
Towards falsifiable interpretability research

Matthew L. Leavitt^{*†}
Facebook AI Research
Menlo Park, California
ito@fb.com

Ari S. Morcos^{*}
Facebook AI Research
Menlo Park, California
arimorcos@fb.com

Abstract

Methods for understanding the decisions of and mechanisms underlying deep neural networks (DNNs) typically rely on building intuition by emphasizing sensory or semantic features of individual examples. For instance, methods aim to visualize the components of an input which are “important” to a network’s decision, or to measure the semantic properties of single neurons. Here, we argue that interpretability research suffers from an over-reliance on intuition-based approaches that risk—and in some cases have caused—illusory progress and misleading conclusions. We identify a set of limitations that we argue impede meaningful progress in interpretability research, and examine two popular classes of interpretability methods—saliency and single-neuron-based approaches—that serve as case studies for how overreliance on intuition and lack of falsifiability can undermine interpretability research. To address these concerns, we propose a strategy to address these impediments in the form of a framework for strong, falsifiable interpretability research. We encourage researchers to use their intuitions as a starting point to develop and test concrete, falsifiable hypotheses, and hope that our framework yields robust, evidence-based interpretability methods that generate meaningful advances in our understanding of DNNs.

1 Introduction

The discovery of instances which confirm a theory means very little if we have not tried, and failed, to discover refutations. For if we are uncritical we shall always find what we want.

-Karl Popper

The Poverty of Historicism [Popper, 2002]

The complexity and high-dimensionality of deep neural networks (DNNs) makes them challenging to understand. One approach for understanding such systems is to extrapolate from intuitions about simpler, lower-dimensional systems. Mapping abstract quantities into concrete sensory forms, such as by visualization, can also be a powerful cognitive aid. However, looks are often deceiving and unchecked extrapolation can easily lead one astray.

The Gaussian (normal) distribution provides a vivid example of the pitfalls of extrapolating intuitions about low-dimensional systems to higher dimensions. A one-dimensional Gaussian distribution is perhaps the most familiar data distribution. Its iconic bell-like silhouette is ubiquitous across textbooks, technical publications, and lay reports of polls and statistics. Its one-dimensional shape extrapolates intuitively into two-dimensions: similar to an idealized mountain on a relief map. One might then expect that the “shape” of a Gaussian remains consistent as its dimensionality increases, with the majority of its mass lying close to the origin, with mass falling off like a bell

^{*}Equal contribution

[†]Work performed as part of the Facebook AI Residency program

curve as distance from the origin increases. Counterintuitively, however, most of the mass of a high-dimensional Gaussian is concentrated in a thin band at a constant distance from the origin. Thus the closest analogy for the shape of a high-dimensional Gaussian distribution is that of a soap bubble, or spherical shell: a fundamentally different structure than that of low-dimensional Gaussians [Blum et al., 2020, Huszár, 2017, Vershynin, 2018].

Intuition is the foundation upon which comprehensive understanding is built. But, as the example of the Gaussian distribution shows, unverified intuition can be misleading. The basis of one’s intuition must be verified at every iteration of extrapolation. While intuition helps reveal what the important questions are, those questions need to be addressed with strong, falsifiable hypotheses. This applies to all science, including deep learning.

The nonlinearity of DNNs renders them resistant to being described and understood using analytical methods, leading researchers to rely instead on empirical approaches. A breadth of such approaches have been developed, typically referred to using the umbrella terms “interpretability” or “explainability” [Yosinski et al., 2015, Bau et al., 2017, Olah et al., 2018, Donnelly and Roegiest, 2019, Radford et al., 2017, Erhan et al., 2009, Zeiler and Fergus, 2014, Karpathy et al., 2016, Amjad et al., 2018, Dhamdhare et al., 2019, Olah et al., 2020, 2017, Zhou et al., 2015, 2019, Gunning, 2017, Samek et al., 2017, Došilović et al., 2018, Lakkaraju et al., 2017, Vellido et al., 2012, Cadamuro et al., 2016, Montavon et al., 2018, Guidotti et al., 2018, Lipton, 2018, Hohman et al., 2019, Roscher et al., 2020]. Many of these methods rely on building intuition by emphasizing sensory or semantic features of individual examples, for instance by visualizing the components of an input which are “important” to a network’s decision [Springenberg et al., 2015, Selvaraju et al., 2017, Fong and Vedaldi, 2017, Baehrens et al., 2010, Simonyan et al., 2014, Zeiler and Fergus, 2014, Shrikumar et al., 2017, Sundararajan et al., 2017, Smilkov et al., 2017, Dabkowski and Gal, 2017, Ancona et al., 2018], or characterizing semantic properties of individual DNN neurons [Amjad et al., 2018, Meyes et al., 2019, Morcos et al., 2018b, Zhou et al., 2015, 2018, Bau et al., 2017, Karpathy et al., 2016, Na et al., 2019, Radford et al., 2017, Rafegas et al., 2019]. However, these practices make two assumptions that we argue are too frequently unverified in interpretability research: the assumption that a method faithfully represents meaningful properties of the model it is applied to, and the assumption that individual instances or examples faithfully represent the phenomena they are being used to illustrate. And their lack of verification is symptomatic of a more general lack of empirical rigour that impedes interpretability research.

In this position paper we discuss a set of impediments to robust interpretability research: a lack of concrete, falsifiable hypotheses; the abundance and seductive allure of visualization; under-verification of the principles upon which interpretability methods are based; a lack of quantifiable approaches; and a lack user research for human verification (Section 2). We next review two case studies in the literature which exemplify the pitfalls of pursuing un(der)verified intuitions in deep learning, and perform a “post-mortem” to identify and explain the impediments that led to misleading conclusions (Sections 3 and 4). We then articulate a strategy to address these impediments in the form of a framework for robust, falsifiable interpretability research (Section 5). And we finish by applying our framework to strengthen an example hypothesis (Section 6).

The motivation for this paper is not unprecedented; pleas have previously been made for greater rigor in the conduct and communication of machine learning research more broadly [Lipton and Steinhardt, 2018, Forde and Paganini, 2019], and in interpretability in particular [Doshi-Velez and Kim, 2017, Lipton, 2018]. Both Lipton [2018] and Doshi-Velez and Kim [2017], motivated by a lack of consensus and rigor regarding the definition, aims, and practices of interpretability research, sought to clarify these factors. In the intervening time since these publications, the technical shortcomings of the approaches in our case studies have precipitated, and many of them have been amended or alternative approaches have been developed. We think it’s important not just to amend these technical shortcomings, but also to analyze them so that we can generalize across cases to identify the impediments that allowed these technical shortcomings to arise in the first place, and determine how we can conduct interpretability research in a manner that avoids these impediments. As such, the recommendations in this paper focus more specifically on interpretability practices and less on the definitions and aims that were key parts of Lipton [2018] and Doshi-Velez and Kim [2017]. Our core aim is to articulate concrete ways of “doing good science” in the context of interpretability research, by way of referencing both exemplary techniques and cautionary tales.

We emphasize that we are not arguing that intuition should be dispensed with altogether; intuition is essential for building understanding, and all else being equal, an intuitive method is strictly preferable to an unintuitive method. We are instead arguing that *unverified* intuition and other aspects of interpretability research can facilitate misapprehension. Accordingly, the goal of this paper is to prevent misdirected effort, better realize the potential of useful approaches, and facilitate more impactful research through the application of scientific rigor.

2 Impediments to falsifiable interpretability research

Here we describe a set of impediments to interpretability research whose resolution would improve the robustness and reliability of our methods and the knowledge they generate. These impediments are discussed in general terms in this section, and in greater practical detail in the case studies presented in the subsequent two sections.

