
Universality of Winning Tickets: A Renormalization Group Perspective

William T. Redman¹ Tianlong Chen² Zhangyang Wang² Akshunna S. Dogra^{3,4}

Abstract

Foundational work on the Lottery Ticket Hypothesis has suggested an exciting corollary: winning tickets found in the context of one task can be transferred to similar tasks, possibly even across different architectures. This has generated broad interest, but methods to study this universality are lacking. We make use of renormalization group theory, a powerful tool from theoretical physics, to address this need. We find that iterative magnitude pruning, the principal algorithm used for discovering winning tickets, is a renormalization group scheme, and can be viewed as inducing a flow in parameter space. We demonstrate that ResNet-50 models with transferable winning tickets have flows with common properties, as would be expected from the theory. Similar observations are made for BERT models, with evidence that their flows are near fixed points. Additionally, we leverage our framework to study winning tickets transferred across ResNet architectures, observing that smaller models have flows with more uniform properties than larger models, complicating transfer between them.

1. Introduction

The lottery ticket hypothesis (LTH) for deep neural networks (DNNs) proposes that DNNs contain sparse subnetworks that can be trained in isolation and reach performance that is equal to, or better than, that of the full DNN in the same number of training iterations (Frankle & Carbin, 2019; Frankle et al., 2020a). These subnetworks are called winning lottery tickets. The LTH has provided paradigm shifting

insight into the success of dense DNNs, suggesting that the emergence of winning tickets, with increasing DNN size, plays a key role, reminiscent of the maxim “more is different” (Anderson, 1972).

In recent years, researchers have found an intriguing corollary: winning tickets found in the context of one task can be transferred to related tasks (Desai et al., 2019; Mehta, 2019; Morcos et al., 2019; Soelen & Sheppard, 2019; Chen et al., 2020; Gohil et al., 2020; Chen et al., 2021a; Sabatelli et al., 2021), possibly even across architectures (Chen et al., 2021c). Their existence has recently been proven (Burkholz et al., 2022), further supporting the idea that they are a general phenomenon. In addition to having applications of practical interest, these results imply that winning tickets can be used to study how tasks and architectures are “similar”. However, there currently exist few tools with which to study this universality, and there exists no way to know, without directly performing transfer experiments, which previously studied tasks a given winning ticket can be transferred to. These are, in part, due to the fact that there is a general lack of theoretical work on iterative magnitude pruning (IMP) [although see Tanaka et al. (2020) and Elesedy et al. (2021)], the most commonly used method for finding winning tickets.

This is in striking analogy to the state of statistical physics in the early-to-mid-20th century. Empirical evidence suggested that disparate systems, governed by seemingly different underlying physics, exhibited the *same*, universal properties near their phase transitions. While heuristic methods provided insight (Kadanoff, 1966), a full theory from first principles was not realized until the development of the renormalization group (RG) (Wilson, 1971a;b; 1975).

RG theory has not only provided a framework for explaining universal behavior near phase transitions, but also a scheme for grouping systems by that behavior. This classification, introduced via the notion of universality classes, has allowed for a detailed understanding of materials (Tougaard, 1997; Winter & Mours, 1997; Durin & Zapperi, 2000; Bonilla et al., 2010). In addition, it provides knowledge as to what previously classified materials a newly studied substance behaves like (see Appendix A for additional background on RG theory in physics).

The analogy between universality in RG and LTH theories

¹Interdepartmental Graduate Program in Dynamical Neuroscience, University of California, Santa Barbara. ²Department of Electrical and Computer Engineering, University of Texas at Austin. ³Department of Mathematics, Imperial College London. ⁴EPSRC CDT in Mathematics of Random Systems: Analysis, Modelling and Simulation. Correspondence to: William T. Redman <wredman@ucsb.edu>.

again emerges when considering the recent work of Rosenfeld et al. (2020). It was found that, when the density (i.e. the percentage of parameters remaining) of a DNN being pruned via IMP is within a certain range, $d_L < d < d_C$, the DNN’s error scales according to a power-law,

$$e \sim (d_C - d)^{-\gamma} = \Delta d^{-\gamma}. \quad (1)$$

Power-law scaling is well known to emerge in critical phenomena that the RG is used to study. For instance, when the temperature of a classical spin system (e.g. two-dimensional Ising model) is near, but below the critical temperature ($t < t_C$), many observables (e.g. magnetization) exhibit power-law scaling,

$$m \sim (t_C - t)^{-\beta} = \Delta t^{-\beta}. \quad (2)$$

These similarities hint at a more fundamental connection between the universality present in winning tickets and in physical systems (see Fig. 1 and Table 1). Indeed, we show that *IMP is an RG scheme* (i.e. IMP satisfies the properties required to be considered an RG operator). By analyzing large-scale winning ticket transfer experiments (Chen et al., 2020; 2021a), we find that models which allow for successful transfer have similar behavior in their IMP flow (i.e. the way in which the parameters of the DNN change after each round of IMP), in agreement with the theory. In addition, we observe that the IMP flow is static for tickets found on pre-trained transformers (BERT), whereas it is dynamic for winning tickets found on pre-trained and randomly initialized convolutional neural networks (ResNet-50). We further leverage our RG framework to interpret recent results on transferring winning tickets across differing architectures (Chen et al., 2021c), finding that the residual blocks of smaller ResNet architectures are more uniform in their sensitivity to pruning than those of larger ResNet architectures. This is consistent with the experimental result that tickets found on smaller architectures did not transfer as well to larger ones.

Taken together, we believe our work shows that RG theory is not only an appropriate language with which to study IMP (given that IMP *is* an RG scheme), but is also a useful one that provides a new perspective on experimental results. We hope that our work encourages further collaboration between those studying machine learning and those studying the RG and statistical physics.

Table 1. Analogous quantities in RG and IMP theory.

RG	IMP
Spins (s_i)	Unit activations (a_i)
Coupling constants (k_i)	Parameters (θ_i)
Hamiltonian ($\mathcal{H}[\mathbf{s}, \mathbf{k}]$)	Loss function ($\mathcal{L}[\mathbf{a}, \boldsymbol{\theta}]$)

2. Related Work

2.1. Transfer of Winning Tickets

The idea that winning tickets found on one task can be successfully transferred to other tasks was first proposed by Frankle & Carbin (2019) in their original description of the LTH. Ensuing work confirmed this, showing that winning tickets are able to perform well on a variety of tasks beyond the one that they were originally discovered on (Desai et al., 2019; Mehta, 2019; Morcos et al., 2019; Soelen & Sheppard, 2019; Gohil et al., 2020; Sabatelli et al., 2021). This has suggested that such tasks share similar underlying properties, in-line with intuition that they are “related”, and that winning tickets are able to perform universal computations. Winning tickets from DNNs pre-trained on large, complex data sets have shown transferability to a number of “downstream” tasks (Chen et al., 2020; 2021a), further supporting the idea that the solution to a challenging task contains solutions to simpler tasks. Finally, a method for transforming winning tickets to fit different architectures has recently been developed (Chen et al., 2021c), highlighting the existence of common structure in winning tickets across families of DNN architectures.

To-date, the primary approach for studying winning ticket transfer, on a fixed architecture, is to first find winning tickets for tasks A and B . The winning ticket found on task A is then applied to task B (and vice versa), and if the performance is similar to the original winning ticket, then the two tasks are said to admit transfer. While straightforward, such an approach makes it computationally expensive to assess the transferability of many tasks (e.g. tasks A, B, C, \dots), as each combination has to be assessed. Additionally, the components of the winning tickets that are common across tasks are not readily identified.

In this work, we develop an approach based on RG theory to quantitatively characterize winning tickets found on different tasks, and different architectures. This characterization is succinct, and winning tickets with the same properties are transferable. Additionally, it identifies some aspects of transferable winning tickets that are common across tasks and architectures.

