Mechanistically Interpreting a Transformer-based 2-SAT Solver: An Axiomatic Approach

Nils Palumbo[1], **Ravi Mangal**[2], **Zifan Wang**[3], **Saranya Vijayakumar**[2],
Corina Păsăreanu[2], **Somesh Jha**[1]

[1]University of Wisconsin-Madison, [2]Carnegie Mellon University, [3]Center for AI Safety

ABSTRACT

Mechanistic interpretability aims to reverse engineer the computation performed by a neural network in terms of its internal components. Although there is a growing body of research on mechanistic interpretation of neural networks, the notion of a *mechanistic interpretation* itself is often ad-hoc. Inspired by the notion of abstract interpretation from the program analysis literature that aims to develop approximate semantics for programs, we give a set of axioms that formally characterize a mechanistic interpretation as a description that approximately captures the semantics of the neural network under analysis in a compositional manner. We use these axioms to guide the mechanistic interpretability analysis of a Transformer-based model trained to solve the well-known 2-SAT problem. We are able to reverse engineer the algorithm learned by the model—the model first parses the input formulas and then evaluates their satisfiability via enumeration of different possible valuations of the Boolean input variables. We also present evidence to support that the mechanistic interpretation of the analyzed model indeed satisfies the stated axioms[1].

1 Introduction

Neural networks are notoriously opaque. Mechanistic interpretability seeks to reverse-engineer the computation described by a neural network in a human-interpretable form. Typically, the resulting *mechanistic interpretation* takes the form of a circuit (i.e., program) operating over features (i.e., symbols). The features refer to human-interpretable properties of the input (such as edges, textures, or shapes for vision models and part-of-speech tags, named entities, or semantic relationships for language models) that the model represents internally in its representations spaces whereas the circuit refers to a chain of human-interpretable operations, encoded by the model in its layers, such that each operation processes and transforms the features [39]. While mechanistic interpretability has been used to analyze aspects of vision as well as language models, it has found most success in analyzing small models trained to solve specific algorithmic tasks such as modular addition [40]. Although such small models are not necessarily of practical use, analyzing them has been fruitful for developing the techniques needed to mechanistically interpret.

Despite the broad interest in mechanistic interpretability, we observe that the notion of a *mechanistic interpretation* has so far only been defined in an ad-hoc manner. In other words, if an analyst claims

---

[1]Our implementation can be found at `https://github.com/nilspalumbo/sat-mi`.

to have a mechanistic interpretation, there is no agreed-upon standard to judge the analyst's claims. In this work, we propose a collection of axioms that characterize a mechanistic interpretation.[2] Our view on mechanistic interpretability is inspired by the notion of *abstract interpretation* from the program analysis literature [15, 16] that seeks to approximate the semantics of programs in order to make their analysis computationally feasible. Similarly, our axioms are meant to capture the property that the mechanistic interpretation should well approximate the semantics of the neural network under analysis. Further, our axioms also capture the intuition that such an approximation should respect the compositional structure of the neural network—not only should the input-output behaviour of the mechanistic interpretation be as close as possible to the original model, but every step in the reverse-engineered algorithm should be as close as possible to the computation performed by the internal components of the model (i.e., neurons and layers). Our axioms clarify the responsibilities of an analyst; in addition to providing a mechanistic interpretation, an analyst also needs to present evidence to support the fact that the provided interpretation satisfies these axioms.

We demonstrate our approach to mechanistic interpretability on a small neural network with a Transformer-based architecture [54] trained to solve a Boolean satisfiability problem. We choose this as a challenging yet well-understood problem that serves to highlight new techniques and our axiomatic evaluation criteria for mechanistic interpretability. Further, there are many other problems in computer science that reduce to it. To keep the analysis tractable, we focus on the 2-SAT problem with a fixed number of clauses and variables; the problem can be solved in polynomial time (while the more general 3-SAT problem is NP-complete).

Through our analysis, we are able to reverse engineer the algorithm learned by the model—the model parses the formula into a clause-level representation in its initial layers and then uses the later layers to evaluate the formula satisfiability via enumeration of different possible valuations of the Boolean variables. Since we only consider 2-SAT formulas with ten clauses and five variables, such an enumerative approach is computationally feasible and our model has the capacity to encode it. We use novel variants of attention pattern analysis to interpret the first Transformer layer of the model. For the second (i.e., final) Transformer layer, attention pattern analysis does not suffice and we use automated learning of functions (decision trees) to derive an interpretation. We also present evidence that the mechanistic interpretation extracted for this model indeed satisfies our axioms.

2   Related Work

Work on interpreting neural networks can be broadly divided between techniques that aim to find the input features with the largest effect on model behavior (*Input Interpretability*) such as gradient-based [32, 50–52] and activation-based [23, 42, 44, 49, 57] attributions, and those that aim to interpret the internal reasoning performed by a model (*Internal Interpretability*). We focus on the latter.

Mechanistic Intepretability. In this work, we use mechanistic interpretability to refer to analyses that seek to understand the model behavior *completely*, i.e., they seek to reverse engineer, in a human-understandable, and hence simplified form, the full algorithm that the neural network implements. Such analyses have typically focused on analyzing toy models trained for algorithmic tasks such as modular addition [40, 67], finding greatest common divisors [10], n-digit integer addition [47], and finite group operations [12]. The standards for judging the validity of the mechanistic interpretations

---

[2]In the same sense that group axioms characterize the notion of a group in abstract algebra.

produced by such analyses are ad-hoc and do not take the compositional structure of the neural network into account. We axiomatize the standards for making such judgment in this paper.

Mechanistic interpretability has also been used to refer to analyses that only seek to understand specific aspects of the model behavior. Such analyses, often referred to as *circuit analyses*, isolate circuits, i.e., subgraphs of the neural network computational graph, that are responsible for the behavior of interest [9, 14, 33, 35, 43, 58, 63]. The circuits are validated by measuring the causal effect of ablating the circuit using techniques such as activation patching [24, 28, 56] on a metric that measures the relevant model behavior. Such analyses are not a subject of our study.

Related to circuit analysis is emerging work on causal abstraction analysis [24, 25, 62, 63] that extracts causal models to explain certain aspects of neural network behavior and uses techniques similar to activation patching for validating the causal abstractions.

Probing. Probing involves training a separate supervised classifier to predict human-understandable features of the data from the model's internal activations [2, 68] and has found wide applications for understanding the representations learnt by language models [6, 8, 18, 26, 34, 41, 53, 59]. Even though a probe might be successful in predicting features from internal activations, it need not imply a causal relationship between these features and the model output [5]. It has also been questioned whether the model activations genuinely represent the feature of interest or the probe itself learns these features by picking up on spurious correlations [29]. However, recent works have demonstrated that for Transformer-based language models probes can be used to steer the model's output behavior [34, 68].

Related to probing is the work on *concept-based reasoning* [17, 31, 64, 65] that extracts representations of high-level concepts as geometric shapes in the neural network representation space (i.e., the activation spaces of inner layers of neural networks) and quantifies the impact of these high-level concepts on model outputs via attribution-based techniques.

Attention Pattern Analysis. There is a large body of work that attempts to understand the behavior of Transformer-based models by analyzing the *attention patterns* computed by the self-attention layer [1, 3, 11, 21, 27, 37, 38, 46, 55, 66]. However, there is an active ongoing debate about the validity of attention patterns as a tool for interpreting model behavior [4, 7, 30, 45, 61].

The surveys by Millière and Buckner [39] and Räuker et al. [48] are excellent resources for a broader and more detailed description of techniques for interpreting the internals of neural networks.

3 Axioms of Mechanistic Interpretation

Neural networks can be seen as programs in a purely functional language, which we denote λ_T, with basic operations corresponding to the commonly used neural network *layers* such as $Embed$, $Unembed$, Lin (i.e., linear), $ReLU$, $Self\text{-}Attention$, $Convolution$, $Residual$, and so on. Since λ_T is purely functional, all these operations are side-effect free—they simply transform inputs into outputs. The syntax of λ_T is as follows ($x \in Var$):

$t ::= Embed(x) \mid Unembed(x) \mid Lin(x) \mid ReLU(x) \mid Self\text{-}Attention(x) \mid Convolution(x) \mid$
$\quad Residual(x,t) \mid \ldots \mid t \circ t$

The *decomposition* of a model $t \in \lambda_T$ is a list of programs in λ_T such that the composition of these programs is syntactically equivalent to t. For example, consider the program $t := Lin \circ ReLU \circ Lin(x)$. One possible decomposition of t is the list $[Lin, ReLU, Lin]$. Given a decomposition d,

we use $d[i]$ to refer to the i^{th} component of d, $d[: i]$ as a shorthand for $d[i-1] \circ d[i-2] \circ \ldots \circ d[1]$ and refer to it as a *prefix* of t of length $i - 1$ (with respect to d), and $d[i :]$ as the shorthand for $d[len(d)] \circ d[len(d)-1] \circ \ldots \circ d[i]$ and referred to as a *suffix* of t.

Mechanistic Interpretation. The mechanistic interpretation of a neural network $t \in \lambda_T$ is a program in a different, purely functional language λ_H, that is human interpretable. The basic constructs of λ_H will tend to be more abstract than the ones supported by λ_T; λ_H might also be less expressive than λ_T to aid human-interpretability. While the specific design of λ_H and the question of how to judge whether a program is human-interpretable is domain-specific and subjective (and therefore difficult to formalize), we focus here on the question of how to judge whether a program $h \in \lambda_H$ is a valid mechanistic interpretation of the neural network $t \in \lambda_T$ under analysis.

At the very least, the input-output behavior of h should be similar to t. Ideally, the input-output behaviors of the two programs should coincide on every input but this requirement is much too strong in practice. Assuming that the inputs are drawn from a distribution $\mathcal{D}$, we formalize the weaker requirement that the outputs of the two programs should be equal with a high probability as an axiom. We additionally require that the behaviors of h and t are equivalent at the level of individual components (obtained via suitably decomposing these programs) and that replacing a component of t with the corresponding component of h has a negligible effect on the neural network's output.

Definition 3.1 (Mechanistic Interpretation). Given a model $t \in \lambda_T$ of type $X \to Y$ and a decomposition d_t of t, an ϵ-accurate mechanistic interpretation of t, with respect to decomposition d_t and a distribution $\mathcal{D}$ over X, is a program $h \in \lambda_H$ with a decomposition d_h such that $len(d_t) = len(d_h)$ and Axioms 1, 2, 3, and 4 hold.

Note that the definition is parametric in decomposition d_t of t as well as the input distribution $\mathcal{D}$. In practice, we will often consider layer-wise decompositions of t.

Abstraction and Concretization. The first four axioms use the notion of α and γ functions which we describe here. The intuition is that a neural network t and a corresponding mechanistic interpretation h operate on different types of data representations. While t operates over real-valued vectors, to aid interpretability, h is intended to operate over human-interpretable features or symbols. Accordingly, α (or *abstraction*) functions map the real-valued activations (which we also refer to as *concrete representations*) computed by t to the corresponding features or symbols (which we refer to as *abstract representations*) that h operates over. Given a decomposition d_t of t, we use α_i to refer to the α function mapping concrete representations computed by $d_t[: i + 1]$, i.e., the prefix of length i. The α functions are similar to *probes* [2] for extracting features from a model's internal activations. The γ (or *concretization*) functions map in the opposite direction—they map abstract features or symbols operated on by h to corresponding real-valued representations of those features in t's representation space. The α and γ functions in our axioms are directly inspired by the abstraction and concretization operations from the abstract interpretation literature used to map values between the original semantics and the abstract semantics of a program [15, 16].

Axioms. Axiom 1 bounds the probability that the abstract and the concrete representations computed by the same-length prefix of h and t, respectively, do not coincide (after applying α functions to map between the representations). Put differently, the mechanistic interpretation prefixes do not introduce too much error.