The primary impediment is an underemphasis on clear, specific, testable, and most critically, falsifiable hypotheses. This includes a lack of concrete null and alternative outcomes in the formulation of interpretability-guided hypotheses. We note that an interpretability method itself doesn't need to be used to test a hypothesis directly; it can provide an intuition that forms the basis of a hypothesis that is testable with other methods.

Not only are interpretability methods underutilized for hypothesis formulation, but *the methods themselves are too often left insufficiently scrutinized*. Specifically, they're not verified as being important for DNN function—for example the assumption that some property of a single neuron is consequential to the behavior of the network. We believe that this issue is exacerbated by the innate visual appeal of many methods, which can provide a potentially false sense of comprehension, serving as an accelerant for unverified approaches.

The abundance and seductive allure of visualization can be another impediment to robust interpretability research. Finding compelling, intuitive ways of representing information is critical for building understanding [Tufté and Graves-Morris, 1983, Victor, 2014], and many interpretability methods produce rich, stimulating visualizations [Mordvintsev et al., 2015a,b, Gatys et al., 2016, Olah et al., 2018, 2017, 2020, Gatys et al., 2015, Zeiler and Fergus, 2014]. But the illustrative power of visualization is a double-edged sword: an evocative graphic can elicit a strong feeling of comprehension regardless of whether the graphic faithfully represents the phenomenon it is attempting to depict. This makes interpretability a domain in which it is particularly easy to fall prey to unverified intuitions, and emphasizes the importance of faithful visualization.

Lack of quantification also impedes interpretability research. While not all methods can be quantified, it's hard to move beyond intuition without quantification. Quantification relates directly to the issue of verification, because quantification is often essential for verification. Anecdotal inspection of visualizations is not a method that scales to the number of scenarios that need to be evaluated to robustly verify a phenomenon. The number of models, data samples, hyperparameter configurations, and so on—the very factors that make deep learning effective—make this untenable. Furthermore, visual inspection presents risks. It can lead to cherry picking—intentionally or unintentionally—samples or instances that over-emphasize certain features or properties. It can also serve as a Rorschach test for a researcher, reinforcing their biases and adding unnecessary researcher degrees of freedom [Simmons et al., 2011].³

The final impediment regards the importance of user research for human verification. Despite some interpretability methods being designed with the stated intent to provide actionable explanations to humans, many studies neglect to conduct *controlled experiments using human subjects* to evaluate the quality and utility of these methods.

We identified five impediments to robust interpretability research: a lack of clear, falsifiable hypotheses; the abundance and seductive allure of visualization; methods based on unverified principles; under-quantification; and insufficient human utility verification. We now present two case studies that

³Researcher degrees of freedom are described by Simmons et al. [2011] as follows: “In the course of collecting and analyzing data, researchers have many decisions to make: Should more data be collected? Should some observations be excluded? Which conditions should be combined and which ones compared? Which control variables should be considered? Should specific measures be combined or transformed or both?” All of these factors can encourage misleading conclusions.

illustrate how these impediments can generate misleading conclusions in interpretability research, and how rigorous, hypothesis-driven studies helped identify and resolve these misleading conclusions.

3 Case study: The misdirection of saliency

“Saliency” methods (also known as feature attribution methods) refer to a collection of approaches that attempt to provide the user with insight as to “why” a DNN made a specific decision (i.e. performed an instance of inference) by identifying the features or parts of the input that are most relevant (or “salient”) to the model’s output, typically by using some version of the gradient of the output with regard to the input [Springenberg et al., 2015, Selvaraju et al., 2017, Fong and Vedaldi, 2017, Baehrens et al., 2010, Simonyan et al., 2014, Zeiler and Fergus, 2014, Shrikumar et al., 2017, Sundararajan et al., 2017, Smilkov et al., 2017, Dabkowski and Gal, 2017, Ancona et al., 2018, Adebayo et al., 2018]. Saliency methods generate rich, visually-appealing annotations or transformations of input samples. For example, a saliency method applied to a network trained for object recognition might generate a heatmap across an example image of a cat that highlights the cat’s ears, leading to the conclusion that ears are the most salient aspect of the image for the network’s determination of “catness”. Such a rich yet straightforward visualization strongly appeals to intuition. Indeed, the widespread study and application of saliency methods speaks to their intuitive appeal [Springenberg et al., 2015, Selvaraju et al., 2017, Fong and Vedaldi, 2017, Baehrens et al., 2010, Simonyan et al., 2014, Zeiler and Fergus, 2014, Shrikumar et al., 2017, Sundararajan et al., 2017, Smilkov et al., 2017, Dabkowski and Gal, 2017, Ancona et al., 2018, Yamashita et al., 2018, Rajpurkar et al., 2017, Wang et al., 2017, 2020, Puri et al., 2019, Mott et al., 2019, Wang et al., 2016b, Iyer et al., 2018, Greydanus et al., 2018, Nikulin et al., 2019].

But recent studies have cast doubt on their reliability [Alqaraawi et al., 2020, Kindermans et al., 2017, Sundararajan et al., 2017, Binder et al., 2016, Shrikumar et al., 2017, Chandrasekaran et al., 2017, Adebayo et al., 2018, Wang et al., 2019, Atrey et al., 2019]. Sundararajan et al. [2017], Binder et al. [2016], Shrikumar et al. [2017] identified that many studies of saliency methods lack a clear baseline for comparison (i.e. null hypothesis), a necessary condition for robust interpretability research. The notice of its absence presaged larger issues with saliency methods.

For saliency methods to be useful, they need to reflect meaningful properties of both the data and the networks to which they’re applied. For example, if an input and a baseline (i.e. control image) differ in one feature and the network has different predictions for the input and baseline, then the different feature should be identified as salient; likewise if a DNN does not depend on some feature, that feature should have zero salience. Sundararajan et al. [2017] identified that many saliency methods fail to exhibit this pair of traits, which they term “sensitivity”. The lack of sensitivity indicates that some saliency methods are prone to false negative and false positive assessments of the relevance of input features.

Saliency methods also need to reflect invariances of the networks to which they’re applied. Take the example of networks with different architectures or parameters but identical output behavior across a dataset, which we term to be functionally equivalent. Sundararajan et al. [2017] state that if two networks are functionally equivalent in this manner, a saliency method should generate identical results for the two networks. They refer to this as implementation invariance, and showed that many saliency methods do not exhibit this property. We are more circumspect about the necessity of implementation invariance for the following reason: There are potentially many ways to achieve identical input-output mappings, and the utility of being able to distinguish between implementations is context-dependent. Does the researcher care about differences between models’ intermediate layers? If so, implementation invariance is undesirable. If not, and the inputs and outputs are all that matters, then implementation invariance is desirable. But regardless of whether or not implementation invariance is desired in a particular context, it is essential to establish whether or not a particular saliency method is implementation invariant, and if it is not, how it is affected by implementation differences. The concept of invariance can also be extended to the input data: saliency methods should be invariant to input transformations that do not affect model predictions or weights, a property that [Kindermans et al., 2017] found is not obeyed by many saliency methods. These results indicate that many saliency methods may not faithfully reflect properties of the systems they are used to understand.

Perhaps most importantly, saliency methods are intended to help understand DNNs. Accordingly, they are only useful for this purpose if they provide information about the data with regards to the model, not just the data in isolation. [Adebayo et al. \[2018\]](#) examined a suite of saliency methods and found that many of them produced nearly identical saliency maps *even when all the model's parameters were reinitialized or the models were trained on label-shuffled data*. They also used an edge detector as an additional control, because edge detectors are not specific to either models or training data, and found that the outputs of the saliency methods were in many cases very similar to those of the edge detector. Because gradient-based methods naturally emphasize edges, and because edges are inherently salient to humans [[Henderson, 2003](#), [Geisler, 2008](#), [Bansal et al., 2013](#)], this means that saliency methods will likely generate maps which appear consistent with human experience, regardless of whether they reflect meaningful properties of the model, undermining their primary purpose of providing insight about a model's decision. Put another way, these methods provided the answer we wanted to hear regardless of whether or not it was true.

