

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TOWARDS THE THREE-PHASE DYNAMICS OF GENERALIZATION POWER OF A DNN

Anonymous authors

Paper under double-blind review

## ABSTRACT

This paper addresses the core challenge in the field of symbolic generalization<sup>1</sup>, i.e., how to define, quantify, and track the dynamics of generalizable and non-generalizable interactions encoded by a DNN throughout the training process. Specifically, this work builds upon the recent theoretical achievement in explainable AI (Ren et al., 2024), which proves that the detailed inference patterns of DNNs can be strictly rewritten as a small number of AND-OR interaction patterns. Based on this, we propose an efficient method to quantify the generalization power of each interaction, and we discover a distinct three-phase dynamics of the generalization power of interactions during training. In particular, the early phase of training typically removes noisy and non-generalizable interactions and learns simple and generalizable interactions. The second and the third phases tend to capture increasingly complex interactions that are harder to generalize. Experimental results verify that the learning of non-generalizable interactions is the direct cause for the gap between the training and testing losses.

## 1 INTRODUCTION

Despite the rapid advancement of deep learning, a sophisticated theoretical understanding of the generalization power of deep neural networks (DNNs) remains elusive. In practice, the widely employed techniques for improving generalization are predominantly empirical, such as chain of thought (CoT) (Wei et al., 2022), data cleaning (Brown et al., 2020; Silcock et al., 2022), and large language model (LLM) alignment via reinforcement learning (Ouyang et al., 2022; Rafailov et al., 2023; Shao et al., 2024). Consequently, a gap persists between these practical methodologies and theoretical analyses of generalization (Foret et al., 2020; Li et al., 2018; Xiao et al., 2020).

Therefore, there is a clear trend towards more sophisticated and principled analysis of a DNN’s generalization power, which is emerged in recent years. **To this end, can the generalization power of a DNN be attributed to the representation quality of the inference patterns in the DNN?**

**Background: symbolic generalization.** The emerging field of symbolic generalization<sup>1</sup> has served as a pioneering response to the above question and has garnered considerable attention (Liu et al., 2023; Ren et al., 2023b;c; Zhou et al., 2024) (see Section 4 for a survey). Within this field, (Ren et al., 2024) have proven a counterintuitive phenomenon: **the complex inference logic of a DNN can be accurately and comprehensively explained by a small set of symbolic interactions.** As illustrated both in Figure 1 and Figure 7, an interaction represents an AND logic or an OR logic among input variables that exerts a specific numerical influence on the model’s output. For instance, a large language model (LLM) in Figure 1 encodes an AND interaction between  $S = \{\text{earth, revolves, sun}\}$ , which is only triggered when all the three words in  $S$  are presented in the model input and exerts a numerical effect  $I_T^{\text{and}} = 0.39$  that increases the LLM’s output confidence in generating the token “year.” Crucially, **(Ren et al., 2024) have proven such AND-OR interaction effects can accurately predict the DNN’s outputs on exponentially massive samples, which ensures the faithfulness of regarding these interactions as the primitive inference patterns encoded by the DNN.**

**Our research.** In this paper, we address a fundamental and long-standing challenge in the field of symbolic generalization<sup>1</sup>: **how to rigorously quantify the generalization power of different**

<sup>1</sup>The entire theory system of symbolic generalization contains more than 20 papers. Representative studies in this direction have been surveyed in Section 4.

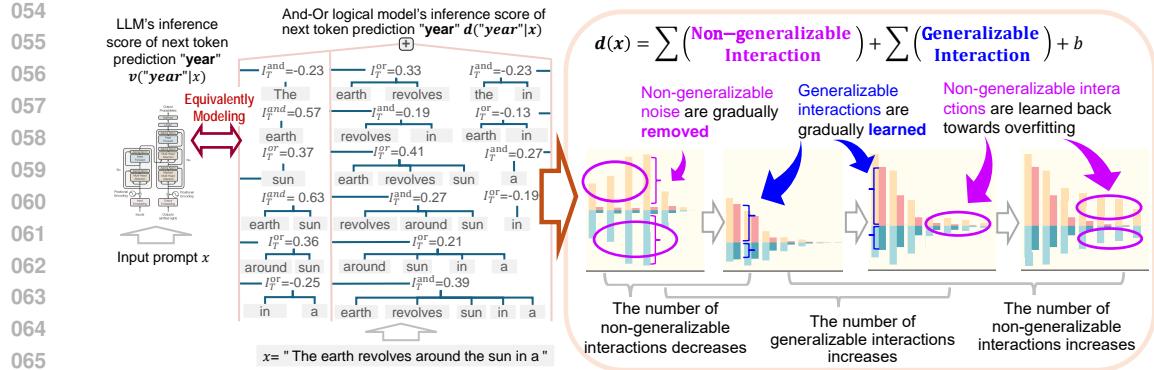


Figure 1: (Left) It has been proven that a small set of AND-OR interaction patterns are sufficient to mathematically represent all intricate inference patterns used by a DNN for inference. (Right) For each DNN, the change of its generalization power can be explained by the distinct three-phase dynamics of the generalization power of interactions encoded by it. Please see **Figure 7 in Appendix F** for more AND-OR logical models that explain LLMs' outputs on different prompts.

**interactions in a DNN and trace their evolution throughout the training process.** As Figure 7 shows, since the output of a DNN can be reformulated as the sum of all interaction effects, if most interactions learned by the DNN can also frequently appear in (be transferred to) unseen testing samples, then the DNN would exhibit a high testing accuracy, and vice versa. In other words, **the generalization power of interactions determines the generalization of the entire DNN.**

To this end, previous studies in the field of symbolic generalization remain limited to qualitative analysis of the complexity (Zhou et al., 2024) and robustness (Ren et al., 2023c) of interactions. People still lack a direct method for measuring the exact generalization power of each individual interaction, which has been a long-standing challenge for symbolic generalization for years.

Therefore, in this paper, we propose an efficient method to quantify the generalization power of each individual interaction. Instead of the computationally prohibitive approach of checking an interaction across all test samples, our method trains a baseline DNN on testing samples. All interactions, which can be transferred to the baseline DNN, are considered generalizable to testing samples. It is because each transferable interaction is also learned by the baseline DNN from testing samples as a primitive inference pattern.

The quantification of generalizable interactions allows us to **uncover the distinctive three-phase dynamics of generalization power of interactions through entire training process** (see Figure 3):

- Phase 1: Given a fully initialized DNN, in the early epochs of training, a large number of non-generalizable interactions are removed from the DNN, while a few generalizable interactions are gradually learned. These generalizable interactions are often simple and involve a small number of input variables, leading to the enhancement of the DNN's generalization power.
- Phase 2: The DNN continues to learn more interactions, but the newly learned interactions exhibit increasing complexity (*i.e.*, learning interactions between more input variables). These more complex interactions often have poorer generalization power. Thus, although the testing loss of the DNN continues to decrease, it tends to saturate due to the decreasing generalization power of the newly learned interactions.
- Phase 3: The DNN continues to learn additional interactions, but these newly learned interactions are often complex and difficult to apply to unseen data. Very few of these late-stage learned interactions represent simple, broadly generalizable patterns. Therefore, we can consider this phase revealing the overfitting of the DNN. During this phase, the gap between training and testing loss continues to widen. And the testing loss may even increase.

To further investigate how these non-generalizable interactions influence the DNN's performance, we propose to remove the effects of non-generalizable interactions from the DNN's output. We find that **the removal of non-generalizable interactions significantly reduces the gap between the training**

108 **loss and the testing loss, thereby effectively alleviating overfitting.** Notably, follow-up studies  
 109 in Appendix K have also shown that penalizing non-generalizable interactions during the training  
 110 process will improve the DNN’s performance.  
 111

## 112 2 METHODOLOGY

### 113 2.1 PRELIMINARIES: AND-OR INTERACTIONS

116 Consider an input sample  $\mathbf{x} = [x_1, x_2, \dots, x_n]^T$  comprising  $n$  input variables, where  $N =$   
 117  $\{1, 2, \dots, n\}$  denotes the index set. Let  $v(\mathbf{x}) \in \mathbb{R}$  denote the scalar output of the DNN. While  
 118  $v(\mathbf{x})$  can be defined in various ways<sup>2</sup>, we adopt the standard formulation widely used by (Ren et al.,  
 119 2024; Zhou et al., 2024; Liu et al., 2023) to let  $v(\mathbf{x})$  represent the classification confidence score for  
 120 the ground-truth label  $\mathbf{y}^*$  as

$$121 \quad v(\mathbf{x}) = \log \left( \frac{p(\mathbf{y}^* | \mathbf{x})}{1 - p(\mathbf{y}^* | \mathbf{x})} \right). \quad (1)$$

124 **Problem setting:** A central challenge in the field of symbolic generalization<sup>1</sup> is to elucidate the  
 125 complex inference patterns of a DNN (Li & Zhang, 2023b; Ren et al., 2024; Zhou et al., 2024). The  
 126 basic idea is to utilize a logical model  $d$  to explain all primitive inference patterns in the DNN  $v$ .  
 127 The faithfulness of this explanation is contingent upon satisfying two fundamental requirements.  
 128 **(1) Fidelity requirement:** The logical model  $d$  must faithfully replicate the DNN’s outputs over a  
 129 comprehensively massive input set  $\mathcal{Q}$ . **(2) Conciseness requirement:** The logical model  $d$  must  
 130 contain parsimonious patterns to allow for succinct explanations, as specified below:

$$131 \quad \forall \mathbf{x}' \in \mathcal{Q}, d(\mathbf{x}') = v(\mathbf{x}'), \text{ subject to } \text{Complexity}(d) \leq K, \quad (2)$$

132 where  $K$  is the maximum complexity allowed for logical model  $d$  to be considered concise.  
 133

134 Concretely, the logical model  $d$  constructed below is designed to encode a set of AND-OR interaction  
 135 patterns. *Please see a video demonstration for symbolic generalization in supplementary material.*  
 136 *Furthermore, please see Figure 7 in Appendix F for more logical models that explain LLMs.*

$$137 \quad d(\mathbf{x}') = \sum_{T \in \Omega_{\text{and}}} \underbrace{I_T^{\text{and}} \cdot \delta_{\text{and}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers AND relation} \\ \text{between input variables in } T \end{array} \right)}_{\text{an AND interaction between input variables in } T} + \sum_{T \in \Omega_{\text{or}}} \underbrace{I_T^{\text{or}} \cdot \delta_{\text{or}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers OR relation} \\ \text{between input variables in } T \end{array} \right)}_{\text{an OR interaction between input variables in } T} + b, \quad (3)$$

140 where  $\mathbf{x}'$  denotes the input  $\mathbf{x}$  where some input variables are masked<sup>3</sup> and,  $b$  is a scalar bias.  
 141

142 *Each AND interaction* is represented by binary function  $\delta_{\text{and}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers AND relation} \\ \text{between input variables in } T \end{array} \right) \in \{0, 1\}$ . It  
 143 returns 1 if all variables in  $T$  are not masked in  $\mathbf{x}'$ ; otherwise it returns 0. *Each OR interaction* is  
 144 represented by a binary function  $\delta_{\text{or}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers OR relation} \\ \text{between input variables in } T \end{array} \right) \in \{0, 1\}$ . It returns 1 if any variables in  
 145  $T$  are not masked in  $\mathbf{x}'$ ; otherwise it returns 0.  $I_T^{\text{and}}$  and  $I_T^{\text{or}}$  are scalar weights.  
 146

147 **First, the fidelity requirement is fulfilled by the universal matching property established in**  
 148 **Theorem 1.** This theorem demonstrates that a logical model with specific settings of  $I_T^{\text{and}}$  and  $I_T^{\text{or}}$  is  
 149 able to accurately predict all outputs of the DNN across an exponential number of  $2^n$  masked input  
 150 states. The comprehensively massive sample set is then constructed as  $\mathcal{Q} = \{\mathbf{x}_S \mid S \subseteq N\}$ , which  
 151 contains all  $2^n$  masked instances of the input sample  $\mathbf{x}$ .  $\mathbf{x}_S$  represents a masked input  $\mathbf{x}$ , in which  
 152 input variables in  $N \setminus S$  are masked<sup>3</sup>.

153 **Theorem 1 (universal matching property**, proven by (Chen et al., 2024)). *Given DNN  $v$  and an*  
 154 *input sample  $\mathbf{x}$ , let us set the scalar weights as  $\forall T \subseteq N$ ,  $I_T^{\text{and}} = \sum_{L \subseteq T} (-1)^{|T|-|L|} o_L^{\text{and}}$ ,  $I_T^{\text{or}} =$*   
 155  *$-\sum_{L \subseteq T} (-1)^{|T|-|L|} o_{N \setminus L}^{\text{or}}$  where  $o_L^{\text{and}} = 0.5 \cdot v(\mathbf{x}_L) + \gamma_L$  and  $o_L^{\text{or}} = 0.5 \cdot v(\mathbf{x}_L) - \gamma_L$ , and set*  
 156  *$b = v(\emptyset)$* <sup>4</sup>.  $\{\gamma_L\}$  are set of learnable parameters. The logical model can then accurately predict the

158 <sup>2</sup>E.g.,  $v(\mathbf{x})$  can be set to the feature dimension of the ground-truth category prior to the softmax operation.

159 <sup>3</sup>The masking of input variable means substituting this variable with a baseline value. We set the baseline  
 160 value of an input variable as the average value of this variable across different input samples. For NLP models,  
 161 we use a specific embedding in (Cheng et al., 2024; Ren et al., 2024) to mask input tokens.

162 <sup>4</sup> $v(\emptyset)$  represents the network output when all input variables in  $\mathbf{x}$  are masked.

162 *network outputs as follows, regardless of how  $\mathbf{x}$  is randomly masked:*

$$164 \quad \forall \mathbf{x}' \in \mathcal{Q}, \quad d(\mathbf{x}') = v(\mathbf{x}'). \quad (4)$$

165 **Computation of interactions.** Building upon the framework established in (Chen et al., 2024; Ren  
 166 et al., 2023b), the parameters  $\{\gamma_L\}$  are then optimized by minimizing a LASSO-regularized loss  
 167  $\min_{\{\gamma_L\}} \|\mathbf{I}_{\text{and}}\|_1 + \|\mathbf{I}_{\text{or}}\|_1$  to obtain the sparsest interaction patterns, where we vectorize interactions  
 168 as  $\mathbf{I}_{\text{and}} = [I_{T_1}^{\text{and}}, \dots, I_{T_{2^n}}^{\text{and}}]^T$ ,  $\mathbf{I}_{\text{or}} = [I_{T_1}^{\text{or}}, \dots, I_{T_{2^n}}^{\text{or}}]^T \in \mathbb{R}^{2^n}$ . Please refer to Appendix A for the  
 169 pseudocode for the computation of interactions and technical details.  
 170

171 **Second, the conciseness requirement is fulfilled by the sparsity<sup>5</sup> property of interactions.** Particular-  
 172 ly, (Ren et al., 2024) have further proven that under three common conditions<sup>6</sup>, DNNs can  
 173 encode only a small number ( $\mathcal{O}(n^p/\tau) \ll 2^{n+1}$ ) of AND-OR interactions with salient effects, while  
 174 all other interactions have almost zero effects. Specifically,  $\tau$  is a tiny threshold for saliency of  
 175 interactions, and empirically,  $p \in [0.9, 1.2]$ .