Axiom 1 (Prefix Equivalence).

$$\forall i \in [len(d)]. \Pr_{x \sim \mathcal{D}}[\alpha_i \circ d_t[i] \circ d_t[:i](x) = d_h[i] \circ d_h[:i] \circ \alpha_0(x)] \geq 1 - \epsilon$$

In contrast, Axiom 2 requires that none of the individual components of the mechanistic interpretation h introduce too much error. Axiom 2 does not imply Axiom 1 since the errors introduced by each component of the mechanistic interpretation can compound in the worst case [20].

Axiom 2 (Component Equivalence).

$$\forall i \in [len(d)]. \Pr_{x \sim \mathcal{D}}[\alpha_i \circ d_t[i] \circ d_t[:i](x) = d_h[i] \circ \alpha_{i-1} \circ d_t[:i](x)] \geq 1 - \epsilon$$

Axioms 3 and 4 are similar except that they consider equivalence of the output. Axiom 3 requires that replacing the prefixes of t with the corresponding prefixes of h (after applying appropriate α and γ functions to map between the representations) has limited effect on t's output. Axiom 4 requires the same when individual components of t are replaced by corresponding components of h.

Axiom 3 (Prefix Replaceability).

$$\forall i \in [len(d)]. \Pr_{x \sim \mathcal{D}}[t(x) = d_t[i+1:] \circ \gamma_i \circ d_h[i] \circ d_h[:i] \circ \alpha_0(x)] \geq 1 - \epsilon$$

Axiom 4 (Component Replaceability).

$$\forall i \in [len(d)]. \Pr_{x \sim \mathcal{D}}[t(x) = d_t[i+1:] \circ \gamma_i \circ d_h[i] \circ \alpha_{i-1} \circ d_t[:i](x)] \geq 1 - \epsilon$$

We propose two additional axioms, namely, Axioms 5 and 6, that are more informal in nature and we do not consider them for the analysis presented in this paper. These axioms are presented as a goal for future work and described in Appendix A.

Checking the Axioms. The proposed axioms can be checked statistically. Given a test dataset (which need not be labeled), we can estimate the error rate ϵ for each axiom by simply computing the proportion of inputs on which the conditions specified by the axioms are satisfied. Confidence intervals for our estimates of ϵ can be computed using any method for the parameter p of a Binomial distribution (Clopper-Pearson method [13] in our experiments). Validating the axioms is thus cheap and feasible in practice. These axioms represent a minimal set that we believe are necessary albeit may not be sufficient in all cases for characterizing a mechanistic interpretation. In future work, we plan to address this by considering analysis of models trained for other problems.

Relationship to Nanda et al. [40]'s Analysis. The seminal work of Nanda et al. [40] on mechanistically interpreting the complete behavior of a toy model serves as an interesting case study to demonstrate the relationship between the existing standards for judging the validity of a mechanistic interpretation and our proposed axioms. Nanda et al. [40] present four different pieces of evidence in support of the validity of their interpretation. The first three pieces of evidence (presented in Sections 4.1, 4.2, and 4.3 in their paper) resemble our Axiom 2 that compares each individual component of the mechanistic interpretation with the corresponding component of the model. The evidence they present in Section 4.4 of their paper is similar to our Axiom 4 that checks the effect on the output of the model when individual model components are replaced by the corresponding components from the mechanistic interpretation. Notably, they do not present the evidence required by our Axioms 1 and 3. This amounts to not considering the compositional structure of the model and, therefore, the possible compounding of error introduced by each component of the mechanistic interpretation.

4 Data and Model Details

The 2-SAT Problem. Given a Boolean formula, i.e., a conjunction of disjunctive clauses over exactly two literals, the 2-SAT problem is to determine whether there is an assignment to the formula's variables that makes the formula evaluate to *true*; the formula is said to be satisfiable, or SAT. For example, formula $(x_0 \vee x_1) \wedge (x_1 \vee \neg x_2)$ is satisfiable, with a satisfying assignment $x_0 = true, x_1 = false, x_2 = false$ (written also as $x_0, \neg x_1, \neg x_2$). If a formula has no satisfying assignment, it is said to be unsatisfiable, or UNSAT.

Dataset. We construct a dataset of randomly generated 2-SAT formulas with ten clauses and up to five variables, eliminating syntactic duplicates. We use a solver (Z3 [19]) to check satisfiability for each formula. We built a balanced dataset of 10^6 SAT and 10^6 UNSAT instances; we used 60% (also balanced) for training, and the rest for testing. For model analysis, we construct a separate dataset with 10^5 SAT and 10^5 UNSAT instances split as before in to train and test sets.

Each formula is represented as a string. For example, $(x_0 \vee x_1) \wedge (x_1 \vee \neg x_2)$ is represented as "$(x_0 x_1)(x_1 \neg x_2) : s$" where s indicates that the formula is SAT (u would indicate UNSAT). We tokenize these formulas by considering each x_i, its negation $\neg x_i$, and each of the symbols in the set $\{(,), :, s, u\}$ as a separate token. The colon token (':') is the final token in the input and we consider the next token predicted by the model as the output; hence, we refer to this token as the *readout* token and refer to its position as the readout position.

Model Architecture and Training. We use a two-layer ReLU decoder-only Transformer with token embeddings of dimension d = 128, learned positional embeddings, one attention head of dimension d in the first layer, four attention heads of dimension d/4 = 32 in the second, and n = 512 hidden units in the MLP (multi-layer perceptron). We use full-batch gradient descent with the AdamW optimizer, setting the learning rate $\gamma = 0.001$ and weight decay parameter $\lambda = 1$. We perform extensive epochs of training (1000 epochs), recognizing the combinatorial nature of SAT problems and the need for thorough exploration of the solution space. We obtained a model with 99.76% accuracy on test data. All our experiments were run on an NVIDIA A100 GPU.

5 Mechanistic Interpretability Analysis

The network can be naturally decomposed into its layers and we describe the analysis of each in the following two sections. As our goal is to understand the algorithm used by the model to determine satisfiability, we refrain from analyzing behavior which does not affect the final prediction. Thus, we focus on the model's prediction at the readout position.

5.1 First Layer as Parser

To reverse-engineer the first layer's behavior, we start by examining the attention patterns, i.e., the attention scores (by default, post-softmax unless stated otherwise) calculated by both layers on given input formulas. Figure 1 shows the attention scores on a formula from our test set. We clearly see patterns emerge—most tokens in the first layer attend heavily to the first token, while tokens in position $4i + 2, 0 \leq i < 10$, that we refer to as the *second literal positions*, primarily attend to themselves and the previous token, i.e., to the two literals in each clause.[3] Attention from the

---

[3] Recall that our formulas are of the form, $(x_0 x_1)(x_1 \neg x_2) : s$, so token positions $4i + 1$ and $4i + 2$ correspond to the literals of the i^{th} clause in the formula.
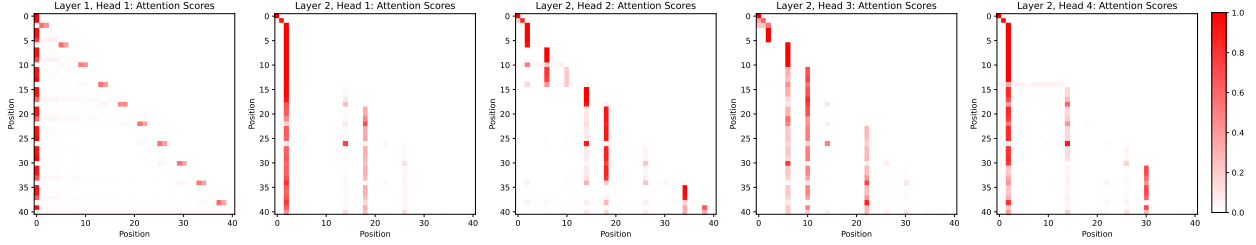
Figure 1: Attention scores for all heads on the first sample of the test set. x-axis represents the source (key) positions and y-axis represents the destination (query) positions for the attention mechanism.
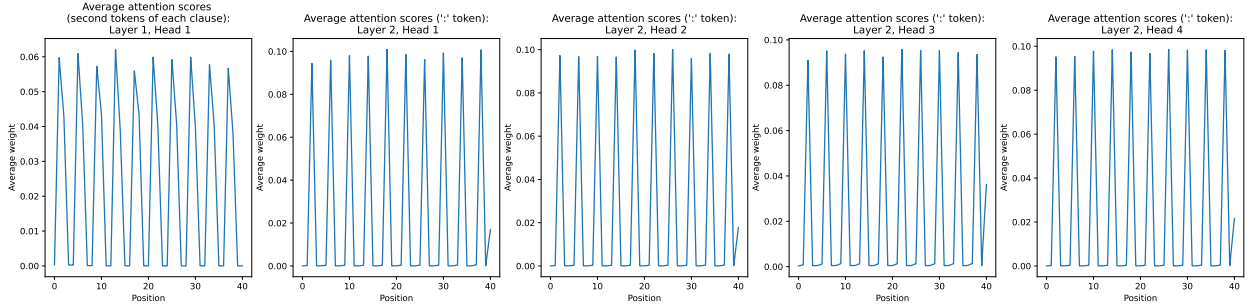

Figure 2: Average attention scores, grouped by source token position, for all heads, calculated over the test set. In the first layer, we further average across destination token positions restricted to positions $4i + 2$, where $0 \leq i < 10$ is the clause index. For the second layer, we only consider the readout token as the destination.

readout token is nearly uniform. This suggests that the first layer parses each clause, storing the parsed clause information in the embeddings computed for the second literal positions.

In the second layer, the attentions heavily focus on these second literal positions, suggesting that the model uses the parsed clause information contained in these embeddings to classify. This allows us to restrict our analysis of the first layer to the second literal positions and the readout position. To further validate this decision, we compute the average attention score from each head on each token across the test set (see Figure 2). The results show strong periodicity with period 4, following the pattern seen in Figure 1.

Distributional Attention Analysis. To further understand the parsing behavior noticed in Figure 1, we perform a more detailed analysis: instead of analyzing the attention patterns on individual formulas, we compute the attention scores in expectation assuming a uniform distribution over literals at each position $4i + 1$ and $4i + 2, 0 \leq i < 10$. We calculate the expected pre-softmax attention scores and then apply softmax. The calculation of these expected scores is described in Appendix B.1 and the expected attention scores are shown in Figure 3, which provides further evidence of the parsing behavior seen earlier in Figures 1 and 2. A decomposed analysis of the attention scores that accounts for positional factors as well a worst-case analysis of the attention scores that consider the minimum possible attention score assigned to embeddings in position $4i + 1$ and $4i + 2$ by the embedding in position $4i + 2$ further confirms the parsing hypothesis for the first layer (presented in Appendix B.1).

Interpretation. Based on our analysis of the attention patterns for the first layer, we hypothesize that it is parsing the input sequence of tokens into a list of clauses (Listing 1 in Appendix B.1). Note that we make this hypothesis without analyzing the MLP in the first layer. We show in Section 5.1.1
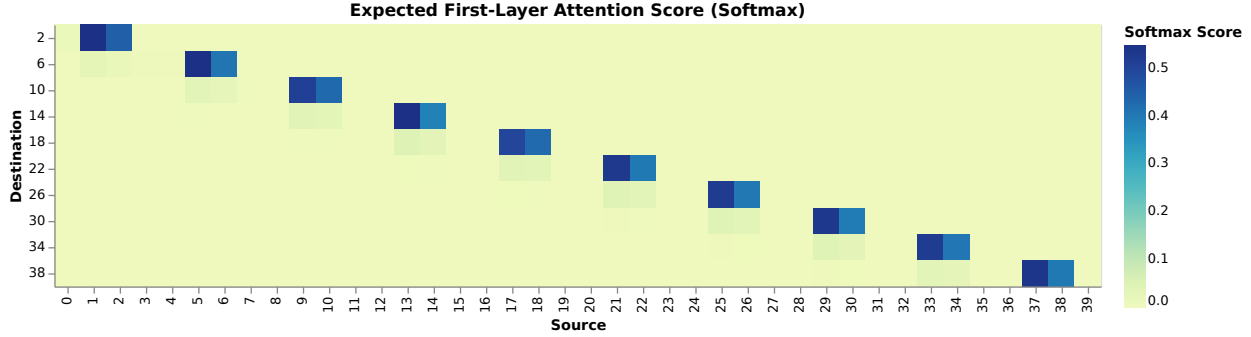
Figure 3: Expected attention scores by position for the second token of each clause.

that we can successfully validate this hypothesis using our axioms. We get lucky here with the first layer—attention patterns turn out to be a less effective tool for analyzing the second layer and we are compelled to analyze the much harder-to-interpret MLP computation.