Saliency methods are intended to provide humans with an understanding of “why” a DNN made a decision. Accordingly, providing humans with the results of saliency methods should improve their ability to reason about and/or predict the behavior of DNNs. [Chandrasekaran et al. \[2017\]](#) used a Visual Question Answering (VQA) task and found that providing humans with saliency maps for the image and/or words in a data sample do not improve subjects' predictions of the DNN's behavior. A subsequent study [[Alqaraawi et al., 2020](#)] found that saliency maps only weakly improve participants' abilities to anticipate a network's output for new images. These studies call into question the efficacy of saliency for achieving its stated goal and highlight the necessity of verifying the human side of interpretability methods. It is possible that saliency methods lack utility for aiding human predictions because they generate outputs that appear meaningful to humans but are unrelated to the models, indicating that rigorous methodological controls and human verification are two sides of the same coin.

It is unfortunate that many saliency methods have not held up to scrutiny, especially given their application to extremely sensitive settings such as medical imaging [[Seo, 2019](#), [Amit et al., 2017](#), [Wang et al., 2017, 2018](#), [Zhang et al., 2017](#), [Mahapatra et al., 2016](#), [Rajpurkar et al., 2017](#), [Yamashita et al., 2018](#), [Karagyris et al., 2019](#)], which present a significant safety risk if deployed uncritically. Thankfully, there has been meaningful progress in assessing and developing robust saliency methods. [Kindermans et al. \[2017\]](#), [Adebayo et al. \[2018\]](#), [Arras et al. \[2019\]](#), [Yeh et al. \[2019\]](#), [Hooker et al. \[2019\]](#), [Sanchez-Lengeling et al. \[2020\]](#) in particular provide guidance for verifying interpretability methods⁴ and/or develop robust methods, and the work of [Chandrasekaran et al. \[2017\]](#), [Doshi-Velez and Kim \[2017\]](#), [Alqaraawi et al. \[2020\]](#) provide examples of how to evaluate human-relevance.

4 Case study: Single units and (sub)optimal examples

While saliency seeks to explain how a model made a decision with regard to a single input, our next case study examines a group of interpretability methods that aim to build intuitions for the broader mechanisms by which DNNs function [[Zhou et al., 2015](#), [Rafegas et al., 2019](#), [Erhan et al., 2009](#), [Zeiler and Fergus, 2014](#), [Simonyan et al., 2014](#), [Yosinski et al., 2015](#), [Nguyen et al., 2016a](#), [Olah et al., 2017, 2018, 2020](#), [Amjad et al., 2018](#), [Meyes et al., 2019](#), [Morcos et al., 2018b](#), [Zhou et al., 2015, 2018](#), [Bau et al., 2017](#), [Karpathy et al., 2016](#), [Na et al., 2019](#), [Radford et al., 2017](#), [Rafegas et al., 2019](#), [Olah et al., 2020](#), [Nguyen et al., 2016b](#), [Dosovitskiy and Brox, 2016](#), [Hohman et al., 2019](#), [Mahendran and Vedaldi, 2014](#)]. DNNs are extremely high-dimensional systems, and as the opening example with the Gaussian distribution demonstrated, it is challenging to mentally navigate high-dimensional spaces. This challenge is a strong motivator for humans to abandon high-dimensional thinking in favor of the familiarity and comfort of low-dimensional signals. This can in turn encourage the properties of individual units to be deployed as an explanation of “what” a network is “doing”, which assumes the unit is both representative of other units and responsible for some delineated sub-component of the network's behavior. Furthermore, many of these methods make claims based solely on visualization or qualitative properties. We start by examining how our falsifiability framework applies to visualizations.

⁴It bears mentioning that some studies identifying limitations in saliency methods used their discoveries to design newer, more robust saliency methods, which were then in turn shown to have significant limitations. In our view, this demonstrates the critical roles of iteration and replication in scientific research.

Visualization is a common method for building intuition about the properties of single units. In the most basic approach, one simply identifies the input sample(s) (e.g. images) or sample subregions that maximally activate a given neuron [Zhou et al., 2015, Rafegas et al., 2019]. Another approach uses optimization to generate samples that maximize unit activations [Erhan et al., 2009, Zeiler and Fergus, 2014, Simonyan et al., 2014, Yosinski et al., 2015, Nguyen et al., 2016a, Olah et al., 2017, 2018, 2020, Nguyen et al., 2016b, Dosovitskiy and Brox, 2016, Mahendran and Vedaldi, 2014]. While these methods may be useful for building intuition, they can also encourage three potentially misleading assumptions: that the visualization is representative of the neuron’s behavior; that the neuron is responsible for a clearly delineated portion of the task or the network’s behavior; and that the neuron’s behavior is representative of the network’s behavior.

To illustrate how visualization can provide misleading intuitions about a neuron’s properties, we summarize Huszár [2017]’s lucid example: If a neuron responds to images of Gaussian noise, a method that generates the “optimal” input for that neuron may end up generating a gray square, which is very different from the samples it most strongly responds to: white noise images which look like static.⁵

In addition to visualization methods potentially being misleading, they are often under-utilized for informing and evaluating strong, falsifiable hypotheses. While demonstrating that an optimal sample exists or that an optimal sample exhibits certain qualities are technically falsifiable hypotheses, these are not strong hypotheses. Stronger hypotheses would relate to network function instead of just proof of existence. It seems that visualization methods have a dual deficit when it comes to testing: their “representativeness” is not always tested, nor are they used to guide or test hypotheses.

Non-visualization-based methods for characterizing the behavior of individual neurons—typically by characterizing the distribution of the neuron’s activity over data classes or values of a data feature—are myriad and ubiquitous [Amjad et al., 2018, Meyes et al., 2019, Morcos et al., 2018b, Zhou et al., 2015, 2018, Bau et al., 2017, Karpathy et al., 2016, Na et al., 2019, Radford et al., 2017, Rafegas et al., 2019]. Because the practice is imported from neuroscience [Sherrington, 1906, Adrian, 1926, Granit, 1955, Hubel and Wiesel, 1959, Barlow, 1972, Kandel et al., 2000], we will abide by their lexicon and collectively refer to these practices as measuring “selectivity”. Selectivity-based methods have a similar allure as visualization, but the appeal is semantic instead of sensory. One notable study identified a neuron selective for sentiment in an LSTM network trained on a word prediction task [Radford et al., 2017]. Another found individual neurons selective for cats faces and human bodies [Radford et al., 2017]. [Olah et al., 2017, 2018, 2020] propose single unit selectivity as a building block for understanding DNNs. The influence of these studies demonstrates the widespread appeal of selectivity [Amjad et al., 2018, Meyes et al., 2019, Morcos et al., 2018b, Zhou et al., 2015, 2018, Bau et al., 2017, Karpathy et al., 2016, Na et al., 2019, Radford et al., 2017, Rafegas et al., 2019, Olah et al., 2017, 2018, 2020]. But neither semantic abstraction nor quantification completely avoid the pitfalls of visualization, because they can still encourage the assumptions that a neuron is responsible for a clearly delineated portion of the task or the network’s function, or is representative of a network’s behavior.

