

# 000 001 QUANTIFYING THE ACCURACY-INTERPRETABILITY 002 TRADE-OFF IN CONCEPT-BASED SIDECHANNEL 003 MODELS 004 005

006 **Anonymous authors**  
007 Paper under double-blind review  
008  
009  
010  
011  
012

## ABSTRACT

013 Concept Bottleneck Models (CBNMs) are deep learning models that provide inter-  
014 pretability by enforcing a bottleneck layer where predictions are based exclusively  
015 on human-understandable concepts. However, this constraint also restricts infor-  
016 mation flow and often results in reduced predictive accuracy. Concept Sidechannel  
017 Models (CSMs) address this limitation by introducing a sidechannel that bypasses  
018 the bottleneck and carry additional task-relevant information. While this improves  
019 accuracy, it simultaneously compromises interpretability, as predictions may rely  
020 on uninterpretable representations transmitted through sidechannels. Currently,  
021 there exists no principled technique to control this fundamental trade-off. In  
022 this paper, we close this gap. First, we present a unified probabilistic concept  
023 sidechannel meta-model that subsumes existing CSMs as special cases. Build-  
024 ing on this framework, we introduce the Sidechannel Independence Score (SIS),  
025 a metric that quantifies a CSM’s reliance on its sidechannel by contrasting predic-  
026 tions made with and without sidechannel information. We propose SIS regular-  
027 ization, which explicitly penalizes sidechannel reliance to improve interpretabil-  
028 ity. Finally, we analyze how the expressivity of the predictor and the reliance  
029 of sidechannel jointly shape interpretability, revealing inherent trade-offs across  
030 different CSM architectures. Empirical results show that state-of-the-art CSMs,  
031 when trained solely for accuracy, exhibit low representation interpretability, and  
032 that SIS regularization substantially improves their interpretability, intervenabil-  
033 ity, and the quality of learned interpretable task predictors. Our work provides  
034 both theoretical and practical tools for developing CSMs that balance accuracy  
035 and interpretability in a principled manner.  
036  
037

## 1 INTRODUCTION

038  
039 Concept-based models (CBMs) have emerged as a promising direction for interpretable deep learn-  
040 ing (Poeta et al., 2023; Koh et al., 2020; Espinosa Zarlenga et al., 2022; Mahinpei et al., 2021).  
041 CBMs achieve interpretability by incorporating high-level, human-understandable *concepts* explic-  
042 itely within their architecture. A well-known example CBM is the Concept Bottleneck Model (Koh  
043 et al., 2020) (CBNM), which first predicts concepts (e.g. *whiskers*, *tail*) using a neural network, and  
044 then maps these predicted concepts to the target task using a linear layer (e.g. *cat*). This architecture  
045 makes predictions inherently explainable: one can trace decisions back to predicted concepts (e.g.  
046 ‘the model predicted *cat* because it sees *whiskers* and a *tail*’), and the linear layer provides additional  
047 transparency by revealing the influence of each concept on the final prediction.  
048

049 Despite their appeal, early CBMs such as CBNMs suffered from a significant drop in task accuracy  
050 compared to black-box models. This is because concepts in CBNMs form an *information bottleneck*:  
051 the model must rely solely on concepts to perform the task. In practice, it is often infeasible to design  
052 a concept set that fully captures all the information needed for accurate predictions, which makes a  
053 performance gap inevitable.

To address this limitation, recent works have augmented CBMs with an additional *sidechannel* that  
transmits extra information beyond concepts (Mahinpei et al., 2021; Espinosa Zarlenga et al., 2022;

054 Barbiero et al., 2023; Sawada & Nakamura, 2022).<sup>1</sup> This sidechannel can take different forms, such  
 055 as embeddings (Yuksekgonul et al., 2022) or one-hot masks (Debot et al., 2024), depending on the  
 056 CBM in question. By leveraging this sidechannel, CBMs can close the accuracy gap with black-box  
 057 models, often achieving similar performance irrespective of the used concept set.

058 However, this gain in accuracy comes at the cost of some interpretability. Unlike in CBNMs, predictions  
 059 in sidechannel CBMs are no longer solely determined by concepts, but also by uninterpretable  
 060 units of informations (e.g. an embedding). The more a model relies on information captured by its  
 061 sidechannel, the more accurate but less interpretable it becomes. In the extreme, the task predictor  
 062 could rely completely on the sidechannel, bypassing the concepts entirely.

063 While some works have acknowledged this tension in a limited way by designing techniques that try  
 064 to maximize concept usage (Kalampalikis et al., 2025; Havasi et al., 2022; Shang et al., 2024), we  
 065 show that this is a weaker criterion than minimizing sidechannel reliance. The interpretability cost  
 066 is never quantified or explicitly optimized, and most evaluations and training objectives still focus  
 067 on accuracy, leaving two key gaps in the literature: the lack of a *metric* to measure interpretability  
 068 cost in sidechannel CBMs, and the lack of a principled way to *optimize* for this cost during training.

069 In this work, we aim to close these gaps. Our contributions are the following:  
 070

- 071 **Unified view:** We show that existing sidechannel CBMs can all be interpreted as different  
 072 parametrizations of a single *meta-model*, which we formalize as a high-level probabilistic  
 073 graphical model (PGM).
- 074 **Interpretability metric:** Using this PGM, we provide a natural evaluation method for  
 075 sidechannel CBMs by equipping them with a *bottleneck mode*, where predictions only de-  
 076 pend on the concepts. Based on bottleneck mode, we subsequently introduce the *Sidechan-  
 077 nel Independence Score* (SIS), which quantifies this interpretability cost.
- 078 **Training objective:** We propose *SIS regularization*, a method that explicitly optimizes  
 079 sidechannel CBMs for interpretability by penalizing reliance on the sidechannel.
- 080 **Interpretability discussion:** We discuss how different notions of interpretability in CSMs  
 081 arise when considering both the expressivity of the predictor and its reliance on the  
 082 sidechannel.

083 Through contributions (1), (2) and (4), we provide a clearer understanding of the structure and trade-  
 084 offs of sidechannel CBMs. Through (3), we introduce practical ways for developing models that are  
 085 not only accurate but also interpretable.

## 087 2 BACKGROUND

090 **Concept Bottleneck Models.** Concept Bottleneck Models (CBNMs) (Koh et al., 2020) are CBMs  
 091 that consist of two functions: a *concept predictor* ( $X \rightarrow C$ ) that maps some low-level input features  
 092  $X$  (e.g. an image) to high-level, human-understandable concepts  $C$  (e.g. *whiskers*, *tail*), and a *task  
 093 predictor* ( $C \rightarrow Y$ ) that maps the concepts to some target task  $Y$  (e.g. *cat*). CBNMs are trained  
 094 by directly supervising both concepts and task with the goal of aligning each concept to a human  
 095 interpretation and obtaining high task performance. This supervision comes either from concept and  
 096 task labels in the dataset or from vision-language models, the latter removing the need for expensive  
 097 human annotations (Oikarinen et al., 2023). Typically, CBNMs use a neural network as concept  
 098 predictor and a linear layer or neural network as task predictor. A key downside of CBNMs is that  
 099 their accuracy for predicting the task  $Y$  is limited by the employed concepts  $C$ , as they form an  
 100 information bottleneck.

101 **Concept Sidechannel Models.** Concept Sidechannel Models (CSMs) are CBMs that address the  
 102 information bottleneck issue by predicting  $Y$  not only using  $C$  but also using some additional in-  
 103 formation (Espinosa Zarlenga et al., 2022; Sawada & Nakamura, 2022; Yuksekgonul et al., 2022;  
 104 Barbiero et al., 2023). This additional information comes in different forms for different CSMs.  
 105 Some examples of this are an embedding predicted from  $X$  (Mahinpei et al., 2021) and a one-hot  
 106 mask predicted from  $X$  (Debot et al., 2024). **The central idea behind CSMs is primarily to achieve  
 107 (near) black-box accuracy, and secondarily to be as interpretable as possible.**

<sup>1</sup>In the literature, this sidechannel is also often referred to as a "residual" or "sidepath".

108 **Notation.** We write random variables in upper case (e.g.  $p(X)$ ) and their assignments in lower  
 109 case (e.g.  $p(X = x)$ ). When it is clear from the context, we will abbreviate assignments (e.g.  
 110  $p(x)$  means  $p(X = x)$ ). Conditional distributions are written using a vertical bar (e.g.  $p(Y|X)$  or  
 111  $p(Y = y|X = x)$ ).

### 113 3 METHOD

116 In this section, we present our proposed method. We begin by clarifying the distinction between  
 117 representation interpretability, which refers to the interpretability of the model’s intermediate rep-  
 118 resentations, and functional interpretability, which refers to the interpretability of the prediction  
 119 process (Section 3.1). Subsequently, we introduce a probabilistic *CSM meta-model*, which provides  
 120 a unified framework for representing all CSMs (Section 3.2). We then demonstrate how this meta-  
 121 model enables two modes of inference: the *default* mode and the *bottleneck* mode, in which the  
 122 sidechannel is deactivated (Section 3.3). The delineation of these two modes naturally motivates  
 123 the introduction of a novel metric, the Sidechannel Independence Score (SIS), which quantifies the  
 124 distance between the modes and, consequently, the dependence of the original CSM on its sidechan-  
 125 nel (Section 3.4). We then show how this distance can be employed as a regularization criterion for  
 126 CSMs, providing explicit control over the accuracy–interpretability trade-off in CSMs. Finally, in  
 127 Section 3.5, we discuss how the sidechannel not only supplements task-relevant information miss-  
 128 ing in the concepts, but also enhances concept-to-task expressivity for some CSMs, highlighting a  
 129 trade-off between representation interpretability and functional interpretability.

#### 130 3.1 REPRESENTATION INTERPRETABILITY VS FUNCTIONAL INTERPRETABILITY

131 Based on existing informal notions within the CBM community (Barbiero et al., 2025), we make  
 132 two distinctions in interpretability for CBMs: *representation interpretability* and *functional inter-  
 133 pretability*. We first consider the classical form of interpretability used in machine learning.

134 **Definition 3.1 (Functional Interpretability of CBMs).** A *CBM* is functionally interpretable if and  
 135 only if the mapping from the *CBM*’s concepts (and sidechannel) to the task is an interpretable  
 136 function.

