
Learning with Hyperspherical Uniformity

Weiyang Liu^{1,2,*} Rongmei Lin^{3,*} Zhen Liu^{4,*} Li Xiong³ Bernhard Schölkopf² Adrian Weller^{1,5}

¹University of Cambridge ²MPI-IS Tübingen ³Emory University ⁴Université de Montréal ⁵Alan Turing Institute

Abstract

Due to the over-parameterization nature, neural networks are a powerful tool for nonlinear function approximation. In order to achieve good generalization on unseen data, a suitable inductive bias is of great importance for neural networks. One of the most straightforward ways is to regularize the neural network with some additional objectives. ℓ_2 regularization serves as a standard regularization for neural networks. Despite its popularity, it essentially regularizes one dimension of the individual neuron, which is not strong enough to control the capacity of highly over-parameterized neural networks. Motivated by this, hyperspherical uniformity is proposed as a novel family of relational regularizations that impact the interaction among neurons. We consider several geometrically distinct ways to achieve hyperspherical uniformity. The effectiveness of hyperspherical uniformity is justified by theoretical insights and empirical evaluations.

1 Introduction

As one of the most effective ways to control the capacity of over-parameterized neural networks, regularization serves an important role to prevent overfitting and improve generalization. Current prevailing weight regularizations can be divided into two major categories: *Individual regularization* and *relational regularization*. Taking n d -dimensional neurons $\mathbf{w}_1, \dots, \mathbf{w}_n \in \mathbb{R}^d$ from one layer of the neural network as an example, we typically can have the following regularization:

$$\mathcal{L}_{\text{reg}} = \underbrace{\lambda_I \cdot \sum_{i=1}^n h(\mathbf{w}_i)}_{\text{Individual Regularization}} + \underbrace{\lambda_R \cdot g(\mathbf{w}_1, \dots, \mathbf{w}_n)}_{\text{Relational Regularization}}$$

where h defines a function that independently describes some properties of the individual weight \mathbf{w}_i (*e.g.*, ℓ_p

norm), and g is a function that characterizes the interaction between different weights $\mathbf{w}_1, \dots, \mathbf{w}_n$ (*e.g.*, orthogonality). Individual regularization (*e.g.*, weight decay) is typically used by default in neural networks, while relational regularization tends to be overlooked. Individual regularization essentially regularizes only one dimension of the weights (from the view of spherical coordinate systems [7]), which is not strong enough for highly over-parameterized neural networks. In contrast, relational regularization can be viewed as regularizing $d - 1$ dimensions of the weights, and encodes stronger inductive bias with relational information. Individual and relational regularizations usually serve complementary roles to each other and can be used simultaneously. In this paper, we study the importance of a suitable relational regularization for over-parameterized models (*e.g.*, neural networks) and explore a novel family of the relational regularizations – *hyperspherical uniformity*.

Hyperspherical uniformity characterizes the diversity of vectors on a unit hypersphere. Different from orthogonality where perpendicular vectors are defined to be diverse, hyperspherical uniformity encourages vectors to be spaced apart with as large an angle as possible such that these vectors can be uniformly distributed over the hypersphere. In order to promote hyperspherical uniformity with an explicit regularization, we formulate several distinct learning objectives that are conceptually appealing and geometrically interpretable. Specifically, we consider minimum hyperspherical energy (MHE) [46], maximum hyperspherical separation (MHS), maximum hyperspherical polarization (MHP), minimum hyperspherical covering (MHC), and maximum Gram determinant (MGD). Different regularization objectives yield distinct geometric interpretations and optimization dynamics. Moreover, we draw inspiration from statistical uniformity testing on the hypersphere and provide a novel and unified view on understanding these learning objectives.

The motivation to encourage hyperspherical uniformity lies in three aspects. First, we argue that hyperspherical uniformity leads to better optimization and generalization. [83] proves that optimizing one-hidden-layer neural networks with hyperspherically uniform

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

neurons has no spurious local minima. [46, 45] empirically show that promoting hyperspherical uniformity of neurons can effectively improve the generalization of neural networks. Interestingly, hyperspherical uniformity implicitly regularizes the neurons to be close to the initialization, partially implementing Occam’s razor to keep neural networks as simple as possible (see Section 7). Second, hyperspherical uniformity can remove neuron redundancy and encourage the neurons to be diverse on the hypersphere. In the light of [71] that shows deeply learned neurons are highly redundant, hyperspherical uniformity can serve as a useful regularization to remove such redundancy. Third, hyperspherical uniformity has a clear geometric interpretation and theoretical merits. There exists a close connection between hyperspherical uniformity and orthogonality. The effectiveness of hyperspherical uniformity can also be justified from multiple theoretical viewpoints. Our contribution can be summarized as follows:

- We introduce a general property – hyperspherical uniformity as a regularizer for neural networks.
- To achieve hyperspherical uniformity, we propose several well-performing approaches (MHC, MHS, MHP, MGD) with distinct geometric interpretations.
- We establish a connection between hyperspherical uniformity and orthogonality, showing that hyperspherical uniformity is a more general property.
- We provide some insights and discussions on the geometric and spectral properties and regularization effects of hyperspherical uniformity.
- We apply hyperspherical uniformity to a number of applications and demonstrate superior performance over existing regularizations such as orthogonality.

2 Related Work

Relational regularizations. There are quite a number of relational regularizations that have been used in neural networks, such as orthogonality regularization [4, 52, 67, 32, 15], unitary constraint [34, 81, 3], decorrelation [67, 16, 85], spectral regularization [86], low-rank regularization [76], angular constraint [84, 44], etc. Most of these relational regularizations are either directly based on orthogonality or based on some notions related to orthogonality (*e.g.*, correlation). Quite differently, hyperspherical uniformity encourages neurons to be uniformly distributed over the hypersphere.

Hyperspherical learning. [51, 52, 49, 48, 50, 80, 79, 20, 19, 60, 14, 47] propose to learn representations on hypersphere and show that angular information in neural networks, in contrast to magnitude, preserves the key semantics and is very crucial to generalization. [46] regularizes the diversity of neurons on hypersphere by minimizing their pairwise energy. [45] explores how projections can help to better minimize such energy.

3 Connection between Hyperspherical Uniformity and Orthogonality

Before we discuss specific methods to achieve hyperspherical uniformity, we first reveal an interesting connection between hyperspherical uniformity and orthogonality with the following theoretical statement:

Theorem 1. For $\epsilon, \delta > 0$ and $d \geq \max(4, \delta^{-2}\epsilon^{-1})$, every orthonormal basis $\mathbf{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_d\}$ in \mathbb{R}^d is ϵ - δ -uniform distributed on the unit hypersphere $\mathbb{S}(\mathbb{R}^d)$, *i.e.*, for every Borel set $\mathbf{A} \subseteq \mathbb{S}(\mathbb{R}^d)$, we have

$$\mathbb{P}\left(\left|\frac{\text{Card}(\mathbf{V} \cap \mathbf{R}(\mathbf{A}))}{\text{Card}(\mathbf{V})} - u(\mathbf{A})\right| \leq \delta\right) \geq 1 - \epsilon \quad (1)$$

where $\text{Card}(\cdot)$ denotes the cardinality of a set, $\mathbf{R}(\mathbf{A})$ is a random rotation of \mathbf{A} , and u denotes the uniform probability measure over $\mathbb{S}(\mathbb{R}^d)$.

Theorem 1 can be obtained from [26] and is related to concentration of measure [56] and Raz’s lemma [66]. It shows that uniformly sampling a basis from the orthogonal group in high dimensions gives an approximate uniform distribution on $\mathbb{S}(\mathbb{R}^d)$, implying that any orthogonal basis is approximately uniformly distributed over the hypersphere in high dimensions. This result bridges hyperspherical uniformity and orthogonality, implying hyperspherical uniformity is more general.

