Published in final edited form as:

Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit. 2020 June ; 2020: 11432–11440.
doi:10.1109/cvpr42600.2020.01145.

Generating Accurate Pseudo-labels in Semi-Supervised Learning and Avoiding Overconfident Predictions via Hermite Polynomial Activations

Vishnu Suresh Lokhande^{†,1}, Songwong Tasneeyapant¹, Abhay Venkatesh¹, Sathya N. Ravi^{¶,2}, Vikas Singh^{||,1}

¹University of Wisconsin-Madison

²University of Illinois at Chicago

Abstract

Rectified Linear Units (ReLUs) are among the most widely used activation function in a broad variety of tasks in vision. Recent theoretical results suggest that despite their excellent practical performance, in various cases, a substitution with basis expansions (e.g., polynomials) can yield significant benefits from both the optimization and generalization perspective. Unfortunately, the existing results remain limited to networks with a couple of layers, and the practical viability of these results is not yet known. Motivated by some of these results, we explore the use of Hermite polynomial expansions as a substitute for ReLUs in deep networks. While our experiments with supervised learning do not provide a clear verdict, we find that this strategy offers considerable benefits in semi-supervised learning (SSL) / transductive learning settings. We carefully develop this idea and show how the use of Hermite polynomials based activations can yield improvements in pseudo-label accuracies and sizable financial savings (due to concurrent runtime benefits). Further, we show via theoretical analysis, that the networks (with Hermite activations) offer robustness to noise and other attractive mathematical properties. Code is available on //GitHub.

1. Introduction

Analyzing the optimization or the loss landscape of deep neural networks has emerged as a promising means to understand the behavior and properties of various neural network architectures [4]. One reason is that insights into how the loss function behaves geometrically is closely tied to the types of optimization schemes that may be needed [33], why specific ideas work whereas others do not [25], and how or whether the corresponding model may generalize to unseen data [34]. Notice that we must leverage such “alternative” strategies as a window into these models’ behavior because in deep learning, most models of interest are highly non-linear and non-convex. As a result, one finds that extending mature ideas, which work well for analysis in classical settings (e.g., linear models), is quite a bit harder for most deep architectures of interest.

[†] lokhande@cs.wisc.edu. [¶] sathya@uic.edu. ^{||} vsingh@biostat.wisc.edu.

Why study activation functions?

While there are many ways we may study the optimization landscape of deep models, a productive line of recent results [6] proposes analyzing the landscape (i.e., the behavior of the loss as a function of the network parameters) via the activation functions of the neural network. This makes a lot of sense because the activation function is one critical place where we introduce non-linearity into the network, and their omission significantly simplifies any analysis. Activation functions greatly influence the functional space which a neural network can represent [15], often the first step in a more formal study of the model's behavior. For example, universal finite-sample expressivity of certain architectures has been shown by fixing the activations to be ReLU functions [9]. In other words, if we use ReLU as activations, such an architecture can be shown to represent any function if the model has more parameters than the sample size. The scope of such work is not just theoretical – for instance, the above results were used to derive a much simpler architecture consisting of only residual convolutional layers and ReLU activations. Further, the estimation scheme needed was also much simpler and required no batch normalization, dropout, or max pooling. In summary, the choice of activations enabled understanding the loss landscape and enabled simplifications.

More on activations functions and loss landscape.

The authors in [31] showed that conditions that prevent presence of spurious valleys on the loss landscape can be identified, via the use of smooth activations. Independently, [28] demonstrated that the use of quadratic activation functions enables efficient localization of global minima in certain classes of deep models. Closely related to smooth and quadratic activations, the use of polynomial non-linearity as an activation has also been studied in the last few years. For example, [22] studied deep ReLU networks by using a polynomial non-linearity as an approximation for ReLU: this enabled a much cleaner analysis of the empirical risk landscape. More recently, [15] analyzed the functional space of the networks with the help of polynomial activation based on techniques from algebraic geometry. Earlier, for a one hidden layer network, [6] investigated optimizing the population risk of the loss using stochastic gradient descent and showed that one could avoid spurious local minima by utilizing an orthogonal basis expansion for ReLUs. A key takeaway from this work is that the optimization would behave better if the landscape was nicely behaved — this is enabled via the basis expansion. The foregoing results and analyses, especially the use of basis expansion, is interesting and a starting point for the development described here. A common feature of the body of work summarized above is that they rely on networks with polynomial activations (polynomial networks) to analyze the loss landscape. This choice helps make the mathematical exploration easier.

Where is the gap in the literature?

Despite the interesting theoretical results summarized above, relatively less is known whether such a strategy or its variants are a good idea for the architectures in computer vision and broadly, in AI, today. We can, in fact, ask a more practically focused question: are there specific tasks in computer vision where such a strategy offers strong empirical advantages? This is precisely the gap this paper is designed to address.

The main contributions include: **(a)** We describe mechanisms via which activation functions based on Hermite polynomials can be utilized within deep networks instead of ReLUs, with only minor changes to the architecture. **(b)** We present evidence showing that while these adjustments are not significantly advantageous in supervised learning, our scheme *does* yield sizable advantages in semi-supervised learning speeding up convergence. Therefore, it offers clear benefits in compute time (and cost) needed to attain a certain pseudo-label accuracy, which has direct cost implications. **(c)** We give technical results analyzing the mathematical behavior of such activations, specifically, robustness results showing how the activation mitigates overconfident predictions for (out of distribution) samples.

2. Brief Review of Hermite polynomials

We will use an expansion based on Hermite polynomials as a substitute for ReLU activations. To describe the construction clearly, we briefly review the basic properties.

Hermite polynomials are a class of orthogonal polynomials which are appealing both theoretically as well as in practice, e.g., radio communications [2]. Here, we will use Hermite polynomials [23] defined as,

$$H_i(x) = (-1)^i e^{x^2} \frac{d^i}{dx^i} e^{-x^2}, i > 0; H_0(x) = 1 \quad (1)$$

In particular, we use normalized Hermite polynomials which are given as $h_i = \frac{H_i}{\sqrt{i!}}$. Hermite polynomials are often used in the analysis of algorithms for nonconvex optimization problems [20]. While there are various mathematical properties associated with Hermite polynomials, we will now discuss the most important property for our purposes.

Hermite polynomials as bases.

Classical results in functional analysis show that the (countably infinite) set $\{H_i\}_{i=0}^{\infty}$ defined in (1) can be used as bases to represent smooth functions [24]. Formally, let $L^2(\mathbb{R}, e^{-x^2/2})$ denote the set of integrable functions w.r.t. the Gaussian measure,

$$L^2(\mathbb{R}, e^{-x^2/2}) = \left\{ f: \int_{-\infty}^{+\infty} f(x)^2 e^{-x^2/2} dx < \infty \right\},$$

It turns out that $L^2(\mathbb{R}, e^{-x^2/2})$ is a Hilbert space with the inner product defined as follows (see [6] for more details)

$$\langle f, g \rangle = \mathbb{E}_{x \sim \mathcal{N}(0, 1)}[f(x)g(x)] \quad (2)$$

The normalized Hermite polynomials form an *orthonormal basis* in $L^2(\mathbb{R}, e^{-x^2/2})$ in the sense that $\langle h_i, h_j \rangle = \delta_{ij}$. Here δ_{ij} is the Kronecker delta function and h_i, h_j are any two normalized Hermite polynomials.

