

564 A Related Work

565 The techniques in this paper were inspired by prior research in multiple areas, including neural
566 architecture and activation function search, as well as research on the FIM.

567 **Neural Architecture Search** In neural architecture search [NAS: 12, 53, 56], the goal is to design
568 a neural network architecture automatically. NAS approaches typically focus on optimizing the
569 type and location of the layers and the connections between them, but often use standard activation
570 functions like ReLU. This work is complementary to NAS approaches, because it uses standard
571 architectures but optimizes the design of the activation function.

572 **Zero-cost NAS Proxies** Recently, zero-cost NAS proxies have received increased attention [39,
573 46, 52]. These approaches aim to accelerate neural architecture search by using cheap surrogate
574 calculations in place of expensive full training of architectures. This paper adopts a similar approach,
575 using FIM eigenvalues and activation function outputs to predict which activation functions are likely
576 to be most promising before dedicating resources to evaluating them.

577 **Activation Function Search** Methods for automatically discovering activation functions include
578 reinforcement learning [45], evolutionary computation [3, 5, 6, 33], and gradient-based methods
579 [1, 5, 18, 40, 50]. This paper builds upon existing work, focusing on efficient search and on
580 understanding the properties that make activation functions effective.

581 **Other Uses of the FIM** This paper used FIM eigenvalues to predict the performance of different
582 activation functions. The FIM is an important quantity in machine learning with several uses. One
583 important example is optimal experiment design [13], where experiments are designed to be optimal
584 with respect to some criterion. The criteria vary, but are often functions of the eigenvalues of the FIM,
585 such as the maximum or minimum eigenvalue, or the trace of the FIM (sum of the eigenvalues) or
586 determinant of the FIM (product of the eigenvalues). Instead of choosing one optimality criterion and
587 only considering one summary statistic, this paper keeps all of the eigenvalues of the FIM and learns
588 an optimal distribution experimentally.

589 Past work has also used the eigenvalues of the FIM to determine suitable values of the batch size
590 or learning rate for neural networks [14, 15, 20, 27, 32]. The FIM provides insights to the learning
591 dynamics of SGD [25] and the dynamics of signal propagation at different layers in networks with
592 and without batch normalization layers [23]. The FIM has also been used to develop second-order
593 optimization algorithms for neural networks [19, 35, 36]. Applying it to activation function design is
594 thus a compelling further opportunity.

595 B Activation Function Search Spaces

596 The activation functions in this paper were implemented as computation graphs from the PANGAEA
597 search space [5]. The space includes unary and binary operators, in addition to existing activation
598 functions [7, 11, 28, 41, 45]. This approach allows specifying families of functions in a compact
599 manner. It is thus possible to focus the search on a space where good functions are likely to be
600 located, and also to search it comprehensively.

601 **Benchmark Datasets** The benchmark datasets introduced in Section 2 contain every activation
602 function of the three-node form $\text{binary}(\text{unary}(x), \text{unary}(x))$ using the operators in Table 3. The
603 result is 5,103 activation functions, of which 2,913 are unique. This space is visualized in Figure 4.

604 For Act-Bench-CNN and Act-Bench-ResNet, the accuracies are the median from three runs. For
605 Act-Bench-ViT, the results are from single runs due to computational costs.

606 **New Tasks** The experiments in Section 6 utilized a larger search space. Specifically, it was based
607 on the following four-node computation graphs: $\text{binary}(\text{unary}(\text{unary}(x)), \text{unary}(x))$,
608 $\text{binary}(\text{unary}(x), \text{unary}(\text{unary}(x)))$, $n\text{-ary}(\text{unary}(x), \text{unary}(x), \text{unary}(x))$,
609 $\text{unary}(\text{binary}(\text{unary}(x), \text{unary}(x)))$, and $\text{unary}(\text{unary}(\text{unary}(x)))$. The
610 unary and binary nodes used the operators in Table 3, and the n -ary node used the sum, product,

Table 3: Activation function search spaces were defined through computation graphs consisting of basic unary and binary operators as well as existing activation functions [5].