A number of recent studies examining class selectivity—a form of selectivity defined as how variable a neuron’s responses are across different classes of data samples—have demonstrated that the properties of single neurons may not robustly account for the behavior of networks Morcos et al. [2018b] used single unit ablation and other techniques to demonstrate that a network’s test set generalization is negatively correlated (or uncorrelated) with the class selectivity of its units, suggesting that highly-selective units are no more important (or perhaps even less important) than confusing, non-selective units. Amjad et al. [2018] confirmed these results for single unit ablation, but also performed cumulative ablation analyses where multiple units are ablated, the results of which suggested that selectivity is mildly beneficial. This dichotomy presents one example of how it can be difficult to extrapolate from single neurons to networks, in this case because of redundancy across units.

Similar results have been shown for models trained on NLP tasks. Dalvi et al. [2019] found that ablating units selective for linguistic features causes greater performance deficits than ablating less-selective units. Donnelly and Roegiest [2019] found that ablating Radford et al. [2017]’s “sentiment neuron” has minimal impact on the network’s ability to detect and label sentiment. These results certainly cast doubt on the assumptions that individual neurons contribute to clearly delineated

⁵We also direct readers to Olah et al. [2017] and Nguyen et al. [2016b], both of which are motivated by the potential of single instances to generate misleading conclusions.

portions of the network’s behavior and that individual neurons are representative of the network, but they are also limited by their reliance on the method of single unit ablation.

As detailed in [Leavitt and Morcos \[2020a\]](#), single unit ablation in trained networks can only address the effects of inactivating a neuron in a network whose training process assumed that the neuron was present. This presents two key limitations to investigating the causal role of selectivity: single unit ablation is unable to address whether the presence of selectivity is beneficial (i.e. is it sufficient), and it is unable to address whether networks need to learn selectivity to function properly. With these limitations in mind, [Leavitt and Morcos \[2020a\]](#) conducted a study in which they regularized against class selectivity in the loss function—sidestepping the shortcomings of single unit ablation—to examine the causal role of class selectivity in network function. They found that class selectivity is largely unnecessary for network function, and reducing it can even improve network performance. They also found in a follow-up study that reducing class selectivity increases robustness to naturalistic corruptions of input data [[Leavitt and Morcos, 2020b](#)]. These results demonstrate that networks do not need to learn class selectivity, and that in some cases learning class selectivity can impair network function, indicating that examining class selectivity in single neurons may be a red herring with regards to understanding network function.

The studies we detailed support the notion that the tractable, semantic allure of single neurons’ low-dimensional representations can encourage an overreliance on intuition. And while these studies question the functional relevance of class selectivity, we view them more broadly as a bulwark against the intuitive appeal of focusing on single neuron properties. This suggests that researchers should focus on properties that exist across neurons—distributed, high-dimensional representations—and highlights the need to develop tools to make these properties more tractable and accessible to intuition. [Agrawal et al. \[2014\]](#), [Wang et al. \[2016a\]](#), [Alain and Bengio \[2018\]](#), [Morcos et al. \[2018a\]](#), [Fong and Vedaldi \[2018\]](#) are examples of studies that attempt to better understand DNNs by examining distributed representations.

5 Best practices for robust, falsifiable interpretability research

We have attempted to illustrate how the appeal of intuition and a lack of verification and falsifiability have resulted in misleading conclusions and illusory progress in two domains of interpretability research: saliency and single neuron-based approaches. What follows is a high-level framework for best practices to help ensure robust, falsifiable interpretability research. These practices are guided by the impediments to interpretability research presented in Section 2, and by the techniques and interventions that revealed the limitations of saliency and single unit-based approaches. We provide practical examples and references whenever possible. We would also like to note that it is not our intent to present this framework as a novel discovery, but simply as a clearly-articulated, practically-applicable collection of best practices.

Our first, and most important recommendation, is to make and test clear, specific, testable, and falsifiable hypotheses. This not only demonstrates or improves the impact of a method, but also serves as a sanity check for the efficacy of that method. However, this requires being able to dissociate whether a result is due to a property of the method or a property of the system being examined. For example, determining whether a saliency method is dependent on the properties of a model, or whether it functions more like an edge detector. It is also important to note that merely proving the existence of something rarely tells us much about whether that phenomenon is relevant to network function. Stronger hypotheses relate to causal mechanisms of network function, e.g., “perturbing the representation in dimension X should change the output in dimension Y”. [Bau et al. \[2019\]](#) is a good example of this: their method, termed GAN dissection, identifies sets of units in GANs whose manipulation controls the presence of specific image features in the generated images, such as windows and people. A weaker alternative would have been, for example, to show that there exist certain subsets of units whose activity is correlated with the presence of certain features in the generated images. Furthermore, some hypotheses may be untestable using existing approaches. For example, single neuron ablation cannot address the sufficiency of class selectivity for network performance, nor whether class-selective solutions are necessary. The inability to address these causal hypotheses is what motivated [Leavitt and Morcos \[2020a\]](#) to develop their class selectivity regularizer.

Next, we caution researchers to be vigilant of the seductive allure of visualization. One’s skepticism should be proportional to the feeling of intuitiveness. Put simply: one should question their “wow”. If individual examples are being used to build intuition, they should be properly contextualized relative to other samples. For example, by providing a meaningful baseline for comparison [e.g. edge detectors for saliency methods [Adebayo et al., 2018](#), [Hooker et al., 2019](#)], or by examining where the examples lie in the distribution(s) over the quantities of interest. This latter recommendation requires a correspondence between the visualizations and a quantifiable feature of the data or model, which leads to our next recommendation.

We encourage researchers to quantify their methods. Not only does inspecting visualizations not scale to the size of deep learning models and datasets, but it can be dangerously prone to experimenter bias and adds researcher degrees of freedom³. [Adebayo et al. \[2018\]](#)’s combination of two simple methods—randomization and the Spearman correlation—were sufficient to generate the bulk of their conclusions about the validity of different saliency methods.⁶ And the challenge of quantifying saliency methods motivated [Kim et al. \[2018\]](#) to develop their Concept Activation Vector approach for interpretability; likewise, [Carlini et al. \[2019\]](#), [Jiang et al. \[2020\]](#), [Agarwal and Hooker \[2020\]](#) all propose and/or evaluate methods for quantifying the difficulty or prototypicality of data samples. While not all methods can be quantified, it is challenging to move beyond intuition without quantification: an unquantifiable hypothesis risks being an unfalsifiable hypothesis.

Our final recommendation is to remember the “human” in “human explainability”. Keep in mind whether a method is intended to inform a human, and verify that it does so. Deep learning researchers may not feel comfortable conducting human behavioral research, but [Chandrasekaran et al. \[2017\]](#), [Doshi-Velez and Kim \[2017\]](#), [Alqaraawi et al. \[2020\]](#), [Berardino et al. \[2017\]](#) illustrate straightforward approaches that yielded clear insights.

6 Building better hypotheses

In this section, we demonstrate how hypotheses can be strengthened with respect to interpretability. We start with a very weak example hypothesis, identify its weaknesses, and then use this information to improve it. We iteratively apply this process until we obtain a strong hypothesis, identifying the strengths and weaknesses of each hypothesis along the way. For the sake of relevance our example hypotheses address the role of feature-selective neurons in DNN function, but our aim is to provide a generalizable example for how to build a strong scientific hypothesis.

Very weak hypothesis: *Feature-selective neurons are the foundation of DNN function.*

The most significant weakness of this hypothesis is that it is not falsifiable. “Foundation” is vague, causing the hypothesis to lack a concrete alternative and testable predictions. Consequently, the best attempt to articulate an alternative hypothesis yields something as vague as the original: “feature-selective neurons are not the foundation of DNN function”. What would it mean for feature-selective neurons not to be the “foundation” of DNN function? Determining the validity of this hypothesis seems to be more dependent on semantics than science.