Recently, the authors in [6] showed that the lower order terms in the Hermite polynomial series expansion of ReLU have a different rate of convergence on the optimization landscape than the higher order terms for one hidden layer networks. This leads to a natural question: are these properties useful to accelerate the performance of gradient-based methods for deep networks as well?

Hermite Polynomials as Activations.

Let $x = (x_1, \dots, x_n)$ be an input to a neuron in a neural network and y be the output. Let $w = (w_1, w_2, \dots, w_n)$ be the weights associated with the neuron. Let σ be the non-linear activation applied to $w^T x$. Often we set $\sigma = \text{ReLU}$. Here, we investigate the scenario where $\sigma(x) = \sum_{i=0}^d c_i h_i(x)$, denoted as σ_{hermite} , where h_i 's are as defined previously and c_i 's are trainable parameters. As [6] suggests, we initialized the parameters c_i 's associated with hermites to be $c_i = \hat{\sigma}_i$, where $\hat{\sigma}_i = \langle \text{ReLU}, h_i \rangle$. Hermite polynomials as activations on a single layer neural network with a single hidden unit can be seen in Figure 1.

3. Sanity checks: What do we gain or lose?

Replacing ReLUs with other activation functions reviewed in Section 1 has been variously attempted in the literature already, for supervised learning. While improvements have been reported in specific settings, ReLUs continue to be relatively competitive. Therefore, it seems that we should not expect improvements in what are toy supervised learning experiments. However, as we describe below, the experiments yield useful insights regarding settings where Hermes will be particularly useful.

A) Two-layer architectures.

We start with the CIFAR10 dataset and a simple two-layer network (more details in the appendix § 7.8). Basically, we ask if this modification will lead to better or worse performance in accuracy and runtime (over ReLU activation) with all other parameters (e.g., learning rates) fixed. Positives: Here, we observed *faster* training loss convergence (over the initial epochs) compared to ReLU in general. Negatives: When we assessed the performance as a function of the number of Hermite polynomials d , we see a tradeoff where the speeds first improve with increasing d and then worsen when d is as high as 8, indicating that too many bases, especially for relatively shallow architectures are not ideal.

B) Deep Autoencoders.

Encouraged by the shallow network experiments, we moved to a standard benchmark for neural network optimization runtimes: the deep autoencoders setup from [33]. We replicate the experiments reported in [33] for the MNIST dataset: we use 4 Hermite polynomials as an activation instead of the original sigmoid. Positives: We see from Table 1 that Hermite

activations not only converge faster but also achieve lower test loss. We also evaluated how activation functions such as ReLU, SeLU or ELU perform in this setting. As shown in Figure 2, we find that Hermite still offer runtime improvements here with a comparable validation loss. We find that SeLU does not seem to be as stable as the other activation functions. Therefore, in the remainder of this paper, we study Hermite activations compared to ReLUs as a baseline. Of course, it is quite possible that some other activations may provide slightly better performance compared to ReLUs in specific settings. But the choice above simplifies the presentation of our results and is consistent with the general message of our work: while research on polynomial activations has so far remained largely theoretical, we intend to show that they are easily deployable and offer various practical benefits.

Adjustments needed for deeper architectures?

When implemented naively, Hermite activations do not work well for deeper architectures directly, which may be a reason that they have not been carefully explored so far. Basically, with no adjustments, we encounter a number of numerical issues that are not easy to fix. In fact, [8] explicitly notes that higher-order polynomials tend to make the activations unbounded making the training unstable. Fortunately, a trick mentioned in [1] in the context of quadratic functions, addresses the problem. The solution is to add a softsign function in (3), which has a form similar to tanh however, it approaches its maximal (and minimal) value slower than tanh.

$$\text{Softsign}(x) = \frac{x}{1 + |x|} \quad (3)$$

C) ResNet18.

With the above modification in hand, we can use our activation function within Resnet18 using Pre-activation Blocks [10, 11]. In the preactivation block, we found that having the second softsign function *after* the weight layer is useful. The slightly modified preactivation block of ResNets is shown in Figure 3. We train ResNets with Hermite activations on CIFAR10 to assess the general behavior of our substitution. We use SGD as an optimizer and follow the data augmentation techniques and the learning rate schedules as in [10, 3]. We perform cross-validation for the hyper-parameters. After training, we obtain the loss curve and the training set accuracies for Hermite activations and ReLUs (see Fig. 4). Positives: Figure 4 shows that the loss values and trainset accuracies converge at a much faster rate for Hermite than ReLUs. While the test set accuracy at the *final* epoch is higher for ReLU than Hermite, the test set accuracies for networks using Hermite activations *make much quicker progress in the initial epochs*. These results hold even when varying the learning rates of ReLU networks (see appendix § 7.4). Negatives: We also tune the choice of the number of Hermite polynomials d for $d \in \{0, 2, 4, 6, 8, 10\}$. The setting $d = 0$ is the case where we only use a softsign as an activation without any Hermite activations. Figure 4b shows the results of this experiment. From the plot, we observe a trend similar to the two-layer network above, where the convergence speeds first improves and then reduces as we increase the number of Hermite polynomials. The close overlap of the trend lines between d

$= 4$ to $d = 10$ suggest that tuning for d is not very expensive. Hence, in all our experiments we mostly set $d = 4$. The setting $d = 0$ performs worse than when d is at least one, suggesting that the performance benefits are due to Hermite activations with softsign (and not the softsign function on its own).

Interesting take away from experiments?

Let us consider the negative results first where we find that a large number of bases is not useful. This makes sense where some flexibility in terms of trainable parameters is useful, but too much flexibility is harmful. Therefore, we simply need to set d to a small (but not too small) constant. On the other hand, the positive results from our simple experiments above suggest that networks with Hermite activations make rapid progress in the early to mid epoch stages – an early riser property – and the performance gap becomes small later on. This provides two potential strategies. If desired, we could design a hybrid optimization scheme that exploits this behavior. One difficulty is that initializing a ReLU based network with weights learned for a network with Hermite activations (and retraining) may partly offset the benefits from the quicker progress made in the early epochs. What will be more compelling is to utilize the Hermite activations end to end, but identify scenarios where this “early riser” property is critical and directly influences the final goals or outcomes of the task. It turns out that recent developments in semi-supervised learning satisfy exactly these conditions where leveraging the early riser property is immensely beneficial.

4. Semi-Supervised Learning (SSL)

SSL, Transductive learning and pseudo-labeling.

