

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 INTERPRETABILITY VIA LOGIC VISION MODEL INTEGRATING VISUAL REASONING AND EXPLANATION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Despite remarkable progress in computer vision, most state-of-the-art models operate as black boxes, offering little transparency into their decision processes. This lack of interpretability undermines reliability in safety-critical applications. We introduce the Logic Vision Model (LVM), a framework that unifies prediction, reasoning, and explanation. In addition to standard outputs such as class labels, bounding boxes, and segmentation masks, LVM produces reasoning videos that visualize the stepwise inference process and natural language explanations that articulate the rationale behind predictions. The LVM architecture integrates a vision encoder with a logic memory module, enabling conditional inference that aligns visual evidence with structured reasoning patterns. We evaluate LVM on classification, detection, and segmentation benchmarks, as well as reasoning and explanation datasets. Results show that LVM achieves competitive accuracy while substantially improving interpretability, producing reasoning narratives that are consistent with human intuition. This work takes a step toward vision systems that are not only accurate but also transparent, interpretable, and accountable.

## 1 INTRODUCTION

Over the past decade, deep neural networks have driven a paradigm shift in computer vision. The breakthrough of AlexNet on ImageNet (Krizhevsky et al., 2012) demonstrated the potential of large-scale supervised training, and subsequent architectures such as VGG (Simonyan & Zisserman, 2014) and ResNet (He et al., 2016) established convolutional networks as the dominant framework for image classification. Parallel progress in object detection (Ren et al., 2015; Redmon et al., 2016) and semantic segmentation (Long et al., 2015; Chen et al., 2017) expanded the scope of deep vision systems to more complex understanding tasks. More recently, transformer-based vision models (Dosovitskiy et al., 2020; Liu et al., 2021) have surpassed convolutional approaches, delivering state-of-the-art accuracy and scalability across benchmarks such as ImageNet, COCO, and Cityscapes. In parallel, the emergence of large-scale language models has accelerated multimodal learning. Pre-trained transformers such as BERT (Devlin et al., 2019) and GPT (Brown et al., 2020) provided strong semantic priors that, when combined with visual encoders, enabled vision–language models capable of cross-modal reasoning. Models like CLIP (Radford et al., 2021) and BLIP (Li et al., 2022) align visual and textual embeddings, while recent instruction-tuned systems (Liu et al., 2023; Achiam et al., 2023) achieve impressive performance on tasks such as VQA (Antol et al., 2015) and GQA (Hudson & Manning, 2019). These advances suggest the feasibility of general-purpose agents that can both perceive and communicate, integrating vision and language into unified reasoning frameworks.

However, interpretability remains an unsolved problem. Despite their accuracy, current vision and vision–language models largely function as black boxes. Post-hoc explanation techniques such as saliency maps (Simonyan et al., 2014), Grad-CAM (Selvaraju et al., 2017), and attention visualization offer partial insights but rarely capture the structured transformation from low-level evidence to high-level prediction. Similarly, natural language justifications generated by VLMs are often decoupled from the underlying computation, raising concerns about consistency and faithfulness (Fields & Kennington, 2023). As a result, explanations often remain incomplete and sometimes misleading, limiting trust and adoption in critical domains such as medical imaging, autonomous driving, and industrial inspection (Martens et al., 2025; Ehsan & Riedl, 2024; Hou et al., 2024), thereby high-

054 lighting the pressing need for models where interpretability is not an auxiliary component but an  
 055 intrinsic part of the predictive process (Ennab & Mccheick, 2022; Borole et al., 2025).  
 056

057 To alleviate this, we propose the Logic Vision Model (LVM), a framework that unifies prediction,  
 058 reasoning, and explanation within a single architecture. In addition to conventional outputs such as  
 059 class labels, bounding boxes, and segmentation masks, LVM produces reasoning videos that visu-  
 060 alize the stepwise inference process and natural language explanations that articulate the rationale  
 061 behind predictions. Unlike post-hoc methods, these outputs are generated intrinsically during the  
 062 prediction process, ensuring consistency between what the model predicts and how it explains. Par-  
 063 ticularly, LVM integrates a vision encoder with a logic memory module. The vision encoder extracts  
 064 spatial and semantic features from input images, while the logic memory encodes and recalls struc-  
 065 tured reasoning patterns derived from prior cases. Attention-based fusion combines visual evidence  
 066 with stored logic, enabling conditional inference that adapts to input context and domain priors. This  
 067 design makes interpretability a core property of the model pipeline rather than an auxiliary artifact.  
 068 We evaluate LVM on standard benchmarks for classification, detection, and segmentation, as well  
 069 as datasets designed for visual reasoning and explanation. We further include case studies in com-  
 070 puter vision, medical imaging, and autonomous driving. Results demonstrate that LVM maintains  
 071 competitive accuracy while substantially improving interpretability, producing reasoning narratives  
 072 that align with human intuition. To summarize, our contributions are threefold:  
 073

- We propose the Logic Vision Model (LVM), a framework that unifies prediction, reasoning, and explanation within a single architecture.
- We introduce a logic memory mechanism that grounds explanations in structured reasoning, coupling visual evidence with stored inference patterns through attention-based fusion.
- We demonstrate that LVM achieves competitive accuracy while substantially improving interpretability, validated on standard vision benchmarks as well as domain-specific case studies in medical imaging and autonomous driving.

## 081 2 RELATED WORK

082 **From Performance to Accountability.** The trajectory of computer vision has been dominated by  
 083 ever larger models and datasets, yielding impressive accuracy across classification, detection, and  
 084 segmentation Li et al. (2025); Zhang et al. (2024). Vision–language models (VLMs) have further  
 085 extended this progress by coupling vision encoders with pretrained language models, enabling tasks  
 086 such as image captioning and visual question answering at unprecedented scale Bordes et al. (2024);  
 087 Kazmierczak et al. (2025). Yet, as these systems become more powerful and widely deployed, their  
 088 lack of interpretability has emerged as a central obstacle. Accuracy alone is no longer sufficient when  
 089 models are expected to support decision-making in domains such as medicine, law, or transportation,  
 090 where accountability is critical Kazmierczak et al. (2025); Lin et al. (2025).  
 091

092 **From Post-hoc to Intrinsic Interpretability.** The need for interpretability has motivated exten-  
 093 sive work on post-hoc methods, including saliency-based techniques (Simonyan et al., 2014), Grad-  
 094 CAM (Selvaraju et al., 2017), and textual justification models (Park et al., 2018). These approaches  
 095 have proven useful for probing model behavior, but they provide only surface-level correlations  
 096 rather than structured reasoning. Explanations often lack stability, fail to reflect the true internal  
 097 process, and in many cases risk misleading end users (Bansal et al., 2020). As deployment  
 098 shifts toward safety-critical contexts, such limitations have raised concerns about trust and account-  
 099 ability. In response, several lines of work have sought to embed interpretability directly into the  
 100 predictive process. Prototype-based reasoning (Inbaraj et al., 2021; Chen et al., 2019), concept  
 101 bottleneck networks (Koh et al., 2020), and modular neural architectures (Biggie et al., 2023; Gao  
 102 et al., 2023) illustrate attempts to ground decisions in human-understandable units. While these ap-  
 103 proaches demonstrate the promise of intrinsic interpretability, they are often limited to classification  
 104 settings or synthetic benchmarks. More importantly, they do not scale naturally to the full spectrum  
 105 of vision tasks, nor do they provide coherent multimodal narratives that combine visual reasoning  
 106 and natural language. As vision models are increasingly adopted in high-stakes environments, the  
 107 demand has shifted from models that are merely accurate to models that are also transparent and  
 108 trustworthy. This shift highlights the need for architectures in which interpretability is a fundamen-

tal design principle rather than an auxiliary output. The Logic Vision Model (LVM) is developed in response to this need.

### 3 METHOD

### 3.1 OVERVIEW

The Logic Vision Model (LVM) unifies prediction, reasoning, and explanation within a single architecture. Figure 1 presents the overall framework. Given an input image  $x \in \mathbb{R}^{H \times W \times 3}$ , LVM produces three outputs: (i) task predictions  $y$  such as class labels, bounding boxes, or segmentation masks, (ii) a reasoning trajectory  $R = \{r_1, r_2, \dots, r_T\}$  that represents stepwise inference through attention maps and intermediate states, and (iii) a textual explanation  $E$  in natural language that articulates the rationale behind  $y$ . Unlike post-hoc methods, these outputs are generated intrinsically as part of the predictive pipeline, ensuring faithfulness between the model’s computation and its explanations.