More interestingly, enforcing hyperspherical uniformity for $2d+2$ vectors in \mathbb{S}^d leads to a cross-polytope [87]. Fig. 1 presents a 3-dimensional cross-polytope in \mathbb{S}^2 (with 6 vectors). We assume that there exists a unit-vector set with $d+1$ vectors in total: $\{\mathbf{v}_1, \dots, \mathbf{v}_{d+1} \in \mathbb{S}^d\}$. Then we

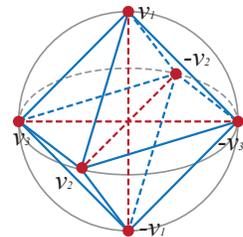


Figure 1: Cross-polytope.

construct a new set with $2d+2$ vectors by adding all vectors with opposite direction to the original set: $\{\mathbf{v}_1, \dots, \mathbf{v}_{d+1}, -\mathbf{v}_1, \dots, -\mathbf{v}_{d+1} \in \mathbb{S}^d\}$. Promoting hyperspherical uniformity for these $2d+2$ vectors (in the new set) is actually equivalent to promoting pairwise orthogonality among the original $d+1$ vectors (in the original set). This equivalence builds a strong connection between hyperspherical uniformity and orthogonality. We will discuss this further in Section 4.6.

4 Towards Hyperspherical Uniformity

We design multiple learning objectives to achieve hyperspherical uniformity and discuss their close connections. MHE was first proposed in our previous work [46], while the others are part of the contributions in this work.

4.1 Minimum Hyperspherical Energy

Inspired by the Thomson problem [78] where one seeks to find an equilibrium state with minimum potential

energy that distributes N electrons on a unit sphere as evenly as possible, MHE [46, 45] encourages hyperspherical diversity and defines the following learning objective for n d -dimensional vectors (*e.g.*, neurons) $\hat{\mathbf{W}}_n = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}$ in the same layer:

$$\min_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \{E_s(\hat{\mathbf{W}}_n) := \sum_{i=1}^n \sum_{j=1, j \neq i}^n K_s(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j)\} \quad (2)$$

where $\hat{\mathbf{w}}_i := \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|}$ is the i -th vector projected onto the unit hypersphere $\mathbb{S}^{d-1} = \{\hat{\mathbf{w}} \in \mathbb{R}^d \mid \|\hat{\mathbf{w}}\| = 1\}$. $K_s(\cdot, \cdot)$ models the interaction between two vectors, and we will typically consider the following Riesz s -kernel function:

$$K_s(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j) = \begin{cases} \rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j)^{-s}, & s > 0 \\ \log(\rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j)^{-1}), & s = 0 \\ -\rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j)^{-s}, & s < 0 \end{cases} \quad (3)$$

where $\rho(\cdot, \cdot)$ is defined to measure the geodesic similarity on the unit hypersphere. In general, we can use either $\rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j) = \|\hat{\mathbf{w}}_i - \hat{\mathbf{w}}_j\|_2$ (*i.e.*, standard Riesz s -kernel) or hyperspherical geodesic distance (*i.e.*, angular distance) $\rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j) = \arccos(\hat{\mathbf{w}}_i^\top \hat{\mathbf{w}}_j)$. Moreover, it is known that the minimizer of this pairwise energy sum asymptotically corresponds to the uniform distribution on the hypersphere [46, 41, 11]. Since the objective of MHE is non-convex and differentiable, we typically optimize it with gradient descent. Different s usually yields slightly different regularization effects and optimization dynamics. MHE has been extensively studied and shown to be effective in many applications [46].

4.2 Maximum Hyperspherical Separation

MHS adopts a maximum geodesic separation criterion by maximizing the *separation distance* (*i.e.*, the smallest pairwise distance in a vector set, denoted as $\vartheta(\hat{\mathbf{W}}_n)$) which is equal to the smallest geodesic distance between any two vectors in the set $\hat{\mathbf{W}}_n = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}$:

$$\max_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \{\vartheta(\hat{\mathbf{W}}_n) := \min_{i \neq j} \rho(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j)\} \quad (4)$$

which is formulated as a max-min optimization problem. It is originally called Tammes problem [77] (or sphere packing problem) where one packs a given number of circles on the surface of a sphere such that the minimum distance between circles can be maximized. Similar to MHE, $\rho(\cdot, \cdot)$ can be either Euclidean distance or geodesic distance on the unit hypersphere. Based on the following proposition, we can obtain that MHS is in fact a limiting case of MHE when $s \rightarrow \infty$.

Proposition 1. *Let $n \in \mathbb{N}, n \geq 2$ be fixed and (\mathbb{S}^{d-1}, ρ) be a compact metric space. Then we have that*

$$\lim_{s \rightarrow \infty} (\varepsilon_s(\mathbb{S}^{d-1}, n))^{\frac{1}{s}} = \frac{1}{\delta_n^\rho(\mathbb{S}^{d-1})} \quad (5)$$

where we define $\varepsilon_s(\mathbb{S}^{d-1}, n) := \min_{\hat{\mathbf{W}}_n \subset \mathbb{S}^{d-1}} E_s(\hat{\mathbf{W}}_n)$ and $\delta_n^\rho(\mathbb{S}^{d-1}) := \max_{\hat{\mathbf{W}}_n \subset \mathbb{S}^{d-1}} \vartheta(\hat{\mathbf{W}}_n)$.

Compared to MHE that has a global regularization effect, MHS focuses more on the local separation since it only takes the minimal geodesic distance into consideration. Specifically, MHS only updates two vectors with smallest angular distance at each iteration, while MHE updates all vectors in each iteration. When the vectors are relatively diverse on the hypersphere, MHS tends to have stronger regularization effects than MHE. This is because the MHS gradient resulted from the two closest vectors has only one direction component and will not be cancelled out. In contrast, the MHE gradient that comes from all the vectors has many direction components and may be largely cancelled out. More intuitively, MHS only activates the repulsive force from the closest vectors at a time while MHE simultaneously activates the repulsive force from all pairwise vectors.

Optimizing MHS is straightforward and efficient. We first need to use a ranking operator to rank all the pairwise distances and obtain the vectors with the smallest distance (*i.e.*, maximal similarity). Then we can simply maximize the minimal distance by updating these two closest vectors via gradient ascent.

4.3 Maximum Hyperspherical Polarization

MHP arises from a practical problem: if $K_s(\mathbf{v}, \hat{\mathbf{w}}_i)$ denotes the amount of a substance received at \mathbf{v} due to an injector of the substance located at $\hat{\mathbf{w}}_i$, what is the smallest number of injectors and their corresponding optimal locations on the hypersphere so that a prescribed minimal amount of the substance can reach every point on the hypersphere? Specifically, MHP maximizes the following s -polarization $P_s(\hat{\mathbf{W}}_n)$ of a n -vector set $\hat{\mathbf{W}}_n = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}$:

$$\max_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \{P_s(\hat{\mathbf{W}}_n) := \min_{\mathbf{v} \in \mathbb{S}^{d-1}} \sum_{i=1}^n K_s(\mathbf{v}, \hat{\mathbf{w}}_i)\} \quad (6)$$

which is a max-min problem and amounts to identifying the optimal location of “poles” for the potential function. Similar to MHE, the potential can be modeled by Riesz s -kernel (*e.g.*, $K(\mathbf{v}, \hat{\mathbf{w}}_i) = \rho(\mathbf{v}, \hat{\mathbf{w}}_i)^{-s}$). We show an intrinsic relationship between MHE and MHP:

Proposition 2. *For every $n \in \mathbb{N}, n \geq 2$ and a compact metric space (\mathbb{S}^{d-1}, ρ) , we have that*

$$\mathcal{P}_s(\mathbb{S}^{d-1}, n) \geq \frac{\varepsilon_s(\mathbb{S}^{d-1}, n+1)}{n+1} \geq \frac{\varepsilon_s(\mathbb{S}^{d-1}, n)}{n-1} \quad (7)$$

where $\mathcal{P}_s(\mathbb{S}^{d-1}, n) := \max_{\hat{\mathbf{W}}_n \subset \mathbb{S}^{d-1}} P_s(\hat{\mathbf{W}}_n)$.

Next, we show an interesting example where the maximal s -polarization problem can be easily solved.

