

Boosted Generative Models

Aditya Grover, Stefano Ermon

Department of Computer Science
Stanford University
{adityag, ermon}@cs.stanford.edu

Abstract

We propose a novel approach for using unsupervised boosting to create an ensemble of generative models, where models are trained in sequence to correct earlier mistakes. Our meta-algorithmic framework can leverage any existing base learner that permits likelihood evaluation, including recent deep expressive models. Further, our approach allows the ensemble to include discriminative models trained to distinguish real data from model-generated data. We show theoretical conditions under which incorporating a new model in the ensemble will improve the fit and empirically demonstrate the effectiveness of our black-box boosting algorithms on density estimation, classification, and sample generation on benchmark datasets for a wide range of generative models.

1 Introduction

A variety of deep generative models have shown promising results on tasks spanning computer vision, speech recognition, natural language processing, and imitation learning (Poon and Domingos 2011; Oord, Kalchbrenner, and Kavukcuoglu 2016; Kingma and Welling 2014; Goodfellow et al. 2014; Zhao, Song, and Ermon 2017; Li, Song, and Ermon 2017). These parametric models differ from each other in their ability to perform various forms of tractable inference, learning algorithms, and objectives. Despite significant progress, existing generative models cannot fit complex distributions with a sufficiently high degree of accuracy, limiting their applicability and leaving room for improvement.

In this paper, we propose a technique for ensembling (imperfect) generative models to improve their overall performance. Our meta-algorithm is inspired by boosting, a technique used in supervised learning to combine weak classifiers (e.g., decision stumps or trees), which individually might not perform well on a given classification task, into a more powerful ensemble. The boosting algorithm will attempt to learn a classifier to correct for the mistakes made by reweighting the original dataset, and repeat this procedure recursively. Under some conditions on the weak classifiers' effectiveness, this procedure can drive the (training) error to zero (Freund, Schapire, and Abe 1999). Boosting can also be thought as a feature learning algorithm, where at each round a new feature is learned by training a classifier on a

reweighted version of the original dataset. In practice, algorithms based on boosting perform extremely well in machine learning competitions (Caruana and Niculescu-Mizil 2006).

We show that a similar procedure can be applied to generative models. Given an initial generative model that provides an imperfect fit to the data distribution, we construct a second model to correct for the error, and repeat recursively. The second model is also a generative one, which is trained on a reweighted version of the original training set. Our meta-algorithm is general and can construct ensembles of any existing generative model that permits (approximate) likelihood evaluation such as fully-observed belief networks, sum-product networks, and variational autoencoders. Interestingly, our method can also leverage powerful discriminative models. Specifically, we train a binary classifier to distinguish true data samples from "fake" ones generated by the current model and provide a principled way to include this discriminator in the ensemble.

A prior attempt at boosting density estimation proposed a *sum-of-experts* formulation (Rosset and Segal 2002). This approach is similar to supervised boosting where at every round of boosting we derive a reweighted additive estimate of the boosted model density. In contrast, our proposed framework uses multiplicative boosting which multiplies the ensemble model densities and can be interpreted as a *product-of-experts* formulation. We provide a holistic theoretical and algorithmic framework for multiplicative boosting contrasting with competing additive approaches. Unlike prior use cases of product-of-experts formulations, our approach is *black-box*, and we empirically test the proposed algorithms on several generative models from simple ones such as mixture models to expressive parametric models such as sum-product networks and variational autoencoders.

Overall, this paper makes the following contributions:

1. We provide theoretical conditions for additive and multiplicative boosting under which incorporating a new model is guaranteed to improve the ensemble fit.
2. We design and analyze a flexible meta-algorithmic boosting framework for including both generative and discriminative models in the ensemble.
3. We demonstrate the empirical effectiveness of our algorithms for density estimation, generative classification, and sample generation on several benchmark datasets.

2 Unsupervised boosting

Supervised boosting provides an algorithmic formalization of the hypothesis that a sequence of weak learners can create a single strong learner (Schapire and Freund 2012). Here, we propose a framework that extends boosting to unsupervised settings for learning generative models. For ease of presentation, all distributions are with respect to any arbitrary $\mathbf{x} \in \mathbb{R}^d$, unless otherwise specified. We use upper-case symbols to denote probability distributions and assume they all admit absolutely continuous densities (denoted by the corresponding lower-case notation) on a reference measure $d\mathbf{x}$. Our analysis naturally extends to discrete distributions, which we skip for brevity. Please refer to the companion technical report for the proofs of all results in the following two sections (Grover and Ermon 2017).

Formally, we consider the following maximum likelihood estimation (MLE) setting. Given some data points $X = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^m$ sampled i.i.d. from an unknown distribution P , we provide a model class \mathcal{Q} parameterizing the distributions that can be represented by the generative model and minimize the Kullback-Liebler (KL) divergence with respect to the true distribution:

$$\min_{Q \in \mathcal{Q}} D_{KL}(P \| Q). \quad (1)$$

In practice, we only observe samples from P and hence, maximize the log-likelihood of the observed data X . Selecting the model class for maximum likelihood learning is non-trivial; MLE w.r.t. a small class can be far from P , whereas a large class poses the risk of overfitting in the absence of sufficient data, or even underfitting due to difficulty in optimizing non-convex objectives that frequently arise due to the use of latent variable models, neural networks, etc.