137 What one considers an interpretable function is subjective to the human user. According to Rudin  
 138 et al. (2022), standard cases of interpretable functions include linear layers, logic rules and small  
 139 decision trees. For CSMs, only considering this form of interpretability is insufficient as CSMs also  
 140 make their task prediction using uninterpretable representations (i.e. the sidechannel). Therefore,  
 141 we must also consider a form of interpretability that depends on the use of such uninterpretable  
 142 representations: *representation interpretability*.

143 **Definition 3.2 (Representation Interpretability of CBMs).** A prediction is said to be representation  
 144 interpretable if and only if it is derived exclusively from units of information that are themselves  
 145 interpretable. A concept-based model is *more representation interpretable if a larger fraction of its*  
 146 *task-level predictions are representation-interpretable*.

147 We make the following assumption regarding interpretable units:

148 **Assumption 3.1 (Interpretable Units of Information in CBMs).** In a *CBM*, the concepts  $C$  constitute  
 149 the only interpretable units, since they are explicitly supervised to align with some human under-  
 150 standing.<sup>2</sup> All other internal variables are uninterpretable, irrespective of their form, as they lack  
 151 explicit human alignment (e.g. sidechannels).

152 **Proposition 3.1 (Representation Interpretability of CBNMs and CSMs).** Under Assumption 3.1:

- 153 1. A *CBNM* is fully representation-interpretable irrespective of the used task predictor (e.g.  
 154 linear layer or neural network): every prediction  $\hat{y}$  depends only on the concepts.
- 155 2. A *CSM* is partially representation-interpretable: a prediction  $\hat{y}$  is representation-  
 156 interpretable if and only if  $\hat{y}$  does not depend on the sidechannel.

157 <sup>2</sup>This assumes appropriate measures have been taken to prevent concept leakage (Marconato et al., 2022),  
 158 which may otherwise misalign the concepts.

162 **Example 3.1.** Consider a CSM with a single sidechannel neuron and a linear layer mapping the  
 163 concepts and neuron to the task. If the class label  $y$  is directly encoded in the sidechannel neuron, the  
 164 model can predict  $y$  by relying solely on this neuron. This model is not representation interpretable  
 165 (since predictions only use the sidechannel) but functionally interpretable (since the linear layer  
 166 clearly shows how  $y$  is obtained). Conversely, if the sidechannel neuron carries no information (a  
 167 dummy value), but the model uses a neural network on the concepts and neuron, then predictions are  
 168 representation interpretable (since they depend only on concepts) but not functionally interpretable  
 169 (due to the black-box predictor).

170 Current CSMs focus on achieving functional interpretability while maintaining high accuracy, and  
 171 do not consider representation interpretability, despite its importance as illustrated above. **There-**  
 172 **fore, our focus lies on evaluating and improving representation interpretability in CSMs.**  
 173

174 3.2 THE META-MODEL OF CSMs  
 175

176 A wide range of CSMs have been developed in recent years. While they differ in many ways, their  
 177 high-level structure is very similar. These models can be understood as different instantiations of  
 178 a common underlying high-level model, which we call the *CSM meta-model*. This CSM meta-  
 179 model consists of three functions: a concept predictor  $\phi_c : X \rightarrow C$ , a sidechannel predictor  $\phi_z : X \rightarrow Z$ , and a task predictor  $\phi_y : C, Z \rightarrow Y$ . Both the functions  $(\phi_c, \phi_z, \phi_y)$  and the variables  
 180  $(C, Z, Y)$  can differ across CSMs. For example, concepts may be continuous ( $C = \mathbb{R}^{n_C}$ ) or discrete  
 181 ( $C \in \{0, 1\}^{n_C}$ ). Sidechannels also vary in form, with examples including embeddings ( $Z \in \mathbb{R}^{|Z|}$ )  
 182 (Mahinpei et al., 2021) and a one-hot mask ( $Z \in \{0, 1\}^{|Z|}$ ) (Debot et al., 2024).  
 183

184 Without loss of generality, the meta-model can be expressed as  
 185 a probabilistic graphical model (PGM) (Figure 1). The PGM  
 186 factorizes the joint distribution as:

$$p(y, c, z, x) = p(x) \cdot p(c|x) \cdot p(z|x) \cdot p(y|c, z) \quad (1)$$

187 with general task inference defined as  
 188

$$p(y|x) = \sum_{c, z} p(c|x) \cdot p(z|x) \cdot p(y|c, z) \quad (2)$$

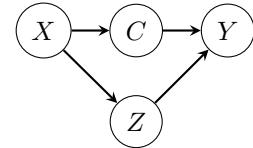


Figure 1: CSM Meta-model

190 Here, the summation is taken over all possible assignments of the concepts  $c$  and the sidechannel  $z$ .  
 191 For continuous variables, the summations are replaced with integrals.  
 192

193 Specific CSMs then correspond to particular parametrizations of this PGM. Even models without  
 194 explicit probabilistic semantics can be cast into this framework (such as Concept Embedding Models  
 195 (Espinosa Zarlenga et al., 2022)).<sup>3</sup> Two examples that we will work out in detail are CRM and CMR.  
 196

197 **Example 3.2** (Concept Residual Model (CRM)) (Mahinpei et al., 2021)). In CRM, concepts are  
 198 represented as delta distributions:  $p(c|x) = \delta(\hat{c} - c)$  where  $\hat{c} = g(x)$  with  $g$  a neural network.  
 199 The sidechannel  $z$  is an embedding also represented by a delta distribution, i.e.  $p(z|x) = \delta(\hat{z} - z)$   
 200 where  $\hat{z} = f(x)$  with  $f$  a neural network. The task predictor  $p(y|c, z)$  is a neural network. Due to  
 201 the sifting property of the Dirac delta, the inference expression simplifies to a single evaluation of  
 202 the task predictor:  $p(y|x) = p(y|\hat{c}, \hat{z})$ . CRM is not functionally interpretable due to the neural task  
 203 predictor.  
 204

205 **Example 3.3** (Concept Memory Reasoner (CMR)) (Debot et al., 2024)). In CMR, concepts are  
 206 modelled as Bernoulli random variables:  $p(c|x) = g(x)$  with  $g$  a neural network with sigmoid  
 207 activation. The sidechannel  $z$  is a categorical distribution representing a one-hot mask:  $p(z|x) =$   
 208  $f(x)$  with  $f$  a neural network with softmax activation. The task predictor  $p(y|c, z)$  possesses a set  
 209 of learned logic rules. At inference time, it evaluates only the rule indicated by the sidechannel  
 210  $z$  on the concepts. The resulting inference expression is:  $p(y|x) = \sum_{c \in \{0, 1\}^{n_C}} \sum_{z=1}^{n_R} p(c|x) \cdot$   
 211  $p(z|x) \cdot p(y|c, z)$ , where  $n_R$  is a hyperparameter signalling CMR's number of learned rules. CMR  
 212 is functionally interpretable due to its rule-based task predictor.

213 In Appendix A, we explain additionally for the following CSMs how they parametrize our meta-  
 214 model: Concept Embedding Models (Espinosa Zarlenga et al., 2022), Concept Bottleneck Models

215 <sup>3</sup>Concepts and sidechannels can be represented as delta distributions, so that the summation in Equation 2  
 216 reduces to a single evaluation due to the Dirac delta's sifting property.

216 with Additional Unsupervised Concepts (Sawada & Nakamura, 2022), Hybrid Post-Hoc Concept  
 217 Bottleneck Models (Yuksekgonul et al., 2022), and Deep Concept Reasoner (Barbiero et al., 2023).  
 218

219 **3.3 DEFAULT AND BOTTLENECK MODES FOR CSMs**  
 220

221 Consider a CSM where the sidechannel  $z$  is replaced with a value independent of the input  $x$ . In  
 222 this case, the task prediction depends only on  $x$  through the concepts, i.e.  $p(y|c, x) = p(y|c)$ . This  
 223 effectively disables the sidechannel, ensuring that predictions are fully representation interpretable.  
 224 This can be achieved by replacing the distribution  $p(z|x)$  with a distribution  $p(z)$ .

225 We can therefore define two operation modes of a CSM:

226 **Default Mode:** 
$$p_{\theta, \phi, \psi}(y|x) = \sum_{c, z} p_{\theta}(c|x) \cdot \boxed{p_{\phi}(z|x)} \cdot p_{\psi}(y|c, z) \quad (3)$$
  
 227

228 **Bottleneck Mode:** 
$$p_{\theta, \gamma, \psi}(\bar{y}|x) = \sum_{c, z} p_{\theta}(c|x) \cdot \boxed{p_{\gamma}(z)} \cdot p_{\psi}(y|c, z) \quad (4)$$
  
 229

230 where the subscripts  $\{\theta, \phi, \psi, \gamma\}$  indicate learnable parameters. Notice that the two modes differ  
 231 *exclusively* on whether their predictions use  $p_{\phi}(z|x)$  or  $p_{\gamma}(z)$ . We can always make predictions  
 232 using either mode, with bottleneck mode guaranteeing full representation interpretability. Note that  
 233 default mode corresponds to the standard inference of CSMs (Equation 2).  
 234

235 A natural question is how to obtain  $p(z)$ . A simple approach is to marginalize over the input dis-  
 236 tribution and approximate this using the training dataset  $\mathcal{D}$ , since the true data distribution  $p(X)$   
 237 is typically unknown:  $p_{\gamma}(z) = \sum_x p(x) \cdot p_{\phi}(z|x) \approx \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} p_{\phi}(z|x)$ . Here, bottleneck mode  
 238 has the same parameters as default mode (so  $\gamma \equiv \phi$ ). We will discuss an alternative approach in  
 239 Appendix B where instead a prior  $p(z)$  is explicitly learned.  
 240

241 **3.4 MEASURING AND OPTIMIZING FOR REPRESENTATION INTERPRETABILITY**  
 242

243 The definition of the two modes defined in the previous section allows use to devise a metric to  
 244 measure how often a CSM relies on its sidechannel by checking the agreement between predictions  
 245 from default and bottleneck mode.