	Unary		Binary
0	$\text{erf}(x)$	$\text{ReLU}(x)$	$x_1 + x_2$
1	$\text{erfc}(x)$	$\text{ELU}(x)$	$x_1 - x_2$
x	$\text{sinh}(x)$	$\text{SELU}(x)$	$x_1 \cdot x_2$
$-x$	$\text{cosh}(x)$	$\text{Swish}(x)$	x_1/x_2
$ x $	$\text{tanh}(x)$	$\text{Softplus}(x)$	$x_1^{x_2}$
x^{-1}	$\text{arcsinh}(x)$	$\text{Softsign}(x)$	$\max\{x_1, x_2\}$
x^2	$\text{arctan}(x)$	$\text{HardSigmoid}(x)$	$\min\{x_1, x_2\}$
e^x	$e^x - 1$	$\text{bessel_i0e}(x)$	
$\sigma(x)$	$\log(\sigma(x))$	$\text{bessel_i1e}(x)$	

611 maximum, and minimum operators. Together, these computation graphs create a search space with
 612 1,023,516 functions, of which 425,896 are unique. This space is visualized in Figure 9.

613 C Fisher Information Matrix Details

614 In order to calculate the FIM, this paper uses the K-FAC approach [19, 35, 36]. This technique is
 615 summarized in this Appendix, with notation similar to that of Grosse and Martens [19].

616 **Preliminaries** A feedforward neural network maps an input $\mathbf{a}_0 = \mathbf{x}$ to an output $\mathbf{a}_L = f(\mathbf{x}; \boldsymbol{\theta})$
 617 through a series of L layers. Each layer $l \in \{1, \dots, L\}$ is comprised of a weight matrix \mathbf{W}_l , a bias
 618 vector \mathbf{b}_l , and an element-wise activation function ϕ_l . With $\bar{\mathbf{W}}_l = (\mathbf{b}_l \quad \mathbf{W}_l)$ and $\bar{\mathbf{a}}_l = (1 \quad \mathbf{a}_l^\top)^\top$,
 619 each layer implements the transformation

$$\mathbf{s}_l = \bar{\mathbf{W}}_l \bar{\mathbf{a}}_{l-1}, \quad (4)$$

$$\mathbf{a}_l = \phi_l(\mathbf{s}_l). \quad (5)$$

620 Let $\boldsymbol{\theta} = (\text{vec}(\bar{\mathbf{W}}_1)^\top \dots \text{vec}(\bar{\mathbf{W}}_L)^\top)^\top$ represent the vector of all network parameters. Parameter-
 621 ized by $\boldsymbol{\theta}$ and given inputs \mathbf{x} drawn from a training distribution $Q_{\mathbf{x}}$, the neural network defines the
 622 conditional distribution $R_{\mathbf{y}|f(\mathbf{x};\boldsymbol{\theta})}$. The Fisher information matrix associated with this model is

$$\mathbf{F} = \mathbb{E}_{\substack{\mathbf{x} \sim Q_{\mathbf{x}} \\ \mathbf{y} \sim R_{\mathbf{y}|f(\mathbf{x};\boldsymbol{\theta})}}} [\nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{y}, f(\mathbf{x}; \boldsymbol{\theta})) \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{y}, f(\mathbf{x}; \boldsymbol{\theta}))^\top]. \quad (6)$$

623 As usual in deep learning, the loss function $\mathcal{L}(\mathbf{y}, \mathbf{z})$ represents the negative log-likelihood associated
 624 with $R_{\mathbf{y}|f(\mathbf{x};\boldsymbol{\theta})}$ and quantifies the discrepancy between the model’s prediction $\mathbf{z} = f(\mathbf{x}; \boldsymbol{\theta})$ and the
 625 true label \mathbf{y} . The network is trained to minimize the loss by updating its parameters according to the
 626 gradient $\nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{y}, f(\mathbf{x}; \boldsymbol{\theta}))$.

627 **Approximations** For ease of notation, write $D\mathbf{v} = \nabla_{\mathbf{v}} \mathcal{L}(\mathbf{y}, f(\mathbf{x}; \boldsymbol{\theta}))$. Recalling that $\boldsymbol{\theta} =$
 628 $(\text{vec}(\bar{\mathbf{W}}_1)^\top \dots \text{vec}(\bar{\mathbf{W}}_L)^\top)^\top$, the FIM can be expressed as an $L \times L$ block matrix:

$$\mathbf{F} = \begin{pmatrix} \mathbb{E} [\text{vec}(D\bar{\mathbf{W}}_1) \text{vec}(D\bar{\mathbf{W}}_1)^\top] & \dots & \mathbb{E} [\text{vec}(D\bar{\mathbf{W}}_1) \text{vec}(D\bar{\mathbf{W}}_L)^\top] \\ \vdots & \ddots & \vdots \\ \mathbb{E} [\text{vec}(D\bar{\mathbf{W}}_L) \text{vec}(D\bar{\mathbf{W}}_1)^\top] & \dots & \mathbb{E} [\text{vec}(D\bar{\mathbf{W}}_L) \text{vec}(D\bar{\mathbf{W}}_L)^\top] \end{pmatrix}. \quad (7)$$

