
Sparsifying networks by traversing Geodesics

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The geometry of weight spaces and functional manifolds of neural networks play
2 an important role towards ‘understanding’ the intricacies of ML. In this paper,
3 we attempt to solve certain open questions in ML, by viewing them through the
4 lens of geometry, ultimately relating it to the discovery of points or paths of
5 equivalent function in these spaces. We propose a mathematical framework to
6 evaluate geodesics in the functional space, to find high-performance paths from
7 a dense network to its sparser counterpart. Our results are obtained on VGG-11
8 trained on CIFAR-10 and MLP’s trained on MNIST. Broadly, we demonstrate that
9 the framework is general, and can be applied to a wide variety of problems, ranging
10 from sparsification to alleviating catastrophic forgetting.

11 1 Introduction

12 The geometry of weight manifolds and functional spaces represented by artificial neural networks is
13 an important window to ‘understanding’ machine learning. Many open questions in machine learning,
14 when viewed through the lens of geometry, can be related to finding points or paths of equivalent
15 function in these spaces.

16 Two prominent examples are (i) Enabling networks to learn multiple tasks sequentially while avoiding
17 catastrophic forgetting[1] and (ii) to discover high-performance paths from a dense neural network to
18 its sparser counterparts (sparsification) [2, 3]. Both these questions, although appearing very different,
19 can be adequately solved by formulating paths in the weight manifold.

20 In this paper, we propose a mathematical framework grounded in differential geometry to find novel
21 solutions to these problems in machine learning. We formalize the “search” for a suitable network
22 as a dynamic movement on a curved pseudo-Riemannian manifold [4]. Further, we demonstrate
23 that geodesics, minimum length paths, on the weight manifold provide high performance paths that
24 the network can traverse to maintain performance while ‘searching-out’ for networks that satisfy
25 additional objectives. Specifically we develop a procedure based on the geodesic equation for
26 finding sets of path connected networks that achieve high performance while also satisfying a second
27 objective like sparsification. Broadly, our work provides procedures that will help discover an array
28 of high-performing neural network architectures for the task of interest.

29 2 Mathematical framework

30 **Preliminaries:** We represent a feed-forward neural network as a smooth, C^∞ function $f(\mathbf{x}, \mathbf{w})$, that
31 maps an input vector, $\mathbf{x} \in \mathbb{R}^k$, to an output vector, $f(\mathbf{x}, \mathbf{w}) = \mathbf{y} \in \mathbb{R}^m$. The function, $f(\mathbf{x}, \mathbf{w})$, is
32 parameterized by a vector of weights, $\mathbf{w} \in \mathbb{R}^n$, that are typically set in training to solve a specific
33 task. We refer to $W = \mathbb{R}^n$ as the *weight space* (W) of the network, and we refer to $\mathcal{F} = \mathbb{R}^m$ as
34 the *functional manifold* [5]. In addition to f , we will sometimes be interested in considering a loss




Figure 1: **Geometric framework for analyzing neural network resilience** (A) Three networks (N_1, N_2, N_3) in weights space W and their relative distance in functional space and loss space. Damage is analyzed by asking how movement in weight space changes functional performance and loss through introduction of a pullback metric g .

35 function, $L : \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}$, that provides a scalar measure of network performance for a given task
 36 (Figure 1).

37 **Construction of metric tensor (g):** We use differential geometry, rooted in a functional notion
 38 of distance, to analyse how infinitesimal movements in the weight space (W) impact functional
 39 performance of the network. Specifically, we construct a local distance metric, g , that can be applied
 40 at any point in W to measure the functional impact of an arbitrary network perturbation.