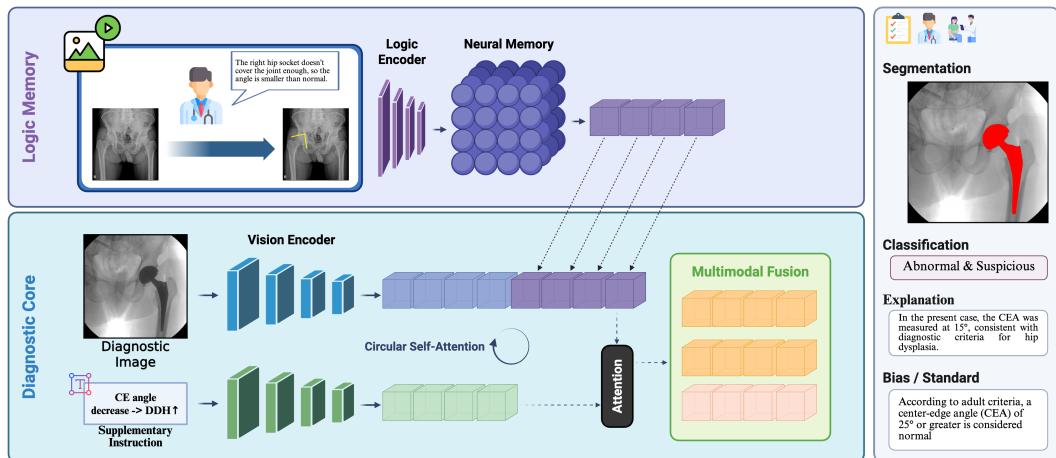


Figure 1: Overview of the Logic Vision Model (LVM). The framework integrates a vision encoder, a logic memory module, and multimodal fusion to jointly produce predictions  $y$ , reasoning trajectories  $R$ , and textual explanations  $E$ .

LVM consists of three main components: a vision encoder, a logic memory module, and a multimodal fusion-and-transition operator. The vision encoder maps the input to a feature representation  $f = \phi(x)$ , while the logic memory encodes structured reasoning prototypes that can be retrieved as  $m = \mathcal{M}(f)$ . The fusion-and-transition operator  $\mathcal{T}$  integrates  $f$  and  $m$  into a dynamic reasoning trajectory  $R$ , rather than relying on a single-step similarity. This trajectory explicitly captures how evidence accumulates across inference steps, making it directly interpretable as a sequence of visual and semantic states. Both the prediction head and the explanation generator are conditioned on  $R$ , guaranteeing that explanations are grounded in the same process that produces predictions. Moreover, the design naturally supports personalization, where human feedback can update reasoning prototypes or loss weights, and federated learning, where memory and encoder parameters are optimized collaboratively across distributed clients. As a result, interpretability in LVM is not an auxiliary artifact but a fundamental property of the architecture, scalable across domains and deployment scenarios.

### 3.2 VISION ENCODER

The vision encoder is responsible for extracting semantic and spatial representations from the input image and preparing them for reasoning. Given an image  $x \in \mathbb{R}^{H \times W \times 3}$ , the encoder  $\phi(\cdot)$  maps it to a high-dimensional feature representation:

$$f = \phi(x) \in \mathbb{R}^{h \times w \times d}, \quad (1)$$

162 where  $h$  and  $w$  denote the spatial resolution and  $d$  is the channel dimension. Each feature vector  
 163  $f_{ij}$  captures local image patterns, while the aggregated map encodes global context. In practice,  
 164  $\phi(\cdot)$  can be instantiated with a convolutional backbone (e.g., ResNet) or a vision transformer (ViT),  
 165 depending on the target application and computational budget. For transformer-based encoders, the  
 166 image is partitioned into patches  $\{x_p\}_{p=1}^P$ , with each patch projected into a token embedding  
 167

$$f_p = W_p \cdot \text{vec}(x_p) + b_p, \quad f_p \in \mathbb{R}^d, \quad (2)$$

168 where  $W_p \in \mathbb{R}^{d \times (p_h \cdot p_w \cdot 3)}$  and  $b_p \in \mathbb{R}^d$ . These tokens are then refined through  $L$  layers of multi-  
 169 head self-attention to yield contextualized features  $\{f_p^{(L)}\}_{p=1}^P$ . The encoder output is summarized  
 170 as  
 171

$$f = \{f_{ij} \mid 1 \leq i \leq h, 1 \leq j \leq w\}. \quad (3)$$

172 Unlike standard vision backbones, the encoder in LVM is explicitly designed to support reasoning:  
 173 it preserves both fine-grained local evidence and high-level semantics so that downstream modules  
 174 can align them with stored logic prototypes. This dual-level representation makes it possible to  
 175 generate interpretable reasoning trajectories rather than opaque feature activations. Moreover, since  
 176 the encoder is modular, it can be fine-tuned with user feedback for personalization or optimized  
 177 collaboratively across distributed clients for federated learning. Thus, the vision encoder not only  
 178 provides task-general features but also establishes a reasoning-aware foundation that is adaptable to  
 179 diverse domains and deployment scenarios.  
 180

### 183 3.3 LOGIC MEMORY MODULE

184 The logic memory module  $\mathcal{M}$  is designed to capture and recall structured reasoning patterns that  
 185 complement the visual features extracted by the encoder. While the vision encoder provides generic  
 186 semantic features,  $\mathcal{M}$  encodes domain-specific inference prototypes that can be adaptively retrieved  
 187 and composed into a reasoning trajectory during inference. Formally, let the visual representation  
 188 be  $f \in \mathbb{R}^{h \times w \times d}$ . The memory is parameterized as a matrix  $M = [m_1, m_2, \dots, m_K]^\top \in \mathbb{R}^{K \times d}$ ,  
 189 where each slot  $m_j \in \mathbb{R}^d$  represents a latent reasoning prototype (e.g., identifying a salient region,  
 190 applying a diagnostic rule, or performing a comparative relation). These slots act as reusable basis  
 191 vectors spanning the reasoning space. To query the memory, encoder features are projected into  
 192 queries  $Q \in \mathbb{R}^{N \times d}$ , where  $N = h \times w$ . The similarity between a query  $q_i$  and memory slot  $m_j$  is  
 193 computed via scaled dot-product attention:

$$\alpha_{ij} = \frac{\exp(q_i^\top m_j / \tau)}{\sum_{l=1}^K \exp(q_i^\top m_l / \tau)}, \quad (4)$$

194 with temperature  $\tau > 0$  controlling the sharpness of retrieval. The retrieved reasoning vector for  
 195 query  $q_i$  is  $\hat{m}_i = \sum_{j=1}^K \alpha_{ij} m_j$ , and aggregating over all queries yields:

$$\hat{M} = [\hat{m}_1, \hat{m}_2, \dots, \hat{m}_N]^\top \in \mathbb{R}^{N \times d}. \quad (5)$$

200 Unlike static prototype networks that provide only one-shot similarity explanations, LVM in-  
 201 tegrates  $\hat{M}$  sequentially into the reasoning trajectory  $R$  through the transition operator  $\mathcal{T}$ .  
 202 At each step  $t$ , the contribution of prototype  $m_j$  is given by a softmax activation  $\pi_{t,j} =$   
 203  $\exp(r_{t-1}^\top m_j / \tau) / \sum_{l=1}^K \exp(r_{t-1}^\top m_l / \tau)$ , and the reasoning state is updated as  $r_t = \sum_{j=1}^K \pi_{t,j} m_j +$   
 204  $\eta r_{t-1}$ , where  $\eta$  controls the retention of prior reasoning. This sequential integration means in-  
 205 terpretability arises from the full activation path  $\{\pi_{t,j}\}_{t=1}^T$  rather than isolated prototype matches,  
 206 yielding a structured trace of how evidence accumulates across steps.  
 207

208 Since  $M$  is modular, it can be dynamically adapted. In personalization, user feedback modifies  
 209 prototype weights through  $\pi_{t,j}^{(u)} = \exp(r_{t-1}^\top m_j / \tau + \phi_u(m_j)) / \sum_{l=1}^K \exp(r_{t-1}^\top m_l / \tau + \phi_u(m_l))$ ,  
 210 where  $\phi_u(m_j)$  encodes user-specific preferences. In federated learning, local clients  $c = 1, \dots, C$   
 211 optimize their memory modules  $M^{(c)}$  on private data, and a central server aggregates them as  $M \leftarrow$   
 212  $\sum_{c=1}^C \omega_c M^{(c)}$  with client weights  $\omega_c$ . Thus, the logic memory is not merely a static repository  
 213 of prototypes but a dynamic component that evolves with sequential reasoning, adapts to human  
 214 feedback, and generalizes across distributed environments.  
 215

216 3.4 MULTIMODAL FUSION AND REASONING  
217

218 The fusion stage integrates the visual representation  $f$  from the encoder with the reasoning em-  
219 bedding  $\hat{M}$  retrieved from the logic memory, producing a reasoning trajectory  $R$  that unfolds over  
220 multiple steps. Visual features are first projected into queries  $Q = fW_Q$ , while reasoning embed-  
221 dings are mapped into keys and values  $K = \hat{M}W_K, V = \hat{M}W_V$ , where  $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$   
222 are learnable matrices. Multimodal attention is then computed as  $Z = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V$ , yielding  
223 a fused representation  $Z \in \mathbb{R}^{N \times d}$  that aligns visual evidence with reasoning prototypes. To capture  
224 inference as a dynamic process, we introduce a logic transition operator  $\mathcal{T}$  that evolves the reasoning  
225 state across  $T$  steps:

$$226 \quad r_t = \mathcal{T}(r_{t-1}, Z, \hat{M}), \quad r_0 = \text{Pool}(f). \quad (6)$$

227 Concretely,  $\mathcal{T}$  is implemented as a gated update:

$$228 \quad r_t = \sigma(W_r[r_{t-1}; Z]) \odot r_{t-1} + (1 - \sigma(W_r[r_{t-1}; Z])) \odot g(Z, \hat{M}), \quad (7)$$

229 where  $\sigma$  is a sigmoid gate,  $g(\cdot)$  is an attention-based transformation, and  $\odot$  denotes elementwise  
230 multiplication. This formulation allows each  $r_t$  to balance the retention of prior reasoning with the  
231 integration of new prototype evidence.