629 Note that $D\bar{\mathbf{W}}_l = D\mathbf{s}_l \bar{\mathbf{a}}_{l-1}^\top$, and recall that $\text{vec}(\mathbf{u}\mathbf{v}^\top) = \mathbf{v} \otimes \mathbf{u}$. Each block of the FIM can be
 630 written as

$$\mathbf{F}_{i,j} = \mathbb{E} [\text{vec}(D\bar{\mathbf{W}}_i) \text{vec}(D\bar{\mathbf{W}}_j)^\top] \quad (8)$$

$$= \mathbb{E} [\text{vec}(D\mathbf{s}_i \bar{\mathbf{a}}_{i-1}^\top) \text{vec}(D\mathbf{s}_j \bar{\mathbf{a}}_{j-1}^\top)^\top] \quad (9)$$

$$= \mathbb{E} [(\bar{\mathbf{a}}_{i-1} \otimes D\mathbf{s}_i)(\bar{\mathbf{a}}_{j-1} \otimes D\mathbf{s}_j)^\top] \quad (10)$$

$$= \mathbb{E} [(\bar{\mathbf{a}}_{i-1} \otimes D\mathbf{s}_i)(\bar{\mathbf{a}}_{j-1}^\top \otimes D\mathbf{s}_j^\top)] \quad (11)$$

$$= \mathbb{E} [\bar{\mathbf{a}}_{i-1} \bar{\mathbf{a}}_{j-1}^\top \otimes D\mathbf{s}_i D\mathbf{s}_j^\top]. \quad (12)$$




Figure 9: Low-dimensional UMAP representation of the 425,896 function search space. The activation functions are embedded according to their outputs; each point represents a unique function. The larger points represent activation functions that were evaluated during the searches; they are colored according to their validation accuracy. Although the space is vast, the searches require only tens of evaluations to discover good activation functions.

631 Two approximations are necessary in order to make representation of the FIM practical. First, assume
 632 that different layers have uncorrelated weight derivatives. The FIM can then be approximated as
 633 a block diagonal matrix, with $\mathbf{F}_{i,j} = \mathbf{0}$ if $i \neq j$. Second, if one approximates the pre-activation
 634 derivatives $\mathcal{D}\mathbf{s}_l$ and activations $\bar{\mathbf{a}}_{l-1}^\top$ as independent, then the diagonal blocks of the FIM can be
 635 further decomposed into the Kronecker product of two smaller matrices:

$$\mathbf{F}_{l,l} = \mathbb{E} [\bar{\mathbf{a}}_{l-1} \bar{\mathbf{a}}_{l-1}^\top \otimes \mathcal{D}\mathbf{s}_l \mathcal{D}\mathbf{s}_l^\top] \approx \mathbb{E} [\bar{\mathbf{a}}_{l-1} \bar{\mathbf{a}}_{l-1}^\top] \otimes \mathbb{E} [\mathcal{D}\mathbf{s}_l \mathcal{D}\mathbf{s}_l^\top]. \quad (13)$$

636 Let $\mathbf{\Omega}_l = \mathbb{E} [\bar{\mathbf{a}}_l \bar{\mathbf{a}}_l^\top]$ and $\mathbf{\Gamma}_l = \mathbb{E} [\mathcal{D}\mathbf{s}_l \mathcal{D}\mathbf{s}_l^\top]$. The approximate empirical FIM is then written as

$$\hat{\mathbf{F}} = \begin{pmatrix} \mathbf{\Omega}_0 \otimes \mathbf{\Gamma}_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{\Omega}_{L-1} \otimes \mathbf{\Gamma}_L \end{pmatrix}. \quad (14)$$

637 **Layer-Specific Implementation** The above example illustrates FIM approximation for a simple
 638 feedforward network. However, most modern architectures contain several different kinds of layers.
 639 Some layers like pooling, normalization, or dropout layers do not have trainable weights, and therefore
 640 these layers are not included in the FIM [24, 48].

641 Each diagonal entry $\mathbf{\Omega}_{l-1} \otimes \mathbf{\Gamma}_l$ corresponds to one layer with weights. The calculation differs slightly
 642 depending on the layer type, but otherwise the example above can be straightforwardly extended to
 643 more complicated networks. Calculations for three common layer types are presented below.