The boosting intuition is to greedily increase model capacity by learning a sequence of weak intermediate models $\{h_t \in \mathcal{H}_t\}_{t=0}^T$ that can correct for mistakes made by previous models in the ensemble. Here, \mathcal{H}_t is a predefined model class (such as \mathcal{Q}) for h_t . We defer the algorithms pertaining to the learning of such intermediate models to the next section, and first discuss two mechanisms for deriving the final estimate q_T from the individual density estimates at each round, $\{h_t\}_{t=0}^T$.

2.1 Additive boosting

In additive boosting, the final density estimate is an arithmetic average of the intermediate models:

$$q_T = \sum_{t=0}^T \alpha_t \cdot h_t$$

where $0 \leq \alpha_t \leq 1$ denote the weights assigned to the intermediate models. The weights are re-normalized at every round to sum to 1 which gives us a valid probability density estimate. Starting with a base model h_0 , we can express the density estimate after a round of boosting recursively as:

$$q_t = (1 - \hat{\alpha}_t) \cdot q_{t-1} + \hat{\alpha}_t \cdot h_t$$

where $\hat{\alpha}_t$ denotes the normalized weight for h_t at round t . We now derive conditions on the intermediate models that guarantee “progress” in every round of boosting.

Theorem 1. Let $\delta_{KL}^t(h_t, \hat{\alpha}_t) = D_{KL}(P \| Q_{t-1}) - D_{KL}(P \| Q_t)$ denote the reduction in KL-divergence at the t^{th} round of additive boosting. The following conditions hold:

1. *Sufficient:* If $\mathbb{E}_P \left[\log \frac{h_t}{q_{t-1}} \right] \geq 0$, then $\delta_{KL}^t(h_t, \hat{\alpha}_t) \geq 0$ for all $\hat{\alpha}_t \in [0, 1]$.
2. *Necessary:* If $\exists \hat{\alpha}_t \in (0, 1]$ such that $\delta_{KL}^t(h_t, \hat{\alpha}_t) \geq 0$, then $\mathbb{E}_P \left[\frac{h_t}{q_{t-1}} \right] \geq 1$.

The sufficient and necessary conditions require that the expected log-likelihood and likelihood respectively of the current intermediate model, h_t are better-or-equal than those of the combined previous model, q_{t-1} under the true distribution when compared using density ratios. Next, we consider an alternative formulation of multiplicative boosting for improving the model fit to an arbitrary data distribution.

2.2 Multiplicative boosting

In multiplicative boosting, we factorize the final density estimate as a geometric average of $T + 1$ intermediate models $\{h_t\}_{t=0}^T$, each assigned an exponentiated weight α_t :

$$q_T = \frac{\prod_{t=0}^T h_t^{\alpha_t}}{Z_T}$$

where the partition function $Z_T = \int \prod_{t=0}^T h_t^{\alpha_t} d\mathbf{x}$. Recursively, we can specify the density estimate as:

$$\tilde{q}_t = h_t^{\alpha_t} \cdot \tilde{q}_{t-1} \quad (2)$$

where \tilde{q}_t is the unnormalized estimate at round t . The base model h_0 is learned using MLE. The conditions on the intermediate models for reducing KL-divergence at every round are stated below.

Theorem 2. Let $\delta_{KL}^t(h_t, \alpha_t) = D_{KL}(P \| Q_{t-1}) - D_{KL}(P \| Q_t)$ denote the reduction in KL-divergence at the t^{th} round of multiplicative boosting. The following conditions hold:

1. *Sufficient:* If $\mathbb{E}_P[\log h_t] \geq \log \mathbb{E}_{Q_{t-1}}[h_t]$, then $\delta_{KL}^t(h_t, \alpha_t) \geq 0$ for all $\alpha_t \in [0, 1]$.
2. *Necessary:* If $\exists \alpha_t \in (0, 1]$ such that $\delta_{KL}^t(h_t, \alpha_t) \geq 0$, then $\mathbb{E}_P[\log h_t] \geq \mathbb{E}_{Q_{t-1}}[\log h_t]$.

In contrast to additive boosting, the conditions above compare expectations under the true distribution with expectations under the *model distribution* in the previous round, Q_{t-1} . The equality in the conditions holds for $\alpha_t = 0$, which corresponds to the trivial case where the current intermediate model is ignored in Eq. (2). For other valid α_t , the non-degenerate version of the sufficient inequality guarantees progress towards the true data distribution. Note that the intermediate models increase the overall capacity of the ensemble at every round. As we shall demonstrate later, we find models fit using multiplicative boosting to outperform their additive counterparts empirically suggesting the conditions in Theorem 2 are easier to fulfill in practice.

From the necessary condition, we see that a “good” intermediate model h_t assigns a better-or-equal log-likelihood under the true distribution as opposed to the model distribution, Q_{t-1} . This condition suggests two learning algorithms for intermediate models which we discuss next.

Algorithm 1 GenBGM($X = \{\mathbf{x}_i\}_{i=1}^m, T$ rounds)

Initialize $d_0(\mathbf{x}_i) = 1/m$ for all $i = 1, 2, \dots, m$.
Obtain base generative model h_0 .
Set (unnormalized) density estimate $\tilde{q}_0 = h_0$
for $t = 1, 2, \dots, T$ **do**
- Choose β_t and update d_t using Eq. (4).
- Train generative model h_t to maximize Eq. (3).
- Choose α_t .
- Set density estimate $\tilde{q}_t = h_t^{\alpha_t} \cdot \tilde{q}_{t-1}$.
end for
Estimate $Z_T = \int \tilde{q}_T \, d\mathbf{x}$.
return $q_T = \tilde{q}_T / Z_T$.