Proposition 3. *For the case of $s = -2$ and $n \geq 2$, a vector configuration on \mathbb{S}^{d-1} : $\hat{\mathbf{W}} = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n\}$ is an optimal solution for the maximal s -polarization problem if and only if $\sum_{i=1}^n \hat{\mathbf{w}}_i = \mathbf{0}$.*

Proposition 3 inspires us to propose a relaxed MHP (R-MHP) regularization objective: $\min_{\hat{\mathbf{W}}_n} \|\sum_{i=1}^n \hat{\mathbf{w}}_i\|$. Although it only corresponds to a specific maximal (-2) -polarization solution, R-MHP can still serve as an interesting relaxed variant of the original MHP regularization for neurons. Geometrically, R-MHP can be viewed as constraining the mass center of all the points to the origin. Note that, R-MHP is a necessary condition to achieve hyperspherical uniformity, and therefore is unable to guarantee hyperspherical uniformity. Moreover, there exists a trivial solution for R-MHP to achieve the optimum, where every two vectors are paired to have the opposite directions. Typically the data fitting loss in training can prevent the neurons from falling into such a trivial solution. Alternatively we can make the number of neurons to be odd to eliminate the existence of such a trivial solution, since the neurons can no longer be paired in this case.

To optimize the standard MHP, we propose to unroll the inner minimization with a few gradient descent steps and then embed it back to the outer maximization, which is conceptually similar to [18, 23, 58, 2].

4.4 Minimum Hyperspherical Covering

MHC minimizes the following *covering radius* $\alpha(\hat{\mathbf{W}}_n)$ (also known as *mesh norm*) of a n -vector set $\hat{\mathbf{W}}_n$:

$$\min_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \left\{ \alpha(\hat{\mathbf{W}}_n) := \max_{\mathbf{v} \in \mathbb{S}^{d-1}} \min_{1 \leq i \leq n} \rho(\mathbf{v}, \hat{\mathbf{w}}_i) \right\} \quad (8)$$

where the covering radius $\alpha(\hat{\mathbf{W}}_n)$ denotes the maximum geodesic distance from a vector in \mathbb{S}^{d-1} to the nearest vector in $\hat{\mathbf{W}}_n$, and can also be viewed as the geodesic radius of the largest hyperspherical cap that contains no points from $\hat{\mathbf{W}}_n$. Covering radius has interesting applications in numerical integration on the sphere [74]. We further show that MHC is the limiting case of the maximal s -polarization (*i.e.*, MHP) as $s \rightarrow \infty$.

Proposition 4. *For a compact metric space (\mathbb{S}^{d-1}, ρ) , we can have the following equation:*

$$\lim_{s \rightarrow \infty} (\mathcal{P}_s(\mathbb{S}^{d-1}, n))^{1/s} = \frac{1}{\eta_n^\rho(\mathbb{S}^{d-1})} \quad (9)$$

where $\eta_n^\rho(\mathbb{S}^{d-1}) := \min_{\hat{\mathbf{W}}_n \subset \mathbb{S}^{d-1}} \alpha(\hat{\mathbf{W}}_n)$ is the minimal n -point covering radius in \mathbb{S}^{d-1} .

Proposition 4 shows that MHC is a special case of MHP. However, it is highly nontrivial to effectively optimize the MHC objective since it involves two inner optimizations. One simple and straightforward way is to alternatively optimize this min-max-min problem. Alternatively, we also propose to use a smooth minimum operator is to approximate and relax the most inner minimization in order to make it differentiable, such that we can then use the unrolling technique to

solve the relaxed optimization. Specifically, we can relax Eq. 8 to the following simpler problem:

$$\min_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \max_{\mathbf{v} \in \mathbb{S}^{d-1}} -\frac{1}{\gamma} \log \sum_{i=1}^n \exp(-\gamma \rho(\mathbf{v}, \hat{\mathbf{w}}_i)) \quad (10)$$

whose inner maximization can be unrolled with a few gradient ascent steps, similar to MHP. Larger γ gives better approximation to the minimum operator.

4.5 Maximum Gram Determinant

Inspired by numerical integration [74] and interpolation [35], hyperspherical uniformity can be achieved by maximizing the determinant of the kernel Gram matrix of the n -vector set $\hat{\mathbf{W}}_n = \{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}$:

$$\max_{\{\hat{\mathbf{w}}_1, \dots, \hat{\mathbf{w}}_n \in \mathbb{S}^{d-1}\}} \log \det \left(\mathbf{G} := (K(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j))_{i,j=1}^n \right) \quad (11)$$

where $\det(\mathbf{G})$ denotes the determinant of the kernel gram matrix $\mathbf{G} \in \mathbb{R}^{n \times n}$, and $K(\mathbf{u}, \mathbf{v})$ can be expanded with $\sum_{i=1}^{\infty} \psi_i(\mathbf{u})\psi_i(\mathbf{v})$. $\{\psi_i\}_{i=1}^{\infty}$ is an orthogonal basis of $\mathcal{H}_K(\mathbb{S}^{d-1})$, the reproducing kernel Hilbert space of K . By approximating the kernel with the first n terms, then we can have that $\hat{\mathbf{G}} = \Psi_n \Psi_n^\top \approx \mathbf{G}$ in which

$$\Psi_n = \begin{bmatrix} \psi_1(\hat{\mathbf{w}}_1) & \cdots & \psi_n(\hat{\mathbf{w}}_1) \\ \vdots & \ddots & \vdots \\ \psi_1(\hat{\mathbf{w}}_n) & \cdots & \psi_n(\hat{\mathbf{w}}_n) \end{bmatrix}. \quad (12)$$

Therefore, maximizing $\det(\mathbf{G})$ can be approximated by maximizing $|\det(\Psi_n)|$ which is also known as *extremal systems* [74] or *Fekete points* [55, 6, 35]. Note that, the construction of Fekete points is independent of the choice of basis. We can therefore use Gaussian kernel:

$$K(\mathbf{u}, \mathbf{v}) = \exp\left(-\sum_{i=1}^d \epsilon^2 (u_i - v_i)^2\right) \quad (13)$$

where $\epsilon > 0$ is a scale parameter. It is also possible to use some other kernel functions, but we stick to the Gaussian kernel for simplicity. Geometrically, the kernel Gram determinant is also closely related to n -dimensional volume of the parallelotope formed by $\hat{\mathbf{W}}_n$.

4.6 Theoretical Results on Optimality

Computing the optimal solutions to these optimizations is highly challenging and often infeasible [11]. Fortunately, under conditions on the dimensionality and the number of vectors, we can characterize the optimum.

Theorem 2 (Simplex optimum for $2 \leq n \leq d+2$). *Let $f : (0, 4] \rightarrow \mathbb{R}$ be a convex and decreasing function defined at $t = 0$ by $\lim_{v \rightarrow 0^+} f(v)$. If $2 \leq n \leq d+2$, then we have that the vertices of regular $(n-1)$ -simplices inscribed in \mathbb{S}^d with centers at the origin minimize the MHE objective on the hypersphere \mathbb{S}^d ($d \geq 2$) with the kernel as $K_s(\hat{\mathbf{w}}_i, \hat{\mathbf{w}}_j) = f(\|\hat{\mathbf{w}}_i - \hat{\mathbf{w}}_j\|^2)$. Moreover, f is strictly convex and strictly decreasing, then these are the only energy minimizing n -point configurations.*

Theorem 2 indicates that the vertices of a regular $(d + 1)$ -simplex (*i.e.*, $(d + 1)$ -dimensional convex hull of $d + 2$ distinct vectors with equal pairwise distances) are universally optimal. Universal optimality [17] refers to a finite subset $\mathcal{C} \subset \mathbb{S}^d$ that minimizes potential energy $\sum_{\{\mathbf{x} \neq \mathbf{y}\} \in \mathbb{S}^d} f(\|\mathbf{x} - \mathbf{y}\|^2)$ among all configurations of $|\mathcal{C}|$ points on \mathbb{S}^d for every completely monotonic potential function f . Particularly, the vertices of a regular $(d + 1)$ -simplex are the minimizer of MHE with the Riesz s -kernel (for $s \geq -2$), MHS, and MHC.

Theorem 3 (Cross-polytope optimum for $n = 2d + 2$). *The vertices of the $\frac{n}{2}$ -cross-polytope form a universally optimal n -point configuration on \mathbb{S}^d .*

Theorem 3 can be obtained from [11, 17, 38, 87] and it indicates the cross-polytope point configurations are the minimizer of MHE and MHS. A cross-polytope can be constructed by the convex hull of unit vectors pointing along each Cartesian coordinate axis. For example, the convex hull of 6 points: $(0, 0, \pm 1)$, $(0, \pm 1, 0)$, and $(\pm 1, 0, 0)$ is a 3-cross-polytope. Most importantly, this result builds a bridge that intrinsically connects hyperspherical uniformity and orthogonality. Orthogonality works in a relatively restricted setting and can be viewed as a special case of hyperspherical uniformity.

