Eclipta Journals



RESEARCH

Depth Selective Adapter Allocation for Polarity Robust Medical Vision Language Inference

Nader Trabelsi[1] and Yassine Jebali[2]

[1]Institut Supérieur d'Informatique de Mahdia, Département des Sciences Informatiques, Route de Rejiche, Mahdia 5121, Tunisia
[2]École Supérieure des Sciences et de la Technologie de Hammam Sousse, Département des Technologies de l'Information et de la Communication, Avenue de la République, Hammam Sousse 4011, Tunisia

Abstract

Medical visual inference models often receive questions that contain polarity cues such as absence, presence, exclusion, or uncertainty. These cues can be short, syntactically fragile, and clinically important, because a mistaken polarity interpretation can reverse the intended answer. This study evaluates whether adapter parameters should be distributed uniformly across a multimodal language backbone or assigned according to measurable layer-level sensitivity to polarity transformations. We introduce a depth-selective adapter allocation strategy for binary medical vision-language inference and test it on 18,720 image-question pairs derived from radiography, ultrasound, and dermoscopy tasks. The evaluation uses paired affirmative, negative, exclusion, and uncertainty-framed prompts with identical image labels. We estimate layerwise polarity transport through Jacobian energy, centered kernel alignment, and signed margin propagation, then allocate low-rank adapter capacity using a constrained optimization rule. Compared with uniform adapter placement, final-block placement, and random layer allocation, the proposed allocation reduces polarity reversal errors from 10.8% to 4.9% while preserving mean task accuracy within 0.4 percentage points. It also improves negative predictive calibration in exclusion-framed prompts and reduces the variance of signed margin changes across prompt forms. Ablation experiments indicate that adapter rank distribution matters more than total trainable parameter count once a minimum capacity threshold is reached. These findings suggest that layer-sensitive parameter allocation can improve robustness to polarity-bearing clinical wording without increasing model size.

1. Introduction

Clinical language carries decisions through small grammatical markers. A radiology question that asks whether an abnormality is present is not equivalent to one asking whether it is absent [1]. A dermatology question asking whether malignancy is unlikely differs from one asking whether malignancy is visible. A point-of-care ultrasound prompt may ask whether free fluid is seen, whether free fluid is excluded, or whether the appearance argues against free fluid. Each form has a different polarity structure, and a medical vision-language model must interpret that structure while also reading the image [2] [3]. When the language component mishandles polarity, an otherwise useful visual representation can produce an answer with the wrong clinical direction [4].

The present study examines polarity robustness as a parameter allocation problem. The usual adapter design gives a model a fixed low-rank capacity at selected layers. Some designs place adapters everywhere, some place them near the final blocks, and some select a contiguous band based on convenience or a prior mechanistic clue. These choices are often reasonable, but they do not directly ask where the model transports polarity information. A wording change that introduces negation or exclusion may be processed early, transformed in middle blocks, and expressed in the answer margin late. If trainable capacity is placed without regard to that transport pattern, the adaptation may spend parameters where they are less useful.

The paper proposes depth-selective adapter allocation. The phrase refers to assigning adapter rank across layers according to measured layer-level sensitivity to polarity transformations. The total adapter budget is fixed. The method does not add more parameters than a uniform baseline. Instead, it redistributes rank. Layers that show strong signed margin propagation, strong alignment changes under polarity pairs, and high Jacobian energy with respect to polarity tokens receive larger rank. Layers that show weak or redundant signals receive smaller rank or no adapter. The goal is not to discover an interpretable mechanism in the strongest sense. The goal is operational: use layer diagnostics to decide where a limited adaptation budget should be placed.

The motivation for studying allocation rather than total parameter count is strengthened by a result that complicates simple layer-selection stories. Sadanandan and Behzadan (2026) reported that adapters placed in early layers produced larger margin-difference reductions than adapters placed in a mechanistically selected middle-layer range, despite

