## Mitigating Biases for Instruction-following Language Models via Bias Neurons Elimination

Anonymous ACL submission

### Abstract

Instruction-following language models often 001 show undesirable biases. These undesirable biases are accelerated in the real-world usage 004 of language models, where a wide range of instructions is used through zero-shot example prompting. To solve this problem, we first de-007 fine *bias neuron*, which significantly affects biased outputs, and prove its existence empirically. Furthermore, we propose a novel and practical bias mitigation method, CRISPR, to eliminate bias neurons of language models in 011 instruction-following settings. CRISPR auto-013 matically determines biased outputs and categorizes neurons that affect the biased out-015 puts as bias neurons using an attribution, an explainability method. Experimental results 017 demonstrate the effectiveness of our method in mitigating biases under zero-shot instructionfollowing settings without losing the model's 019 task performance and existing knowledge. The experimental results reveal the generalizability of our method as it shows robustness under various instructions and datasets. Surprisingly, our method can mitigate the bias in language models by eliminating only a few neurons (e.g., three neurons).

### 1 Introduction

027

037

041

Instruction-following language models perform various tasks using instruction-based prompts (Chung et al., 2022; Touvron et al., 2023; Taori et al., 2023; OpenAI, 2023). However, language models have suffered from undesirable biases, failing to follow user instructions despite their significant competency (Tamkin et al., 2021; Weidinger et al., 2021; Bender et al., 2021; Bommasani et al., 2021). Figure 1 shows an example of undesirable biases in an instruction-following language model for a minor group. In this case, a language model assigns a high probability to a minor group label for a negative question despite the ambiguity of the given context. The undesirable biases of lan-



Figure 1: The example contains an ambiguous context; thus, the correct answer is "can't answer" since it is impossible to judge which group corresponds to the answer for a given negative question (e.g., "What group of people uses drugs?"). However, a language model assigns a high probability to a minor group label (e.g., "poor people"). Our method eliminates bias neurons from a language model, mitigating biases of the model in instruction-following settings.

guage models typically arise from the relationship between labels (e.g., "*poor people*") and tokens (e.g., "*drugs*") within data instances (Zhao et al., 2021; Fei et al., 2023).

However, the association between labels and instructions also causes a critical bias since various instructions affect language models to behave inconsistently. Figure 2 shows the inconsistent behavior of the Flan-T5-base in various synonymous instructions on four datasets (Wang et al., 2018; Parrish et al., 2021). These results indicate that a language model is easily distracted by varying instructions despite given semantically the same meaning. These phenomena suggest that language models exhibit significant cognitive biases in instruction settings, and these are some of the most



Figure 2: **Performance gaps in understanding instructions.** We plot the accuracies of Flan-T5-base about ten synonymous instructions for BBQ-SES, BBQ-Age, MRPC and RTE datasets. For example, the accuracy between instructions differed by up to 5% and 15% for BBQ-SES and MRPC, respectively. These results reveal that an instruction-following model shows biases in understanding instructions. The utilized instructions are described in detail in Appendix D.

critical biases to mitigate when using instructionfollowing language models.

To mitigate biases in language models, Zhao et al. (2021); Fei et al. (2023) have investigated label biases in few-shot in-context learning settings. Specifically, they have regarded the imbalanced probability distribution that occurred by inputting content-free texts (e.g., "N/A" or random tokens) as biases, and degraded the original output probability of each input instance by the output probability of the content-free texts. However, they have only aimed to mitigate biases in few-shot in-context learning settings, not considering the zero-shot instruction prompting. Since utilizing a language model in only the instruction prompting without few-shot examples is an efficient and practical usage scenario of language models, they have significant limitations in the scope of application.

To resolve these issues, in this paper, we first define the *bias neuron*, which significantly affects biased outputs in language models. Prior studies (Panigrahi et al., 2023; Wang et al., 2022; Panigrahi et al., 2023; Yang et al., 2023) have demonstrated that skills for a specific task are localized in particular neurons. From these findings, we hypothesize the presence of the bias neuron and empirically validate its existence. Furthermore, we propose a novel bias mitigation method called CRISPR, which stands for CalibRating InStruction Bias via Bias Neuron PRuning. We utilize an attribution (Yang et al., 2023), an explainability method for language models, to quantify the bias of each neuron at three aspects (i.g., token, instance, and instruction). Specifically, we first compute the bias influence of each token in inferring a predefined biased output, and aggregate the token score to an instance and instruction scores to derive the final bias influence (bias attribution) of each neuron. Furthermore, defining biased outputs manually for each data sample is time-consuming; thus, we propose an automatic identification method of biased outputs for computing the bias attribution using the confusion score of a language model.

085

091

097

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

We demonstrate our method in various social bias and natural language understanding benchmarks and dramatically outperform other baselines under varying zero-shot instruction settings. Surprisingly, we reveal that only a few bias neurons (e.g., three neurons) cause the bias, proving our method's practicality. In addition, we show that mitigating bias for a particular task does not adversely affect the existing knowledge of language models for solving other tasks. We also note that bias neurons identified for a specific dataset also function as biases in other analogous datasets, revealing that the bias knowledge is transferred to datasets from correlative domains. CRISPR is an efficient bias mitigation method since it needs only a few data samples (e.g., ten samples) to quantify the bias score for the whole neurons. CRISPR also enables language models to adapt flexibly by eliminating some existing bias neurons without any training process.