We consider the class of algorithms that solve SSL by generating pseudo-labels for the unlabelled data. These *Pseudo-label (PL) based SSL* methods serve as a way to execute transductive learning [14, 26], where label inference on the given test dataset is more important for the user as compared to using the trained model later, on other unseen samples from other data sources. Stand-alone PL based methods can be used for generic SSL problems also – where the generation of pseudo-labels are merely a means to an end. The literature suggests that in general they perform competitively, but some specialized SSL methods may provide some marginal benefits. In the other direction, general (non PL based) SSL methods can also be used for obtaining pseudo-labels. But PL based SSL generates pseudo-labels concurrently with (in fact, to facilitate) estimating the parameters of the model. So, if the eventual goal is to obtain labels for unlabeled data at hand, these methods are a good fit – they explicitly assign (potentially) high accurate labels for the unlabeled data at the end of training.

Improved deep classifiers by assigning PL.

It is well known that poor (random) labels are harder to fit (train) [34], or equivalently, high quality (pseudo) labels accelerates the training procedure. In other words, highly accurate labels $\rightarrow$ fast training. Interestingly, [5] showed that the converse statement, i.e., fast training $\rightarrow$ accuracy of labels, is *also* true using empirical evidence, during training. SaaS (Speed as a Supervisor) is used to *estimate* the pseudo-labels

Algorithm 1

Input: Labeled data (x_i, y_i) , unlabeled data (z_i) , number of classes k , #-outer (inner) epochs $M_O(M_I)$, loss function $L = L_{CE}(x_i, y_i) + L_{CE}(z_i, y_i) + Reg_E(z_i, y_i)$, learning rates $\eta_w, \eta_P^p, \eta_P^d$. Initial Pseudo-labels for unlabeled data chosen as: $y_i = e_i$ with probability $1/k$ where e_i is the one-hot vector at i -th coordinate.

for $O = 0, 1, 2, \dots, M_O$ **do**
 Reinitialize the network parameters w^0
 $P_u = 0$
for $I = 0, 1, 2, \dots, M_I$ **do**
 (Primal) SGD Step on w : $w^{t+1} \leftarrow w^t - \eta_w \nabla L$
 (Primal) SGD Step on ΔP_u : $\Delta P_u \leftarrow \Delta P_u - \eta_P^p \nabla L$
end for
 (Dual) SGD Step on P_u : $P_u \leftarrow P_u - \eta_P^d \Delta P_u$

end for

Output: Classification model w .

using “speed” as a surrogate. That is, for a classifier such as ResNet, SaaS takes advantage of the fact that the *loss decreases at a faster rate for correct labels* (compared to random labels) during training. Furthermore, the *rate of loss reduction decreases as the percentage of incorrect labels in the dataset increases*. The SaaS framework. There are two phases in SaaS. In the primal phase (inner loop), SaaS seeks to find a set of pseudo-labels that decreases the loss function over a *small number of epochs*, the most. In the dual phase (outer loop), the “pseudo-label” optimization is carried out by computing a posterior distribution over the unlabeled data. Specifically, the algorithm consists of two loops: (i) in the outer loop, we optimize over the posterior distribution of the pseudo-labels, and (ii) in the inner loop, we retrain the network with fixed pseudo-labels. After every inner loop, we reinitialize the network, and compute the rate of change of the loss value with which the posterior can be computed. A flowchart is shown in Figure 5. We can easily use a ResNet/DenseNet model in the primal phase. There are two different loss functions that are optimized during training: the cross-entropy loss (L_{CE}) over the labeled data, the unlabeled data and an entropy regularizer (Reg_E) on the pseudo-labels. A pseudocode is provided in Algorithm 1.

Pseudo-labels with Hermites.

Recall that networks with Hermite activations manifest the “early riser” property. This turns out to be ideal in the setting described above and next, we show how this property can be exploited for transductive/semi-supervised learning. Intuitively, the early riser property implies that the training loss decreases at a much faster rate in the initial epochs. This is expected, since the optimization landscape, when we use Hermites are, by definition, *smoother* than ReLUs, since *all* the neurons are *always* active during training with probability 1.

Setting up.

For our experiments, we used a ResNet-18 architecture (with a preactivation block) [11] architecture to run SaaS [5] with one crucial difference: ReLU activations were replaced with Hermite activations. All other hyperparameters needed are provided in Table 2. We will call this version, Hermite-SaaS. To make sure that our findings are broadly applicable, we conducted experiments with four datasets commonly used in the semi-supervised learning literature: SVHN, CIFAR10, SmallNORB, and MNIST.

4.1. Faster Convergence.

SaaS tries to identify labels on which the training loss decreases at a faster rate. Our experiments show that the use of Hermite provides models on which the loss function can be optimized easily, thus stabilizing the two phase procedure. Figure 7 shows the result of our experiment on CIFAR10 dataset. Notice that the periodic jumps in the loss value is expected. This is because the big jumps correspond to the termination of an inner epoch, where the pseudo-labels are updated and weights are reinitialized. From Figure 7, we observe that Hermite-SaaS provides a smoother landscape during training, accelerating the training process overall.

4.2. Computational Benefits.

When the accuracy of pseudo-labels is assessed against the actual labels for SVHN dataset, we obtain an error rate of 5.79 over the pseudo-label error rate of 6.22 in [5]. This indicates that the quality of pseudo-labels can be significantly improved using Hermite-SaaS. Moreover, we also find that the number of epochs needed to reach a specific pseudo-label accuracy is also significantly lower. Table 3 shows two common metrics used in classification: (i) “max gap in accuracy” (Max) measures the maximum difference in pseudo-label accuracy over epochs during training; and (ii) “pseudo-label accuracy” (Max PL ACC) measures the maximum pseudo-label accuracy attained during the course of training. Both these measures form a proxy to assess the the quality of pseudo-labels. In addition, Figure 6 shows that the number of epochs needed to reach a specific pseudo-label accuracy is significantly lower when using Hermite-SaaS. This behavior for Hermite-SaaS holds over all the four datasets.

4.3. Financial Savings.

ReLU takes less time per iteration since in our Hermite-SaaS procedure, we must perform a few more matrix-vector multiplications. However, our experimental results indicate that the lower *per iteration* time of ReLU is negligible as compared to the total number of iterations. To provide a better context for the amount of savings possible using Hermite-SaaS, we performed an experiment to assess financial savings. We use AWS p3.16x large cluster for training. We calculate the cost (to train a network) by running the algorithm to reach a minimum level of pseudo-label accuracy, using the default pricing model given by AWS. We can clearly see from Table 3, that we get significant cost savings if we use Hermite-SaaS. Note that we found cases where ReLU-SaaS could not reach the pseudo-label accuracy that

is achieved by Hermite-SaaS: in these cases, we report a conservative lower bound for cost savings.

4.4. Better Generalization.

From the work of [34], we know that more accurate labels provide better generalization. With more accurate pseudo-labels from Hermite-SaaS, we expect a standard supervised classifier, trained using the pseudo-labels, to provide better accuracy on an unseen test set. We indeed observe this behavior in our experiments. The left sub-table in Table 4 shows the performance of Hermite-SaaS on SVHN dataset against popular SSL methods reported in [5]. We also compare against known PL based SSL methods for CIFAR10 dataset in right sub-table of Table 4. We use ResNet18 with a preactivation block for the supervised training phase, although other networks can be used.