176 The above two requirements enable us to construct a **concise logical model with a few salient**  
 177 **AND-OR interactions in**  $\Omega_{\text{and}}^{\text{Salient}} = \{S \mid |I_S^{\text{and}}| > \tau\}$  and  $\Omega_{\text{or}}^{\text{Salient}} = \{S \mid |I_S^{\text{or}}| > \tau\}$ , which can  
 178 *faithfully match network outputs on all samples in  $\mathcal{Q}$ . The theoretical analysis in Appendix D and lots*  
 179 *of empirical results from diverse DNNs (including LLMs) in both Appendix F and (Li & Zhang, 2023b;*  
 180 *Ren et al., 2024; Zhou et al., 2024) support the the scientific rigor of the symbolic generalization.*

181 **The order of an interaction**  $S$  represents the complexity of this interaction. The order is defined as  
 182 the number of input variables involved in the interaction  $\text{Order}_S = |S|$ .  
 183

## 184 2.2 QUANTIFYING GENERALIZATION POWER OF INTERACTIONS

### 186 2.2.1 DEFINITION AND QUANTIFICATION

188 The above fidelity requirement and conciseness requirements have become foundational pillars for  
 189 the emerging direction of symbolic generalization (Li & Zhang, 2023b; Ren et al., 2023c; Zhou  
 190 et al., 2024). Crucially, studies in this direction introduce a paradigm shift in understanding the  
 191 generalization behavior of DNNs. This is grounded in the observation that a DNN’s predictive output  
 192 can be mathematically decomposed into a linear combination of AND-OR interaction effects (see  
 193 Figure 1). Consequently, a DNN’s generalization power can be considered to be determined by the  
 194 generalization power of the interactions in the DNN.

195 **Previous definition of the generalization power of interactions.** Let us first revisit the definition  
 196 of generalization power of an interaction in (Zhou et al., 2024). Given a salient AND interaction  
 197  $S$  extracted from an input  $\mathbf{x}$  subject to  $|I_S^{\text{and}}| > \tau$ , if this interaction  $S$  frequently occurs in the  
 198 testing set and consistently making an effect on the classification of a category, then it is deemed  
 199 to be generalizable; otherwise, it fails to generalize to testing samples. This criterion extends  
 200 analogously to OR interactions. For instance, in a bird species classification task, suppose an AND  
 201 interaction  $S = \{\text{red features, long beak}\}$  consistently appears in both training and testing datasets  
 202 and consistently increases the confidence in predicting the “*Flamingo*” class. In this case, the  
 203 interaction  $S$  demonstrates strong generalization power to unseen testing samples.

204 **Efficient quantification of the generalization power of interactions.** However, the above definition  
 205 does not provide an efficient method for quantifying the generalization power of interactions. For  
 206 example, in natural language processing tasks, this would necessitate an exhaustive search for specific  
 207 interactions across numerous testing samples, leading to prohibitive computational costs.

208 Therefore, we propose an approximate and efficient method **that first quantifies the exact generaliz-  
 209 ability of each individual interaction.** Instead of performing an exhaustive search, we evaluate the  
 210 transferability of individual interactions encoded by the DNN to a baseline DNN, denoted by  $v^{\text{base}}$   
 211 and trained on the testing samples. Since all representations in the baseline DNN  $v^{\text{base}}$  are acquired  
 212 from the testing samples, all interactions, which are transferable to  $v^{\text{base}}$ , can be regarded as primitive

213 <sup>5</sup>In the field of symbolic generalization, the sparsity is defined as the state that almost all interactions have  
 214 negligible values with only very few interactions having salient values. Please see Appendix B for details.

215 <sup>6</sup>(Ren et al., 2024) have demonstrated that the sparsity of interactions can be ensured through three common  
 216 conditions for smooth inferences in DNNs on randomly masked samples. See Appendix B for more details.

216 inference patterns inherent in these testing samples, thus strongly suggesting their generalization  
 217 power. Please refer to Appendix C for ablation study.

219 To elucidate this approach, let us consider an extracted<sup>7</sup> salient AND interaction  $S$ , s.t.  $|I_S^{\text{and}}| > \tau$ .  
 220 To this end, the interaction  $S$  is deemed to represent a generalizable pattern within the testing set  
 221 if the baseline DNN  $v^{\text{base}}$  concurrently identifies the AND interaction  $S$  to have salient effect (*i.e.*,  
 222  $|I_{S,v^{\text{base}}}^{\text{and}}| > \tau$ ) and the AND interaction  $S$  exerts a consistent directional influence on the classification  
 223 of  $\mathbf{x}$  (*i.e.*,  $I_{S,v^{\text{base}}}^{\text{and}} \cdot I_S^{\text{and}} > 0$ ). The same principle applies to OR interactions. Accordingly, the  
 224 generalization power of an AND/OR interaction is assessed by the following binary metrics:

$$\begin{aligned} \mathcal{G}_{S,v^{\text{base}}}^{\text{and}} &= \mathbf{1}(|I_{S,v^{\text{base}}}^{\text{and}}| > \tau) \cdot \mathbf{1}(I_S^{\text{and}} \cdot I_{S,v^{\text{base}}}^{\text{and}} > 0) \in \{0, 1\}, \\ \mathcal{G}_{S,v^{\text{base}}}^{\text{or}} &= \mathbf{1}(|I_{S,v^{\text{base}}}^{\text{or}}| > \tau) \cdot \mathbf{1}(I_S^{\text{or}} \cdot I_{S,v^{\text{base}}}^{\text{or}} > 0) \in \{0, 1\}, \end{aligned} \quad (5)$$

228 where  $\mathbf{1}(\cdot)$  is a binary indicator function that outputs 1 only when the condition is satisfied.

### 230 2.2.2 VERIFYING PREVIOUS FINDINGS ON THE GENERALIZATION POWER

232 The above metrics  $\mathcal{G}_{S,v^{\text{base}}}^{\text{and}}$  and  $\mathcal{G}_{S,v^{\text{base}}}^{\text{or}}$  allow us to further verify and elucidate decisive factors behind  
 233 the classical empirical finding in symbolic generalization<sup>1</sup> (Zhou et al., 2024), *i.e.*, low-order interac-  
 234 tions generally exhibit stronger generalization power compared to high-order interactions, which was  
 235 observed by (Zhou et al., 2024) in experiments<sup>8</sup> and lacked rigorous quantification and analysis.

236 Therefore, although the above claim regarding the low generalization power of high-order interactions  
 237 has been widely cited by (Chen et al., 2024; Cheng et al., 2024), this assertion still requires more  
 238 rigorous verification by identifying the generalization power of each specific salient interaction.

239 Specifically, to assess the generalization power of interactions<sup>7</sup> at distinct orders<sup>9</sup>, we utilize four  
 240 distinct metrics: (1)  $\mathbb{I}_{\text{total}}^{(m),+}$  denotes the aggregate effect strength of all salient interactions of the  $m$ -th  
 241 order with positive effect; (2)  $\mathbb{I}_{\text{total}}^{(m),-}$  denotes the aggregate effect strength of all salient interactions of  
 242 the  $m$ -th order with negative effect; (3)  $\mathbb{I}_{\text{gen}}^{(m),+}$  denotes the aggregate effect strength of all generalizable  
 243 salient interactions with positive effect; (4)  $\mathbb{I}_{\text{gen}}^{(m),-}$  denotes the aggregate effect strength of all  
 244 generalizable salient interactions with negative effect. The metrics are calculated as

$$\mathbb{I}_{\text{total}}^{(m),-} = \sum_{\substack{\text{type} \in \\ \{\text{and, or}\}}} \sum_{\substack{S \in \Omega_{\text{type}}^{\text{Salient}}: \\ |S|=m}} \max(I_S^{\text{type}}, 0), \quad \mathbb{I}_{\text{total}}^{(m),+} = \sum_{\substack{\text{type} \in \\ \{\text{and, or}\}}} \sum_{\substack{S \in \Omega_{\text{type}}^{\text{Salient}}: \\ |S|=m}} \min(I_S^{\text{type}}, 0), \quad (6)$$

$$\mathbb{I}_{\text{gen}}^{(m),-} = \sum_{\substack{\text{type} \in \\ \{\text{and, or}\}}} \sum_{\substack{S \in \Omega_{\text{type}}^{\text{Salient}}: \\ |S|=m}} \max(I_S^{\text{type}} \cdot \mathcal{G}_S^{\text{type}}, 0), \quad \mathbb{I}_{\text{gen}}^{(m),+} = \sum_{\substack{\text{type} \in \\ \{\text{and, or}\}}} \sum_{\substack{S \in \Omega_{\text{type}}^{\text{Salient}}: \\ |S|=m}} \min(I_S^{\text{type}} \cdot \mathcal{G}_S^{\text{type}}, 0). \quad (7)$$

253 To this end, we conducted experiments to analyze the strength of generalizable and non-generalizable  
 254 interactions in comparison across different interaction orders. We trained VGG-19 and ResNet-101  
 255 on the Tiny-ImageNet dataset, and VGG-16 on the CUB-200-2011 dataset. We also trained Bert-large  
 256 and Bert-medium on the SST-2 dataset and Bert-large on the AG News dataset. Please refer to  
 257 Appendix E for further experimental details. Figure 2 demonstrates that high-order interactions  
 258 consistently exhibited lower generalization power than low-order interactions in different cases.

## 260 3 EVALUATION

### 261 3.1 THREE-PHASE DYNAMICS OF GENERALIZATION POWER OF INTERACTIONS

263 Based on the new generalization power metrics  $\mathbb{I}_{\text{total}}^{(m),+}$ ,  $\mathbb{I}_{\text{total}}^{(m),-}$ ,  $\mathbb{I}_{\text{gen}}^{(m),+}$ ,  $\mathbb{I}_{\text{gen}}^{(m),-}$  proposed in Section  
 264 2.2, this subsection investigates the evolution of generalization power of interactions throughout the

266 <sup>7</sup>Interaction extraction is introduced in both Section 2.1 and the pseudocode in Appendix A.

267 <sup>8</sup>Evidence from two experiments substantiates this claim: (1) It is observed that, when more label noise is  
 268 introduced, overfitted DNNs tend to model more high-order interactions compared to standard DNNs trained on  
 269 clean data. (2) It is demonstrated that high-order interactions typically show worse robustness to input noise.

270 <sup>9</sup>As explained in Section 2.1, the order of an interaction refers to the number of input variables it involves.

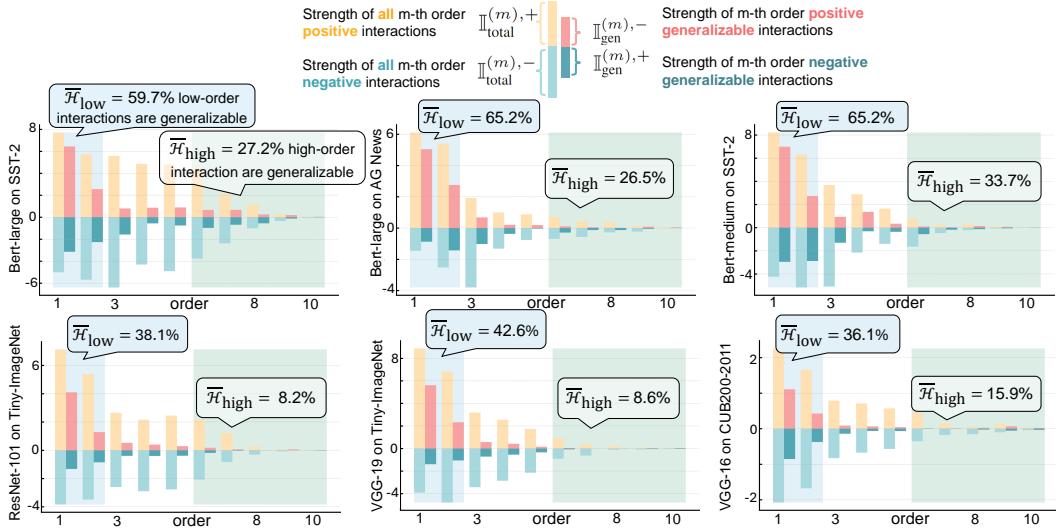


Figure 2: Comparison of generalization power between low-order interactions and high-order interactions. High-order interactions generally exhibit weaker generalization power than low-order interactions. Considering up to 10th order, we report the ratio of generalizable interactions of low orders as  $\bar{\mathcal{H}}_{\text{low}} = \sum_{m=1}^2 (|\mathbb{I}_{\text{gen}}^{(m),+}| + |\mathbb{I}_{\text{gen}}^{(m),-}|) / \sum_{m=1}^2 (|\mathbb{I}_{\text{total}}^{(m),+}| + |\mathbb{I}_{\text{total}}^{(m),-}|)$  and the ratio of generalizable interactions of high orders as  $\bar{\mathcal{H}}_{\text{high}} = \sum_{m=6}^{10} (|\mathbb{I}_{\text{gen}}^{(m),+}| + |\mathbb{I}_{\text{gen}}^{(m),-}|) / \sum_{m=6}^{10} (|\mathbb{I}_{\text{total}}^{(m),+}| + |\mathbb{I}_{\text{total}}^{(m),-}|)$ .

entire training process, which is recognized as a long-standing challenge in the field of symbolic generalization<sup>1</sup>. Specifically, we propose the following metric  $\bar{\mathcal{H}}$  to measure the average generalization power of all salient interactions on an input sample  $\mathbf{x}$  as follows:

$$\bar{\mathcal{H}} = \sum_m \left( |\mathbb{I}_{\text{gen}}^{(m),+}| + |\mathbb{I}_{\text{gen}}^{(m),-}| \right) / \sum_m \left( |\mathbb{I}_{\text{total}}^{(m),+}| + |\mathbb{I}_{\text{total}}^{(m),-}| \right). \quad (8)$$

We measure the number of interactions  $\bar{N}$  as the number of top-ranked salient interactions that account for 80.0% of the total interaction strength. Specifically, the number of AND interactions, denoted as  $\bar{N}_{\text{and}}$ , can be determined as follows:

$$\bar{N}_{\text{and}} = \underset{\Omega \subseteq 2^N}{\text{argmin}} |\Omega| \quad \text{s.t.} \quad \sum_{S \in \Omega} |I_S^{\text{and}}| \geq 0.8 \cdot \sum_{S \subseteq N} |I_S^{\text{and}}|. \quad (9)$$

where  $2^N = \{S \subseteq N\}$ . The number of OR interactions, denoted as  $\bar{N}_{\text{or}}$ , can be quantified through the same approach. Therefore, the overall interaction number  $\bar{N}$  is calculated as

$$\bar{N} = \bar{N}_{\text{and}} + \bar{N}_{\text{or}}. \quad (10)$$

**Three-phase dynamics.** Figure 3 shows that for most DNNs, the distribution of interactions follows a three-phase dynamics in terms of the generalization power  $\bar{\mathcal{H}}$  and the interaction number  $\bar{N}$ , which can be summarized as:

**Finding 1:** During the training of all DNNs, the number of interactions always exhibits a regularity of first decreasing and then increasing.

**Finding 2:** The average generalization power of interactions always demonstrates a regularity of first increasing and then decreasing.

**Finding 3:** The three-phase dynamics of interactions is temporally aligned with and well explains decisive factors behind the curve of the training-testing loss gap. Detailed dynamics is as follows.

• **Before Training:** Given a DNN with fully initialized parameters, although all interactions extracted from the initialized DNN represent random noise patterns, statistically, the interactions exhibit a specific spindle-shaped distribution. In words, the model primarily encodes middle-order interactions, while very high-order and very low-order interactions are rarely modeled. Furthermore, the effects of the positive interactions and the negative interactions often tend to cancel out. This phenomenon has been demonstrated in previous study (Cheng et al., 2025).