644 **Dense Layers** For dense layers, the matrices $\mathbf{\Omega}_{l-1}$ and $\mathbf{\Gamma}_l$ can be readily computed with one
 645 forward and backward pass through the network using a mini-batch of data. The eigenvalues are then
 646 computed using standard techniques.

647 **Convolutional Layers** Convolutional layers require special consideration to calculate $\mathbf{\Omega}_{l-1}$ and
 648 $\mathbf{\Gamma}_l$. For a given layer, let M represent the batch size, \mathcal{T} the set of spatial locations (typically a
 649 two-dimensional grid), Δ the set of spatial offsets from the center of the filter, and I and J the number
 650 of output and input maps, respectively. The activations are represented by the $M \times |\mathcal{T}| \times J$ array
 651 \mathbf{A}_{l-1} . The weights are represented by the $I \times |\Delta| \times J$ array \mathbf{W}_l which is interpreted as an $I \times |\Delta| J$
 652 matrix. The expansion operator $\llbracket \cdot \rrbracket$ extracts patches around each spatial location and flattens them
 653 into vectors that become the rows of a matrix: $\llbracket \mathbf{A}_{l-1} \rrbracket$ is a $M|\mathcal{T}| \times J|\Delta|$ matrix.

654 Similar to the feedforward networks, the bias (if used) can be prepended to the weights matrix as
 655 $\bar{\mathbf{W}}_l = (\mathbf{b}_l \quad \mathbf{W}_l)$ and a homogeneous column of ones to the expanded activations as $\llbracket \mathbf{A}_{l-1} \rrbracket_H =$
 656 $(\mathbf{1} \quad \llbracket \mathbf{A}_{l-1} \rrbracket)$. This construction allows the forward pass to be written as

$$\mathbf{S}_l = \llbracket \mathbf{A}_{l-1} \rrbracket_H \bar{\mathbf{W}}_l^\top, \quad (15)$$

$$\mathbf{A}_l = \phi(\mathbf{S}_l), \quad (16)$$

657 from which the factors are computed as

$$\mathbf{\Omega}_l = \mathbb{E} [\llbracket \mathbf{A}_l \rrbracket_H^\top \llbracket \mathbf{A}_l \rrbracket_H], \quad (17)$$

$$\mathbf{\Gamma}_l = \frac{1}{|\mathcal{T}|} \mathbb{E} [\mathcal{D}\mathbf{S}_l^\top \mathcal{D}\mathbf{S}_l]. \quad (18)$$

658 **Depthwise Convolutional Layers** Depthwise convolutional layers utilize separate kernels for each
 659 channel. In this case, $\llbracket \mathbf{A}_{l-1} \rrbracket$ is a $M|\mathcal{T}|J \times |\Delta|$ matrix. Otherwise, the factors $\mathbf{\Omega}_{l-1}$ and $\mathbf{\Gamma}_l$ are
 660 calculated in the same way as they are for standard convolutional layers.

661 **Eigenvalue Calculation** Because $\hat{\mathbf{F}}$ is a block-diagonal matrix, its eigenvalues are simply the
 662 combined eigenvalues of each block: $\lambda(\hat{\mathbf{F}}) = \{\lambda(\hat{\mathbf{F}}_l)\}_{l=1}^L$. The eigenvalue calculation for one block
 663 $\hat{\mathbf{F}}_l = \mathbf{\Omega}_{l-1} \otimes \mathbf{\Gamma}_l$ is further simplified by first computing the eigenvalues $\lambda(\mathbf{\Omega}_{l-1})$ and $\lambda(\mathbf{\Gamma}_l)$ for each
 664 Kronecker factor separately and then returning all pairwise products from the two sets of eigenvalues.
 665 For numerical stability, the eigenvalues can first be log-scaled and then all pairwise sums from the
 666 two sets are returned. Calculating the eigenvalues requires one forward and backward pass through
 667 the network with a mini-batch of data. The computational cost is therefore relatively cheap, especially
 668 compared with the cost of fully training a network from scratch.

669 It is possible for the FIM eigenvalues to be invalid. For example, if the forward propagated activations
670 or backward propagated gradients explode or vanish, then the diagonal entries $\Omega_{l-1} \otimes \Gamma_l$ may be
671 undefined. Such invalid values result from activation functions that are unstable. Therefore, invalid
672 FIM eigenvalues provide a good way to filter out bad activation functions.

673 D Features and Surrogate Details

674 This section describes how the activation function features were implemented and how the surrogate
675 was constructed.