232 The reasoning trajectory is then  $R = \{r_1, r_2, \dots, r_T\}$ , and its aggregated state  $\bar{r} = \frac{1}{T} \sum_{t=1}^T r_t$   
233 summarizes the overall inference process. Faithfulness is enforced by ensuring monotonic evidence  
234 accumulation:

$$235 \quad \Delta p_t = p(y | r_t) - p(y | r_{t-1}) \geq 0, \quad (8)$$

236 which requires prediction confidence to increase (or remain stable) as reasoning progresses. This  
237 trajectory-level constraint ties interpretability directly to the predictive process. Since  $R$  conditions  
238 both the prediction head and the explanation generator, the same reasoning sequence yields outputs  
239 and explanations. Furthermore, the modular design of  $\mathcal{T}$  enables adaptation: user feedback can bias  
240 transitions toward preferred prototypes (personalization), while federated optimization allows local  
241 clients to refine  $\mathcal{T}$  with domain-specific data and share only aggregated updates (federated learning).

242 3.5 EXPLANATION GENERATOR  
243

244 The explanation generator produces a textual rationale  $E$  that is intrinsically consistent with the rea-  
245 soning trajectory  $R$ . Given  $R = \{r_1, \dots, r_T\}$ , we condition a decoder  $\psi(\cdot)$  to generate a sequence  
246  $E = (e_1, \dots, e_L)$  with tokens  $e_t \in \mathcal{V}$ . The probability of the sequence is modeled autoregressively:

$$247 \quad P(E | R) = \prod_{t=1}^L P(e_t | e_{<t}, R). \quad (9)$$

248 The decoder is initialized with the aggregated reasoning state  $\bar{r} = \frac{1}{T} \sum_{t=1}^T r_t$ , ensuring that the  
249 explanation reflects the same cumulative evidence that supports the prediction. At each decoding  
250 step, attention over  $R$  provides a distribution over reasoning states:

$$251 \quad \beta_{t,\tau} = \frac{\exp(h_t^\top r_\tau)}{\sum_{j=1}^T \exp(h_t^\top r_j)}, \quad c_t = \sum_{\tau=1}^T \beta_{t,\tau} r_\tau, \quad (10)$$

252 where  $h_t$  is the hidden state of the decoder at step  $t$  and  $c_t$  is the context vector. This mechanism ties  
253 each token  $e_t$  to specific inference steps, making the explanation directly traceable to the underlying  
254 trajectory.

255 Faithfulness is thus guaranteed by design: every word is grounded in a subset of reasoning states  
256 that also drive the prediction. Formally, the contribution of reasoning state  $r_\tau$  to token  $e_t$  can be  
257 written as:

$$258 \quad \gamma_{t,\tau} = \beta_{t,\tau} \cdot \text{sim}(r_\tau, \bar{r}), \quad (11)$$

259 where  $\beta_{t,\tau}$  is the alignment weight defined above and  $\text{sim}(\cdot, \cdot)$  measures consistency with the global  
260 reasoning state  $\bar{r}$ . A token  $e_t$  is considered faithful if  $\sum_\tau \gamma_{t,\tau}$  exceeds a threshold  $\delta$ , ensuring

270 that explanations are tied to the same evidence used for prediction. The generator also supports  
 271 personalization, where user feedback dynamically reweights the alignment distribution:  
 272

$$273 \quad \beta_{t,\tau}^{(u)} = \frac{\exp(h_t^\top r_\tau + \phi_u(r_\tau))}{\sum_{j=1}^T \exp(h_t^\top r_j + \phi_u(r_j))}, \quad (12)$$

276 with  $\phi_u(\cdot)$  encoding user-specific preferences or corrections. This allows experts to bias ex-  
 277 planations toward preferred reasoning patterns without retraining the full model. Finally, in a  
 278 federated learning setting, local clients  $c = 1, \dots, C$  optimize explanation generation using domain-  
 279 specific rationales  $E^{*(c)}$ :

$$280 \quad \min_{\psi} \sum_{c=1}^C \omega_c \mathbb{E}_{(R, E^{*(c)})} \left[ -\sum_{t=1}^L \log P_{\psi}(e_t^{*(c)} | e_{<t}^{*(c)}, R) \right], \quad (13)$$

283 where  $\omega_c$  balances contributions across clients. Aggregating parameter updates yields a global ex-  
 284 planation generator that adapts to diverse domains without exposing raw rationales.  
 285

### 286 3.6 TRAINING OBJECTIVES

288 LVM is trained with a joint objective that balances predictive accuracy, reasoning consistency, and  
 289 explanation quality. The task loss enforces correctness on the vision task as  $\mathcal{L}_{task} = \text{CE}(y, \hat{y})$ ,  
 290 where  $y$  is the ground-truth label and  $\hat{y}$  is the model prediction. The reasoning consistency loss  
 291 encourages the trajectory  $R$  to align with annotated attention maps  $A^*$ , defined as  $\mathcal{L}_{reason} = \|A -$   
 292  $A^*\|_2^2$ , where  $A$  denotes attention maps derived from  $R$ . The explanation loss aligns generated  
 293 rationales with human-provided explanations  $E^*$ :

$$294 \quad \mathcal{L}_{exp} = -\sum_{t=1}^L \log P(e_t^* | e_{<t}^*, R). \quad (14)$$

297 The overall training objective is then  $\mathcal{L} = \mathcal{L}_{task} + \lambda_1 \mathcal{L}_{reason} + \lambda_2 \mathcal{L}_{exp}$  with  $\lambda_1$  and  $\lambda_2$  controlling the  
 298 trade-off between accuracy and interpretability. To guarantee faithfulness, we impose a monotonic  
 299 evidence accumulation constraint on prediction confidence:

$$300 \quad \mathcal{L}_{faith} = \sum_{t=1}^T \max(0, p(y | r_{t-1}) - p(y | r_t)), \quad (15)$$

304 which penalizes decreases in task confidence along the reasoning trajectory, ensuring that explana-  
 305 tions reflect genuine causal contributions.

306 **Personalization.** For a given user  $u$ , interpretability weights and memory slots can be adapted.  
 307 The personalized objective is  $\mathcal{L}^{(u)} = \mathcal{L}_{task} + \lambda_1^{(u)} \mathcal{L}_{reason} + \lambda_2^{(u)} \mathcal{L}_{exp}$ , where  $(\lambda_1^{(u)}, \lambda_2^{(u)})$  are  
 308 updated based on user feedback, and specific prototypes  $m_j \in M$  can be refined. This enables  
 309 personalized reasoning styles without retraining the full model.