3 Boosted generative models

In this section, we design and analyze meta-algorithms for multiplicative boosting of generative models. Given any base model which permits (approximate) likelihood evaluation, we provide a mechanism for boosting this model using an ensemble of generative and/or discriminative models.

3.1 Generative boosting

Supervised boosting algorithms such as AdaBoost typically involve a reweighting procedure for training weak learners (Freund and Schapire 1995). We can similarly train an ensemble of generative models for unsupervised boosting, where every subsequent model performs MLE w.r.t a reweighted data distribution D_t :

$$\max_{h_t} \mathbb{E}_{D_t} [\log h_t] \quad (3)$$

$$\text{where } d_t \propto \left(\frac{p}{q_{t-1}} \right)^{\beta_t} \quad (4)$$

and $\beta_t \in [0, 1]$ is the reweighting coefficient at round t . Note that these coefficients are in general different from the model weights α_t that appear in Eq. (2).

Proposition 1. *If we can maximize the objective in Eq. (3) optimally, then $\delta_{KL}^t(h_t, \alpha_t) \geq 0$ for any $\beta_t \in [0, 1]$ with the equality holding for $\beta_t = 0$.*

While the objective in Eq. (3) can be hard to optimize in practice, the target distribution becomes easier to approximate as we reduce the reweighting coefficient. For the extreme case of $\beta_t = 0$, the reweighted data distribution is simply uniform. There is no free lunch however, since a low β_t results in a slower reduction in KL-divergence leading to a computational-statistical trade-off.

The pseudocode for the corresponding boosting meta-algorithm, referred to as GenBGM, is given in Algorithm 1. In practice, we only observe samples from the true data distribution, and hence, approximate p based on the empirical data distribution which is defined to be uniform over the dataset X . At every subsequent round, GenBGM learns an intermediate model that maximizes the log-likelihood of data sampled from a reweighted data distribution.

Algorithm 2 DiscBGM($X = \{\mathbf{x}_i\}_{i=1}^m, T$ rounds, f -div)

Initialize $d_0(\mathbf{x}_i) = 1/m$ for all $i = 1, 2, \dots, m$.
Obtain base generative model h_0 .
Set (unnormalized) density estimate $\tilde{q}_0 = h_0$
for $t = 1, 2, \dots, T$ **do**
- Generate negative samples from q_{t-1}
- Optimize r_t to maximize RHS in Eq. (5).
- Set $h_t = [f']^{-1}(r_t)$.
- Choose α_t .
- Set density estimate $\tilde{q}_t = h_t^{\alpha_t} \cdot \tilde{q}_{t-1}$.
end for
Estimate $Z_T = \int \tilde{q}_T \, d\mathbf{x}$.
return $q_T = \tilde{q}_T / Z_T$.

3.2 Discriminative boosting

A base generative model can be boosted using a discriminative approach as well. Here, the intermediate model is specified as the density ratio obtained from a binary classifier. Consider the following setup: we observe an equal number of samples drawn i.i.d. from the true data distribution (w.l.o.g. assigned the label $y = +1$) and the model distribution in the previous round Q_{t-1} (label $y = 0$).

Definition 1. *Let $f : \mathbb{R}^+ \rightarrow \mathbb{R}$ be any convex, lower semi-continuous function satisfying $f(1) = 0$. The f -divergence between P and Q is defined as, $D_f(P||Q) = \int q \cdot f(p/q) \, d\mathbf{x}$.*

Notable examples include the Kullback-Liebler (KL) divergence, Hellinger distance, and the Jensen-Shannon (JS) divergence among many others. The binary classifier in discriminative boosting maximizes a variational lower bound on any f -divergence at round t :

$$D_f(P||Q_{t-1}) \geq \sup_{r_t \in \mathcal{R}_t} (\mathbb{E}_P[r_t] - \mathbb{E}_{Q_{t-1}}[f^*(r_t)]) \quad (5)$$

where f^* denotes the Fenchel conjugate of f and $r_t : \mathbb{R}^d \rightarrow \text{dom}_{f^*}$ parameterizes the classifier. Under mild conditions on f (Nguyen, Wainwright, and Jordan 2010), the lower bound in Eq. (5) is tight if $r_t^* = f'(p/q_{t-1})$.

Hence, a solution to Eq. (5) can be used to estimate density ratios. The density ratios naturally fit into the multiplicative boosting framework and provide a justification for the use of objectives of the form Eq. (5) for learning intermediate models as formalized in the proposition below.

Proposition 2. *For any given f -divergence, let r_t^* denote the optimal solution to Eq. (5) in the t^{th} round of boosting. Then, the model density at the end of the boosting round matches the true density if we set $\alpha_t = 1$ and $h_t = [f']^{-1}(r_t^*)$ where $[f']^{-1}$ denotes the inverse of the derivative of f .*

The pseudocode for the corresponding meta-algorithm, DiscBGM is given in Algorithm 2. At every round, we train a binary classifier to optimize the objective in Eq. (5) for a chosen f -divergence. As a special case, the negative of the cross-entropy loss commonly used for binary classification is also a lower bound on an f -divergence. While Algorithm 2 is applicable for any f -divergence, we will focus on cross-entropy henceforth to streamline the discussion.