5.1.1 Checking the Axioms

Since the input to the neural model is a sequence of token IDs, α_0 here is the function that translates token IDs in to tokens. To derive α_1 and γ_1, we first compute a representative embedding for each possible clause ψ. Drawing from the observation that the output of attention on the second token of each clause is largely a function of the two tokens in the clause alone (ignoring all other past tokens), we compute the embedding output by layer 1 for clause ψ in each position i by masking attention such that it is only applied to the left and right literals in the clause. The final representative embedding for ψ is the average of its embeddings in each position i. α_1 maps the list of embeddings output by the first layer to a list of clauses by comparing (via cosine similarity) the embeddings in the second-variable positions $4i + 2$ with the representative clause embeddings and returning the clauses with the most-similar representative embeddings. γ_1 maps a list of clauses to a list of embeddings. For positions other than $4i + 2$, γ_1 returns the mean embedding in that position across the training set. For position $4i + 2$, it returns the representative embedding for the clause in position i.

Prefix Equivalence. For each clause position i, α_1 outputs a clause $\phi := l \vee r$, where l and r are literals, and, as such, we consider the clauses $r \vee l$ and $l \vee r$ equal. With this notion of equality, we obtain an ϵ[4] of approximately 0.0000374 (corresponding to perfect matching on the test set) between the abstract states output by $\alpha_1 \circ d_t[1]$ and $d_h[1] \circ \alpha_0$. When we enforce order consistency, we obtain a much worse ϵ of 0.849; however, as the abstract state is order independent and the second layer's behavior does not vary significantly with the order of variables in a clause, an order-dependent notion of equality is not necessary here.

Component Equivalence. As this is the first layer, component equivalence and prefix equivalence are identical, hence the results are the same as above.

Prefix Replaceability. We obtain $\epsilon \approx 0.0418$, i.e. substituting the first layer of the model with $\gamma_1 \circ d_h[1] \circ \alpha_0$ affects the final output 4.2% of the time. In particular, this quantifies the effect of replacing the actual embedding of ':' with the mean embedding across the dataset and the effect of ignoring attention to previous clauses when computing the representative clause embeddings.

---

[4]We always report ϵ in terms of the upper bound of a one-sided 95% Clopper-Pearson confidence interval [13].
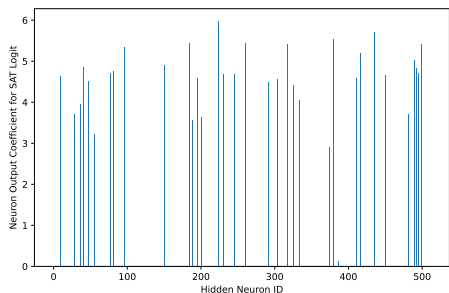
8

Figure 4: Output coefficients of hidden neurons computed using the composition of the weight matrix for the output layer of the MLP and the unembedding matrix projected to the SAT logit.
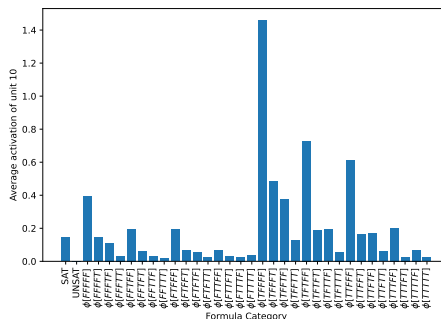

Figure 5: Average activation of neuron 10 from SAT formulas, UNSAT formulas, and formulas satisfiable with particular assignments to the variables.

Component Replaceability. As this is the first layer, component and prefix replaceability are identical, hence the results are the same as above.

5.2 Second Layer as Evaluator

From the first layer's analysis, we know that the inputs to the second layer can be described, in an abstract sense, as the list of clauses in the formula. As the model's output is a label SAT or UNSAT, the interpretation of the second layer is a function that checks satisfiability given a list of clauses. This means that the core logic of the model is encoded primarily in the second layer.

Identifying the Key Pathway. We observe that the unembeding vector for UNSAT is nearly the negative of that of SAT, and hence, their logits are likewise negatives of each other. Also, for all formulas in the dataset, either the SAT or the UNSAT logit output at the readout position is much larger than the logits for all other tokens. Hence, the effect of components on classification can be fully described by their effects on the SAT logit. Furthermore, we observe that the effect of the attention mechanism on the SAT logit is consistently supressive and varies little between samples. Hence, the core pathway necessary to identify SAT flows through the MLP at the readout position. Unlike the first layer, we cannot identify an interpretation directly from the behavior of attention. See Appendix B.2 for more details.

Sparsity in the MLP Hidden Layer. We begin our analysis of the MLP by noting that, somewhat surprisingly, it exhibits sparsity in the hidden neurons. We conclude this by first observing that the second layer of the MLP and the unembedding operation are both linear, so their composition is again linear. From this perspective, then, the MLP's effect on the classification is fully captured by a weighted sum of the post-activation outputs of the hidden neurons. We refer to these weights as *neuron output coefficients* for the SAT logit. This is the sense in which we observe sparsity; to see this clearly, see Figure 4. There are only 35 neurons which have a non-negligible absolute coefficient ($> 10^{-6}$); one of these has a significantly lower coefficient than the others along with consistently low activation values on the samples in our analysis training set resulting in a negligble effect on the SAT logit. Hence, we restrict our analysis to the remaining 34 neurons.

Evaluation in the Hidden Neurons. Note furthermore that all the non-negligible coefficients are positive; hence, the effect of each of the relevant neurons is *purely* to promote SAT. This suggests that the neurons may each recognize some subset of SAT formulas instead of recognizing

9

characteristics of UNSAT formulas. Figure 5 shows that the typical activation of one such neuron on SAT formulas is highly dependent on the assignment to the variables which satisfies the formula. This suggests that the MLP may implement a very simple algorithm: that of exhaustive evaluation. In particular, if we treat the formula ϕ as a Boolean function f_ϕ, ϕ is SAT if and only if f_ϕ is not uniquely false. Hence, we can identify SAT formulas by a brute-force technique which computes the truth table of the corresponding Boolean functions and checks whether any entry is True.

We'll see that this is indeed the case; hence, we will refer to these neurons as *evaluating neurons*. However, the behavior of the neurons is somewhat more complex than the most natural form, in which each neuron specializes in identifying formulas satisfiable with one of the 32 possible assignments. In particular, as seen in Figure 5, some neurons are affected by multiple assignments to the variables. The notation $\phi[...]$ means that the corresponding Boolean function f_ϕ evaluates to True when the variables are assigned the shown truth values. For instance, $\phi[TTFFT]$ indicates that formula ϕ is satisfiable with x_0, x_1, x_4 set to True and x_2, x_3 set to False.

A natural hypothesis is then that each of the evaluating neurons evaluates ϕ on some set of assignments, with some overlap between neurons; we refer to such interpretations of neuron activation behavior as *disjunction-only* since they take the form $h ::= \phi[...] \mid h \vee h$. However, we observe that the actual behavior is less intuitive: in some cases, satisfiability with an assignment can *decrease* activation on some evaluating neurons. Hence, we also consider more general Boolean expressions in terms of the $\phi[...]$'s as neuron interpretations. We find that our axiomatic analysis provides a clear way to quantify the tradeoff between interpretability and fidelity of the different interpretations.

As decision trees can represent arbitrary Boolean functions, we use standard decision tree learning to learn the more general interpretations of neurons; specifically, we learn classifiers trained to predict whether a neuron's activation is above or below a threshold (0.5 in our experiments) given all the features $\phi[...]$ for a formula ϕ. Deriving the simpler disjunction-only interpretation is easier; we do so by identifying the assignments a such that the empirical mean, calculated over our analysis training set, of a neuron's post-activation value on formulas satisfied by a is above the threshold. For instance, in Figure 5, these are the assignments that cause the post-activation score to be above the threshold.

Finally, we observe that the output coefficients of the evaluating neurons for the SAT logit are consistently high (> 2.9 in all cases), which means that a high activation on any individual evaluating neuron forces prediction of SAT; see Appendix B.2 for more details. In this sense, the action of MLP's output layer and the unembedding layer is conceptually close to logical OR operation. We now have enough information to describe our interpretation of the second layer.

Interpretation. Based on our analysis, we decompose the interpretation of the second layer in to two components. We hypothesize that the second layer up to the MLP's hidden layer evaluates the input formula with different assignments and outputs the results of these evaluations as a list of Booleans. The neuron interpretations describe the precise combination of assignments considered for computing each Boolean output (see Tables 1 and 2 of Appendix C for two different neuron interpretations). We evaluated the completeness of the neuron interpretations, i.e., their ability to correctly evaluate all 2-SAT formulas, using Z3. The rest of the model, namely, the MLP's output layer and the unembedding layer, performs a logical OR over the Booleans output by the previous component and outputs True if the formula is SAT. Listing 2 and 3 in Appendix B.2 describe the mechanistic interpretation of the second layer and the entire model, respectively, using Python syntax.

5.2.1 Checking the Axioms: Hidden Layer

The concrete state at the MLP's hidden layer is a pair consisting of the residual from attention and the post-activation outputs of the hidden neurons; α_2 maps these to a Boolean vector representing whether or not sufficiently high activation (> 0.5) has occurred at each of the 34 evaluating neurons. γ_2 simply returns a constant value for the residual (the mean on the analysis training set); the MLP's hidden neurons have zero activation except for the evaluating neurons with True values in the corresponding position in the abstract state vector, which we assign an arbitrary large activation (2 in our experiments). Note that this is higher than the threshold for α_2; we observe that this amplification is necessary for a reliable interpretation; see below.

Prefix Equivalence. We obtain $\epsilon \approx 0.182$ when using the decision tree neuron interpretations, compared with approximately 0.309 for the disjunction-only interpretations. This means that the disjunction-only interpretations are much worse at predicting the intermediate state of the model. This is the only place where this simplifying choice has a cost with respect to the validity of the mechanistic interpretation; aside from component equivalence below, the disjunction-only interpretations differ little from the decision trees in all other respects. This illustrates the need to check the axiomatic properties at component level for the model under analysis.

Component Equivalence. We again obtain $\epsilon \approx 0.182$ for the decision tree interpretations and 0.309 for the disjunction-only interpretations. The similarity in the ϵ values for prefix and component equivalence is because the interpretation of the first layer perfectly matches the behavior of neural network, up to literal order, on the analysis test set.

Prefix Replaceability. In both cases, substituting the first two concrete layers with the first two abstract layers has minimal effect on the model's output behavior: specifically, we get $\epsilon \approx 0.0128$ for the decision tree interpretations, and $\epsilon \approx 0.00290$ for the disjunction-only interpretations. If we drop the amplification step discussed above for γ_2, we significantly affect the model's predictions: we get an ϵ of approximately 0.249 (i.e. over 24% of samples are affected) for the decision tree interpretations and approximately 0.135 for the disjunction-only interpretations.

Component Replaceability. As above, we obtain $\epsilon \approx 0.0128$ for the decision tree interpretations, and $\epsilon \approx 0.00290$ for the disjunction-only interpretations.

5.2.2 Checking the Axioms: Output Layer

For our analysis, we consider the model to be composed with a function which returns True if and only if the top logit at the readout position is SAT. The output type of this composed model and the mechanistic interpretation are identical; hence α_3 and γ_3 are both the identity function.

Prefix Equivalence. Since prefix equivalence incorporates the mechanistic interpretation of the previous component, we need to consider both the decision tree and the disjunction-only version of the previous component. We obtain very close matching in both cases, with $\epsilon \approx 0.0128$ for the decision tree interpretations and ≈ 0.00290 for the disjunction-only interpretations.