Weak hypothesis: *If feature selectivity is important for DNN function, then we should find feature-selective neurons.*

This hypothesis improves on the prior example by making a specific, falsifiable prediction: “we should find feature-selective neurons”. Unfortunately, proving the existence of some phenomenon is typically a weak outcome that lacks explanatory power in the absence of sufficient context and/or strong pre-existing evidence against its existence. This hypothesis also provides no basis for evaluating the strength or validity of its prediction: Is the baseline just the number of feature-selective neurons expected to exist by chance? And is it meaningful to simply find a quantity of feature-selective neurons that exceeds the baseline? Is a baseline used at all? Finally, this hypothesis lacks alternative potential explanations for the presence of feature-selective neurons: are there scenarios in which feature-selective neurons could exist even if feature selectivity is *not* important for DNN function?

Average hypothesis: *If feature selectivity is necessary to maximize test accuracy, ablating feature-selective single neurons should cause a decrease in test accuracy.*

⁶They also used more sophisticated metrics, but primarily to buttress the findings obtained with the Spearman correlation.

This hypothesis improves on the prior examples by addressing causality: The term “important” has been replaced with “necessary”, which specifies the causal role of feature-selective neurons. It also describes a specific experiment that is capable of addressing causality (ablation) and the predicted outcome of that experiment, which is clearly falsifiable and made in concrete terms (“DNN function” becomes “decrease in test accuracy”).

Strong hypothesis: *If feature selectivity in single neurons is necessary to maximize test accuracy, ablating neurons selective for a specific feature (e.g. curves) should cause a decrease in test accuracy proportional to the strength of the neuron’s feature selectivity for samples that contain that feature (e.g. round objects). Alternatively, if networks rely more on feature selectivity across neurons than on feature selectivity in individual neurons, then zero-ing activity in feature-selective directions (i.e. a linear combination of units that represents curves) that are not axis-aligned should cause a decrease in test accuracy that is proportional to the strength of feature selectivity and exceeds the decrease from ablating only single units.*

This formulation is substantially more comprehensive than any of the prior hypotheses. It dissociates necessity and sufficiency, and makes specific falsifiable predictions for these two factors. It also presents multiple competing hypotheses: the importance of single unit vs. multi-unit representations. The two experiments—ablating feature-selective single neurons and ablating feature-selective non-axis-aligned activity—allow the dissociation of necessity and sufficiency of these two forms of representations. A baseline for comparing the importance of the two forms of representation is also provided: the test accuracy decrease from ablating off-axis activity needs to exceed the test accuracy decrease from ablating single units for networks to be considered “more reliant” on multi-unit representations.

7 Conclusion

We detailed four impediments to robust interpretability research arising from an overemphasis on building intuition: the abundance and seductive allure of visualization; under-verification of the principles upon which interpretability methods are based; a lack of quantifiable approaches; and under-utilization of user research for human verification. We then examined two classes of interpretability methods—saliency and single neuron-based approaches—in which these impediments yielded misleading results and undermined meaningful progress. We then synthesized these impediments into a framework for robust, falsifiable interpretability research, and finished by applying this framework to iteratively improve an example hypotheses.

We would like to reiterate that the goal of this paper is to prevent misdirected effort, better realize the potential of useful approaches, and facilitate more impactful research through the application of scientific rigor. Intuition is an indispensable first step, but it has just as much power to mislead as it does to enlighten. By using these intuitions to develop and test clear, falsifiable hypotheses, we hope that researchers will converge on interpretability methods which are both intuitive *and* stand up to scrutiny to reveal deep insights about DNNs.

Acknowledgements

We would like to thank Been Kim, Sara Hooker, and Jessica Thompson for their insightful feedback.

References

- Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity Checks for Saliency Maps. arXiv:1810.03292 [cs, stat], October 2018. URL <http://arxiv.org/abs/1810.03292>. arXiv: 1810.03292.
- E. D. Adrian. The impulses produced by sensory nerve endings. *The Journal of Physiology*, 61(1): 49–72, March 1926. ISSN 0022-3751. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1514809/>.
- Chirag Agarwal and Sara Hooker. Estimating Example Difficulty using Variance of Gradients. arXiv:2008.11600 [cs], August 2020. URL <http://arxiv.org/abs/2008.11600>. arXiv: 2008.11600.

- Pulkit Agrawal, Ross Girshick, and Jitendra Malik. Analyzing the Performance of Multilayer Neural Networks for Object Recognition. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science, pages 329–344. Cham, 2014. Springer International Publishing. ISBN 978-3-319-10584-0. doi: 10.1007/978-3-319-10584-0_22.
- Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. [arXiv:1610.01644 \[cs, stat\]](https://arxiv.org/abs/1610.01644), November 2018. URL <http://arxiv.org/abs/1610.01644>. arXiv: 1610.01644.
- Ahmed Alqaraawi, Martin Schuessler, Philipp Weiß, Enrico Costanza, and Nadia Berthouze. Evaluating Saliency Map Explanations for Convolutional Neural Networks: A User Study. [arXiv:2002.00772 \[cs\]](https://arxiv.org/abs/2002.00772), February 2020. URL <http://arxiv.org/abs/2002.00772>. arXiv: 2002.00772.
- Guy Amit, Omer Hadad, Sharon Alpert, Tal Tlusty, Yaniv Gur, Rami Ben-Ari, and Sharbell Hashoul. Hybrid Mass Detection in Breast MRI Combining Unsupervised Saliency Analysis and Deep Learning. In Maxime Descoteaux, Lena Maier-Hein, Alfred Franz, Pierre Jannin, D. Louis Collins, and Simon Duchesne, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2017*, Lecture Notes in Computer Science, pages 594–602. Cham, 2017. Springer International Publishing. ISBN 978-3-319-66179-7. doi: 10.1007/978-3-319-66179-7_68.
- Rana Ali Amjad, Kairen Liu, and Bernhard C. Geiger. Understanding Individual Neuron Importance Using Information Theory. April 2018. URL <https://arxiv.org/abs/1804.06679v3>.
- Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. Towards better understanding of gradient-based attribution methods for Deep Neural Networks. February 2018. URL <https://openreview.net/forum?id=Sy21R9JAW>.
- Leila Arras, Ahmed Osman, Klaus-Robert Müller, and Wojciech Samek. Evaluating Recurrent Neural Network Explanations. [arXiv:1904.11829 \[cs, stat\]](https://arxiv.org/abs/1904.11829), June 2019. URL <http://arxiv.org/abs/1904.11829>. arXiv: 1904.11829.
- Akanksha Atrey, Kaleigh Clary, and David Jensen. Exploratory Not Explanatory: Counterfactual Analysis of Saliency Maps for Deep Reinforcement Learning. September 2019. URL <https://openreview.net/forum?id=rk13m1BFDB>.
- David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *Journal of Machine Learning Research*, 11(Jun):1803–1831, 2010.
- Aayush Bansal, Adarsh Kowdle, Devi Parikh, Andrew Gallagher, and Larry Zitnick. Which Edges Matter? In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, June 2013.
- H B Barlow. Single Units and Sensation: A Neuron Doctrine for Perceptual Psychology? *Perception*, 1(4):371–394, December 1972. ISSN 0301-0066. doi: 10.1068/p010371. URL <https://doi.org/10.1068/p010371>.
- David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network Dissection: Quantifying Interpretability of Deep Visual Representations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3319–3327, Honolulu, HI, July 2017. IEEE. ISBN 978-1-5386-0457-1. doi: 10.1109/CVPR.2017.354. URL <http://ieeexplore.ieee.org/document/8099837/>.
- David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. GAN Dissection: Visualizing and Understanding Generative Adversarial Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Alexander Berardino, Valero Laparra, Johannes Ballé, and Eero Simoncelli. Eigen-Distortions of Hierarchical Representations. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*,