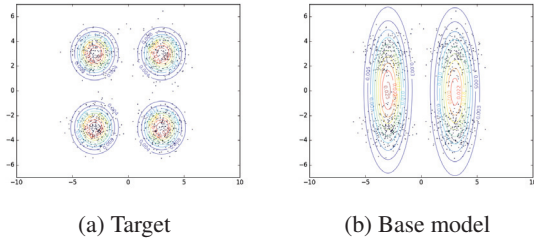


Figure 1: The mixture of Gaussians setup showing (a) true density and (b) base (misspecified) model.

Corollary 1. Consider the (negative) cross-entropy objective maximized by a binary classifier:

$$\sup_{c_t \in \mathcal{C}_t} \mathbb{E}_P[\log c_t] + \mathbb{E}_{Q_{t-1}}[\log(1 - c_t)]. \quad (6)$$

If a binary classifier c_t trained to optimize Eq. (6) is Bayes optimal, then the model density after round t matches the true density if we set $\alpha_t = 1$ and $h_t = c_t/1 - c_t$.

In practice, a classifier with limited capacity trained on a finite dataset will not generally be Bayes optimal. The above corollary, however, suggests that a good classifier can provide a ‘direction of improvement’, in a similar spirit to gradient boosting for supervised learning (Freund and Schapire 1995). Additionally, if the intermediate model distribution h_t obtained using the above corollary satisfies the conditions in Theorem 2, it is guaranteed to improve the fit.

The weights $\alpha_t \in [0, 1]$ can be interpreted as our confidence in the classification estimates, akin to the step size used in gradient descent. While in practice we heuristically assign weights to the intermediate models, the greedy optimum value of these weights at every round is a critical point for δ_{KL}^t (defined in Theorem 2). For example, in the extreme case where c_t is uninformative, *i.e.*, $c_t \equiv 0.5$, then $\delta_{KL}^t(h_t, \alpha_t) = 0$ for all $\alpha_t \in [0, 1]$. If c_t is Bayes optimal, then δ_{KL}^t attains a maxima when $\alpha_t = 1$ (Corollary 1).

3.3 Hybrid boosting

Intermediate models need not be exclusively generators or discriminators; we can design a boosting ensemble with any combination of generators and discriminators. If an intermediate model is chosen to be a generator, we learn a generative model using MLE after appropriately reweighting the data points. If a discriminator is used to implicitly specify an intermediate model, we set up a binary classification problem.

3.4 Regularization

In practice, we want boosted generative models (BGM) to generalize to data outside the training set X . Regularization in BGMs is imposed primarily in two ways. First, every intermediate model can be independently regularized by incorporating explicit terms in the learning objective, early stopping based on validation error, heuristics such as dropout, etc. Moreover, restricting the number of rounds of boosting is another effective mechanism for regularizing BGMs. Fewer rounds of boosting are required if the intermediate models are sufficiently expressive.

Table 1: Average test NLL for mixture of Gaussians.

Model	NLL (in nats, with std. error)
Base model	4.69 ± 0.01
Add model	4.64 ± 0.02
GenBGM	4.58 ± 0.10
DiscBGM-NCE	4.42 ± 0.01
DiscBGM-HD	4.35 ± 0.01

4 Empirical evaluation

Our experiments are designed to demonstrate the superiority of the proposed boosting meta-algorithms on a wide variety of generative models and tasks. A reference implementation of the boosting meta-algorithms is available at <https://github.com/ermongroup/bgm>.

4.1 Multiplicative vs. additive boosting

A common pitfall with learning parametric generative models is model misspecification with respect to the true underlying data distribution. For a quantitative and qualitative understanding of the behavior of additive and multiplicative boosting, we begin by considering a synthetic setting for density estimation on a mixture of Gaussians.

Experimental setup. The true data distribution is a equi-weighted mixture of four Gaussians centered symmetrically around the origin, each having an identity covariance matrix. The contours of the underlying density are shown in Figure 1a. We observe 1,000 training samples drawn independently from the data distribution (shown as black dots in Figure 2), and the task is to learn this distribution. The test set contains 1,000 samples from the same distribution. We repeat the process 10 times for statistical significance.

As a base (misspecified) model, we fit a mixture of two Gaussians to the data; the contours for an example instance are shown in Figure 1b. We compare multiplicative and additive boosting, each run for $T = 2$ rounds. For additive boosting (Add), we extend the algorithm proposed by Rosset and Segal (2002) setting $\hat{\alpha}_0$ to unity and doing a line search over $\hat{\alpha}_1, \hat{\alpha}_2 \in [0, 1]$. For Add and GenBGM, the intermediate models are mixtures of two Gaussians as well. The classifiers for DiscBGM are multi-layer perceptrons with two hidden layers of 100 units each and ReLU activations, trained to maximize f -divergences corresponding to the negative cross-entropy (NCE) and Hellinger distance (HD) using the Adam optimizer (Kingma and Welling 2014).

The test negative log-likelihood (NLL) estimates are listed in Table 1. Qualitatively, the contour plots for the estimated densities after every boosting round on a sample instance are shown in Figure 2. Multiplicative boosting algorithms outperform additive boosting in correcting for model misspecification. GenBGM initially leans towards maximizing coverage, whereas both versions of DiscBGM are relatively more conservative in assigning high densities to data points away from the modes.