Component Equivalence. We obtain a very close match with $\epsilon \approx 0.00433$.

Prefix Replaceability. As this is the output layer of the network and α_3 and γ_3 are the identity, this is the same as prefix equivalence.

Component Replaceability. Same as component equivalence for the reason above.

6 Conclusion

We presented a set of axioms that characterize a mechanistic interpretation and enable judgement of the interpretation's validity with respect to the neural network under analysis. Using the axioms as a guide, we analyzed a Transformer-based model trained to solve the 2-SAT problem. Our axioms provide an automated and quantitative way of evaluating the quality of a mechanistic interpretation (via the ϵ values). Not only can the ϵ's serve as useful progress measures (in the sense of Nanda et al. [40]) to understand the training dynamics of models but can also help in the development of techniques for automated mechanistic interpretability analyses by serving as a useful evaluation metric. We intend to explore these directions next in conjunction with further development and formalization of the interpretability techniques used in this work.

References

[1] S. Abnar and W. Zuidema. Quantifying attention flow in transformers. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.385. URL https://aclanthology.org/2020.acl-main.385.

[2] G. Alain and Y. Bengio. Understanding intermediate layers using linear classifier probes, 2017. URL https://openreview.net/forum?id=ryF7rTqgl.

[3] Z. Allen-Zhu and Y. Li. Physics of language models: Part 1, context-free grammar. *arXiv preprint arXiv:2305.13673*, 2023.

[4] J. Bastings and K. Filippova. The elephant in the interpretability room: Why use attention as explanation when we have saliency methods? In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, 2020.

[5] Y. Belinkov. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219, Mar. 2022. doi: 10.1162/coli_a_00422. URL https://aclanthology.org/2022.cl-1.7.

[6] N. Belrose, Z. Furman, L. Smith, D. Halawi, I. Ostrovsky, L. McKinney, S. Biderman, and J. Steinhardt. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*, 2023.

[7] A. Bibal, R. Cardon, D. Alfter, R. Wilkens, X. Wang, T. François, and P. Watrin. Is attention explanation? an introduction to the debate. In S. Muresan, P. Nakov, and A. Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.269. URL https://aclanthology.org/2022.acl-long.269.

[8] C. Burns, H. Ye, D. Klein, and J. Steinhardt. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=ETKGuby0hcs.

[9] N. Cammarata, S. Carter, G. Goh, C. Olah, M. Petrov, L. Schubert, C. Voss, B. Egan, and S. K. Lim. Thread: Circuits. *Distill*, 2020. doi: 10.23915/distill.00024. https://distill.pub/2020/circuits.

[10] F. Charton. Can transformers learn the greatest common divisor? *arXiv preprint arXiv:2308.15594*, 2023.

[11] H. Chefer, S. Gur, and L. Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 782–791, June 2021.

[12] B. Chughtai, L. Chan, and N. Nanda. A toy model of universality: reverse engineering how networks learn group operations. ICML'23. JMLR.org, 2023.

[13] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. ISSN 00063444. URL http://www.jstor.org/stable/2331986.

[14] A. Conmy, A. N. Mavor-Parker, A. Lynch, S. Heimersheim, and A. Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=89ia77nZ8u.

[15] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 238–252, 1977.

[16] P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of logic and computation*, 2(4):511–547, 1992.

[17] J. Crabbé and M. van der Schaar. Concept activation regions: A generalized framework for concept-based explanations. *Advances in Neural Information Processing Systems*, 35: 2590–2607, 2022.

[18] H. Cunningham, A. Ewart, L. Riggs, R. Huben, and L. Sharkey. Sparse autoencoders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*, 2023.

[19] L. De Moura and N. Bjørner. Z3: An efficient smt solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.

[20] N. Dziri, X. Lu, M. Sclar, X. L. Li, L. Jiang, B. Y. Lin, S. Welleck, P. West, C. Bhagavatula, R. L. Bras, J. D. Hwang, S. Sanyal, X. Ren, A. Ettinger, Z. Harchaoui, and Y. Choi. Faith and fate: Limits of transformers on compositionality. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=Fkckkr3ya8.

[21] J. Ebrahimi, D. Gelda, and W. Zhang. How can self-attention networks recognize dyck-n languages? In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4301–4306, 2020.

[22] N. Elhage, N. Nanda, C. Olsson, T. Henighan, N. Joseph, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, N. DasSarma, D. Drain, D. Ganguli, Z. Hatfield-Dodds, D. Hernandez, A. Jones, J. Kernion, L. Lovitt, K. Ndousse, D. Amodei, T. Brown, J. Clark, J. Kaplan, S. McCandlish, and C. Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. https://transformer-circuits.pub/2021/framework/index.html.

[23] R. C. Fong and A. Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3449–3457, 2017.

[24] A. Geiger, H. Lu, T. Icard, and C. Potts. Causal abstractions of neural networks. *Advances in Neural Information Processing Systems*, 34:9574–9586, 2021.

[25] A. Geiger, Z. Wu, C. Potts, T. Icard, and N. D. Goodman. Finding alignments between interpretable causal variables and distributed neural representations. *arXiv preprint arXiv:2303.02536*, 2023.

[26] W. Gurnee, N. Nanda, M. Pauly, K. Harvey, D. Troitskii, and D. Bertsimas. Finding neurons in a haystack: Case studies with sparse probing. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=JYs1R9IMJr.

[27] Y. Hao, L. Dong, F. Wei, and K. Xu. Self-attention attribution: Interpreting information interactions inside transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12963–12971, 2021.

[28] S. Heimersheim and N. Nanda. How to use and interpret activation patching. *arXiv preprint arXiv:2404.15255*, 2024.

[29] J. Hewitt and P. Liang. Designing and interpreting probes with control tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, 2019.

[30] S. Jain and B. C. Wallace. Attention is not Explanation. In J. Burstein, C. Doran, and T. Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1357. URL https://aclanthology.org/N19-1357.

[31] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pages 2668–2677. PMLR, 2018.

[32] K. Leino, S. Sen, A. Datta, M. Fredrikson, and L. Li. Influence-directed explanations for deep convolutional networks. In *2018 IEEE International Test Conference (ITC)*, pages 1–8. IEEE, 2018.

[33] M. A. Lepori, T. Serre, and E. Pavlick. Uncovering intermediate variables in transformers using circuit probing. *arXiv e-prints*, pages arXiv–2311, 2023.

[34] K. Li, A. K. Hopkins, D. Bau, F. Viégas, H. Pfister, and M. Wattenberg. Emergent world representations: Exploring a sequence model trained on a synthetic task. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=DeG07_TcZvT.

[35] T. Lieberum, M. Rahtz, J. Kramár, G. Irving, R. Shah, and V. Mikulik. Does circuit analysis interpretability scale? evidence from multiple choice capabilities in chinchilla. *arXiv preprint arXiv:2307.09458*, 2023.

[36] D. Lindner, J. Kramar, S. Farquhar, M. Rahtz, T. McGrath, and V. Mikulik. Tracr: Compiled transformers as a laboratory for interpretability. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 37876–37899. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/771155abaae744e08576f1f3b4b7ac0d-Paper-Conference.pdf.

[37] B. Liu, J. T. Ash, S. Goel, A. Krishnamurthy, and C. Zhang. Transformers learn shortcuts to automata. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=De4FYqjFueZ.

[38] K. Lu, Z. Wang, P. Mardziel, and A. Datta. Influence patterns for explaining information flow in bert. *Advances in Neural Information Processing Systems*, 34:4461–4474, 2021.

[39] R. Millière and C. Buckner. A philosophical introduction to language models - part ii: The way forward, 2024.

[40] N. Nanda, L. Chan, T. Lieberum, J. Smith, and J. Steinhardt. Progress measures for grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=9XFSbDPmdW.

[41] N. Nanda, A. Lee, and M. Wattenberg. Emergent linear representations in world models of self-supervised sequence models. In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 16–30, 2023.

[42] C. Olah, L. Schubert, and A. Mordvintsev. Feature visualization. *Distill*, 2017. URL https://distill.pub/2017/feature-visualization/.

[43] C. Olsson, N. Elhage, N. Nanda, N. Joseph, N. DasSarma, T. Henighan, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, D. Drain, D. Ganguli, Z. Hatfield-Dodds, D. Hernandez, S. Johnston, A. Jones, J. Kernion, L. Lovitt, K. Ndousse, D. Amodei, T. Brown, J. Clark, J. Kaplan, S. McCandlish, and C. Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html.

[44] V. Petsiuk, A. Das, and K. Saenko. Rise: Randomized input sampling for explanation of black-box models. In *BMVC*, 2018.

[45] D. Pruthi, M. Gupta, B. Dhingra, G. Neubig, and Z. C. Lipton. Learning to deceive with attention-based explanations. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*,

pages 4782–4793, Online, July 2020. Association for Computational Linguistics. doi: 10. 18653/v1/2020.acl-main.432. URL https://aclanthology.org/2020.acl-main.432.

[46] Y. Qiang, D. Pan, C. Li, X. Li, R. Jang, and D. Zhu. Attcat: Explaining transformers via attentive class activation tokens. *Advances in Neural Information Processing Systems*, 35: 5052–5064, 2022.

[47] P. Quirke et al. Understanding addition in transformers. *arXiv preprint arXiv:2310.13121*, 2023.

[48] T. Räuker, A. Ho, S. Casper, and D. Hadfield-Menell. Toward transparent ai: A survey on interpreting the inner structures of deep neural networks. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 464–483. IEEE, 2023.

[49] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128(2):336–359, Oct 2019. ISSN 1573-1405. doi: 10.1007/s11263-019-01228-7. URL http://dx.doi.org/10.1007/s11263-019-01228-7.

[50] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps, 2013.

[51] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg. Smoothgrad: removing noise by adding noise, 2017.

[52] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, 2017. URL https://api.semanticscholar.org/CorpusID:16747630.

[53] I. Tenney, D. Das, and E. Pavlick. BERT rediscovers the classical NLP pipeline. In A. Korhonen, D. Traum, and L. Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1452. URL https://aclanthology.org/P19-1452.

[54] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[55] J. Vig and Y. Belinkov. Analyzing the structure of attention in a transformer language model. In T. Linzen, G. Chrupała, Y. Belinkov, and D. Hupkes, editors, *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 63–76, Florence, Italy, Aug. 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4808. URL https://aclanthology.org/W19-4808.

[56] J. Vig, S. Gehrmann, Y. Belinkov, S. Qian, D. Nevo, S. Sakenis, J. Huang, Y. Singer, and S. Shieber. Causal mediation analysis for interpreting neural nlp: The case of gender bias. *arXiv preprint arXiv:2004.12265*, 2020.

[57] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 24–25, 2020.

[58] K. R. Wang, A. Variengien, A. Conmy, B. Shlegeris, and J. Steinhardt. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=NpsVSN6o4ul.

[59] X. Wang, S. Mao, N. Zhang, S. Deng, Y. Yao, Y. Shen, L. Liang, J. Gu, and H. Chen. Editing conceptual knowledge for large language models. *ArXiv*, abs/2403.06259, 2024. URL https://api.semanticscholar.org/CorpusID:268358212.

[60] G. Weiss, Y. Goldberg, and E. Yahav. Thinking like transformers. In *International Conference on Machine Learning*, pages 11080–11090. PMLR, 2021.

[61] S. Wiegreffe and Y. Pinter. Attention is not not explanation. In K. Inui, J. Jiang, V. Ng, and X. Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1002. URL https://aclanthology.org/D19-1002.

[62] Z. Wu, K. D'Oosterlinck, A. Geiger, A. Zur, and C. Potts. Causal proxy models for concept-based model explanations. In *International conference on machine learning*, pages 37313–37334. PMLR, 2023.

[63] Z. Wu, A. Geiger, C. Potts, and N. D. Goodman. Interpretability at scale: Identifying causal mechanisms in alpaca. *arXiv preprint arXiv:2305.08809*, 2023.