676 **Calculating FIM Eigenvalues** The FIM eigenvalues were calculated for each activation function
677 as discussed in Section 3. The eigenvalues were log-scaled for numerical stability. By definition, the
678 number of eigenvalues is the same as the number of weights in the neural network. To save space,
679 the eigenvalues were binned to histograms. For a layer l with $|\theta_l|$ weights, $\lfloor |\theta_l|/100 \rfloor$ equally sized
680 bins from -100 to 100 were used. One histogram was computed for each layer in a network, and
681 all of the histograms were concatenated together into a single feature vector for a given activation
682 function. In this manner, the total dimensionality was 13,692 for All-CNN-C, 16,500 for ResNet-56,
683 and 11,013 for MobileViTv2-0.5.

684 **Calculating Activation Function Outputs** The activation function outputs $y = f(x)$ were calcu-
685 lated for each activation function f by sampling $n = 1,000$ values $x \sim \mathcal{N}(0, 1)$ and truncating to the
686 range $[-5, 5]$. The same random inputs were used for all activation functions.

687 **Per-Layer FIM Eigenvalues** In Figure 5, the eigenvalues for the entire network are shown for
688 completeness. However, the UMAP representations shown in Figure 4 were produced by keeping the
689 eigenvalues at each layer separate and computing a weighted distance between them (according to
690 Equation 2). As pointed out in the main text, FIM eigenvalues are informative but noisy features. In
691 preliminary experiments, keeping the eigenvalues separate at each layer reduced some of this noise,
692 resulting in a more informative Figure 4 and consequently improving the performance of the search
693 algorithms.

694 **FIM Eigenvalue Features** Preliminary experiments aimed to predict activation function perfor-
695 mance using common features in the literature, including maximum eigenvalue, minimum eigenvalue,
696 sum of the eigenvalues, and product of the eigenvalues [13]. More recently proposed features, such
697 as (second moment) / (first moment)², were also considered [44]. Ultimately, learning the relevant
698 features from the entire eigenvalue distribution was found to be the most flexible and powerful
699 approach.

700 **UMAP Settings** UMAP exposes a number of parameters that can be used to customize its be-
701 havior [37]. The `metric` parameter determines how distances are computed between points, the
702 `n_neighbors` parameter adjusts the tradeoff between the local and global structure of the data, and
703 the `min_dist` parameter controls the minimum distance between points in the embedding space.

704 The plots in Figure 4 were produced by computing the distances between FIM eigenvalues and acti-
705 vation function outputs. For the FIM eigenvalues UMAP(`metric='manhattan', n_neighbors=3,`
706 `min_dist=0.1`) was used, and for the activation function outputs UMAP(`metric='euclidean',`
707 `n_neighbors=15, min_dist=0.1`) was used. The distance metrics were chosen to implement
708 Equations 2 and 3.

709 In preliminary experiments, decreasing `n_neighbors` from the default of 15 down to 3 for the
710 FIM eigenvalues qualitatively improved the embedding for the combined features. The com-
711 bined features were visualized with a union model, i.e. `umap_combined = umap_fim_eigs +`
712 `umap_fn_outputs` [37].

713 E Experiment Details

714 This section specifies the details for the experiments in the main text of the paper. Several variations
715 to the approach presented in the main text were also evaluated in preliminary experiments. The




Figure 10: UMAP projections of FIM eigenvalues using the default hyperparameter of $n_neighbors=15$. The embedding is informative but also noisy. Using $n_neighbors=3$, as shown in the main text, improved performance.

716 approach turned out to be robust to most of them, but the results also justify the choices used for the
 717 main experiments.

718 **Training Details** For CIFAR-10 and CIFAR-100, balanced validation sets were created by sampling
 719 5,000 images from the training set. Full training details and hyperparameters are listed in Table 4.

720 **Search Implementation** In order to predict performance for an unevaluated activation function, the
 721 function outputs and FIM eigenvalues must first be computed. Thus, the searches in Section 6 were
 722 implemented in three steps. First, activation function outputs for all 425,896 activation functions in
 723 the search space were calculated. This computation is inexpensive and easily parallelizable. Second,
 724 eight workers operated in parallel to sample activation functions uniformly at random from the search
 725 space and calculate their FIM eigenvalues. Third, once the number of activation functions with
 726 FIM eigenvalues calculated reached 5,000, seven of the workers began the search by evaluating the
 727 functions with the highest predicted performance. The eighth worker continued calculating FIM
 728 eigenvalues for new functions so that their performance could be predicted during the search. This
 729 setup allowed taking best advantage of the available compute for the regression-type search methods.