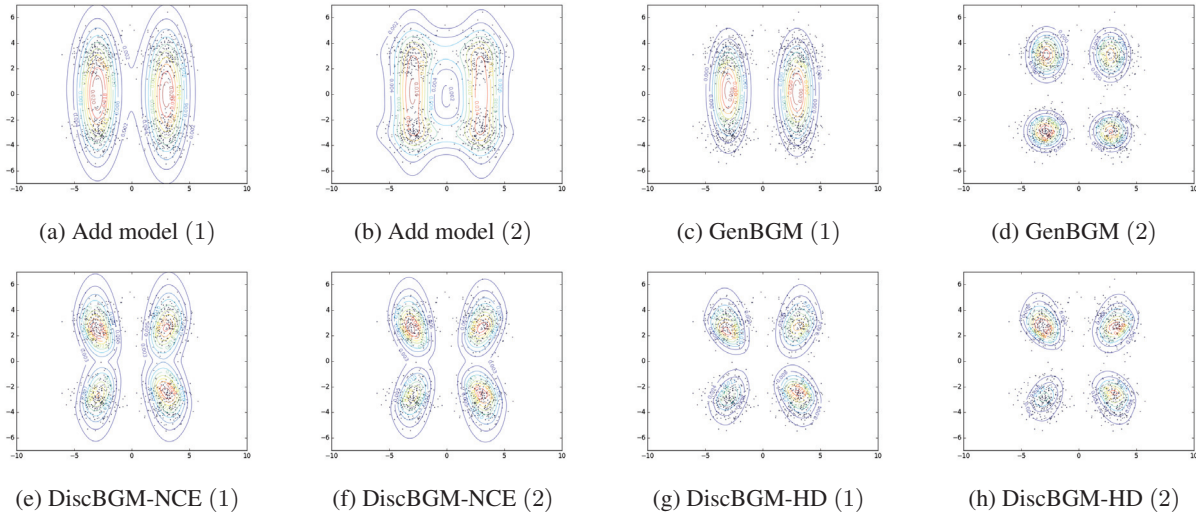


Figure 2: Multiplicative boosting algorithms such as GenBGM (c-d) and DiscBGM with negative cross-entropy (e-f) and Hellinger distance (g-h) outperform additive boosting (a-b) in correcting for model misspecification. Numbers in parenthesis indicate boosting round t .

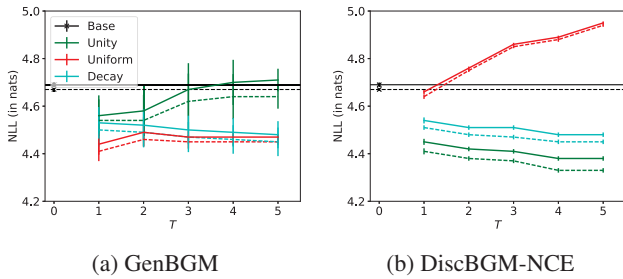


Figure 3: Train (dashed curves) and test (bold curves) NLL (in nats) for weighting heuristics on mixture of Gaussians. T is the number of rounds of boosting. The base model is shown as a black cross at $T = 0$.

Heuristic model weighting strategies. The multiplicative boosting algorithms require as hyperparameters the number of rounds of boosting and weights assigned to the intermediate models. For any practical setting, these hyperparameters are specific to the dataset and task under consideration and should be set based on cross-validation. While automatically setting model weights is an important direction for future work, we propose some heuristic weighting strategies. Specifically, the *unity* heuristic assigns a weight of 1 to every model in the ensemble, the *uniform* heuristic assigns a weight of $1/(T+1)$ to every model, and the *decay* heuristic assigns a weight of $1/2^t$ to the t^{th} model in the ensemble.

In Figure 3, we observe that the performance of the algorithms is sensitive to the weighting strategies. In particular, DiscBGM produces worse estimates as T increases for the “uniform” (red) strategy. The performance of GenBGM also degrades slightly with increasing T for the “unity” (green) strategy. Notably, the “decay” (cyan) strategy achieves stable performance for both the algorithms. Intuitively, this

heuristic follows the rationale of reducing the step size in gradient based stochastic optimization algorithms, and we expect this strategy to work better even in other settings. However, this strategy could potentially result in slower convergence as opposed to the unity strategy.

Density estimation on benchmark datasets. We now evaluate the performance of additive and multiplicative boosting for density estimation on real-world benchmark datasets (Van Haaren and Davis 2012). We consider two generative model families: mixture of Bernoullis (MoB) and sum-product networks (Poon and Domingos 2011). While our results for multiplicative boosting with sum-product networks (SPN) are competitive with the state-of-the-art, the goal of these experiments is to perform a robust comparison of boosting algorithms as well as demonstrate their applicability to various model families.

We set $T = 2$ rounds for additive boosting and GenBGM. Since DiscBGM requires samples from the model density at every round, we set $T = 1$ to ensure computational fairness such that the samples can be obtained efficiently from the base model sidestepping running expensive Markov chains. Model weights are chosen based on cross-validation. The results on density estimation are reported in Table 2. Since multiplicative boosting estimates are unnormalized, we use importance sampling to estimate the partition function.