2.2. Renormalization Group Theory and Machine Learning

The success of RG theory in studying emergent phenomena in physical systems has led to interest in applying it to machine learning. For instance, Restricted Boltzmann Machines (RBMs) have been proposed to have an exact mapping to the RG (Mehta & Schwab, 2014) [although see comments by Lin et al. (2017) and the response by Schwab & Mehta (2016)]. RBMs trained on data generated from Ising models having been found to perform RG-like coarse

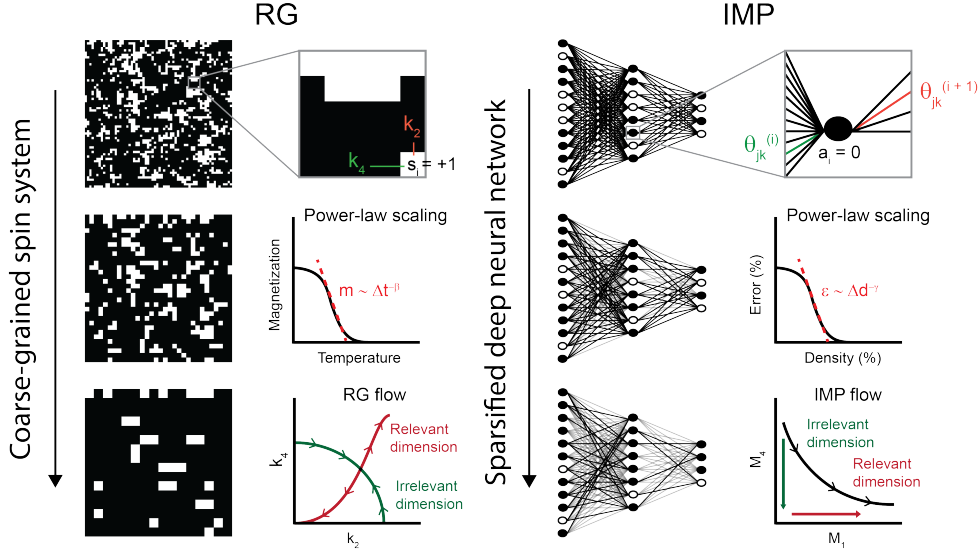


Figure 1. Summary of the similarities between the RG and IMP. Both the RG and IMP are applied iteratively to coarse-grain systems, revealing “relevant” features. Certain observables are known to have regimes where they follow power-law scaling. In the case of the RG, this scaling, and its universality, are associated with properties of the flow that the RG induces in the space of coupling constants (see Sec. 3.1). The nature of the flow that IMP induces has not been previously studied, but none-the-less exists (see Sec. 4). Table 1 outlines the analogous quantities of each theory.

graining operations (Mehta & Schwab, 2014; Koch et al., 2020) [although potentially with different properties from the RG (Iso et al., 2018; Koch et al., 2020)]. Additionally, an RG theory has recently been developed for studying and classifying the output of trained DNNs (Roberts et al., 2021). Lastly, theoretical work has been performed to bridge the language of RG theory to perspectives that are more familiar to machine learning practitioners, such as information theory (Koch-Janusz & Ringel, 2018; Gordon et al., 2021) and principal component analysis (Bradde & Bialek, 2017).

To our knowledge, RG theory has not yet been applied to the study of sparse machine learning. As we show, there is a direct connection between IMP and the RG (Fig. 1 and Table 1), implying that RG theory is an appropriate language with which to study IMP.

3. The Renormalization Group and Iterative Magnitude Pruning

The RG operator, \mathcal{R} , can be viewed as a method for “coarse-graining”. That is, each application of \mathcal{R} replaces local degrees of freedom with a composite of their values¹. An example of this for the two-dimensional (2D) Ising model is shown in the left-hand column of Fig. 1, where neighborhoods of four spins are iteratively replaced by their mode.

¹Note that here we consider the block spin, or real space, RG. The momentum space RG has a different, but related, interpretation.

The formal way to study \mathcal{R} is to consider its action on a given Hamiltonian (i.e. energy function). For classical spin systems (e.g. 2D Ising model), \mathcal{H} has the general form

$$\mathcal{H}(\mathbf{s}, \mathbf{k}) = - \sum_i k_1 s_i - \sum_{\langle i,j \rangle} k_2 s_i s_j - \dots, \quad (3)$$

where the s_i are the spins of the system (e.g. $s_i \in \{-1, +1\}$), $\langle \cdot, \cdot \rangle$ represents sites on the lattice that are nearest neighbors, and the k_i are the strengths of the different coupling constants (e.g. k_2 is the strength of the nearest neighbor coupling).

Due to the fact that coarse-graining amalgamates spins, the spin system resulting from an application of \mathcal{R} can be viewed as equivalent to the original, but with a new set of coupling constants. Therefore, the coarse-grained system has a different Hamiltonian, which is given by

$$\mathcal{R}\mathcal{H}(\mathbf{s}, \mathbf{k}) = \mathcal{H}(\mathbf{s}', \mathcal{T}\mathbf{k}) = \mathcal{H}(\mathbf{s}', \mathbf{k}'), \quad (4)$$

where \mathbf{s}' is the new set of spins and \mathbf{k}' is the new set of couplings determined by the operator $\mathcal{T} : \mathbb{R}^K \rightarrow \mathbb{R}^K$. Here K is the maximum number of couplings considered, usually introduced to keep the considered operators finite.

\mathcal{R} can be applied iteratively, defining a flow through the function space of Hamiltonians via Eq. 4, with an associated flow in the space of coupling constants. The latter is commonly referred to as the RG flow. While \mathcal{T} is often a complicated, non-linear function, it can be linearized near fixed points of the flow (Goldenfeld, 2005). The flow will

grow in the direction of the eigenvectors \mathbf{v}_i with eigenvalues $\lambda_i > 1$ and will shrink along the direction of the eigenvectors \mathbf{v}_i with $\lambda_i < 1$. These eigenvectors are called **relevant** and **irrelevant**, respectively, and are defined in relation to a given fixed point. A schematic of these, and the general RG flow, is presented in Fig. 1. By examining the effect of applying \mathcal{R} to a given system, and its flow in the space of coupling constants, it is possible to find which components of the system are necessary for certain macroscopic behavior (i.e. those that are relevant), and which are not. Systems belonging to a particular universality class have the same relevant directions.

Interestingly, many of the properties discussed above can be analogously found in the context of IMP. To see this, consider a DNN with loss function $\mathcal{L}(\mathbf{a}, \boldsymbol{\theta})$, where \mathbf{a} is the unit activations and $\boldsymbol{\theta}$ is the DNN parameters. Let \mathcal{I} represent a single application of the IMP process. The DNN pruned via \mathcal{I} is given by the relation

$$\mathcal{I}\mathcal{L}(\mathbf{a}, \boldsymbol{\theta}) = \mathcal{L}(\mathbf{a}', \mathcal{T}\boldsymbol{\theta}) = \mathcal{L}(\mathbf{a}', \boldsymbol{\theta}'), \quad (5)$$

where the new set of parameters, $\boldsymbol{\theta}'$, are given by an operator \mathcal{T} . This new set of parameters leads to a new, coarse-grained set of activations, \mathbf{a}' . The similarities between Eqs. 4 and 5 are striking.

We note here that, in the case of IMP, \mathcal{T} is the composition of a masking operator, \mathcal{M} , and a refining operator \mathcal{F} . That is, $\mathcal{T} = \mathcal{F} \circ \mathcal{M} : \mathbb{R}^N \rightarrow \mathbb{R}^N$, where N is the number of parameters in the DNN. \mathcal{M} is defined via the pruning procedure implemented (e.g. magnitude pruning). While we are considering IMP here, because of its connection to the LTH, there are many other pruning procedures [e.g. Hessian based pruning (LeCun et al., 1989)], all of which have their own \mathcal{M} . Similarly, \mathcal{F} is defined via the refinement procedure used, making it dependent on the choice of optimizer and whether or not the DNN parameters are left alone or “rebound” to their value at a previous point during training (Frankle et al., 2020a; Renda et al., 2020).

For n rounds of \mathcal{I} , the resulting DNN is given by

$$\mathcal{I}^n \mathcal{L}(\mathbf{a}^{(0)}, \boldsymbol{\theta}^{(0)}) = \mathcal{L}(\mathbf{a}^{(n-1)}, \mathcal{T}^n \boldsymbol{\theta}^{(0)}) = \mathcal{L}(\mathbf{a}^{(n-1)}, \boldsymbol{\theta}^{(n-1)}). \quad (6)$$

This defines a trajectory in parameter space, $\boldsymbol{\theta}^{(0)} \rightarrow \boldsymbol{\theta}^{(1)} \rightarrow \dots \rightarrow \boldsymbol{\theta}^{(n-1)}$, which we will refer to as the **IMP flow**. A schematic illustration of what the IMP flow might look like is given in Fig. 1 (the axes of which will be explained in Sec. 4). Just as in the case of RG theory, the IMP flow is determined by the eigenvectors of \mathcal{T} , growing or shrinking exponentially by the magnitude of the associated eigenvalues along each direction.

3.1. Connection Between the RG and LTH Frameworks

Before we formally show that IMP is an RG scheme, we discuss how the standard perspective that has emerged from studying the LTH is related to RG theory.