4.5. Noise Tolerance.

SSL approaches, especially PL methods, suffer from confirmation bias [30] which is when the network starts to predict incorrect labels with higher confidence and resists change in the course of training. A recent result describes how ReLU networks [12] tend to provide high confidence predictions even when the test data is different from the training data: this can be used in adversarial settings [21]. In other words, [30] and [12] taken together hint at the possibility that ReLUs may, at least, partly contribute to the confirmation bias in PL based SSL methods. While evaluating this question comprehensively is an interesting topic on its own, we present some interesting preliminary results. We can assess the question partly by asking whether Hermite-SaaS is more tolerant to noise than ReLU-SaaS: approached by artificially injecting noise into the labels. In this section, we first evaluate this issue on the theoretical side (for one hidden-layer Hermite networks) where we find that Hermite networks do not give (false) high confidence predictions when the test data are different from the train data. We later demonstrate noise tolerance of Hermite-SaaS.

Accurate confidence of predictions in one-hidden layer hermite networks.—In order to prove our main result, we first prove a generic perturbation bound in the following lemma. Consider a 2 layer network with an input layer, hidden layer and an output layer, each with multiple units. Denote $f_k(x)$ and $f_l(x)$ to be two output units with x denoting the input vector. Let w_j be the weight vector between input and the j^{th} and a_{jk} be the weight connecting j^{th} hidden unit to the k^{th} output unit.

Lemma 1. *Consider the output unit of the network, $f_k(x) = \sum_j a_{kj} \sum_{i=0}^d c_i h_i(w_j^T x)$, where c_i 's are the Hermite coefficients and d is the maximum degree of the hermite polynomial considered. Then,*

$$|f_l(x) - f_k(x)| \leq C d \alpha \beta$$

Here, C is a constant proportional to the ℓ_∞ norms of c_i 's;

$$\alpha = \max_{lk} \sum_j |a_{lj} - a_{kj}|; \beta = \max \left(\|w\|_p^d \|x\|_q^d, \|w\|_p \|x\|_q \right),$$

such that $1/p + 1/q = 1$.

Now, we use the perturbation bound from Lemma 3 to show the following result (proof in the appendix § 7.1) that characterizes the behavior of a network if the test example is “far” from examples seen during training.

Theorem 2. Let $f_k(x) = \sum_j a_{kj} \sum_{i=0}^{\infty} c_i h_i(w_j^T x)$ be a one-hidden layer network with the sum of infinite series of Hermite polynomials as an activation function. Here, $k = 1, 2, \dots, K$ are the different classes. Define $w_J = \min w_j^T x$. Let the data x be mean normalized. If $\epsilon > 0$, the Hermite coefficients $c_i = (-1)^i$ and

$$\|x\| \geq \frac{1}{\|w_J\|} \log_{\log(1+K\epsilon)} \frac{\alpha}{\epsilon}$$

then, we have that the predictions are approximately (uniformly) random. That is,

$$\frac{1}{K} - \epsilon \leq \frac{e^{f_k(x)}}{\sum_{l=1}^K e^{f_l(x)}} \leq \frac{1}{K} + \epsilon \quad \forall k \in \{1, 2, \dots, K\}$$

Proof. Sketch. Note that the form of coefficients is important for this theorem. In particular, we use the exponential generating functions expansion of Hermites and exploit the form of normalization due to the softmax layer to provide a lower bound for the confidence of prediction for an arbitrary class. For this event to occur with high probability, we show that the test data has to be at least a certain distance far away from training data. $\square$

As the data is mean normalized, any test example x with high $\|x\|$ implies that it is far from the training data. For large $\|x_{\text{test}}\|$, this result shows that the predictions are fairly random – a desirable property – clearly, we should not predict confidently if we have not seen such an example earlier.

Hermite-SaaS is more noise tolerant.—To further assess if Hermite-SaaS reduces confirmation bias, we experiment by injecting noise in the labels provided for SSL training. In particular, we chose the label for a subset of images, uniformly at random. We conduct experiments with 30% label noise levels on CIFAR10. After estimating the pseudo-labels we trained a model in a supervised manner using Resnet18. Our results show that Hermite-SaaS based models obtain similar or a higher test set accuracy of about 80%. This is encouraging, but we also observe faster convergence (95 epochs) compared to a ReLU-SaaS model (at least 600 epochs). In other words, Hermite activations based training yields models/estimators with low variance suggesting that they may behave well in the presence of outliers. We also investigate how Hermite activations behave in presence of a standard robust

learning method with noisy labels, specifically [29]. We observe that the Hermite version boosts the performance of [29] both in terms of accuracy and rate of convergence (see appendix § 7.7).

5. Why HermiteS provide faster convergence?

We discussed how the noise tolerance properties of HermiteS help with faster convergence of Hermite-SaaS. Here, we show how incorporating Hermite activations within a network makes the loss landscape smoother relative to ReLUs. Smoother landscapes implying faster convergence is not controversial [7]. One difference between ReLUs and HermiteS is the nonsmooth behavior: for ReLU networks, standard first order methods require $\mathcal{O}(1/\epsilon^2)$ (versus $\mathcal{O}(1/\epsilon)$ iterations for Hermite nets) to find a local minima. This is the reason why it is not enough to just replace HermiteS with (a sufficiently large number of) ReLUs even though ReLU networks are universal approximators. We provide three pieces of empirical evidences to support our claim that HermiteS provide smoother landscape.

(a) Lower Loss Values.

We examine the loss landscape, along the SGD trajectory, following the directions outlined in [25]. The authors there showed that BN generates smoother landscapes: in Fig. 8, we show that BN+HermiteS generate even smoother landscapes implying much faster training. In Fig. 8 (left), we plot training loss $L(w - \eta \nabla L)$ for different η values for ResNet18 architecture. Hermite network generates a better loss landscape (lower magnitude) than ReLU network. **(b) Maximum Beta smoothness.** In Fig. 8 (middle), we show the maximum difference in the ℓ_2 norm of gradients over the distance moved in that direction. Hermite networks have a lower variation of gradient norm change than ReLU networks, indicating that the gradients on Hermite Landscape are more stable implying faster convergence. **(c) Higher effective learning rate.** In Fig. 8 (right), we plot deviation of weights from initialization and observe an increase in this metric with HermiteS.

6. Conclusion

In this paper, we studied the viability and potential benefits of using a finite Hermite polynomial bases as activation functions, as a substitute for ReLUs. The lower order Hermite polynomials have nice mathematical properties from the optimization point of view, although little is known in terms of their practical applicability to networks with more than a few layers. We observed from our experiments that simply replacing ReLU with an expansion in terms of Hermite polynomials can yield significant computational benefits, and we demonstrate the utility of this idea in a computationally intensive semi-supervised learning task. Under the assumption that the training is being performed on the cloud (published pricing structure), we show sizable financial savings are possible. On the mathematical side, we also showed that Hermite based networks have nice noise stability properties that appears to be an interesting topic to investigate, from the robustness or adversarial angles. Furthermore, since Hermite-nets avoid over-confident predictions on newer test samples, it would be interesting to investigate the benefits of using Hermite-nets to solve (variants of) meta learning problems.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Research supported by NIH R01 AG062336, NSF CAREER award RI#1252725 and American Family Insurance. We thank one of the reviewers for pointing out a promising connection to meta-learning that will be pursued in follow-up work. We are grateful to Glenn Fung for discussions and pointing out a nice paper on smooth SVMs by Yuh-Jye Lee and Olvi Mangasarian [18] which dealt with smoothing the “plus” function. Rest in peace, Olvi.