- pages 3530–3539. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6944-eigen-distortions-of-hierarchical-representations.pdf>.
- Alexander Binder, Grégoire Montavon, Sebastian Bach, Klaus-Robert Müller, and Wojciech Samek. Layer-wise Relevance Propagation for Neural Networks with Local Renormalization Layers. [arXiv:1604.00825 \[cs\]](https://arxiv.org/abs/1604.00825), April 2016. URL [http://arxiv.org/abs/1604.00825](https://arxiv.org/abs/1604.00825). arXiv: 1604.00825.
- Avrim Blum, John Hopcroft, and Ravi Kannan. *Foundations of Data Science*. Cambridge University Press, 1 edition, January 2020. ISBN 978-1-108-75552-8 978-1-108-48506-7. doi: 10.1017/9781108755528. URL <https://www.cambridge.org/core/product/identifier/9781108755528/type/book>.
- Gabriel Cadamuro, Ran Gilad-Bachrach, and Xiaojin Zhu. Debugging machine learning models. In *ICML Workshop on Reliable Machine Learning in the Wild*, 2016.
- Nicholas Carlini, Ulfar Erlingsson, and Nicolas Papernot. *Prototypical Examples in Deep Learning: Metrics, Characteristics, and Utility*. 2019. URL <https://openreview.net/forum?id=r1xyx3R9tQ>.
- Arjun Chandrasekaran, Deshraj Yadav, Prithvijit Chattopadhyay, Viraj Prabhu, and Devi Parikh. It Takes Two to Tango: Towards Theory of AI’s Mind. [arXiv:1704.00717 \[cs\]](https://arxiv.org/abs/1704.00717), October 2017. URL [http://arxiv.org/abs/1704.00717](https://arxiv.org/abs/1704.00717). arXiv: 1704.00717.
- Piotr Dabkowski and Yarin Gal. Real Time Image Saliency for Black Box Classifiers. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6967–6976. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7272-real-time-image-saliency-for-black-box-classifiers.pdf>.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, Anthony Bau, and James Glass. What Is One Grain of Sand in the Desert? Analyzing Individual Neurons in Deep NLP Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6309–6317, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.33016309. URL <https://aaai.org/ojs/index.php/AAAI/article/view/4592>.
- Kedar Dhamdhere, Mukund Sundararajan, and Qiqi Yan. How Important is a Neuron. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SylKoo0cKm>.
- Jonathan Donnelly and Adam Roegiest. On Interpretability and Feature Representations: An Analysis of the Sentiment Neuron. In Leif Azzopardi, Benno Stein, Norbert Fuhr, Philipp Mayr, Claudia Hauff, and Djoerd Hiemstra, editors, *Advances in Information Retrieval, Lecture Notes in Computer Science*, pages 795–802, Cham, 2019. Springer International Publishing. ISBN 978-3-030-15712-8. doi: 10.1007/978-3-030-15712-8_55.
- Finale Doshi-Velez and Been Kim. Towards A Rigorous Science of Interpretable Machine Learning. [arXiv:1702.08608 \[cs, stat\]](https://arxiv.org/abs/1702.08608), March 2017. URL [http://arxiv.org/abs/1702.08608](https://arxiv.org/abs/1702.08608). arXiv: 1702.08608.
- Alexey Dosovitskiy and Thomas Brox. Inverting Visual Representations with Convolutional Networks. [arXiv:1506.02753 \[cs\]](https://arxiv.org/abs/1506.02753), April 2016. URL [http://arxiv.org/abs/1506.02753](https://arxiv.org/abs/1506.02753). arXiv: 1506.02753.
- Filip Karlo Došilović, Mario Brčić, and Nikica Hlupić. Explainable artificial intelligence: A survey. In *2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pages 0210–0215, May 2018. doi: 10.23919/MIPRO.2018.8400040.
- Dumitru Erhan, Yoshua Bengio, Aaron C. Courville, and Pascal Vincent. Visualizing Higher-Layer Features of a Deep Network. 2009.

- Ruth Fong and Andrea Vedaldi. Net2Vec: Quantifying and Explaining How Concepts are Encoded by Filters in Deep Neural Networks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8730–8738, Salt Lake City, UT, June 2018. IEEE. ISBN 978-1-5386-6420-9. doi: 10.1109/CVPR.2018.00910. URL <https://ieeexplore.ieee.org/document/8579008/>.
- Ruth C. Fong and Andrea Vedaldi. Interpretable Explanations of Black Boxes by Meaningful Perturbation. In 2017 IEEE International Conference on Computer Vision (ICCV), pages 3449–3457, Venice, October 2017. IEEE. ISBN 978-1-5386-1032-9. doi: 10.1109/ICCV.2017.371. URL <http://ieeexplore.ieee.org/document/8237633/>.
- Jessica Zosa Forde and Michela Paganini. The Scientific Method in the Science of Machine Learning. arXiv:1904.10922 [cs, stat], April 2019. URL <http://arxiv.org/abs/1904.10922>. arXiv: 1904.10922.
- L. A. Gatys, A. S. Ecker, and M. Bethge. Image Style Transfer Using Convolutional Neural Networks. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2414–2423, June 2016. doi: 10.1109/CVPR.2016.265. ISSN: 1063-6919.
- Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A Neural Algorithm of Artistic Style. arXiv:1508.06576 [cs, q-bio], September 2015. URL <http://arxiv.org/abs/1508.06576>. arXiv: 1508.06576.
- Wilson S. Geisler. Visual Perception and the Statistical Properties of Natural Scenes. Annual Review of Psychology, 59(1):167–192, January 2008. ISSN 0066-4308, 1545-2085. doi: 10.1146/annurev.psych.58.110405.085632. URL <http://www.annualreviews.org/doi/10.1146/annurev.psych.58.110405.085632>.
- Ragnar Granit. Receptors and sensory perception. Receptors and sensory perception. Yale University Press, New Haven, CT, US, 1955.
- Samuel Greydanus, Anurag Koul, Jonathan Dodge, and Alan Fern. Visualizing and Understanding Atari Agents. In International Conference on Machine Learning, pages 1792–1801, July 2018. URL <http://proceedings.mlr.press/v80/greydanus18a.html>. ISSN: 1938-7228 Section: Machine Learning.
- Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A Survey of Methods for Explaining Black Box Models. ACM Computing Surveys, 51(5):93:1–93:42, August 2018. ISSN 0360-0300. doi: 10.1145/3236009. URL <https://doi.org/10.1145/3236009>.
- David Gunning. Explainable artificial intelligence (XAI). Defense Advanced Research Projects Agency (DARPA), nd Web, 2, 2017.
- John M. Henderson. Human gaze control during real-world scene perception. Trends in Cognitive Sciences, 7(11):498–504, November 2003. ISSN 1364-6613. doi: 10.1016/j.tics.2003.09.006. URL <http://www.sciencedirect.com/science/article/pii/S1364661303002481>.
- Fred Hohman, Minsuk Kahng, Robert Pienta, and Duen Horng Chau. Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. IEEE Transactions on Visualization and Computer Graphics, 25(8):2674–2693, August 2019. ISSN 1941-0506. doi: 10.1109/TVCG.2018.2843369. Conference Name: IEEE Transactions on Visualization and Computer Graphics.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A Benchmark for Interpretability Methods in Deep Neural Networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-Buc, E. Fox, and R. Garnett, editors, Advances in Neural Information Processing Systems 32, pages 9734–9745. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9167-a-benchmark-for-interpretability-methods-in-deep-neural-networks.pdf>.
- D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat’s striate cortex. The Journal of Physiology, 148(3):574–591, 1959. ISSN 1469-7793. doi: 10.1113/jphysiol.1959.sp006308. URL <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1959.sp006308>.