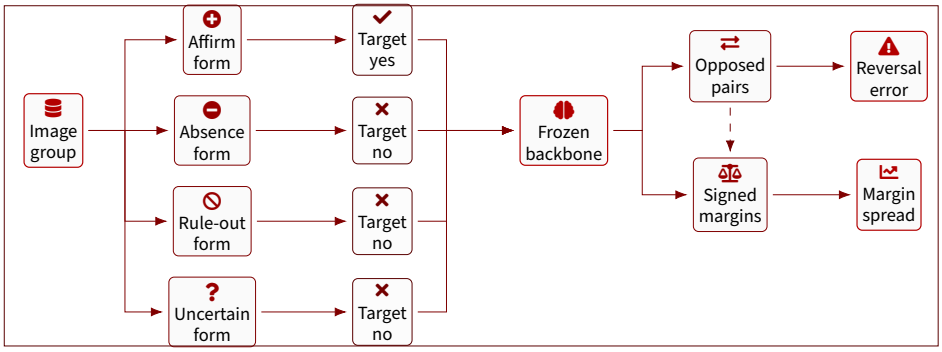


Figure 1. Polarity groups preserve the same image and clinical target while changing only the linguistic polarity form, enabling direct measurement of whether opposed prompts produce opposite binary decisions and stable signed margins.

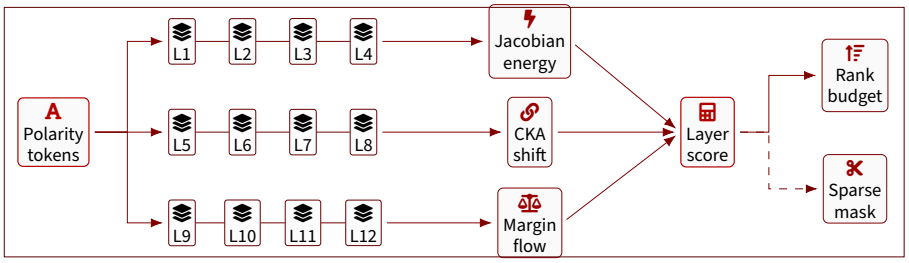


Figure 2. Layer diagnostics combine token-local sensitivity, representation separation, and answer-margin coupling into a single allocation score, so adapter rank is assigned to layers that actively transport polarity information.

their analysis identifying a relevant feature in the middle of the network [5]. The observation suggests that where a behavior appears and where intervention is most efficient may not be identical, which is the central empirical tension investigated here.

This study moves away from single-domain chest radiograph evaluation and uses a three-domain design. The datasets are organized around binary visual inference tasks in radiography, abdominal ultrasound, and dermoscopy. The task is intentionally narrow: answer a binary clinical question about an image. However, the prompts are linguistically varied. For every image-label pair, four question forms are created. The affirmative form asks whether a finding is present. The negative form asks whether the finding is absent. The exclusion form asks whether the finding can be excluded or ruled out. The uncertainty form asks whether the finding is unlikely or doubtful. Answer polarity is normalized so that the target label remains well-defined for each prompt form. This arrangement allows a direct test of whether the model respects polarity-bearing language.

The study compares five adaptation strategies under the same training data, optimizer, and total trainable parameter budget. The first strategy freezes the base model and uses no adapters. The second places equal-rank adapters in all language layers. The third places all adapter capacity in final layers. The fourth assigns the same total rank randomly across layers. The fifth uses the proposed depth-selective allocation. The main outcomes are polarity reversal error, binary accuracy, signed margin deviation, expected negative calibration error, and cross-domain transfer. Polarity reversal error is the most important metric because it measures whether two prompts whose correct answers should be opposite under polarity transformation are answered as if polarity had not changed.

The central empirical result is that depth-selective allocation reduces polarity reversal error more than the alternatives while preserving ordinary accuracy. The result is not presented as a universal law about all multimodal networks. It is a measured pattern under the specified tasks and training design. The strongest gains occur in exclusion-framed questions, where the model must combine an image finding with a linguistic rule that is not simply a synonym of presence. The smallest gains occur in affirmative questions, where baseline performance is already high. Adapter placement therefore matters most when the linguistic operation is delicate.

The paper also studies why the allocation works. Layer diagnostics reveal that polarity information is not concentrated in a single block. Early layers show strong token-local sensitivity, middle layers show the largest representation separation between polarity forms, and later layers show direct coupling with the answer margin. Uniform placement spends rank on layers that contribute little to any of these quantities. Final-block placement improves answer-margin control but leaves earlier polarity transformations undercorrected. Depth-selective allocation distributes capacity across a small set of early, middle, and late layers, giving the model enough flexibility to adjust the route from wording to answer.