In this picture, the success of winning tickets is attributed to the surviving sparse DNN being able to find a “good” local minimum in the loss landscape (Frankle & Carbin, 2019; Frankle et al., 2020a). This may be possible because the training of DNNs via stochastic gradient descent (SGD) is rapidly confined to a low-dimensional subspace (Gur-Ari et al., 2018; Li et al., 2022), implying that DNNs only “feel” changes to a small number of parameters during much of training. Parameters that are not in this low-dimensional subspace can, therefore, be removed with minimal impact. If a sparse DNN is initialized in this subspace (as late rewinding aims to do), then it may be possible for training to find the same, or related, local minima as the full DNN (Evcı et al., 2020; Maene et al., 2021; Zhang et al., 2021).

In this way, the parameters that do not lie in the low-dimensional subspace are irrelevant, and repeated applications of IMP are expected to remove them. On the other hand, the relevant parameters are those that span this subspace, and their removal changes the local minimum that the sparse DNN converges to. Therefore, certain observable functions of the DNN, such as error, are expected to only be sensitive to the relevant, but not irrelevant, parameters.

If two models share the same low-dimensional subspace, then they should be able to transfer winning tickets. Note that this transferability is highly non-trivial. Indeed, it was only with the development of the LTH, and subsequent experimental work, that this was considered to be a possibility (Frankle & Carbin, 2019; Mehta, 2019; Morcos et al., 2019). It is, in general, difficult to find these subspaces and accurately compare them across experiments.

RG theory provides us with tools for finding the relevant and irrelevant directions, which can be compared across models. In particular, RG theory says that if two models have the same eigenvectors of \mathcal{T} with eigenvalues > 1 , then they have the same relevant parameters. This means that, once the eigenvectors and eigenvalues are computed, we can potentially know whether winning tickets are transferable between two models, without having to run any additional experiments. Note that the converse, two models having distinct relevant directions, does not directly imply that they cannot transfer winning tickets. However, it does imply that they are differently affected by IMP, suggesting different underlying properties of the models.

We end by remarking that numerical algorithms for computing RG critical exponents, such as those that make use of Monte Carlo methods (Ma, 1976; Swendsen, 1979; Shenker & Tobochnik, 1980), are sensitive to the exact distribution

of spins. Therefore, even though the RG theory outlined in this section does not explicitly depend on the amount of coarse-graining used at each application of the RG (although see Eq. 14 for how this impacts the λ), in practice multiple iterations of the RG with small amounts of coarse-graining give the best results. This is similar in spirit to results showing that winning tickets found using multiple rounds of IMP with moderate sparsity outperform tickets found with a single, large sparsity round of IMP (i.e. “one-shot” pruning) (Frankle & Carbin, 2019).

3.2. IMP as an RG Scheme

To make the connection between IMP and RG more precise, we show that IMP fits the definition of an RG scheme. To do this, we consider the projection operator, \mathcal{P} , that is associated with the RG operator. In the case of classical spin systems, the projection operator maps the spins, s_i , to a coarse-grained spin system, s'_I , such that it satisfies

$$\text{Tr}_{\{s_i\}} \mathcal{P}(s_i, s'_I) \exp[\mathcal{H}(s_i, \mathbf{k})] = \exp[\mathcal{H}(s'_I, \mathbf{k}')], \quad (7)$$

where $\text{Tr}_{\{s_i\}}$ is the trace operator over the values that the s_i can take (e.g. ± 1) (Goldenfeld, 2005). For the 2D Ising model, it is standard to take

$$\mathcal{P}(s_i, s'_I) = \prod_I \delta \left[s'_I - \text{sign} \left(\sum_{j \in I} s_j \right) \right], \quad (8)$$

where δ is the Kronecker delta function and $\text{sign}(\cdot)$ is $+1$ if the argument is positive and -1 if it is negative (Goldenfeld, 2005). Eq. 8 formally defines the mapping from $s_i \rightarrow s'_I$ as

$$s'_I = \text{sign} \left(\sum_{j \in I} s_j \right). \quad (9)$$

The projection operator is not unique, but must satisfy the following three properties (Goldenfeld, 2005):

- ❶ $\mathcal{P}(s_i, s'_I) \geq 0$;
- ❷ $\mathcal{P}(s_i, s'_I)$ respects the symmetry of the system;
- ❸ $\sum_{\{s'_I\}} \mathcal{P}(s_i, s'_I) = 1$.

To find a projection operator associated with \mathcal{I} , we start by finding a mapping between the activations of all the units, \mathbf{a} , before and after an application of IMP. This is because \mathbf{a} , and not $\boldsymbol{\theta}$, is the analogous quantity to \mathbf{s} (Table 1). Without loss of generality, we can consider the activation of unit j in layer i as being defined by

$$a_j^{(i)} = h \left[\sum_k g_k(\mathbf{a}, \boldsymbol{\theta}) \right], \quad (10)$$

where h is the activation function (e.g. ReLu, sigmoid) and the g_k are functions that determine how the different parameters and activations of other units affect $a_j^{(i)}$. For instance, in a feedforward DNN, the impact of the bias of unit j in layer i is given by $g_0 = \theta_j^{(i)}$, and the weighted input from the previous layer is given by $g_1 = \sum_{k=1}^{N^{(i-1)}} \theta_{jk}^{(i)} a_k^{(i-1)}$. Here $N^{(i-1)}$ is the number of units in layer $i-1$.

As discussed above, IMP changes the parameters $\boldsymbol{\theta}$ by the operator \mathcal{T} , which is defined by a composition of \mathcal{M} and \mathcal{F} . Therefore, the activation of unit j in layer i , after applying \mathcal{I} , is given by

$$a_j'^{(i)} = h \left[\sum_k g_k(\mathbf{a}, \mathcal{F} \circ \mathcal{M} \boldsymbol{\theta}) \right], \quad (11)$$

and thus, the projection operator associated with \mathcal{I} is

$$\mathcal{P} \left(a_j^{(i)}, a_j'^{(i)} \right) = \prod_{j=1}^N \delta \left\{ a_j'^{(i)} - h \left[\sum_k g_k(\mathbf{a}, \mathcal{F} \circ \mathcal{M} \boldsymbol{\theta}) \right] \right\}. \quad (12)$$

We can easily verify that the projection operator defined by Eq. 12 satisfies all three properties needed for an RG projection operator. First, property #1 is satisfied, as Eq. 12 is a product of Kronecker delta functions. Property #2 is satisfied, as IMP only removes parameters, so the form of $a_j^{(i)}$ given in Eq. 11 and, thus, the loss function, will remain intact until layer collapse (i.e. when all the weights from one layer to another are removed). Finally, for property #3 to be satisfied, we can fix the test and training sample ordering for each epoch (as is done when the random seed is fixed). In such a case, both the mask and refining operators in Eq. 12 are deterministic and $\mathcal{P} \left(a_j^{(i)}, a_j'^{(i)} \right)$ will be unique.

Having found that Eq. 12 satisfies the three properties necessary for an RG projection operator, we have therefore shown that \mathcal{I} meets the criteria for being an RG operator. To the best of our knowledge, this has not been previously identified and provides new insight into why IMP has found success in discovering winning tickets (Frankle & Carbin, 2019; Frankle et al., 2020a) and an uncovering more general DNN phenomena (Frankle et al., 2020b).

4. IMP Eigenfunctions

In order to study the IMP flow, we need to find eigenfunctions of \mathcal{T} from Eq. 5. By looking at the eigenvalues corresponding to those eigenfunctions, we can determine the relevant and irrelevant directions.

For spin systems in statistical physics, the RG flow is studied in the space of coupling constants (Fig. 1). Each coupling constant is usually assumed to be the same for all spins,

greatly reducing the dimensionality of the underlying space. While most DNNs do not set all parameters of a given type to the same value, we can none-the-less estimate the relative ‘‘influence’’ the parameters of a given layer or residual block have on the full DNN, by considering the functions

$$M_i(n) = \frac{\sum_{j=1}^{N^{(i)}} |m_j^{(i)}(n) \cdot \theta_j^{(i)}(n)|}{\sum_{k=1}^N |m_k(n) \cdot \theta_k(n)|}, \quad (13)$$

which describe the percentage of the total remaining parameter magnitude that remains in layer or residual block i after n applications of IMP. For example, $M_1(2) = 0.25$ implies that, after 2 rounds of IMP, layer 1 has 25% of all the remaining weight of the DNN. The numerator in Eq. 13 sums over all $N^{(i)}$ of the parameters restricted to layer or residual block i , $\theta^{(i)}$. Multiplying by the corresponding elements of the pruning mask, $\mathbf{m}^{(i)} \in \{0, 1\}^{N^{(i)}}$, makes sure that only the non-pruned parameters are considered in the sum. The denominator sums over all N parameters of the DNN, multiplied again by the corresponding mask values, making it equal to the total magnitude of all non-pruned parameters.