[64] C.-K. Yeh, B. Kim, S. Arik, C.-L. Li, T. Pfister, and P. Ravikumar. On completeness-aware concept-based explanations in deep neural networks. *Advances in neural information processing systems*, 33:20554–20565, 2020.

[65] C.-K. Yeh, B. Kim, and P. Ravikumar. Human-centered concept explanations for neural networks. In *Neuro-Symbolic Artificial Intelligence: The State of the Art*, pages 337–352. IOS Press, 2021.

[66] Y. Zhang, A. Backurs, S. Bubeck, R. Eldan, S. Gunasekar, and T. Wagner. Unveiling transformers with LEGO: A synthetic reasoning task, 2023. URL https://openreview.net/forum?id=1jDN-RfQfrb.

[67] Z. Zhong, Z. Liu, M. Tegmark, and J. Andreas. The clock and the pizza: Two stories in mechanistic explanation of neural networks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=S5wmbQc1We.

[68] A. Zou, L. Phan, S. Chen, J. Campbell, P. Guo, R. Ren, A. Pan, X. Yin, M. Mazeika, A.-K. Dombrowski, S. Goel, N. Li, M. J. Byun, Z. Wang, A. Mallen, S. Basart, S. Koyejo, D. Song, M. Fredrikson, J. Z. Kolter, and D. Hendrycks. Representation engineering: A top-down approach to ai transparency, 2023.

A Axioms of Mechanistic Interpretation

Axioms 5 and 6 are stated more informally. Unlike the earlier axioms requiring the observed behaviors of the mechanistic interpretation and the neural network to coincide, these axioms are concerned with the *compilability* of the mechanistic interpretation into a neural network. Both axioms assume the existence of a semantics-preserving compiler from λ_H to λ_T (denoted by $compiler_{\lambda_T}^{\lambda_H}$). While Axiom 5 requires the compiled version of h to have the same structure and parameters as t (i.e., syntactic equivalence), Axiom 6 requires the compiled version to only have the same structure as t. The requirements imposed by these last two axioms are hard to establish in practice, and we do not consider them for the analysis presented in this paper. We present these axioms as a goal for future work. Recent work on compilers from the RASP language [60] to Transformer models is a promising step in this direction [36].

Axiom 5 (Strong Mechanistic Derivability). h is mechanistically derived from t if there exists a semantics-preserving compilation $compiler_{\lambda_T}^{\lambda_H}(h)$ of h in λ_T such that $compiler_{\lambda_T}^{\lambda_H}(h)$ and t are syntactically equal.

Axiom 6 (Weak Mechanistic Derivability). h is mechanistically derivable from t if there exists a semantics-preserving compilation $compiler_{\lambda_T}^{\lambda_H}(h)$ of h in λ_T such that $compiler_{\lambda_T}^{\lambda_H}(h)$ has the same *architecture* (with the same number of parameters) as t.

B More Details on Mechanistic Interpretability Analysis

B.1 First Layer as Parser

Decomposed Attention Analysis. The parsing behavior in the first layer cannot be fully described by token-to-token attention, and hence, we further decompose the standard QK-decomposition of Elhage et al. [22] to account for positional factors as well. In a Transformer, each token's initial embedding is the sum of a positional embedding and a token embedding; hence, rather than viewing attention as from destination embeddings to source embeddings, we can equivalently express the pre-softmax scores as the sum of four sets of preferences—token-to-token attention, token-to-position attention, position-to-token attention, and position-to-position attention. In particular we can decompose the first-layer's self-attention mechanism as follows.[5] Given token t_{src} in position p_{src} to token t_{dst} in position t_{dst}, the pre-softmax score is

$$
\begin{aligned}
& \left(t_{\text{src}}^T W_E^T + p_{\text{src}}^T W_{POS}^T \right) W_K^T W_Q \left(W_E t_{\text{dst}} + W_{POS} p_{\text{dst}} \right) \\
= & \left(t_{\text{src}}^T W_E^T + p_{\text{src}}^T W_{POS}^T \right) W_{QK} \left(W_E t_{\text{dst}} + W_{POS} p_{\text{dst}} \right) \\
= & t_{\text{src}}^T W_E^T W_{QK} W_E t_{\text{dst}} + t_{\text{src}}^T W_E^T W_{QK} W_{POS} t_{\text{src}} + \\
& p_{\text{src}}^T W_{POS}^T W_{QK} W_E t_{\text{dst}} + p_{\text{src}}^T W_{POS}^T W_{QK} W_{POS} p_{\text{dst}} \\
= & t_{\text{src}}^T W_{QK}^{\text{tok}\to\text{tok}} t_{\text{dst}} + t_{\text{src}}^T W_{QK}^{\text{tok}\to\text{pos}} p_{\text{src}} + p_{\text{src}}^T W_{QK}^{\text{pos}\to\text{tok}} t_{\text{dst}} + p_{\text{src}}^T W_{QK}^{\text{pos}\to\text{pos}} p_{\text{dst}},
\end{aligned}
$$

where W_E and W_{POS} are the token and positional embedding matrices, W_{QK} is the QK matrix, t_{src} and t_{dst} are the source and destination tokens, while p_{src} and p_{dst} are source and destination positions.

Our observations in Figures 1 and 2 suggest that the tokens relevant to the final classification decision, are in position $4i + 2$ for $0 \le i < 10$, which, given the structure of the inputs to the model, are the second tokens of each clause of the form x_i or $\neg x_i$ for $0 \le i < 5$. If we restrict our view

---

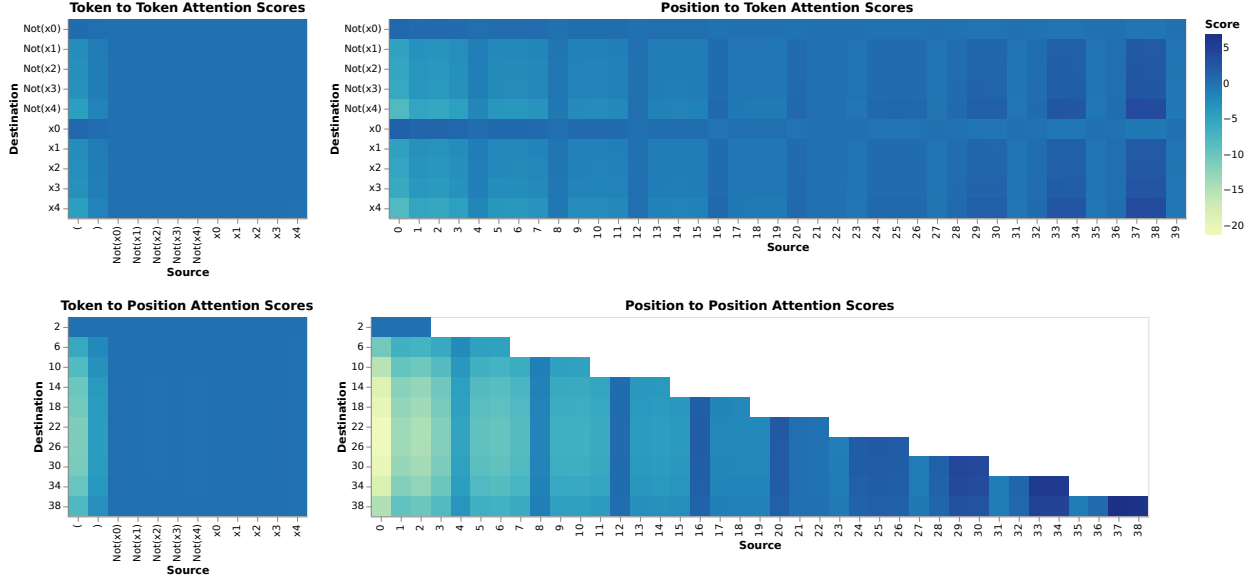[5]We use the QK circuit notation of Elhage et al. [22]

Figure 6: Decomposed QK circuit of the first layer's attention mechanism; the subset where the destination token is the second variable of a clause is shown.

to destination tokens at this position, we can express the effects of the four components of the first layer's QK circuit as shown in Figure 6. With the somewhat odd exception of the first clause, and, in particular, a first clause with second literal x_0 or $\neg x_0$, attention strongly de-prioritizes the parentheses but has weak preferences for the source token otherwise. Outside of destination tokens x_0 and $\neg x_0$, preferences for source position are fairly consistent, and hence, the core effect is a combination of position-to-position preferences (see the bottom right of Figure 6) and a preference against punctuation (the parentheses tokens '(' and ')'). This can be seen in the token-to-position preferences in the bottom left, with an exception in the case of token 2, which may behave differently as all tokens to the left belong to the first clause, limiting the importance of learning more specific attention patterns. A general preference against punctuation can be seen in the token-to-position preferences (top left) as well. For the first several clauses, positional preferences encode a preference for the positions where parentheses occur, open parentheses in particular; however, the preference against parentheses shown in the token-to-position attention scores (corresponding to $W_{QK}^{\text{tok}\rightarrow\text{pos}}$) counteract that effect to result in a net effect of attention to the first literal of the clause.

Distributional Attention Analysis. To illustrate how these four components of the QK circuit work together on typical formulas, we can use the values of the four QK matrices to compute the attention scores in expectation given a uniformly distributed choice of literals. To compute the attention scores in expectation, we only consider destination positions p ($p = 4i + 2$ for some i). Call the token in that position, t, and the full embedding passed to attention, e. To compute the expected pre-softmax attention score on position $p' \leq p$ (defining t' and e' similarly), we first check if p' contains punctuation. If $p' \equiv 0 \pmod 4$, t' is '(' and if $p' \equiv 3 \pmod 4$, t' is ')'. For such t

19

and t' we can use our decomposition into the four sets of preferences to calculate an expected score:

$$\mathbb{E}_{t,t'} e'^T W_{QK} e = \mathbb{E}_{t',t} \left[t'^T W_{QK}^{\text{tok}\to\text{tok}} t + t'^T W_{QK}^{\text{tok}\to\text{pos}} p + p'^T W_{QK}^{\text{pos}\to\text{tok}} t + p'^T W_{QK}^{\text{pos}\to\text{pos}} p \right]$$

$$= \mathbb{E}_{t} t'^T W_{QK}^{\text{tok}\to\text{tok}} t + t'^T W_{QK}^{\text{tok}\to\text{pos}} p + \mathbb{E}_{t} p'^T W_{QK}^{\text{pos}\to\text{tok}} t + p'^T W_{QK}^{\text{pos}\to\text{pos}} p$$

$$= \frac{1}{|lit|} \sum_{t \in lit} W_{QK}^{\text{tok}\to\text{tok}}{}_{t',t} + t'^T W_{QK}^{\text{tok}\to\text{pos}} p + $$

$$\frac{1}{|lit|} \sum_{t \in lit} W_{QK}^{\text{pos}\to\text{tok}}{}_{p',t} + p'^T W_{QK}^{\text{pos}\to\text{pos}} p$$

where we overload notation and use t, t' as tokens, indices, and the corresponding one-hot representations, and similarly for the positions p and p' and where lit is the set of literals.

If p' does not contain punctuation, then we know that t' is a variable (x_i or $\neg x_i$ for some i), and similarly for t. Using that, we can use our decomposition into the four sets of preferences to calculate an expected score:

$$\mathbb{E}_{t,t'} e'^T W_{QK} e = \mathbb{E}_{t',t} \left[t'^T W_{QK}^{\text{tok}\to\text{tok}} t + t'^T W_{QK}^{\text{tok}\to\text{pos}} p + p'^T W_{QK}^{\text{pos}\to\text{tok}} t + p'^T W_{QK}^{\text{pos}\to\text{pos}} p \right]$$

$$= \mathbb{E}_{t',t} t'^T W_{QK}^{\text{tok}\to\text{tok}} t + \mathbb{E}_{t'} t'^T W_{QK}^{\text{tok}\to\text{pos}} p + \mathbb{E}_{t} p'^T W_{QK}^{\text{pos}\to\text{tok}} t + p'^T W_{QK}^{\text{pos}\to\text{pos}} p$$

$$= \frac{1}{|lit|^2} \sum_{t,t' \in lit} W_{QK}^{\text{tok}\to\text{tok}}{}_{t',t} + \frac{1}{|lit|} \sum_{t' \in lit} W_{QK}^{\text{tok}\to\text{pos}}{}_{t',p} + $$

$$\frac{1}{|lit|} \sum_{t \in lit} W_{QK}^{\text{pos}\to\text{tok}}{}_{p',t} + p'^T W_{QK}^{\text{pos}\to\text{pos}} p$$

After taking the softmax of the resulting values, the result is shown in Figure 3.