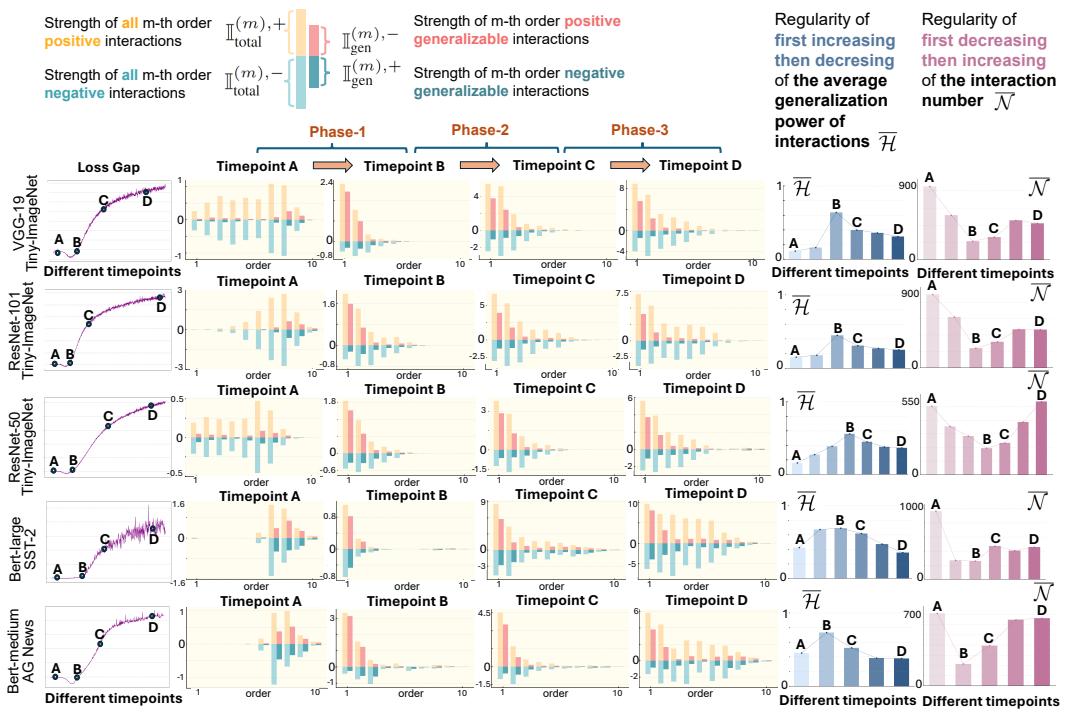


Figure 3: We visualize the change of interactions through the entire training process and track the change of the average generalization power of interactions  $\bar{\mathcal{H}}$  and the number of interactions  $\bar{N}$ . The dynamics of quantity and the average generalization power of interactions explains the change of the loss gap. Please see Appendix E for the training loss and the testing loss.

At this stage, we observe that most interactions cannot transfer to the baseline DNN. This corroborates that interactions encoded in an initialized DNN are noise patterns.

- **Phase 1:** In the early epochs of training, the average generalization power of interactions  $\bar{\mathcal{H}}$  usually increases, and the number of generalizable interactions  $\bar{\mathcal{N}}$  usually decreases. The training primarily eliminates noise in the spindle-shaped distribution while simultaneously learning a few low-order interactions. During this phase, the DNN mainly learns these low-order interactions, which tend to generalize better to the baseline DNN as they capture simple, meaningful patterns. Meanwhile, high-order interactions initially modeled by the DNN are gradually eliminated.

According to Equation (15), we can conclude that the classification confidence depends jointly on the number of salient interactions and the average generalization power of these interactions. Therefore, the emergence of generalizable low-order interactions and the elimination of non-generalizable high-order interactions characterize the healthiest stage of training.

- **Phase 2:** As training progresses, the average generalization power of interactions  $\bar{\mathcal{H}}$  gradually decreases, and the number of interactions  $\bar{\mathcal{N}}$  usually increases. Both generalizable (mainly low-order) interactions and non-generalizable (mainly mid-to-high-order) interactions are learned by the DNN simultaneously. Since the initial noise interactions have been removed after the first phase, in the second phase the DNN progressively learns interactions of increasing order. Notably, we find that high-order interactions are generally less likely to be generalized to testing samples than lower-order interactions according to Section 2.2.2. Thus, two phenomena co-occur in the second phase: (1) the DNN continues to learn additional interactions; (2) the generalization power of newly learned interactions steadily decreases.

Consequently, the second phase faces a dilemma. Although the continuous learning of new interactions provides meaningful features for classification, the diminishing generalization power of these interactions introduces a risk of overfitting. This observation is further corroborated by the training curves: The rate of decrease of the training loss outpaces that of the testing loss, resulting in a widening gap that indicates potential risk of overfitting of the model to the training data. This phe-



Figure 4: The change of the cross-entropy loss when we remove non-generalizable interactions from the network output  $v_{\text{removed}}(\mathbf{x})$ . The removal of non-generalizable interactions significantly increases the training loss, which restores the training loss that was inappropriately depressed by overfitted representations. In comparison, the testing loss either decreases significantly or does not change a lot.

nomenon may arise because non-generalizable interactions contribute exclusively to the classification of training samples and fail to aid in the classification of unseen testing samples.

• **Phase 3:** During the final training stage, DNNs primarily acquire non-generalizable (mainly mid-to-high-order) interactions, while it learns minimal generalizable (predominantly low-order) interactions. As a matter of fact, low-order interactions have already been thoroughly captured during the preceding two phases. Meanwhile, these emergent high-order interactions typically demonstrate poor generalization to unseen samples.

We hypothesize that these high-order interactions represent potential sources of overfitting. While generalizable interactions facilitate effective classification of unseen testing samples, non-generalizable interactions widen the gap between training and testing loss.

### 3.2 INFLUENCE OF NON-GENERALIZABLE INTERACTIONS ON THE CLASSIFICATION LOSS

In addition to the three-phase dynamics of interactions' generalization power, in this subsection, we propose an alternative perspective to examine if the non-generalizable<sup>10</sup> interactions are truly responsible for the overfitting of the DNN. The basic idea is that if we remove the effects of non-generalizable interactions from the scalar output of the DNN, then we can test **whether the removal of non-generalizable interactions will reduce the training-testing loss gap**. To accomplish this, let us rewrite the scalar output of the DNN in Equation (3) as (see Appendix H for proof)

$$\begin{aligned}
 v(\mathbf{x}) = & \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_{\text{type}}(\mathbf{x} \text{ triggers type relation between input variables in } T) & *_{\text{generalizable interactions}} \\
 & + \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_{\text{type}}(\mathbf{x} \text{ triggers type relation between input variables in } T) & *_{\text{non-generalizable interactions}} \\
 & + \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_{\text{type}}(\mathbf{x} \text{ triggers type relation between input variables in } T) + b & *_{\text{negligible non-salient interactions}},
 \end{aligned} \tag{11}$$

where  $\Omega_{\text{type}}^{\text{Non-Salient}} = \{S : |I_S^{\text{or}}| \leq \tau\}$  represents the set of non-salient interactions.

<sup>10</sup>When interaction  $S$  does not satisfy the generalization condition in Equation (5), *i.e.*, when interaction  $S$  does not have a salient effect or does not exert a consistent directional influence. And we have  $G_{S, v_{\text{base}}}^{\text{type}} = 0$ .

432 Then, the probability of classifying the input  $\mathbf{x}$  to the ground-truth label  $\mathbf{y}^*$  after the removal of  
 433 non-generalizable interactions can be derived from Equation (1) as follows:  
 434

$$435 \quad p_{\text{removed}}(\mathbf{y}^*|\mathbf{x}) = \text{sigmoid}(v_{\text{removed}}(\mathbf{x})) \\ 436 \quad \text{subject to } v_{\text{removed}}(\mathbf{x}) = v(\mathbf{x}) - \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}). \quad (12)$$

437 Then, given the revised classification probability  $p_{\text{removed}}(\mathbf{y}^*|\mathbf{x})$ , we can compute the new cross-  
 438 entropy loss as:  $L_{\text{removed}} = -\sum_{\mathbf{x}} \log p_{\text{removed}}(\mathbf{y}^*|\mathbf{x})$ .  
 439

440 We conducted experiments to train VGG-19, ResNet-50, and ResNet-101 on the Tiny-ImageNet  
 441 dataset, and Bert-large and Bert-medium on the SST-2 dataset, and Bert-large the AG News dataset.  
 442 We computed the standard cross-entropy loss using the original classification probability  $p(\mathbf{y}^*|\mathbf{x})$ ,  
 443 as well as a modified cross-entropy loss using the revised probability  $p_{\text{removed}}(\mathbf{y}^*|\mathbf{x})$ . Please see  
 444 Appendix J for details. As shown in Figure 4, **the removal of non-generalizable interactions**  
 445 **significantly reduced the training-testing loss gap**. This removal primarily increased the training  
 446 loss while significantly decreasing or having minimal impact on the testing loss. These results  
 447 suggested that the removed non-generalizable interactions were highly correlated with the overfitting  
 448 of the DNN, **as they exclusively contributed to training samples without generalizing to unseen**  
 449 **data**.  
 450

## 451 4 RELATED WORK: THEORY SYSTEM OF SYMBOLIC GENERALIZATION

452 *Symbolic explanation: a seemingly impossible task.* Studies (Deng et al., 2022; Li & Zhang, 2023b;  
 453 Ren et al., 2023b) show DNN inference can be reformulated via symbolic interactions between inputs.  
 454 Though counterintuitive, (Ren et al., 2023a) proved a logical model based on such interactions can  
 455 approximate DNN outputs on masked inputs. Further, (Ren et al., 2024) found DNNs encode only a  
 456 limited set of salient interactions under common conditions<sup>6</sup>. Complementing this, (Chen et al., 2024)  
 457 developed a method to extract generalizable interaction patterns consistent across models. These  
 458 works form a mathematical basis for interpreting DNNs via variable interactions.  
 459

460 *Characterizing DNN representational capacity.* Interaction theory also serves as a tool to analyze  
 461 DNN representation quality. (Ren et al., 2021) showed adversarial attacks target complex interactions,  
 462 with robustness decaying exponentially as interaction complexity grows. (Ren et al., 2023c) proved  
 463 Bayesian networks (Pearl, 1985) inherently struggle to model complex interactions. (Deng et al.,  
 464 2021) identified a DNN limitation: they capture very simple or highly complex interactions but falter  
 465 with intermediate ones. (Zhang et al., 2021) found dropout improves generalization by modulating  
 466 interactions. (Zhou et al., 2024) verified simpler interactions generalize better. (Liu et al., 2023) gave  
 467 theoretical evidence that DNNs find complex interactions harder to learn.  
 468

469 However, towards the core issue of the generalization power of interactions, previous studies in the  
 470 field of symbolic generalization all failed to quantify the generalization power of a specific interaction.  
 471 To this end, we propose to use a baseline DNN to identify generalizable interactions, which breaks  
 472 through the computational efficiency bottleneck. In this way, we discover a distinctive three-phase  
 473 dynamics in how generalization power evolves during training. Significantly, our research connects  
 474 DNN’s overfitting to the learning of non-generalizable interactions. Our discoveries establish a clear  
 475 causal relationship between interaction patterns and a DNN’s generalization power.  
 476

## 477 5 CONCLUSION

478 In this paper, we focus on the core issues in the field of symbolic generalization, *i.e.*, how to quantify  
 479 generalization power of each individual interaction in a DNN and how to utilize the evolution of  
 480 generalization power of interactions to diagnose the training process of a DNN. Specifically, we  
 481 propose an efficient method to quantify the generalization power of each individual interaction learned  
 482 by the DNN, so that we can use generalization power of all interactions to interpret the generalization  
 483 power of the entire DNN. Our analysis reveals a three-phase dynamics in the generalization power  
 484 of interactions, which well aligns the change of the DNN’s training-testing loss gap. In the beginning,  
 485 most non-generalizable interactions are eliminated, and the network primarily models simple,  
 486 generalizable interactions. In the intermediate phase, the network begins to capture increasingly less

486 generalizable interactions, and by the final phase, it predominantly learns entirely non-generalizable  
487 ones. Our experiments demonstrate that the learning of non-generalizable interactions is a direct  
488 cause for the gap between the training loss and the testing loss.

489 **Practical value.** Our experiments have demonstrated that removing non-generalizable interactions  
490 can reduce the training-testing loss gap. Besides, follow-up studies in Appendix K have also shown  
491 that penalizing the strength of non-generalizable interactions  $|I_S^{\text{NoneG}}|$  along with the minimization of  
492 the classification loss during the training process substantially improves the DNN’s performance.

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540 REFERENCES  
541

542 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
543 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
544 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

545 Hua Cai, Shuang Zhao, Liang Zhang, Xuli Shen, Qing-Song Xu, Weilin Shen, Zihao Wen, and Tianke  
546 Ban. Unilaw-r1: A large language model for legal reasoning with reinforcement learning and  
547 iterative inference. *ArXiv*, abs/2510.10072, 2025. URL <https://api.semanticscholar.org/CorpusID:282057924>.

548

549 Lu Chen, Siyu Lou, Benhao Huang, and Quanshi Zhang. Defining and extracting generalizable  
550 interaction primitives from DNNs. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=OCqyFVFNeF>.

551

552

553 Lei Cheng, Junpeng Zhang, Qihan Ren, and Quanshi Zhang. Revisiting Generalization Power of a  
554 DNN in Terms of Symbolic Interactions. *arXiv e-prints*, art. arXiv:2502.10162, February 2025.  
555 doi: 10.48550/arXiv.2502.10162.

556

557 Xu Cheng, Lei Cheng, Zhaoran Peng, Yang Xu, Tian Han, and Quanshi Zhang. Layerwise change of  
558 knowledge in neural networks. In *Forty-first International Conference on Machine Learning*, 2024.  
559 URL <https://openreview.net/forum?id=7zEoinErzQ>.

560

561 Huiqi Deng, Qihan Ren, Hao Zhang, and Quanshi Zhang. Discovering and explaining the representa-  
562 tion bottleneck of dnns. *International Conference on Learning Representations*, 2021.

563

564 Huiqi Deng, Qihan Ren, Hao Zhang, and Quanshi Zhang. Discovering and explaining the representa-  
565 tion bottleneck of dnns. *International Conference on Learning Representations*, 2022.

566

567 Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization  
568 for efficiently improving generalization. *Learning, Learning*, Oct 2020.

569

570 Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape  
571 of neural nets. *Advances in neural information processing systems*, 31, 2018.

572

573 Mingjie Li and Quanshi Zhang. Defining and quantifying and-or interactions for faithful and concise  
574 explanation of dnns. *arXiv preprint arXiv:2304.13312*, 2023a.

575

576 Mingjie Li and Quanshi Zhang. Does a neural network really encode symbolic concepts? In  
577 *International conference on machine learning*, pp. 20452–20469. PMLR, 2023b.

578

579 Dongrui Liu, Huiqi Deng, Xu Cheng, Qihan Ren, Kangrui Wang, and Quanshi Zhang. Towards the  
580 difficulty for a deep neural network to learn concepts of different complexities. In *Thirty-seventh  
581 Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=mZ3hnyL9bS>.

582

583 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
584 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
585 low instructions with human feedback. *Advances in Neural Information Processing Systems*,  
586 35:27730–27744, 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/hash/b1efde5ae831c26027e15d81b4b94ij.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/hash/b1efde5ae831c26027e15d81b4b94ij.pdf).