If the M_i are eigenfunctions of the IMP operator, then they will scale exponentially with respect to the number of iterations \mathcal{T} has been applied. That is, $M_i(n+1) = \mathcal{T}M_i(n) = \lambda_i M_i(n) = \lambda_i^{n+1} M(0)$, where λ_i is the eigenvalue governing its exponential growth/decay. We can find λ_i by the simple inversion, $\lambda_i = M_i(n+1)/M_i(n)$. As long as $M_i(n) \neq 0$, this is well defined and we can approximate λ_i by averaging over the computed values from each pair of data points. When we compute the λ_i of the computer vision models we study in Sec. 5, the standard error of the mean (SEM) is $< 5\%$, supporting the idea that it is appropriate to consider the M_i as eigenfunctions.

Because the degree of coarse-graining at each step (i.e. the amount the DNN gets sparsified, $c \in [0, 1]$, with each round of IMP) affects the magnitude of the eigenvalues, the RG community often considers the quantity σ ,

$$\lambda_i \sim c^{\sigma_i}, \quad (14)$$

which is invariant to the choice of c (Goldenfeld, 2005). Taking $\log_c(\lambda_i)$ gives σ_i . Because we are interested in comparing across models that prune using different, and possibly even variable, c values, we will report σ_i (see Appendix B for details on how we determined c). Note that, relevant directions, which have $\lambda_i > 1$, have $\sigma_i > 0$, and irrelevant directions, which have $\lambda_i < 1$, have $\sigma_i < 0$.

The M_i are not necessarily the only functions one could consider. Given that the sparsity ratio in each layer or residual block appears to play a crucial role in the success of sparse models (Su et al., 2020; Tanaka et al., 2020; Frankle et al., 2021), we also evaluated the functions $P_i(n) = \sum_{j=1}^{N^{(i)}} m_j^{(i)}(n) / \sum_{k=1}^N m_k(n)$, which correspond

to the percent of non-pruned parameters in each layer or residual block. We found that these functions could also be reasonably considered eigenfunctions of \mathcal{T} , and the results we got using them to study the IMP flow were consistent with those we got when using the M_i ². Future work will have to determine what set of functions is best to consider.

5. Results

5.1. Computer Vision Transfer

We start by examining the IMP flow of computer vision models, where the LTH and winning ticket transfer have been most thoroughly characterized. The goal of this analysis is to examine models that are known to allow for the successful transfer of winning tickets and determine whether they have the same relevant and irrelevant directions. If they do, this would support the idea that the theory developed in Sec. 3 properly captures important aspects of sparse machine learning and can be used to identify transferable tickets, without additional experiments.

Because ResNet-50 has four residual blocks, IMP induces a four-dimensional flow in the space of the M_i (Eq. 13). Example 2D slices of these are plotted in Fig. 2 of Appendix C for ResNet-50 trained on CIFAR-10, and CIFAR-100 from random initialization, with 5% rewind. This data comes from experiments performed by Chen et al. (2021a).

Computing the σ_i (using the procedure described in Sec. 4) for ResNet-50 trained on CIFAR-10 and CIFAR-100 from random initialization, with 5% rewind revealed that CIFAR-10 and CIFAR-100 have similar distributions (Table 2, top row). In particular, both have the first three residual blocks being relevant and the last block as being irrelevant. From the theory discussed in Sec. 3, these results imply that both models are in similar universality classes and that winning tickets can be transferred between them. Morcos et al. (2019) indeed confirmed this, finding that winning tickets could be successfully transferred between CIFAR-10 and CIFAR-100 for ResNet-50s.

This similar distribution of σ_i may be due to the fact that the first few residual blocks learn low level statistics of the data (Neyshabur et al., 2020), which scale in specific, non-trivial ways for natural images (Ruderman & Bialek, 1994; Saremi & Sejnowski, 2013). The particular sensitivity of the first two residual blocks (revealed by their large σ_i magnitudes), which are the furthest from the output layer, is in contrast to standard spin systems. In the case of the

²Note that this is for trivial reasons for the pre-trained computer vision and natural language processing models we consider in Secs. 5.1 and 5.2, as the masks for all the downstream tasks come from the same pre-trained model. However, even in cases when we did not consider pre-trained models, P_i gave qualitatively similar σ_i as M_i did.

Table 2. Mean computed σ_i , corresponding to the eigenfunctions M_i , for ResNet-50 trained from random initialization (top row) and pre-trained (PT) on ImageNet. Results from one experiment each.

TASK	σ_1	σ_2	σ_3	σ_4
CIFAR-10	0.20	0.14	0.07	-0.68
CIFAR-100	0.14	0.20	0.04	-0.05
PT CIFAR-10	0.15	0.13	0.02	-0.10
PT CIFAR-100	0.15	0.22	0.10	-0.16

2D Ising model, the RG removes long-range interactions (i.e. couplings between next next-nearest neighbors, etc.), which are typically assumed, for physical reasons, to be weak. Systems with relevant long-range interactions have been found to have unique properties, such as the existence of a phase transition in 1D (Dyson, 1969).

Recent work has found that using DNN parameters from models that have been pre-trained on complex tasks allows for substantial transfer (Chen et al., 2020; 2021a). We therefore examined the effect of pre-training using ImageNet (Huh et al., 2016). We find that, while the winning tickets found using the pre-trained parameter values identify the same relevant and irrelevant residual blocks as with the randomly initialized case, CIFAR-10’s σ_4 is pushed considerably closer to CIFAR-100’s σ_4 (Table 2, bottom row). This suggests that the success transferring winning tickets has found after pre-training comes, in part, from constraining the IMP flow, making the σ_i more similar across downstream tasks. Note that this is in agreement with intuition that pre-training induces biases. We find similar results when computing the σ_i of ResNet-50, pre-trained on ImageNet, with no task specific refining (Table 5 of Appendix D), although differences in experimental set-up prevent direct comparisons between σ_i values.

5.2. Natural Language Processing Transfer

Winning tickets and their transfer have also been studied in the context of natural language processing (NLP) models (Yu et al., 2020; Chen et al., 2020; Prasanna et al., 2020). We computed the σ_i on tickets that were found by applying IMP to pre-trained BERT models on ten downstream NLP tasks (Rajpurkar et al., 2016; Wang et al., 2018), which were known to allow for ticket transfer. Data comes from experiments performed by Chen et al. (2020).

Similar to the computer vision transfer results of Sec. 5.1, we find that tasks that have transferable winning tickets have very similar σ_i (Table 3). That the σ_i , for different tasks, so nearly match each other provides further support for the idea that pre-training confines the IMP flow.

However, there is an important difference between the ResNet-50 and BERT results. Namely, all the BERT σ_i (other than σ_1) are closely clustered around 0, being one to two orders of magnitude smaller than those of ResNet-50. This suggests that, unlike pre-trained ResNet-50, pre-trained BERT has a considerably more static IMP flow, possibly being near a fixed point.

Previous work on transformers applied to computer vision tasks have shown that self-attention heads in consecutive layers can become extremely similar (Touvron et al., 2021; Gong et al., 2021; Zhou et al., 2021). If BERT layers are likewise learning highly overlapping features, then it would be natural to find it at an IMP fixed point. Future work should investigate how methods that have been developed to break this self-similarity (Touvron et al., 2021; Gong et al., 2021; Zhou et al., 2021) affect the IMP eigenvalues and the transferability of tickets between tasks.

5.3. Elastic Lottery Ticket Hypothesis

Recent work on the ‘‘Elastic Lottery Ticket Hypothesis’’ (E-LTH) has extended the notion of transferability by finding it possible to transform winning tickets found on one model to another with a different architecture (Chen et al., 2021c). In particular, it was found that, for architectures in the same family (e.g. ResNets), trained on the same task, it was possible to either squeeze (by removing residual blocks) or stretch (by replicating residual blocks) winning tickets so that they could be transferred.

Three results from this study are especially noteworthy (Chen et al., 2021c):

- ❶ The smallest ResNets considered (ResNet-14 and ResNet-20) had the weakest transferability properties;
- ❷ The more unique residual blocks replicated, the better the performance;
- ❸ Removing the later residual blocks, in the case of shrinking, or replicating the earlier residual blocks, in the case of stretching, led to the best results.

While preliminary hypotheses on the origin of these results were developed by Chen et al. (2021c), involving dynamical systems and unrolled estimation, detailed understanding was left to future work. We examined whether the RG framework developed here could provide further insight into these results.

We computed the σ_i for each ‘‘normal’’ residual block of ResNet-14/20/32/44/56, trained on CIFAR-10 [data from Chen et al. (2021c)]. As discussed in the original paper, the down-sampling residual block in each stage was not transformed between winning tickets, so they were not con-

Table 3. Mean computed σ_i , corresponding to the eigenfunctions M_i , for pre-trained BERT on ten downstream tasks NLP tasks (Rajpurkar et al., 2016; Wang et al., 2018). Results from one experiment each.