41 To construct a metric mathematically, we fix the input, \mathbf{x} , into a network and ask how the output of the
 42 network, $f(\mathbf{w}, \mathbf{x})$, moves on the functional manifold, \mathcal{F} , given an infinitesimal weight perturbation,
 43 $d\mathbf{u}$, in W where $\mathbf{w}_d = \mathbf{w}_t + d\mathbf{u}$. For an infinitesimal perturbation $d\mathbf{u}$,

$$f(\mathbf{x}, \mathbf{w}_t + d\mathbf{u}) \approx f(\mathbf{x}, \mathbf{w}_t) + \mathbf{J}_{\mathbf{w}_t} d\mathbf{u}, \quad (1)$$

44 where $\mathbf{J}_{\mathbf{w}_t}$ is the Jacobian of $f(\mathbf{x}, \mathbf{w})$ for a fixed \mathbf{x} , $J_{i,j} = \frac{\partial f_i}{\partial w_j}$, evaluated at \mathbf{w}_t . We measure the
 45 change in functional performance given $d\mathbf{u}$ as the mean squared error

$$d(\mathbf{w}_t, \mathbf{w}_d) = |f(\mathbf{x}, \mathbf{w}_t) - f(\mathbf{x}, \mathbf{w}_d)|^2 = d\mathbf{u}^T (\mathbf{J}_{\mathbf{w}_t}^T \mathbf{J}_{\mathbf{w}_t}) d\mathbf{u} \quad (2)$$

$$= d\mathbf{u}^T \mathbf{g}_{\mathbf{w}_t} d\mathbf{u}, \quad (3)$$

$$(4)$$

46 where $\mathbf{g}_{\mathbf{w}_t} = \mathbf{J}_{\mathbf{w}_t}^T \mathbf{J}_{\mathbf{w}_t}$ is the metric tensor evaluated at the point $\mathbf{w}_t \in W$. The metric tensor \mathbf{g} is
 47 an $n \times n$ symmetric matrix that defines an inner product and local distance metric, $\langle d\mathbf{u}, d\mathbf{u} \rangle_{\mathbf{w}} =$
 48 $d\mathbf{u}^T \mathbf{g}_{\mathbf{w}} d\mathbf{u}$, on the tangent space of the manifold, $T_{\mathbf{w}}(W)$ at each $\mathbf{w} \in W$.

49 **Global paths in functional space:** Globally, we can use the metric to determine the functional
 50 performance change across a path connected set of networks. Mathematically, the metric changes as
 51 we move in W due to the curvature of the ambient space that reflects changes in the vulnerability of
 52 a network to weight perturbation (Figure 1c). As a network moves along a path, $\gamma(t) \in W$ from a
 53 given trained network $\gamma(0) = \mathbf{w}_t$ to a sparse-hyperplane (hyperplane that consists neural networks
 54 of desired sparsity) such that $\gamma(1) = \mathbf{w}_s$, we can analyze the integrated impact of sparsification on
 55 network performance along $\gamma(t)$ by using the metric (g) to calculate the length of the path $\gamma(t)$ as:

$$L(\gamma) = \int_0^1 \left\langle \frac{d\gamma(t)}{dt}, \frac{d\gamma(t)}{dt} \right\rangle_{\gamma(t)} dt, \quad (5)$$

56 where $\langle \frac{d\gamma(t)}{dt}, \frac{d\gamma(t)}{dt} \rangle_{\gamma(t)} = \frac{d\gamma(t)}{dt}^T \mathbf{g}_{\gamma(t)} \frac{d\gamma(t)}{dt}$ is the infinitesimal functional change accrued while
 57 traversing path $\gamma(t) \in W$. As our objective is to find a path $\gamma(t) \in W$ that minimizes the total
 58 function length ($L(\gamma)$), we evaluate the geodesic between the dense network and its sparser counter
 59 part.

60 3 Geodesics

61 In this section, we evaluate the geodesic path, defined as the minimum functional path-length in
 62 weight space between any pair of networks. We also demonstrate that the formulation of high-
 63 performance paths holds the key to solving seemingly unrelated open-questions: (i) catastrophic
 64 forgetting and (ii) network sparsification.