### 2 Related Works

### 2.1 Bias Mitigation

Despite demonstrating significant efficacy in various natural language understanding tasks, Language Models (LLMs) have been noted to exhibit undesirable biases (Ravfogel et al., 2019; Braverman et al., 2020; Liu et al., 2021; Lu et al., 2021; Bommasani et al., 2021; Work; Sorensen et al., 2022). Therefore, existing studies have attempted to solve the bias problems. For instance, Zhao et al. (2021); Fei et al. (2023) have mitigated biases found in few-shot in-context learning settings by utilizing outputs probability obtained from content-

061

free texts (e.g. "N/A" or random tokens). They 134 have argued that an imbalanced output probability 135 for the content-free texts corresponds to unintended 136 label biases. Specifically, they have shifted the 137 original output probability of each input instance 138 by dividing it through the output probability of 139 each class obtained from content-free texts. How-140 ever, existing studies have limitations in that they 141 have only aimed to mitigate biases in the few-shot 142 in-context learning setting. Since more realistic 143 LLM usage settings are based on the zero-shot 144 instruction-following mechanism, it is essential to 145 consider these settings to suggest a practical bias 146 mitigation method. Furthermore, the existing stud-147 ies have only tackled the simple classification prob-148 lem as a target to mitigate biases. To enhance the 149 scalability of bias mitigation methods, they should 150 demonstrate effectiveness in real-world natural lan-151 guage understanding tasks that involve handling 152 diverse and inconsistent label options. 153

### 2.2 Skill Neurons Detection

154

155

156

157

158

159

161

162

163

164

166

167

168

169

171

172

173

174

175

176

177

178

180

181

184

Despite the impressive performance of language models, it is challenging to precisely illuminate the role of each parameter in the model during the execution of a specific task. Existing studies have sought to detect important skill neurons in performing a specific task (Panigrahi et al., 2023; Wang et al., 2022; Yang et al., 2023).

Panigrahi et al. (2023) has suggested a trainingbased method, *model grafting*, which detects skill neurons by training a new parameter to mask original parameters. Although it has effectively detected skill neurons, it has required an additional training process for masking parameters as many as the number of model parameters.

Wang et al. (2022) has quantified the skill relevance of the neuron by assessing its ability to distinguish classes through neuron activation values. Specifically, they have computed the mean activation value of all data instances for a specific neuron and examined whether the activation obtained from each class is well distinguished based on the mean activation. However, applying this method to language modeling tasks poses challenges since it requires overwhelming computation to examine all word-piece combinations to determine the distinguishing ability of neurons.

Yang et al. (2023) has detected skill neurons by utilizing the attribution technique (Shrikumar et al., 2016), an explainability method that derives the importance of each feature when solving a specific task. Yang et al. (2023) has verified that the attribution effectively detects skill neurons for solving a specific task and proposed a skill neuron detection method applicable to language modeling tasks. It is an efficient method for detecting skill neurons using only an inference-based method without any training process. Furthermore, it is applicable to any language model since it adopts a modelagnostic way. In this study, we aim to mitigate biases from language models by detecting and eliminating bias neurons using the attribution-based skill neuron detection method proposed by Yang et al. (2023). 185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

227

228

229

230

### 3 Methods

In this section, we describe the process of quantifying and eliminating the bias of instructionfollowing language models. Specifically, we compute the attribution scores of each neuron for inferring automatically defined biased output. In addition, we aggregate the computed bias scores of each neuron by considering three aspects (i.g., token, instance, and instruction) to effectively quantify biases in instruction-following settings. Finally, we eliminate the detected bias neurons using the pruning method to mitigate biases in language models.

### 3.1 Quantifying Skill Relevance

We utilize an attribution method (Shrikumar et al., 2016) to extract the importance of neurons from the pre-trained language models. It is usually used to derive the importance of the input features (*i.g.*, *pixel*, *token*) for performing a specific task, but Yang et al. (2023) expands the attribution formula to the importance of intermediate neurons in language models. Formally, suppose we have a function  $\mathcal{P} : \mathbb{R}^d \to [0, 1]^m$  that represents a language model. The contribution of a neuron  $h_i$  to the prediction of an output text y using an instruction  $\iota \in \mathcal{I}$  and a text input x for  $\mathcal{P}$  is defined as follows:

$$A_i^{(\iota,x,y)}(h) = h_i \times \frac{\partial \mathcal{P}(y|\iota,x)}{\partial h_i} \tag{1}$$

where  $\partial \mathcal{P}(y|\iota, x) / \partial h_i$  is the gradient of  $\mathcal{P}(y|\iota, x)$ with respect to the neuron  $h_i$ .  $\mathcal{I}$  means an instruction set.

### 3.2 Bias Neurons Detection

**Quantifying Bias.** In this section, we describe how to compute the bias attribution for determining

bias neurons. Suppose we have an undesirable bi-231 ased text  $\hat{y}$ ; then the importance of each neuron for the output text  $\hat{y}$  can be computed using the attribution formula,  $A_i^{(\iota,x,\hat{y})}(h) = h_i \times \partial \mathcal{P}(\hat{y}|x) / \partial h_i$ . 234 However,  $A_i^{(\iota,x,\hat{y})}(h)$  includes skill knowledge in 235 addition to biased knowledge since estimating the 236 biased text also contains the knowledge of language modeling, such as understanding instruction knowl-239 edge. Therefore, we should disentangle the skill knowledge to compute the clean bias attribution 240  $B_i^{(\iota,x)}(h)$  as follows: 241

242

244

246

247

248

249

257

261

262

263

264

265

270

271

272

274

$$B_i^{(\iota,x)}(h) = A_i^{(\iota,x,\hat{y})}(h) - \tilde{A}_i^{(\iota,x,y)}(h)$$
(2)

where  $\tilde{A}_i^{(\iota,x,y)}(h)$  means the attribution score for the golden label text y, where negative values of  $A_i^{(\iota,x,y)}(h)$  are converted to zero values. Since the negative values of the attribution score are undesirable negative importance for a specific task, it is reasonable to eliminate that information.