TASK	σ_1	σ_2	σ_3	σ_4	σ_5	σ_6	σ_7	σ_8	σ_9	σ_{10}	σ_{11}	σ_{12}
MNLI	-0.054	-0.005	-0.001	0.002	0.017	0.018	0.013	-0.010	-0.004	0.013	0.007	0.004
QQP	-0.056	-0.004	-0.001	0.001	0.017	0.018	0.013	-0.010	-0.003	0.014	0.007	0.005
STS-B	-0.058	-0.004	-0.002	0.001	0.016	0.018	0.013	-0.010	-0.003	0.014	0.008	0.006
WNLI	-0.058	-0.004	-0.002	0.001	0.016	0.018	0.013	-0.010	-0.003	0.015	0.008	0.006
QNLI	-0.057	-0.004	-0.001	0.001	0.017	0.018	0.013	-0.010	-0.003	0.014	0.007	0.005
MRPC	-0.058	-0.004	-0.002	0.001	0.016	0.018	0.013	-0.010	-0.002	0.015	0.008	0.006
RTE	-0.058	-0.004	-0.001	0.001	0.016	0.018	0.014	-0.010	-0.003	0.014	0.008	0.006
SST-2	-0.058	-0.004	-0.001	0.001	0.017	0.018	0.013	-0.010	-0.003	0.014	0.007	0.006
CoLA	-0.058	-0.004	-0.002	0.001	0.016	0.018	0.013	-0.010	-0.002	0.014	0.007	0.006
SQuAD v1.1	-0.052	-0.005	-0.001	0.002	0.017	0.018	0.014	-0.009	-0.003	0.012	0.005	0.004

Table 4. Mean computed σ_i , corresponding to the eigenfunctions M_i , for the first stage of various ResNet architectures. Results are mean across three experiments.

Architecture	σ_i
ResNet-14	0.23
ResNet-20	0.23, 0.26
ResNet-32	-0.04, 0.19, 0.18, 0.26
ResNet-44	-0.03, -0.00, 0.20, 0.19, 0.19, 0.23
ResNet-56	0.06, -0.10, 0.25, 0.20, 0.18, 0.21, 0.11, 0.24

sidered. The σ_i for the residual blocks in the first stage, of all the architectures considered, are presented in Table 4.

We find that, while the smaller ResNets (ResNet-14 and ResNet-20) have only relevant residual blocks, the larger ResNets have at least one non-relevant block, and have a broader distribution of positive σ_i values. Therefore, using ResNet-14 or ResNet-20 as the source ticket necessarily leads to the transformed ticket having residual blocks with larger relative weighting than the target ticket (because of the larger σ_i). This is likely related to the idea that smaller ResNet models are not sufficiently expressive. In addition, we find that replicating few unique residual blocks (i.e. replicating only block 1 or blocks 1 and 2) can lead to an over representation of either relevant or irrelevant blocks. This again leads to a mismatch in the relative weighting of each residual block between the transformed and target tickets, likely making transfer less effective.

Finally, we find that most of the residual blocks with $\sigma_i < 0$ are in the early part of the first stage. Therefore, removing the first several blocks when shrinking a winning ticket (e.g. dropping the first four blocks to go from ResNet-56 to ResNet-32), or replicating the last several blocks when stretching a winning ticket (e.g. adding the last two blocks to go from ResNet-32 to ResNet-44), will again result in the transformed source ticket having a different structure than the target ticket.

Note that for the third stage, the nature of the distribution of σ_i is different (Table 7 of Appendix E). In particular, Stage 3 has only marginal and non-relevant blocks. However, there is consistent evidence that the distributions of σ_i in the second and third stages (Tables 6 and 7 of Appendix E) are similarly related to the three results from the original E-LTH paper (Chen et al., 2021c) we highlighted.

6. Discussion

Inspired by similarities between the current state of sparse machine learning and the state of statistical physics in the early-to-mid-20th century, we found that iterative magnitude pruning (IMP), the principal method used to discover winning tickets, is a renormalization group (RG) scheme. Given that the development of the RG led to a first principled understanding of universal behavior near phase transitions, as well as a way in which to characterize materials by such behavior, we reasoned that viewing IMP from an RG perspective may prove useful when studying the universality of winning tickets (Desai et al., 2019; Mehta, 2019; Morcos et al., 2019; Soelen & Sheppard, 2019; Chen et al., 2020; Gohil et al., 2020; Chen et al., 2021a; Sabatelli et al., 2021) and interpreting the general success IMP has found as a tool for interrogating DNNs (Frankle et al., 2020b).

By viewing IMP as inducing a flow in DNN parameter space (Eq. 6 and Fig. 2 of Appendix C), we showed that ResNet-50 models have trajectories with common properties across different computer vision tasks (Table 2). Winning tickets are known to be transferable between these models (Morcos et al., 2019), suggesting that the theory developed in Sec. 3 captures important properties of sparse machine learning models and may be useful in identifying transferability, without requiring additional experiments.

Winning tickets found on ResNet-50 models pre-trained on ImageNet led to similar identification of relevant and irrelevant residual blocks. However, the values of σ_4 , which had differed most significantly in the case of no pre-training,

were substantially more alike across tasks (Table 2). This suggests that pre-training constrains the IMP flow, enabling better transfer. Studying winning tickets found on natural language processing tasks, using pre-trained BERT models, again identified σ_i distributions that were very similar across downstream tasks (Table 3). Unlike the σ_i of the pre-trained ResNet-50, the σ_i of the pre-trained BERT were tightly clustered around 0, suggesting that the models are near a fixed point of the IMP flow. This is consistent with work that has found that vision transformers can develop self-attention heads that are nearly identical across layers (Touvron et al., 2021; Gong et al., 2021; Zhou et al., 2021).

Finally, we applied the RG framework to recent work that extended the notion of winning ticket universality by transforming tickets between different architectures (Chen et al., 2021c). We found that the distributions of σ_i in the smallest ResNet architectures were more uniform than the distributions of the largest ResNet architectures. Therefore, using them as source tickets leads to a mismatch in structure with target tickets. This is in-line with the observation that the smallest ResNet architectures had the worst transfer properties. We additionally found that the σ_i values offered interpretations of other experimental results found by Chen et al. (2021c).

6.1. Future Directions

A wealth of literature has been developed around the RG, including a number of numerical and theoretical tools that go beyond what we have used here. Given the connection between IMP and the RG, made explicit for the first time in this work, we believe that there is considerable potential for collaboration between the two fields. Beyond the ideas already noted, possible directions of future study include:

- Bringing the RG framework to other, non-IMP sparsification methods;
- Computing the σ_i associated with models that are used outside of computer vision and natural language processing, such as in the context of reinforcement learning and lifelong learning, where the LTH has been studied (Yu et al., 2020; Chen et al., 2021b);
- Computing and classifying systems by their critical exponents (e.g. γ in Eq. 1). We attempted to do this, but we were limited by having only a few independent seeds and a scaling function with multiple free parameters (Rosenfeld et al., 2020) (see Appendix F). Using more advanced methods, such as finite scaling (Goldenfeld, 2005), and increasing the number of independent seeds may allow for better results;
- Developing novel pruning methods. The RG has a complimentary, momentum space perspective, to the

real space perspective presented in this manuscript (Appendix A). This framework allows for coarse-graining at continuous, instead of discrete, scale. Whether this is possible in the context of DNN pruning is an exciting question;

- Identifying, via the RG framework, the minimal density a winning ticket can have and still be transferred to a given task.

Acknowledgements

We would like to thank Jonathan Rosenfeld for his advice on performing the fits described by Eq. 15 (Appendix F).

W.T.R. is partially supported by a UC Chancellor’s Fellowship. T.C. is supported by an IBM PhD fellowship. Z.W. is supported by a US Army Research Office Young Investigator Award (W911NF2010240). A.S.D.’s research is supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation (EP/S023925/1). A.S.D. is funded by the President’s PhD Scholarships at Imperial College London.

References

- Anderson, P. W. More is different. *Science*, 177(4047): 393–396, 1972.
- Bonilla, C. M., Herrero-Albillos, J., Bartolomé, F., García, L. M., Parra-Borderías, M., and Franco, V. Universal behavior for magnetic entropy change in magnetocaloric materials: An analysis on the nature of phase transitions. *Phys. Rev. B*, 81:224424, Jun 2010.
- Bradde, S. and Bialek, W. Pca meets rg. *Journal of Statistical Physics*, 167(3):462–475, 2017.
- Burkholz, R., Laha, N., Mukherjee, R., and Gotovos, A. On the existence of universal lottery tickets. In *International Conference on Learning Representations*, 2022.
- Chen, T., Frankle, J., Chang, S., Liu, S., Zhang, Y., Wang, Z., and Carbin, M. The lottery ticket hypothesis for pre-trained bert networks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 15834–15846. Curran Associates, Inc., 2020.
- Chen, T., Frankle, J., Chang, S., Liu, S., Zhang, Y., Carbin, M., and Wang, Z. The lottery tickets hypothesis for supervised and self-supervised pre-training in computer vision models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16306–16316, June 2021a.