65 **Network sparsification:** We pose the problem of finding a $p\%$ sparse counterpart (w_s^p) of a dense
 66 neural network (w_t) as finding a geodesic from (w_{t1}) to a $p\%$ sparse hyperplane (W_s^p). The sparse
 67 hyper-plane is the set of all networks that are $p\%$ sparse. As the geodesic minimizes the total distance
 68 traversed on the functional space, it ensures that the sparse network obtained will be both, functionally
 69 similar to the original network and high-performing.

70 **Catastrophic forgetting:** In order to obtain a single network well-trained on two sequential tasks,
 71 we train two networks on the two tasks independently (w_{t1} , w_{t2}) and pose the problem of finding a
 72 network that performs well on both the tasks, as a geodesic between the two trained networks (w_{t1}
 73 and w_{t2}) on the functional space. The functional manifold used could pertain to either of the tasks
 74 (ensuring that a single dataset/task is used at a time).

75 **Geodesics mathematical machinery:** To find geodesics on W , we can solve the geodesic equation
 76 given by:

$$\frac{d^2 w^\eta}{dt^2} + \Gamma_{\mu\nu}^\eta \frac{dw^\mu}{dt} \frac{dw^\nu}{dt} = 0 \quad (6)$$

77 where, w^j defines the j 'th basis vector of the weights space W , $\Gamma_{\mu\nu}^\eta$ specifies the Christoffel symbols
 78 ($\Gamma_{\mu\nu}^\eta = \sum_r \frac{1}{2} g_{\eta r}^{-1} (\frac{\partial g_{r\mu}}{\partial x^\nu} + \frac{\partial g_{r\nu}}{\partial x^\mu} - \frac{\partial g_{\mu\nu}}{\partial x^r})$) on the manifold. The Christoffel symbols capture infinitesimal
 79 changes in the metric tensor (\mathbf{g}) along a set of directions in the manifold. They are computed by
 80 setting the covariant derivative of the metric tensor along a path specified by $\gamma(t)$ to zero.

81 We, specifically, compute geodesic paths, $\gamma(t)$, so that $\gamma(0) = \mathbf{w}_{t1}$ and $\gamma(1) \in W_s$ where W_s is
 82 the sparsity hyper-plane (for network sparsification) and $\gamma(1) = w_{t2}$ (for the catastrophic forgetting
 83 problem). As the computation of the Christoffel symbols is both memory and computationally
 84 intensive, we propose an optimization algorithm to evaluate the ‘approximate’ geodesic in the
 85 manifold.

86 Given a trained network, our procedure updates the weights of the network to optimize performance
 87 given a direction of sparsification. To discover a geodesic path $\gamma(t)$, we begin at a trained network
 88 and iteratively solve for the tangent vector, $\theta(w)$, at every point, $\mathbf{w} = \gamma(t)$, along the path, starting
 89 from \mathbf{w}_{t1} and terminating at the sparse hyperplane, W_s^p (for sparsification) or network (w_{t2}) trained
 90 for task-2 (for catastrophic forgetting). We specifically solve

$$\text{argmin}_{\theta(\mathbf{w})} \langle \theta(\mathbf{w}), \theta(\mathbf{w}) \rangle_{\mathbf{w}} - \beta \theta(\mathbf{w})^T v_w \text{ subject to: } \theta(\mathbf{w})^T \theta(\mathbf{w}) \leq 0.01. \quad (7)$$

91 The tangent vector $\theta(\mathbf{w})$ is obtained by simultaneously optimizing two objective functions: (1)
 92 minimizing the increase in functional distance along the path measured by the metric tensor (\mathbf{g}_w)
 93 [min: $(\langle \theta(\mathbf{w}), \theta(\mathbf{w}) \rangle_{\mathbf{w}}) = (\theta(\mathbf{w})^T \mathbf{g}_w \theta(\mathbf{w}))$] and (2) maximizing the dot-product between the tangent
 94 vector and v_w , vector pointing in the desired direction [max: $(\theta(\mathbf{w})^T v_w)$] to enable movement towards
 95 the sparse hyperplane or network w_{t2} .