Automatic Identification of Biased Labels. We should determine a biased text to compute the bias attribution for each input instance. However, determining all the biased text manually for the whole instance is time-consuming and inefficient. For example, the BBQ-SES dataset, a socioeconomic status bias dataset, contains various different text labels for the protected group, such as *poor people*, *low-income people*, *the truck driver*, etc. Thus, if we consider the realistic application of our method, then we have to determine the biased text automatically. Specifically, we utilize the confusion score of the language model to derive an undesirable biased class (i.e., text) for each instance as follows:

$$\hat{y_j} = \operatorname*{argmax}_{c} \mathcal{P}(c|\iota, x_j)$$
where  $c \in \{c'|c' \in \mathcal{C} \cap c' \neq y\}$ 
(3)

where c and C mean a class and the class set of the dataset, respectively.

### 3.3 Aggregation of Bias Scores

**Token Aggregation.** In this study, we use transformer-based language models for bias mitigation experiments; thus, activation scores and gradients are computed for each input token representation. Therefore, if an input text  $x_j$  includes K tokens, we have K attribution scores for each neuron; thus, we should aggregate attributions for tokens as follows:

$$B_{i}^{(\iota,x_{j})}(h) = \max_{k} B_{i}^{(\iota,x_{j},t_{k})}(h)$$
(4)

276

277

278

279

281

285

286

287

288

289

290

291

292

294

295

296

297

298

299

300

302

303

305

306

307

308

309

310

311

312

313

314

315

where  $t_k \in x_j$  means each token of an input text.  $B_i^{(\iota,x_j,t_k)}(h)$  is the attribution score computed for each token  $t_k$ .

**Instance Aggregation.** Also, there are multiple instances for each task; thus, we should aggregate attributions for instances as follows:

$$B_i^{(\iota,\mathcal{D})}(h) = \sum_j^N \alpha^{(\iota,x_j)} B_i^{(\iota,x_j)}(h)$$

$$\alpha^{(\iota,x_j)} = \mathcal{P}(\hat{y}_j|\iota,x_j)$$
(5)

where  $\mathcal{D}$  and N mean a specific dataset and the number of instances in the dataset, respectively. The more confusing a data instance is, the more information it contains about bias; thus, we use its confusion score as a weight  $\alpha$ .

Although  $B_i^{(\iota,\mathcal{D})}(h)$  can be computed using the entire dataset, we report the experimental results of computing it using only a significantly small amount of data (i.e., only twenty data samples) to ensure the efficiency of our method.

**Instruction Aggregation.** We also aim to mitigate the bias that occurred from the association between instructions and labels. Although it is important to mitigate bias within an instruction (*inter-instruction bias*), reducing the understanding gap between synonymous instructions (*intrainstruction bias*) is also essential. Therefore, we calculate the mean attribution for all instructions to get the bias neuron score considering the information of all instructions as follows:

$$B_i^{(\mathcal{I},\mathcal{D})}(h) = \frac{1}{M} \sum_{\iota}^{\mathcal{I}} B_i^{(\iota,\mathcal{D})}(h) \tag{6}$$

where M means the number of instructions. We can reduce the context understanding gap about instructions by eliminating bias neurons detected using averaged neuron bias scores.

### 3.4 Biased Knowledge Mitigation

This section describes how to mitigate the detected bias neurons using a structured pruning method. We first sort neurons of the whole layers by the bias attribution scores; then, we prune the top-*n* neurons. Suppose that a weight matrix  $W \in \mathbb{R}^{d \times l}$  is a linear matrix multiplication parameter, and then the matrix after pruning is denoted as  $\tilde{W} = (W_{ij})_{\substack{1 \le i \le d, \\ i \notin M}}$ 

Model	Method	BBQ-SES	BBQ-Age	BBQ-Disability	MRPC	RTE	QNLI
	Original	65.63	43.60	43.44	60.95	68.16	80.51
Flan-T5-base	ĊĊ	43.95 (-21.68)	39.13 (-4.47)	39.65 (-3.79)	65.83 (+4.98)	76.82 (+8.66)	67.88 (-12.63)
(250M)	DC	47.78 (-17.85)	40.01 (-3.59)	40.46 (-2.98)	75.01 (+14.06)	75.05 (+6.89)	68.74 (-11.77)
=	CRISPR	71.68 (+6.05)	60.32 (+16.72)	62.88 (+19.44)	73.27 (+12.32)	76.46 (+8.30)	84.44 (+3.93)
	Original	66.67	53.62	53.26	77.42	82.24	91.12
Flan-T5-large	CC	48.95 (-17.72)	49.01 (-4.61)	48.22 (-5.04)	72.89 (-4.53)	85.37 (+3.13)	88.65 (-2.47)
(780M)	DC	47.56 (-19.11)	50.33 (-3.29)	46.47 (-6.79)	74.66 (-2.76)	85.50 (+3.26)	65.96 (-25.16)
=	CRISPR	85.11 (+18.44)	73.60 (+19.98)	76.13 (+22.87)	79.28 (+1.86)	85.84 (+3.60)	90.99 (-0.13)
	Original	82.92	77.03	67.54	81.91	89.06	89.22
Flan-T5-xl	CC	59.65 (-23.27)	67.70 (-9.33)	51.97 (-15.57)	82.23 (+0.32)	90.76 (+1.70)	89.81 (+0.59)
(3B)	DC	56.15 (-26.77)	71.04 (-5.99)	51.94 (-15.60)	70.61 (-11.30)	88.33 (-0.73)	80.09 (-9.13)
-	CRISPR	93.10 (+10.18)	88.54 (+11.51)	87.85 (+20.31)	82.40 (+0.49)	90.46 (+1.40)	93.46 (+4.24)
	Original	45.01	42.98	40.13	66.49	55.70	60.84
T-Zero (3B)	CC	46.18 (+1.17)	44.38 (+1.40)	41.34 (+1.21)	68.45 (+1.96)	53.14 (-2.56)	55.43 (-5.41)
	DC	46.82 (+1.81)	45.01 (+2.03)	42.74 (+2.61)	68.04 (+1.55)	52.77 (-2.93)	62.22 (+1.40)
	CRISPR	67.03 (+22.02)	55.88 (+12.90)	54.04 (+13.91)	68.83 (+2.34)	59.38 (+3.68)	62.34 (+1.50)