587

588 J. Pearl. *Bayesian Networks: A Model of Self-activated Memory for Evidential Reasoning*. Report  
589 (University of California, Los Angeles. Computer Science Dept.). UCLA Computer Science De-  
590 partment, 1985. URL <https://books.google.com.hk/books?id=1sfMOgAACAAJ>.

591

592 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and  
593 Chelsea Finn. Direct preference optimization: Your language model is secretly a reward  
594 model. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.),  
595 *Advances in Neural Information Processing Systems*, volume 36, pp. 53728–53741. Curran Asso-  
596 ciates, Inc., 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf).

594 Jie Ren, Die Zhang, Yisen Wang, Lu Chen, Zhanpeng Zhou, Yiting Chen, Xu Cheng, Xin  
 595 Wang, Meng Zhou, Jie Shi, and Quanshi Zhang. Towards a unified game-theoretic view  
 596 of adversarial perturbations and robustness. In A. Beygelzimer, Y. Dauphin, P. Liang, and  
 597 J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL  
 598 <https://openreview.net/forum?id=fMaIxda5Y6K>.

599 Jie Ren, Mingjie Li, Qirui Chen, Huiqi Deng, and Quanshi Zhang. Defining and quantifying the  
 600 emergence of sparse concepts in dnns. In *Proceedings of the IEEE/CVF conference on computer*  
 601 *vision and pattern recognition*, pp. 20280–20289, 2023a.

602 Jie Ren, Zhanpeng Zhou, Qirui Chen, and Quanshi Zhang. Can we faithfully represent absence states  
 603 to compute shapley values on a dnn? In *International Conference on Learning Representations*,  
 604 2023b.

605 Qihan Ren, Huiqi Deng, Yunuo Chen, Siyu Lou, and Quanshi Zhang. Bayesian neural networks avoid  
 606 encoding complex and perturbation-sensitive concepts. In *International Conference on Machine*  
 607 *Learning*, pp. 28889–28913. PMLR, 2023c.

608 Qihan Ren, Jiayang Gao, Wen Shen, and Quanshi Zhang. Where we have arrived in proving the  
 609 emergence of sparse symbolic concepts in ai models. *International Conference on Learning*  
 610 *Representations*, 2024.

611 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 612 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the Limits of  
 613 Mathematical Reasoning in Open Language Models. *arXiv e-prints*, art. arXiv:2402.03300,  
 614 February 2024. doi: 10.48550/arXiv.2402.03300.

615 Emily Silcock, Luca D’Amico-Wong, Jinglin Yang, and Melissa Dell. Noise-robust de-duplication at  
 616 scale. *International Conference on Learning Representations*, 2022.

617 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
 618 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
 619 *neural information processing systems*, 35:24824–24837, 2022.

620 Lechao Xiao, Jeffrey Pennington, and Samuel Schoenholz. Disentangling trainability and generaliza-  
 621 tion in deep neural networks. In *International Conference on Machine Learning*, pp. 10462–10472.  
 622 PMLR, 2020.

623 Hao Zhang, Sen Li, Yinchao Ma, Mingjie Li, Yichen Xie, and Quanshi Zhang. Interpreting and boost-  
 624 ing dropout from a game-theoretic view. *International Conference on Learning Representations*,  
 625 2021.

626 Huilin Zhou, Hao Zhang, Huiqi Deng, Dongrui Liu, Wen Shen, Shih-Han Chan, and Quanshi Zhang.  
 627 Explaining generalization power of a dnn using interactive concepts. In *Proceedings of the AAAI*  
 628 *Conference on Artificial Intelligence*, volume 38, pp. 17105–17113, 2024.

629  
 630  
 631  
 632  
 633  
 634  
 635  
 636  
 637  
 638  
 639  
 640  
 641  
 642  
 643  
 644  
 645  
 646  
 647

648 APPENDIX  
649

650  
651 This appendix provides detailed information supporting our main paper. Appendix A describes  
652 the details of computing interactions. Appendix B discusses three common conditions for the  
653 sparsity of interactions. Appendix C provides ablation study on the number of baseline DNNs.  
654 Appendix E contains experimental details. Appendix H presents proof of Equation 12 in the main  
655 paper. Appendix J details our experiments on removing non-generalizable interactions from model  
656 outputs, including sample selection methods in multiple datasets. Appendix L describes the computing  
657 resources and experimental time consumed for computing in this study.  
658

659 A COMPUTATION OF INTERACTIONS  
660661 A.1 SETTING OF INTERACTION COMPUTATION  
662

663 For an input sample  $\mathbf{x}$  with  $n$  distinct input variables, the total number of subsets to evaluate scales  
664 as  $2^n$ , rendering exhaustive computation intractable when each pixel (in vision) or each token (in  
665 language) is treated as an individual variable. To mitigate this, we adopt the protocol of prior work  
666 (Li & Zhang, 2023b; Liu et al., 2023) by restricting the interaction analysis to exactly 10 variables, as  
667 detailed below.

668 **Image task.** We conduct experiments on image classification using convolutional networks (ResNet-  
669 50, ResNet-101, VGG-19). From a chosen intermediate feature map of each model (the first Relu  
670 layer), we partition the spatial dimensions into an  $5 \times 5$  grid and sample 10 patches to form the set  
671  $N$ . We set the baseline value of each input variable to its empirical mean across the dataset. For any  
672 subset  $S \subseteq N$ , a masked input  $\mathbf{x}_S$  is generated by replacing all patches in  $N \setminus S$  with  $b$ , thereby  
673 enabling efficient estimation of interaction strengths.

674 **Natural language task.** For text classification task, we employ BERT-Medium and BERT-Large  
675 architectures. For each input sentence, stop words, punctuation, and non-alphanumeric symbols  
676 removed. In this way, we then randomly select 10 tokens with semantic meanings to constitute  
677  $N$ . To construct each masked sample, we follow the masking embedding scheme of (Cheng et al.,  
678 2024; Ren et al., 2024): for BERT models, every token in  $N \setminus S$  is substituted with the pre-trained  
679 [MASK] embedding vector (ID=103), yielding  $\mathbf{x}_S$  and facilitating computation of higher-order token  
680 interactions.

681  
682 **Algorithm 1** Extract Interactions

---

683 **Input:** Trained DNN  $v(\cdot)$ ; input sample  $\mathbf{x}$  (features indexed  $1, \dots, n$ ; let  $N = \{1, \dots, n\}$  be the  
684 set of these indices).

685 **Constants assumed: number of top interactions  $k$ .**

686 **Output:** Interaction weights  $I_S^{\text{and}}, I_S^{\text{or}}$  for all  $S \subseteq N$ ; bias  $b$ ; sets of top- $k$  salient interactions  
687  $\Omega_{\text{and}}^{\text{Salient}}, \Omega_{\text{or}}^{\text{Salient}}$ .

688 **for all**  $S \subseteq N$  **do**  
689     Generate masked input  $\mathbf{x}_S$  and the corresponding model output  $v(\mathbf{x}_S)$ .  
690 **end for**  
691 **for all**  $S \subseteq N$  **do**  
692     Determine  $I_S^{\text{and}}, I_S^{\text{or}}$  by solving  $\min_{\{\delta_S, \gamma_S\}} \sum_{S \subseteq N} (|I_S^{\text{and}}| + |I_S^{\text{or}}|)$   
693 **end for**  
694     Determine the bias term  $b$  as the network's output for the empty set of features:  $b \leftarrow v(\emptyset)$ .  
695 **for all**  $S \subseteq N$  **do**  
696     Filter  $|I_S^{\text{and}}|$  and  $|I_S^{\text{or}}|$  to identify the most salient  $k$  interactions to construct  $\Omega_{\text{and}}^{\text{Salient}}$  and  $\Omega_{\text{or}}^{\text{Salient}}$   
697 **end for**  
698     **return** the computed interactions  $I_S^{\text{and}}, I_S^{\text{or}}$ , the bias  $b$ , and the sets of salient interactions  
699  $\Omega_{\text{and}}^{\text{Salient}}, \Omega_{\text{or}}^{\text{Salient}}$ .

---

702 A.2 PRESUCODE FOR COMPUTATION OF INTERACTIONS  
703

704 Additionally, drawing upon insights from (Chen et al., 2024; Ren et al., 2024; Li & Zhang, 2023a),  
705 the presence of noise  $\sigma_S$  in the model output  $v(\mathbf{x})$  can significantly impede the accurate computation  
706 of interactions. Following the approach outlined in (Li & Zhang, 2023a) to address this challenge,  
707 we remove the noise  $\sigma_S$  from the output  $v(\mathbf{x})$  to enable the extraction of AND-OR interactions.  
708 More specifically, the decomposition  $v(\mathbf{x}_S) = o_S^{\text{and}} + o_S^{\text{or}} + \sigma_S$  is reformulated such that  $o_S^{\text{and}} =$   
709  $0.5 \cdot (v(\mathbf{x}_S) - \sigma_S) + \gamma_S$  and  $o_S^{\text{or}} = 0.5 \cdot (v(\mathbf{x}_S) - \sigma_S) - \gamma_S$ . Consequently, the parameters  $\{\gamma_S\}$   
710 and  $\{\sigma_S\}$  are optimized within a LASSO-like loss framework. This process is designed to enforce  
711 sparsity and mitigate the influence of noise on the model output, as detailed in pseudocode 1.  
712

713 B THREE COMMON CONDITIONS FOR THE SPARSITY OF INTERACTIONS  
714

715 In the field of symbolic generalization, the sparsity is defined as the state that almost all interactions  
716 have negligible values with only very few interactions having salient values. And (Ren et al., 2024)  
717 have proven the sparsity of interactions under three common conditions as follows.

718 **Condition 1: finite-Order interactions**  
719

720 The network's output, denoted by a function  $v$ , can be fully characterized by interactions involving at  
721 most  $M$  input variables. This implies that any interaction effects  $I(S)$  for subsets of input variables  
722  $S \subseteq \{1, \dots, n\}$  where the size of the subset  $|S| > M$  are zero:  $I(S) = 0$  for all  $S$  such that  $|S| >$   
723  $M$ . Equivalently, in a Taylor expansion of the function  $v$ , all terms corresponding to mixed partial  
724 derivatives of order  $M + 1$  or higher vanish. Specifically, for any point  $b \in \mathbb{R}^n$  and any non-negative  
725 integers  $\kappa_1, \dots, \kappa_n$  such that  $\sum_{i=1}^n \kappa_i \geq M + 1$ , we have:  $\frac{\partial^{\kappa_1 + \dots + \kappa_n} v}{\partial x_1^{\kappa_1} \dots \partial x_n^{\kappa_n}} \Big|_{x=b} = 0$ . Such a constraint  
726 limits model complexity to lower-order effects, aligning with observations in some large models  
727 where very high-order interaction strengths are minimal, thereby directly supporting the premise of  
728 sparse high-order interactions.  
729

730 **Condition 2: monotonicity of average network output**  
731

732 Let  $v(x_S)$  be the network's output when only the features in subset  $S$  are revealed (and others are  
733 masked or set to a baseline  $x_\emptyset$ ). Define the average increase in network output when exactly  $m$   
734 features are revealed as:  $\bar{u}(m) = \mathbb{E}_{S \subseteq \{1, \dots, n\}: |S|=m} [v(x_S) - v(x_\emptyset)]$ . This condition requires that the  
735 average network output is monotonically non-decreasing with the number of revealed features. That  
736 is, for any  $m' \leq m$ :  $\bar{u}(m') \leq \bar{u}(m)$ . Average monotonic behavior suggests feature interactions have  
737 underlying simplicity. Models dominated by many complex high-order interactions would unlikely  
738 show such stability. This predictable response pattern indicates sparse, lower-order interactions drive  
739 outputs rather than numerous erratic high-order effects, indirectly showing the interaction structure  
740 isn't arbitrarily complex.

741 **Condition 3: Polynomial lower bound on average output under masking**  
742

743 There exists a positive constant  $p > 0$  such that for any number of revealed features  $m'$  and  $m$   
744 with  $m' \leq m$ , the average network output  $\bar{u}(m')$  (as defined under Condition 2) is lower-bounded  
745 as follows:  $\bar{u}(m') \geq \left(\frac{m'}{m}\right)^p \bar{u}(m)$ . The polynomial bound on output degradation when masking  
746 features suggests an interaction structure without dense high-order terms. If complex high-order  
747 interactions were critical, feature removal would cause more drastic drops. This supports the sparsity  
748 hypothesis for impactful high-order interactions.  
749

750 C ABLATION STUDY TOWARDS THE NUMBER OF BASELINE DNNs  
751

752 When we delve deeper into the problem of interaction generalization power, we observe that there  
753 is no clear boundary between generalizable and non-generalizable interactions. Consequently, the  
754 quantification of interaction generalization power is caught in the following dilemma.  
755

- (1) On one hand, we can only guarantee that all interactions modeled by a baseline DNN are learned  
from the testing samples. However, we cannot ensure that the baseline DNN models all possible

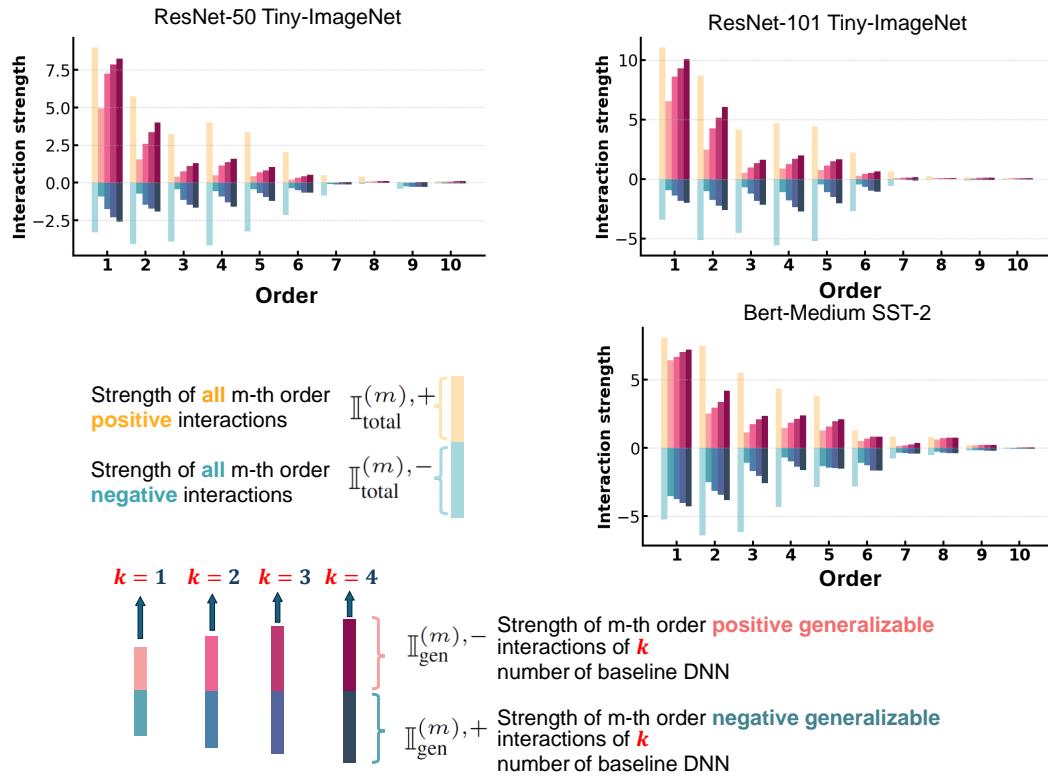


Figure 5: Ablation results of distribution of generalizable interactions with varying numbers of baseline DNNs. The distributions of interactions across different orders demonstrate a positive relationship wherein the strength of interactions increases proportionally with increasing values of  $k$

interactions. This issue necessitates training multiple baseline DNNs to cover as many different generalizable interactions as possible.