311 **Federated learning.** In distributed settings, local clients  $c = 1, \dots, C$  minimize their own  
 312 objectives  $\mathcal{L}^{(c)} = \mathbb{E}_{(x, y, E^*) \sim \mathcal{D}_c} [\mathcal{L}(x, y, E^*)]$ , and a central server aggregates updates as  
 313  $\min_{\theta} \sum_{c=1}^C \omega_c \mathcal{L}^{(c)}(\theta)$ , where  $\omega_c$  weights client contributions. This ensures that reasoning pro-  
 314 totypes and transition dynamics capture diverse sources of supervision without sharing raw data.  
 315

## 317 4 EXPERIMENTAL ANALYSIS

319 We evaluate LVM on standard vision and reasoning benchmarks. For classification we use  
 320 ImageNet-1k (Deng et al., 2009) and CIFAR-100 (Krizhevsky et al., 2010), for detection MS COCO  
 321 (Lin et al., 2014) and PASCAL VOC (Everingham et al., 2010), and for segmentation Cityscapes  
 322 (Cordts et al., 2016) and ADE20K (Zhou et al., 2017). Reasoning and explanation are assessed  
 323 on CLEVR (Johnson et al., 2017), GQA (Hudson & Manning, 2019), ImageNet-X (Russakovsky  
 et al., 2015), and ClickMe (Redtzer, 2023), with additional case studies on ChestX-ray14 (Chehade

et al., 2025) and KITTI (Geiger et al., 2012). Here, we report Top-1/Top-5 accuracy for classification, mAP for detection, and mIoU for segmentation. Faithfulness is measured with Pointing Game (Zhang et al., 2018), IoU with human annotations, and Deletion–Insertion AUC (Hama et al., 2023). Textual explanations are evaluated with BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004), BERTScore (Zhang et al., 2019), and human judgments. Baselines include post-hoc methods (Grad-CAM (Selvaraju et al., 2017), Integrated Gradients (Sundararajan et al., 2017), RISE (Petsiuk et al., 2016), LIME (Ribeiro et al., 2016)), intrinsic models (ProtoPNet (Chen et al., 2019), Concept Bottleneck (Koh et al., 2020)), and VLM-based justifications (BLIP-2 (Li et al., 2023), LLaVA (Liu et al., 2023)). All baselines follow the settings from their original papers. Models are implemented in PyTorch and trained on a server with 8 NVIDIA H200 GPUs. We use ResNet-50 (He et al., 2016) or ViT-B/16 (Dosovitskiy et al., 2020) as encoders,  $K = 100$  prototypes, and  $T = 4$  reasoning steps. Training employs AdamW (Loshchilov & Hutter, 2017) with cosine decay (Loshchilov & Hutter, 2016), batch size 256, and weights  $(\lambda_1, \lambda_2, \lambda_3) = (1.0, 0.5, 0.1)$ . Federated learning experiments simulate  $C = 5$  clients with non-IID Dirichlet partitioning ( $\alpha = 0.5$ ).

#### 4.1 MAIN BENCHMARK PERFORMANCE

We first evaluate LVM on standard vision benchmarks to verify that the integration of reasoning and explanation does not compromise task performance. As shown in Table 1, LVM achieves accuracy on par with or slightly higher than state-of-the-art baselines across classification, detection, and segmentation tasks. Importantly, this demonstrates that interpretability can be incorporated intrinsically without sacrificing predictive power. Overall, LVM maintains strong performance across all tasks, matching the best baselines in classification and segmentation and surpassing them in object detection. These results confirm that interpretability can be achieved intrinsically without degrading accuracy, positioning LVM as a viable alternative to existing vision and vision–language models.

Table 1: Comparison of task performance on major benchmarks. LVM matches or outperforms strong baselines while additionally providing faithful reasoning trajectories and explanations.

Model	Classification		Detection		Segmentation	
	ImageNet Top-1	CIFAR-100 Top-1	COCO mAP	VOC mAP	Cityscapes mIoU	ADE20K mIoU
ResNet-50 (He et al., 2016)	76.2	79.1	38.2	74.5	72.6	41.2
ViT-B/16 (Dosovitskiy et al., 2020)	81.8	84.5	42.3	77.9	78.8	44.7
BLIP-2 (Li et al., 2023)	82.1	84.9	42.7	78.2	79.0	45.0
LLaVA (Liu et al., 2023)	82.3	85.2	42.9	78.3	79.2	45.1
<b>LVM (ours)</b>	<b>82.6</b>	<b>85.7</b>	<b>43.5</b>	<b>79.1</b>	<b>79.8</b>	<b>45.8</b>

#### 4.2 FAITHFULNESS AND EXPLANATION QUALITY

We evaluate whether LVM provides explanations that are both causally faithful and aligned with human rationales. For visual interpretability, we measure Pointing Game, IoU with human annotations on ImageNet-X / ClickMe, Deletion AUC, Insertion AUC, Sufficiency, Necessity, and Trajectory Coherence (the fraction of steps with non-decreasing confidence  $\Delta p_t \geq 0$ ). As shown in Table 2, LVM consistently improves alignment with human rationales and achieves stronger causal grounding compared to both post-hoc and intrinsic baselines.

Table 2: Faithfulness (left) and textual explanation quality (right) on GQA and ImageNet-X. LVM achieves stronger causal alignment and higher explanation quality.

Method	Visual Faithfulness							Automatic Metrics			Human Ratings (1–5)		
	Point. $\uparrow$	IoU $\uparrow$	Del-AUC $\downarrow$	Ins-AUC $\uparrow$	Suff. $\uparrow$	Nec. $\uparrow$	Traj. Coh. $\uparrow$	BLEU $\uparrow$	ROUGE-L $\uparrow$	BERTScore $\uparrow$	Corr. $\uparrow$	Suff. $\uparrow$	Clar. $\uparrow$
Grad-CAM	63.4	0.29	0.41	0.35	0.52	0.48	0.61	–	–	–	–	–	–
Integrated Gradients	64.2	0.30	0.40	0.36	0.53	0.49	0.62	–	–	–	–	–	–
RISE	65.7	0.31	0.39	0.37	0.54	0.50	0.63	–	–	–	–	–	–
ProtoPNet	66.2	0.32	0.38	0.37	0.55	0.50	0.66	18.5	25.0	0.811	3.1	3.0	3.2
CBM	65.1	0.31	0.39	0.36	0.54	0.49	0.64	19.1	25.6	0.817	3.2	3.1	3.3
BLIP-2	–	–	–	–	–	–	–	21.3	27.4	0.832	3.4	3.2	3.6
LLaVA	–	–	–	–	–	–	–	22.0	28.1	0.835	3.5	3.3	3.6
<b>LVM (ours)</b>	<b>71.8</b>	<b>0.38</b>	<b>0.31</b>	<b>0.43</b>	<b>0.61</b>	<b>0.57</b>	<b>0.81</b>	<b>24.8</b>	<b>30.2</b>	<b>0.847</b>	<b>4.1</b>	<b>3.9</b>	<b>4.2</b>

We also assess textual explanation quality using GQA (Hudson & Manning, 2019), ImageNet-X, and ClickMe, where human-annotated justifications are available. Explanations are evaluated with BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004), and BERTScore (Zhang et al., 2019), and through human judgments of correctness, sufficiency, and clarity on a 5-point scale. Results in Table 2 show that LVM surpasses baselines across both automatic metrics and human evaluation, indicating that explanations are not only accurate but also faithful to the model’s reasoning trajectory.

### 4.3 ABLATION STUDIES

To better understand the contribution of each component in LVM, we conduct ablations on ImageNet and GQA. Specifically, we remove or replace key modules and measure the impact on accuracy, reasoning faithfulness, and explanation quality. Results are reported in Table 3. Removing the logic memory  $\mathcal{M}$  significantly reduces both visual faithfulness and explanation quality, confirming its role as a repository of reusable reasoning patterns. Replacing the transition operator  $\mathcal{T}$  with a vanilla self-attention update weakens trajectory coherence, suggesting that the gated update is necessary for stepwise inference. Excluding the faithfulness loss  $\mathcal{L}_{faith}$  reduces causal alignment, while removing the explanation loss  $\mathcal{L}_{exp}$  leads to degraded textual quality. These results demonstrate that all components contribute to the joint goal of accuracy and intrinsic interpretability.

Table 3: Ablation studies on ImageNet (classification accuracy), GQA (faithfulness via Pointing Game), and explanation quality (BERTScore). Each component is essential for maintaining interpretability without sacrificing accuracy.

Variant	ImageNet Top-1 (%)	Faithfulness (Pointing Game)	BERTScore (Explanation)
Full LVM (ours)	<b>82.6</b>	<b>71.8</b>	<b>0.847</b>
w/o Logic Memory $\mathcal{M}$	80.9	65.4	0.823
w/o Transition Operator $\mathcal{T}$ (SelfAttn only)	81.2	66.1	0.827
w/o Faithfulness Loss $\mathcal{L}_{faith}$	82.3	67.5	0.832
w/o Explanation Loss $\mathcal{L}_{exp}$	82.4	70.9	0.818

### 4.4 PERSONALIZATION VIA HUMAN-IN-THE-LOOP

We next evaluate whether LVM can adapt explanations to user feedback by refining prototype activations in the logic memory. On a subset of ImageNet-X and GQA, we simulate expert feedback by providing alternative rationales for 10% of the training samples. Personalization is implemented by reweighting memory slots  $m_j$  according to user-provided preferences. Table 4 reports explanation quality before and after personalization. While task accuracy remains stable, alignment with user rationales improves significantly. This shows that LVM explanations can be tailored to individual experts or domains without retraining the entire model. We observe a +0.08 improvement in IoU alignment with user attention and a +0.015 gain in BERTScore, along with a notable increase in human ratings of sufficiency and clarity. These results confirm that LVM can incorporate human-in-the-loop feedback to provide domain-adapted explanations while preserving predictive accuracy.

Table 4: Effect of personalization on explanation alignment. Metrics are computed against user-provided rationales on ImageNet-X and GQA.