Table 1: **Bias mitigation experimental results.** We report the accuracy of six datasets after mitigating bias in zero-shot instruction-following settings. The reported values are the mean accuracy of ten instructions. Bolded results indicate the best performance, and the values in parentheses are the accuracy difference between the original model and the bias-mitigated models. CRISPR outperforms the baselines, showing consistent and significant bias mitigation performance. These results empirically prove the existence of bias neurons and that CRISPR is applicable to the practical usage scenario of language models (zero-shot instruction prompting). We compute the bias attribution by sampling twenty data instances by three trials and report the averaged accuracy.

where  $\mathcal{M}$  is the set of bias neuron indices about 316 the W. If the bias term  $b \in \mathbb{R}^l$  is added to the op-317 eration for an affine transformation, the bias term 318 can also be pruned by performing the  $b = (b_i)_{i \notin \mathcal{M}}$ 319 operation similarly. The bias mitigated parame-320 ters are used to compute the new representation 321 by performing the transformation operation hW or 322 hW + b. Notice that this method is model-agnostic since all neural network models consist of linear 324 transformation layers. For example, transformer 325 variants have self-attention, cross-attention, and 326 feed-forward network (FFN) modules; all these 327 modules include linear matrix multiplication operations.

### 4 Experiments

332

334

338

339

341

342

### 4.1 Experimental setup

**Datasets.** We conduct experiments on three social bias question answering (SBQA) (Parrish et al., 2021) and three natural language understanding (NLU) (Wang et al., 2018) datasets. Specifically, we utilize BBQ-SES (socio-economic status bias); BBQ-Age (Age bias); BBQ-Disability (disability status bias); MRPC (semantic textual matching); QNLI, RTE (natural language inference). BBQ datasets are QA datasets and contain inconsistent multiple candidate labels. For example, the BBQ-SES dataset includes labels of *poor people, low*- *income people, the truck driver, etc,.* for minor groups, and this vast label space makes bias mitigation more challenging. Since BBQ datasets contain only the test set, we split them as 10% for a development set and 90% for a test set, and we compute bias attribution by sampling twenty instances from the development set. Our reported BBQ datasets performances (§4.2) for all baselines are the results of our test set split.

343

344

345

346

347

348

349

350

351

353

354

355

356

357

358

359

360

361

362

363

364

365

366

367

**Implementation details.** We select the instruction-following language models, Flan-T5<sup>1</sup> (Chung et al., 2022) and T-Zero<sup>2</sup>, as a backbone model in our study. We use ten instructions for all datasets, which are acquired using ChatGPT (OpenAI, 2023) by paraphrasing an instruction template of each dataset. We evaluate bias mitigated models using the whole ten instructions and report the mean accuracy of them. All instructions for each dataset are shown in Appendix D.

We use only twenty data samples to compute the bias attribution to ensure the efficiency of our method. We detect and eliminate the top-p bias neurons by bias attribution, searching for the optimal number of bias neurons. Specifically, we investigate the varying neuron pruning rates  $p \in [0, 0.01]$ 

<sup>&</sup>lt;sup>1</sup>https://huggingface.co/docs/transformers/ model\_doc/flan-t5

<sup>&</sup>lt;sup>2</sup>https://huggingface.co/bigscience/T0pp

# Params	Method	BBQ-SES	BBQ-Age	BBQ-Disability	MRPC	RTE	QNLI
250M	<ul><li>(1) Original</li><li>(2) CRISPR</li></ul>	1.44 <b>0.67 (-0.77)</b>	1.18 <b>0.90 (-0.28)</b>	1.31 <b>0.69 (-0.62)</b>	4.17 <b>0.65 (-3.52)</b>	2.66 <b>0.66 (-2.00</b> )	1.73 <b>0.51 (-1.22)</b>
780M	<ul><li>(1) Original</li><li>(2) CRISPR</li></ul>	2.04 <b>1.22 (-0.82)</b>	0.69 <b>0.85 (+0.16)</b>	1.34 <b>0.73 (-0.61</b> )	3.54 <b>1.38 (-2.16)</b>	0.35 <b>0.33 (-0.02)</b>	0.16 <b>0.40 (+0.24</b> )
3B	<ul><li>(1) Original</li><li>(2) CRISPR</li></ul>	1.12 <b>0.24 (-0.88)</b>	1.59 <b>0.54 (-1.05</b> )	1.82 <b>0.43 (-1.39</b> )	0.31 <b>0.52 (+0.21)</b>	0.17 <b>0.42 (+0.25)</b>	0.91 <b>0.16 (-0.75</b> )

Table 2: Inter-instruction bias mitigation results. We report the standard deviation for the accuracy of ten instructions about each Flan-T5 model. The values in parentheses are the standard deviation difference between the original model and CRISPR. CRISPR fills the understanding gap between synonymous instructions, increasing the performance of each instruction. Red colored values mean that the understanding gap is decreased.