- Chen, T., Zhang, Z., Liu, S., Chang, S., and Wang, Z. Long live the lottery: The existence of winning tickets in lifelong learning. In *International Conference on Learning Representations*, 2021b.
- Chen, X., Cheng, Y., Wang, S., Gan, Z., Liu, J., and Wang, Z. The elastic lottery ticket hypothesis. *arXiv preprint*, arXiv:2103.16547, 2021c.
- Desai, S., Zhan, H., and Aly, A. Evaluating lottery tickets under distributional shifts. *arXiv preprint*, arXiv:1910.12708, 2019.
- Durin, G. and Zapperi, S. Scaling exponents for barkhausen avalanches in polycrystalline and amorphous ferromagnets. *Phys. Rev. Lett.*, 84:4705–4708, May 2000.
- Dyson, F. J. Existence of a phase-transition in a one-dimensional ising ferromagnet. *Communications in Mathematical Physics*, 12(2):91–107, 1969.
- Elesedy, B., Kanade, V., and Teh, Y. W. Lottery tickets in linear models: An analysis of iterative magnitude pruning. *arXiv preprint*, arXiv:2007.08243, 2021.
- Evcı, U., Ioannou, Y. A., Keskin, C., and Dauphin, Y. Gradient flow in sparse neural networks and how lottery tickets win. *arXiv preprint*, arXiv:2010.03533, 2020.
- Feigenbaum, M. J., Kadanoff, L. P., and Shenker, S. J. Quasiperiodicity in dissipative systems: a renormalization group analysis. *Physica D: Nonlinear Phenomena*, 5(2-3):370–386, 1982.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3259–3269. PMLR, 13–18 Jul 2020a.
- Frankle, J., Schwab, D. J., and Morcos, A. S. The early phase of neural network training. In *International Conference on Learning Representations*, 2020b.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations*, 2021.
- Gohil, V., Narayanan, S. D., and Jain, A. [Re] One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers. *ReScience C*, 6(2):#4, May 2020.
- Goldenfeld, N. *Lectures on Phase Transitions and the Renormalization Group*. Levant Books, 2005.
- Gong, C., Wang, D., Li, M., Chandra, V., and Liu, Q. Vision transformers with patch diversification. *arXiv preprint*, arXiv: 2104.12753, 2021.
- Gordon, A., Banerjee, A., Koch-Janusz, M., and Ringel, Z. Relevance in the renormalization group and in information theory. *Phys. Rev. Lett.*, 126:240601, Jun 2021.
- Gur-Ari, G., Roberts, D. A., and Dyer, E. Gradient descent happens in a tiny subspace. *arXiv preprint*, arXiv:1812.04754, 2018.
- Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R., and Sethna, J. P. Universally sloppy parameter sensitivities in systems biology models. *PLoS computational biology*, 3(10):e189, 2007.
- Huh, M., Agrawal, P., and Efros, A. A. What makes imagenet good for transfer learning? *arXiv preprint*, arXiv:1608.08614, 2016.
- Iso, S., Shiba, S., and Yokoo, S. Scale-invariant feature extraction of neural network and renormalization group flow. *Phys. Rev. E*, 97:053304, May 2018.
- Jones, E., Oliphant, T., Peterson, P., et al. SciPy: Open source scientific tools for Python, 2001–. URL <http://www.scipy.org/>.
- Jörg, T. and Katzgraber, H. G. Universality and universal finite-size scaling functions in four-dimensional ising spin glasses. *Physical Review B*, 77(21):214426, 2008.
- Kadanoff, L. P. Scaling laws for ising models near T_c . *Physics Physique Fizika*, 2:263–272, Jun 1966.
- Katzgraber, H. G., Körner, M., and Young, A. Universality in three-dimensional ising spin glasses: A monte carlo study. *Physical Review B*, 73(22):224432, 2006.
- Koch, E. D. M., Koch, R. D. M., and Cheng, L. Is deep learning a renormalization group flow? *IEEE Access*, 8: 106487–106505, 2020.
- Koch-Janusz, M. and Ringel, Z. Mutual information, neural networks and the renormalization group. *Nature Physics*, 14(6):578–582, 2018.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In Touretzky, D. S. (ed.), *Advances in Neural Information Processing Systems 2, [NIPS Conference, Denver, Colorado, USA, November 27-30, 1989]*, pp. 598–605. Morgan Kaufmann, 1989.

- Li, T., Tan, L., Huang, Z., Tao, Q., Liu, Y., and Huang, X. Low dimensional trajectory hypothesis is true: Dnns can be trained in tiny subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2022. doi: 10.1109/TPAMI.2022.3178101.
- Lin, H. W., Tegmark, M., and Rolnick, D. Why does deep and cheap learning work so well? *Journal of Statistical Physics*, 168(6):1223–1247, 2017.
- Ma, S.-k. Renormalization group by monte carlo methods. *Physical Review Letters*, 37(8):461, 1976.
- Machta, B. B., Chachra, R., Transtrum, M. K., and Sethna, J. P. Parameter space compression underlies emergent theories and predictive models. *Science*, 342(6158):604–607, 2013.
- Maene, J., Li, M., and Moens, M.-F. Towards understanding iterative magnitude pruning: Why lottery tickets win. *arXiv preprint*, arXiv:2106.06955, 2021.
- Mehta, P. and Schwab, D. J. An exact mapping between the variational renormalization group and deep learning. *arXiv preprint*, arXiv:1410.3831, 2014.
- Mehta, R. Sparse transfer learning via winning lottery tickets. *arXiv preprint*, arXiv:1905.07785, 2019.
- Mezić, I. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1):309–325, 2005.
- Morcos, A., Yu, H., Paganini, M., and Tian, Y. One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Newman, M. E. and Watts, D. J. Renormalization group analysis of the small-world network model. *Physics Letters A*, 263(4-6):341–346, 1999.
- Neyshabur, B., Sedghi, H., and Zhang, C. What is being transferred in transfer learning? In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 512–523. Curran Associates, Inc., 2020.
- Prasanna, S., Rogers, A., and Rumshisky, A. When BERT Plays the Lottery, All Tickets Are Winning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 3208–3229, Online, November 2020. Association for Computational Linguistics.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics.
- Raju, A., Clement, C. B., Hayden, L. X., Kent-Dobias, J. P., Liarte, D. B., Rocklin, D. Z., and Sethna, J. P. Normal form for renormalization groups. *Phys. Rev. X*, 9:021014, Apr 2019.
- Redman, W. T. Renormalization group as a koopman operator. *Phys. Rev. E*, 101:060104, Jun 2020.
- Renda, A., Frankle, J., and Carbin, M. Comparing rewinding and fine-tuning in neural network pruning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- Reynolds, P. J., Stanley, H., and Klein, W. A real-space renormalization group for site and bond percolation. *Journal of Physics C: Solid State Physics*, 10(8):L167, 1977.
- Roberts, D. A., Yaida, S., and Hanin, B. The principles of deep learning theory. *arXiv preprint*, arXiv:2106.10165, 2021.
- Rosenfeld, J. S., Frankle, J., Carbin, M., and Shavit, N. On the predictability of pruning across scales. *arXiv preprint*, arXiv:2006.10621, 2020.
- Ruderman, D. L. and Bialek, W. Statistics of natural images: Scaling in the woods. *Phys. Rev. Lett.*, 73:814–817, Aug 1994.
- Sabatelli, M., Kestemont, M., and Geurts, P. On the transferability of winning tickets in non-natural image datasets. In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 4: VISAPP*, pp. 59–69. INSTICC, SciTePress, 2021.
- Saremi, S. and Sejnowski, T. J. Hierarchical model of natural images and the origin of scale invariance. *Proceedings of the National Academy of Sciences*, 110(8):3071–3076, 2013.
- Schwab, D. J. and Mehta, P. Comment on “why does deep and cheap learning work so well?” [arxiv:1608.08225], 2016.
- Shenker, S. H. and Tobochnik, J. Monte carlo renormalization-group analysis of the classical heisenberg model in two dimensions. *Physical Review B*, 22(9):4462, 1980.