7.: Appendix

7.1. Proof of Theorems

Assume that the data is mean normalized (mean 0 and identity covariance matrix). In order to prove our main result, we first prove a generic perturbation bound in the following lemma. We analyze hermite polynomials on a simpler setting, a one-hidden-layer network, to get an intuitive feeling of our argument. Consider a one-hidden-layer network with an input layer, a hidden layer, and an output layer, each with multiple units. Let x denote the input vector, and $f_k(x)$ and $f_l(x)$ to be two output units. Denote w_j to be the weight vector between input and the j^{th} hidden unit, and a_{lk} to be the weight connecting l^{th} hidden unit to the k^{th} output unit. The network can be visualized as in Figure 9.

In the lemma 3, we bound the distance between two output units $f_k(x)$ and $f_l(x)$. This gives us an upperbound on how any two output units differ as a function of the norm of the (input) data, $\|x\|$.

Lemma 3. *Consider the output unit of the network, $f_k(x) = \sum_j a_{kj} \sum_{i=0}^d c_i h_i(w_j^T x)$ where c_i 's are the Hermite coefficients and d is the maximum degree of the hermite polynomial considered. Then,*

$$\|f_l(x) - f_k(x)\| \leq C d \alpha \beta$$

Here, C is a constant that depends on the coefficients of the Hermite polynomials and

$$\alpha = \max_{l,k} \sum_j |a_{lj} - a_{kj}|; \beta = \max \left(\|w\|_p^d \|x\|_q^d, \|w\|_p \|x\|_q \right),$$

such that $1/p + 1/q = 1$.

Proof.

$$f_l(x) = \sum_j a_{lj} \left(\sum_i c_i h_i(w_j^T x) \right) \text{ and}$$

$$f_k(x) = \sum_j a_{kj} \left(\sum_i c_i h_i(w_j^T x) \right)$$

$$|f_l(x) - f_k(x)| = \left| \sum_j (a_{lj} - a_{kj}) \left(\sum_i c_i h_i(w_j^T x) \right) \right| \quad (4)$$

$$\leq \sum_j |a_{lj} - a_{kj}| \left\| \sum_i c_i h_i(w_j^T x) \right\| \quad (5)$$

h_i is the normalized hermite polynomial $\left(h_i = \frac{H_i}{\sqrt{i!}} \right)$. $|h_i(z)| = \left| \frac{H_i(z)}{\sqrt{i!}} \right| \leq C_1 \frac{|z|^i}{\sqrt{i!}}$, here, C_1 is a function of the coefficients of a given hermite polynomial.

Let's bound $\left| \left(\sum_i c_i h_i(w_j^T x) \right) \right|$,

$$\begin{aligned} \left\| \sum_i c_i h_i(w_j^T x) \right\| &\leq \|c_i\|_p \|h_i\|_q \text{ where } \frac{1}{p} + \frac{1}{q} = 1 \\ &\leq \left[\sum_i c_i^p \right]^{1/p} \left[\sum_i h_i^q \right]^{1/q} \\ &\leq d^{1/p} C_2 \left[\sum_{i=0}^d \frac{C_1 |w_j^T x|^{iq}}{\sqrt{i!}^q} \right]^{1/q} \quad \forall i, c_i \leq C_2 \\ &\leq d^{1/p} C_2 \left[\sum_{i=0}^d C_1 |w_j^T x|^{iq} \right]^{1/q} \end{aligned}$$

Replace C_1, C_2 with C

$$\begin{aligned} &\leq d^{1/p} C d^{1/q} \max(|w_j^T x|, |w_j^T x|^d) \\ &\leq C d \max(|w_j^T x|, |w_j^T x|^d) \end{aligned}$$

Using, $J = \operatorname{argmax} \|w_j\|$ and $1/p + 1/q = 1$

$$\leq C d \max(\|w_J\|_{1/p} \|x\|_{1/q}, \|w_J\|_{1/p}^d \|x\|_{1/q}^d)$$

Substituting the above result back into Equation 5, we get,

$$\begin{aligned}
& |f_l(x) - f_k(x)| \\
& \leq \left| \sum_j (a_{lj} - a_{kj}) \left(\sum_i c_i h_i(w_j^T x) \right) \right| \\
& \leq Cd \max(\|w_J\| \|x\|, \|w_J\|^d \|x\|^d) \max_{kl} \sum_j |a_{lj} - a_{kj}| \\
& = Cda\beta
\end{aligned}$$

Here, $\alpha = \max_{kl} \sum_j |a_{lj} - a_{kj}|$ and $\beta = \max(\|w_J\| \|x\|, \|w_J\|^d \|x\|^d)$ $\square$

The above lemma can be seen as a perturbation bound since we can simply interpret $f_k = f_l + g_l$ for some other function, (hopefully with nicer properties) g_l . Our main theorem is simply an application of the above lemma to a special case. This allows us to quantify the noise resilience of hermite polynomials. In particular, we show that if the test data is far from the train data, then hermite polynomials trained deep networks give low confidence predictions. This property allows us to detect outliers during inference, especially in mission critical applications since it allows a human in a loop system [32].

Theorem 4. Let $f_k(x) = \sum_j a_{kj} \sum_{i=0}^{\infty} c_i h_i(w_j^T x)$, be a one-hidden-layer network with the sum of infinite series of hermite polynomials as an activation function. Here, $k = 1, 2, \dots, K$ are the different classes. Define $w_J = \min w_j^T x$. Let the data x be mean normalized. If $\epsilon > 0$, the Hermite coefficients $c_i = (-1)^i$ and

$$\|x\| \geq \frac{1}{\|w_J\|} \log \frac{\alpha}{\log(1 + K\epsilon)}$$

then, we have that the predictions are approximately (uniformly) random. That is,

$$\frac{1}{K} - \epsilon \leq \frac{e^{f_k(x)}}{\sum_{l=1}^K e^{f_l(x)}} \leq \frac{1}{K} + \epsilon \quad \forall k \in \{1, 2, \dots, K\}$$

Proof. Observe that,

$$\frac{e^{f_k(x)}}{\sum_{l=1}^K e^{f_l(x)}} = \frac{1}{\sum_{l=1}^K e^{f_l(x) - f_l(x)}}$$

Using the fact that $|x| \leq x$ and $-|x| \leq x$, we observe,

$$\frac{1}{\sum_{l=1}^K e^{|f_l(x) - f_k(x)|}} \leq \frac{1}{\sum_{l=1}^K e^{f_l(x) - f_k(x)}} \leq \frac{1}{\sum_{l=1}^K e^{-|f_l(x) - f_k(x)|}} \quad (6)$$