When the base model is MoB, the Add model underperforms and is often worse than even the baseline model for the best performing validated non-zero model weights. GenBGM consistently outperforms Add and improves over the baseline model in a most cases (4/6 datasets). DiscBGM performs the best and convincingly outperforms the baseline, Add, and GenBGM on all datasets. For results on SPNs, the boosted models all outperform the baseline. GenBGM again edges out Add models (4/6 datasets), whereas DiscBGM

Table 2: Experimental results for density estimation. Negative log-likelihoods reported in nats. Lower is better with best performing models in **bold**. Overall, multiplicative boosting outperforms additive boosting and baseline models specified as Mixture of Bernoullis (MoB, **middle** columns) and Sum Product Networks (SPN, **right** columns).

Dataset	#vars	MoB Base	Add	GenBGM	DiscBGM	SPN Base	Add	GenBGM	DiscBGM
Accidents	111	42.23	43.19	41.43	34.51	31.08	29.92	29.55	28.09
Retail	135	11.27	12.24	11.20	10.91	14.94	11.27	11.21	10.88
Pumsbstar	163	55.67	55.91	50.66	34.93	26.70	25.00	25.00	23.69
DNA	180	99.42	100.37	99.23	98.45	92.60	86.93	87.79	86.63
Kosarek	190	11.72	12.57	12.41	11.13	12.71	10.97	10.73	10.67
Ad	1556	63.13	63.73	63.19	54.79	19.19	18.12	18.14	17.82

Table 3: Experimental results for classification. Prediction accuracy for predicting one variable given the rest. Higher is better with best performing models in **bold**. Multiplicative boosting again outperforms additive boosting and baseline models specified as Mixture of Bernoullis (MoB, **middle** columns) and Sum Product Networks (SPN, **right** columns).

Dataset	#test	MoB Base	Add	GenBGM	DiscBGM	SPN Base	Add	GenBGM	DiscBGM
Accidents	283,161	0.8395	0.8393	0.8473	0.9043	0.9258	0.9266	0.9298	0.9416
Retail	595,080	0.9776	0.9776	0.9776	0.9792	0.9780	0.9790	0.9789	0.9791
Pumsb-star	399,676	0.8461	0.8501	0.8819	0.9267	0.9599	0.9610	0.9611	0.9636
DNA	213,480	0.7517	0.7515	0.7531	0.7526	0.7799	0.7817	0.7828	0.7811
Kosarek	1,268,250	0.9817	0.9816	0.9818	0.9831	0.9824	0.9838	0.9838	0.9838
Ad	763,996	0.9922	0.9923	0.9818	0.9927	0.9982	0.9981	0.9982	0.9982

models outperform all other models on all datasets. These results demonstrate the usefulness of boosted expressive model families, especially the DiscBGM approach, which performs the best, while GenBGM is preferable to Add.

4.2 Applications of generative models

Classification. Here, we evaluate the performance of boosting algorithms for classification. Since the datasets above do not have any explicit labels, we choose one of the dimensions to be the label (say y). Letting $\mathbf{x}_{\bar{y}}$ denote the remaining dimensions, we can obtain a prediction for y as,

$$p(y = 1 | \mathbf{x}_{\bar{y}}) = \frac{p(y = 1, \mathbf{x}_{\bar{y}})}{p(y = 1, \mathbf{x}_{\bar{y}}) + p(y = 0, \mathbf{x}_{\bar{y}})}$$

which is efficient to compute even for unnormalized models. We repeat the above procedure for all the variables predicting one variable at a time using the values assigned to the remaining variables. The results are reported in Table 3. When the base model is a MoB, we observe that the Add approach could often be worse than the base model whereas GenBGM performs slightly better than the baseline (4/6 datasets). The DiscBGM approach consistently performs well, and is only outperformed by GenBGM for two datasets for MoB. When SPNs are used instead, both Add and GenBGM improve upon the baseline model while DiscBGM again is the best performing model on all but one dataset.

Sample generation. We compare boosting algorithms based on their ability to generate image samples for the binarized MNIST dataset of handwritten digits (LeCun, Cortes, and Burges 2010). We use variational autoencoders (VAE) as the base model (Kingma and Welling 2014). While any sufficiently expressive VAE can generate impressive examples, we design the experiment to evaluate the model complexity approximated as the number of learnable parameters.

Ancestral samples obtained by the baseline VAE model are shown in Figure 4a. We use the evidence lower bound (ELBO) as a proxy for approximately evaluating the marginal log-likelihood during learning. The conventional approach to improving the performance of a latent variable model is to increase its representational capacity by adding hidden layers (Base + depth) or increasing the number of hidden units in the existing layers (Base + width). These lead to a marginal improvement in sample quality as seen in Figure 4b and Figure 4c.

In contrast, boosting makes steady improvements in sample quality. We start with a VAE with much fewer parameters and generate samples using a hybrid boosting GenDiscBGM sequence VAE→CNN→VAE (Figure 4d). The discriminator used is a convolutional neural network (CNN) (LeCun and Bengio 1995) trained to maximize the negative cross-entropy. We then generate samples using independent Markov chain Monte Carlo (MCMC) runs. The boosted sequences generate sharper samples than all baselines in spite of having similar model capacity.