730 The experiments on ImageNet required substantially more compute than the experiments on CIFAR-
 731 100. For this reason, all eight workers evaluated activation functions once the number of functions
 732 with FIM eigenvalues reached 7,000.

733 Computing FIM eigenvalues took approximately 26 seconds, 84 seconds, and 37 seconds per
 734 activation function for All-CNN-C, ResNet-56, and MobileViTv2-0.5, respectively. This cost is not
 735 trivial, but it is well worth it, as the experiments in the main paper show.

736 **Unique Activation Functions** Different computation graphs can result in the same activation
 737 function (e.g. $\max\{x, 0\}$ and $\max\{0, x\}$). In the benchmark dataset and in the larger search space
 738 of Section 6, repeated activation functions were filtered out. 1,000 inputs were sampled $\mathcal{N}(0, 1)$
 739 and truncated to $[-5, 5]$. Two activation functions were considered the same if their outputs were
 740 identical.

741 **Improving the Combined UMAP Projection** Figure 10 displays a projection of FIM eigenvalues
 742 using default UMAP hyperparameters. The plots show the eigenvalues organized in multiple distinct
 743 one-dimensional manifolds. Again, FIM eigenvalues are noisy features; there are some clusters of
 744 activation functions achieving similar performance, but there are also regions where performance
 745 varies widely. As mentioned in the main text, this issue was addressed by reducing the UMAP
 746 parameter $n_neighbors$ to 3. This change reduced the connectivity of the low-dimensional FIM
 747 eigenvalue representation, resulting in a space with many distinct clusters (as seen in Figure 4).

748 On its own, this setting did not improve the search on the benchmark datasets. However, it did
 749 improve performance when the FIM eigenvalues were combined with activation function outputs
 750 (as was discussed in Section 4). The reason is that the UMAP model for the activation function
 751 outputs did not decrease $n_neighbors$, and so the combined UMAP model relied more on the

Table 4: Training details and hyperparameter values used in the experiments.

All-CNN-C on CIFAR-10 and CIFAR-100	
Batch Size	128
Dropout	0.5
Epochs	25 for Act-Bench-CNN and search (Figure 7), 50 for full evaluation (Table 1)
Image Size	32×32
Learning Rate	Linear warmup to 0.1 for five epochs, then linear decay
Mean/Std. Normalization	Yes
Momentum	0.9
Optimizer	SGD
Random Crops	32×32 crops of images padded with four pixels on all sides
Random Flips	Yes
Weight Decay	$1e^{-4}$
Weight Initialization	AutoInit [4]
ResNet-56 on CIFAR-10 and CIFAR-100	
Batch Size	128
Dropout	0.0
Epochs	25 for Act-Bench-ResNet and search (Figure 7), 50 for full evaluation (Table 1)
Image Size	32×32
Learning Rate	Linear warmup to 0.1 for five epochs, then linear decay
Mean/Std. Normalization	No
Momentum	0.9
Optimizer	SGD
Random Crops	32×32 crops of images padded with five pixels on all sides
Random Flips	Yes
Weight Decay	$1e^{-4}$
Weight Initialization	AutoInit [4]
MobileViTv2-0.5 on Imagenette and ImageNet	
Batch Size	256
CutMix Alpha [54]	1.0
Epochs	105
Evaluation Center Crop	95%
Image Size	160×160
Learning Rate	Linear warmup from $1e^{-4}$ to $4e^{-3}$ for five epochs, then cosine decay to $1e^{-6}$
Mixup Alpha [55]	0.1
Optimizer	AdamW [34]
RandAugment [9]	Magnitude six, applied twice
Random Resized Crop [49]	Minimum 8% of the original image
Weight Decay	$0.02 \times$ current learning rate
ResNet-50 on ImageNet	
Batch Size	256
CutMix Alpha [54]	1.0
Epochs	105
Evaluation Center Crop	95%
Image Size	160×160
Learning Rate	Linear warmup from $1e^{-4}$ to $2e^{-3}$ for five epochs, then cosine decay to $1e^{-6}$
Mixup Alpha [55]	0.1
Optimizer	AdamW [34]
RandAugment [9]	Magnitude six, applied twice
Random Resized Crop [49]	Minimum 8% of the original image
Weight Decay	$0.02 \times$ current learning rate
Weight Initialization	AutoInit [4]

752 activation function outputs than it did on the FIM eigenvalues. As Figure 4 shows, the activation
753 function outputs are reliable but sometimes project good activation functions to distinct regions in
754 the search space. Introducing extra connectivity into the fuzzy topological representation via the
755 FIM eigenvalues was sufficient to address this issue, bringing good activation functions to common
756 regions of the space.