It is easy to verify that both simplex and cross-polytope form optimal point configurations in \mathbb{S}^1 and \mathbb{S}^2 . Theorem 2 and Theorem 3 essentially show that such a conclusion generalizes intuitively to high dimensional space. Moreover, [17] shows that there is no universally optimal point set on \mathbb{S}^d that contains more points than the regular simplex and fewer points than the cross-polytope (*i.e.*, $d + 2 < n < 2d + 2$). For $n > 2d + 2$, characterizing optimum becomes much more involved [11].

5 A Statistical Perspective from Uniformity Testing on Hypersphere

This section aims to gain more interesting insights by casting an alternative view on hyperspherical uniformity from a statistical uniformity testing perspective. Since we aim to promote hyperspherical uniformity, we can draw inspirations from the formulation of the test statistic in uniformity testing. Given *i.i.d.* samples $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{S}^{d-1}$ of a unit random vector \mathbf{u} , the assessment of the presence of uniformity on the hypersphere is formalized as the testing of the null hypothesis $\mathcal{H}_0 : \mathbf{P} = \text{Uniform}(\mathbb{S}^{d-1})$ against $\mathcal{H}_1 : \mathbf{P} \neq \text{Uniform}(\mathbb{S}^{d-1})$, where \mathbf{P} denotes the probability distribution of \mathbf{u} and $\text{Uniform}(\mathbb{S}^{d-1})$ denotes the uniform distribution on \mathbb{S}^{d-1} .

The intuitive idea of the Sobolev test is to map the hypersphere \mathbb{S}^{d-1} into the Hilbert space $L^2(\mathbb{S}^{d-1}, \mu)$ of square-integrable functions on \mathbb{S}^{d-1} by a function $t : \mathbb{S}^{d-1} \rightarrow L^2(\mathbb{S}^{d-1}, \mu)$ such that, if \mathbf{u} is uniformly

distributed on the hypersphere, then the mean of $t(\mathbf{u})$ will be zero. Specifically, we denote ϵ_k ($p_{d,k} = \dim \epsilon_k$) as the space of eigenfunctions $\mathbb{S}^{d-1} \rightarrow \mathbb{R}$ corresponding to the k -th non-zero eigenvalue of the Laplacian, there exists a well-defined mapping $t_k : \mathbb{S}^{d-1} \rightarrow \epsilon_k$ which can be written as $t_k(\mathbf{u}) := \sum_{i=1}^{p_{d,k}} g_{i,k}(\mathbf{u}) g_{i,k}$ ($\{g_{i,k}\}_{i=1}^{p_{d,k}}$ constructs an orthonormal basis of ϵ_k). We let $\{v_k\}_{k=1}^\infty$ be a sequence such that $\sum_{k=1}^\infty v_k^2 p_{d,k} < \infty$, and then the function $\mathbf{u} \rightarrow t(\mathbf{u}) := \sum_{k=1}^\infty v_k t_k(\mathbf{u})$ is a mapping from \mathbb{S}^{d-1} to the Hilbert space $L^2(\mathbb{S}^{d-1}, \mu)$ of square-integrable real functions on \mathbb{S}^{d-1} *w.r.t.* μ which is the uniform measure on \mathbb{S}^{d-1} . The Sobolev test [5, 24] rejects \mathcal{H}_0 for large values of the following test statistic:

$$S_n := \frac{1}{n} \sum_{i,j} \sum_{k=1}^\infty v_k^2 \langle t_k(\mathbf{u}_i), t_k(\mathbf{u}_j) \rangle \quad (14)$$

where $\langle f, g \rangle := \int_{\mathbb{S}^{d-1}} f(\mathbf{u}) g(\mathbf{u}) d\mu(\mathbf{u})$ denotes the inner product on $L^2(\mathbb{S}^{d-1}, \mu)$. [61] gives an explicit form for $\langle t_k(\mathbf{u}_i), t_k(\mathbf{u}_j) \rangle$ in \mathbb{S}^{d-1} (assuming $d > 2$):

$$\langle t_k(\mathbf{u}_i), t_k(\mathbf{u}_j) \rangle = \left(1 + \frac{2k}{d-2}\right) C_k^{(d-2)/2}(\mathbf{u}_i^\top \mathbf{u}_j) \quad (15)$$

where $C_k^\alpha(\cdot)$ denotes the Gegenbauer polynomial of index α and order k . The asymptotic distribution of S_n under \mathcal{H}_0 is the infinite linear combination of independent chi-squared distribution $\sum_{k=1}^\infty v_k^2 \mathcal{X}_{p_{d,k}}^2$.

Connection to MHE. Let $v_k = 0$ when k is even and $v_k = (\pi k)^{-1}$ when k is odd. As one of the Sobolev test, Ajne test [61, 1] uses the following test statistic:

$$A_n = \frac{n}{4} - \frac{1}{n\pi} \sum_{1 \leq i < j \leq n} \arccos(\mathbf{u}_i^\top \mathbf{u}_j) \quad (16)$$

which rejects \mathcal{H}_0 for large values. It is well connected to MHE in the sense that both MHE and Eq. (16) are formulated based on pairwise relations. More specifically, Minimizing Eq. (16) *w.r.t.* $\{\mathbf{u}_i\}_{i=1}^n$ is in fact equivalent to MHE with Riesz s -kernel where $s = -1$.

Connection to R-MHP. Rayleigh test [65], which is a special case of the Sobolev test (with $v_1 = 1$ and $v_k = 0$ for $k \geq 2$), gives the test statistic in \mathbb{S}^{d-1} : $R_n = nd \|\bar{\mathbf{u}}\|^2$ where $\bar{\mathbf{u}} := \frac{1}{n} \sum_{i=1}^n \mathbf{u}_i$. R_n is asymptotically distributed as \mathcal{X}_d^2 under \mathcal{H}_0 . Therefore, $\|\bar{\mathbf{u}}\|$ is approaching to zero, which exactly matches the objective function of the R-MHP.

Connection to MHS and MHC. Spacing tests on \mathbb{S}^1 are constructed from the gaps between the ordered samples (*i.e.*, $\theta_{i+1} > \theta_i, \forall i$): $d_i := \theta_{i+1} - \theta_i$ where $i = 1, \dots, n-1$ and $d_n := 2\pi - (\theta_n - \theta_1)$. A special spacing test, called Range test [64], has the following statistic: $T_n := 2\pi - \max_i d_i$ which rejects \mathcal{H}_0 with low values. Then we have the following proposition that builds connection among range test, MHS and MHC:

Proposition 5. *Maximizing the statistic of range test on \mathbb{S}^1 is equivalent to MHS and MHC on \mathbb{S}^1 .*

In summary, there is a close connection between statistical uniformity testing on the hypersphere and our proposed objectives towards hyperspherical uniformity. Revisiting classic uniformity tests can not only help us to gain more insights, but also inspire more useful ways to promote hyperspherical uniformity.

6 Discussions

Removing collinearity. Promoting hyperspherical uniformity may lead to the neuron collinearity problem, which causes redundancy and is therefore undesirable in neural networks [68, 71]. In order to address this, we use the virtual neuron trick in [46, 53]. We construct a set of virtual neurons that always have the opposite directions to the original neurons and then regularize both original and virtual neurons together. Specifically, we assume a set of original neurons $\{\mathbf{w}_i\}_{i=1}^n$ and the virtual opposite neurons are $\{\mathbf{v}_i = -\mathbf{w}_i\}_{i=1}^n$. Finally we will apply hyperspherical uniformity to simultaneously regularize all the neurons $\{\mathbf{w}_1, \dots, \mathbf{w}_n, \mathbf{v}_1, \dots, \mathbf{v}_n\}$. The virtual neuron trick is not applicable to R-MHP.

Why hyperspherical uniformity. A popular choice out of many existing relational regularizations is the orthogonality regularization [13, 52, 4, 32, 33, 63]. Despite its popularity, orthogonality yields a few drawbacks. When the number of neurons exceeds the neuron dimension, promoting orthogonality among neurons will become problematic [46]. In contrast, hyperspherical uniformity avoids such a problem while also being more general. It works well in all circumstances.