5 Discussion and related work

In this work, we revisited boosting, a class of meta-algorithms developed in response to a seminal question: *Can a set of weak learners create a single strong learner?* Boosting has offered interesting theoretical insights into the fundamental limits of supervised learning and led to the development of algorithms that work well in practice (Schapire 1990; Freund, Schapire, and Abe 1999; Friedman 2002; Caruana and Niculescu-Mizil 2006). Our work provides a foundational framework for unsupervised boosting with connections to prior work discussed below.

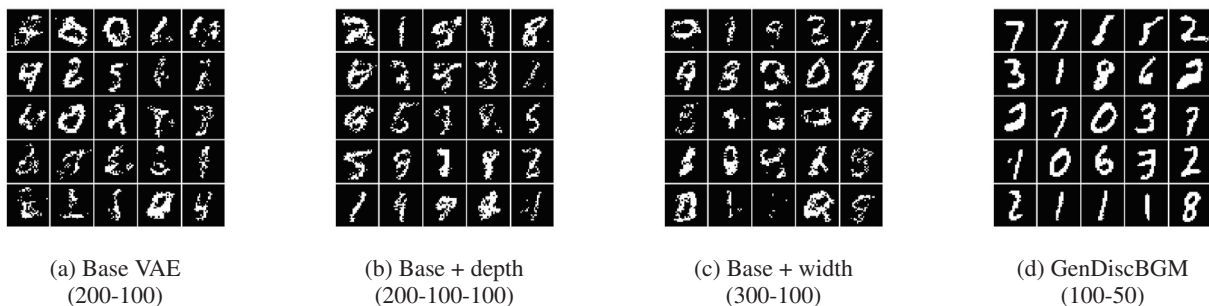


Figure 4: The boosted model (d) demonstrates how ensembles of weak learners can generate sharper samples, compared to naively increasing model capacity (a-c). Note that we show samples of binarized digits and not mean values for the pixels. VAE hidden layer architecture given in parenthesis.

Sum-of-experts. Rosset and Segal (2002) proposed an algorithm for density estimation using Bayesian networks similar to gradient boosting. These models are normalized and easy to sample, but are generally outperformed by multiplicative formulations for correcting for model misspecification, as we show in this work. Similar additive approaches have been used for improving approximate posteriors for specific algorithms for variational inference (Guo et al. 2016; Miller, Foti, and Adams 2017) and generative adversarial networks (Tolstikhin et al. 2017). For a survey on variations of additive ensembling for unsupervised settings, refer to the survey by Bourel and Ghattas (2012).

Product-of-experts. Our multiplicative boosting formulation can be interpreted as a product-of-experts approach, which was initially proposed for feature learning in energy based models such as Boltzmann machines. For example, the hidden units in a restricted Boltzmann machine can be interpreted as weak learners performing MLE. If the number of weak learners is fixed, they can be efficiently updated in parallel but there is a risk of learning redundant features (Hinton 1999; 2002). Weak learners can also be added incrementally based on the learner’s ability to distinguish observed data and model-generated data (Welling, Zemel, and Hinton 2002). Tu (2007) generalized the latter to boost arbitrary probabilistic models; their algorithm is a special case of DiscBGM with all α ’s set to 1 and the discriminator itself a boosted classifier. DiscBGM additionally accounts for imperfections in learning classifiers through flexible model weights. Further, it can include any classifier trained to maximize any f -divergence.

Related techniques such as noise-contrastive estimation, ratio matching, and score matching methods can be cast as minimization of Bregman divergences, akin to DiscBGM with unit model weights (Gutmann and Hirayama 2011). A non-parametric algorithm similar to GenBGM was proposed by Di Marzio and Taylor (2004) where an ensemble of weighted kernel density estimates are learned to approximate the data distribution. In contrast, our framework allows for both parametric and non-parametric learners and uses a different scheme for reweighting data points than proposed in the above work.

Unsupervised-as-supervised learning. The use of density ratios learned by a binary classifier for estimation was first proposed by Friedman, Hastie, and Tibshirani (2001) and has been subsequently applied elsewhere, notably for parameter estimation using noise-contrastive estimation (Gutmann and Hyvärinen 2010) and sample generation in generative adversarial networks (GAN) (Goodfellow et al. 2014). While GANs consist of a discriminator distinguishing real data from model generated data similar to DiscBGM for a suitable f -divergence, they differ in the learning objective for the generator (Nowozin, Cseke, and Tomioka 2016). The generator of a GAN performs an adversarial minimization of the same objective the discriminator maximizes, whereas DiscBGM uses the likelihood estimate of the base generator (learned using MLE) and the density ratios derived from the discriminator(s) to estimate the model density for the ensemble.

Limitations and future work. In the multiplicative boosting framework, the model density needs to be specified only up to a normalization constant at any given round of boosting. Additionally, while many applications of generative modeling such as feature learning and classification can sidestep computing the partition function, if needed it can be estimated using techniques such as Annealed Importance Sampling (Neal 2001). Similarly, Markov chain Monte Carlo methods can be used to generate samples. The lack of implicit normalization can however be limiting for applications requiring fast log-likelihood evaluation and sampling.

In order to sidestep this issue, a promising direction for future work is to consider boosting of normalizing flow models (Dinh, Krueger, and Bengio 2014; Dinh, Sohl-Dickstein, and Bengio 2017; Grover, Dhar, and Ermon 2018). These models specify an invertible multiplicative transformation from one distribution to another using the change-of-variables formula such that the resulting distribution is self-normalized and efficient ancestral sampling is possible. The GenBGM algorithm can be adapted to normalizing flow models whereby every transformation is interpreted as a weak learner. The parameters for every transformation can be trained greedily after suitable reweighting resulting in a self-normalized boosted generative model.

