-2.160
Toys	62	-2.219	-2.216
Average	-	-3.185	-3.157

C.3 Error Bars

Lastly, we report the main results from Table 1 and Table 2 with error bars, computed over 3 separate runs with different random seeds.

Table 7: We report the values from in Table 1 and Table 2 with error bars, computed over 3 separate runs. We denote statistical significance with a $*$.

Name	Variables	SPN	HyperSPN
NLTCS	16	-6.02 \pm 0.00	-6.01* \pm 0.00
MSNBC	17	-6.04 \pm 0.00	-6.05 \pm 0.01
KDDCup2k	64	-2.14 \pm 0.00	-2.13* \pm 0.00
Plants	69	-13.41 \pm 0.03	-13.27* \pm 0.06
Audio	100	-40.14 \pm 0.01	-39.74* \pm 0.03
Jester	100	-52.99 \pm 0.02	-52.74* \pm 0.02
Netflix	100	-57.18 \pm 0.01	-56.62* \pm 0.01
Accidents	111	-35.55 \pm 0.07	-35.40 \pm 0.11
Retail	135	-10.90 \pm 0.01	-10.89 \pm 0.03
PumSB-star	163	-31.08 \pm 0.12	-31.07 \pm 0.02
DNA	180	-98.42* \pm 0.05	-98.79 \pm 0.09
Kosarek	190	-10.88 \pm 0.01	-10.90 \pm 0.03
MSWeb	294	-10.14 \pm 0.02	-9.90* \pm 0.02
Book	500	-34.84 \pm 0.03	-34.86 \pm 0.02
EachMovie	500	-52.85 \pm 0.11	-51.62* \pm 0.22
WebKB	839	-159.68 \pm 0.21	-157.69* \pm 0.63
Reuters-52	889	-90.15 \pm 0.49	-86.12* \pm 0.13
20Newsgrp	910	-154.36 \pm 0.04	-152.49* \pm 0.67
BBC	1058	-262.77 \pm 0.10	-254.44* \pm 0.29
Ad	1556	-54.82 \pm 0.97	-28.58* \pm 0.20
Apparel	100	-9.32 \pm 0.01	-9.28 \pm 0.02
Bath	100	-8.54 \pm 0.01	-8.52 \pm 0.02
Bedding	100	-8.59 \pm 0.01	-8.57* \pm 0.00
Carseats	34	-4.72 \pm 0.01	-4.65* \pm 0.01
Diaper	100	-9.99 \pm 0.07	-9.91* \pm 0.02
Feeding	100	-11.35 \pm 0.01	-11.30* \pm 0.00
Furniture	32	-4.54 \pm 0.02	-4.32 \pm 0.13
Gear	100	-9.21 \pm 0.01	-9.21 \pm 0.01
Gifts	16	-3.43 \pm 0.02	-3.40* \pm 0.00
Health	62	-7.49 \pm 0.01	-7.41* \pm 0.02
Media	58	-7.86 \pm 0.01	-7.83 \pm 0.03
Moms	16	-3.48 \pm 0.00	-3.47 \pm 0.03
Safety	36	-4.48 \pm 0.01	-4.34* \pm 0.01
Strollers	40	-5.25 \pm 0.01	-5.00* \pm 0.03
Toys	62	-7.83 \pm 0.02	-7.79 \pm 0.03
Average	-	-36.87 \pm 0.07	-35.55* \pm 0.08