Increasing effective width of neural networks. Empirical evidences in [68, 71, 28, 29] show that naively training a neural network typically leads to severe neuron redundancy and [88] shows that increasing the width of neural networks can significantly improve generalization. Hyperspherical uniformity can increase the effective width and improve the representation efficiency by penalizing hyperspherically similar neurons.

Decoupling neuron norm and angle. The norm of neurons is typically regularized by individual regularizations, while the relational regularizations are designed for neuron directions. Orthonormality does not decouple the two components, since it also regularizes neuron norm to be close to 1, which serves a redundant role to weight decay. In contrast, hyperspherical uniformity fully decouples the two regularization components.

Optimization difficulty. Although all the proposed regularization objectives share the same goal, there is a substantial difference when actually performing gradient descent with them. MHP and MHC define max-min (or min-max) problems which require alternative update or unrolled formulation to solve, so optimizing them may lead to much more bad local minima and therefore is not as stable as MHE, MHS and MGD.

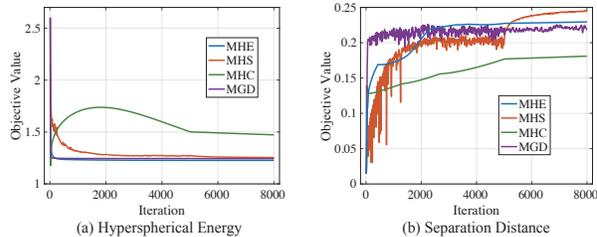


Figure 2: Comparison of local separatedness and global uniformity among MHE, MHS, MHC and MGD. (a) Hyperspherical energy vs. Iteration. (b) Separation distance vs. Iteration. Experimental details are in Appendix L. The figure is best viewed in color.

7 Properties and Insights

7.1 Geometric Properties

We discuss the geometric properties of the proposed regularizations by first connecting them with the hyperspherical uniformity. [46] has established that minimizing MHE ultimately leads to hyperspherical uniformity. This is achieved by showing that unit point mass at each vector in $\hat{\mathbf{W}}_n$ asymptotically approaches to the spherical measure on \mathbb{S}^{d-1} . [37, 9] have shown that asymptotically best-packing points on a rectifiable set are uniformly distributed. MHS is essentially a special case of this result, indicating that MHS also asymptotically leads to hyperspherical uniformity. [10, 8, 22] and [37] prove the same argument holds for MHP and MHC, respectively. [55] proves the asymptotic equidistribution of Fekete points on the hypersphere, implying that MGD also leads to hyperspherical uniformity.

Although all the regularization objectives asymptotically leads to hyperspherical uniformity, there are still some delicate differences in terms of the geometric properties. For example, MHS focuses on more on *local separatedness* by ensuring any nearest two vectors to be far away, while MHE characterizes the *global uniformity* with the sum of pairwise energies. Therefore, we can use the separation distance (*i.e.*, $\vartheta(\hat{\mathbf{W}}_n)$ in Eq. (4)) as a measure of local separatedness and use the hyperspherical energy (*i.e.*, $E_s(\hat{\mathbf{W}}_n)$ in Eq. (2)) as a measure of global uniformity. We connect MHE and MHC with the following geometric properties.

Theorem 4. *For $d-2 \leq s < d-1$, there is a constant $\lambda_{s,d} > 0$ such that for $n \geq 2$ and any solution $\hat{\mathbf{W}}_n^*$ that attains the optima of MHE, we have $\vartheta(\hat{\mathbf{W}}_n^*) \geq \frac{\lambda_{s,d}}{n^{1/(d-1)}}$.*

Theorem 5. *When $0 < s < d-1$ and n approaches infinity, every optimal MHS solution on \mathbb{S}^{d-1} is asymptotically an optimal solution for MHE.*

Theorem 4 gives a lower bound of separation distance for any MHE solution. It further indicates that optimal MHE configurations are well-separated and minimizing MHE objective can partially maximize the MHS objective. Theorem 5 shows that optimal MHS configurations will asymptotically have the same global uniformity (*i.e.*, hyperspherical energy) as optimal MHE

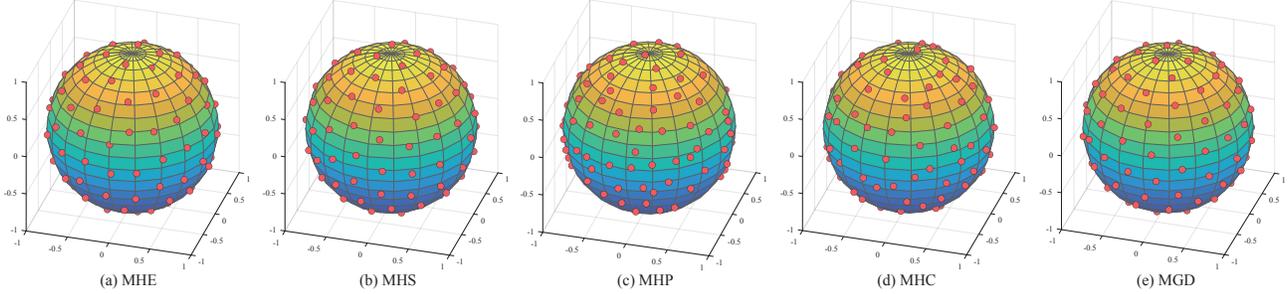


Figure 3: Visualization of regularization effects for MHE, MHS, MHP, MHC and MGD. We naively optimize all the proposed uniformity regularizations with gradient descent on 3D spherical data. Experimental details are given in Appendix L. The figure is best viewed in color.

configurations. These two results well connect the geometric properties between the MHE and MHS.

In fact, it requires involved analyses to fully study the geometric properties of these objectives [9, 11], which is out of our scope. In order to gain more insights, we conduct a simulation in Fig. 2 to empirically compare MHE, MHS, MHC and MGD. Specifically, we initialize 200 neurons and separately optimize MHE, MHS, MHC or MGD using gradient descent with these neurons. We compute the separation distance for all regularizations during training to compare their local separatedness, and also compute the hyperspherical energy to compare the global uniformity. We see that MHS and MHE achieve the best local separatedness and global uniformity, respectively. This is because MHS and MHE are directly optimizing the measures. MHE, MHS and MGD have similar global uniformity since they achieve similar hyperspherical energy. MHE, MGD and MHS reach a good balance between global uniformity and local separatedness, justifying their effectiveness. In contrast, MHC has worse geometric properties due to the complex objective. The results verify that although these hyperspherical uniformity regularizations share the same goal, the encoded regularization effects and inductive biases can be vastly different.

7.2 Spectral Properties

While orthogonality guarantees that all singular values are equal to 1, here we discuss the spectral properties of hyperspherical uniformity. The following theorem implies that hyperspherical uniformity will constrain the largest singular value from being too large and the smallest singular value from being too small.

Theorem 6. *Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be i.i.d. random vectors uniformly distributed on the unit hypersphere \mathbb{S}^{d-1} . If the ratio $\frac{n}{d}$ converges to a constant $\lambda \in (0, 1)$, asymptotically we have for $\mathbf{W} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\} \in \mathbb{R}^{d \times n}$ that*

$$\begin{aligned} \lim_{n \rightarrow \infty} \sigma_{\max}(\mathbf{W}) &\leq (\sqrt{d} + \sqrt{\lambda d}) \cdot \left(\max_i \frac{1}{\|\mathbf{u}_i\|_2} \right) \\ \lim_{n \rightarrow \infty} \sigma_{\min}(\mathbf{W}) &\geq (\sqrt{d} - \sqrt{\lambda d}) \cdot \left(\min_i \frac{1}{\|\mathbf{u}_i\|_2} \right) \end{aligned} \quad (17)$$

where $\sigma_{\max}(\cdot)$ and $\sigma_{\min}(\cdot)$ denote the largest and the

smallest singular value, respectively. $\mathbf{u}_i, \forall i$ are i.i.d. Gaussian variables with mean 0 and variance 1.

Theorem 6 is the direct application of a classic result on the limits of the largest and smallest singular values of Gaussian matrix [72]. The terms $\max_i 1/\|\mathbf{u}_i\|_2$ and $\min_i 1/\|\mathbf{u}_i\|_2$ in Eq. (17) come from the condition that $\mathbf{v}_1, \dots, \mathbf{v}_n$ lie on the unit hypersphere \mathbb{S}^{d-1} .