(2) On the other hand, given a group of baseline DNNs  $v_1^{\text{base}}, v_2^{\text{base}}, \dots, v_K^{\text{base}}$ , each generating a set of salient AND-OR interactions  $\Omega_{\text{and/or},k}^{\text{Salient}}$ , determining the generalization power of a target interaction becomes a fuzzy problem. If an interaction can only be generalized to one baseline DNN out of all  $K$  baseline DNNs, whether this interaction is truly a generalizable one remains an open problem.

Therefore, we conducted ablation experiments to illustrate the generalizable interactions that were identified by using different numbers of baseline DNNs. Given  $k$  baseline DNNs, a generalizable AND/OR interaction  $S$  was defined by its ability to generalize to any one baseline DNN, as follows.

$$\hat{G}_S^{\text{type}}(k) = \mathbf{1} \left( \underbrace{G_{S,v_1^{\text{base}}}^{\text{type}} = 1 \vee G_{S,v_2^{\text{base}}}^{\text{type}} = 1 \vee \dots \vee G_{S,v_k^{\text{base}}}^{\text{type}} = 1}_{S \text{ generalized to any one of the baseline DNNs}} \right) \in \{0, 1\} \quad (13)$$

where  $\text{type} \in \{\text{and, or}\}$  corresponds to AND/OR interactions, and  $\mathbf{1}(\cdot)$  is a binary indicator function that outputs 1 if the condition is satisfied and 0 otherwise. Figure 5 shows the distribution of generalizable interactions and that of non-generalizable interactions under different numbers of baseline DNNs for different models. As evident from the figure, the proportion of generalizable interactions increases consistently with the number of baseline DNNs across all interaction orders. This observation aligns with our theoretical understanding that employing more baseline DNNs enhances the detection capability for generalizable interactions. These results empirically validate our approach of using multiple baseline DNNs to comprehensively capture the spectrum of generalizable interactions, thereby addressing the dilemma outlined earlier. Notably, it is evident that selecting  $k = 2$  for our experiments offers a favorable balance between effectiveness and computational efficiency.

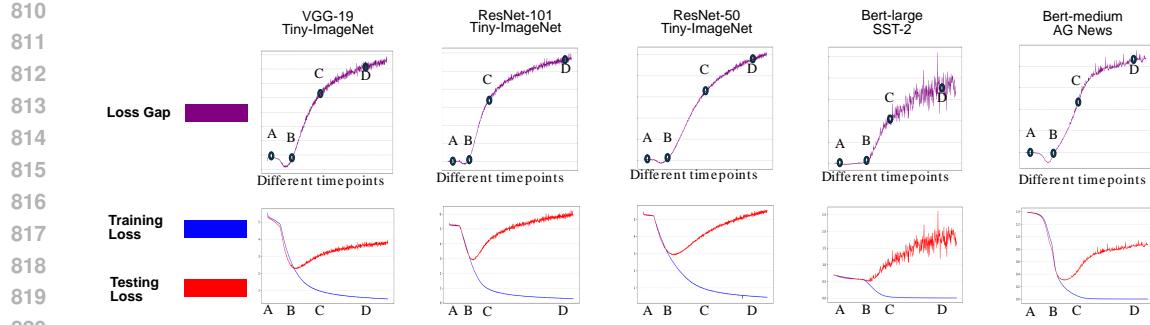


Figure 6: Training loss and testing loss for models.

## D LOGICAL MODEL ANALYSIS

The fidelity and conciseness in 2.1 requirements enable us to construct a concise logical model with only a few salient AND-OR interactions as below:

$$\forall \mathbf{x}' \in \mathcal{Q}, \quad v(\mathbf{x}') \approx d'(\mathbf{x}'), \quad (14)$$

$$d'(\mathbf{x}') = \sum_{T \in \Omega_{\text{Salient}}^{\text{and}}} I_T^{\text{and}} \cdot \delta_{\text{and}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers AND relation} \\ \text{between input variables in } T \end{array} \right) + \sum_{T \in \Omega_{\text{Salient}}^{\text{or}}} I_T^{\text{or}} \cdot \delta_{\text{or}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers OR relation} \\ \text{between input variables in } T \end{array} \right) + b. \quad (15)$$

## E EXPERIMENTAL DETAILS

### E.1 ARCHITECTURES AND DATASETS

In our comprehensive experimental framework, we employed a diverse range of deep learning architectures across multiple benchmark datasets to systematically analyze the strength of generalizable and non-generalizable interactions across varying interaction orders. For computer vision tasks, we utilized several convolutional neural network architectures: VGG-19 and ResNet-101/Resnet-50 were trained on the Tiny-ImageNet dataset, which contains 200 classes with 500 training images per class, downsampled to 64×64 pixels. Additionally, we trained VGG-16 on the CUB-200-2011 dataset (Caltech-UCSD Birds), which consists of 11,788 images across 200 bird species. For natural language processing tasks, we leveraged transformer-based models, specifically BERT-large (24 layers, 1024 hidden dimensions) and BERT-medium (8 layers, 512 hidden dimensions), which were trained on the Stanford Sentiment Treebank (SST-2) dataset containing binary sentiment classifications of movie reviews. Furthermore, we extended our NLP experiments by training both BERT-large and BERT-medium variants on the AG News dataset, which comprises news articles categorized into four classes.

For the training image preprocessing, the CUB-200-2011 dataset underwent bounding box processing to focus on the bird regions within each image, eliminating background noise that could potentially affect classification performance. The training images were then processed using data augmentation techniques including random resized cropping to 224×224 pixels, random horizontal flipping, tensor conversion, and normalization using ImageNet statistics. Similarly, Tiny-ImageNet images were preprocessed using data augmentation techniques, ensuring consistent input dimensions across all vision models despite the original differences in image resolutions between datasets.

To facilitate the training of both an original DNN and a corresponding baseline DNN, we randomly sampled and split the original training datasets to create two balanced subsets with equivalent data distributions. These equally distributed training subsets were used separately to train the original and baseline DNNs with different initialized parameters, maintaining a controlled experimental environment and ensuring that any observed differences between models could be attributed to our specific experimental manipulations rather than to variations in training data distribution.

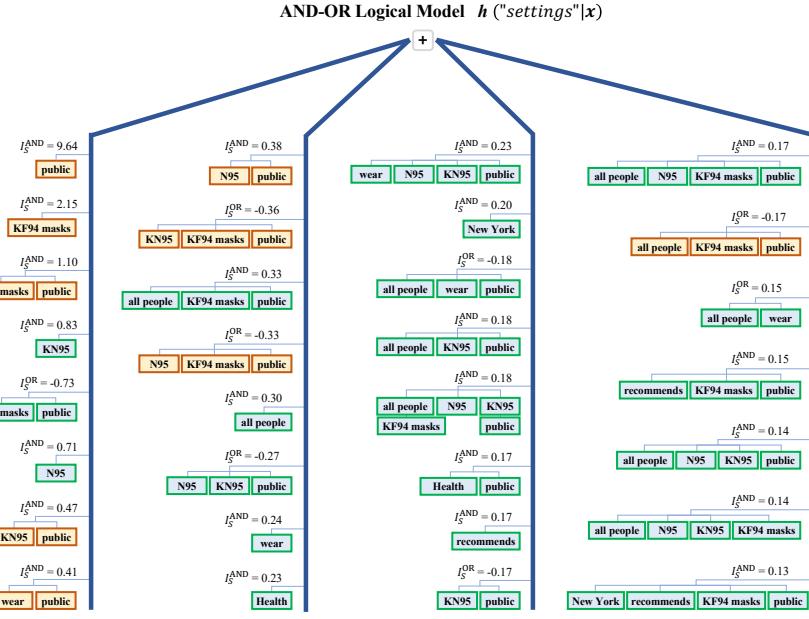
### E.2 TRAINING LOSS AND TESTING LOSS

We report the training and testing loss for the model used in Section 3.3 as shown in Figure 6.

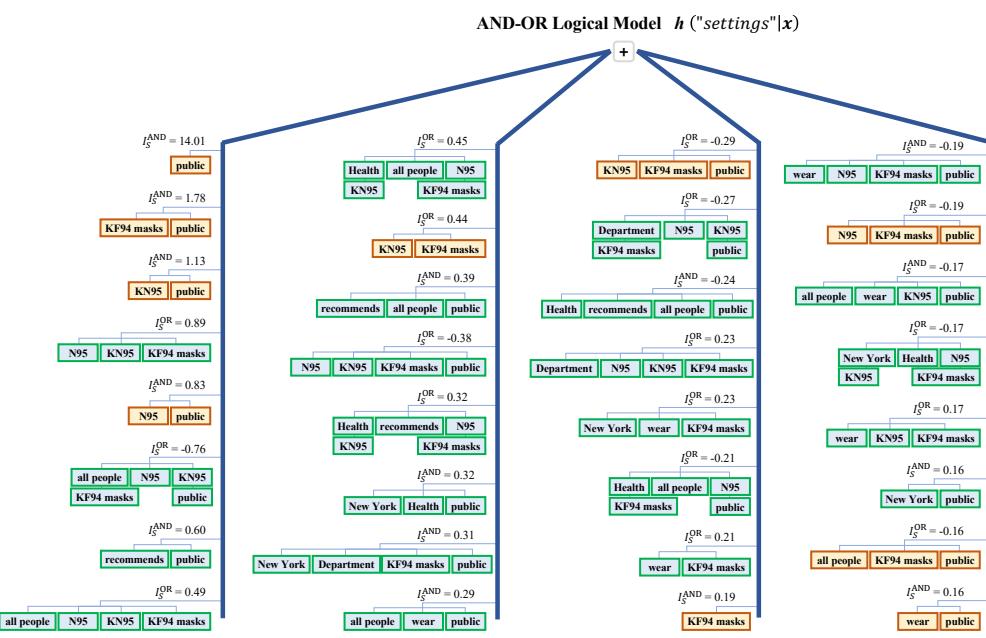
864 F AND-OR GRAPH  
865  
866  
867  
868  
869

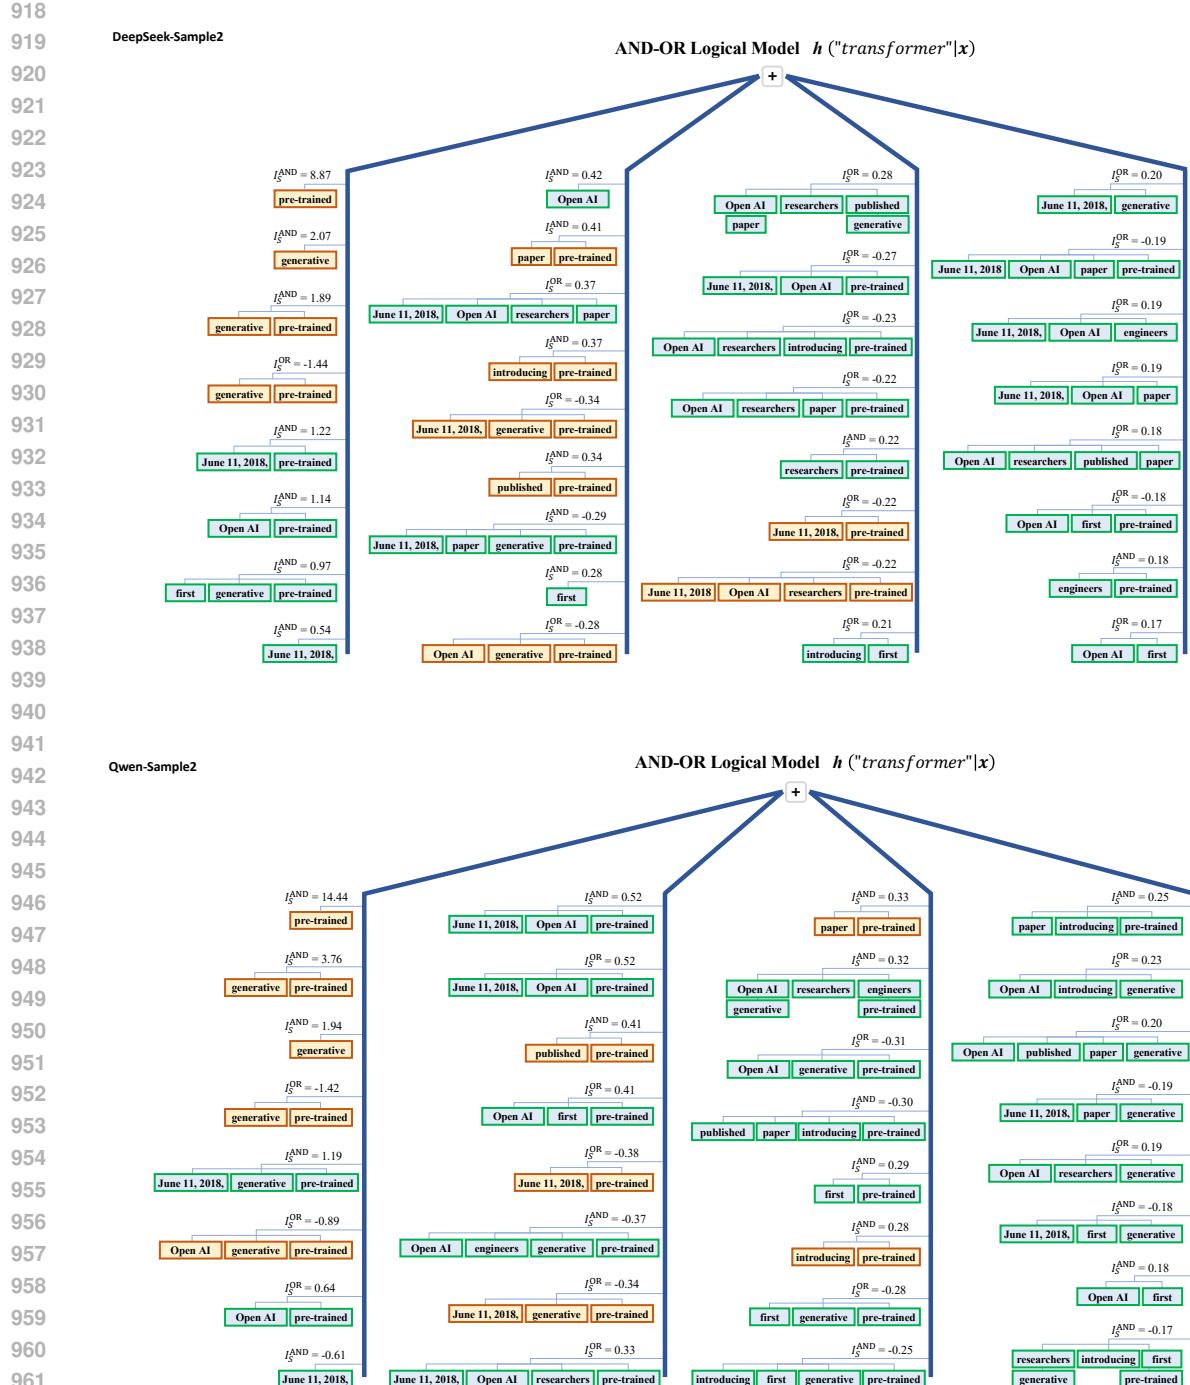
870 In this section we provide additional AND-OR logical models to explain the outputs of DeepSeek-R1-  
871 Distill-Llama-8B and Qwen2.5-7B across different prompts . The results detailed in Figure 7 show  
872 that the AND-OR logical model can explain LLMs’ outputs on any randomly occluded samples.  
873

874 DeepSeek-Sample1



896 Qwen-Sample1





964     Figure 7: The AND-OR logical model successfully explains model outputs on arbitrarily occluded  
965     samples. The used distinct prompt-target pairs are : (1) Sample 1: “New York Department of Health  
966     recommends that all people should wear N95, KN95, or KF94 masks in all public” → “settings”.  
967     Input variables: “New York”, “Health”, “recommends”, “all people”, “wear”, “N95”, “KN95”, “KF94  
968     masks”, “public”, “Health”; (2) Sample 2: “On June 11, 2018, OpenAI researchers and engineers  
969     published a paper introducing the first generative pre-trained” → “transformer”. Input variables:  
970     “June 11, 2018”, “OpenAI”, “researchers”, “engineers”, “published”, “paper”, “introducing”, “first”,  
971     “generative”, “pre-trained”.