246 **Definition 3.3** (Sidechannel Independence Score - SIS). *Let  $y_x$  and  $\bar{y}_x$  denote predictions from  
 247 default and bottleneck mode, respectively, obtained by thresholding the corresponding distributions  
 248  $p(y|x)$  and  $p(\bar{y}|x)$  at a given threshold probability. We define the Sidechannel Independence Score  
 249 (SIS) as:*

250 
$$SIS = \mathbb{E}_{x \sim p(X)} [\mathbb{1}[y_x = \bar{y}_x]]$$

251 The SIS measures the frequency with which the model’s prediction changes when the sidechannel is  
 252 removed, i.e. how dependent it is on the sidechannel rather than concepts. In practice, the SIS cannot  
 253 be computed exactly since the true data distribution  $p(X)$  is unknown. Instead, we approximate it  
 254 empirically using a dataset  $\mathcal{D}$ :  $\widehat{SIS} = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \mathbb{1}[y_x = \bar{y}_x]$ . Assuming the dataset is i.i.d.,  
 255 we can use Hoeffding’s inequality to provide formal guarantees about the model’s SIS (and thus  
 256 representation interpretability) on unseen data. In particular,  $p(|\widehat{SIS} - SIS| \geq \epsilon) \leq 2e^{-2|\mathcal{D}|\epsilon^2}$ .  
 257 For instance, if we find an empirical  $\widehat{SIS} = 60\%$  on a test set with size  $|\mathcal{D}| = 1000$ , then the  
 258 95% confidence interval of  $SIS$  is  $[58\%, 62\%]$ . This metric is pragmatic in the sense that a human  
 259 can easily interpret it and decide on whether they consider the model representation-interpretable  
 260 enough to trust it.  
 261

262 Most CSMs are currently trained to maximize accuracy (see Section 4). Our meta-model enables  
 263 explicit optimization for representation interpretability by introducing a loss term that penalizes  
 264 discrepancies between predictions from default mode ( $p(y|x)$ ) and bottleneck mode ( $p(\bar{y}|x)$ ). Any  
 265 suitable divergence, such as total variation distance or symmetric Kullback-Leibler divergence, can  
 266 be used. This approach integrates seamlessly into any CSM. For example, when maximizing the  
 267 likelihood of training data, the objective becomes:

268 
$$\arg \max_{\phi, \psi, \gamma, \theta} \left[ \sum_{(x, c, y) \in \mathcal{D}} (\log p_{\phi, \psi}(y|c, x) + \alpha \cdot \log p_{\theta}(c|x) - \beta \cdot \text{DIV}(p_{\phi, \psi}(y|c, x) || p_{\gamma, \psi}(\bar{y}|c, x))) \right]$$
  
 269

270 where  $\alpha$  and  $\beta$  are hyperparameters, and  $\text{DIV}(\cdot||\cdot)$  denotes a chosen divergence measure. We refer  
 271 to this additional term as **SIS regularization**. It can also be incorporated into alternative training  
 272 schemes, such as sequential or joint CBM training (Koh et al., 2020).

273 Furthermore, instead of marginalizing over the input  $x$  to obtain a prior  $p(z)$ , one can introduce  
 274 a learnable prior. This reduces computational cost, simplifies optimization, and in some cases in-  
 275 creases the expressivity of the model in bottleneck mode. For more details, we refer to Appendix  
 276 B. Importantly, this method is applicable to any CSM, regardless of its parametrization within our  
 277 meta-model PGM (e.g. choice of sidechannel  $z$  or form of  $p(y|c, z)$ ).

### 279 3.5 EXPRESSIVITY OF CSMs IN BOTTLENECK MODE

281 In default mode, CSMs are typically universal classifiers: they are as expressive as neural networks  
 282 for classification. This universality comes from the sidechannel, which allows them to learn any  
 283 mapping from input to task ( $X \rightarrow Y$ ), irrespective of the employed concept set.

284 When in bottleneck mode, however, their ability to learn input-to-task mappings is limited by the  
 285 quality of the concept set. Because the concept set is dependent on the dataset, a general comparison  
 286 between CSMs cannot be made. Instead, we can analyze their expressivity in learning mappings  
 287 from concepts to task ( $C \rightarrow Y$ ). Some CSMs are more expressive than others in this regard, for  
 288 instance:

- 290 • CRM (Mahinpei et al., 2021) applies a neural network to the concepts. This is fully expres-  
 291 sive but not functionally interpretable.
- 292 • CBM-AUC (Sawada & Nakamura, 2022) applies a linear layer to the concepts, which is  
 293 functionally interpretable but not expressive.
- 294 • CMR (Debot et al., 2024) falls in between, as it applies a set of learned logic rules to the  
 295 concepts. Its expressivity depends on the number of learned rules, and increases with this  
 296 capacity.

298 For more examples, we refer to Appendix A.

300 This difference in  $C \rightarrow Y$  expressivity between default and bottleneck mode for some CSMs implies  
 301 that they do not only use the sidepath to use information related to  $Y$  that cannot be found in  $C$  but  
 302 also to improve their  $C \rightarrow Y$  expressivity.

303 This yields two important insights. First, for achieving the same accuracy on expressive tasks  $Y$ ,  
 304 a functionally interpretable but inexpressive CSM (e.g. CBM-AUC) must rely more heavily on its  
 305 sidechannel than a non-functionally interpretable but expressive CSM (e.g. CRM). Thus, the former  
 306 may achieve higher functional interpretability but lower representation interpretability than the lat-  
 307 ter. Second, even if the concepts are sufficient for the task (i.e. a perfect predictor exists given a  
 308 sufficiently expressive model), a functionally interpretable but inexpressive CSM will still need to  
 309 use its sidechannel to achieve high accuracy. We show this empirically in Section 5, and illustrate it  
 310 with the following simple example.

311 **Example 3.4.** Consider two binary concepts  $c_1$  and  $c_2$  and a task defined as their logical XOR  
 312 ( $y := c_1 \oplus c_2$ ). Suppose a CSM uses a linear layer to map the concepts  $c$  and a sidechannel  $z$  (some  
 313 neurons) to the task  $y$ . In bottleneck mode, the CSM cannot predict  $y$  accurately, since it is not  
 314 linearly separable. In default mode, the model can encode features such as  $c_1 \wedge \neg c_2$  and  $\neg c_1 \wedge c_2$   
 315 into the sidechannel. The linear layer can then compute the logical OR of these neurons, thereby  
 316 solving the XOR task. In this way, the sidechannel effectively extends the concept bottleneck with  
 317 combinations of concepts to increase the model’s expressivity.

## 318 4 RELATED WORK

321 While explicitly measuring and optimizing representation interpretability has not been studied di-  
 322 rectly, several related research directions pursue overlapping but distinct goals. Most of these works  
 323 focus on encouraging models to use concepts extensively, whereas our focus lies in minimizing  
 324 reliance on the sidechannel. Note that the latter entails the former, but goes even further.

324 A major line of research in CBMs emphasizes *intervenability* (Espinosa Zarlenga et al., 2022; 2023;  
 325 Havasi et al., 2022): when concept predictions are replaced with their ground truth, downstream task  
 326 accuracy should improve as much as possible. Such interventions are designed to simulate human  
 327 expert interaction at decision time. However, intervenability is influenced by many factors, for in-  
 328 stance: (1) the extent to which concepts are used, (2) whether interventions remain in-distribution,  
 329 and (3) architectural design choices. Of these, only (1) partially relates to representation inter-  
 330 pretability. Importantly, a model may achieve high intervenability by heavily relying on concepts  
 331 while still encoding substantial task-relevant information in the sidechannel. In this case, interven-  
 332 ability remains high even though representation interpretability may be low.

333 Some works attempt to *disentangle* the sidechannel from the concepts (Zabounidis et al., 2023):  
 334 make it complementary to the concepts, not re-encode the same information. This approach im-  
 335 proves representation interpretability to some degree by preventing concept-related information  
 336 from being in the sidechannel, but does not reduce sidechannel usage beyond this. Moreover, enfor-  
 337 cing disentanglement substantially harms accuracy in CSMs that rely on sidechannels for expressivity  
 338 (Barbiero et al., 2023; Debot et al., 2024; Sawada & Nakamura, 2022), even though disentanglement  
 339 is not strictly necessary for achieving representation interpretability (see Appendix C).

340 Other works introduce methods to maximize *concept utilization* during prediction (Kalampalikis  
 341 et al., 2025; Shang et al., 2024). These approaches encourage reliance on concepts but do not  
 342 directly address sidechannel reduction (see Section 5). Furthermore, they are tailored to CSMs with  
 343 embedding-based sidechannels, leaving open how they might generalize to alternative architectures.  
 344 Finally, they may require structural modifications to the model (e.g. a factorizable task predictor  
 345 (Shang et al., 2024)), whereas our approach does not require this. **Similarly, Zhang et al. (2024)**  
 346 introduce an embedding-based CSM with a factorized task predictor and a regularization approach,  
 347 which *does* encourage sidechannel reduction, but is specific to their task predictor.

348 Finally, Havasi et al. (2022) propose a metric known as the *completeness score*, which estimates  
 349 how fully concepts capture predictive information for the task. This is computed by learning a  
 350 distribution  $q(z|c)$  after training and relies on mutual information. While informative, it suffers  
 351 from two limitations. First, mutual-information-based quantities are less intuitive for humans than  
 352 accuracy-based metrics like SIS. Second, it is not a measure for representation interpretability, as  
 353 it may be high even though representation interpretability is poor (see Appendix C for an in-depth  
 354 discussion and a comparison between their  $q(z|c)$  and our  $p(z)$ ). **Zhang et al. (2024) also introduce**  
 355 **a metric for interpretability, but this can be only be computed for factorizable task predictors and**  
 356 **similarly does not measure representation interpretability.**

## 357 5 EXPERIMENTS

### 360 5.1 SETUP

362 This section provides essential information about the experiments (see Appendix D for details). We  
 363 report results averaged over three seeds, and we give standard-deviations as shaded areas. For pareto  
 364 curves, we only give the results for the first seed (see Appendix D for the remaining ones).

365 **Datasets.** We use **three** standard concept-based datasets: CelebA (Liu et al., 2015), with 200k  
 366 celebrity faces annotated with facial attribute concepts (e.g. *blond hair, beard*); **CUB** (Welinder  
 367 et al., 2010), **where the task is to classify birds**; and MNIST-Addition (Manhaeve et al., 2018),  
 368 where the task is to predict the sum of two digit images. CelebA’s concepts are insufficient for the  
 369 task; MNIST-Addition’s are sufficient but has an expressive task. **CUB’s concepts are sufficient,**  
 370 **has an inexpressive task, but has difficult concept prediction. For the results on CUB, we refer to**  
 371 **Appendix D.2.**

372 **Models.** We use the following CSMs in our experiments: Concept Residual Models (CRM) (Mahin-  
 373 pei et al., 2021), Concept Embedding Models (CEM) (Espinosa Zarlenga et al., 2022), Deep Concept  
 374 Reasoner (DCR) (Barbiero et al., 2023), and Concept Memory Reasoner (CMR) (Debot et al., 2024).  
 375 We also define Linear Residual Model (LRM) as CRM but with a linear layer as task predictor. DCR,  
 376 CMR, and LRM are *functionally interpretable* but *inexpressive* in bottleneck mode (see Section 3.5);  
 377 CRM and CEM are *not functionally interpretable* but *expressive* in bottleneck mode. As **baselines**,  
 we consider approaches that maximize concept usage (Shang et al., 2024; Kalampalikis et al., 2025).