6 Conclusion

We presented a general-purpose framework for boosting generative models by explicit factorization of the model likelihood as a product of simpler intermediate model densities. These intermediate models are learned greedily using discriminative or generative approaches, gradually increasing the overall model’s capacity. We demonstrated the effectiveness of these models over baseline models and additive boosting for the tasks of density estimation, classification, and sample generation. Extensions to semi-supervised learning (Kingma et al. 2014) and structured prediction (Sohn, Lee, and Yan 2015) are exciting directions for future work.

Acknowledgements

We are thankful to Neal Jean, Daniel Levy, and Russell Stewart for helpful critique. This research was supported by a Microsoft Research PhD fellowship in machine learning for the first author, NSF grants #1651565, #1522054, #1733686, a Future of Life Institute grant, and Intel.

References

- Bourel, M., and Ghattas, B. 2012. Aggregating density estimators: An empirical study. *arXiv preprint arXiv:1207.4959*.
- Caruana, R., and Niculescu-Mizil, A. 2006. An empirical comparison of supervised learning algorithms. In *International Conference on Machine Learning*.
- Di Marzio, M., and Taylor, C. C. 2004. Boosting kernel density estimates: A bias reduction technique? *Biometrika* 91(1):226–233.
- Dinh, L.; Krueger, D.; and Bengio, Y. 2014. NICE: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2017. Density estimation using Real NVP. In *International Conference on Learning Representations*.
- Freund, Y., and Schapire, R. E. 1995. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*.
- Freund, Y.; Schapire, R.; and Abe, N. 1999. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence* 14(771-780):1612.
- Friedman, J.; Hastie, T.; and Tibshirani, R. 2001. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin.
- Friedman, J. H. 2002. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38(4):367–378.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*.
- Grover, A., and Ermon, S. 2017. Boosted generative models. *arXiv preprint arXiv:1702.08484*.
- Grover, A.; Dhar, M.; and Ermon, S. 2018. Flow-GAN: Combining maximum likelihood and adversarial learning in generative models. In *AAAI Conference on Artificial Intelligence*.
- Guo, F.; Wang, X.; Fan, K.; Broderick, T.; and Dunson, D. B. 2016. Boosting variational inference. *arXiv preprint arXiv:1611.05559*.
- Gutmann, M., and Hirayama, J.-i. 2011. Bregman divergence as general framework to estimate unnormalized statistical models. In *Uncertainty in Artificial Intelligence*.
- Gutmann, M., and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Artificial Intelligence and Statistics*.
- Hinton, G. E. 1999. Products of experts. In *International Conference on Artificial Neural Networks*.
- Hinton, G. E. 2002. Training products of experts by minimizing contrastive divergence. *Neural computation* 14(8):1771–1800.
- Kingma, D. P., and Welling, M. 2014. Auto-encoding variational Bayes. In *International Conference on Learning Representations*.
- Kingma, D. P.; Mohamed, S.; Rezende, D. J.; and Welling, M. 2014. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*.
- LeCun, Y., and Bengio, Y. 1995. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks* 3361(10):1995.
- LeCun, Y.; Cortes, C.; and Burges, C. J. 2010. MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist>.
- Li, Y.; Song, J.; and Ermon, S. 2017. InfoGAIL: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*.
- Miller, A. C.; Foti, N.; and Adams, R. P. 2017. Variational boosting: Iteratively refining posterior approximations. In *International Conference on Machine Learning*.
- Neal, R. M. 2001. Annealed importance sampling. *Statistics and Computing* 11(2):125–139.
- Nguyen, X.; Wainwright, M. J.; and Jordan, M. I. 2010. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory* 56(11):5847–5861.
- Nowozin, S.; Cseke, B.; and Tomioka, R. 2016. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*.
- Oord, A. v. d.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. In *International Conference on Machine Learning*.
- Poon, H., and Domingos, P. 2011. Sum-product networks: A new deep architecture. In *Uncertainty in Artificial Intelligence*.
- Rosset, S., and Segal, E. 2002. Boosting density estimation. In *Advances in Neural Information Processing Systems*.
- Schapire, R. E., and Freund, Y. 2012. *Boosting: Foundations and algorithms*. MIT press.
- Schapire, R. E. 1990. The strength of weak learnability. *Machine learning* 5(2):197–227.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*.
- Tolstikhin, I.; Gelly, S.; Bousquet, O.; Simon-Gabriel, C.-J.; and Schölkopf, B. 2017. AdaGAN: Boosting generative models. In *Advances in Neural Information Processing Systems*.
- Tu, Z. 2007. Learning generative models via discriminative approaches. In *Computer Vision and Pattern Recognition*.
- Van Haaren, J., and Davis, J. 2012. Markov network structure learning: A randomized feature generation approach. In *AAAI Conference on Artificial Intelligence*.
- Welling, M.; Zemel, R. S.; and Hinton, G. E. 2002. Self supervised boosting. In *Advances in Neural Information Processing Systems*.
- Zhao, S.; Song, J.; and Ermon, S. 2017. Learning hierarchical features from deep generative models. In *International Conference on Machine Learning*.