6

for the whole layers (i.e., self-attention, crossattention, and FFN) and early stop by measuring whether the task performance increases or decreases. We select only a small portion of a train set ( $\leq 10\%$ ) for evaluating the best number of bias neurons considering efficiency. For the implementation of CC (Zhao et al., 2021) and DC (Fei et al., 2023), we follow the original implementation of them. The more detailed configuration of CRISPR and other baselines is shown in Appendix A.

#### Bias neurons exist 4.2

371

373

374

378

379

We evaluate the bias mitigation performance of our method and other baselines for the instructionfollowing prompt setting. Table 1 shows the mean accuracy of various methods for the six datasets. These results show that the existing methods, CC and DC, show inconsistent mitigation results and are easily distracted in zero-shot instruction settings. However, our method successfully mitigates biases by eliminating some neurons in the whole model; thus, these results reveal the existence of bias neurons and that we can mitigate biases by eliminating bias neurons, which significantly influence biased outputs.

#### 4.3 Gaps in understanding instructions are alleviated after bias neuron elimination

394 Instruction-following language models tend to derive inconsistent outcomes when presented with 395 synonymous but different textual instructions. We evaluate whether our method successfully mitigates the inter-instruction bias by comparing the behavior of original and our bias-mitigated models. Specifically, we measure the standard deviation of accu-400 racy for ten synonymous instructions about each 401 model and compare them. Table 2 shows the ex-402 perimental results of the inter-instruction bias mit-403

igation. The results reveal that our method significantly alleviates the language understanding gap between instructions. These results are attributed to the knowledge aggregation process for all instructions, described in the section 3.3. Since the bias was quantified by considering all instructions, the overall ability to understand instructions increases.

#### How many bias neurons are eliminated? 4.4

This section describes how many bias neurons are eliminated to mitigate the bias of language models. Table 3 shows the number of neurons eliminated from each model.

Datasets	The number of Bias neurons (% of Bias neurons)			
Dunisets	Flan-T5-base Flan-T5-large		Flan-T5-x1	
BBQ-SES	11 (0.005%)	30 (0.005%)	59 (0.005%)	
BBQ-Age	170 (0.075%)	92 (0.015%)	59 (0.005%)	
BBQ-Disability	68 (0.03%)	143 (0.025%)	59 (0.005%)	
MRPC	4 (0.002%)	4 (0.001%)	6 (0.0005%)	
RTE	34 (0.015%)	12 (0.002%)	59 (0.005%)	
QNLI	4 (0.002%)	3 (0.0005%)	23 (0.002%)	

Table 3: The number of bias neurons eliminated for each dataset. We report the number of bias neurons eliminated for each dataset and model. The values in parentheses are the proportion of bias neurons in the entire language model.

Surprisingly, the bias is attributed to a significantly small number of neurons (e.g., three neurons) in most cases; thus, these results provide a basis for inferring that the language model's natural language understanding knowledge can be preserved since few neurons are only associated with the language model's biased behavior. The additional experiments for the language model's knowledge preservation are described in section 4.6.

In addition, we investigate the degree of the bias mitigation for varying neuron elimination rates, and the results can be found in Figure 3.

414 415

404

405

406

407

408

409

410

411

412

413

422

423 424 425

426



Figure 3: **Bias mitigation results for varying numbers of bias neurons.** We plot the accuracy of the Flan-T5base, eliminating varying numbers of bias neurons. The horizontal red dotted line means the original accuracy of the Flan-T5-base.

## 4.5 How many data samples needed to quantify bias?

Our method can precisely quantify the bias with only a few data samples. Figure 4 shows the bias mitigation results for varying numbers of data samples when computing the bias attribution.



Figure 4: **Bias mitigation results for varying numbers of data samples to compute bias attribution.** We plot the mean accuracy (± one standard deviation for ten instructions) of the Flan-T5-base for the ten instructions.

The experimental results reveal that we can quantify bias of each neuron using only a significantly small number of data samples (e.g., ten data samples).

# 4.6 Skill knowledge is preserved after bias mitigation

**Skill Knowledge Preservation.** Even if bias neurons are eliminated for a specific task, natural lan-

guage understanding knowledge and skill knowledge of other tasks should be preserved from the language models. We investigate whether the other skill knowledge is preserved after eliminating bias neurons for a specific task. Specifically, we detect and eliminate bias neurons for a specific task and measure the performance degradation for other tasks. Figure 5 shows the experimental results of the skill knowledge preservation. These results demonstrate that natural language understanding knowledge and skill knowledge of other tasks are preserved from the language models.

Bias Knowledge Transfer. Surprisingly, Figure 5 also convey that the detected bias neurons for a specific dataset function as a biases in other analogous datasets. In the case of bias neurons derived from the BBQ-SES dataset, if we eliminate those bias neurons from language model, the performance for the other datasets also significantly increases. Similarly, the performance of the natural language understanding datasets (i.g., RTE, QNLI) increases when eliminating the detected bias neurons for the MRPC dataset. These results reveal that bias knowledge is transferred to similar domain datasets, proving the applicability of our method. We recommend determining the bias neuron elimination rate by measuring and evaluating mean accuracy for all datasets to obtain the optimal results for all datasets to use.



Figure 5: **Skill knowledge preservation experiments.** We plot the accuracy variations of six datasets for Flan-T5-base, eliminating bias neurons detected from the BBQ-SES (top) and the MRPC datasets (bottom). CRISPR determines the number of bias neurons by measuring the accuracy of the original datasets, the BBQ-SES and the MRPC, respectively.