Setting	Top-1 Acc. (%)	IoU w/ User Masks $\uparrow$	BERTScore $\uparrow$	Human Rating (1-5) $\uparrow$
Pre-personalization	82.6	0.34	0.831	3.5
Post-personalization	82.5	<b>0.42</b>	<b>0.846</b>	<b>4.1</b>

### 4.5 FEDERATED LEARNING FOR PRIVACY-PRESERVING INTERPRETABILITY

Finally, we evaluate LVM in a federated learning setting where data remain decentralized. We simulate  $C = 5$  clients on ChestX-ray14 and KITTI, partitioned non-IID via a Dirichlet distribution ( $\alpha = 0.5$ ). Clients train locally with the joint objective, and a central server aggregates updates by weighted averaging. As shown in Table 5, federated LVM achieves performance and interpretability

432 close to centralized training, with reasoning trajectories and explanations remaining faithful without  
 433 data sharing.  
 434

435 Table 5: Federated vs. centralized training on ChestX-ray14 and KITTI. LVM preserves accuracy  
 436 and interpretability under federated optimization.

Setting	Task Metric $\uparrow$	Faithfulness (IoU) $\uparrow$	BERTScore $\uparrow$	Trajectory Coherence $\uparrow$
Centralized (ChestX-ray14)	79.5 AUC	0.36	0.842	0.78
Federated (ChestX-ray14)	79.2 AUC	0.35	0.839	0.76
Centralized (KITTI)	73.1 mAP	0.34	0.837	0.74
Federated (KITTI)	72.8 mAP	0.33	0.834	0.73

443 The results show only marginal drops ( $< 0.5$  points) compared to centralized training, while main-  
 444 taining faithful reasoning and explanation quality. This highlights that LVM can scale across dis-  
 445 tributed environments, supporting sensitive domains such as healthcare and autonomous driving  
 446 without compromising privacy.  
 447

#### 448 4.6 CAUSAL INTERVENTION ON LOGIC MEMORY

450 We evaluate the causal role of logic memory by intervening on top- $k$  activated prototypes  $\{m_j\}$ . For  
 451 each input, salient slots are masked or permuted, and the impact is measured as  $\Delta\text{Acc} = \text{Acc}(\hat{y}) -$   
 452  $\text{Acc}(\hat{y}')$  and  $\Delta\text{Exp} = \text{BERTScore}(E, E^*) - \text{BERTScore}(E', E^*)$ . As shown in Table 6, masking  
 453 salient slots reduces accuracy by 2–3 points and explanation alignment by more than 0.02, while  
 454 random slot perturbations have minimal effect. This confirms that LVM predictions and explanations  
 455 are grounded in memory-encoded reasoning rather than post-hoc correlations. These results provide  
 456 direct evidence that the reasoning trajectory  $R$  depends on structured memory prototypes, rather  
 457 than being an artifact of post-hoc alignment. Hence, interpretability in LVM is grounded in causal  
 458 contributions from memory, strengthening trust in the model’s explanations.  
 459

460 Table 6: Causal intervention on memory slots. Masking salient prototypes degrades both prediction  
 461 accuracy and explanation quality, confirming the causal role of logic memory.

Setting	Top-1 Acc. (%)	$\Delta\text{Acc} \downarrow$	BERTScore $\uparrow$	$\Delta\text{Exp} \downarrow$
Original LVM	82.6	–	0.847	–
Mask top- $k$ slots	79.7	-2.9	0.823	-0.024
Permute top- $k$ slots	80.1	-2.5	0.826	-0.021
Mask random slots	82.2	-0.4	0.844	-0.003

## 468 5 DISCUSSION AND CONCLUSION

471 We presented the Logic Vision Model (LVM), which unifies prediction, reasoning, and expla-  
 472 nation in a single framework. Unlike post-hoc methods, LVM intrinsically generates reasoning  
 473 trajectories and textual explanations tied to its predictive process. Interpretability is enforced  
 474 through the joint loss  $\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_1 \mathcal{L}_{\text{reason}} + \lambda_2 \mathcal{L}_{\text{exp}} + \lambda_3 \mathcal{L}_{\text{faith}}$ , with a monotonic con-  
 475 straint  $\Delta p_t = p(y \mid r_t) - p(y \mid r_{t-1}) \geq 0$  ensuring that evidence accumulates faithfully. Ex-  
 476 periments show that LVM achieves competitive accuracy while improving both visual faithfulness  
 477 and textual explanation quality. Moreover, the framework supports personalization, where prototype  
 478 weights are adapted via  $\pi_{t,j}^{(u)} \propto \exp(r_{t-1}^\top m_j / \tau + \phi_u(m_j))$ , and federated optimization, where local  
 479 objectives  $\mathcal{L}^{(c)}$  are aggregated as  $\min_{\theta} \sum_{c=1}^C \omega_c \mathcal{L}^{(c)}(\theta)$ , enabling privacy-preserving deployment.  
 480 While effective, LVM still faces limitations such as fixed reasoning length  $T$  and potential ineffi-  
 481 ciency in federated settings. Future work will explore adaptive trajectories, multimodal extensions,  
 482 communication-efficient federated learning, and causal grounding of memory slots. Overall, LVM  
 483 takes a step toward interpretable vision systems where transparency is a design principle rather than  
 484 an afterthought.<sup>1</sup>

485 <sup>1</sup>Upon acceptance, we will publicly release code, pretrained models, and reasoning visualizations to facili-  
 486 tate further research and adoption.

486 REFERENCES  
487

- 488 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
489 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical  
490 report. *arXiv preprint arXiv:2303.08774*, 2023.
- 491 Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, et al. Vqa: Visual question answering. In *Proceed-  
492 ings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- 493 Somil Bansal, Andrea Bajcsy, Ellis Ratner, Anca D Dragan, and Claire J Tomlin. A hamilton-jacobi  
494 reachability-based framework for predicting and analyzing human motion for safe planning. In  
495 *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7149–7155. IEEE,  
496 2020.
- 497 Harel Biggie, Ajay Narasimha Mopidevi, Dusty Woods, and Christoffer Heckman. Tell me where  
498 to go: A composable framework for context-aware embodied robot navigation. *arXiv preprint  
499 arXiv:2306.09523*, 2023.
- 500 Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C Li, Adrien Bardes, Suzanne  
501 Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, et al. An introduction to  
502 vision-language modeling. *arXiv preprint arXiv:2405.17247*, 2024.
- 503 Piyush Borole, Tongjie Wang, Antonio Vergari, and Ajitha Rajan. Can interpretability and accuracy  
504 coexist in cancer survival analysis? *bioRxiv*, pp. 2025–04, 2025.
- 505 Tom B. Brown, Benjamin Mann, Nick Ryder, et al. Language models are few-shot learners. *Ad-  
506 vances in Neural Information Processing Systems*, 2020.
- 507 Aya Hage Chehade, Nassib Abdallah, Jean-Marie Marion, Mathieu Hatt, Mohamad Oueidat, and  
508 Pierre Chauvet. Advancing chest x-ray diagnostics: A novel cyclegan-based preprocessing ap-  
509 proach for enhanced lung disease classification in chestx-ray14. *Computer Methods and Programs  
in Biomedicine*, 259:108518, 2025.
- 510 Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks  
511 like that: deep learning for interpretable image recognition. *Advances in neural information  
512 processing systems*, 32, 2019.
- 513 Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille.  
514 Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and  
515 fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):  
516 834–848, 2017.
- 517 Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo  
518 Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban  
519 scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern  
520 recognition*, pp. 3213–3223, 2016.
- 521 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
522 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
523 pp. 248–255. Ieee, 2009.
- 524 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
525 bidirectional transformers for language understanding. *NAACL-HLT*, 2019.
- 526 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
527 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An  
528 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint  
529 arXiv:2010.11929*, 2020.
- 530 Upol Ehsan and Mark O Riedl. Explainability pitfalls: Beyond dark patterns in explainable ai.  
531 *Patterns*, 5(6), 2024.
- 532 Mohammad Ennab and Hamid Mccheick. Designing an interpretability-based model to explain the  
533 artificial intelligence algorithms in healthcare. *Diagnostics*, 12(7):1557, 2022.