Let's bound $|f_l(x) - f_k(x)|$,

$$\begin{aligned}
& |f_l(x) - f_k(x)| \\
&= \left| \sum_j a_{kj} \left(\sum_i c_i h_i(w_j^T x) \right) - \sum_j a_{lj} \left(\sum_i c_i h_i(w_j^T x) \right) \right| \\
&= \left| \sum_j (a_{kj} - a_{lj}) \left(\sum_i c_i h_i(w_j^T x) \right) \right| \\
&= \left| \sum_j (a_{kj} - a_{lj}) \left(\sum_i \frac{c_i}{(-1)^n} h_i(w_j^T x) (-1)^n \right) \right|
\end{aligned}$$

From the properties of hermite polynomials,

$$\begin{aligned}
e^{xt - t^2/2} &= \sum_{i=0}^{\infty} h_i(x) t^i \\
e^{-x - 1/2} &= \sum_{i=0}^{\infty} h_i(x) (-1)^i, \text{ when } t = -1
\end{aligned} \tag{7}$$

Choose $c_i = (-1)^i$, then

$$\max_j \sum_i (c_i)^i h_i(w_j^T x) (-1)^i = \max_j e^{-w_j^T x - 1/2} \leq e^{-\min_j w_j^T x}$$

Thus,

$$\begin{aligned}
& |f_l(x) - f_k(x)| \\
&= \left| \sum_j (a_{kj} - a_{lj}) \left(\sum_i h_i(w_j^T x) (-1)^i \right) \right| \\
&\leq e^{-\min_j w_j^T x} \left| \sum_j (a_{kj} - a_{lj}) \right| \\
&\leq e^{-\min_j w_j^T x} \alpha \text{ where } \alpha = \max_{kl} \left| \sum_j (a_{kj} - a_{lj}) \right|
\end{aligned}$$

Let $e^{-\min_j w_j^T x} \alpha \leq \log(1 + K\epsilon)$, then,

$$\begin{aligned}
&\Rightarrow \frac{1}{\sum_{l=1}^K e^{|f_l(x) - f_k(x)|}} \leq \frac{1}{\sum_{l=1}^K e^{f_l(x) - f_k(x)}} \leq \frac{1}{\sum_{l=1}^K e^{-|f_l(x) - f_k(x)|}} \\
&\Rightarrow \frac{1}{K(1 + K\epsilon)} \leq \frac{1}{\sum_{l=1}^K e^{f_l(x) - f_k(x)}} \leq \frac{1 + K\epsilon}{K} \\
&\Rightarrow \frac{1}{K} - \epsilon \leq \frac{1}{\sum_{l=1}^K e^{f_l(x) - f_k(x)}} \leq \frac{1}{K} + \epsilon
\end{aligned}$$

Let $J = \operatorname{argmin}_j (w_j^T x)$. The condition,

$$\begin{aligned}
e^{-w_J^T x} \alpha &\leq \log(1 + K\epsilon), \\
\Rightarrow w_J^T x &\geq \log \frac{\alpha}{\log(1 + K\epsilon)} \\
\Rightarrow \|w_J\| \|x\| &\geq w_J^T x \geq \log \frac{\alpha}{\log(1 + K\epsilon)} \\
\Rightarrow \|x\| &\geq \frac{1}{\|w_J\|} \log \frac{\alpha}{\log(1 + K\epsilon)}
\end{aligned}$$

□

We interpret the above theorem in the following way: whenever $\|x_{test}\|$ is large, the infinity norm of the function/prediction is approximately $1/K$ where K is the number of classes. This means that the predictions are the same as that of random chance – low confidence. The only caveat of the above theorem is that it requires the all of (infinite) hermite polynomials for the result to hold. However, all of our experiments indicate that such a property holds true empirically, as well.

7.2. Computational Benefits on AWS p2.xlarge

In the paper, we have seen the cost benefits of using Hermite Polynomials on AWS p3.16xlarge instance. In this section, we will examine the performance on AWS p2.xlarge instance as indicated in Table 5. Clearly, as AWS p2.xlarge costs less, the benefits achieved when using hermite polynomials is more significant in the compute time.

7.3. Higher number of active units with Hermite Polynomials

There is consensus that ReLUs suffer from the “dying ReLU” problem. Even ignoring pathological cases, it is not uncommon to find that as much as 40% of the activations could “die”. A direct consequence of this behavior is that the dead neurons will stop responding to variations in error/input. This makes ReLUs not very suitable for Recurrent Network based architectures such as LSTMs [17]. On the other hand [27] shows that having a certain number of active units, the number determined based on the network architecture is a necessary condition for successful training. We demonstrate here that we meet the necessary condition with a large margin.

Experiment—We investigated the number of active units present in a Hermite Network and a ReLU network at the start and the end of training using preactResNet18 model. We determined that the percentage of active units in Hermes are 100%/99.9% (EPOCH 0 / EPOCH 200) whereas in ReLU are 51%/45% (EPOCH 0 / EPOCH 200). We also plot the number of active units at the end of training across the layers in Figure 10. It can be seen in Figure 10 that Hermes have twice the number of Active Units than ReLU.

7.4. Early Riser Property Preserved: ResNets

Setup—We used CIFAR10 data in all the experiments in here with random crop and random horizontal flip for data augmentation schemes. With a batch size of 128 for training and *SGD* as an optimizer, we trained ResNets for 200 epochs. We began with some initial

learning rate and then reduce it by 10 at 82nd and 123rd epoch – this is standard practice. We repeated every experiment 4 times in order to control the sources of variance.

Experiment—The experimental results for ResNet18 and ResNet50 models can be found in Figure 11, Figure 11a and Figure 11b respectively. The results for ResNet150 can be found in Table 6. We observe that the early riser property is preserved across different learning rates for these ResNet models. There is a small increase in the number of parameters proportional to the number of layers in the network.

7.5. Early Riser Property Preserved: DenseNets

Setup—All the experiments here were conducted on CIFAR10 dataset with random crop and random horizontal flip schemes for data augmentation. We used the basic block architecture [13] with 40 layers, a growth rate of 12, and a compression rate of 1.0 in the transition stage.

Experiment—Figure 12 shows the experimental results of using Hermite Polynomials in a DenseNet. We observe that the early riser property is preserved only in the training loss and the training accuracy curves.

7.6. SSL Results on CIFAR10–1K

In the paper, we report the cost benefits on CIFAR10 dataset with 4000 labelled items. Here, we report benefits on Hermite Polynomials on CIFAR10 dataset with 1000 labelled item, another popular experimental setting for CIFAR10 dataset. We find that Hermite-SaaS works equally well for CIFAR10–1K. We achieve a max pseudolabel accuracy of 77.4% compared to 75.6% in ReLU-SaaS. Hermite-SaaS saves cost > \$280 on AWS p3.16xlarge and compute time by 13+ hours on AWS p2.xlarge, compared to ReLU-SaaS when run with (20/135) inner/outer epoch config.