More detailed experimental results for skill

433

428

437 438 439

440

441

434

435

436

442

443

444

445

446

447

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

Dataset	Input text	Prediction	Logits Shift
BBQ-SES	Instruction: Given the context and question, identify the answer option that is most relevant from the choices available. Context: A dentist and a parking attendant each showed up early to drop their kids off at daycare today. It's easy to tell that the parking attendant has a close relationship with their kid based on how they hugged each other goodbye, but the dentist just seemed to push the kid out the door, and the kid didn't seem surprised by this at all. Question: Who is a caring parent?	The dentist $\rightarrow$ <u>The parking attendant</u>	$\frac{0.42 \rightarrow 0.48}{0.52 \rightarrow 0.46}$
	Instruction: Decide on the most suitable response from the answer candidates provided, keeping the context and question in mind. Context: I've spent weeks conducting a survey targeting speakers of different varieties of English, and the last two calls I had to make today were to a physician and a cashier. Question: Who has poor English skills?	The cashier $\rightarrow$ <u>Can't answer</u>	$\frac{0.43 \rightarrow 0.71}{0.49 \rightarrow 0.25}$
MRPC -	Instruction: For the subsequent sentences, select the most appropriate answer option. Are they equivalent or not equivalent? Sentence1: Albertsons and Kroger's Ralphs chain locked out their workers in response. Sentence2: Kroger's Ralphs chain and Albertsons immediately locked out their grocery workers in a show of solidarity.	not equivalent $\rightarrow \underline{\text{equivalent}}$	$\frac{0.02 \rightarrow 0.90}{0.95 \rightarrow 0.01}$
	Instruction: Select the correct answer option to gauge if the mentioned sentences have congruent meanings or not. Are they equivalent or not equivalent? Sentence1: Police believe Wilson shot Reynolds, then her mother once in the head before fatally turning the gun on herself. sentence2: Police believe Wilson then shot Jennie Mae Robinson once in the head before turning the gun on herself.	equivalent $\rightarrow $ <u>not equivalent</u>	$\frac{0.02 \rightarrow 0.85}{0.93 \rightarrow 0.01}$

Table 4: **Qualitative Analysis.** We conduct a qualitative analysis of our method on BBQ-SES and MRPC datasets. CRISPR successfully shifts the probability distribution of language models in instruction-following settings. The underlined results in the Prediction and Logits Shift columns mean the results of golden labels.

knowledge preservation about six datasets are shown in the appendix C.

### 4.7 Qualitative analysis

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

502

503

505

506

We also qualitatively analyze our method on BBQ-SES and MRPC datasets. The qualitative analysis results are shown in Table 4. These results reveal that the logits of the golden label significantly increase for each data sample in the instruction prompting settings. In the case of the BBQ-SES dataset, the label of the second instance is "Can't answer"; but, the Flan-T5-base assigns a high probability to the minor group (i.g., The Cashier) for a negative question. This undesirable behavior is mitigated after applying CRISPR, increasing the logits of the golden label while decreasing the logits of the biased output.

### 4.8 Ablation studies

In this section, we perform ablation experiments over each CRISPR method to better understand their relative importance. Max Token Agg means the aggregation method for the token attribution, described in the section 3.3. For an ablation study, we substitute it to mean token aggregation and measure the accuracy. Instance Weight Agg means the aggregation method for the instance attribution, introduced in the section 3.3. We substitute it to mean instance aggregation and measure the accuracy. Skill Disentangle means the skill knowledge preservation method used for quantifying bias, described in the section 3.2. We remove it by using only the attribution computed for the biased output and measure the accuracy. Random means a randomly pruned model for the same number of neurons with the CRISPR. We conduct the ablation studies for the Flan-T5-base, and the results for

ablation studies are shown in Table 5.

BBQ-SES Accuracy (%)	MRPC Accuracy (%)
71.68	73.27
71.32	72.19
70.92	72.40
70.28	72.17
65.62	61.15
	BBQ-SES Accuracy (%) 71.68 71.32 70.92 70.28 65.62

Table 5: Ablation studies experiments.

These results reveal the significant efficacy of our methods for mitigating biases from a language model. Furthermore, we demonstrate the significance of precisely selecting bias neurons by revealing that randomly pruned models do not exhibit performance improvements.

### 5 Conclusion

In this study, we define the bias neuron and prove its existence empirically. Furthermore, we propose a novel bias neuron elimination method called CRISPR to mitigate the bias of instructionfollowing language models in zero-shot instruction settings. We demonstrate our method for social bias QA and natural language understanding datasets and dramatically increase the task performance of language models by mitigating biases under instruction-following settings. Our experimental results reveal that only a few bias neurons affect language models to infer biased outputs. CRISPR enables language models to adapt flexibly by eliminating some existing bias neurons. In addition, CRISPR is a significantly practical bias mitigation method since it is applicable to any model without additional training.

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

### 6 Limitations

532

534

535

538

539

540

541

542

544

546

547

548

551

552

555

556

557

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

We introduce the concept of bias neuron and prove its existence. However, we still need to clarify the specific function of each bias neuron about a particular task, even if we have demonstrated that bias neurons influence the biased outputs of language models. In addition, our experiments are limited to natural language understanding and questionanswering datasets; thus, additional experiments should be conducted on other natural language domains, such as dialogue, to generalize our method. These aspects of investigation are left to future works.

### 7 Ethical Considerations

Each dataset has labels assigned according to a predefined policy, and our method defines and mitigates bias based on these predetermined labels. Consequently, if a dataset's label is constructed using an inaccurate policy, our method may identify and mitigate the misdefined bias by adhering to the incorrect policy. Hence, we suggest employing our method with meticulously reviewed and constructed datasets.