972  
973  
974  
975  
976  
977  
978



Figure 8: Examples of input variable selection on image data.

## G EXAMPLES OF INPUT VARIABLES SELECTION

986 We provide examples of input variable selection on image data in Figure 8, where the input variables  
987 are highlighted with green boxes. And we provide examples of input variable selection on natural  
988 language data in Figure 9.

**991** **Input sample:** Two parent butterflies with normal wings have offspring with altered wing shapes. What most likely caused this change?  
**992**

993 **Input variables:** "Two", "parent butterflies", "normal", "wings", "offspring"  
994 "altered", "wing", "What most likely", "caused", "change"

<img alt="A hierarchical tree diagram showing the evolution of butterfly traits. The root node is a blue box labeled '+'. The tree branches into two main paths. The left path leads to a blue box labeled 'What most likely' at the bottom. The right path leads to a blue box labeled 'What most likely' at the bottom. The tree structure is as follows: 
 - Root: +
 - Level 1: 
 - Left: What most likely
 - Right: What most likely
 - Level 2: 
 - Left-Left: wings
 - Left-Middle: offspring
 - Left-Right: wing
 - Right-Left: wings
 - Right-Middle: offspring
 - Right-Right: What most likely
 - Level 3: 
 - Left-Left-Left: normal
 - Left-Left-Middle: wings
 - Left-Left-Right: What most likely
 - Left-Middle-Left: What most likely
 - Left-Middle-Middle: wings
 - Left-Middle-Right: What most likely
 - Left-Right-Left: What most likely
 - Left-Right-Middle: wings
 - Left-Right-Right: What most likely
 - Right-Left-Left: wings
 - Right-Left-Middle: offspring
 - Right-Left-Right: wing
 - Right-Middle-Left: What most likely
 - Right-Middle-Middle: wings
 - Right-Middle-Right: What most likely
 - Right-Right-Left: What most likely
 - Right-Right-Middle: wings
 - Right-Right-Right: What most likely
 - Level 4: 
 - Left-Left-Left-Left: normal
 - Left-Left-Left-Middle: wings
 - Left-Left-Left-Right: What most likely
 - Left-Left-Middle-Left: What most likely
 - Left-Left-Middle-Middle: wings
 - Left-Left-Middle-Right: What most likely
 - Left-Left-Right-Left: What most likely
 - Left-Left-Right-Middle: wings
 - Left-Left-Right-Right: What most likely
 - Left-Middle-Left-Left: What most likely
 - Left-Middle-Left-Middle: wings
 - Left-Middle-Left-Right: What most likely
 - Left-Middle-Middle-Left: What most likely
 - Left-Middle-Middle-Middle: wings
 - Left-Middle-Middle-Right: What most likely
 - Left-Middle-Right-Left: What most likely
 - Left-Middle-Right-Middle: wings
 - Left-Middle-Right-Right: What most likely
 - Left-Right-Left-Left: What most likely
 - Left-Right-Left-Middle: wings
 - Left-Right-Left-Right: What most likely
 - Left-Right-Middle-Left: What most likely
 - Left-Right-Middle-Middle: wings
 - Left-Right-Middle-Right: What most likely
 - Left-Right-Right-Left: What most likely
 - Left-Right-Right-Middle: wings
 - Left-Right-Right-Right: What most likely
 - Right-Left-Left-Left: What most likely
 - Right-Left-Left-Middle: wings
 - Right-Left-Left-Right: What most likely
 - Right-Left-Middle-Left: What most likely
 - Right-Left-Middle-Middle: wings
 - Right-Left-Middle-Right: What most likely
 - Right-Left-Right-Left: What most likely
 - Right-Left-Right-Middle: wings
 - Right-Left-Right-Right: What most likely
 - Right-Middle-Left-Left: What most likely
 - Right-Middle-Left-Middle: wings
 - Right-Middle-Left-Right: What most likely
 - Right-Middle-Middle-Left: What most likely
 - Right-Middle-Middle-Middle: wings
 - Right-Middle-Middle-Right: What most likely
 - Right-Middle-Right-Left: What most likely
 - Right-Middle-Right-Middle: wings
 - Right-Middle-Right-Right: What most likely
 - Right-Right-Left-Left: What most likely
 - Right-Right-Left-Middle: wings
 - Right-Right-Left-Right: What most likely
 - Right-Right-Middle-Left: What most likely
 - Right-Right-Middle-Middle: wings
 - Right-Right-Middle-Right: What most likely
 - Right-Right-Right-Left: What most likely
 - Right-Right-Right-Middle: wings
 - Right-Right-Right-Right: What most likely
 - Level 5: 
 - Left-Left-Left-Left-Left: normal
 - Left-Left-Left-Left-Middle: wings
 - Left-Left-Left-Left-Right: What most likely
 - Left-Left-Left-Middle-Left: What most likely
 - Left-Left-Left-Middle-Middle: wings
 - Left-Left-Left-Middle-Right: What most likely
 - Left-Left-Left-Right-Left: What most likely
 - Left-Left-Left-Right-Middle: wings
 - Left-Left-Left-Right-Right: What most likely
 - Left-Left-Middle-Left-Left: What most likely
 - Left-Left-Middle-Left-Middle: wings
 - Left-Left-Middle-Left-Right: What most likely
 - Left-Left-Middle-Middle-Left: What most likely
 - Left-Left-Middle-Middle-Middle: wings
 - Left-Left-Middle-Middle-Right: What most likely
 - Left-Left-Middle-Right-Left: What most likely
 - Left-Left-Middle-Right-Middle: wings
 - Left-Left-Middle-Right-Right: What most likely
 - Left-Left-Right-Left-Left: What most likely
 - Left-Left-Right-Left-Middle: wings
 - Left-Left-Right-Left-Right: What most likely
 - Left-Left-Right-Middle-Left: What most likely
 - Left-Left-Right-Middle-Middle: wings
 - Left-Left-Right-Middle-Right: What most likely
 - Left-Left-Right-Right-Left: What most likely
 - Left-Left-Right-Right-Middle: wings
 - Left-Left-Right-Right-Right: What most likely
 - Left-Middle-Left-Left-Left: What most likely
 - Left-Middle-Left-Left-Middle: wings
 - Left-Middle-Left-Left-Right: What most likely
 - Left-Middle-Left-Middle-Left: What most likely
 - Left-Middle-Left-Middle-Middle: wings
 - Left-Middle-Left-Middle-Right: What most likely
 - Left-Middle-Left-Right-Left: What most likely
 - Left-Middle-Left-Right-Middle: wings
 - Left-Middle-Left-Right-Right: What most likely
 - Left-Middle-Middle-Left-Left: What most likely
 - Left-Middle-Middle-Left-Middle: wings
 - Left-Middle-Middle-Left-Right: What most likely
 - Left-Middle-Middle-Middle-Left: What most likely
 - Left-Middle-Middle-Middle-Middle: wings
 - Left-Middle-Middle-Middle-Right: What most likely
 - Left-Middle-Middle-Right-Left: What most likely
 - Left-Middle-Middle-Right-Middle: wings
 - Left-Middle-Middle-Right-Right: What most likely
 - Left-Middle-Right-Left-Left: What most likely
 - Left-Middle-Right-Left-Middle: wings
 - Left-Middle-Right-Left-Right: What most likely
 - Left-Middle-Right-Middle-Left: What most likely
 - Left-Middle-Right-Middle-Middle: wings
 - Left-Middle-Right-Middle-Right: What most likely
 - Left-Middle-Right-Right-Left: What most likely
 - Left-Middle-Right-Right-Middle: wings
 - Left-Middle-Right-Right-Right: What most likely
 - Left-Right-Left-Left-Left: What most likely
 - Left-Right-Left-Left-Middle: wings
 - Left-Right-Left-Left-Right: What most likely
 - Left-Right-Left-Middle-Left: What most likely
 - Left-Right-Left-Middle-Middle: wings
 - Left-Right-Left-Middle-Right: What most likely
 - Left-Right-Left-Right-Left: What most likely
 - Left-Right-Left-Right-Middle: wings
 - Left-Right-Left-Right-Right: What most likely
 - Left-Right-Middle-Left-Left: What most likely
 - Left-Right-Middle-Left-Middle: wings
 - Left-Right-Middle-Left-Right: What most likely
 - Left-Right-Middle-Middle-Left: What most likely
 - Left-Right-Middle-Middle-Middle: wings
 - Left-Right-Middle-Middle-Right: What most likely
 - Left-Right-Middle-Right-Left: What most likely
 - Left-Right-Middle-Right-Middle: wings
 - Left-Right-Middle-Right-Right: What most likely
 - Left-Right-Right-Left-Left: What most likely
 - Left-Right-Right-Left-Middle: wings
 - Left-Right-Right-Left-Right: What most likely
 - Left-Right-Right-Middle-Left: What most likely
 - Left-Right-Right-Middle-Middle: wings
 - Left-Right-Right-Middle-Right: What most likely
 - Left-Right-Right-Right-Left: What most likely
 - Left-Right-Right-Right-Middle: wings
 - Left-Right-Right-Right-Right: What most likely
 - Right-Left-Left-Left-Left: What most likely
 - Right-Left-Left-Left-Middle: wings
 - Right-Left-Left-Left-Right: What most likely
 - Right-Left-Left-Middle-Left: What most likely
 - Right-Left-Left-Middle-Middle: wings
 - Right-Left-Left-Middle-Right: What most likely
 - Right-Left-Left-Right-Left: What most likely
 - Right-Left-Left-Right-Middle: wings
 - Right-Left-Left-Right-Right: What most likely
 - Right-Left-Middle-Left-Left: What most likely
 - Right-Left-Middle-Left-Middle: wings
 - Right-Left-Middle-Left-Right: What most likely
 - Right-Left-Middle-Middle-Left: What most likely
 - Right-Left-Middle-Middle-Middle: wings
 - Right-Left-Middle-Middle-Right: What most likely
 - Right-Left-Middle-Right-Left: What most likely
 - Right-Left-Middle-Right-Middle: wings
 - Right-Left-Middle-Right-Right: What most likely
 - Right-Left-Right-Left-Left: What most likely
 - Right-Left-Right-Left-Middle: wings
 - Right-Left-Right-Left-Right: What most likely
 - Right-Left-Right-Middle-Left: What most likely
 - Right-Left-Right-Middle-Middle: wings
 - Right-Left-Right-Middle-Right: What most likely
 - Right-Left-Right-Right-Left: What most likely
 - Right-Left-Right-Right-Middle: wings
 - Right-Left-Right-Right-Right: What most likely
 - Right-Middle-Left-Left-Left: What most likely
 - Right-Middle-Left-Left-Middle: wings
 - Right-Middle-Left-Left-Right: What most likely
 - Right-Middle-Left-Middle-Left: What most likely
 - Right-Middle-Left-Middle-Middle: wings
 - Right-Middle-Left-Middle-Right: What most likely
 - Right-Middle-Left-Right-Left: What most likely
 - Right-Middle-Left-Right-Middle: wings
 - Right-Middle-Left-Right-Right: What most likely
 - Right-Middle-Middle-Left-Left: What most likely
 - Right-Middle-Middle-Left-Middle: wings
 - Right-Middle-Middle-Left-Right: What most likely
 - Right-Middle-Middle-Middle-Left: What most likely
 - Right-Middle-Middle-Middle-Middle: wings
 - Right-Middle-Middle-Middle-Right: What most likely
 - Right-Middle-Middle-Right-Left: What most likely
 - Right-Middle-Middle-Right-Middle: wings
 - Right-Middle-Middle-Right-Right: What most likely
 - Right-Middle-Right-Left-Left: What most likely
 - Right-Middle-Right-Left-Middle: wings
 - Right-Middle-Right-Left-Right: What most likely
 - Right-Middle-Right-Middle-Left: What most likely
 - Right-Middle-Right-Middle-Middle: wings
 - Right-Middle-Right-Middle-Right: What most likely
 - Right-Middle-Right-Right-Left: What most likely
 - Right-Middle-Right-Right-Middle: wings
 - Right-Middle-Right-Right-Right: What most likely
 - Right-Right-Left-Left-Left: What most likely
 - Right-Right-Left-Left-Middle: wings
 - Right-Right-Left-Left-Right: What most likely
 - Right-Right-Left-Middle-Left: What most likely
 - Right-Right-Left-Middle-Middle: wings
 - Right-Right-Left-Middle-Right: What most likely
 - Right-Right-Left-Right-Left: What most likely
 - Right-Right-Left-Right-Middle: wings
 - Right-Right-Left-Right-Right: What most likely
 - Right-Right-Middle-Left-Left: What most likely
 - Right-Right-Middle-Left-Middle: wings
 - Right-Right-Middle-Left-Right: What most likely
 - Right-Right-Middle-Middle-Left: What most likely
 - Right-Right-Middle-Middle-Middle: wings
 - Right-Right-Middle-Middle-Right: What most likely
 - Right-Right-Middle-Right-Left: What most likely
 - Right-Right-Middle-Right-Middle: wings
 - Right-Right-Middle-Right-Right: What most likely
 - Right-Right-Right-Left-Left: What most likely
 - Right-Right-Right-Left-Middle: wings
 - Right-Right-Right-Left-Right: What most likely
 - Right-Right-Right-Middle-Left: What most likely
 - Right-Right-Right-Middle-Middle: wings
 - Right-Right-Right-Middle-Right: What most likely
 - Right-Right-Right-Right-Left: What most likely
 - Right-Right-Right-Right-Middle: wings
 - Right-Right-Right-Right-Right: What most likely
 </p>

**Input sample:** One type of animal hatches from an egg, breathes through gills when it is young, and mainly lives on land as an adult. Into which group is this animal classified?

**Input variables:** "animal hatches from", "egg", "breathes through", "gills", "young", "lives", "land", "adult", "Into which group", "classified "

Figure 9: Examples of input variable selection on nlp data.