7.7. Noise injection experiments on Hermites

Setup—We repeat the SaaS experiments by injecting 10% and 30% label noise to CIFAR10 dataset. We utilize the method proposed by [29], let's call it *Tanaka et al.*, as an optional post-processing step after obtaining pseudo-labels via SaaS. Basically, *Tanaka et al.* performs a “correction” to minimize the influence of noise once the pseudolabels have been estimated. The authors in [29] propose an alternating optimization algorithm where the weights of the network and the labels of the data are updated alternatively. We conduct experiments with 10% and 30% label noise levels on the CIFAR10 dataset. After estimating the pseudolabels and/or using the scheme in [29], we trained a model in a supervised manner using Resnet18.

Experiment—Our results summarized in Table 7 show that Hermite-SaaS based models obtain a similar or higher test set accuracy. We also observe that our model converges faster compared to a ReLU-SaaS model. Our experimental results also indicate that post processing techniques (such as [29]) may not always be useful to improve the generalization performance of models.

7.8. Experiments on Shallow Nets

Setup—We used a 3-layer network where the hidden layers have 256 nodes each. We used CIFAR10 dataset for this experiment with normalized pixel values and no other data augmentation schemes. The architecture we worked with is thus [3072, 256, 256, 10]. We ran two experiments, one with ReLU activation function and the other with Hermite activations ($d=5$ polynomials). The loss function was cross-entropy, SGD was the optimization algorithm and we added batch normalization. Learning rate was chosen as 0.1 and the batch size as 128. Figure 13 shows the loss function and the training accuracy curves that we obtained. It was observed that, when ReLU is just replaced with Hermite activations while maintaining everything else the same, the loss function for hermites converges at a faster rate than ReLU atleast for the earlier epochs.

7.9. Cloud Services Details

We utilize AWS EC2 On-Demand instances in our research. Specifically, we use p3.16xlarge instance which costs \$24.48 per hour and p2.xlarge instance which costs \$0.9 per hour.

References

- [1]. Bergstra James, Desjardins Guillaume, Lamblin Pascal, and Bengio Yoshua. Quadratic polynomials learn better image features. Technical report, 1337, 2009.
- [2]. Boyd John P. Asymptotic coefficients of hermite function series. *Journal of Computational Physics*, 54(3):382–410, 1984.
- [3]. Chollet François et al. Keras. <https://keras.io>, 2015.
- [4]. Choromanska Anna, Henaff Mikael, Mathieu Michael, Arous Gérard Ben, and LeCun Yann. The loss surfaces of multilayer networks. In *Artificial Intelligence and Statistics*, pages 192–204, 2015.
- [5]. Cicek Safa, Fawzi Alhussein, and Soatto Stefano. Saas: Speed as a supervisor for semi-supervised learning. In *The European Conference on Computer Vision (ECCV)*, 9 2018.
- [6]. Ge Rong, Lee Jason D, and Ma Tengyu. Learning one-hidden-layer neural networks with landscape design. *arXiv preprint arXiv:1711.00501*, 2017.
- [7]. Ghadimi Saeed and Lan Guanghui. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [8]. Glorot Xavier, Bordes Antoine, and Bengio Yoshua. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323, 2011.
- [9]. Hardt Moritz and Ma Tengyu. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.
- [10]. He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11]. He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian. Identity mappings in deep residual networks In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [12]. Hein Matthias, Andriushchenko Maksym, and Bitter-wolf Julian. Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem. *arXiv preprint arXiv:1812.05720*, 2018.
- [13]. Huang Gao, Liu Zhuang, Van Der Maaten Laurens, and Weinberger Kilian Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

- [14]. Iscen Ahmet, Tolias Giorgos, Avrithis Yannis, and Chum Ondrej. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5070–5079, 2019.
- [15]. Kileel Joe, Trager Matthew, and Bruna Joan. On the expressive power of deep polynomial neural networks. *arXiv preprint arXiv:1905.12207*, 2019.
- [16]. Laine Samuli and Aila Timo. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [17]. Quoc V Le Navdeep Jaitly, and Hinton Geoffrey E. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.
- [18]. Lee Yuh-Jye and Mangasarian Olvi L. Ssvm: A smooth support vector machine for classification. *Computational optimization and Applications*, 20(1):5–22, 2001.
- [19]. Miyato Takeru, Maeda Shin-ichi, Koyama Masanori, and Ishii Shin. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. [PubMed: 30040630]
- [20]. Mossel Elchanan, O'Donnell Ryan, and Oleszkiewicz Krzysztof. Noise stability of functions with low influences: invariance and optimality. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, 2005.
- [21]. Nguyen Anh, Yosinski Jason, and Clune Jeff. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015.
- [22]. Poggio Tomaso and Liao Qianli. Theory II: Landscape of the empirical risk in deep learning. PhD thesis, Center for Brains, Minds and Machines (CBMM), *arXiv*, 2017.
- [23]. Rasiah AI, Togneri R, and Attikiouzel Y. Modelling 1-d signals using hermite basis functions. *IEE Proceedings-Vision, Image and Signal Processing*, 144(6):345–354, 1997.
- [24]. Rudin Walter. *Real and complex analysis*. Tata McGraw-Hill Education, 2006.
- [25]. Santurkar Shibani, Tsipras Dimitris, Ilyas Andrew, and Madry Aleksander. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pages 2483–2493, 2018.
- [26]. Shi Weiwei, Gong Yihong, Ding Chris, Ma Zhiheng Tao Xiaoyu, and Zheng Nanning. Transductive semi-supervised deep learning using min-max features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 299–315, 2018.
- [27]. Shin Yeonjong and Karniadakis George Em. Trainability and data-dependent initialization of over-parameterized relu neural networks. *arXiv preprint arXiv:1907.09696*, 2019.
- [28]. Soltanolkotabi Mahdi, Javanmard Adel, and Lee Jason D. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 65(2):742–769, 2018.
- [29]. Tanaka Daiki, Ikami Daiki, Yamasaki Toshihiko, and Aizawa Kiyoharu. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5552–5560, 2018.
- [30]. Tarvainen Antti and Valpola Harri. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017.
- [31]. Venturi Luca, Bandeira Afonso S, and Bruna Joan. Spurious valleys in two-layer neural network optimization landscapes. *arXiv preprint arXiv:1802.06384*, 2018.
- [32]. Yu Fisher, Seff Ari, Zhang Yinda, Song Shuran, Funkhouser Thomas, and Xiao Jianxiong. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [33]. Zaheer Manzil, Reddi Sashank, Sachan Devendra, Kale Satyen, and Kumar Sanjiv. Adaptive methods for nonconvex optimization. In *Advances in Neural Information Processing Systems*, pages 9793–9803, 2018.
- [34]. Zhang Chiyuan, Bengio Samy, Hardt Moritz, Recht Benjamin, and Vinyals Oriol. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.

Figure 1: Hermite Polynomials as Activations

(leftmost): Incorporating Hermite Polynomials as an activation function in a single hidden unit one hidden layer network. **(middle):** The functional form of the first 5 hermites are shown in the right.

Figure 2: Other Activation Functions.

We observe an increase in the rate of convergence of the validation loss for Hermite when compared to other activations such as ELU, ReLU and SeLU. A standard auto-encoder model as outlined in [33] is used to test on MNIST dataset.

Figure 3: Hermite Polynomials as Activations in ResNets.