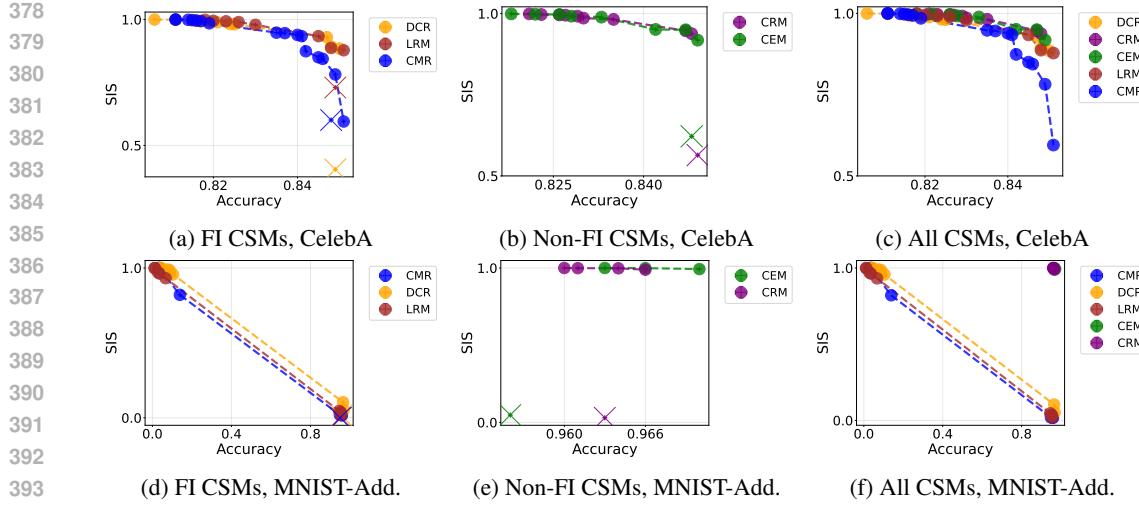


Figure 2: Accuracy vs representation interpretability trade-off in CSMs. Each point is a hyperparameter configuration, keeping only pareto-efficient points. Crosses are the most accurate configuration trained without SIS regularization (not in (c) and (f)). FI means “functionally interpretable”.

## 5.2 RESULTS AND DISCUSSION

We consider the following research questions: **(Interpretability)** How representation interpretable are state-of-the-art CSMs? How much does SIS regularization improve their representation interpretability? What is the trade-off w.r.t. accuracy? How do current methods encouraging concept usage for CRM-like models (Section 4) fare regarding representation interpretability? **(Side-effects)** Does SIS regularization improve intervenability? Does SIS regularization improve the quality of learned interpretable task predictors?

**Current optimization of CSMs results in uninterpretable models (Figure 2, crosses), contrary to when using SIS regularization (Figure 2, circles).** Models that are optimized only for accuracy have a low SIS score, meaning they significantly use the sidechannel. Notably, this is even the case when the sidechannel is completely unnecessary, e.g. when training expressive CSMs on datasets where the concepts are sufficient for the task (Figure 2e). Conversely, when using our SIS regularization, CSMs become significantly more representation interpretable, and the human can choose how much accuracy to trade for how much interpretability.

**When using the sidechannel is unnecessary, SIS regularization ensures CSMs avoid it (Figure 2e).** With sufficient concept sets, expressive CSMs can reach high accuracy without relying on the sidechannel, effectively functioning as concept bottleneck models with the same interpretability.

**Inexpressive CSMs require the sidechannel for expressive tasks (Figure 2d).** Despite MNIST-Add having sufficient concepts, the linear expressivity of these CSMs’ bottleneck mode makes them either very accurate but completely not representation interpretable, or completely representation interpretable but very inaccurate. Interestingly, our analysis shows that CMR seems unable to exploit its non-linear expressivity. In Appendix D, we address this by replacing its rule learner with an existing alternative, enabling near-perfect accuracy and interpretability on this task.

**Existing concept usage approaches yield smaller gains in representation interpretability (Figure 3).** While approaches that encourage higher concept usage for CRM-like CSMs also increase representation interpretability, their improvements are smaller than those achieved with SIS regular-

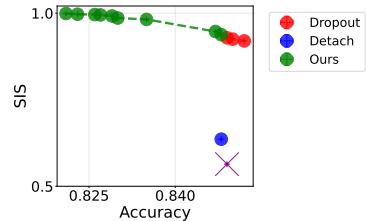


Figure 3: Accuracy vs SIS for CRM on CelebA, comparing with concept usage approaches. The cross is the most accurate CRM trained without any approach.

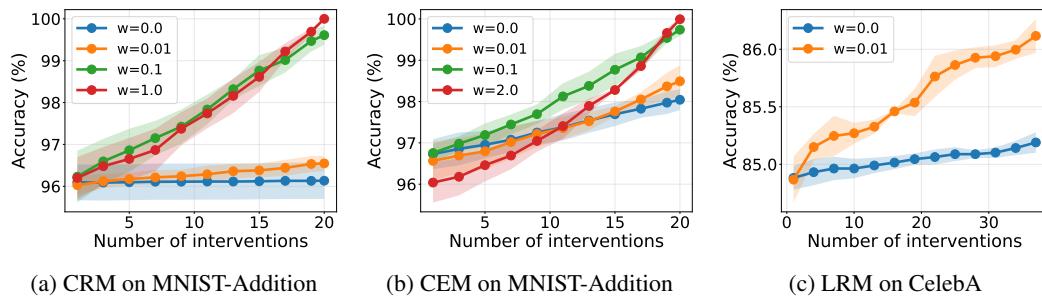
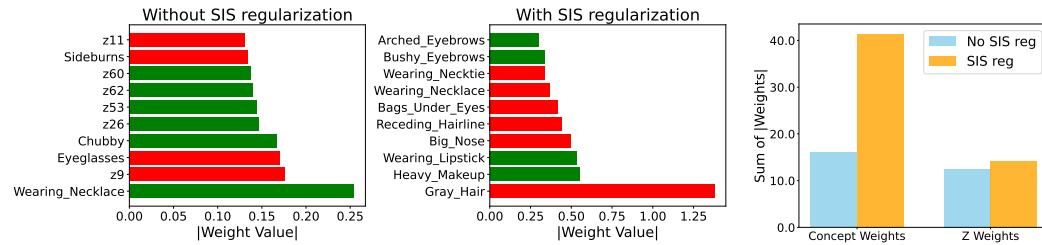


Figure 4: Intervenability in CSMs with ( $w > 0$ ) and without ( $w = 0$ ) SIS regularization for different regularization weights  $w$ . The y-axis denotes accuracy after intervening on a number of concepts denoted by the x-axis.



(a) Top 10 largest weights. Bars represent absolute weight magnitudes, with green and red indicating positive and negative weights, respectively. Weights corresponding to the sidechannel’s neurons are denoted with a ‘z’. (b) Total absolute weight magnitudes assigned to concepts and sidechannel neurons.

Figure 5: Inspection of LRM’s linear layer for predicting the task ‘Young’ for CelebA, comparing LRM’s trained with and without SIS regularization.

ization. We observe that *dropout* (Kalampalikis et al., 2025) enhances SIS only insofar as accuracy is not compromised, and that *detach* (Shang et al., 2024) only slightly improves SIS.

**SIS regularization increases intervenability (Figure 4).** As SIS regularization reduces sidechannel reliance in the CSM, an automatic side-effect and advantage is that the CSM relies more on the concepts to predict the task, making it more responsive to concept interventions. For more curves, see Appendix D.

**SIS regularization improves the task predictor quality of functionally interpretable CSMs (Figure 5).** We compare the most accurate LRM configurations trained with and without regularization on CelebA (84.75% vs 84.94% accuracy). Without SIS regularization, many of the most contributing weights of LRM’s linear task predictor are uninterpretable latent factors from the sidechannel (e.g.  $z11$ ), harming interpretability. In contrast, with SIS regularization, large weights correspond to semantically meaningful concepts (e.g. *Gray Hair*), while reliance on the sidechannel is suppressed (Figure 5a), and the total weight magnitude shifts significantly toward concepts rather than the sidechannel (Figure 5b).

## 6 CONCLUSION

We addressed a fundamental gap in concept-based models: the lack of principled methods to measure and control the trade-off between accuracy and interpretability in concept sidechannel models (CSMs). We proposed a unified probabilistic meta-model that places existing CSMs within a single framework, enabling us to disconnect representation interpretability from functional interpretability. Building on this, we introduced the Sidechannel Independence Score (SIS) as a metric that quantifies a model’s representation interpretability. We demonstrated how SIS can serve as a regularization objective, allowing human users to explicitly control the extent to which models rely on uninterpretable sidechannels. Our experiments reveal that state-of-the-art CSMs, when trained using

486 typical objectives, are not genuinely interpretable, and that SIS regularization produces models that  
 487 are more representation interpretable, more responsive to interventions and more transparent in their  
 488 task predictors. Our analysis also allowed us to find a weakness in a state-of-the-art CSM (CMR),  
 489 being unable to exploit its theoretical expressivity, which we addressed. Our contributions provide  
 490 both theoretical foundations and practical tools for developing interpretable CSMs.

491 **Limitations and future work** Future work could investigate the accuracy-interpretability trade-off  
 492 among datasets beyond vision (e.g. language) and more CSMs, and investigate the effect of SIS  
 493 regularization on the quality of learned rules for rule-based CSMs.

494 **Reproducibility statement.** All our experiments are seeded, and we will make the code publicly  
 495 available upon publication of the paper. Moreover, in Appendix D, we describe in detail the setup  
 496 of each experiment, the implementation of each model, and the training setup.

## 498 REFERENCES

500 Pietro Barbiero, Gabriele Ciravegna, Francesco Giannini, Mateo Espinosa Zarlenga, Lucie Charlotte  
 501 Magister, Alberto Tonda, Pietro Lio', Frederic Precioso, Mateja Jamnik, and Giuseppe Marra.  
 502 Interpretable neural-symbolic concept reasoning. In *ICML*, 2023.