- Soelen, R. V. and Sheppard, J. W. Using winning lottery tickets in transfer learning for convolutional neural networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2019.
- Su, J., Chen, Y., Cai, T., Wu, T., Gao, R., Wang, L., and Lee, J. D. Sanity-checking pruning methods: Random tickets can win the jackpot. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 20390–20401. Curran Associates, Inc., 2020.
- Swendsen, R. H. Monte carlo renormalization group. *Physical Review Letters*, 42(14):859, 1979.
- Tanaka, H., Kunin, D., Yamins, D. L., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33:6377–6389, 2020.
- Tougaard, S. Universality classes of inelastic electron scattering cross-sections. *Surface and Interface Analysis*, 25(3):137–154, 1997.
- Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., and Jégou, H. Going deeper with image transformers. *arXiv preprint*, arXiv: 2103.17239, 2021.
- Vosk, R. and Altman, E. Many-body localization in one dimension as a dynamical renormalization group fixed point. *Phys. Rev. Lett.*, 110:067204, Feb 2013.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Wang, J.-S. and Swendsen, R. H. Monte carlo renormalization-group study of ising spin glasses. *Physical Review B*, 37(13):7745, 1988.
- Wilson, K. G. Renormalization group and critical phenomena. i. renormalization group and the kadanoff scaling picture. *Phys. Rev. B*, 4:3174–3183, Nov 1971a.
- Wilson, K. G. Renormalization group and critical phenomena. ii. phase-space cell analysis of critical behavior. *Phys. Rev. B*, 4:3184–3205, Nov 1971b.
- Wilson, K. G. The renormalization group: Critical phenomena and the kondo problem. *Rev. Mod. Phys.*, 47: 773–840, Oct 1975.
- Wilson, K. G. and Fisher, M. E. Critical exponents in 3.99 dimensions. *Phys. Rev. Lett.*, 28:240–243, Jan 1972.
- Winter, H. H. and Mours, M. *Rheology of Polymers Near Liquid-Solid Transitions*, pp. 165–234. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.
- Yu, H., Edunov, S., Tian, Y., and Morcos, A. S. Playing the lottery with rewards and multiple languages: lottery tickets in rl and nlp. In *International Conference on Learning Representations*, 2020.
- Zhang, Z., Jin, J., Zhang, Z., Zhou, Y., Zhao, X., Ren, J., Liu, J., Wu, L., Jin, R., and Dou, D. Validating the lottery ticket hypothesis with inertial manifold theory. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- Zhou, D., Kang, B., Jin, X., Yang, L., Lian, X., Jiang, Z., Hou, Q., and Feng, J. Deepvit: Towards deeper vision transformer. *arXiv preprint*, arXiv: 2103.11886, 2021.

A. Brief Background on Renormalization Group Theory in Physics

Phase transitions have long fascinated physicists. These transitions can be between phases of matter (e.g. liquid, solid, gas), but more generally refer to regimes where a continuous, systematic change in one parameter (called the control parameter) leads to a divergence in another parameter (called the order parameter), or its derivative. Examples of this include the ferromagnetic phase transition, where a metal goes from being non-magnetic to magnetic as it is cooled (see Fig. 1).

The point at which the divergence occurs is called the critical point. Specific phenomena occur near critical points, such as power-law scaling of certain observables. As discussed in the main text, $m \sim (t_C - t)^{-\beta}$, where m is the magnetization, t_C is the critical temperature, and t is the temperature of the system (Eq. 2). The exponent present in the power-law scaling is called a critical exponent. A major discovery of 20th century physics was that systems with seemingly different underlying physics exhibited power-law scaling with *the same* critical exponents. For example, the liquid gas phase transition also has an order parameter, the difference in density between gas and liquid, which scales as $|\rho_+ - \rho_-| \sim |t_C - t|^{-\beta}$. The experimentally measured values of β , in the case of the ferromagnetic transition, are nearly identical to experimental values of β , in the case of the liquid gas transition (Goldenfeld, 2005).

Various methods were proposed to understand phases transitions, such as mean field theory and Landau theory, however none could properly compute the critical exponents for two and three dimensional systems. The introduction of the renormalization group (RG) (Wilson, 1971a;b; 1975) solved this problem by providing an explicit, albeit challenging, way in which to compute critical exponents and understand their universality.

The formulation of the RG considered in the main text (i.e. the “block spin” or “real space” RG) considers iterative transformations of physical space (i.e. “coarse-graining”) as a kind of dynamical system, whose linearization around fixed points of the RG flow gives the associated critical exponents. Note that this is an approximation to the true nonlinear nature of the RG operator \mathcal{R} , and corrections to these computed critical exponents can be obtained. A general framework with which to do this via normal form theory has recently been proposed (Raju et al., 2019). Beyond this, Redman (2020) identified the RG as a Koopman operator (Mezić, 2005) (i.e. an infinite dimensional linear operator in function space, directly associated with the nonlinear dynamics in state space). If a suitably invariant subspace could be found, then the RG eigenvalues could in principle be computed as eigenvalues of the Koopman operator, without needing to perform any linearization.

In addition to providing insight into the nature of critical exponents and universality classes, the real space RG has been used to analytically and numerically compute them. For example, in the former case, RG methods have been used to study many-body localization (Vosk & Altman, 2013), percolation (Reynolds et al., 1977), the “route to chaos” (Feigenbaum et al., 1982), and small-world networks (Newman & Watts, 1999). In the latter case, Monte Carlo methods have been adapted to the RG (Ma, 1976; Swendsen, 1979; Shenker & Tobochnik, 1980), making it possible to study complicated spin systems, like the Ising spin glass (Wang & Swendsen, 1988; Katzgraber et al., 2006; Jörg & Katzgraber, 2008).

In addition to the real space RG, there is a related framework for using the RG, called the momentum space RG. As the name suggests, instead of removing degrees of freedom in physical space, this approach removes Fourier modes of the system with large momentum, and thus, high energy. Because of the relationship between energy and wavelength, this in effect integrates out short wavelength (i.e. local) degrees of freedom, again coarse-graining the system in physical space. While the end result is similar, the threshold for what constitutes as “large momentum” can be continuously varied. This allows for a more refined coarse-graining that is done using the real space RG. The momentum space RG plays a major role in quantum field theories, and has also been used to compute critical exponents in dimensions less than 4 (Wilson & Fisher, 1972) and extend principal component analysis (Bradde & Bialek, 2017).

We end by noting that the distinction between relevant and irrelevant interactions have been found in many systems, including biological models and feedforward neural networks (Gutenkunst et al., 2007; Machta et al., 2013).

B. Determining c

In the standard RG theory, the amount of coarse-graining that occurs during each application of the RG operator is given by b . This corresponds to the size, in each dimension, of the neighborhood over which the spins get amalgamated. For instance, in the case of the two-dimensional (2D) Ising model example presented in the left most column of Fig. 1, $b = 2$. Note that one way to interpret b is that it corresponds to how many of the spins in the original system get turned into one spin in the coarse-grained system (for each dimension).

Using this idea, we can compute the amount of coarse-graining that occurs after each round of IMP by considering how the number of weights a given unit has (on average) changes. If a DNN is sparsified by an amount $x \in (0, 1)$ each round of IMP, then, for every set of $1/(1-x)$ weights in the original DNN, there will be 1 weight in the sparsified DNN. Therefore, we should define the amount of coarse-graining, c , to be $c = 1/(1-x)$.

In the case of the computer vision experiments (Secs. 5.1 and 5.3), the models were sparsified 20% each round. Therefore, $x = 0.2$ and $c = 1/(1-0.2) = 1/0.8 = 1.25$. Intuitively, this corresponds to the fact that for every 10 weights in the model that is to be pruned, only 8 would remain after a round of IMP.

In the case of the natural language processing experiments (Sec. 5.2), 10% of the original number of weights are pruned each round. This means that, x changes as a function of the number of rounds the DNN has been pruned. Initially, $x(0) = 0.1$ and $c = 1/(1-0.1) = 1.1\bar{1}$. As $x(n)$ increases with increasing n , so too does c increase. To compute c at IMP round n , we again use $c(n) = 1/(1-x(n))$, where $x(n) = 1/(10-n)$. Thus, when $n = 2$, $c(2) = 1/[1-1/(10-2)] = 8/7 = 1.14\dots$

C. IMP Flow

As discussed in the main text, we chose to analyze the IMP flow by considering its action on the functions $M_i(n)$ (Eq. 13). These describe the proportion of parameter magnitude remaining in layer or residual block i after n rounds of IMP. We found that these functions evolved in an exponential manner, described by the constant λ_i . This made us consider them to be eigenfunctions of \mathcal{T} (Eq. 5), with λ_i being eigenvalues that determine whether residual block is relevant, irrelevant, or marginal (defined in the RG sense – see Sec. 3). In Fig. 2, we visualize the flow projected onto two of the eigenfunctions. While the curves do not start from the same point (implying distinct distributions of weights for different tasks), they evolve at a similar rate (determined by the λ_i – see Table 2 for the reported values of $\log_c \lambda_i$).