- 540 Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman.  
 541 The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):  
 542 303–338, 2010.
- 543 Clayton Fields and Casey Kennington. Vision language transformers: A survey. *arXiv preprint*  
 544 *arXiv:2307.03254*, 2023.
- 545 Difei Gao, Lei Ji, Luowei Zhou, Kevin Qinghong Lin, Joya Chen, Zihan Fan, and Mike Zheng  
 546 Shou. Assistgpt: A general multi-modal assistant that can plan, execute, inspect, and learn. *arXiv*  
 547 *preprint arXiv:2306.08640*, 2023.
- 548 Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti  
 549 vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*,  
 550 pp. 3354–3361. IEEE, 2012.
- 551 Naofumi Hama, Masayoshi Mase, and Art B Owen. Deletion and insertion tests in regression  
 552 models. *Journal of Machine Learning Research*, 24(290):1–38, 2023.
- 553 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 554 nition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp.  
 555 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- 556 Junlin Hou, Sicen Liu, Yequan Bie, Hongmei Wang, Andong Tan, Luyang Luo, and Hao Chen.  
 557 Self-explainable ai for medical image analysis: A survey and new outlooks. *arXiv preprint*  
 558 *arXiv:2410.02331*, 2024.
- 559 D. Hudson and C. Manning. Gqa: A new dataset for real-world visual reasoning and compositional  
 560 question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 561 *Recognition*, 2019.
- 562 Xavier Alphonse Inbaraj, Charlyn Villavicencio, Julio Jerison Macrohon, Jyh-Horng Jeng, and Jer-  
 563 Guang Hsieh. A novel machine learning approach for tuberculosis segmentation and prediction  
 564 using chest-x-ray (cxr) images. *Applied Sciences*, 11(19):9057, 2021.
- 565 Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and  
 566 Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual  
 567 reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,  
 568 pp. 2901–2910, 2017.
- 569 Rémi Kazmierczak, Eloïse Berthier, Goran Frehse, and Gianni Franchi. Explainability and vision  
 570 foundation models: A survey. *Information Fusion*, 122:103184, 2025.
- 571 Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and  
 572 Percy Liang. Concept bottleneck models. In *International conference on machine learning*, pp.  
 573 5338–5348. PMLR, 2020.
- 574 Alex Krizhevsky, Geoff Hinton, et al. Convolutional deep belief networks on cifar-10. *Unpublished*  
 575 *manuscript*, 40(7):1–9, 2010.
- 576 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convo-  
 577 lutional neural networks. In *Advances in Neural Information Processing Systems 25*, pp. 1097–  
 578 1105. Curran Associates, Inc., 2012.
- 579 Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-  
 580 training for unified vision-language understanding and generation. In *International conference on*  
 581 *machine learning*, pp. 12888–12900. PMLR, 2022.
- 582 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image  
 583 pre-training with frozen image encoders and large language models. In *International conference*  
 584 *on machine learning*, pp. 19730–19742. PMLR, 2023.
- 585 Zongxia Li, Xiyang Wu, Hongyang Du, Fuxiao Liu, Huy Nghiem, and Guangyao Shi. A survey of  
 586 state of the art large vision language models: Alignment, benchmark, evaluations and challenges.  
 587 *arXiv preprint arXiv:2501.02189*, 2025.

- 594 Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization*  
 595 *branches out*, pp. 74–81, 2004.
- 596
- 597 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr  
 598 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European*  
 599 *conference on computer vision*, pp. 740–755. Springer, 2014.
- 600
- 601 Zihao Lin, Samyadeep Basu, Mohammad Beigi, Varun Manjunatha, Ryan A Rossi, Zichao Wang,  
 602 Yufan Zhou, Sriram Balasubramanian, Arman Zarei, Keivan Rezaei, et al. A survey on mechanis-  
 603 tic interpretability for multi-modal foundation models. *arXiv preprint arXiv:2502.17516*, 2025.
- 604
- 605 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances*  
 606 *in neural information processing systems*, 36:34892–34916, 2023.
- 607
- 608 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.  
 609 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the*  
 610 *IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- 611
- 612 Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic  
 613 segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,  
 614 pp. 3431–3440, 2015.
- 615
- 616 Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv*  
 617 *preprint arXiv:1608.03983*, 2016.
- 618
- 619 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*  
 620 *arXiv:1711.05101*, 2017.
- 621
- 622 David Martens, Galit Shmueli, Theodoros Evgeniou, Kevin Bauer, Christian Janiesch, Stefan Feuer-  
 623 riegel, Sebastian Gabel, Sofie Goethals, Travis Greene, Nadja Klein, et al. Beware of “explanations”  
 624 of ai. *arXiv preprint arXiv:2504.06791*, 2025.
- 625
- 626 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic  
 627 evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association*  
 628 *for Computational Linguistics*, pp. 311–318, 2002.
- 629
- 630 Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Dar-  
 631 rell, and Marcus Rohrbach. Multimodal explanations: Justifying decisions and pointing to the  
 632 evidence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
 633 8779–8788, 2018.
- 634
- 635 Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of  
 636 black-box models. *arxiv* 2018. *arXiv preprint arXiv:1806.07421*, 6, 1806.
- 637
- 638 Alec Radford, Jong Wook Kim, Chris Hallacy, et al. Learning transferable visual models from  
 639 natural language supervision. In *Proceedings of the 38th International Conference on Machine*  
 640 *Learning (ICML)*, 2021.
- 641
- 642 Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified,  
 643 real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern*  
 644 *recognition*, pp. 779–788, 2016.
- 645
- 646 Isabel Redtzer. Click me: thumbnail extraction for fashion videos: An approach for selecting en-  
 647 gaging video thumbnails based on clothing identification, sharpness, and contrast., 2023.
- 648
- 649 Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object  
 650 detection with region proposal networks. *Advances in neural information processing systems*, 28,  
 651 2015.
- 652
- 653 Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ” why should i trust you?” explaining the  
 654 predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference*  
 655 *on knowledge discovery and data mining*, pp. 1135–1144, 2016.

- 648 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng  
649 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei.  
650 ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*  
651 (IJCV), 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- 652 Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
653 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based local-  
654 ization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626,  
655 2017.
- 656 Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image  
657 recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- 658 Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Visualising image classification models  
659 and saliency maps. *Deep Inside Convolutional Networks*, 2(2), 2014.
- 660 Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In  
661 *International conference on machine learning*, pp. 3319–3328. PMLR, 2017.
- 662 Jingyi Zhang, Jiaxing Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A  
663 survey. *IEEE transactions on pattern analysis and machine intelligence*, 46(8):5625–5644, 2024.
- 664 Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. Interpretable convolutional neural networks.  
665 In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8827–  
666 8836, 2018.
- 667 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluat-  
668 ing text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
- 669 Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene  
670 parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and*  
671 *pattern recognition*, pp. 633–641, 2017.
- 672
- 673
- 674
- 675
- 676
- 677
- 678
- 679
- 680
- 681
- 682
- 683
- 684
- 685
- 686
- 687
- 688
- 689
- 690
- 691
- 692
- 693
- 694
- 695
- 696
- 697
- 698
- 699
- 700
- 701

702 **A EXPERIMENTAL DETAILS (EXTENDED)**  
703704 **A.1 BENCHMARKS AND PROTOCOLS**  
705

706 We evaluate classification on ImageNet-1k (Deng et al., 2009) and CIFAR-100, detection on MS  
707 COCO (Lin et al., 2014) and PASCAL VOC (Everingham et al., 2010), and segmentation on  
708 Cityscapes (Cordts et al., 2016) and ADE20K (Zhou et al., 2017). Reasoning/explanation are  
709 assessed on CLEVR (Johnson et al., 2017), GQA (Hudson & Manning, 2019), ImageNet-X (?), and  
710 ClickMe (?). We use Top-1/Top-5 for classification, mAP for detection, and mIoU for segmenta-  
711 tion. Visual faithfulness uses Pointing Game, IoU with human masks, Deletion/Insertion AUC (?),  
712 Sufficiency/Necessity, and trajectory coherence defined as the fraction of steps with non-decreasing  
713 confidence  $\frac{1}{T} \sum_t \mathbb{1}[p(y | r_t) - p(y | r_{t-1}) \geq 0]$ . Textual quality uses BLEU (Papineni et al.,  
714 2002), ROUGE-L (Lin, 2004), BERTScore (Zhang et al., 2019), and human ratings (1–5) on cor-  
715 rectness/sufficiency/clarity.

716 **Table 7: Benchmarks summary.**  
717

718 <b>Dataset</b>	719 <b>Task</b>	720 <b>Primary Metric</b>	721 <b>Scale</b>	722 <b>Notes</b>
720 ImageNet-1k	721 Classification	722 Top-1/Top-5	723 1.28M images	724 Standard val split
720 CIFAR-100	721 Classification	722 Top-1	723 60k images	724 100 classes
720 MS COCO	721 Detection	722 mAP@[.5:.95]	723 118k/5k	724 1x schedule
720 PASCAL VOC	721 Detection	722 mAP@0.5	723 07+12/07-test	724
720 Cityscapes	721 Segmentation	722 mIoU	723 3k/500	724 Fine annotations
720 ADE20K	721 Segmentation	722 mIoU	723 20k/2k	724 150 classes
720 CLEVR	721 Reasoning	722 Acc/Program exec	723 100k	724 Synthetic reasoning
720 GQA	721 Reasoning/Explain	722 Acc + Expl. metrics	723 22M QA pairs	724
720 ImageNet-X	721 Rationale	722 IoU/BERTScore	723 1k images	724 Human masks
720 ClickMe	721 Rationale	722 IoU	723 400k maps	724 Human attention

732 **A.2 BASELINES AND REPRODUCTION**  
733

734 Baselines include post-hoc XAI (Grad-CAM (Selvaraju et al., 2017), Integrated Gradients (Sun-  
735 dararajan et al., 2017), RISE, LIME (Ribeiro et al., 2016)), intrinsic models (ProtoPNet (Chen et al.,  
736 2019), Concept Bottleneck (Koh et al., 2020), Neural Module Networks (?)), and VLM justifications  
737 (BLIP-2 (Li et al., 2023), LLaVA (Liu et al., 2023)). All baselines are reproduced with the original  
738 paper settings.