503 Pietro Barbiero, Giuseppe Marra, Gabriele Ciravegna, David Debot, Francesco De Santis, Michelangelo  
 504 Diligenti, Mateo Espinosa Zarlenga, and Francesco Giannini. Neural interpretable reasoning.  
 505 *arXiv preprint arXiv:2502.11639*, 2025.

506 David Debot, Pietro Barbiero, Francesco Giannini, Gabriele Ciravegna, Michelangelo Diligenti, and  
 507 Giuseppe Marra. Interpretable concept-based memory reasoning. *Advances of neural information  
 508 processing systems 37, NeurIPS 2024*, 2024.

509 Mateo Espinosa Zarlenga, Pietro Barbiero, Gabriele Ciravegna, Giuseppe Marra, Francesco Giannini,  
 510 Michelangelo Diligenti, Zohreh Shams, Frederic Precioso, Stefano Melacci, Adrian Weller,  
 511 Pietro Lio, and Mateja Jamnik. Concept embedding models: Beyond the accuracy-explainability  
 512 trade-off. *Advances in Neural Information Processing Systems*, 35, 2022.

513 Mateo Espinosa Zarlenga, Katie Collins, Krishnamurthy Dvijotham, Adrian Weller, Zohreh Shams,  
 514 and Mateja Jamnik. Learning to receive help: Intervention-aware concept embedding models.  
 515 *Advances in Neural Information Processing Systems*, 36:37849–37875, 2023.

516 Marton Havasi, Sonali Parbhoo, and Finale Doshi-Velez. Addressing leakage in concept bottleneck  
 517 models. *Advances in Neural Information Processing Systems*, 35:23386–23397, 2022.

518 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recogni-  
 519 tion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
 520 770–778, 2016.

521 Nektarios Kalampalikis, Kavya Gupta, Georgi Vitanov, and Isabel Valera. Towards reasonable con-  
 522 cept bottleneck models. *arXiv preprint arXiv:2506.05014*, 2025.

523 Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and  
 524 Percy Liang. Concept bottleneck models. In *International conference on machine learning*, pp.  
 525 5338–5348. PMLR, 2020.

526 Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild.  
 527 In *ICCV*, 2015.

528 Anita Mahinpei, Justin Clark, Isaac Lage, Finale Doshi-Velez, and Weiwei Pan. Promises and  
 529 pitfalls of black-box concept learning models. *arXiv preprint arXiv:2106.13314*, 2021.

530 Robin Manhaeve, Sebastian Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt.  
 531 DeepProbLog: Neural Probabilistic Logic Programming. In *NeurIPS*, pp. 3753–3763, 2018.

532 Emanuele Marconato, Andrea Passerini, and Stefano Teso. Glancenets: Interpretable, leak-proof  
 533 concept-based models. *Advances in Neural Information Processing Systems*, 35:21212–21227,  
 534 2022.

540 Tuomas Oikarinen, Subhro Das, Lam Nguyen, and Lily Weng. Label-free concept bottleneck mod-  
 541 els. In *International Conference on Learning Representations*, 2023.

542

543 Eleonora Poeta, Gabriele Ciravegna, Eliana Pastor, Tania Cerquitelli, and Elena Baralis. Concept-  
 544 based explainable artificial intelligence: A survey. *arXiv preprint arXiv:2312.12936*, 2023.

545 Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. In-  
 546 terpretable machine learning: Fundamental principles and 10 grand challenges. *Statistic Surveys*,  
 547 16:1–85, 2022.

548

549 Yoshihide Sawada and Keigo Nakamura. Concept bottleneck model with additional unsupervised  
 550 concepts. *IEEE Access*, 10:41758–41765, 2022.

551

552 Chenming Shang, Shiji Zhou, Hengyuan Zhang, Xinzhe Ni, Yujiu Yang, and Yuwang Wang. In-  
 553 cremental residual concept bottleneck models. In *Proceedings of the IEEE/CVF Conference on*  
*553 Computer Vision and Pattern Recognition*, pp. 11030–11040, 2024.

554

555 Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Be-  
 556 longie, and Pietro Perona. Caltech-ucsd birds 200. Technical Report CNS-TR-201, Caltech,  
 557 2010. URL [/se3/wp-content/uploads/2014/09/WelinderEtal10\\_CUB-200.pdf](https://se3/wp-content/uploads/2014/09/WelinderEtal10_CUB-200.pdf), <http://www.vision.caltech.edu/visipedia/CUB-200.html>.

558

559 Mert Yuksekgonul, Maggie Wang, and James Zou. Post-hoc concept bottleneck models. *arXiv*  
 560 *preprint arXiv:2205.15480*, 2022.

561

562 Renos Zabounidis, Ini Oguntola, Konghao Zhao, Joseph Campbell, Simon Stepputtis, and Katia  
 563 Sycara. Benchmarking and enhancing disentanglement in concept-residual models. *arXiv preprint*  
*563 arXiv:2312.00192*, 2023.

564

565 Rui Zhang, Xingbo Du, Junchi Yan, and Shihua Zhang. The decoupling concept bottleneck model.  
 566 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

# 594 595 596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612 613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646 647 Supplementary Material

## TABLE OF CONTENTS

<b>A Overview of state-of-the-art CSMs</b>	<b>13</b>
<b>B Learnable prior <math>p(z)</math></b>	<b>14</b>
<b>C Limits of Disentanglement and Concept Completeness Score</b>	<b>14</b>
C.1 Disentangling for representation interpretability . . . . .	14
C.2 Concept Completeness Score vs Sidechannel Independence Score . . . . .	14
<b>D Experiments</b>	<b>15</b>
D.1 Experimental details . . . . .	15
D.2 Additional results . . . . .	18
<b>E LLM usage declaration</b>	<b>20</b>
<b>F Code and licenses</b>	<b>20</b>

648 A OVERVIEW OF STATE-OF-THE-ART CSMs  
649  
650  
651  
652653 *In this section, we give more examples of how state-of-the-art CSMs are parametrizations of our  
654 meta-model.*  
655656 We do not give extensive details about the models, only a high-level overview. We also take the  
657 examples CRM and CMR again as they appear in the main text.658 **Concept Residual Model (CRM) (Mahinpei et al., 2021).** In CRM, concepts are represented as  
659 delta distributions:  $p(c|x) = \delta(\hat{c} - c)$  where  $\hat{c} = g(x)$  with  $g$  a neural network. The sidechannel  
660  $z$  is an embedding also represented by a delta distribution, i.e.  $p(z|x) = \delta(\hat{z} - z)$  where  $\hat{z} = f(x)$   
661 with  $f$  a neural network. The task predictor  $p(y|c, z)$  is a neural network. Due to the sifting property  
662 of the Dirac delta, the inference expression simplifies to a single evaluation of the task predictor:  
663  $p(y|x) = p(y|\hat{c}, \hat{z})$ . CRM is not functionally interpretable due to the neural task predictor.664 **Concept Embedding Models (CEM) (Espinosa Zarlenga et al., 2022).** In CEM, concepts are  
665 represented as delta distributions:  $p(c|x) = \delta(\hat{c} - c)$  where  $\hat{c} = g(x)$  with  $g$  a neural network. The  
666 sidechannel  $z$  consists of 2 embeddings per concept, each represented by a delta distribution, i.e.  
667 for  $i \in \{1..n_C\}$  and  $j \in \{1, 2\}$ ,  $p(z_{i,j}|x) = \delta(\hat{z}_{i,j} - z_{i,j})$  where  $\hat{z}_{i,j} = f_{ij}(x)$  with  $f_{ij}$  a neural  
668 network. The task predictor  $p(y|c, z)$  first mixes the two embeddings of each concept using the  
669 concept scores ( $z_i = c_i \cdot z_{i,1} + (1 - c_i) \cdot z_{i,2}$ ), and the resulting embeddings are concatenated and  
670 fed through a neural network. Due to the sifting property of the Dirac delta, the inference expression  
671 simplifies to a single evaluation of the task predictor:  $p(y|x) = p(y|\hat{c}, \hat{z})$ . CEM is not functionally  
672 interpretable due to the neural task predictor.673 **Concept Memory Reasoner (CMR) (Debot et al., 2024).** In CMR, concepts are modelled as  
674 Bernoulli random variables:  $p(c|x) = g(x)$  with  $g$  a neural network with sigmoid activation. The  
675 sidechannel  $z$  is a categorical distribution representing a one-hot mask:  $p(z|x) = f(x)$  with  $f$  a neu-  
676 ral network with softmax activation. The task predictor  $p(y|c, z)$  encodes a set of learned logic rules.  
677 At inference time, it evaluates only the rule indicated by the sidechannel  $z$  on the concepts. The re-  
678 sulting inference expression is:  $p(y|x) = \sum_{c \in \{0,1\}^{n_C}} \sum_{z=1}^{n_R} p(c|x) \cdot p(z|x) \cdot p(y|c, z)$ , where  $n_R$  is  
679 a hyperparameter signalling CMR’s number of learned rules. CMR is functionally interpretable due  
680 to the rule-based task predictor.681 **Deep Concept Reasoner (DCR) (Barbiero et al., 2023).** In DCR, concepts are modelled as delta  
682 distributions:  $p(c|x) = \delta(\hat{c} - c)$  where  $\hat{c} = g(x)$  with  $g$  a neural network. The sidechannel  $z$   
683 represents a fuzzy logic rule, given by 2 fuzzy values per concept, signalling the relevance and  
684 polarity of each concept in the rule: for  $i \in \{1..n_C\}$  and  $j \in \{1, 2\}$ ,  $p(z_{i,j}|x) = \delta(\hat{z}_{i,j} - z_{i,j})$   
685 where  $\hat{z}_{i,j} = f_{ij}(x)$  with  $f_{ij}$  a neural network with sigmoid activation. The task predictor  $p(y|c, z)$   
686 uses fuzzy logical inference to deduce the task from the fuzzy logic rule and the predicted concepts.  
687 DCR is functionally interpretable due to the rule-based task predictor.688 **Concept Bottleneck Model with Additional Unsupervised Concepts (CBM-AUC) (Sawada &  
689 Nakamura, 2022).** In CBM-AUC, concepts are also delta distributions. The sidechannel consists  
690 of two components: some unsupervised concepts and a weight for each concept and unsupervised  
691 concept. Both are represented by delta distributions. The task predictor computes the dot product of  
692 all concepts and weights. CBM-AUC is functionally interpretable due to the linear task predictor.693 **Hybrid Post-hoc Concept Bottleneck Models (PCBM-h) (Yuksekgonul et al., 2022).** In PCBM-  
694 h, concepts are delta distributions, and the sidechannel is a single value denoted by a delta distri-  
695 bution. The task predictor applies some interpretable predictor (e.g. linear layer) to the concepts,  
696 providing a score  $y_1$ , which is summed with the value predicted by the sidechannel. PCBM-h is  
697 functionally interpretable due to the interpretable predictor followed by a simple sum.698 **Expressivity.** Of these models, CRM and CEM are universal classifiers ( $C \rightarrow Y$ ) in bottleneck  
699 mode, but not functionally interpretable. DCR and CBM-AUC are functionally interpretable, but  
700 have a linear expressivity ( $C \rightarrow Y$ ) in bottleneck mode. PCBM-h is functionally interpretable but  
701 its expressivity depends on the used predictor. CMR is functionally interpretable and its expressivity  
in bottleneck mode is limited only by the number of rules it uses.