96 Our optimization procedure is a quadratic program that trades off, through the hyper-parameter
 97 β , motion towards the damage hyper-plane and the maximization of the functional performance
 98 of the intermediate networks along the path (elaborated in the appendix). The strategy discovers
 99 multiple paths from the trained network \mathbf{w}_{t1} to W_s , sparse hyper-plane, where networks maintain
 100 high functional performance during sparsification. Of the many paths obtained, we can select the path
 101 with the shortest total length (with respect to the metric \mathbf{g}) as the best approximation to the geodesic
 102 in the manifold.

103 **Geodesics for network sparsification:** We apply the geodesic strategy to discover high-performance
 104 paths from a trained network (VGG11 on CIFAR-10) to a sparse hyperplane (blue curve in figure-2A).
 105 Here, the sparse hyperplane is defined as all networks that have zeroed out 50 (out of 60) conv-filters
 106 from layer-1 in VGG11. We compared our strategy with conventional heuristic fine-tuning (figure-
 107 2A) and demonstrate that the geodesic procedure is both rationale and computationally efficient.
 108 Specifically, an iterative prune-train cycle achieved through structured pruning of a single node at a
 109 time, coupled with SGD re-training [6, 3] (Figure 4A) requires 70 training epochs to identify a
 110 sparsification path. However, our geodesic strategy finds paths that quantitatively out-perform the
 111 iterative prune-train procedure and obtains these paths with only 10 training epochs (figure 2A,B).

112 In figure-2C, we apply the geodesic strategy to find a path from a trained multi-layer perceptron
 113 (MLP) to its 70% sparser counterpart. The geodesic discovers a network performing at a test-accuracy
 114 of 96.8% on the 70% sparse hyperplane.