757 **Increasing the Dimension of the UMAP Projections** The UMAP plots show two-dimensional
758 projections of FIM eigenvalues and activation function outputs. Regression algorithms were also
759 trained on five and 10-dimensional projections. These runs resulted in comparable or worse perfor-
760 mance. Therefore, the two-dimensional projections were selected in the paper for simplicity and for
761 consistency between the algorithm implementation and figure visualizations.

762 **Gaussian Process Regression** As an alternative search method, Gaussian process regression (GPR)
763 was evaluated in activation function search. Several different acquisition mechanisms were used,
764 including expected improvement, probability of improvement, maximum predicted value, and upper
765 confidence bound. The approach worked well, but the results were inconsistent across the different
766 acquisition mechanisms. GPR was also more expensive to run compared to the algorithms in the
767 main text (KNR, RFR, SVR), and so those algorithms were used instead for simplicity and efficiency.

768 **Adjusting k in KNR** The initial experiments with the KNR algorithm used $k = 3$. Experimenting
769 with $k = \{1, 5, 8\}$ did not reliably improve performance, so $k = 3$ was kept.

770 **Uniformly Spaced Inputs for Activation Function Outputs** In an alternative implementation,
771 equally spaced inputs from -5 to 5 were given to the activation functions instead of normally
772 distributed inputs. This variation did not noticeably change the quality of the embeddings nor the
773 performance of the search algorithms. Therefore, normal inputs were used for consistency with
774 Equation 3. Figure 3 is the only exception; it used 80 inputs equally spaced from -5 to 5 and increased
775 the UMAP parameter `min_dist` to 0.5 . These settings improved the quality of the reconstructed
776 activation functions in the plot.

777 F Future Work

778 This paper demonstrated that FIM eigenvalues and activation function outputs are efficient and
779 reliable features that can predict performance of activation functions accurately. This finding enabled
780 discovering better activation functions for various tasks, improving the state of the art in machine
781 learning. Because the technique is efficient, it was possible to scale it up to large datasets such as
782 ImageNet. These discoveries inspire several avenues for future research, discussed below.

783 **New Search Spaces** The PANGAEA search space was used in this paper because it is known to
784 work well for deep architectures [5]. In the future it will be interesting to explore search spaces with
785 different unary, binary, and n -ary operators. Beyond computation graphs, it may also be possible to
786 apply techniques in this paper to optimize continuous vector representations of activation functions
787 [1, 40].

788 **Exploration vs. Exploitation** The KNR approach was utilized to search for new activation
789 functions because it performed well on the benchmark datasets (Section 5). In the future, it will
790 be interesting to consider other algorithms and analyze their tradeoffs between exploration and
791 exploitation. For example, in a resource-constrained environment where improvement is needed
792 quickly, a more exploitative approach could be used to find an improved activation function in a short
793 time. On the other hand, if substantial compute is available, an approach that focuses on exploration
794 could be used to discover activation functions that perform well but are maximally different from
795 functions used in modern architectures (Figure 8b). Novelty search [30] could serve as a suitable
796 approach, and such discoveries could further understanding of how neural networks utilize different
797 kinds of activation functions to learn.

798 **Optimizing Multiple Activation Functions** In a typical neural network design, the same activation
799 function is used throughout the network. However, recent work has shown that it may be beneficial
800 to have different activation functions at different locations, and further, that it may be useful to

801 have different activation functions in the early and late stages of training [5]. Indeed, many hybrid
802 architectures use Swish in convolutional layers and ReLU in attention layers [38]. Unfortunately, it is
803 difficult to design these strategies manually, and so practitioners often use a single activation function
804 for simplicity.

805 The techniques proposed in this paper may provide an avenue toward optimizing multiple activation
806 functions in tandem. For example, the features for multiple candidate activation functions could
807 be concatenated into a single feature vector, and this vector could be projected with UMAP to a
808 low-dimensional space where performance prediction is more straightforward.

809 **Optimizing Parametric Activation Functions** Parametric activation functions have learnable
810 parameters that allow them to refine their shape via gradient descent. In some tasks, this extra
811 flexibility results in better performance over fixed activation functions [5]. The techniques introduced
812 in this paper can be readily extended to optimizing the design of parametric activation functions as
813 well. Because the surrogate considers the state of the network and activation function at initialization,
814 it is possible to predict the performance by treating the activation function parameters as fixed to their
815 initial values.