702 **B LEARNABLE PRIOR  $p(z)$**   
703

704 When using SIS regularization, we propose to use a learnable prior for  $p(z)$  as opposed to a marginal-  
705 ization approach. For some models, avoiding the marginalization can improve expressivity. For  
706 instance, for CMR, using the marginalization approach means that the final task prediction is still  
707 using a single logic rule, with uncertainty modelled over which rule to use. As a consequence,  
708 CMR’s expressivity ( $C \rightarrow Y$ ) would still be linear. By instead allowing CMR to use the entire  
709 set of learned rules (which would be impossible by the marginalization approach, as this value for  
710  $z$  is out-of-distribution), CMR’s expressivity becomes bounded only by the number of rules it has  
711 learned.

712 **This ensures the SIS regularization is computationally quite cheap, this ensures that the SIS regu-  
713 larization requires only a single forward pass of the model and computing a KL divergence between  
714 the original forward pass and this additional one.**

716 **C LIMITS OF DISENTANGLEMENT AND CONCEPT COMPLETENESS SCORE**  
717718 **C.1 DISENTANGLING FOR REPRESENTATION INTERPRETABILITY**  
719

720 *In this section, we argue that disentangling the sidechannel from the concepts is neither a sufficient,  
721 nor necessary condition for achieving a representation-interpretable CSM. Moreover, it may even  
722 needlessly harm the CSM’s accuracy. This highlights that representation interpretability (and SIS)  
723 capture(s) a distinct notion of interpretability not already addressed by disentanglement.*

724 Zabounidis et al. (2023) propose disentangling the sidechannel from the concepts. This im-  
725 proves representation interpretability by preventing concept-related information from leaking into  
726 the sidechannel. However, as noted in Section 4, this is not sufficient: task-relevant information that  
727 is unrelated to the concepts is still free to appear in the sidechannel.

728 Importantly, disentanglement is not strictly necessary for a CSM to be representation interpretable.  
729 The simplest example is when the task predictor relies only on the concepts for each prediction. In  
730 this case, the prediction is representation interpretable regardless of what information the sidechan-  
731 nel encodes.

732 A more interesting example is the following, where disentanglement is clearly unnecessary and even  
733 hurts the model’s accuracy. Suppose the dataset is such that the concepts are sufficient for solving  
734 the task. That is, they capture all task-relevant information. Consider CMR, a CSM where the  
735 sidechannel is a one-hot mask. The task predictor contains a set of logic rules. In default mode,  
736 the predictor uses the sidechannel to select a single rule to apply to the concepts; in bottleneck  
737 mode, it applies all rules simultaneously. For some inputs, the prediction made in default mode  
738 (via rule selection) may coincide with the prediction in bottleneck mode (via all rules), meaning  
739 the prediction is representation interpretable. However, if the sidechannel were disentangled from  
740 the concepts, rule selection could not exploit any task-relevant information (since all of it resides in  
741 the concepts). In that case, accuracy would collapse, as the rule selection would effectively become  
742 random.

743 **C.2 CONCEPT COMPLETENESS SCORE VS SIDECHANNEL INDEPENDENCE SCORE**  
744

745 *In this section, we argue that the concept completeness score (Havasi et al., 2022) does not capture  
746 representation interpretability, comparing it to our SIS metric.*

747 The concept completeness score (CCS)  $\tau$  is defined by Havasi et al. (2022) as

$$749 \tau = \frac{I(y; c)}{I(y; c, x)} \approx \frac{H(y) + \mathbb{E}_{(x, c) \sim \mathcal{D}}[\log \mathbb{E}_{z \sim q(z|c)}[p(y|c, z)]]}{H(y) + \mathbb{E}_{(x, c) \sim \mathcal{D}}[\log \mathbb{E}_{z \sim p(z|x)}[p(y|c, z)]]} \quad (5)$$

750 where  $H(\cdot)$  denotes entropy. This score is used to estimate how much task-relevant information  
751 is present in the concepts. Note that the numerator uses  $q(z|c)$ : CCS will be high whenever pre-  
752 dictions based on  $z$  derived from  $c$  resemble those based on  $z$  derived from  $x$ . Crucially, this can  
753 happen even if the task predictor heavily relies on the sidechannel, so long as all the information in  
754 the sidechannel is also present in the concepts. Intuitively: *the concept completeness score measures*

756 *how much information more than concepts is encoded in the sidechannel; SIS measures how much*  
 757 *the sidechannel is used.* CCS can be very high for a completely uninterpretable prediction, which is  
 758 not the case for SIS. Concept-related information encoded in an uninterpretable form (e.g. embed-  
 759 dings) is not interpretable, since humans cannot identify what it represents, despite CCS being high.  
 760 Furthermore, concepts are explicitly supervised to align with human-understandable interpretations,  
 761 whereas sidechannels are not. Thus, even if a sidechannel has a simple form (e.g. a logic rule, as in  
 762 DCR (Barbiero et al., 2023)) and is predicted from the same information as the concepts, we argue  
 763 it remains uninterpretable. Consider the following two examples.

764 Consider an example with DCR. Suppose the concepts capture all task-relevant information, and the  
 765 model uses five concepts  $c_1, c_2, \dots, c_5$ . For an input  $x$  with ground-truth label  $y = \text{True}$ , assume all  
 766 concepts are predicted as *True*, and the sidechannel outputs the rule  $y \leftarrow c_1 \wedge c_2 \wedge \dots \wedge c_5$ . DCR’s  
 767 task predictor applies this rule to the predicted concepts, correctly predicting  $y = \text{True}$ . For this  
 768 individual prediction,  $\tau$  would be high, since the rule is derived from concept-related information.  
 769 However, it is opaque to the human why the model used this particular rule. Why not  $y \leftarrow c_1 \wedge$   
 770  $c_2$ , which would yield the same correct label, or even  $y \leftarrow c_1$ , or  $y \leftarrow \text{True}$ ? In the last case,  
 771 the concepts are bypassed entirely, and the prediction is made solely by the sidechannel: a neural  
 772 network, hence uninterpretable. The core issue is that the sidechannel is not explicitly *aligned* with  
 773 human-understandable semantics, unlike the concepts. If it were, then it would be a concept.

774 Consider an example with LRM. Suppose the concepts capture all task-relevant information, and the  
 775 model uses some concepts and a sidechannel  $z$  that is a single neuron. These are passed to a linear  
 776 layer for the task prediction. If the linear layer has learned every weight to be zero except the weight  
 777 on the sidechannel, then the task prediction probability is determined entirely by the sidechannel,  
 778 i.e. by the underlying neural network that predicts  $z$ . In this case,  $\tau$  will be 1, since  $z$  is derived from  
 779 concept-related information. However, the interpretability the linear layer provides is completely  
 780 lost. A human examining the weights will only see that the task decision (completely) depends  
 781 on  $z$ , but has no way to understand how the individual concepts contribute. While the linear layer  
 782 provides functional interpretability, it is useless because the model is completely not representation  
 783 interpretable: the linear layer purely uses an uninterpretable representation.

784 Similarly, the example in Section 3.5 would also have a high completeness score.

785 **Using  $p(z)$  versus  $q(z | c)$ .** One could consider defining bottleneck mode using  $q(z | c)$  instead  
 786 of  $p(z)$ , where the former is parameterized by a neural network. This substitution would increase  
 787 expressivity: CSMs that currently act as linear classifiers ( $C \rightarrow Y$ ) in bottleneck mode under  $p(z)$   
 788 would become universal classifiers ( $C \rightarrow Y$ ) under  $q(z | c)$ , since  $q$  is parameterized by a neural  
 789 network. However, as we explained above, this added flexibility comes at the cost of interpretability.  
 790 With  $p(z)$ , the mapping from  $C \rightarrow Y$  in bottleneck mode is interpretable for functionally  
 791 interpretable CSMs, whereas with  $q(z | c)$  a neural network is put between concepts and task pre-  
 792 dictions, reducing interpretability.

## 793 D EXPERIMENTS

### 794 D.1 EXPERIMENTAL DETAILS

795 **Datasets.** In MNIST-Addition (Manhaeve et al., 2018), each input consists of two MNIST images  
 796 each representing some digit. The task is to predict the sum of the 2 digits. The concepts denote  
 797 which digit is in each image. In CelebA (Liu et al., 2015), the concepts are face attributes such  
 798 as *Blond Hair* and *Wears Make-up*. As tasks, we take the concepts *Male*, *Young* and *Attractive*,  
 799 dropping them from the concept set.

800 **Metrics.** For CelebA, we use regular accuracy and SIS as defined in the main text. For MNIST-  
 801 Addition, we use *subset* versions of these metrics due to the large imbalance and mutually exclusive  
 802 nature of both concepts and tasks.

803 **Seeds.** We use seeds 1, 2 and 3.

804 **Hard concepts.** To avoid the problem of concept leakage which harms interpretability (Marconato  
 805 et al., 2022), we employ hard concepts (as opposed to soft concepts) by thresholding the concept  
 806 predictions at 50% before passing them to the task predictor, which is common for CBMs.