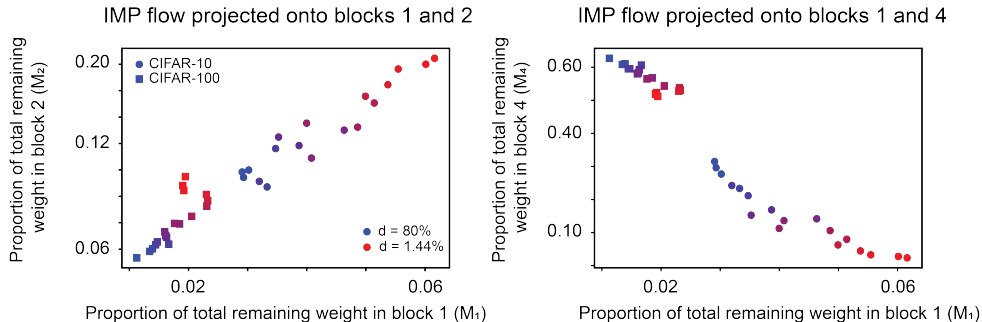


Figure 2. **IMP flow.** Two-dimensional projections of the IMP flow for ResNet-50 trained from random initialization, with 5% rewind, applied to two different tasks: CIFAR-10 (circles), and CIFAR-100 (squares). Blue corresponds to a single application of IMP (density equal to 80%) and red corresponds to 17 applications of IMP (density equal to 1.44%).

D. Additional Computer Vision Transfer Results

In addition to the computer vision transfer results that we present in the main text (Table 2), we also computed the σ_i for a ResNet-50 pre-trained on ImageNet with no fine-tuning (Table 5). The exact way in which IMP was applied was slightly different from the experiments presented in the main text (namely, all layers of each residual block could be pruned, including the first one). Given these differences, we decided to report the results separately from those of the main text. However, we reach a similar conclusion as we did for the other ResNet-50 experiments (Table 2). Namely that the first three residual blocks are relevant, and the last one is irrelevant.

E. Additional E-LTH Results

In the main text on the experimental results surrounding the ‘‘Elastic Lottery Ticket Hypothesis’’ (E-LTH), we reported the σ_i corresponding to the eigenfunctions of the residual blocks in the first stage. The models considered by Chen et al. (2021c) had two other stages, whose σ_i we report here in Tables 6 and 7.

While the distributions of σ_i are different than that of Stage 1, we again find that ResNet-14 and ResNet-20 have residual

Table 5. Mean computed σ_i , corresponding to the eigenfunctions M_i , for ResNet-50 pre-trained on ImageNet with no fine-tuning. Results from one experiment.

NETWORK	σ_1	σ_2	σ_3	σ_4
PRE-TRAINED IMAGENET	0.45	0.30	0.08	-0.24

blocks that are more uniform in their sensitivity to pruning (i.e. have σ_i with more similar values) than the larger models.

Table 6. Computed σ_i , corresponding to the eigenfunctions M_i , for the second stage of various ResNet architectures. Results are mean across three experiments.

ResNet architecture	σ_i
ResNet-14	0.10
ResNet-20	0.08, 0.07
ResNet-32	0.11, 0.11, 0.04, -0.01
ResNet-44	0.15, 0.11, 0.10, 0.09, 0.02, 0.04
ResNet-56	0.10, 0.06, 0.06, 0.06, 0.02, -0.01, -0.06, 0.03

Table 7. Computed σ_i , corresponding to the eigenfunctions M_i , for the third stage of various ResNet architectures. Results are mean across three experiments.

ResNet architecture	σ_i
ResNet-14	-0.21
ResNet-20	-0.16, -0.20
ResNet-32	-0.06, -0.09, -0.13, -0.16
ResNet-44	-0.09, -0.13, -0.16, -0.20, -0.15, -0.08
ResNet-56	-0.11, -0.12, -0.13, -0.15, -0.18, -0.12, -0.04, -0.03

F. Critical Exponents

Following the recent work on scaling laws for IMP developed by [Rosenfeld et al. \(2020\)](#), we fit the error, ϵ , defined as 100% minus the top-1 accuracy, as a function of density, d , defined as percent of weights remaining, via the functional form

$$\begin{aligned} \hat{\epsilon}(d, \epsilon_{np}, \epsilon^\dagger, \gamma, p) &= \epsilon_{np} \left\| \frac{d - ip \left(\frac{\epsilon^\dagger}{\epsilon_{np}} \right)^{\frac{1}{\gamma}}}{d - ip} \right\|^\gamma \\ &= \epsilon_{np} \left[\frac{\left(d^2 + p^2 \left(\frac{\epsilon^\dagger}{\epsilon_{np}} \right)^{2/\gamma} \right)^{\gamma/2}}{(d^2 + p^2)} \right], \end{aligned} \quad (15)$$

where $i = \sqrt{-1}$. Here ϵ_{np} is interpreted as the error associated with not pruning, ϵ^\dagger as the asymptotic error upon maximal pruning, γ as the sensitivity of the combination of network architecture, task, activation function, and optimizer to pruning, and p as controlling how the transition from no change in error to power-law scaling takes place. Importantly, γ can be viewed as a critical exponent under the RG framework (Eq. 1). Because systems in the same universality class have the same critical exponents, we examined whether we could find evidence of universality in the computed value of γ for various different computer vision models.

We found the fits to be sensitive to ϵ_{np} and, in some cases, there was not a clear single value for ϵ_{np} . Therefore, we included ϵ_{np} as a free parameter (which [Rosenfeld et al. \(2020\)](#) did not do), but we tightly bounded the value so as to

Table 8. Critical exponent, γ , for ResNet-50 with either no pre-training and 5% rewinding, or pre-training using ImageNet or simCLR. Data from Chen et al. (2021a). Results are from fitting Eq. 15 to mean of two or three seeds.

Data-set	No pre-train	ImageNet	simCLR	MoCo
CIFAR-10	$\gamma = -0.14$	$\gamma = -0.10$	$\gamma = -0.16$	$\gamma = -0.21$
CIFAR-100	$\gamma = -0.11$	$\gamma = -0.15$	$\gamma = -0.06$	$\gamma = -0.24$
SVHN	$\gamma = -0.07$	$\gamma = -0.06$	$\gamma = -0.03$	$\gamma = -0.29$
ImageNet	-	$\gamma = -0.42$	-	-
simCLR	-	-	$\gamma = -1.01$	-

Error as a function of density scaling

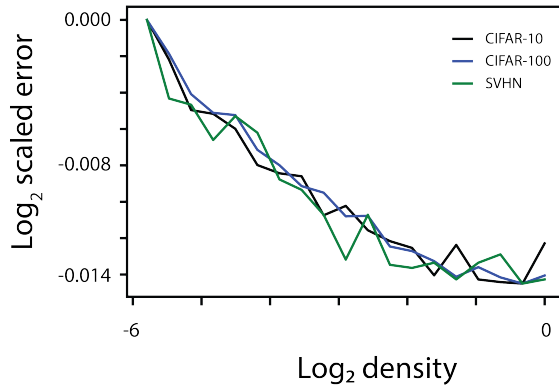


Figure 3. **Error, as a function of density, scaling of CIFAR-10/100 and SVHN.** Overlaying the error, as a function of density, curves for ResNet-50, trained from random initialization, on CIFAR-10/100 or SVHN, with 5% rewind. The overlaid curves shows similar, albeit noisy, behavior.

minimize instability in adding another free parameter. To numerically compute the fits, we used scipy’s curve fitting function: `scipy.optimize.curve_fit` (Jones et al., 2001–).

The computed values of γ for ResNet-50 evaluated on CIFAR-10, CIFAR-100, and SVHN are presented in Table 8. The effect of pre-training using various methods is also given. The γ values vary, sometimes considerably, across tasks and method of pre-training (e.g. CIFAR-10 and SVHN pre-trained via simCLR). However, we note that the data came from only two or three seeds, making the fits susceptible to noise and making the quantification of error in the computed γ difficult. In addition, the fitting function of Eq. 15 has multiple free parameters (including an additional one we added), again making the fit susceptible to noise.

Taking these possible complications into account, we overlaid the different error as a function of density scaling curves (an example of which is given in Fig. 3). We found that they qualitatively matched each other more than the computed γ may have made it seem. We imagine that trying different fitting procedures [especially making use of methods developed in the context of statistical physics, such as finite scaling (Goldenfeld, 2005)], as well as using more independent random seeds, will enable more accurate and robust approximations of γ .