## 1026 H PROOF OF EQUATION (11)

1028 We want to derive Equation (11):

$$\begin{aligned}
 1030 \quad v(\mathbf{x}) &= \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_{\text{type}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers type relation} \\ \text{between input variables in } S \end{array} \right) \\
 1031 \\
 1032 \quad &+ \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_{\text{type}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers type relation} \\ \text{between input variables in } S \end{array} \right) \\
 1033 \\
 1034 \quad &+ \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_{\text{type}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers type relation} \\ \text{between input variables in } S \end{array} \right) + b
 \end{aligned}$$

1035 Let  $v(\mathbf{x}')$  be the output of the DNN for a masked input  $\mathbf{x}'$ . From Theorem 1 (universal matching property), we know that the logical model  $d(\mathbf{x}')$  precisely captures the DNN's output, so  $v(\mathbf{x}') = d(\mathbf{x}')$ . Equation (3) states:

$$1042 \quad d(\mathbf{x}') = \sum_{T \in \Omega_{\text{and}}} I_T^{\text{and}} \cdot \delta_{\text{and}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers AND relation} \\ \text{between input variables in } T \end{array} \right) + \sum_{T \in \Omega_{\text{or}}} I_T^{\text{or}} \cdot \delta_{\text{or}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers OR relation} \\ \text{between input variables in } T \end{array} \right) + b. \quad (16)$$

1045 For consistency with Equation (11)'s summation index  $S$ , we rewrite Equation (16) using  $S$  as the index:

$$1047 \quad v(\mathbf{x}') = \sum_{S \in \Omega_{\text{and}}} I_S^{\text{and}} \cdot \delta_S^{\text{and}}(\mathbf{x}') + \sum_{S \in \Omega_{\text{or}}} I_S^{\text{or}} \cdot \delta_S^{\text{or}}(\mathbf{x}') + b, \quad (17)$$

1049 where  $\delta_S^{\text{type}}(\mathbf{x}')$  is a shorthand for  $\delta_{\text{type}} \left( \begin{array}{l} \mathbf{x}' \text{ triggers type relation} \\ \text{between input variables in } S \end{array} \right)$ . The term  $v(\mathbf{x})$  on the LHS of 1050 Equation (11) can be interpreted as  $v(\mathbf{x}')$  where  $\mathbf{x}'$  is the specific input (potentially unmasked  $\mathbf{x}$ ) for 1051 which the score is being decomposed. For generality, we proceed with  $v(\mathbf{x}')$ . 1052

1053 Equation (17) can be written more compactly as:

$$1055 \quad v(\mathbf{x}') = \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + b. \quad (18)$$

1057 The set of all interactions for a given type,  $\Omega_{\text{type}}$ , can be partitioned into salient interactions  $\Omega_{\text{type}}^{\text{Salient}}$  1058 and non-salient interactions  $\Omega_{\text{type}}^{\text{Non-Salient}}$ . These are defined as:

- 1061 •  $\Omega_{\text{type}}^{\text{Salient}} = \{S \mid |I_S^{\text{type}}| > \tau\}$
- 1062 •  $\Omega_{\text{type}}^{\text{Non-Salient}} = \{S \mid |I_S^{\text{type}}| \leq \tau\}$

1064 Thus,  $\Omega_{\text{type}} = \Omega_{\text{type}}^{\text{Salient}} \cup \Omega_{\text{type}}^{\text{Non-Salient}}$ , and these two sets are disjoint.

1065 Substituting this partition into the sum:

$$\begin{aligned}
 1068 \quad v(\mathbf{x}') &= \sum_{\text{type} \in \{\text{and, or}\}} \left( \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') \right) + b \\
 1069 \\
 1070 \quad &= \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') \\
 1071 \\
 1072 \quad &+ \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + b.
 \end{aligned}$$

1076 The second term here matches the "negligible non-salient interactions" part of Equation (11):  
 1077  $\sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}')$ .

1078 Now, let's focus on the first term, which sums over salient interactions:  $T_{\text{salient}} =$   
 1079  $\sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}')$ . We introduce the generalizability indicator  $G_S^{\text{type}}$  (defined

1080 as Equation (5) to indicate whether a interaction is generalizable), where  $G_S^{\text{type}} \in \{0, 1\}$ . For any term  
 1081  $A_S^{\text{type}} = I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}')$  in the sum, we can use the identity  $A_S^{\text{type}} = A_S^{\text{type}} \cdot G_S^{\text{type}} + A_S^{\text{type}} \cdot (1 - G_S^{\text{type}})$ .  
 1082 So,  $I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') = I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_S^{\text{type}}(\mathbf{x}')$ . Substituting this back  
 1083 into the sum  $T_{\text{salient}}$ :

$$\begin{aligned} 1085 \quad T_{\text{salient}} &= \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} (I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_S^{\text{type}}(\mathbf{x}')) \\ 1086 \\ 1087 \quad &= \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') \\ 1088 \\ 1089 \quad &+ \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_S^{\text{type}}(\mathbf{x}'). \\ 1090 \\ 1091 \\ 1092 \end{aligned}$$

1093 The first part of this expression corresponds to the “generalizable interaction” and the second part  
 1094 corresponds to the “non-generalizable interactions” in Equation (11).

1095 Combining all components, we obtain:

$$\begin{aligned} 1096 \quad v(\mathbf{x}') &= \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot G_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') \\ 1097 \\ 1098 \quad &+ \sum_{\substack{\text{type} \in \\ \{\text{and, or}\}}} \sum_{S \in \Omega_{\text{type}}^{\text{Salient}}} I_S^{\text{type}} \cdot (1 - G_S^{\text{type}}) \cdot \delta_S^{\text{type}}(\mathbf{x}') \\ 1099 \\ 1100 \quad &+ \sum_{\text{type} \in \{\text{and, or}\}} \sum_{S \in \Omega_{\text{type}}^{\text{Non-Salient}}} I_S^{\text{type}} \cdot \delta_S^{\text{type}}(\mathbf{x}') + b. \\ 1101 \\ 1102 \\ 1103 \\ 1104 \end{aligned}$$

1105 This matches equation 11. If Equation (11) is intended for the unmasked input  $\mathbf{x}$ , we simply set  
 1106  $\mathbf{x}' = \mathbf{x}$  throughout this derivation, and the LHS  $v(\mathbf{x})$  becomes consistent.

## 1108 I DEFINE THE THREE-PHASE MORE QUANTITATIVELY

1110 Adopting a more rigorous quantitative approach helps to better delineate the boundaries between the  
 1111 three phases. Therefore, we have formulated more precise, quantitative definitions for these phases  
 1112 below.

1113 We characterize the learning dynamics using several key metrics. Let  $\delta_g(t)$  and  $\delta_n(t)$  be the incre-  
 1114 mental strengths of new **generalizable** and **non-generalizable** interactions learned at time  $t$ . We  
 1115 also track the average interaction generalization power,  $\bar{\mathcal{H}}(t)$ , and the number of learned interactions,  
 1116  $\bar{\mathcal{N}}(t)$ .

1117 We can now characterize the three phases:

- 1119 • **Phase 1 (Interval:  $[t_0, t_1]$ )**

1120 The model primarily learns generalizable patterns.  $\bar{\mathcal{H}}(t)$  increases while  $\bar{\mathcal{N}}(t)$  decreases.  
 1121 This phase concludes when  $\bar{\mathcal{H}}(t)$  reaches its peak:

$$1123 \quad t_1 = \arg \max_t \bar{\mathcal{H}}(t).$$

- 1125 • **Phase 2 (Interval:  $(t_1, t_2]$ )**

1126 The model begins to learn non-generalizable interactions along with generalizable interac-  
 1127 tions.  $\bar{\mathcal{H}}(t)$  decreases and  $\bar{\mathcal{N}}(t)$  increases. Learning of generalizable patterns still dominates,  
 1128 satisfying the condition

$$1129 \quad \delta_g(t) \geq r \cdot \delta_n(t).$$

1130 This phase ends at time  $t_2$ , the first point where the strengths balance:

$$1131 \quad \delta_g(t_2) = r \cdot \delta_n(t_2),$$

1132 given the monotonic increase of the ratio  $\frac{\Delta_n(t)}{\Delta_g(t)}$ .

1134 • **Phase 3 (Interval:  $(t_2, t_3]$ )**  
 1135 The acquisition of new interactions is now dominated by non-generalizable ones, with their  
 1136 strengths satisfying  
 1137 
$$\delta_g(t) < r \cdot \delta_n(t).$$
  
 1138 Moreover,  $\bar{\mathcal{N}}(t)$  continues to increase and  $\bar{\mathcal{H}}(t)$  continues to decrease.  
 1139

1140 **J EXPERIMENT OF REMOVING THE NON-GENERALIZABLE INTERACTIONS**  
 1141 **FROM THE MODEL OUTPUT**  
 1142

1143 **J.1 SAMPLE SELECTION FOR LOSS COMPUTATION**

1144 The loss compare in Section 3.2, encompassing both standard and modified loss calculations after  
 1145 removing non-generalizable interactions, were derived from an average over 100 samples per dataset.  
 1146 For the Tiny-ImageNet dataset, these 100 samples were obtained by first randomly selecting 10  
 1147 classes, and then randomly drawing 10 image instances from each chosen class. This sample selection  
 1148 process was applied independently to both the training and testing sets to gather the respective  
 1149 samples for loss calculation. For SST-2 and AG News datasets, 100 samples were randomly drawn  
 1150 from within all available classes, aiming for a proportional representation or a balanced distribution  
 1151 across classes (e.g., attempting to sample  $m = 100/C$  instances from each of the  $C$  classes, with  
 1152 necessary adjustments made for class imbalances or varying class sizes).  
 1153

1154 **J.2 FUTURE WORK AND POTENTIAL DIRECTIONS**  
 1155

1156 The demonstration that non-generalizable interactions are a key driver of overfitting, as discussed in  
 1157 Section 3.2, suggests several compelling research avenues. Future efforts could focus on developing  
 1158 novel training paradigms—encompassing targeted regularization, interaction-aware network pruning,  
 1159 or refined early stopping criteria—all designed to proactively prevent or penalize the formation  
 1160 of these detrimental interactions. Another promising direction involves moving beyond simple  
 1161 removal, exploring methods to dynamically modulate or down-weight the influence of identified non-  
 1162 generalizable interactions during training or inference, potentially guided by a deeper understanding  
 1163 of their emergence.  
 1164

1165 **K EXPERIMENT OF PENALIZING NON-GENERALIZABLE INTERACTIONS IN**  
 1166 **DNN**  
 1167

1168 To bridge the theoretical discussion in our paper with practical implementation, we conducted a  
 1169 follow-up experiment. The core idea was to leverage a dedicated validation set to identify and  
 1170 subsequently penalize non-generalizable interactions and classification loss during the end-to-end  
 1171 training process.  
 1172

1173 Our approach consisted of two key steps. First, to **quantify the generalization power of interactions**,  
 1174 we trained a baseline Deep Neural Network (DNN) on a separate validation set. This model allowed  
 1175 us to quantify which of the learned feature interactions were non-generalizable. Second, to **design**  
 1176 **the penalty term**, we designed a specialized loss function based on the quantification from the  
 1177 previous step. During the main model’s training, this function directly penalized the absolute values  
 1178 of the interactions identified as non-generalizable, denoted as  $|I_S^{\text{NoneG}}|$ , along with minimization of  
 1179 the classification loss.  
 1180

1181 We evaluated this approach on the CIFAR-10 dataset using a VGG-11 architecture, and the results  
 1182 were encouraging. The method led to a clear improvement in the quality of learned interactions:  
 1183 The proportion of generalizable interactions, as measured by  $\bar{\mathcal{H}}$  defined in Section 3.1, increased  
 1184 significantly from 28% to 39%. Correspondingly, the proportion of non-generalizable interactions  
 1185 decreased from 72% to 61%.

1186 **Regarding the impact on model accuracy**, the proposed method consistently improves performance  
 1187 across different architectures and datasets. For the VGG-16 model on the Tiny-ImageNet dataset, test  
 1188 accuracy increased by **2.5% (46.3% to 48.8%)**. Similarly, for the BERT-Tiny model on the SST-2  
 1189 sentiment classification task, accuracy improved by **4.4% (72.2% to 76.6%)**. The algorithm also

1188 demonstrates significant gains in few-shot learning scenarios. For instance, in an object recognition  
 1189 task with only 200 training samples, our method achieved a **6.0% (76.6% to 82.6%)** improvement in  
 1190 classification accuracy using a ResNet-18 model.

1191 In summary, this experiment demonstrated that our proposed method effectively reduced the learning  
 1192 of non-generalizable interactions and improved the neural network’s performance, showing particular  
 1193 promise in situations with limited training data.

## 1195 L COMPUTATIONAL RESOURCES AND PROCEDURAL DURATIONS

1198 All experiments were conducted on a dedicated server equipped with four NVIDIA GeForce RTX  
 1199 3080 Ti GPUs and powered by Intel(R) Xeon(R) Gold 6146 CPU.

1200 The model training durations varied depending on the specific experiment and model complexity:

- 1202 • **Ablation study:** The training for ablation studies, involving up to five baseline Deep Neural  
 1203 Networks (DNNs), represented the most time-consuming part. Due to the necessity of  
 1204 training these models sequentially or on separate GPUs to ensure resource isolation, the  
 1205 cumulative training time for these experiments was approximately 50 hours.
- 1206 • **Small to medium models:** For other experiments involving small to medium-sized models  
 1207 such as ResNet-50, individual training runs were typically completed within 12 hours per  
 1208 model.
- 1209 • **Larger models:** Training larger models, including ResNet-101 and BERT-large, for the  
 1210 remaining experimental setups generally concluded within 30 hours per model.

1211 Regarding specific experimental procedures:

- 1213 • **Interaction extraction:** The process of extracting each individual interaction took approxi-  
 1214 mately 60 seconds.
- 1215 • **Generalizable interaction assessment:** The evaluation of interaction generalization was  
 1216 comparatively swift. Assessing the generalizable power for a set of 100 samples typically  
 1217 completed within a range of 8 to 14 seconds.

1219 These resources and timeframes allowed for comprehensive experimentation and validation of the  
 1220 proposed methods.

## 1222 M REDEFINITION OF THE GENERALIZATION POWER METRIC

1224 In this section, we provide the redefinition of the generalizability metric for interactions in 5. Instead  
 1225 of the binary metrics, we use a continuous metric within the range  $[0, 1]$  to quantify the generalization  
 1226 power of interactions between the target model  $v$  and the baseline model  $v^{base}$ .

1228 For each salient interaction  $S$  (an AND interaction or an OR interaction), the generalization power is  
 1229 redefined as the similarity between the interaction effect in the target model and that in the baseline  
 1230 model as follows:

$$1232 \mathcal{G}_{S, v^{base}}^{\text{and}} = \mathbf{1} \left( I_S^{\text{and}} \cdot I_{S, v^{base}}^{\text{and}} > 0 \right) \cdot \frac{\min \left( |I_S^{\text{and}}|, |I_{S, v^{base}}^{\text{and}}| \right)}{\max \left( |I_S^{\text{and}}|, |I_{S, v^{base}}^{\text{and}}| \right)}, \quad (19)$$

$$1236 \mathcal{G}_{S, v^{base}}^{\text{or}} = \mathbf{1} \left( I_S^{\text{or}} \cdot I_{S, v^{base}}^{\text{or}} > 0 \right) \cdot \frac{\min \left( |I_S^{\text{or}}|, |I_{S, v^{base}}^{\text{or}}| \right)}{\max \left( |I_S^{\text{or}}|, |I_{S, v^{base}}^{\text{or}}| \right)}. \quad (20)$$

1240 Here,  $\mathbf{1}(\cdot)$  is an indicator function, which returns 1 if the condition holds, and returns 0 otherwise.