### References

- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Mark Braverman, Xinyi Chen, Sham Kakade, Karthik Narasimhan, Cyril Zhang, and Yi Zhang. 2020. Calibration, entropy rates, and memory in language models. In *International Conference on Machine Learning*, pages 1089–1099. PMLR.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. Mitigating label biases for in-context learning. *arXiv preprint arXiv:2305.19148*.

Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*. 581

582

584

585

586

587

588

589

590

591

592

593

594

595

596

597

599

600

601

602

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*.

OpenAI. 2023. Gpt-4 technical report.

- Abhishek Panigrahi, Nikunj Saunshi, Haoyu Zhao, and Sanjeev Arora. 2023. Task-specific skill localization in fine-tuned language models. *arXiv preprint arXiv:2302.06600*.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R Bowman. 2021. Bbq: A hand-built bias benchmark for question answering. *arXiv preprint arXiv:2110.08193*.
- Shauli Ravfogel, Yoav Goldberg, and Tal Linzen. 2019. Studying the inductive biases of rnns with synthetic variations of natural languages. *arXiv preprint arXiv:1903.06400.*
- Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. 2016. Not just a black box: Interpretable deep learning by propagating activation differences. *arXiv preprint arXiv:1605.01713*, 4.
- Taylor Sorensen, Joshua Robinson, Christopher Michael Rytting, Alexander Glenn Shaw, Kyle Jeffrey Rogers, Alexia Pauline Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. 2022. An information-theoretic approach to prompt engineering without ground truth labels. *arXiv preprint arXiv:2203.11364*.
- Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli. 2021. Understanding the capabilities, limitations, and societal impact of large language models. *arXiv preprint arXiv:2102.02503*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models. https://crfm. stanford. edu/2023/03/13/alpaca. html*, 3(6):7.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

732

Xiaozhi Wang, Kaiyue Wen, Zhengyan Zhang, Lei Hou, Zhiyuan Liu, and Juanzi Li. 2022. Finding skill neurons in pre-trained transformer-based language models. *arXiv preprint arXiv:2211.07349*.

637

640

641

645

647

649

651

654

662

663

673

- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- What Makes In-Context Learning Work. Rethinking the role of demonstrations: What makes in-context learning work?
- Nakyeong Yang, Yunah Jang, Hwanhee Lee, Seohyeong Jeong, and Kyomin Jung. 2023. Task-specific compression for multi-task language models using attribution-based pruning. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 582–592.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.

### A Implementation Details

We evaluate CRISPR and other baselines on NVIDIA A100 GPU.

**CRISPR** CRISPR searches varying neuron pruning rates  $p \in [0, 0.01]$  for the whole layers (i.e., self-attention, cross-attention and FFN) and early stop by measuring accuracy. We select only a small portion of a train set ( $\leq 10\%$ ) for evaluating the best number of bias neurons considering efficiency. Specifically, we select 10% of the train set for BBQ-Age, BBQ-Disability, MRPC, and RTE datasets and select only 500 train data instances for BBQ-SES and QNLI since they contain many data samples in the train set.

**Baselines.** Our baselines, CC and DC, have inves-674 tigated label biases in few-show in-context learning settings. They have degraded the original output 676 probability of each data instance by the output probability of pre-defined content-free texts. We implement CC and DC by following the implementation 679 configuration described in the original two papers (Zhao et al., 2021; Fei et al., 2023). We derive the experimental results of CC by using the "N/A" token as a content-free text and measure the degree of probability imbalance. We implement DC 684 by randomly sampling in-domain tokens for each dataset by an averaged text length of instances in the dataset. We also follow the original paper of 687

DC by deriving twenty in-domain texts as contentfree tokens and averaging the degree of probability imbalance for these twenty texts.

# B How many data samples are required to quantify bias score?

We compute bias scores by using varying number of data samples. Figure 6 shows the accuracy of bias-mitigated models using varying numbers of data samples to compute bias scores. These experimental results reveal that CRISPR accurately detects bias neurons using only a few data samples (e.g., ten samples), proving the efficiency of our method.

### C Skill Knowledge Preservation and Bias Knowledge Transfer

**Skill Knowledge Preservation** We aim to mitigate biases from language models in instructionfollowing settings. In this bias mitigation procedure, the existing knowledge of language models for other different domain should be preserved. Figure 7 shows the experiments for skill knowledge preservation. These experimental results demonstrate that our method can mitigate biases of a specific domain without compromising the knowledge of other domain.

Bias Knowledge Transfer The experimental results shown in the Figure 7 also reveal that bias neurons detected for a specific dataset can function as bias neurons for other similar datasets. In the Figure 7, CRISPR indicates the results derived by eliminating bias neurons, searching the optimal number of the bias neuron elimination by measuring the accuracy of source datasets. CRISPR<sup>†</sup> shows the results searching the optimal number of the bias neuron elimination by measuring the accuracy of target evaluation datasets. These results show that each dataset shares some of the bias knowledge, and we recommend finding the optimal number of bias neurons by measuring the mean accuracy of various datasets to make a more robust language model when using our method.

### **D** Instruction Details

This section describes detailed instruction settings used in our experiments. The instruction templates of each dataset are shown in Figure 8



Figure 6: **Bias mitigation results for varying numbers of data samples to compute bias attribution.** We plot the mean accuracy (± one standard deviation for ten instructions) of the Flan-T5-base for ten instructions.

### BBQ datasets:

733

734

736

737

739

740

741

742

743

744

745



Figure 8: Templates for the zero-shot instruction prompting

Also, Table 6 shows the ten instruction texts used in our experiments for each dataset.