7.3 Connections to Neuron Initialization

We also discuss how hyperspherical uniformity is connected to the neuron initialization and its potential connection to Occam’s razor. First, almost all initialization schemes [42, 30, 25] for neural networks initializes the directions of neurons in the same layer to be uniformly distributed on the hypersphere, which is proved in Appendix J. Other than the zero-mean Gaussian distribution used in [42, 30, 25], we may alternatively characterize the hyperspherical uniformity with Cauchy-like distribution [75], which may inspire alternative initializations. Promoting hyperspherical uniformity in training aims to regularize the parameters to be similar to the initialization in terms of the directional support (since the pairwise relationship among neurons on the hypersphere is encouraged to be close to the initialization), which is essentially making the neural network to be as “simple” as possible. In contrast to the weight decay that implements Occam’s razor in terms of magnitude (*i.e.*, discrepancy to 0), promoting hyperspherical uniformity can essentially be viewed as implementing Occam’s razor in the angular space (*i.e.*, discrepancy to hyperspherical uniformity).

7.4 Regularization Effects

In order to intuitively understand these different objectives towards hyperspherical uniformity, we visualize their regularization effects by minimizing them with gradient descent on a 3D sphere. The results in Fig. 3 show that MHE, MHS and MGD yield visually superior regularization effects for hyperspherical uniformity, while MHP and MHC give suboptimal and weaker regularization effects partially due to their complex optimization objectives. Because both MHP and MHC involve an inner optimization problem, their objective

landscapes are highly non-convex and also more difficult to optimize. In contrast, MHE, MHS and MGD have much strong regularization effects on the 3-sphere.

8 Applications

Our goal is to show the performance gain by applying hyperspherical uniformity to different applications instead of achieving state-of-the-art performance. For fairness, we always stick to clean baselines without bells and whistles and ensure the experimental setup is the same for all compared methods. All the detailed experimental settings are specified in Appendix L.

8.1 Discriminative Learning

Multi-layer perceptrons. We first compare all the hyperspherical uniformity regularizations on MNIST with a 3-layer MLP. We use Xavier initialization [25] to initialize neuron weights in MLP. Results (error rates) in Table 1 show that all the proposed methods outperform the baseline by a considerable margin, indicating that hyperspherical uniformity is generally useful to improve the generalization for MLP. Among all the proposed regularizations, MGD works the best and reduces the error of the baseline by more than 23%.

Convolutional neural networks. We also compare all the hyperspherical uniformity regularizations on both plain VGG-like CNN [73] and ResNet [31]. Specific architecture configurations are in Appendix L.

CIFAR-100. We first perform experiments on CIFAR-100 with the plain CNN-9 and ResNet-32. The experimental settings mostly follows [46] for fair comparison. Table 2 gives the error rates. The results show that most of the hyperspherical uniformity regularizations perform significantly and consistently better than the baseline, the orthogonal regularization and a state-of-the-art orthogonality-based regularization called SRIP [4]. Surprisingly, R-MHP, as a very simple regularization method, also outperforms both the baseline and orthogonal regularization, implying the effectiveness of hyperspherical uniformity in general. Although MHP and MHC are less stable to optimize with gradient descent, they still perform reasonably well compared to the baseline, implying that promoting hyperspherical uniformity is generally beneficial to generalization. Mostly notably, MGD performs the best on CNN-9, while MHS performs the best on ResNet-18.

Method	Error
Baseline	2.14
Orthogonal	1.95
MHE	1.85
MHS	1.72
MHP	1.92
R-MHP	1.99
MHC	1.88
MGD	1.64

Table 1: MLP (%).

Method	CNN-9	ResNet-32
Baseline	28.13	22.87
Orthogonal	26.94	22.36
SRIP [4]	25.92	22.02
MHE	25.94	21.82
MHS	25.43	20.97
MHP	25.92	21.24
R-MHP	26.02	22.19
MHC	25.62	21.88
MGD	25.32	21.06

Table 2: CNN on CIFAR-100 (%).

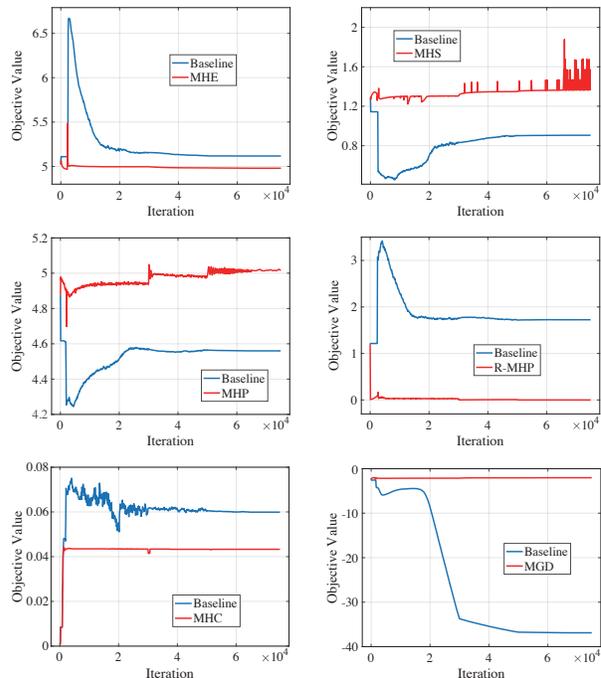


Figure 4: Regularization dynamics (objective value vs. iteration) of MHE, MHS, MHP, R-MHP, MHC and MGD.

It partially suggests that encouraging local separation may be more important to generalization than promoting global uniformity. One of the possible reasons could be that MHE updates all the neurons simultaneously by accumulating all the pairwise interactions while MHS only focus on updating the two most similar neurons.

Regularization dynamics. Then we look into the effectiveness of the hyperspherical uniformity regularizations for minimizing (or maximizing) the objective values during training. We perform classification on CIFAR-100 with plain CNN-9, and plot the objective value curve of each hyperspherical uniformity regularization during training. The results are given in Fig. 4. Note that, MHS, MHP and MGD maximize the regularization objective, while MHE, R-MHP and MHC minimize the objective. From Fig. 4, we can see that all the hyperspherical uniformity achieve better objective values than the baseline, which indicates that all the proposed regularizations have well served the purpose.

ImageNet. Finally, we conduct the experiments on ImageNet-2012 [69] with ResNet-18 to further evaluate the performance of hyperspherical uniformity on large-scale datasets. Since our purpose is to compare all the regularizations, we use the same simple data augmentation scheme as in AlexNet [39]. Detailed settings are given in Appendix L. Table 3 shows the Top-1 error on ImageNet test set. One can see from the results that most of

Method	Error
Baseline	32.95
Orthogonal	32.65
SRIP [4]	32.53
MHE	32.45
MHS	32.06
MHP	32.32
R-MHP	32.71
MHC	32.28
MGD	32.16

Table 3: ImageNet (%).

the hyperspherical uniformity regularizations can effectively improve the generalization of ResNet-18 on ImageNet. Among all, MGD achieves the best performance and outperforms the baseline by nearly 2%. Considering the fact that we are merely adding a regularization without changing the network architecture, the improvement is actually very significant.

Graph networks.

In order to show that hyperspherical uniformity is diversely useful, we also perform experiments on graph convolution networks (GCN) [36].

Method	Citeseer	Cora	Pubmed
Baseline	70.3	81.3	79.0
Orthogonal	70.4	81.5	78.8
MHE	71.5	82.0	79.0
MHS	71.7	82.3	79.2
MHP	71.3	81.5	79.0
MHC	71.2	81.6	79.0
MGD	71.8	82.3	79.2

Table 4: Graph networks (%).

We use the same 2-layer GCN as [36] and evaluate on Citeseer, Cora and Pubmed data [70]. Regularizing GCN is conceptually similar to MLP. Specifically, the forward model of GCN is $\mathbf{Z} = \text{Softmax}(\tilde{\mathbf{A}} \cdot \text{ReLU}(\tilde{\mathbf{A}} \cdot \mathbf{X} \cdot \mathbf{W}_0) \cdot \mathbf{W}_1)$ where $\tilde{\mathbf{A}} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}}$. \mathbf{A} is the adjacency matrix of the graph, $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ (\mathbf{I} is an identity matrix), and $\tilde{\mathbf{D}} = \sum_j \mathbf{A}_{ij}$. $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the feature matrix of n nodes in the graph (feature dimension is d). \mathbf{W}_1 is the weights of the classifiers. \mathbf{W}_0 is the weight matrix of size $d \times h$ where h is the dimension of the hidden space. We treat each column vector of \mathbf{W}_0 as a neuron, so there are h neurons in total. We apply the hyperspherical uniformity regularizations on \mathbf{W}_0 and report the testing accuracy in Table 4. From the results, we observe all the hyperspherical uniformity regularizations outperform the baseline by a considerable margin, indicating that promoting hyperspherical uniformity is very helpful to the generalization of GCN. On Citeseer, MHS achieves the best accuracy and outperforms the baseline by 1.6%. On Cora and Pubmed, MGD performs the best. The results validate the universality and superiority of hyperspherical uniformity.