1241 The ratio term  $\frac{\min(|I_S|, |I_{S, v^{base}}|)}{\max(|I_S|, |I_{S, v^{base}}|)} \in [0, 1]$  measures the similarity of the interaction effects encoded by

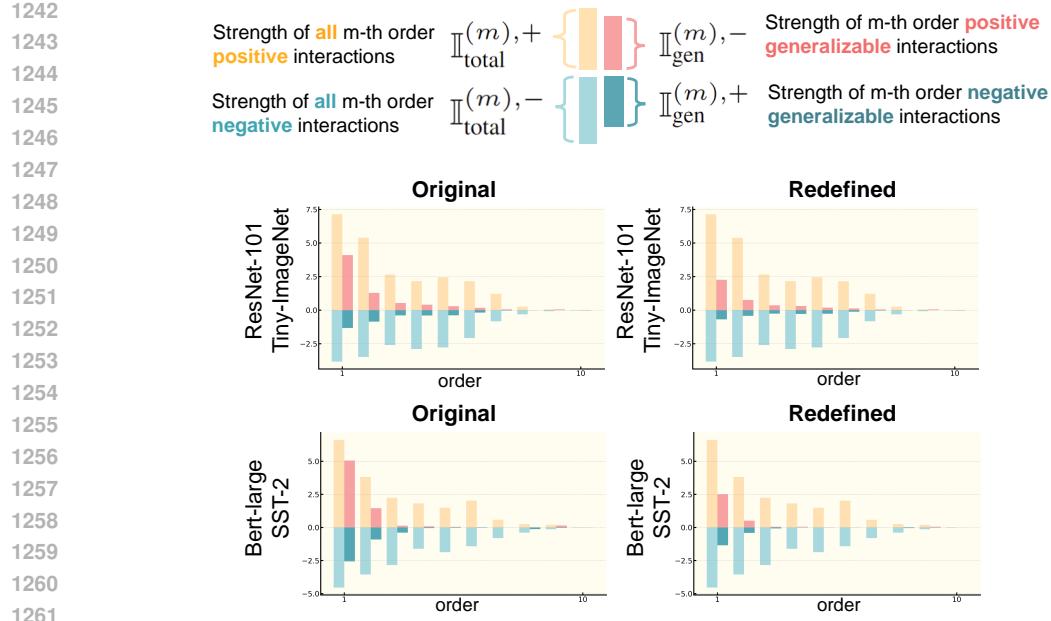


Figure 10: Analysis of generalization power across interaction orders. (Left) The distribution of interactions calculated with the original metric ( $\mathcal{G}$ ) titled with “**Original**”. (Right) The distribution of interactions calculated with the redefined metric ( $\mathcal{G}'$ ) titled with “**Redefined**”.

the two models, which reflects generalizability of this interaction. Only when the same interaction  $S$  (or OR interaction  $S$ ) is simultaneously extracted by both DNNs and exhibits exactly the same effect, we consider it to be 100.0% generalizable.

**Normalization.** Considering that the outputs of different DNNs often exhibit different ranges of perturbation, it is necessary to normalize the interaction effects to ensure fair comparison. We normalize  $I_S^{\text{and}}$  and  $I_S^{\text{or}}$  as:  $I_S^{\text{and}} \leftarrow \frac{I_S^{\text{and}}}{\mathbb{E}_{\mathbf{x}}[|v(N) - v(\emptyset)|]}$ , and  $I_S^{\text{or}} \leftarrow \frac{I_S^{\text{or}}}{\mathbb{E}_{\mathbf{x}}[|v(N) - v(\emptyset)|]}$ , where  $\mathbb{E}_{\mathbf{x}}[|v(N) - v(\emptyset)|]$  represents the expected output variation of the model.

Using this redefined metric  $\mathcal{G}_{S,v^{\text{base}}}^{\text{and}}$  and  $\mathcal{G}_{S,v^{\text{base}}}^{\text{or}}$ , we analyzed the distribution of generalizable interactions across different interaction orders on ResNet-101 trained with Tiny-ImageNet and Bert-large trained with SST-2. As shown in Figure 10, the new results aligned closely with results based on the old metric. Because the new metric is more stringent, the strength of generalizable interactions is reduced. However, the distribution of generalizable interactions over different orders based on the new metric is still similar to that based on the old metric.

## N EVALUATION OF THE QUANTIFICATION OF GENERALIZATION POWER

We conducted a new experiment to explicitly evaluate the quantification of generalization power. Specifically, we uniformly partitioned all data into three sets (A, B, and C) to perform cross-validation. We trained the target DNN on the dataset A, and trained two baseline DNNs on the dataset B and the dataset C, respectively. Here we regarded the baseline model trained on dataset B as representing the interactions on the validation set and regarded the baseline model trained on dataset C as representing the interactions on the test set. Then, given each input sample, we followed Equation (7) in Section 2.2 in the manuscript to measure the distribution (i.e.,  $\mathbb{I}_{\text{gen}_B}^{(m),+}$  and  $\mathbb{I}_{\text{gen}_B}^{(m),-}$ ) of interactions that could generalize to the validation set (i.e., the set B) and the distribution (i.e.,  $\mathbb{I}_{\text{gen}_C}^{(m),+}$  and  $\mathbb{I}_{\text{gen}_C}^{(m),-}$ ) of interactions that can generalize to the test set (i.e., the set C).<sup>11</sup>

<sup>11</sup>Specifically, extending the definition in Equation (7), we utilized  $\mathbb{I}_{\text{gen}_B}^{(m),+}$  to denote the aggregate effect strength of all generalizable salient interactions with positive effect extracted from the baseline DNN trained on

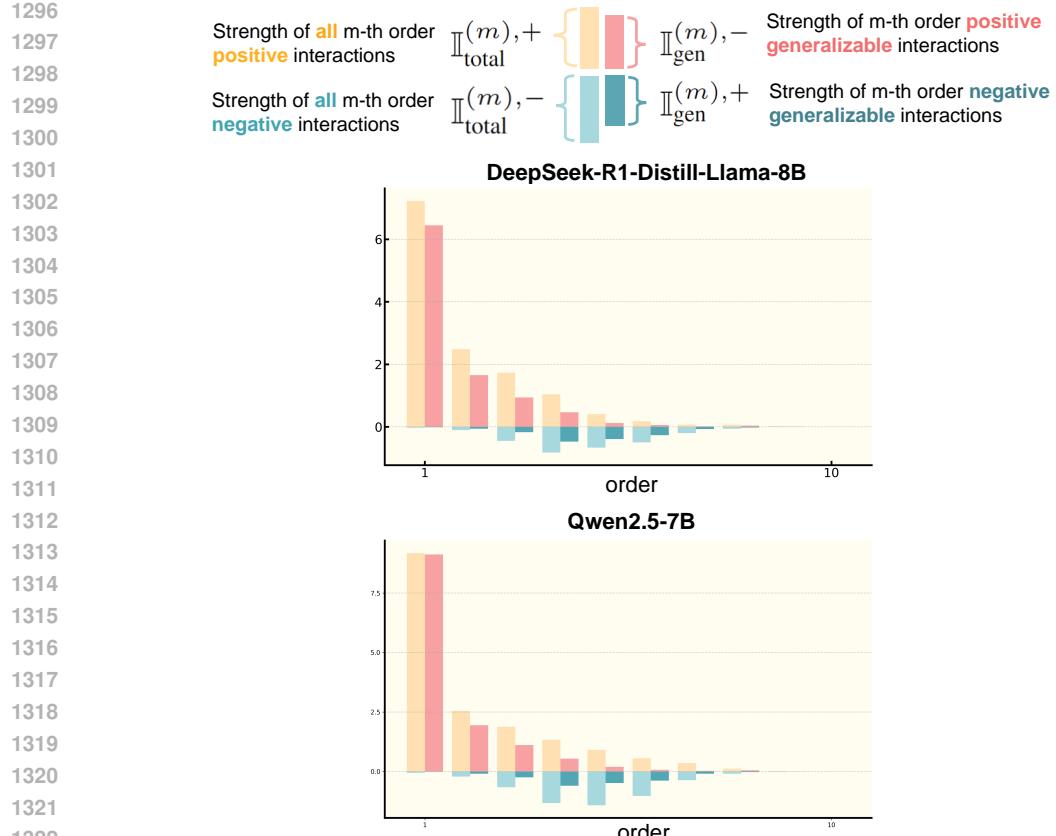


Figure 11: Distribution of interactions measured between Qwen2.5-7B (baseline) and DeepSeek-R1-Distill-Llama-8B (target) using SQuAD data following Equation 7 in Section 2.2. The plot illustrates that the two LLMs tend to encode low-order, simple interactions that are more transferable.

Table 1: Jaccard Similarity ( $J$ ) of generalizable interaction distributions between validation set (B) and test set (C).

Model and Dataset	Jaccard Similarity ( $J$ )
ResNet-50 on Tiny-ImageNet	71.4%
Bert-large on SST-2	74.7%

The validation was conducted to test whether the interactions, which could generalize to the validation set, could also generalize to the test set. If feasible, the faithfulness of our method was validated. Specifically, we computed the Jaccard similarity between the two distributions as follows:

$$J = \frac{\sum_{m=1}^n \left[ \min(\left| \mathbb{I}_{\text{gen}_B}^{(m),+} \right|, \left| \mathbb{I}_{\text{gen}_C}^{(m),+} \right|) + \min(\left| \mathbb{I}_{\text{gen}_B}^{(m),-} \right|, \left| \mathbb{I}_{\text{gen}_C}^{(m),-} \right|) \right]}{\sum_{m=1}^n \left[ \max(\left| \mathbb{I}_{\text{gen}_B}^{(m),+} \right|, \left| \mathbb{I}_{\text{gen}_C}^{(m),+} \right|) + \max(\left| \mathbb{I}_{\text{gen}_B}^{(m),-} \right|, \left| \mathbb{I}_{\text{gen}_C}^{(m),-} \right|) \right]}.$$

dataset B,  $\mathbb{I}_{\text{gen}_B}^{(m),-}$  to denote the aggregate effect strength of all generalizable salient interactions with negative effect extracted from the baseline DNN trained on dataset B,  $\mathbb{I}_{\text{gen}_C}^{(m),+}$  to denote the aggregate effect strength of all generalizable salient interactions with positive effect extracted from the baseline DNN trained on dataset C and  $\mathbb{I}_{\text{gen}_C}^{(m),-}$  to denote the aggregate effect strength of all generalizable salient interactions with negative effect extracted from the baseline DNN trained on dataset C.

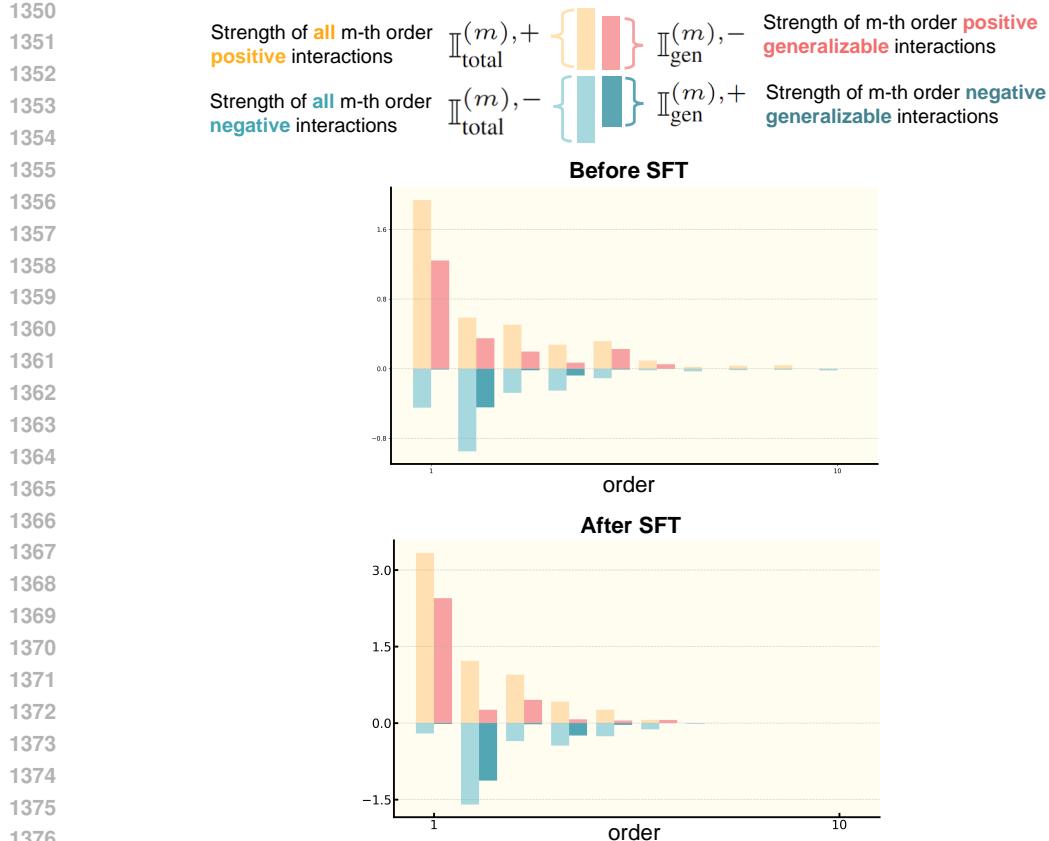


Figure 12: Distribution of interactions before and after the SFT process.

The experimental results in Table 1 showed that the DNNs exhibited high Jaccard similarity of generalizable interaction distribution between the test set and the validation set. This supported the faithfulness of our method of identifying generalizable interactions.

## O EXPERIMENT ON TRANSFORMER-BASED LLMs

Extensively, we have conducted an experiment of analyzing the distribution of generalizable interactions on several open source LLMs based on the transformer architecture, including Qwen2.5-7B and DeepSeek-R1-Distill-Llama-8B, which are well-known large language models and have been widely applied.

Notably, we acknowledge the significant challenge of retraining a foundational model from scratch to serve as a baseline. Therefore, considering that different LLMs released by various organizations/companies are typically trained on different and massive datasets, we provisionally utilized one LLM as the baseline DNN to evaluate another.

Specifically, we selected Qwen2.5-7B model as the baseline model and DeepSeek-R1-Distill-Llama-8B model as the target model. We used data from the SQuAD dataset and followed Equation (7) from Section 2.2 in the manuscript to measure the distribution of interactions that can generalize to the baseline DNN.

As illustrated in Figure 11, the new experimental results suggested that the two large models tended to encode low-order, simple generalizable interactions. This indicated that these well-trained LLMs might focus on learning simple patterns during training, and such simple interactions could potentially be more transferable between models. A systematic investigation of transferability across a larger set of LLMs and datasets is left to future work.

---

**P ANALYZE LLM FINE-TUNING WITH OUR METHOD**

For the SFT process, we used the Unilaw-R1-Data (Cai et al., 2025) to fine-tune the Qwen2.5-7B-Instruct model, which was originally trained on 18T tokens of data. Experimental results showed that during the SFT process, low-order interactions increased, while high-order interactions decreased. The proportion of generalizable interactions changed from 45.2% to 52.3%. The current experimental results had shown that the SFT process was still within the first phase. Figure 12 shows the distribution of interactions before and after the SFT process. We will continue to finetune the LLM and keep reporting latest distribution of interactions to provide dynamics in the second and third phases in future. The above experiments show that our method remains applicable to analyzing LLM training.

**Q THE USE OF LARGE LANGUAGE MODELS (LLMs)**

In the preparation of this article, Large Language Models (LLMs) were utilized solely for the purpose of language polishing and identifying grammatical inaccuracies. The core research, analysis, findings, and intellectual contributions remain entirely those of the authors. The LLMs acted as an assistive tool to improve the clarity and fluency of the academic prose, without influencing the substantive content or the logical arguments presented.

1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457