Worst-case Attention Analysis. We can also show that the parsing behavior occurs by a worst-case analysis of the attention scores, i.e. what is the minimum weight from attention to the first token of the clause and to the clause as a whole? This allows us to dispense with any distributional assumptions, and to validate that the inconsistent preferences for particular literals (i.e. the differences in preferences when the destination token is x_0). Given the definition of softmax, we can compute the minimum weight on the first token of clause i by the second token of clause i as follows:

$$\min_{\phi} \text{score}_{4i+1,4i+2}^{\text{softmax}}(\phi)$$

$$= \min_{t \in tok_{4i+2}} \min_{\{\phi | \phi_{4i+2} = t\}} \frac{e^{\text{score}_{4i+1,4i+2}(\phi)}}{\sum_{j \leq 4i+2} e^{\text{score}_{j,4i+2}(\phi)}}$$

$$\leq \min_{t \in tok_{4i+2}} \frac{e^{\min_{\{\phi | \phi_{4i+2} = t\}} \text{score}_{4i+1,4i+2}(\phi)}}{e^{\min_{\{\phi | \phi_{4i+2} = t\}} \text{score}_{4i+1,4i+2}(\phi)} + \sum_{j \leq 4i \vee j = 4i+2} e^{\max_{\{\phi | \phi_{4i+2} = t\}} \text{score}_{j,4i+2}(\phi)}}$$

where tok_i is the set of all possible tokens in position i and where $\text{score}_{s,d}(\phi)$ and $\text{score}_{s,d}^{\text{softmax}}(\phi)$ refer to the pre- and post-softmax attention scores with destination token in position d and source token in position s and where the input formula is ϕ. For the full clause, the approach is similar, except we add the term $e^{\min_{\{\phi | \phi_{4i+2} = t\}} \text{score}_{4i+2,4i+2}(\phi)}$ to the numerator.
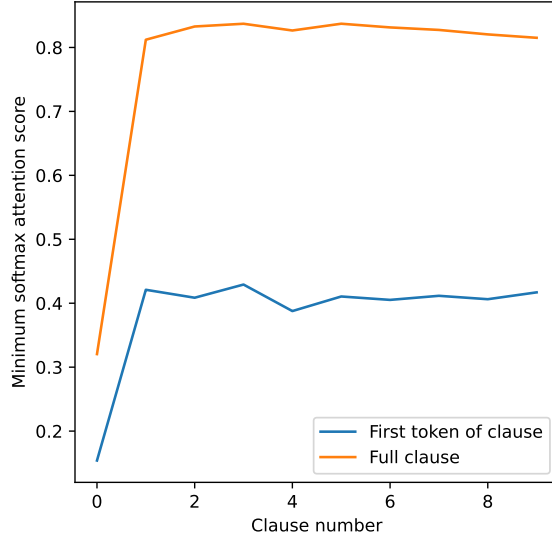
20

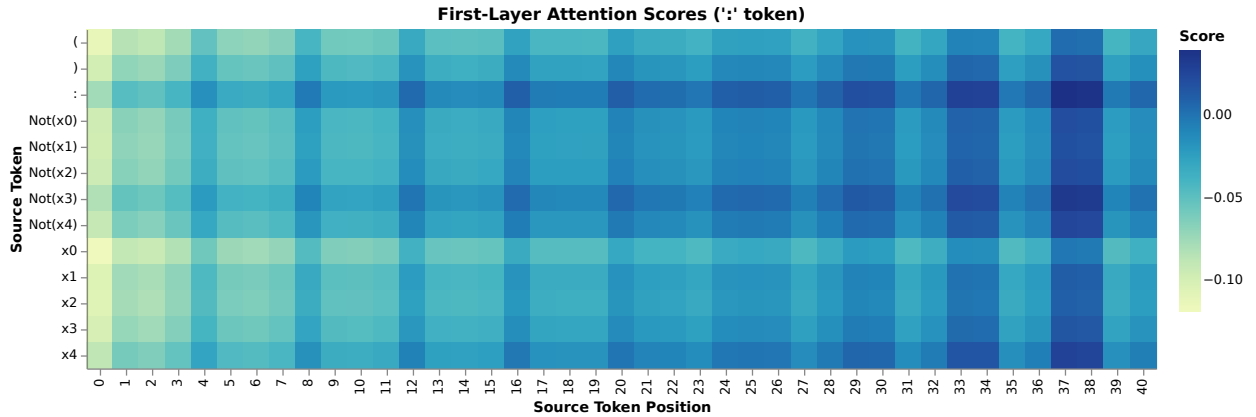Figure 7: Minimum post-softmax attention weights for each clause.

Figure 8: Pre-softmax attention scores by the ':' token to each token, by position and token type.

Now, we can derive the minimal and maximal scores for and s and d using our decomposition of the QK circuit:

$$\min_{\{\phi|\phi_d=t\}} \text{score}_{s,d}(\phi) = W_{QK}^{\text{pos}\to\text{pos}}{}_{s,d} + \min_{t'\in tok_s} W_{QK}^{\text{tok}\to\text{pos}}{}_{t',d} + W_{QK}^{\text{pos}\to\text{tok}}{}_{s,t} + \min_{t'\in tok_s} W_{QK}^{\text{tok}\to\text{tok}}{}_{t',t}$$

and similarly for the maximum (noting that if $s = d$, we must have $t' = t$ as well).

The results are shown in Figure 7, and demonstrate that regardless of the choice of formula, the majority of attention will be placed on the clause as a whole, except in the case of the first clause, in which case any additional attention is to the first '(' token, which contains no useful information.

Now, we can perform similar analysis on the attention of the ':' token. We observe in Figure 1 that the attention of the ':' token in the first-layer is near uniform. To see why this occurs, see Figure 8, which fully describes the first-layer attention of ':'. These are derived from the four QK matrices

21

```python
def parse_clauses(formula: list[tok]) -> list[tuple[lit, lit]]:
    return [(formula[4*i+1], formula[4*i+2]) for i in range(10)]
```

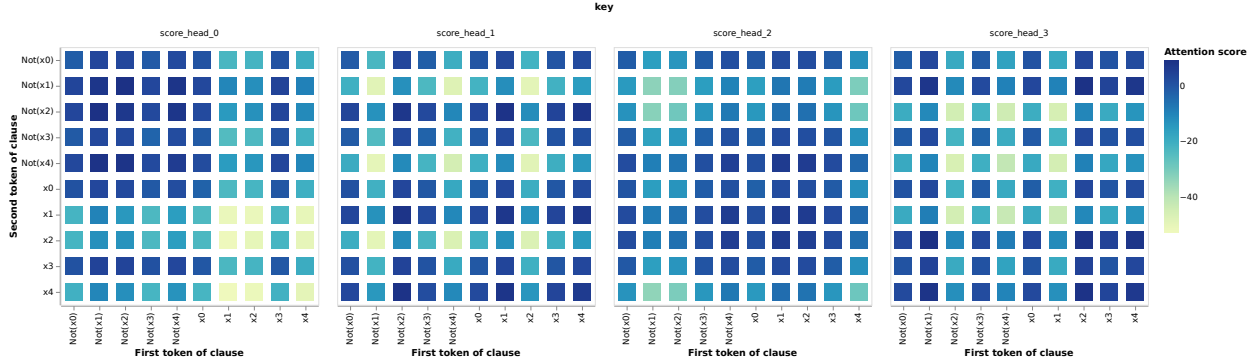Listing 1: First layer's mechanistic interpretation in Python syntax.

Figure 9: Second-layer abstract attention preferences.

using the known position $c = 40$ of the token. The attention scores are consistently very small and have very little variation. As above, we can use this set of preferences to calculate a lower bound on the attention paid by the ':' token to any individual token:

$$\min_{i,\phi} \text{score}^{\text{softmax}}_{i,40}(\phi) \leq \frac{e^{\min_{i,t} \text{score}^{(:)}_{i,t}}}{e^{\min_{i,t} \text{score}^{(:)}_{i,t}} + 40e^{\max_{i,t} \text{score}^{(:)}_{i,t}}}$$

where $\text{score}^{(:)}_{i,t}$ refers to the ":" token's pre-softmax attention to token t in position i. We derive that the minimum attention paid to any token by ':' is approximately 0.0209, and we can similarly show that the maximum attention to any token is ≈ 0.0285; hence, the ":" token's attention can never vary far from the ≈ 0.0244 paid by uniform attention.

Interpretation. Listing 1 shows our extracted mechanistic interpretation of the first layer as Python code.

B.2 Second Layer as Evaluator

Abstract Attention Analysis. We showed in Section 5.1 that the behavior of the first layer is accurately captured by an interpretation which, after the application of γ_1, outputs fixed representative embeddings for each clause at the second literal positions and is constant at all other positions. Prefix replaceability, in particular, enforces this property directly. Furthermore, we must only study attention with the readout token as the destination to explain the output behavior; hence, we can characterize the behavior of attention in the second layer solely by the preferences of each head in the readout token position to the representative clause embeddings in the second literal positions, dramatically reducing dimensionality.

In this way, the key behavior of attention in the second layer is fully described by Figure 9. Each of the four charts shows the preferences on each clause by the corresponding head; the columns correspond to the first literal of the clause and the rows to the second.
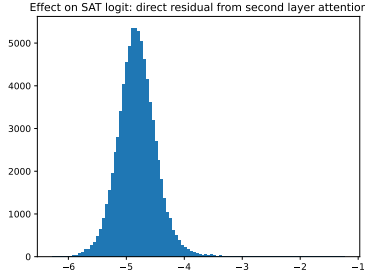
Figure 10: Effect of the residual connection from the second layer's attention mechanism on the SAT logit. x-axis shows the effect on the SAT logit and y-axis is the number of instances in the analysis test dataset.
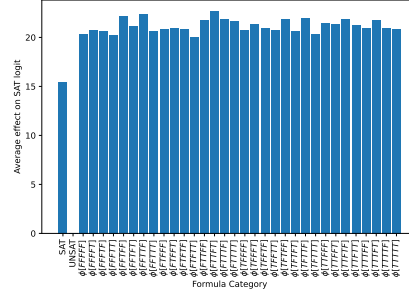

Figure 11: Average total effect on the SAT logit across all relevant neurons from SAT, UNSAT, and formulas satisfiable with particular assignments to the variables.

We see that the behavior of attention is nearly identical for a clause and its equivalent mirror image (i.e. the attention score on $(l \lor r)$ is approximately the same as that on $(r \lor l)$); this is as expected given that, logically, $l \lor r = r \lor l$. We can also see that each head prefers clauses containing certain literals (for instance, head 0 prefers negated literals as well as x_0 and x_3); each clause receives a high score from some head (so no clauses are overlooked by the attention mechanism).

While we might read far more into these patterns, this is a mistake in this case. In this instance, the behavior of attention is, in fact, not relevant to the underlying algorithm and serves only to obscure the underlying behavior. As discussed in Section 5.2, we can much better understand attention in concert with the first layer of the MLP; collectively, these components encode the observed evaluating behavior, as discussed in more detail below.

Identifying the Key Pathway. We observe that the unembedding vector for UNSAT is nearly exactly the negative of the unembedding vector for SAT ($\|W_U^{\text{SAT}} + W_U^{\text{UNSAT}}\| \approx 6.2 \times 10^{-6}$); hence, the SAT logit and the UNSAT logit are almost exactly negatives of each other; furthermore, the model's output token is always SAT or UNSAT. Across all formulas in the analysis dataset, the minimum value of the logit associated with the prediction (either SAT or UNSAT) is slightly over 0, while while the maximum logit on any other token across the dataset is below -9. Hence, we can consider the behavior of the SAT logit exclusively.