Point cloud networks.

We also evaluate the hyperspherical uniformity on the 3D point cloud classification task where each 3D object is represented by a unordered set of points (*i.e.*, 3D coordinates).

Method	Accuracy
Baseline	87.10
MHE	87.44
MHS	87.60
MHP	87.41
R-MHP	87.10
MHC	87.33
MGD	87.61

Table 5: PointNet (%).

PointNet [62] is a neural network designed for processing point clouds. PointNet consists of a group of weight-sharing MLPs and we regularize these MLPs using our hyperspherical uniformity regularizations. For simplicity, we use a vanilla PointNet without T-Net and experiment on the ModelNet-40 dataset [82]. The classification accuracy is given in Table 5. All the hyperspherical uniformity regularizations consistently outperform the baseline PointNet by a significant margin, showing the effectiveness of hyperspherical uniformity in point cloud classification.

Among all the compared regularizations, MGD achieves the best accuracy by outperforming the baseline by 0.73%, validating its universality in improving generalization for different types of neural networks.

8.2 Generative Modeling

Beside the applications in the discriminative learning for classification tasks, we further apply hyperspherical uniformity to improve the generative adversarial network (GAN) [27] in unconditional image generation.

Specifically, we perform image generation on CIFAR-10 with a vanilla GAN. We regularize the discriminator in the vanilla GAN using MHE, MHS, MGD and R-MHP. We adopt the following vanilla GAN as the clean baseline:

Method	Inception Score
Baseline	7.14
SN [57]	7.40
MHE	7.40
MHS	7.61
R-MHP	7.31
MGD	7.49

Table 6: GAN on CIFAR-10.

where $V(G, D) = \mathbb{E}_{\mathbf{x} \sim q_{\text{data}}(\mathbf{x})}[\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})}[\log(1 - D(G(\mathbf{z})))]$ where $D(\cdot)$ denotes the discriminator and $G(\cdot)$ denotes the generator. We do not use spectral normalization for the discriminator and generator. We train the models for 200k iterations on CIFAR-10 using Adam optimizer ($\beta_1 = 0.5, \beta_2 = 0.999$) with learning rate set to $1e-3$ and batch size set to 64. Detailed experimental setup is given in Appendix L. Results in Table 6 show that hyperspherical uniformity can generally improve the generation quality of GANs. Notably, MHS achieves the best inception score 7.61, which outperforms the vanilla GAN and the state-of-the-art spectral normalization (SN) [57] by a considerable margin. MGD also performs better than SN, well verifying that promoting hyperspherical uniformity is beneficial to GANs.



Figure 5: Comparison of randomly generated images from GAN.

We also give some qualitative examples in Fig. 5 to compare the generation quality between baseline and hyperspherical uniformity regularizations. The results show that hyperspherical uniformity can greatly improve GAN by regularizing it to generate visually plausible images with more diverse color and semantic meaning.

9 Concluding Remarks

This paper considers a novel family of regularizations – hyperspherical uniformity for training neural networks. Specifically, we propose several conceptually appealing instances and provide some statistical discussions and theoretical insights. Comprehensive experiments well validate the effectiveness of hyperspherical uniformity.

Acknowledgements

Weiyang Liu and Adrian Weller acknowledge support from the Leverhulme Trust via CFI. Adrian Weller acknowledges support from the David MacKay Newton research fellowship at Darwin College, and The Alan Turing Institute under EPSRC grant EP/N510129/1 and U/B/000074.

References

- [1] Björn Ajne. A simple test for uniformity of a circular distribution. *Biometrika*, 55(2):343–354, 1968. 5
- [2] Marcin Andrychowicz, Misha Denil, Sergio Gomez Colmenarejo, Matthew W Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *NIPS*, 2016. 4
- [3] Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *ICML*, 2016. 2
- [4] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep networks? In *NeurIPS*, 2018. 2, 6, 8
- [5] RJ Beran. Testing for uniformity on a compact homogeneous space. *Journal of Applied Probability*, 5(1):177–195, 1968. 5
- [6] Robert Berman, Sébastien Boucksom, David Witt Nyström, et al. Fekete points and convergence towards equilibrium measures on complex manifolds. *Acta Mathematica*, 207(1):1–27, 2011. 4
- [7] LE Blumenson. A derivation of n-dimensional spherical coordinates. *The American Mathematical Monthly*, 67(1):63–66, 1960. 1
- [8] S Borodachov, D Hardin, A Reznikov, and E Saff. Optimal discrete measures for riesz potentials. *Transactions of the American Mathematical Society*, 370(10):6973–6993, 2018. 6
- [9] S Borodachov, D Hardin, and E Saff. Asymptotics of best-packing on rectifiable sets. *Proceedings of the American Mathematical Society*, 135(8):2369–2380, 2007. 6, 7
- [10] Sergiy V Borodachov and Nattapong Bosuwan. Asymptotics of discrete riesz d-polarization on subsets of d-dimensional manifolds. *Potential Analysis*, 41(1):35–49, 2014. 6
- [11] Sergiy V Borodachov, Douglas P Hardin, and Edward B Saff. *Discrete energy on rectifiable sets*. Springer, 2019. 3, 4, 5, 7
- [12] Johann S Brauchart, Peter D Dragnev, and Edward B Saff. Riesz external field problems on the hypersphere and optimal point separation. *Potential analysis*, 41(3):647–678, 2014. 18
- [13] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016. 6
- [14] Beidi Chen, Weiyang Liu, Zhiding Yu, Anshumali Shrivastava, Animesh Garg, and Animashree Anandkumar. Angular visual hardness. In *ICML*, 2020. 2
- [15] Krzysztof Choromanski, Carlton Downey, and Byron Boots. Initialization matters: Orthogonal predictive state recurrent neural networks. In *ICLR*, 2018. 2
- [16] Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. *arXiv preprint arXiv:1511.06068*, 2015. 2
- [17] Henry Cohn and Abhinav Kumar. Universally optimal distribution of points on spheres. *Journal of the American Mathematical Society*, 20(1):99–148, 2007. 5
- [18] Bo Dai, Hanjun Dai, Niao He, Weiyang Liu, Zhen Liu, Jianshu Chen, Lin Xiao, and Le Song. Coupled variational bayes via optimization embedding. In *NeurIPS*, 2018. 4
- [19] Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018. 2
- [20] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 2
- [21] PD Dragnev and EB Saff. Riesz spherical potentials with external fields and minimal energy points separation. *Potential Analysis*, 26(2):139–162, 2007. 18
- [22] Tamás Erdélyi and Edward B Saff. Riesz polarization inequalities in higher dimensions. *Journal of Approximation Theory*, 171:128–147, 2013. 6
- [23] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 4