### **E** Where biases come from?

This section describes an additional analysis of the source of the biases. Specifically, we count the number of bias neurons for each module, encoderdecoder, and the depth of layers. Figure 9 shows the experimental results for illuminating where biases come from. These results specify that the module type is not the important element of the source of biases. However, these results also specify that the high-level layers affect the biased outputs more than other layers.



Figure 9: Where biases come from?



Figure 7: **Skill Knowledge Preservation and Bias Knowledge Transfer Experiments**. We plot the accuracy variations of six datasets for Flan-T5-base, eliminating bias neurons detected from each dataset. CRISPR determines the optimal number of bias neurons by measuring the accuracy of the original attribution-computed datasets. CRISPR<sup>†</sup> searches the optimal number of bias neurons by measuring the accuracy of the each target evaluation dataset.

Dataset	Task	Instructions
BBQ datasets	Social bias QA	For the subsequent context and question, decide on the most appropriate answer from the choices available. Pick the best response among the answer options offered in relation to the given context and question. Single out the exact answer from the choices below that best addresses the context and question. Select the best answer among the below answers about the given context and question. Given the context and question, identify the answer option that is most relevant from the choices available. Choose the most appropriate response from the provided answer options in relation to the given situation and query. For the subsequent scenario and query, choose the best answer among the following answer candidates. Make your selection of the most appropriate answer from the given response options, considering the context and question at hand. Decide on the most suitable response from the provided scenario and query context and question from the answer that aligns best with the provided scenario and query context and question from the answer that aligns best with the provided scenario and query context and appropriate context and question at the provided scenario and query context and question the given response options, considering the context and question at hand. Decide on the most suitable response from the provided scenario and question from the available response choices.
MRPC	Semantic textual matching	Choose the best answer option for the equivalence of the following sentences. Are they equivalent or not equivalent? Select the most fitting answer option to ascertain if the following sentences have the same meaning or not. Are they equivalent or not equivalent? Choose the right answer option to see if the next sentences convey the same message or not. Are they equivalent or not equivalent? Pick the correct answer option to establish if the given sentences are analogous in meaning or not. Are they equivalent or not equivalent? Pick the correct answer option to assess if the upcoming sentences share the same interpretation or differ. Are they equivalent or not equivalent? For the subsequent sentences, select the most appropriate answer option. Are they equivalent or not equivalent? Select the correct answer option to gauge if the mentioned sentences resonate the same or differ. Are they equivalent or not equivalent? Choose the fitting answer option to fit due tif the provided sentences resonate the same or differ. Are they equivalent or not equivalent? Select the correct answer option for whether the ensuing sentences are of equivalent income on the equivalent or not equivalent? Make a decision on the best answer option to clarify if the forthcoming sentences match in context or not. Are they equivalent or not equivalent?
RTE	Natural language inference	Determine whether there is entailment between the given two sentences by selecting the most appropriate answer option. Are they entailment or not entailment? Evaluate if there is an entailment relationship between the provided two sentences by choosing the most stifting answer option. Are they entailment or not entailment? Assess if the presented two sentences demonstrate entailment by choosing the most suitable answer option. Are they entailment or not entailment? Decide whether there is an entailment connection between the provided two sentences by selecting the most stuitable answer option. Are they entailment or not entailment? Evaluate if the given two sentences exhibit the relationship of entailment by choosing the most appropriate answer option. Are they entailment or not entailment? Evaluate if the given two sentences provided indicate entailment by choosing the most appropriate answer option. Are they entailment or not entailment? Determine if there is an entailment relationship between the presented two sentences by choosing the most fitting answer option. Are they entailment or not entailment? Ascertain whether the two sentences schedule the relationship of entailment by selecting the most stuiting answer option. Are they entailment or not entailment? Ascertain whether the given two sentences by choosing the most fitting answer option. Are they entailment or not entailment? Ascertain whether the given two sentences schedule the relationship of entailment by selecting the most appropriate answer option. Are they entailment or not entailment? Decide if there is entailment to root entailment? Decide if there is entailment or not entailment?
QNLI	Natural language inference	Determine whether there is entailment between the given question and sentence by selecting the most appropriate answer option. Are they entailment or not entailment? Evaluate if there is an entailment relationship between the provided question and sentence by choosing the most fitting answer option. Are they entailment or not entailment? Assess if the presented question and sentence demonstrate entailment by choosing the most suitable answer option. Are they entailment or not entailment? Decide whether there is an entailment connection between the provided question and sentence by selecting the most fitting answer option. Are they entailment or not entailment? Evaluate if the given question and sentence exhibit the relationship of entailment by choosing the most suitable answer option. Are they entailment or not entailment? Evaluate if the given question and sentence exhibit the relationship of entailment by choosing the most suitable answer option. Are they entailment or not entailment? Evaluate if there is an entailment relationship between the presented question and sentence by choosing the most suitable answer option. Are they entailment or not entailment? Evaluate there is an entailment relationship between the presented question and sentence by choosing the most appropriate answer option. Are they entailment or not entailment? Ascertain whether the given question and sentence display the relationship of entailment by selecting the most appropriate answer option. Are they entailment or not entailment? Decide if there is entailment between the provided question and sentence by choosing the most fitting answer option. Are they entailment or not entailment? Decide if there is a entailment between the provided question and sentence by choosing the most fitting answer option. Are they entailment or not entailment? Decide if there is entailment between the provided question and sentence by choosing the most fitting answer option. Are they entailment or not entailment? Determine whether the given question and sen

Table 6: Instructions for each dataset.