Next, we show that the key pathway that determines classification passes through the second layer's MLP. For an input formula, the effect of the second layer's attention mechanism through the residual connection on the SAT logit simply the dot product of the post-attention embedding in the readout token position with W_U^{SAT}; as we can see in Figure 10, this value is fairly consistent across the dataset and always negative; so a positive SAT logit and hence a SAT classification *must* depend on the action of the MLP.

Evaluation in the Hidden Neurons. As discussed in Section 5.2, only 34 hidden neurons have a significant effect on classification; hence, we focus on the behavior of these neurons. As observed in Figure 10, the effect of the residual stream from the second layer attention on the SAT logit reduces the SAT logit by less than 6 units for all formulas in the analysis test dataset. Furthermore, as observed in Section 5.2, the model always predicts SAT or UNSAT, and the SAT and UNSAT logits

are tied (in particular, they are negatives). Hence, the model will predict SAT when the collective effect of these 34 neurons is to increase SAT by > 6 units. We have observed that the expected behavior of these neurons is strongly dependent on *which* assignments to the variables satisfy the formula ϕ.

A natural question to ask is whether, whether these neurons will collectively increase the SAT logit sufficiently to output SAT on SAT formulas regardless of which assignment to the variables satisfies the formulas. Calculating this in the same way as we do for an individual neuron in Figure 5, we obtain Figure 11, which shows that, indeed, the collective effect of the evaluating neurons (recall that we refer to the 34 relevant neurons as evaluating neurons) correctly identifies SAT formulas regardless of *how* the formulas are satisfiable. It may seem odd that the average effect of SAT is significantly below the minimum of the effects for formulas satisfiable with any given assignment, as any SAT formula belongs to one such category; this is an artefact of the fact that no neuron has a large effect for all SAT formulas, and that each specializes in a subset.

Now, each of these neurons has a large positive (> 2.9) output coefficient (recall that the output coefficient of a neuron is the weight in the weighted sum expression constructed by composing the output layer of the MLP and the unembedding matrix projected to the SAT logit), so an activation above 2 units on any of the evaluating neurons is enough to force SAT, ignoring all other neurons; hence, the model outputs SAT if *any* evaluating neuron activates sufficiently strongly. In this sense, the effect of the final output components (the MLP's output layer, the unembedding matrix, and the residual from attention) can be viewed as an OR operation.

However, in reality, individual neurons do not consistently reach activations of ≥ 2 (for example, see Figure 5, in which the average output value for formulas satisfiable with assignment TFFFF is approximately 1.5), and multiple neurons recognize a given SAT formula. Hence, moderate activation may be necessary across several neurons to predict SAT in some cases; this may explain the redundancy in evaluation behavior between neurons seen in Appendix C. For simplicity in our interpretation (in particular, to allow replacing a numerical weighted sum with Boolean operations) we handle this by using varying thresholds between the corresponding α and γ functions for the MLP's hidden layer (in particular, γ_2 outputs an activation of 2 for neurons whose interpretation predicts high activation, while 0.5 is considered a high activation in the case of α_2); as noted in Section 5.2.1, if γ_2 and α_2 both use the lower threshold, replacing the neural network components with our interpretation no longer has a minimal effect on classification.

Interpreting Neurons via Decision Trees. We observe that each assignment to the variables has *some* neuron for which formulas satisfiable with that assignment result in a significant increase to the SAT logit on average, as we'd expect if the model was implementing the natural exhaustive enumeration algorithm for satisfiability checking. In particular, every assignment a has some neuron which increases the SAT logit by at least 2.9 units on average on formulas where a is a satisfying assignment; however, the average-case analysis hides the complexity of the behavior of these neurons. We can see this by noting that for every relevant neuron that on average demonstrates high activation for formulas satisfiable with some assignment a, we are able to find some formula satisfiable with a which fails to result in a sufficient activation for the neuron.

We'll focus on the general decision tree classifiers discussed in Section 5.2 here. We train decision tree classifiers on the thresholded activations of each evaluating neuron on the analysis training set; this allows us to learn an arbitrary Boolean function in the features $\phi[...]$. We observe that these decision trees do indeed reliably predict activation. In particular, even very limited-size decision
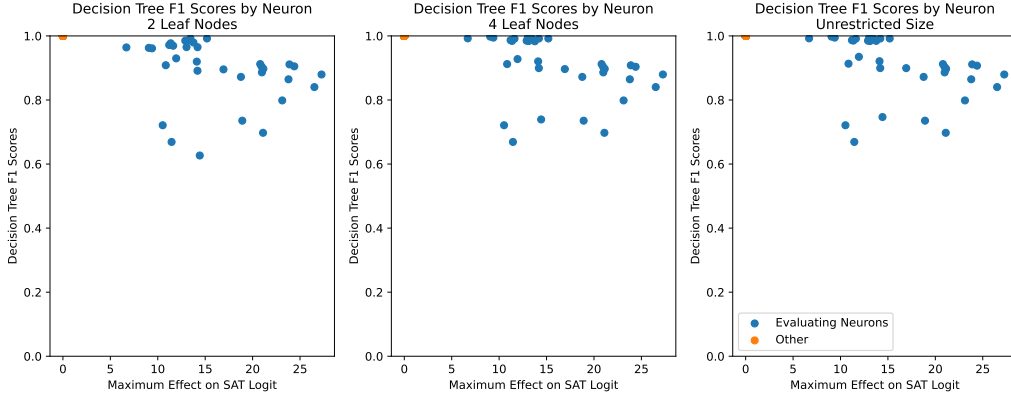
Figure 12: F1 scores of decision tree interpretations classifying high (> 0.5) activation of MLP hidden neurons.
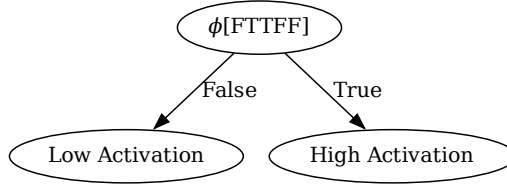

Figure 13: Decision tree interpretation for unit 29 of the second layer MLP's hidden layer, maximum leaf nodes: two.

trees are reasonably strong predictors of the behavior of the neurons; we limit our decision trees to four leaf nodes in our experiments. Figures 13, 14, and 15 illustrate the effect on the decision trees as we lift this constraint. Figure 12 shows that more complex decision trees are not significantly better predictors of the behavior of the neurons.

Implementation of Decision Trees by the Model. To see how the model implements the decision trees, we'll consider the effect of the composition of the second layer MLP's hidden layer with the attention mechanism on the representative embeddings of the clauses, using the abstract attention approach discussed earlier. As the pre-softmax attention score by each head on the second literal of a clause is a fixed value for each clause (recall that γ_1 leaves the embedding of the readout token constant), we can express the effect of attention precisely in terms of the number of occurrences of each clause, as we show next.

For head h and clause $l \vee r$, let the pre-softmax attention score paid by h to the clause be $w_{l \vee r}^h$. Moreover, the output of the OV circuit is likewise a fixed vector for each clause, call it $v_{l \vee r}^h$. Say that our formula contains a list of n clauses ($n = 10$ for all our experiments) where the i^{th} clause is $l_i \vee r_i$. Then, the attention paid to the i^{th} clause by head h is:

$$\text{softmax}([w_{l_j \vee r_j}^h]_{j \in [1..n]})_i = \frac{e^{w_{l_i \vee r_i}^h}}{\sum_{j=1}^n e^{w_{l_j \vee r_j}^h}}$$
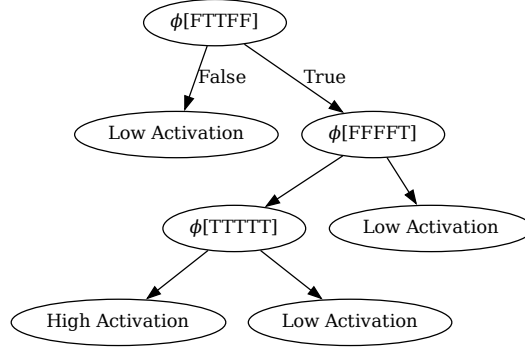
25

Figure 14: Decision tree interpretation for unit 29 of the second layer MLP's hidden layer, maximum leaf nodes: four.

and hence the output of the attention mechanism (call it o) at the readout token is, by rearranging terms,

$$e_: + \sum_{h=1}^{m} \sum_{i=1}^{n} \text{softmax}([w_{l_j \vee r_j}^h]_{j\in[1..n]})_i v_{l_i \vee r_i}^h = e_: + \sum_{h=1}^{m} \frac{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi) v_{l\vee r}^h}{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi)}$$

where $count_{l\vee r}(\phi)$ is the number of occurrences of $l \vee r$ in the formula ϕ and where m is the number of attention heads and where $e_:$ is the fixed representative embedding of the readout token ':'.

Now, consider the action of the MLP's hidden layer on this value: in particular, consider the effect on a specific neuron n, where w_n is the corresponding column of the MLP's input weight matrix and b_n is the corresponding bias value.

The pre-activation score for n will then be:

$$b_n + w_n^T o = b_n + w_n^T \left(e_: + \sum_{h=1}^{m} \frac{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi) v_{l\vee r}^h}{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi)} \right)$$

$$= b_n + w_n^T e_: + \sum_{h=1}^{m} \frac{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi) w_n^T v_{l\vee r}^h}{\sum_{l\in lit, r\in lit} e^{w_{l\vee r}^h} count_{l\vee r}(\phi)}$$

$$= c_n + \sum_{h=1}^{m} \frac{\sum_{l\in lit, r\in lit} c_{l\vee r}^{h,n} count_{l\vee r}(\phi)}{\sum_{l\in lit, r\in lit} d_{l\vee r}^h count_{l\vee r}(\phi)}$$

where we define the following constants: $c_n = b_n + w_n^T e_:$, $d_{l\vee r}^h = e^{w_{l\vee r}^h}$, $c_{l\vee r}^{h,n} = d_{l\vee r}^h w_n^T v_{l\vee r}^h$.

Note that the numerator and denominator of the fraction describing the contribution of head h are both linear functions of the number of ocurrences of each clause in the formula.

We'll briefly describe how $\phi[a]$ and $\neg\phi[a]$ can be implemented with a single neuron for an arbitrary assignment a to the variables in this formulation (in particular, we can ensure activation if and only if the condition that ϕ is satisfiable by a or not satisfiable by a holds), and then analyze the behavior of a neuron with an interpretation of this form.