810     **Training.** We use a training objective similar to sequential training. We maximize the likelihood  
 811 of our data (see Section 3.4), with as difference that we use the predicted concepts  $\hat{c}$  instead of  
 812 the ground truth labels for  $c$  in the task predictor  $p(y|c, z)$ . This makes the model more robust  
 813 and aware of mistakes in the concept prediction. Importantly, we do not allow the gradient from  $y$   
 814 to pass through  $c$  (so we avoid the joint training objective some CBMs employ), as this is known  
 815 to cause task leakage, harming interpretability (Mahinpei et al., 2021). For CEM, we also use its  
 816 *randint* regularization. We always use the AdamW optimized with learning rate 0.001 and train for  
 817 80 epochs, restoring the weights that resulted in the lowest validation loss. Our training/validation  
 818 split is 9/1.

819     **Backbones.** For each model, we use the same backbone (which differs between datasets). For  
 820 CelebA, we train on pretrained ResNet18 embeddings (He et al., 2016), similar to Debot et al.  
 821 (2024). Images are first resized to (224, 224) using bi-linear interpolation. They are then normal-  
 822 ized per channel with means (0.485, 0.456, 0.406) and standard deviations (0.229, 0.224, 0.225).  
 823 By removing the last layer of the pretrained ResNet18, using the resulting model on each image,  
 824 and flattening the output, we obtain an embedding. For MNIST-Addition, we train directly on the  
 825 images. The backbone is a CNN (learned jointly with the rest of the CSM) that consists of the fol-  
 826 lowing layers: a convolution layer with 6 output channels and kernel size 5, a max-pool layer with  
 827 kernel size and stride 2, a ReLu activation, a convolution layer with 16 output channels and kernel  
 828 size 5, a max-pool layer with kernel size and stride 2, a ReLu activation, a flattening layer, a linear  
 829 layer with  $emb\_size//2$  output features, and 3 linear layers each with  $emb\_size//2$  output features  
 830 and ReLu activation (except the last one). The backbone is applied to each MNIST image, and the  
 831 two resulting embeddings are concatenated.  $emb\_size$  is a hyperparameter. If any hyperparameters  
 832 are unmentioned, we use the default values.

833     **Concrete CSMs.** We will describe our implementation of the different CSMs. For each CSM, the  
 834 input first passed through the backbone before being passed to the layers we will describe in what  
 835 follows. Unless explicitly mentioned otherwise, each linear layer has  $emb\_size$  output features.

836     For CRM (Mahinpei et al., 2021), the concept predictor is a neural network consisting of 2 linear  
 837 layers with ReLU activation, and a linear layer with  $n_C$  output features with sigmoid activation. The  
 838 task predictor is a neural network with 3 linear layers and a linear layer with  $n_y$  output features and  
 839 sigmoid activation. The sidechannel is a neural network with 2 linear layers with ReLU activation,  
 840 and a linear layer, outputting an embedding. The prior  $p(z)$  is denoted by a single learnable Torch  
 841 embedding object with  $emb\_size$  weights.

842     For CEM (Espinosa Zarlenga et al., 2022), the sidechannel is a neural network with 3 linear layers  
 843 with  $2 \cdot c\_emb\_size \cdot n_c$  output features and ReLU activation, and a linear layer with  $2 \cdot c\_emb\_size \cdot n_c$   
 844 output features (with  $c\_emb\_size$  a hyperparameter). This is reshaped to  $n_c, 2, c\_emb\_size$ . The  
 845 concept predictor concatenates for each concept its 2 embeddings, and applies a (different) linear  
 846 layer to the resulting embedding with 1 output feature and sigmoid activation. This is the concept  
 847 prediction. The task predictor uses the concept predictions to mix for each concept its 2 embeddings,  
 848 and concatenates the resulting embeddings. The result is passed through 3 linear layers with ReLU  
 849 activation, and a linear layer with  $n_y$  output features and sigmoid activation. The prior  $p(z)$  is  
 850 denoted by a single learnable Torch embedding object with  $2 \cdot c\_emb\_size \cdot n_c$  weights.

851     For DCR (Barbiero et al., 2023), the sidechannel is a neural network with 3 linear layers with  
 852  $2 \cdot c\_emb\_size \cdot n_c$  output features and ReLU activation, and a linear layer with  $2 \cdot c\_emb\_size \cdot n_c$   
 853 output features (with  $c\_emb\_size$  a hyperparameter). This is reshaped to  $n_c, 2, c\_emb\_size$ . The  
 854 concept predictor concatenates for each concept its 2 embeddings, and applies a (different) linear  
 855 layer to the resulting embedding with 1 output feature and sigmoid activation. This is the concept  
 856 prediction. The task predictor uses the concept predictions to mix for each concept its 2 embeddings.  
 857 Each resulting *concept embedding* is passed through a (different) neural network consisting of 3  
 858 linear layers with  $c\_emb\_size$  output features and ReLU activation, and a linear layer with  $2 \cdot n_y$   
 859 output features with sigmoid activation. These should be interpreted as, for each concept, its polarity  
 860 and relevance for each task, as defined by DCR. These are used together with the concept predictions  
 861 to infer the task prediction. We employ DCR’s logic formula (see Barbiero et al. (2023)) using  
 862 the product t-norm. The prior  $p(z)$  is denoted by a single learnable Torch embedding object with  
 863  $2 \cdot c\_emb\_size \cdot n_c$  weights. To avoid needing to finetune DCR’s *temperature* hyperparameter, which  
 864 a human user would do to find a preferred rule parsimony, we modelled DCR’s role and relevance  
 865 with a sigmoid activation instead of their rescaled softmax activation.

864 For CMR (Debot et al., 2024), the concept predictor is a neural network consisting of 2 linear layers with ReLU activation, and a linear layer with  $n_C$  output features with sigmoid activation. The sidechannel is a neural network with 3 linear layers with ReLU activation, and a linear layer with  $n_r \cdot n_y$  output features, where  $n_r$  is CMR’s allowed number of rules to learn (hyperparameter). This is reshaped to  $(n_y, n_r)$  with a softmax on the last dimension (as this is a categorical random variable). The task predictor consists of a Torch embedding object with shape  $n_y, n_r, rule\_emb\_size$  weights, with  $rule\_emb\_size$  a hyperparameter. This is CMR’s rulebook in its latent representation, which is decoded into explicit logic rules by using a neural network on each rule embedding with 3 linear layers with ReLU activation and a linear layer with  $3 \cdot n_c$  output features with softmax activation. The task predictor then uses CMR’s inference formula (see (Debot et al., 2024)) to derive the task prediction from the categorical distribution over rules (given by the sidechannel) and the learned set of rules. In bottleneck mode, CMR applies the entire set of rules to the concepts, each making a  $y$  prediction. The final prediction is the logical OR of the individual predictions. CMR’s hyperparameter for its prototype regularization is set to 0 for the same reason as why we changed DCR’s role and relevance activation, as this hyperparameter would be used by the human to find a preferred rule parsimony.

879 For LRM, we use the same setup as for CRM, with as difference that the task predictor is a linear  
 880 layer with  $n_y$  output features and sigmoid activation.  
 881

882 **Training with and without learnable prior.** Whenever we use SIS regularization, we train with a  
 883 learnable prior. Otherwise, we compute  $p(z)$  by marginalizing out  $x$ , as proposed in the main text.  
 884 For deterministic CSMs that use an embedding as sidechannel (e.g. CRM), this typically means  $p(z)$   
 885 becomes a mixture of delta distributions, with one delta per training instance (as embeddings rarely  
 886 perfectly coincide). As this does not scale, we instead relax this by instead considering a single  $z$ ,  
 887 namely the average of all such embeddings.

888 **CMR\*.** We define CMR\* as an adaptation of CMR by equipping it with rules learned by a decision  
 889 tree. Specifically, we extract the rules decision trees learned on ground truth  $(c, y)$  pairs predicting  
 890 positive classes. We start with decision trees with a maximum depth of 1, increasing it until the  
 891 decision tree has more than 17 such rules, or depth exceeds 40. Then, we take the rules from the  
 892 tree with the highest validation accuracy. We inject these rules into CMR, and allow it to learn 3  
 893 additional rules on its own. This is done through rule interventions, which CMR supports (see Debot  
 894 et al. (2024)).

895 **For obtaining Figure 2**, we perform a **grid search** where  $emb\_size$  is taken  
 896 from  $\{64, 128, 256\}$  and the weight of the SIS regularization is taken from  
 897  $\{0, 0.0001, 0.001, 0.01, 0.05, 0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0\}$ . For CEM and DCR, we use  
 898  $p_{randint} = 0.05$ . **For obtaining Figure 4**, we have 2 phases. In the first phase, we perform  
 899 the same grid search except that we do not employ SIS regularization. For each CSM, we take  
 900 the most accurate configuration on the validation set. Then, in the second phase, for each CSM,  
 901 we train this configuration with different values of the SIS regularization weight, taken from  
 902  $\{0, 0.0001, 0.001, 0.01, 0.05, 0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0\}$ . For each such trained model,  
 903 we compute its intervenability curve by (1) generating a random concept intervention order, (2)  
 904 intervening on increasingly more concepts following this order, (3) each time computing the  
 905 accuracy on the tasks after the intervention. **For obtaining Figure 5**, we take the most accurate  
 906 LRM on the validation set (taken from the first mentioned grid search) with and without SIS  
 907 regularization and inspect the learned weights. **For obtaining Figure 3**, we re-use the results of  
 908 CRM from the first mentioned grid search, and additionally train a *dropout* version inspired by  
 909 Kalampalikis et al. (2025) and a *detach* version inspired by Shang et al. (2024). For the dropout  
 910 version, we extend the grid by considering values for  $p_{dropout}$  from  $\{0.0, 0.2, 0.4, 0.8, 0.8, 1.0\}$ .  
 911 For more details of these two versions, see below. **For obtaining Figure 10**, we first finetune a  
 912 ResNet18 on CUB, where we add a linear layer with a softmax for predicting the classes and a linear  
 913 layer with a sigmoid for predicting the concepts. We use a learning rate of 0.01 (SGD optimizer,  
 914 momentum of 0.9), batch size 128 and train for 200 epochs. This is similar to the configuration  
 915 many other CBM works use (Espinosa Zarlenga et al., 2022). We then drop the classification layers  
 916 and run the CUB images through the ResNet18 to obtain image embeddings, which we use to train  
 917 the CSMs on (instead of on the images). We use the following hyperparameter grid: SIS weight  
 918 within  $\{0, 0.0001, 0.001, 0.01, 0.05, 0.1, 0.5, 1.0, 1.5, 2.0, 3.0, 4.0, 5.0, 7.0, 10.0, 12.0, 20.0, 50.0\}$ ,  
 919  $emb\_size$  512, sidechannel number of hidden layers between 0 and 2, concept predictor number of  
 920 hidden layers either 0 or 1, task predictor number of hidden layers between 0 and 2.