- [24] Evarist Giné. Invariant tests for uniformity on compact riemannian manifolds based on Sobolev norms. *The Annals of statistics*, pages 1243–1266, 1975. 5
- [25] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, 2010. 7, 8
- [26] Sheldon Goldstein, Joel L Lebowitz, Roderich Tumulka, and Nino Zanghi. Any orthonormal basis in high dimension is uniformly distributed over the sphere. *arXiv preprint arXiv:1406.2576*, 2014. 2
- [27] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 9
- [28] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016. 6
- [29] Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. In *NIPS*, 2015. 6
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 7
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8, 22
- [32] Lei Huang, Xianglong Liu, Bo Lang, Adams Wei Yu, Yongliang Wang, and Bo Li. Orthogonal weight normalization: Solution to optimization over multiple dependent stiefel manifolds in deep neural networks. In *AAAI*, 2018. 2, 6
- [33] Kui Jia, Shuai Li, Yuxin Wen, Tongliang Liu, and Dacheng Tao. Orthogonal deep neural networks. *TPAMI*, 2019. 6
- [34] Li Jing, Yichen Shen, Tena Dubcek, John Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, and Marin Soljačić. Tunable efficient unitary neural networks (eunn) and their application to rnns. In *ICML*, 2017. 2
- [35] Toni Karvonen, Simo Särkkä, and Ken’ichiro Tanaka. Kernel-based interpolation at approximate feketé points. *Numerical Algorithms*, pages 1–24, 2020. 4
- [36] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 9
- [37] Andrei Nikolaevich Kolmogorov and Vladimir Mikhailovich Tikhomirov. ε -entropy and ε -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14(2):3–86, 1959. 6
- [38] AV Kolushov and VA Yudin. Extremal dispositions of points on the sphere. *Analysis Mathematica*, 23(1):25–34, 1997. 5
- [39] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 8
- [40] ABJ Kuijlaars, EB Saff, and X Sun. On separation of minimal riesz energy points on spheres in euclidean spaces. *Journal of computational and applied mathematics*, 199(1):172–180, 2007. 18
- [41] Arno Kuijlaars and E Saff. Asymptotics for minimal discrete energy on the sphere. *Transactions of the American Mathematical Society*, 350(2):523–538, 1998. 3
- [42] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012. 7
- [43] Paul Leopardi. Discrepancy, separation and riesz energy of finite point sets on the unit sphere. *Advances in Computational Mathematics*, 39(1):27–43, 2013. 18
- [44] Jianxin Li, Haoyi Zhou, Pengtao Xie, and Yingchun Zhang. Improving the generalization performance of multi-class svm via angular regularization. In *IJCAI*, 2017. 2
- [45] Rongmei Lin, Weiyang Liu, Zhen Liu, Chen Feng, Zhiding Yu, James M. Rehg, Li Xiong, and Le Song. Regularizing neural networks via minimizing hyperspherical energy. In *CVPR*, 2020. 2, 3
- [46] Weiyang Liu, Rongmei Lin, Zhen Liu, Lixin Liu, Zhiding Yu, Bo Dai, and Le Song. Learning towards minimum hyperspherical energy. In *NeurIPS*, 2018. 1, 2, 3, 6, 8
- [47] Weiyang Liu, Rongmei Lin, Zhen Liu, James M. Rehg, Li Xiong, Adrian Weller, and Le Song. Orthogonal over-parameterized training. *arXiv preprint arXiv:2004.04690*, 2020. 2
- [48] Weiyang Liu, Zhen Liu, James Rehg, and Le Song. Neural similarity learning. In *NeurIPS*, 2019. 2

- [49] Weiyang Liu, Zhen Liu, Zhiding Yu, Bo Dai, Rongmei Lin, Yisen Wang, James M Rehg, and Le Song. Decoupled networks. *CVPR*, 2018. 2
- [50] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spherefacer: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 2
- [51] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016. 2
- [52] Weiyang Liu, Yan-Ming Zhang, Xingguo Li, Zhiding Yu, Bo Dai, Tuo Zhao, and Le Song. Deep hyperspherical learning. In *NIPS*, 2017. 2, 6
- [53] Yueming Lyu. Spherical structured feature maps for kernel approximation. In *ICML*, 2017. 6
- [54] Albert W Marshall, Ingram Olkin, and Barry C Arnold. *Inequalities: theory of majorization and its applications*, volume 143. Springer, 1979. 19
- [55] Jordi Marzo and Joaquim Ortega-Cerdà. Equidistribution of feketé points on the sphere. *Constructive Approximation*, 32(3):513–521, 2010. 4, 6
- [56] Vitali D Milman and Gideon Schechtman. *Asymptotic theory of finite dimensional normed spaces: Isoperimetric inequalities in riemannian manifolds*, volume 1200. Springer, 2009. 2
- [57] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 9, 22, 23
- [58] Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *arXiv preprint arXiv:1912.10557*, 2019. 4
- [59] Sean O’Rourke, Van Vu, and Ke Wang. Eigenvectors of random matrices: a survey. *Journal of Combinatorial Theory, Series A*, 144:361–442, 2016. 19
- [60] Sung Woo Park and Junseok Kwon. Sphere generative adversarial network based on geometric moment matching. In *CVPR*, 2019. 2
- [61] MJ Prentice et al. On invariant tests of uniformity for directions and orientations. *The Annals of Statistics*, 6(1):169–176, 1978. 5
- [62] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 9, 22, 23
- [63] Haozhi Qi, Chong You, Xiaolong Wang, Yi Ma, and Jitendra Malik. Deep isometric learning for visual recognition. In *ICML*, 2020. 6
- [64] JS Rao. *Some contributions to the analysis of circular data*. PhD thesis, Indian Statistical Institute, Kolkata, 1969. 5
- [65] Lord Rayleigh. On the problem of random vibrations, and of random flights in one, two, or three dimensions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 37(220):321–347, 1919. 5
- [66] Ran Raz. Exponential separation of quantum and classical communication complexity. In *SOTC*, 1999. 2
- [67] Pau Rodríguez, Jordi Gonzalez, Guillem Cucurull, Josep M Gonfaus, and Xavier Roca. Regularizing cnns with locally constrained decorrelations. *arXiv preprint arXiv:1611.01967*, 2016. 2
- [68] Aruni RoyChowdhury, Prakhar Sharma, Erik Learned-Miller, and Aruni Roy. Reducing duplicate filters in deep neural networks. In *NIPS workshop on Deep Learning: Bridging Theory and Practice*, 2017. 6
- [69] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, pages 1–42, 2014. 8
- [70] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliass-Rad. Collective classification in network data. *AI magazine*, 2008. 9
- [71] Wenling Shang, Kihyuk Sohn, Diogo Almeida, and Honglak Lee. Understanding and improving convolutional neural networks via concatenated rectified linear units. In *ICML*, 2016. 2, 6
- [72] Jack W Silverstein et al. The smallest eigenvalue of a large dimensional wishart matrix. *The Annals of Probability*, 13(4):1364–1368, 1985. 7, 19
- [73] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 8, 22
- [74] Ian H Sloan and Robert S Womersley. Extremal systems of points and numerical integration on the sphere. *Advances in Computational Mathematics*, 21(1-2):107–125, 2004. 4

- [75] Paweł J Szablowski. Uniform distributions on spheres in finite dimension and their generalizations. *Journal of multivariate analysis*, 64(2):103–117, 1998. 7
- [76] Cheng Tai, Tong Xiao, Yi Zhang, Xiaogang Wang, et al. Convolutional neural networks with low-rank regularization. *arXiv preprint arXiv:1511.06067*, 2015. 2
- [77] Pieter Merkus Lambertus Tammes. On the origin of number and arrangement of the places of exit on the surface of pollen-grains. *Recueil des travaux botaniques néerlandais*, 27(1):1–84, 1930. 3
- [78] Joseph John Thomson. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 7(39):237–265, 1904. 2
- [79] Feng Wang, Weiyang Liu, Haijun Liu, and Jian Cheng. Additive margin softmax for face verification. *arXiv preprint arXiv:1801.05599*, 2018. 2
- [80] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Zhifeng Li, Dihong Gong, Jingchao Zhou, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018. 2
- [81] Scott Wisdom, Thomas Powers, John Hershey, Jonathan Le Roux, and Les Atlas. Full-capacity unitary recurrent neural networks. In *NIPS*, 2016. 2
- [82] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 9, 22
- [83] Bo Xie, Yingyu Liang, and Le Song. Diverse neural network learns true target functions. In *AISTATS*, 2017. 1
- [84] Pengtao Xie, Yuntian Deng, Yi Zhou, Abhimanu Kumar, Yaoliang Yu, James Zou, and Eric P Xing. Learning latent space models with angular constraints. In *ICML*, 2017. 2
- [85] Pengtao Xie, Aarti Singh, and Eric P Xing. Uncorrelation and evenness: a new diversity-promoting regularizer. In *ICML*, 2017. 2
- [86] Yuichi Yoshida and Takeru Miyato. Spectral norm regularization for improving the generalizability of deep learning. *arXiv preprint arXiv:1705.10941*, 2017. 2
- [87] Vladimir Alexandrovich Yudin. Minimum potential energy of a point system of charges. *Diskretnaya Matematika*, 4(2):115–121, 1992. 2, 5
- [88] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 6