816 However, it may be possible to extend this idea further. Because the activation function parameters are
817 implemented as neural network weights, each parameter will have a corresponding FIM eigenvalue.
818 These extra eigenvalues will provide the surrogate with additional information that may help predict
819 the performance more accurately.

820 For simplicity, current parametric activation functions usually initialize their parameters either to be
821 1.0 or to approximate some existing activation function, and the initialization is usually the same
822 throughout the network. This method is likely suboptimal; the surrogate introduced in this paper
823 could provide a smarter approach. By adjusting the initial parameter values and observing the change
824 in predicted performance, the surrogate can be used to find better initializations, including different
825 ones at different layers in the network. This contribution could make parametric activation functions
826 even more powerful.

827 **Optimizing Other Aspects of Neural Network Design** By fixing the neural network architecture
828 and varying the activation function, this paper showed that it is possible to use FIM eigenvalues
829 to infer future performance. As the FIM is a fundamental quantity in machine learning, it may
830 be possible to apply a similar strategy to optimize other aspects of neural network design, such as
831 normalization layers, loss functions, or data augmentation strategies [8, 16, 17, 33]. If a meaningful
832 distance metric between such objects can be defined, then UMAP could be used to map them to a
833 low-dimensional space where performance prediction is much simpler.

834 Similarly, one could use the FIM eigenvalues to optimize alternate objectives beyond accuracy.
835 Robustness is a particularly interesting objective, because the FIM can be used to describe a neural
836 network’s robustness to small parameter perturbations. Other objectives, such as interpretability, fair-
837 ness, or inference cost, could also be considered. For example, one could consider a multidimensional
838 regression approach where instead of just predicting accuracy, the surrogate would predict each of
839 these quantities separately. Such a method could present the user with a Pareto front of activation
840 functions involving tradeoffs between these quantities.

841 **Reverse Engineering Activation Functions** UMAP was used to project activation functions to a
842 low-dimensional space, and regression algorithms to predict the performance of activation functions
843 in this space, i.e. to serve as a fitness function for the search. However, it is possible that there is no
844 activation function that maps to the optimum of this fitness landscape. Indeed, because such search
845 spaces are finite, the activation functions do not completely fill them. For example, there are empty
846 regions in Figure 4, corresponding to activation functions outside of the predefined search space.

847 What should be done if an empty region of the embedding space has a higher predicted fitness
848 than any of the candidate activation functions? In the paper, these regions were simply ignored,
849 and the activation function with the highest predicted fitness was used. However, in the future,
850 it may be possible to create activation functions that map to these empty spaces, and in so doing
851 improve performance. One approach could be based on inverse transforms: Given a coordinate in the
852 low-dimensional embedding space, UMAP can apply an inverse transform and return an object that
853 would have mapped to those coordinates. This technique was already used for visualization in Figure

854 **3** Using this approach, UMAP could generate a hypothetical desired FIM eigenvalue distribution, or
855 a list of activation function outputs.

856 There are two challenges to this approach. First, because UMAP is a dimensionality-reduction
857 algorithm, different activation functions can map to the same location in the embedding space. Thus,
858 the mapping from embedding space back to activation functions is not well defined. Second, even if
859 UMAP prescribes a FIM eigenvalue distribution that is predicted to result in good performance, it
860 may be difficult to manually design an activation function to satisfy that distribution.

861 However, a generated list of prescribed activation function outputs is already a good start. From this
862 list, it is possible to construct an activation function that interpolates through these points, either in
863 a piecewise linear fashion, with splines, or using some other standard technique. Even without the
864 corresponding FIM eigenvalues, such an approach could potentially improve the efficiency of novel
865 activation function discovery, and lead to better designs for activation functions in the future.

866 **G Compute Infrastructure**

867 The experiments in this paper were implemented using an AWS `g5.48xlarge` instance with eight
868 NVIDIA A10G GPUs. The total compute cost for the search experiments in Section **6** was 14.49
869 GPU-hours for All-CNN-C on CIFAR-100, 21.67 GPU-hours for ResNet-56 on CIFAR-100, and
870 196.25 GPU-days for MobileViTv2-0.5 on ImageNet. This cost includes the time to train the eight
871 baseline activation functions and then to evaluate 100 additional functions. The instance ran in Oregon
872 (`us-west-2`) and was powered by renewable energy, so the experiments for this paper contributed
873 no carbon emissions.