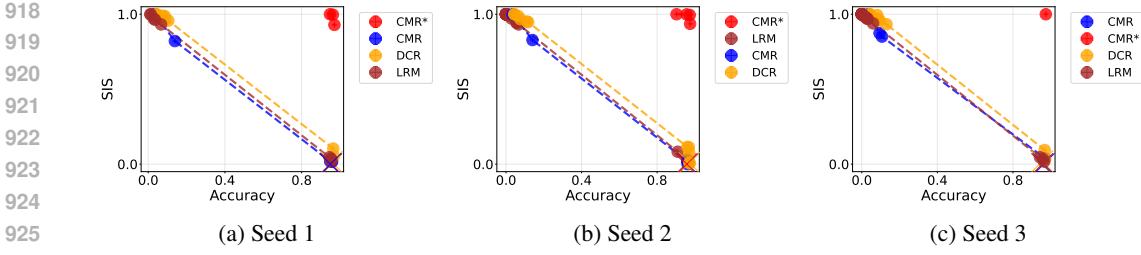


Figure 6: Accuracy vs representation interpretability trade-off in functionally interpretable CSMs, including CMR\*. Different points are different hyperparameter configurations, keeping only pareto-efficient points. Crosses denote the most accurate configuration trained without SIS regularization.

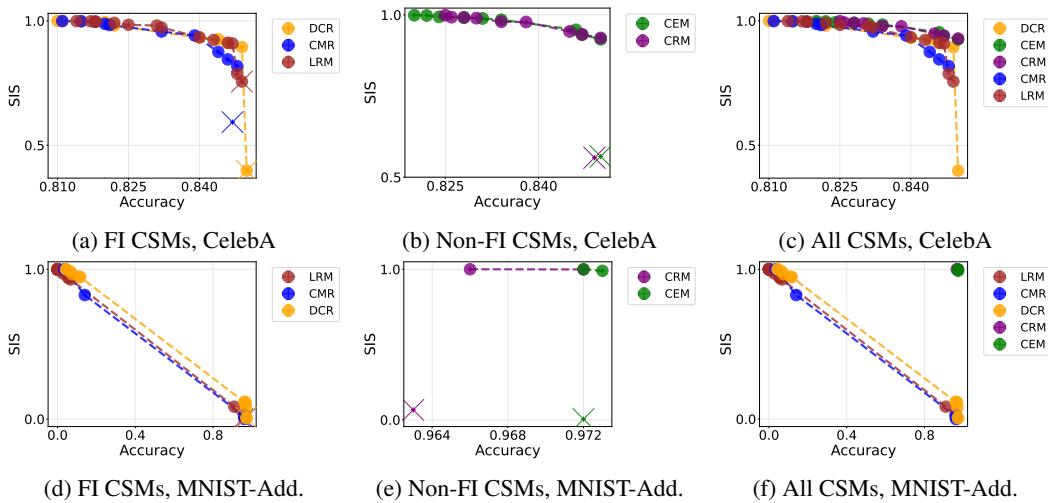


Figure 7: Accuracy vs representation interpretability trade-off in CSMs (seed 2). Different points are different hyperparameter configurations, keeping only pareto-efficient points. Crosses denote the most accurate configuration trained without SIS regularization (excluded from (c) and (f)). "(Non)-FI" denotes the figure only contains (non)-functionally interpretable CSMs.

**Detach version (CRM).** We adapt CRM's architecture to resemble Shang et al. (2024). Concretely, we redefine the task predictor as  $y = f(c) + g(z)$ . For  $g$ , we use a linear layer with  $n_y$  output features and sigmoid activation. For  $f$ , we use a neural network consisting of 3 linear layers with ReLU activation, and a linear layer with  $n_y$  output features and sigmoid activation. The model is trained by optimizing (1) a cross entropy between  $f(c)$  and the label, and (2) a cross entropy between  $f(c).detach() + g(x)$  and the label. The first objective tries to maximize concept usage, while the second one trains the sidechannel.

**Dropout version (CRM).** During training, for each batch, we set the entire sidechannel to 0 with probability  $p_{dropout}$ , which is a hyperparameter. This encourages the task predictor to rely on the concepts when the sidechannel is dropped out.

## D.2 ADDITIONAL RESULTS

With a different rule learner, CMR can achieve high accuracy when allowed to learn enough rules (Figure 6). CMR\* achieves near-perfect accuracy and representation interpretability on MNIST-Addition.

Figures 7 and 8 give additional accuracy-interpretability trade-off results for seed 2 and 3. Figure 9 shows some additional intervenability curves, which also show that SIS regularization improves intervenability (similar to Figure 4 in the main text).

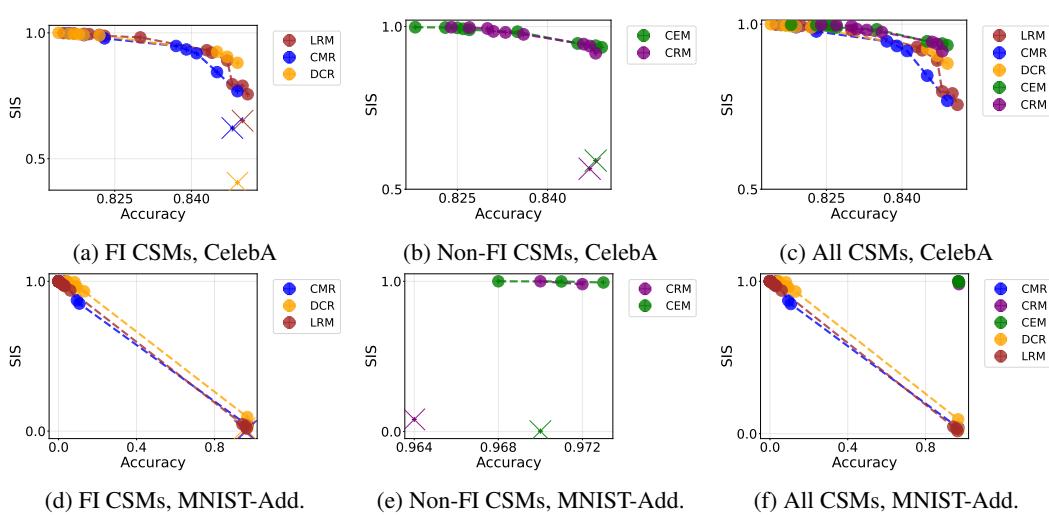


Figure 8: Accuracy vs representation interpretability trade-off in CSMs (seed 3). Different points are different hyperparameter configurations, keeping only pareto-efficient points. Crosses denote the most accurate configuration trained without SIS regularization (excluded from (c) and (f)). ”(Non)-FI” denotes the figure only contains (non)-functionally interpretable CSMs.

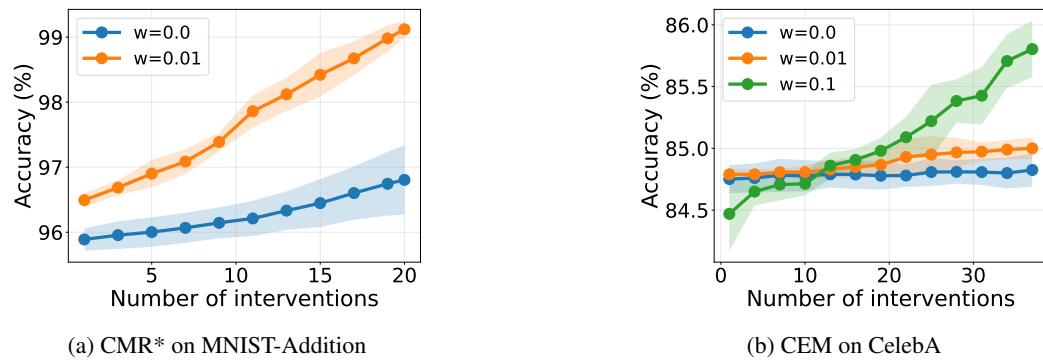


Figure 9: Intervenability in CSMs with ( $w > 0$ ) and without ( $w = 0$ ) SIS regularization for different regularization weights  $w$ . The y-axis denotes accuracy after intervening on a number of concepts denoted by the x-axis.

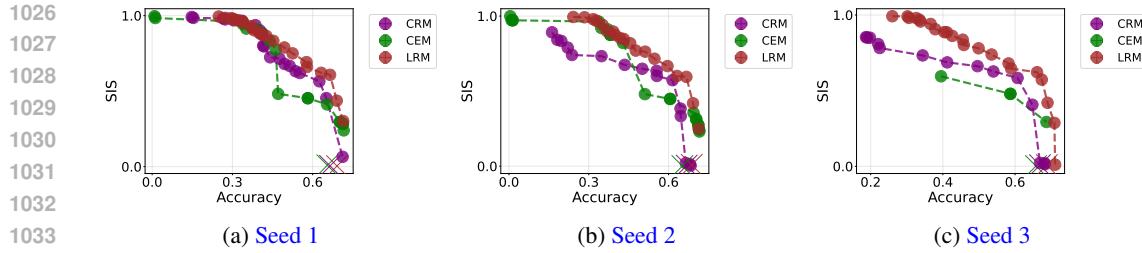


Figure 10: Accuracy vs representation interpretability trade-off on CUB. Different points are different hyperparameter configurations, keeping only pareto-efficient points. Crosses denote the most accurate configuration trained without SIS regularization.

On CUB, we only used CRM, LRM and CEM because CMR and DCR do not naturally support multiclass classification. Figure 10 shows their accuracy-interpretability trade-off, which is similar to the ones in the main text.

## E LLM USAGE DECLARATION

During writing, Large Language Models (LLMs) were used only to polish and improve the clarity of the text.

## F CODE AND LICENSES

Our code will be made publicly available upon acceptance under the Apache license, Version 2.0. We used Python 3.10.12 and the following libraries: PyTorch v2.5.1 (BSD license) (Paszke et al., 2019), PyTorch-Lightning v2.5.0 (Apache license 2.0), scikit-learn v1.5.2 (BSD license) (Pedregosa et al., 2011), PyC v0.0.11 (Apache license 2.0). We used CUDA v12.7 and plots were made using Matplotlib (BSD license). The CelebA dataset is available for non-commercial research purposes only<sup>4</sup> and MNIST is available on the web with the CC BY-SA 3.0 DEED license.

<sup>4</sup><https://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>