- Ferenc Huszár. Gaussian Distributions are Soap Bubbles, November 2017. URL <https://www.inference.vc/high-dimensional-gaussian-distributions-are-soap-bubble/>. Library Catalog: www.inference.vc.
- Rahul Iyer, Yuezhong Li, Huao Li, Michael Lewis, Ramitha Sundar, and Katia Sycara. Transparency and Explanation in Deep Reinforcement Learning Neural Networks. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES '18, pages 144–150, New Orleans, LA, USA, December 2018. Association for Computing Machinery. ISBN 978-1-4503-6012-8. doi: 10.1145/3278721.3278776. URL <https://doi.org/10.1145/3278721.3278776>.
- Ziheng Jiang, Chiyuan Zhang, Kunal Talwar, and Michael C. Mozer. Characterizing Structural Regularities of Labeled Data in Overparameterized Models. arXiv:2002.03206 [cs, stat], June 2020. URL <http://arxiv.org/abs/2002.03206>. arXiv: 2002.03206.
- E R Kandel, J H Schwartz, and Jessica Chao. Principles of neural science. McGraw-Hill, New York, 2000.
- Alexandros Karargyris, Satyananda Kashyap, Joy T. Wu, Arjun Sharma, Mehdi Moradi, and Tanveer Syeda-Mahmood. Age prediction using a large chest X-ray dataset. arXiv:1903.06542 [cs], March 2019. URL <http://arxiv.org/abs/1903.06542>. arXiv: 1903.06542.
- Andrej Karpathy, Justin Johnson, and Li Fei-Fei. Visualizing and Understanding Recurrent Networks. In International Conference on Learning Representations, page 11, 2016.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). arXiv:1711.11279 [stat], June 2018. URL <http://arxiv.org/abs/1711.11279>. arXiv: 1711.11279.
- Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T. Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. The (Un)reliability of saliency methods. arXiv:1711.00867 [cs, stat], November 2017. URL <http://arxiv.org/abs/1711.00867>. arXiv: 1711.00867.
- Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. Interpretable & Explorable Approximations of Black Box Models. arXiv:1707.01154 [cs], July 2017. URL <http://arxiv.org/abs/1707.01154>. arXiv: 1707.01154.
- Matthew L. Leavitt and Ari Morcos. Selectivity considered harmful: evaluating the causal impact of class selectivity in DNNs. arXiv:2003.01262 [cs, q-bio, stat], March 2020a. URL <http://arxiv.org/abs/2003.01262>. arXiv: 2003.01262.
- Matthew L. Leavitt and Ari S. Morcos. On the relationship between class selectivity, dimensionality, and robustness, 2020b.
- Zachary C Lipton. The mythos of model interpretability. Queue, 16(3):31–57, 2018. Publisher: ACM New York, NY, USA.
- Zachary C. Lipton and Jacob Steinhardt. Troubling Trends in Machine Learning Scholarship. arXiv, July 2018. URL <http://arxiv.org/abs/1807.03341>. arXiv: 1807.03341.
- Dwarkanath Mahapatra, Pallab K. Roy, Suman Sedai, and Rahil Garnavi. Retinal Image Quality Classification Using Saliency Maps and CNNs. In Li Wang, Ehsan Adeli, Qian Wang, Yinghuan Shi, and Heung-Il Suk, editors, Machine Learning in Medical Imaging, Lecture Notes in Computer Science, pages 172–179, Cham, 2016. Springer International Publishing. ISBN 978-3-319-47157-0. doi: 10.1007/978-3-319-47157-0_21.
- Aravindh Mahendran and Andrea Vedaldi. Understanding Deep Image Representations by Inverting Them. arXiv:1412.0035 [cs], November 2014. URL <http://arxiv.org/abs/1412.0035>. arXiv: 1412.0035.
- Richard Meyes, Melanie Lu, Constantin Waubert de Puiseau, and Tobias Meisen. Ablation Studies in Artificial Neural Networks. arXiv:1901.08644 [cs, q-bio], February 2019. URL <http://arxiv.org/abs/1901.08644>. arXiv: 1901.08644.

- Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15, February 2018. ISSN 1051-2004. doi: 10.1016/j.dsp.2017.10.011. URL <http://www.sciencedirect.com/science/article/pii/S1051200417302385>.
- Ari Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 5727–5736. Curran Associates, Inc., 2018a. URL <http://papers.nips.cc/paper/7815-insights-on-representational-similarity-in-neural-networks-with-canonical-correlation.pdf>.
- Ari S. Morcos, David G. T. Barrett, Neil C. Rabinowitz, and Matthew Botvinick. On the importance of single directions for generalization. In *International Conference on Learning Representations*, 2018b. URL <https://openreview.net/forum?id=r1iuQjxCZ>.
- Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Deepdream—a code example for visualizing neural networks. *Google Research*, 2(5), 2015a. URL <https://ai.googleblog.com/2015/07/deepdream-code-example-for-visualizing.html>.
- Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks. 2015b. URL <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>.
- Alex Mott, Daniel Zoran, Mike Chrzanowski, Daan Wierstra, and Danilo J. Rezende. Towards Interpretable Reinforcement Learning Using Attention Augmented Agents. arXiv:1906.02500 [cs, stat], June 2019. URL <http://arxiv.org/abs/1906.02500>. arXiv: 1906.02500.
- Seil Na, Yo Joong Choe, Dong-Hyun Lee, and Gunhee Kim. Discovery of Natural Language Concepts in Individual Units of CNNs. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=S1EERs09YQ>.
- Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3387–3395. Curran Associates, Inc., 2016a. URL <http://papers.nips.cc/paper/6519-synthesizing-the-preferred-inputs-for-neurons-in-neural-networks-via-deep-generator-networks.pdf>.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. Multifaceted Feature Visualization: Uncovering the Different Types of Features Learned By Each Neuron in Deep Neural Networks. arXiv:1602.03616 [cs], May 2016b. URL <http://arxiv.org/abs/1602.03616>. arXiv: 1602.03616.
- Dmitry Nikulin, Anastasia Ianina, Vladimir Aliev, and Sergey Nikolenko. Free-Lunch Saliency via Attention in Atari Agents. arXiv:1908.02511 [cs], October 2019. URL <http://arxiv.org/abs/1908.02511>. arXiv: 1908.02511.
- Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature Visualization. *Distill*, 2(11):e7, November 2017. ISSN 2476-0757. doi: 10.23915/distill.00007. URL <https://distill.pub/2017/feature-visualization>.
- Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. The Building Blocks of Interpretability. *Distill*, 3(3):e10, March 2018. ISSN 2476-0757. doi: 10.23915/distill.00010. URL <https://distill.pub/2018/building-blocks>.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom In: An Introduction to Circuits. *Distill*, 5(3):e00024.001, March 2020. ISSN 2476-0757. doi: 10.23915/distill.00024.001. URL <https://distill.pub/2020/circuits/zoom-in>.
- Karl Raimund Popper. *The poverty of historicism*. Psychology Press, 2002.