739 **A.3 IMPLEMENTATION DETAILS**  
740

741 All models use PyTorch, trained on a single server with 8×NVIDIA H200 GPUs. Encoders are  
742 ResNet-50 (He et al., 2016) or ViT-B/16 (?). Logic memory size is  $K=100$ , reasoning steps  $T=4$ .  
743 Optimization uses AdamW (?) with cosine decay (Loshchilov & Hutter, 2016), batch size 256, and  
744  $(\lambda_1, \lambda_2, \lambda_3) = (1.0, 0.5, 0.1)$ . Federated experiments simulate  $C=5$  non-IID clients via Dirichlet  
745 partitioning with  $\alpha=0.5$  and weighted averaging every 5 local epochs.

746 **Table 8: Key hyperparameters.**  
747

751 <b>Backbone</b>	752 <b><math>K</math></b>	753 <b><math>T</math></b>	754 <b>Batch</b>	755 <b>Optimizer</b>	756 <b>Base LR</b>	757 <b><math>(\lambda_1, \lambda_2, \lambda_3)</math></b>
751 ResNet-50	752 100	753 4	754 256	755 AdamW	756 1e-3	757 $(1.0, 0.5, 0.1)$
751 ViT-B/16	752 100	753 4	754 256	755 AdamW	756 5e-4	757 $(1.0, 0.5, 0.1)$

756 **B ADDITIONAL EXPERIMENTS**  
757758 **B.1 CAUSAL INTERVENTION ON MEMORY SLOTS**  
759760 To test whether logic memory contributes causally to predictions and explanations, we intervene  
761 on the most salient prototypes. For each input, the top- $k$  activated slots  $\{m_j\}$  are either masked  
762 (set to zero) or permuted (shuffled across examples), and we measure the change in both accu-  
763 racy and explanation quality. Impact is quantified as  $\Delta\text{Acc} = \text{Acc}(\hat{y}) - \text{Acc}(\hat{y}')$  and  $\Delta\text{Exp} =$   
764  $\text{BERTScore}(E, E^*) - \text{BERTScore}(E', E^*)$ , where  $(\hat{y}, E)$  are the original outputs and  $(\hat{y}', E')$  the  
765 outputs after intervention.766 Table 9: Causal intervention on GQA/ImageNet-X (higher is better unless noted). Masking or  
767 permuting salient prototypes produces notable degradation, while random masking has negligible  
768 effect.

Setting	Acc. (%)	$\Delta\text{Acc} \downarrow$	BERTScore	$\Delta\text{Exp} \downarrow$
Original LVM	82.6	–	0.847	–
Mask top- $k$	79.7	-2.9	0.823	-0.024
Permute top- $k$	80.1	-2.5	0.826	-0.021
Mask random	82.2	-0.4	0.844	-0.003

777 As shown in Table 9, interventions on salient slots cause accuracy to drop by up to 3 points and ex-  
778 planation alignment to fall by 0.02–0.03, whereas random masking produces only marginal changes.  
779 This contrast demonstrates that LVM’s predictions and explanations depend directly on specific pro-  
780 tootypes, rather than being post-hoc correlations. In other words, the reasoning trajectory is not  
781 incidental but causally grounded in memory-encoded inference patterns, reinforcing the model’s  
782 interpretability claims.783 **B.2 COUNTERFACTUAL CONSISTENCY**  
784785 We apply controlled perturbations (texture/background/sketch) and compare top- $k$  prototype sets  
786 and explanations between  $(x, E)$  and  $(x', E')$  via Prototype Stability (Jaccard overlap) and Ex-  
787 planation Consistency ( $\text{BERTScore}(E, E')$ ).  
788789 Table 10: Counterfactual consistency (higher is better).  
790

Perturbation	Prototype Stability	Explanation Consistency
Texture shift	0.72	0.84
Background change	0.75	0.86
Sketch transform	0.70	0.82

797 Stability remains high under mild perturbations, indicating robust prototype selection and rationales.  
798799 **B.3 SANITY CHECKS FOR EXPLANATIONS**  
800801 A faithful explanation should degrade as the model itself is randomized, rather than remaining un-  
802 changed. We progressively randomize either model weights (from shallow to deep layers) or training  
803 labels, and compute Kendall- $\tau$  correlation between randomization strength and degradation of in-  
804 terpretability metrics. If explanations are genuinely tied to the model’s internal reasoning, metrics  
805 such as IoU and BERTScore should decrease monotonically as randomization increases. Conversely,  
806 post-hoc artifacts would remain spuriously stable regardless of underlying model corruption.807 As shown in Table 11, LVM achieves higher  $\tau$  values, meaning its explanations degrade more con-  
808 sistently with model corruption. This behavior rules out spurious stability and confirms that LVM  
809 explanations are intrinsically linked to the predictive process, unlike post-hoc methods that may  
remain unaffected even when the model is randomized.

810  
811 Table 11: Sanity checks with progressive randomization (higher  $\tau$  indicates stronger monotonic  
812 degradation).

Method	$\tau$ (IoU vs. randomization)	$\tau$ (BERTScore vs. randomization)
Grad-CAM	0.41	0.38
ProtoPNet	0.55	0.49
<b>LVM (ours)</b>	<b>0.68</b>	<b>0.62</b>

813  
814  
815  
816  
817  
818  
819 **B.4 CALIBRATION AND SELECTIVE PREDICTION**  
820

821 Interpretability should not only explain predictions but also reflect when the model is confident or  
822 uncertain. To test this, we correlate trajectory-level faithfulness with predictive reliability. Specifically,  
823 we compute Expected Calibration Error (ECE; lower is better) and the area under coverage-  
824 risk curves (higher is better). Coverage-risk curves are obtained by ranking predictions by faithfulness  
825 scores (trajectory coherence  $\frac{1}{T} \sum_t \mathbb{1}[p(y | r_t) - p(y | r_{t-1}) \geq 0]$ ) and reporting accuracy at  
826 varying coverage thresholds. A faithful system should achieve low calibration error and maintain  
827 high accuracy even when selectively abstaining on uncertain cases.

828  
829 Table 12: Calibration and selective prediction (lower ECE, higher AUC).

Method	ECE $\downarrow$	Coverage–Risk AUC $\uparrow$
ResNet+Grad-CAM	0.082	0.61
ProtoPNet	0.074	0.64
<b>LVM (ours)</b>	<b>0.051</b>	<b>0.71</b>

830  
831  
832  
833  
834  
835  
836 As shown in Table 12, LVM achieves lower calibration error and higher AUC than baselines. This  
837 indicates that reasoning trajectories not only improve interpretability but also serve as reliable indicators  
838 of prediction trustworthiness, enabling selective prediction in safety-critical scenarios where  
839 calibrated confidence is essential.

840  
841  
842 **B.5 SCALABILITY AND EFFICIENCY**  
843

844 We study how the number of reasoning steps  $T$  trades off performance and compute. Unless noted,  
845 we fix the backbone (ViT-B/16), input resolution (224), memory size  $K=100$ , and batch size 1;  
846 latency is measured end-to-end in FP16 on a single H200 and averaged over 1,000 images without  
847 I/O. Table 13 shows that accuracy and interpretability improve with larger  $T$ , while latency grows  
848 approximately linearly.

849  
850 Table 13: Reasoning steps vs. performance and latency.