We introduce softsign function to handle the numerical issues from the unbounded nature of Hermite polynomials. W denotes the weight, BN denotes batch normalization, σ is the Hermite activation and SS is the softsign function.

Figure 4: Hermites vs. ReLUs on ResNet18.

(a) Hermites provide faster convergence of train loss and train accuracies than ReLUs.

Hermes have faster convergence in test accuracies over the initial epochs but ReLU has the higher test accuracy at the end of training. (b) As we increase the number of hermite polynomials, the speed of loss convergence increases until $d=6$ and then it starts to reduce.

$d=1$ performs better than $d=0$ where only softsign is used as an activation. The close overlap between $d=4$ to $d=10$ suggests that tuning for d is not very expensive.

Figure 5: SaaS framework illustration.

SaaS runs over two loops: inner loop denoted by I and outer loop denoted by O .

Figure 6: Hermite-SaaS trains faster.

We plot the number of outer epochs M_O vs. the pseudo-label accuracy across 4 datasets. We consistently observe that the minimum number of outer epochs M_O to reach a given value of pseudo-label accuracy is always lower for Hermite-SaaS than ReLU-SaaS.

Figure 7: Convergence of loss functions in SaaS.

Larger spikes correspond to the end of inner loop and are due to weight reinitialization.

Right plot shows that Hermite's accelerate the training, ensuring high-quality pseudo-labels.

Figure 8: Hermite generates smoother landscape than ReLU.

(left): Magnitude of loss and its variation is lower for Hermite. (middle): Gradients are more stable on Hermite loss landscape: lower maximum beta smoothness. (right): Hermite networks have a higher effective learning rate.

Figure 9:
One Hidden Layer Network

Figure 10: Hermite networks have more active units.

This figure displays the number of active units present in different layers of ResNet18 network at the end of training. It can be seen that Hermite Nets have close to 100% active units while in ReLU nets half the number of neurons are dead at the end of training.

Figure 11: Early riser property preserved in ResNets

(a) ResNet18 and (b) ResNet50. We observe that a test accuracy of 90% is achieved in approximately half the number of epochs in Hermite-ResNets over ReLU-ResNets on CIFAR10 dataset over different learning rates. The closest that Hermite gets to ReLU is the case when both their learning rates are 0.01 in ResNet50. In this case, we observe that 90% testset accuracy achieved in 20 epochs for Hermite-ResNet50 and 30 epochs for ReLU-ResNet50.

Figure 12: Early Riser Property in DenseNets.

We observe that the training loss and the training accuracy converge at a faster rate when using Hermite Polynomials in DenseNets.

Figure 13: Experiments on Shallow Nets.

The training loss and training accuracy reduce at a faster rate with Hermite-Polynomials

Table 1:
Deep autoencoders with Hermite Activations give lower test loss

Results from our implementation following directions from [33].

Method	LR	ϵ	Train Loss	Test Loss
Sigmoid-Adam [33]	10^{-3}	10^{-8}	2.97 ± 0.06	7.91 ± 0.14
Sigmoid-Adam [33]	10^{-3}	10^{-3}	1.90 ± 0.08	4.42 ± 0.29
Hermite-Adam (Ours)	10^{-3}	10^{-8}	1.89 ± 0.01	3.16 ± 0.02

Table 2:**SSL experimental details.**

A and N denote the data augmentation techniques, affine transformation and normalization respectively. M_I and M_O denote the inner and outer epochs respectively.

Dataset	# Labeled.	# Unlabeled.	Augmnt.	M_I / M_O
SVHN	1K	72,257	A + N	5 / 75
CIFAR-10	4K	46,000	A + N	10 / 135
SmallNORB	1K	23,300	None	10 / 135
MNIST	1K	59,000	N	1 / 75

Table 3:
Cost effective and accurate pseudo-labels are generated by Hermite-SaaS.

Expenses when training Hermite-SaaS (H) and ReLU-SaaS (R) on AWS and the performance metrics on pseudo-labels generated. We observe that although Hermite-SaaS takes more time per epoch than ReLU-SaaS, the overall gains are better for Hermite-SaaS.

		Max	Max PL ACC	Time/Epoch (sec)	Total Time (hrs)	Cost (\$)	Saved (\$)
SVHN	H	6.1%	94.2%	470	5.6	137	71
	R		93.3%	409	8.5	208	
CIFAR10	H	3.4%	85.5%	348	2.7	66	213
	R		84.4%	304	11	280	
Small NORB	H	5.2%	92.6%	47	0.5	12	13
	R		90.4%	27	1	25	
MNIST	H	4.5%	98.2%	94	1.2	29	-11
	R		98.2%	55	0.3	18	

Table 4:
Hermite-SaaS generalizes better.

(left): Test-set accuracies on SVHN dataset in comparison to the baselines provided in [5]. **(right):** Test-set accuracies on CIFAR10 dataset in comparison to other pseudo-label based SSL methods.

Metdod	SVHN	Metdod	CIFAR10			
	1000	(SSL with PL)	500	1000	2000	4000
VAT+EntMin [19]	3.86	TSSDL [26]	-	21.13	14.65	10.90
MT [30]	3.95					
TSSDL [26]	3.80	Label Prop. [16]	32.4	22.02	15.66	12.69
TE [16]	4.42					
ReLU-SaaS	3.82	ReLU-SaaS	-	-	-	10.94
Hermite-SaaS	3.57 ± 0.04	Hermite-SaaS	29.25	20.77	14.28	10.65

Table 5:
Hermite-SaaS on AWSp2.xlarge instance.

Hermite-SaaS saves compute time when compared to ReLU-SaaS

		Time per Epoch (sec)	Total Time (hours)	Cost (\$)	Time Savings (hours)
SVHN	Hermite	666.8	2.22	2.0	2.6
	ReLU	435.3	4.84	4.35	
CIFAR-10	Hermite	454.2	3.53	3.18	8.5
	ReLU	320.7	12	10.8	
SmallNORB	Hermite	164.8	1.74	1.57	1.33
	ReLU	81.8	3.07	2.76	
MNIST	Hermite	345.4	4.41	3.97	-2
	ReLU	180	2.4	2.16	

Table 6:
Hermite Polynomials in ResNet152.

We observe small increase in the number of parameters. Test accuracy for the hermite model converges in less than half the number of epochs.

Dataset	CIFAR10	Number of Trainable Parameters	Best Test Accuracy	Epochs to reach 90% Test Accuracy
Hermite		58,145,574	95.48%	30
ReLU		58,144,842	94.5%	80

Table 7:
SaaS experiments on Noisy Labelled dataset.

We report the best accuracy (A_{Best}), and number of epochs (N_{Best}) to reach this accuracy. *Tanaka* et al. stands for noisy label processing method proposed in [29].

10% Noise		A_{Best}	N_{Best}
SaaS	Hermite	85	224
	ReLU	84	328
SaaS + <i>Tanaka</i> et al.	Hermite	85	244
	ReLU	84	284
30% Noise		A_{Best}	N_{Best}
SaaS	Hermite	~80	95
	ReLU	~80	600
SaaS + <i>Tanaka</i> et al.	Hermite	~80	61
	ReLU	~80	299