- Nikaash Puri, Sukriti Verma, Piyush Gupta, Dhruv Kayastha, Shripad Deshmukh, Balaji Krishnamurthy, and Sameer Singh. Explain Your Move: Understanding Agent Actions Using Specific and Relevant Feature Attribution. September 2019. URL <https://openreview.net/forum?id=SJgzLkBKPB>.
- Alec Radford, Rafal Jozefowicz, and Ilya Sutskever. Learning to Generate Reviews and Discovering Sentiment. [arXiv:1704.01444 \[cs\]](https://arxiv.org/abs/1704.01444), April 2017. URL <http://arxiv.org/abs/1704.01444>. arXiv: 1704.01444.
- Ivet Rafegas, Maria Vanrell, Luis A. Alexandre, and Guillem Arias. Understanding trained CNNs by indexing neuron selectivity. *Pattern Recognition Letters*, page S0167865519302909, October 2019. ISSN 01678655. doi: 10.1016/j.patrec.2019.10.013. URL <http://arxiv.org/abs/1702.00382>. arXiv: 1702.00382.
- Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng. CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. [arXiv:1711.05225 \[cs, stat\]](https://arxiv.org/abs/1711.05225), December 2017. URL <http://arxiv.org/abs/1711.05225>. arXiv: 1711.05225.
- Ribana Roscher, Bastian Bohn, Marco F. Duarte, and Jochen Garcke. Explainable Machine Learning for Scientific Insights and Discoveries. *IEEE Access*, 8:42200–42216, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.2976199. Conference Name: IEEE Access.
- Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models. [arXiv:1708.08296 \[cs, stat\]](https://arxiv.org/abs/1708.08296), August 2017. URL <http://arxiv.org/abs/1708.08296>. arXiv: 1708.08296.
- Benjamin Sanchez-Lengeling, Jennifer N Wei, Brian K Lee, Emily Reif, Peter Y Wang, Wesley Wei Qian, Kevin McCloskey, Lucy Colwell, and Alexander B Wiltchko. Attribution for graph neural networks. In *Advances in Neural Information Processing Systems*, 2020.
- Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-CAM: Why did you say that? [arXiv:1611.07450 \[cs, stat\]](https://arxiv.org/abs/1611.07450), January 2017. URL <http://arxiv.org/abs/1611.07450>. arXiv: 1611.07450.
- Jae Duk Seo. Visualizing Uncertainty and Saliency Maps of Deep Convolutional Neural Networks for Medical Imaging Applications. [arXiv:1907.02940 \[cs, eess, stat\]](https://arxiv.org/abs/1907.02940), July 2019. URL <http://arxiv.org/abs/1907.02940>. arXiv: 1907.02940.
- Charles S. Sherrington. *The integrative action of the nervous system*. The integrative action of the nervous system. Yale University Press, New Haven, CT, US, 1906. doi: 10.1037/13798-000.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning Important Features Through Propagating Activation Differences. In *International Conference on Machine Learning*, pages 3145–3153, July 2017. URL <http://proceedings.mlr.press/v70/shrikumar17a.html>. ISSN: 1938-7228 Section: Machine Learning.
- Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11):1359–1366, November 2011. ISSN 0956-7976. doi: 10.1177/0956797611417632. URL <https://doi.org/10.1177/0956797611417632>. Publisher: SAGE Publications Inc.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. [arXiv:1312.6034 \[cs\]](https://arxiv.org/abs/1312.6034), April 2014. URL <http://arxiv.org/abs/1312.6034>. arXiv: 1312.6034.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. SmoothGrad: removing noise by adding noise. [arXiv:1706.03825 \[cs, stat\]](https://arxiv.org/abs/1706.03825), June 2017. URL <http://arxiv.org/abs/1706.03825>. arXiv: 1706.03825.
- Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. [arXiv:1412.6806 \[cs\]](https://arxiv.org/abs/1412.6806), April 2015. URL <http://arxiv.org/abs/1412.6806>. arXiv: 1412.6806.

- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. [arXiv:1703.01365 \[cs\]](https://arxiv.org/abs/1703.01365), June 2017. URL <http://arxiv.org/abs/1703.01365>. arXiv: 1703.01365.
- Edward R Tufte and Peter R Graves-Morris. The visual display of quantitative information, volume 2. Graphics press Cheshire, CT, 1983.
- Alfredo Vellido, Jose D Martin-Guerrero, and Paulo J G Lisboa. Making machine learning models interpretable. Computational Intelligence, page 10, 2012.
- Roman Vershynin. High-dimensional probability: An introduction with applications in data science, volume 47. Cambridge university press, 2018.
- Bret Victor. Humane representation of thought: a trail map for the 21st century. In Proceedings of the companion publication of the 2014 ACM SIGPLAN conference on Systems, Programming, and Applications: Software for Humanity, pages 5–5, 2014.
- Jianyu Wang, Zhishuai Zhang, Cihang Xie, Vittal Premachandran, and Alan Yuille. Unsupervised learning of object semantic parts from internal states of CNNs by population encoding. [arXiv:1511.06855 \[cs\]](https://arxiv.org/abs/1511.06855), November 2016a. URL <http://arxiv.org/abs/1511.06855>. arXiv: 1511.06855.
- Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3462–3471, July 2017. doi: 10.1109/CVPR.2017.369. URL <http://arxiv.org/abs/1705.02315>. arXiv: 1705.02315.
- Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, and Ronald M. Summers. TieNet: Text-Image Embedding Network for Common Thorax Disease Classification and Reporting in Chest X-rays. [arXiv:1801.04334 \[cs\]](https://arxiv.org/abs/1801.04334), January 2018. URL <http://arxiv.org/abs/1801.04334>. arXiv: 1801.04334.
- Yulong Wang, Hang Su, Bo Zhang, and Xiaolin Hu. Learning Reliable Visual Saliency for Model Explanations. IEEE Transactions on Multimedia, pages 1–1, 2019. ISSN 1941-0077. doi: 10.1109/TMM.2019.2949872. Conference Name: IEEE Transactions on Multimedia.
- Yuyao Wang, Masayoshi Mase, and Masashi Egi. Attribution-based Saliency Method towards Interpretable Reinforcement Learning. In AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering (1), 2020.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling Network Architectures for Deep Reinforcement Learning. [arXiv:1511.06581 \[cs\]](https://arxiv.org/abs/1511.06581), April 2016b. URL <http://arxiv.org/abs/1511.06581>. arXiv: 1511.06581.
- Rikiya Yamashita, Mizuho Nishio, Richard Kinh Gian Do, and Kaori Togashi. Convolutional neural networks: an overview and application in radiology. Insights into Imaging, 9(4):611–629, August 2018. ISSN 1869-4101. doi: 10.1007/s13244-018-0639-9. URL <https://doi.org/10.1007/s13244-018-0639-9>.
- Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun Suggala, David I Inouye, and Pradeep K Ravikumar. On the (In)fidelity and Sensitivity of Explanations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, Advances in Neural Information Processing Systems 32, pages 10967–10978. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9278-on-the-infidelity-and-sensitivity-of-explanations.pdf>.
- Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding Neural Networks Through Deep Visualization. In ICML Workshop on Visualization for Deep Learning, 2015. URL <http://arxiv.org/abs/1506.06579>. arXiv: 1506.06579.
- Matthew D. Zeiler and Rob Fergus. Visualizing and Understanding Convolutional Networks. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, Computer Vision – ECCV 2014, pages 818–833, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10590-1.

- Zizhao Zhang, Yuanpu Xie, Fuyong Xing, Mason McGough, and Lin Yang. MDNet: A Semantically and Visually Interpretable Medical Image Diagnosis Network. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 3549–3557, Honolulu, HI, July 2017. IEEE. ISBN 978-1-5386-0457-1. doi: 10.1109/CVPR.2017.378. URL <http://ieeexplore.ieee.org/document/8099861/>.
- B. Zhou, D. Bau, A. Oliva, and A. Torralba. Interpreting Deep Visual Representations via Network Dissection. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(9):2131–2145, September 2019. ISSN 1939-3539. doi: 10.1109/TPAMI.2018.2858759.
- Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object Detectors Emerge in Deep Scene CNNs. In International Conference on Learning Representations, April 2015. URL <http://arxiv.org/abs/1412.6856>. arXiv: 1412.6856.
- Bolei Zhou, Yiyou Sun, David Bau, and Antonio Torralba. Revisiting the Importance of Individual Units in CNNs via Ablation. arXiv:1806.02891 [cs], June 2018. URL <http://arxiv.org/abs/1806.02891>. arXiv: 1806.02891.