26

φ[FTTFF]

False   True

Low Activation   φ[FFFFT]

φ[TTTTT]   φ[FTTFT]

High Activation   φ[FTTFT]   Low Activation   φ[TFFFT]

φ[FFTFF]   High Activation   High Activation   Low Activation

Low Activation   φ[TFTFT]

φ[TFTFF]   Low Activation
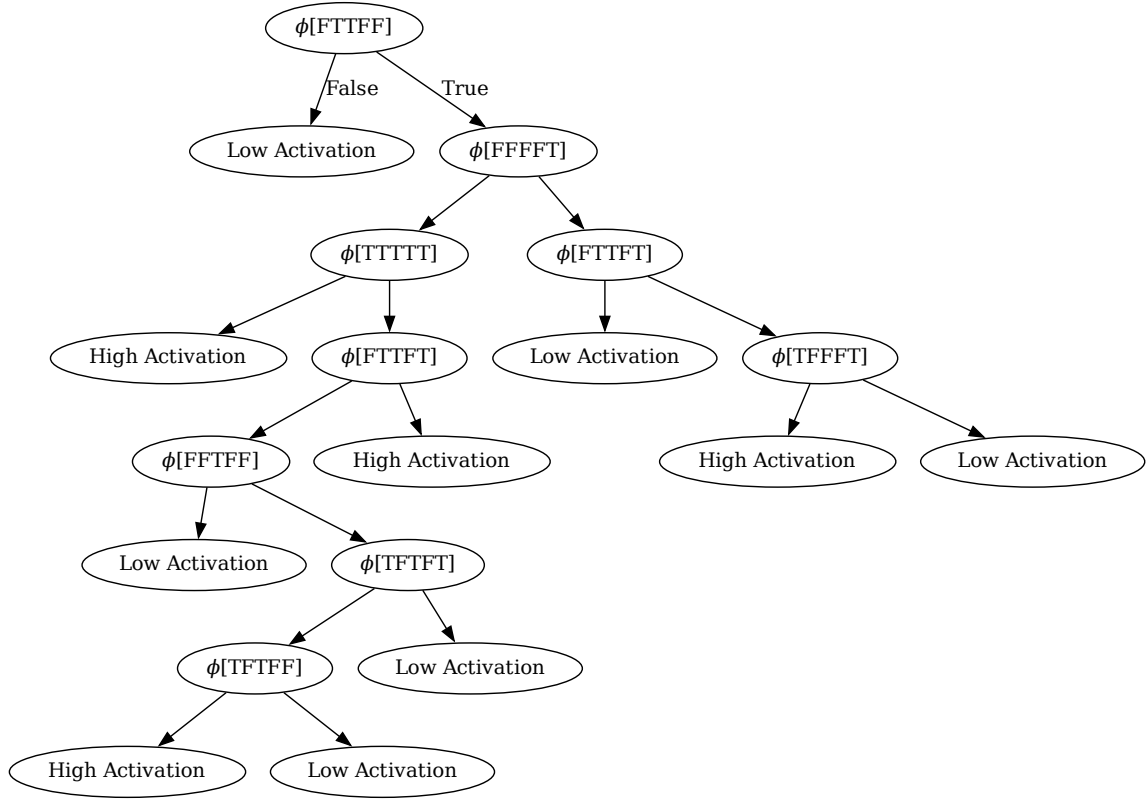
High Activation   Low Activation

Figure 15: Decision tree interpretation for unit 29 of the second layer MLP's hidden layer, unlimited leaf node count.

As the formula ϕ is in conjunctive normal form, $\phi[a]$ is true if and only if none of the clauses $l \vee r$ in ϕ evaluate to false given the assignment a to the variables (in particular, this holds if neither l nor r holds given a). We can then implement $\phi[a]$ in neuron n if we let $d^h_{l \vee r}$ be constant, $c^{h,n}_{l \vee r} = 0$ for $l \vee r$ which are true given assignment a, and let $c^{h,n}_{l \vee r} = -\infty$ for $l \vee r$ which are false given a, and then let c_n be a large positive number. Then, if $\phi[a]$, $count_{l \vee r}(\phi) = 0$ for $l \vee r$ unsatisfiable given a, so the output is c_n. Otherwise, if $\neg\phi[a]$, $count_{l \vee r}(\phi) > 0$ and so the output is $-\infty$. After applying ReLU, the activation of unit n is a large positive number when $\phi[a]$ and zero otherwise. Similarly, we can implement $\neg\phi[a]$ by assigning c_n to a negative number and letting $c^{h,n}_{l \vee r} = \infty$ for $l \vee r$ which are false given a.

Now, suppose that we want to implement some DNF expression e in the $\phi[...]$ without negation of any $\phi[...]$. Say that e contains k conjuncts e_i for $1 \le i \le k$ and that the attention mechanism has at least k heads. We can extend the ideas in the implementation of $\phi[a]$ to one for e as follows: let $c_n = k$; now, consider head h with corresponding conjunct e_i as in the $\phi[a]$ case, let $c^{h,n}_{l \vee r} = 0$ for $l \vee r$ which are not true for every assignment a appearing in e_i. As in that case, ϕ satisfies e_i if and only if no clause violates the constraint (following from the assumption that e_i is a conjunction of $\phi[a_j]$ which appear without negation, as, for an e_i of this form to be satisfied, each clause in ϕ must be true for each assignment a_j). Hence, if we let $c^{h,n}_{l \vee r} = -1$ for $l \vee r$ which violate e_i, let $d^h_{l \vee r} = 1$

27

```python
def evaluate_satisfiability(
    clauses: list[tuple[lit, lit]],
    neuron_interpretations: list[Callable[[list[tuple[lit, lit]]], bool]]
) -> list[bool]:
    return [
        neuron_interp(clauses)
        for neuron_interp in neuron_interpretations]

def predict_satisfiability(satisfiabilities: list[bool]) -> bool:
    return any(satisfiabilities)
```

Listing 2: Second layer's mechanistic interpretation. First, it applies a series of functions, namely, the neuron interpretations, to the parsed clauses from the first layer (`evaluate_satisfiability`) and then, it applies an OR operation (`predict_satisfiability`).

```python
def abstract_model(formula: list[tok]) -> bool:
    return predict_satisfiability(
            evaluate_satisfiability(
                parse_clauses(formula)))
```

Listing 3: Mechanistic interpretation of full model; see Listings 1 and 2 for the component interpretations.

for such $l \lor r$ and $d_{l \lor r}^h$ some negligible value for $l \lor r$ which do not violate the condition, the contribution of head h to the expression is 0 if e_i is satisfied and -1 otherwise (as the numerator will be $-v$ and the denominator will be $v + \epsilon$ where v is the number of clauses violating the condition). Hence, if all conjuncts e_i fail to be satisfied, the output will be 0, else, some e_i must be satisfied, so the output will be at least $k - (k - 1) = 1$, and, hence, the neuron will activate.

Interpretation. Listing 2 shows our derived mechanistic interpretation for the second layer. Listing 3 shows the mechanistic interpretation for the entire model.

C   Neuron Interpretations

Tables 1 and 2 show the decision tree based and disjunction-only neuron interpretations, respectively.

D   Limitations and Impacts

Limitations. A limitation of our work is that the proposed axioms are validated by intuitively motivating them and by comparison against past work [40]. We also demonstrate their usage via our case study of mechanistically interpreting a Transformer-based 2-SAT solver. While we believe that our axioms represent a minimal set that it necessary for characterizing a mechanistic interpretation, it may not be sufficient in all cases. In future work, we plan to address this and further validate our axioms by considering analysis of models trained for other problems. We are also able to state Axioms 5 and 6 only in a somewhat informal fashion.

Additionally, while our axioms are well-suited for mechanistic interpretability analyses that aim to understand the model behavior completely, they are not suitable in their present form for circuit analyses that only seek to understand specific aspects of model behavior. We believe that similar axioms can be formulated for circuit analyses but we leave this for future work.

Table 1: General neuron interpretations derived by decision tree learning on neuron activations. A neuron activates if the corresponding condition in the second column holds.

Neuron	Interpretation (General Case)
10	$\phi[TFFFF]$
29	$\phi[FTTFF] \wedge \neg\phi[FFFFT] \wedge \neg\phi[TTTTT]$
36	$\phi[TTFFF] \wedge \neg\phi[FTFTF]$
41	$\phi[FFFFF]$
48	$(\neg\phi[FTTFF] \wedge \phi[FTFFF]) \vee \phi[FTTFF]$
55	$\phi[TTTFF]$
77	$\phi[FFTTT]$
81	$\phi[FTTTF]$
96	$\phi[TFTTF] \wedge \neg\phi[FFTFF]$
150	$\phi[FTFTF] \wedge \neg\phi[TTFFF]$
185	$\phi[FFFTT]$
189	$\phi[FTTTT] \wedge \neg\phi[TTTFT]$
195	$\phi[TFFTF] \wedge \neg\phi[TTTTT]$
201	$\phi[TTFFT] \wedge \neg\phi[TFTFF]$
224	$\phi[FFFFT]$
231	$\phi[FTFTT] \wedge \neg\phi[FFTTF]$
245	$\phi[FTFFT]$
261	$\phi[TTFTT]$
291	$\phi[TFTTT]$
304	$\phi[FFTTF] \wedge \neg\phi[TFTFF]$
317	$\phi[TFFFT]$
326	$\phi[TTTTF]$
334	$\phi[FTTTF] \wedge \neg\phi[FFFTT]$
374	$\phi[TTFTT] \wedge \neg\phi[TFTTF] \wedge \neg\phi[FFFFF]$
380	$\phi[TFFTT]$
411	$\phi[TTFFT]$
416	$\phi[TTTFF]$
435	$(\phi[TTFTF] \wedge \neg\phi[FTFFF]) \vee (\phi[TTFTF] \wedge \phi[FTFFF] \wedge \phi[TFFTF])$
450	$\phi[FFTFT] \wedge \neg\phi[FTFFF]$
482	$\phi[TTTTT] \wedge \neg\phi[FTTFT]$
490	$(\phi[FTTFT] \wedge \neg\phi[TTTTT]) \vee (\phi[FTTFT] \wedge \phi[TTTTT] \wedge \neg\phi[TTTFT])$
492	$\phi[TFTFT]$
495	$\phi[FFFTF]$
499	$(\phi[FFTFF] \wedge \neg\phi[TFTTF]) \vee (\phi[FFTFF] \wedge \phi[TFTTF] \wedge \phi[FFTTF])$

Our techniques for analyzing and interpreting the 2-SAT model require human effort to implement and cannot be automatically applied for the analysis of other models. An exciting yet challenging direction for future research is to develop techniques for automated mechanistic interpretability analysis. We believe that the interpretability techniques that we use and describe (in Section 5 and Appendix B), such as variants of attention pattern analysis and automated learning of decision trees

as a component of the mechanistic interpretation, can be generalized to aid in the development of automated interpretability techniques.

Impacts. We believe that work on mechanistic interpretability in particular and on internal interpretability of neural networks in general is essential for increasing trust in neural networks and mitigating their risks. In this context, the work presented in this paper is of a foundational nature and advances the interpretability research agenda by clarifying the notion of a mechanistic interpretation. Therefore, we do not foresee a direct path from our work to negative societal impacts.

Table 2: Disjunction-only neuron interpretations derived by finding satisfying assignments that are correlated with a high average neuron activation value over the analysis training set. A neuron activates if the corresponding condition in the second column holds.

Neuron	Interpretation (Disjunction-Only)
10	$\phi[TFFFF] \vee \phi[TFTFF] \vee \phi[TTFFF]$
29	$\phi[FTTFF]$
36	$\phi[TTFFF]$
41	$\phi[FFFFF] \vee \phi[FTFFF]$
48	$\phi[FTFFF] \vee \phi[FTTFF]$
55	$\phi[TTTFF] \vee \phi[TTTFT]$
77	$\phi[FFTTF] \vee \phi[FFTTT] \vee \phi[FTTTT]$
81	$\phi[FTTTF] \vee \phi[FTTTT]$
96	$\phi[TFTTF]$
150	$\phi[FTFTF]$
185	$\phi[FFFTT]$
189	$\phi[FTTTT]$
195	$\phi[TFFTF] \vee \phi[TFTTF] \vee \phi[TTFTF]$
201	$\phi[TTFFT]$
224	$\phi[FFFFT]$
231	$\phi[FTFTT] \vee \phi[FTTTT]$
245	$\phi[FTFFF] \vee \phi[FTFFT] \vee \phi[FTTFT]$
261	$\phi[TTFTT]$
291	$\phi[TFTTF] \vee \phi[TFTTT] \vee \phi[TTTTT]$
304	$\phi[FFTTF]$
317	$\phi[TFFFT]$
326	$\phi[TTFTF] \vee \phi[TTTTF] \vee \phi[TTTTT]$
334	$\phi[FTTTF]$
374	$\phi[TTFTT]$
380	$\phi[TFFTT]$
411	$\phi[TTFFT] \vee \phi[TTTFT]$
416	$\phi[TFTFF] \vee \phi[TTTFF]$
435	$\phi[TTFTF]$
450	$\phi[FFTFF] \vee \phi[FFTFT] \vee \phi[FTTFT]$
482	$\phi[TTTTT]$
490	$\phi[FTTFT]$
492	$\phi[TFTFF] \vee \phi[TFTFT] \vee \phi[TTTFT]$
495	$\phi[FFFTF]$
499	$\phi[FFTFF]$