$T$	Acc. (%)	IoU	BERTScore	Latency (ms)
2	80.3	0.33	0.833	42
4	82.1	0.36	0.842	68
6	82.6	0.38	0.847	95

851  
852  
853  
854  
855  
856  
857 From a computational standpoint, the per-step complexity is dominated by attention between  $N$   
858 visual tokens and  $K$  memory slots; with token/channel size  $d$ , projections and attention scale as  
859  $\mathcal{O}(Nd^2 + NKd)$  per step, so total cost grows as  $\mathcal{O}(T)$  when  $N$ ,  $K$ , and  $d$  are fixed. In practice we  
860 cache  $(K, V)$  from the memory retrieval and the fused representation  $Z$  so that only the transition  
861 operator  $\mathcal{T}$  is applied at each step; this makes latency increase close to linear in  $T$ , as reflected in Ta-  
862 ble 13. Memory footprint scales as  $\mathcal{O}(Nd + Kd + NK)$  due to token/slot embeddings and attention  
863 weights; with  $K=100$  and  $N$  in the order of a few hundred patches, VRAM remains bounded for the

864 reported settings. To exploit diminishing returns, we also consider early-exit policies driven by the  
 865 trajectory statistics. A confidence stabilization rule terminates when  $|p(y | r_t) - p(y | r_{t-1})| < \epsilon$  for  
 866  $S$  consecutive steps, and a coherence rule stops when  $\frac{1}{t} \sum_{i=1}^t \mathbb{1}[p(y | r_i) - p(y | r_{i-1}) \geq 0] \geq \tau$ . In  
 867 our setting, modest thresholds (e.g.,  $\epsilon$  in the  $10^{-3}$ – $10^{-2}$  range,  $S=2$ ,  $\tau \approx 0.8$ ) preserve accuracy and  
 868 explanation quality while skipping late steps on easy inputs. Because  $Z$  and memory keys/values  
 869 are cached, early exit reduces latency nearly proportionally to the number of omitted steps. Overall,  
 870 performance saturates around  $T=6$  while  $T=4$  offers a favorable cost–benefit point. For real-time  
 871 or interactive deployments we recommend  $T=4$  with early exit, whereas offline analysis can use  
 872  $T=6$  to maximize interpretability metrics (IoU and BERTScore) with manageable compute.  
 873

## 874 B.6 GENERALIZATION AND TRANSFER

875 To examine whether reasoning prototypes learned by LVM are reusable across domains, we test  
 876 transfer to ImageNet-S under different freezing strategies. In the first setting, we freeze the logic  
 877 memory  $M$  and tune only the vision encoder and prediction head. In the second, we freeze the  
 878 encoder and fine-tune the memory slots. Finally, we allow joint fine-tuning of both modules.  
 879

880  
 881 Table 14: Transfer to ImageNet-S under different freezing strategies.  
 882

883 Setup	884 Acc. (%)	885 IoU	886 BERTScore
885 Encoder tuned, memory frozen	886 75.4	887 0.31	888 0.832
886 Memory tuned, encoder frozen	887 72.9	888 0.29	889 0.829
887 Joint fine-tuning	888 76.1	889 0.32	890 0.835

890 As shown in Table 14, freezing the memory while adapting the encoder yields only a small drop  
 891 compared to joint fine-tuning, whereas freezing the encoder hurts more. This suggests that proto-  
 892 types stored in memory are relatively domain-general and reusable, while the encoder benefits from  
 893 adaptation to new data distributions. Thus, the memory acts as a library of transferable reasoning  
 894 patterns that can generalize across tasks.  
 895

## 896 B.7 HUMAN-IN-THE-LOOP ADAPTATION

897 We further test whether LVM can be refined through expert feedback. In this setting, annotators re-  
 898 view a small fraction of samples (10%) and provide corrections to either predictions or explanations.  
 899 Feedback is used to reweight or refine the most relevant prototypes  $m_j$ , without retraining the full  
 900 model. We evaluate performance after one and two feedback rounds.  
 901

902  
 903 Table 15: Effect of human feedback rounds on GQA.  
 904

905 Rounds	906 Acc. (%)	907 IoU	908 BERTScore
908 0	909 82.6	910 0.34	911 0.847
909 1	910 83.4	911 0.38	912 0.852
910 2	911 84.1	912 0.40	913 0.856

914 Table 15 shows consistent gains in both accuracy and interpretability with only minimal supervi-  
 915 sion. One round of feedback improves accuracy by nearly one point and IoU by 0.04, while two  
 916 rounds provide further gains. This demonstrates that LVM can be adapted efficiently in real-world  
 917 deployments, where domain experts can guide prototype reweighting to align with domain-specific  
 918 reasoning styles.  
 919

918 B.8 TEMPORAL SMOOTHNESS OF REASONING VIDEOS  
919

920 Beyond static attention quality, interpretability also depends on how reasoning evolves across steps.  
921 We measure temporal smoothness of attention maps using total variation  $\sum_t \|A_t - A_{t-1}\|_1$ , where  
922 lower values indicate smoother transitions and more coherent reasoning videos. This metric reflects  
923 whether the model gradually accumulates evidence rather than oscillating between unrelated regions.

924 Table 16: Temporal smoothness of trajectories (lower TV is better).  
925

Method	TV ↓
Grad-CAM	0.42
ProtoPNet	0.36
<b>LVM (ours)</b>	<b>0.28</b>

933 As shown in Table 16, LVM achieves the lowest temporal variation, indicating that its reasoning  
934 trajectory evolves in a stable, human-like manner. In qualitative inspection, this corresponds to  
935 smoother attention videos, where the model consistently focuses and refines evidence rather than  
936 abruptly shifting.

937 B.9 HYPERPARAMETER SENSITIVITY  
938

939 We also study robustness to hyperparameters by varying the number of prototypes  $K$  and the inter-  
940 pretability loss weights  $(\lambda_1, \lambda_2)$ . These control the richness of the reasoning library and the relative  
941 importance of interpretability in optimization.

942 Table 17: Sensitivity to  $K$  and interpretability weights.  
943

Config	Acc. (%)	IoU	BERTScore	Coherence
$K=50, (0.5, 0.5)$	82.1	0.34	0.842	0.73
$K=100, (1.0, 1.0)$	82.6	0.38	0.847	0.81
$K=200, (2.0, 2.0)$	82.8	0.39	0.849	0.83

951 Results in Table 17 show that larger prototype sets and stronger interpretability weights steadily  
952 improve IoU, BERTScore, and trajectory coherence, with only minor changes in accuracy. This  
953 suggests that interpretability can be enhanced without sacrificing task performance, and that LVM  
954 remains stable across a wide hyperparameter range.

955 B.10 PRIVACY-PRESERVING FEDERATED EXTENSIONS  
956

958 In federated settings, privacy preservation is often a key requirement. To test whether interpretability  
959 can be retained under differential privacy (DP), we add Gaussian noise during aggregation and vary  
960 the privacy budget  $\epsilon$ . Lower  $\epsilon$  provides stronger privacy but injects more noise into parameter  
961 updates.

962 Table 18: Federated learning with differential privacy on ChestX-ray14.  
963

Privacy	AUC (↑)	IoU (↑)	BERTScore (↑)	Coherence (↑)
No DP	79.6	0.36	0.841	0.78
$\epsilon=5$	78.8	0.35	0.837	0.76
$\epsilon=1$	77.2	0.33	0.828	0.72

970 As shown in Table 18, moderate privacy budgets ( $\epsilon=5$ ) retain most of the accuracy and interpretabil-  
971 ity, while extreme noise ( $\epsilon=1$ ) leads to noticeable drops in both task performance and explanation

972 quality. This highlights a fundamental privacy–interpretability trade-off, but also suggests that practical  
 973 privacy guarantees can be achieved without severely harming the faithfulness of reasoning  
 974 trajectories.

### 976 B.11 CONCEPT GROUNDING OF PROTOTYPES

978 For interpretability to be useful, prototypes in the memory should correspond to meaningful and  
 979 human-recognizable concepts. We therefore visualize nearest neighbors for each  $m_j$  and compute  
 980 CLIP-based similarity to textual concept labels. Two quantitative measures are reported: *concept*  
 981 *purity* (the proportion of nearest neighbors belonging to the same class) and *nameability* (average  
 982 human rating from 1–5 on whether the prototype corresponds to a recognizable concept).

983  
 984 Table 19: Prototype grounding comparison.

Metric	ProtoPNet	<b>LVM</b>
Concept Purity	0.62	<b>0.74</b>
Nameability (1–5)	3.1	<b>4.0</b>

990 Results in Table 19 show that LVM prototypes achieve higher purity and are more easily nameable  
 991 by humans. This suggests that the logic memory encodes reasoning units aligned with semantic  
 992 concepts, making the explanations not only faithful but also more accessible to end users.

### 994 B.12 FAILURE ANALYSIS

996 Finally, we analyze when and why LVM fails. We focus on reasoning steps where confidence  
 997 decreases, i.e.,  $p(y | r_t) - p(y | r_{t-1}) < 0$ , and inspect the associated prototypes. Failures are  
 998 categorized into cases with or without spurious prototype activations (irrelevant or misleading slots  
 999 being activated early in the trajectory).

1000  
 1001 Table 20: Failure breakdown by spurious activation.

Condition	Error Rate (%)	Mean failure step
No spurious activation	12.3	3.7
With spurious activation	28.7	2.1

1008 Table 20 shows that errors are concentrated when spurious prototypes are activated early, leading to  
 1009 incorrect reasoning trajectories that cascade into wrong predictions. This analysis suggests clear di-  
 1010 rections for improvement, such as refining prototype selection, regularizing early steps of reasoning,  
 1011 or incorporating causal constraints to suppress misleading activations.