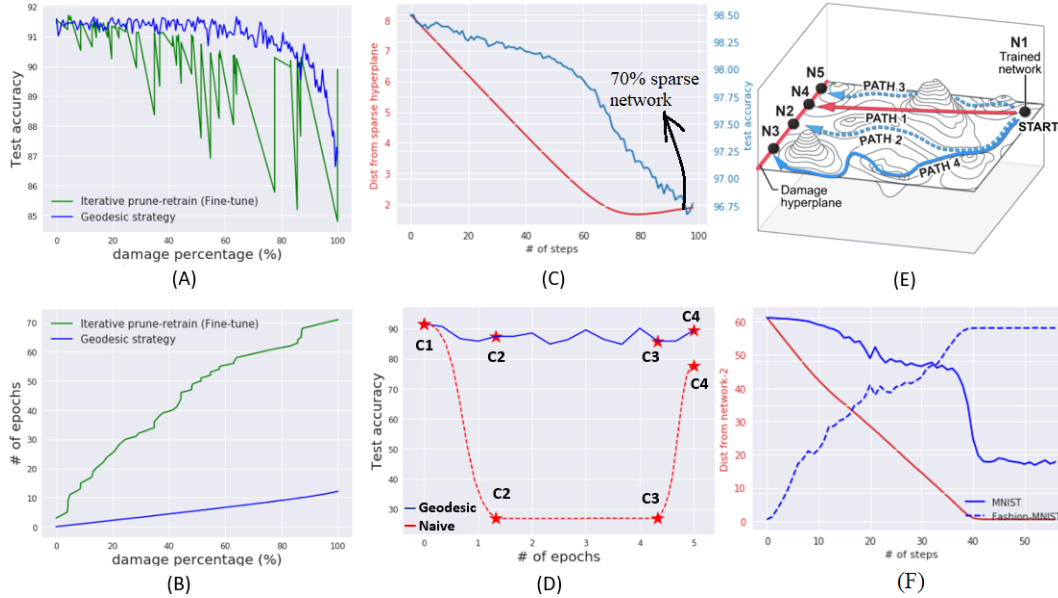


Figure 2: Geodesic paths allow damage compensation through weight adjustment: (A) Test accuracy (A) and (B) number of network update epochs (B) for geodesic recovery (blue) vs fine-tuning (green) while 50 (out of 60) conv-filters are deleted from layer1 in VGG11. Geodesic recovery requires ≤ 10 total update epochs. (C) Traversing the geodesic from an MLP (trained on MNIST) to its 70% sparser counterpart. The blue line captures the test-accuracy of the networks along the path and red-line captures the euclidean distance of the network from the sparse hyperplane (W_s^P). (D) Geodesic strategy (blue) allows networks to dynamically transition between configurations: C1, trained VGG11 network ; C2, 50 conv-filters removed from C1; C3, 1000 additional nodes removed from classifier-layers in C2; C4, 30 conv-filters in conv-layer1 restored to C3. Dynamic transitioning enabled within 5 epochs. Naive strategy is red). (E) A depiction of multiple sparsification paths on the loss landscape from trained network (N1) to networks on the sparse (damage) hyper-plane (N2, N3, N4, N5). The z-axis is network-loss, while the x,y axes are neural net weights. (F) Traversing the geodesic from network-1 trained on MNIST to network-2 trained on Fashion-MNIST along the MNIST functional manifold. The solid blue line is the test-accuracy of the networks (along the geodesic) on MNIST, while the dashed blue-line is the test-accuracy of the networks on Fashion-MNIST, and the red-line is the distance between from network-2. At the point of intersection of the solid and dashed blue line, we find networks that perform at 80% for both the tasks.

115 Additionally, the same geodesic strategy enables us to dynamically shift networks between different
 116 weight configurations (eg from a dense to sparse or vice-versa) while maintaining performance
 117 (Figure 2D). The rapid shifting of networks is relevant for networks on neuromorphic hardware to
 118 ensure that the real-time functionality of the hardware isn't compromised while transitioning between
 119 different power configurations.

120 **Geodesics for alleviating catastrophic forgetting:** In this section, we present our preliminary
 121 results on the geodesic framework formulated to alleviate catastrophic forgetting.

122 We chose two datasets (MNIST and Fashion MNIST) and train our network on both tasks one after
 123 the other. On doing so, we notice that the network initially performs at an accuracy of 99% on MNIST
 124 (after training on MNIST alone), but drops its performance to 10% after being trained with Fashion
 125 MNIST. This is a clear demonstration of the network catastrophically forgetting.

126 To alleviate this issue, we train 2 networks (on MNIST and Fashion MNIST). We find a geodesic path
 127 from w_{t1} to w_{t2} on the functional manifold (MNIST dataset). As we traverse the geodesic, we obtain
 128 a set of networks along the path that perform well on both MNIST and Fashion-MNIST (without
 129 having to be trained on both datasets simultaneously!). In figure-2F, we evaluated the geodesic path
 130 between the 2 trained networks and show that some networks along the geodesic have an accuracy of
 131 80% on both the MNIST and Fashion MNIST dataset, without requiring rigorous training using both
 132 datasets simultaneously.

133 **References**

- 134 [1] German I Parisi et al. “Continual lifelong learning with neural networks: A review”. In: *Neural*
135 *Networks* 113 (2019), pp. 54–71.
- 136 [2] Christos Louizos, Max Welling, and Diederik P Kingma. “Learning Sparse Neural Networks
137 through L_0 Regularization”. In: *arXiv preprint arXiv:1712.01312* (2017).
- 138 [3] Jonathan Frankle and Michael Carbin. “The lottery ticket hypothesis: Finding sparse, trainable
139 neural networks”. In: *arXiv preprint arXiv:1803.03635* (2018).
- 140 [4] Shun-Ichi Amari. “Differential geometry of curved exponential families—curvatures and infor-
141 mation loss”. In: *The Annals of Statistics* (1982), pp. 357–385.
- 142 [5] Detlef H Mache, József Szabados, and Marcel G de Bruin. *Trends and Applications in Con-*
143 *structive Approximation*. Vol. 151. Springer Science & Business Media, 2006.
- 144 [6] Song Han et al. “Learning both weights and connections for efficient neural network”. In:
145 *Advances in neural information processing systems*. 2015, pp. 1135–1143.