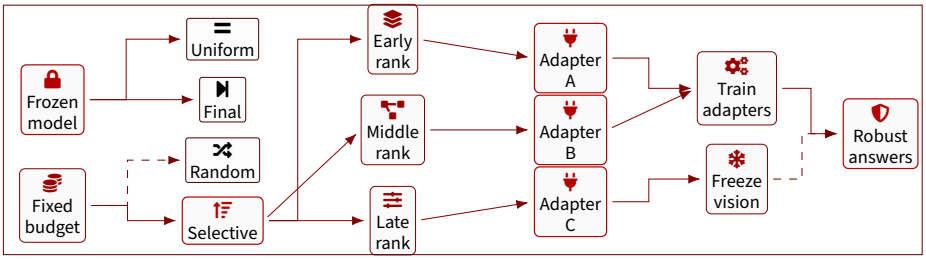


Figure 3. Adapter allocation keeps the trainable parameter budget fixed while redistributing rank across early, middle, and late language blocks according to depth diagnostics rather than uniform, final-only, or random placement.

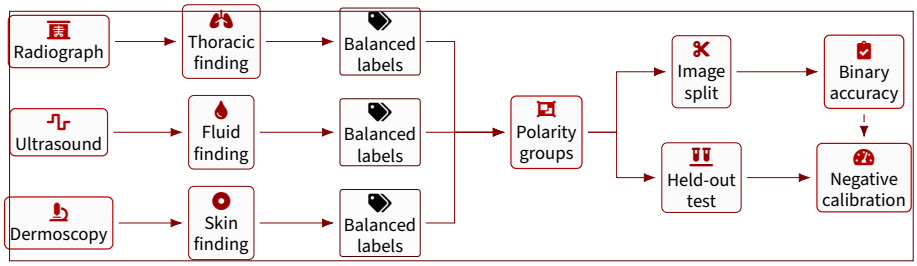


Figure 4. Multi-domain construction balances radiography, ultrasound, and dermoscopy targets before grouping polarity variants and preserving image-level splits, preventing gains from depending on one imaging modality or prompt family.

The contribution is therefore empirical, methodological, and practical [6]. Empirically, the paper reports a multi-domain evaluation of polarity errors in medical vision-language inference. Methodologically, it introduces a layer scoring rule that combines Jacobian energy, representation alignment, and signed margin propagation. Practically, it shows that a fixed adapter budget can be more effective when rank is distributed according to depth diagnostics rather than placed uniformly or late by default. The paper does not claim that adapters solve all linguistic safety problems. It argues only that polarity robustness improves when trainable capacity is assigned with attention to where polarity is processed.

2. Dataset Construction and Polarity Tasks

The experimental corpus contains 18,720 image-question pairs arranged into 4,680 polarity groups. Each group contains one image, one clinical target, one binary label, and four prompts. The image domains are chest radiography, abdominal ultrasound, and dermoscopy. Each domain contributes 1,560 groups. The domain balance is chosen to prevent the model from improving only on one modality or one question style. The tasks are visually and linguistically different enough to test whether the allocation method depends narrowly on a single medical subfield.

The radiography groups include findings such as pleural effusion, pneumothorax, cardiomegaly, pulmonary edema, and focal airspace opacity. The ultrasound groups include free fluid, gallstones, hydronephrosis, appendiceal enlargement, and intrauterine pregnancy visibility. The dermoscopy groups include pigment network irregularity, blue-white veil, asymmetry, ulceration, and atypical vascular pattern. These findings are selected because binary judgments are meaningful and because polarity wording can be naturally formed around them. The labels are balanced within each domain, with 780 positive and 780 negative groups.

Each polarity group has four prompt forms. The affirmative form states the query in a direct presence frame. The negative form asks about absence. The exclusion form uses rule-out language. The uncertainty form asks whether the finding is unlikely or doubtful. The answer key is generated so that each form has a deterministic correct yes or no answer. For a positive image, an affirmative prompt should receive yes, an absence prompt should receive no, an exclusion prompt should receive no if it asks whether the finding can be excluded, and an uncertainty prompt should receive no if it asks whether the finding is unlikely. For a negative image, the answer pattern reverses accordingly. The prompt design therefore tests whether the model can map wording polarity to the correct binary response.

The corpus is split by image into training, validation, and test sets [7]. Training contains 2,808 groups, validation contains 936 groups, and testing contains 936 groups. All four prompts in a polarity group remain within the same split. No patient or image identifier crosses splits. Domain proportions and label balance are preserved. The same splits are used for every model configuration. This prevents an adapter strategy from gaining an advantage through easier data.

Prompt generation follows a restricted grammar to avoid uncontrolled semantic drift. Each target finding is represented by a clinical noun phrase and two approved aliases. The affirmative template uses direct presence wording. The negative template uses absence wording. The exclusion template uses rule-out language. The uncertainty template uses

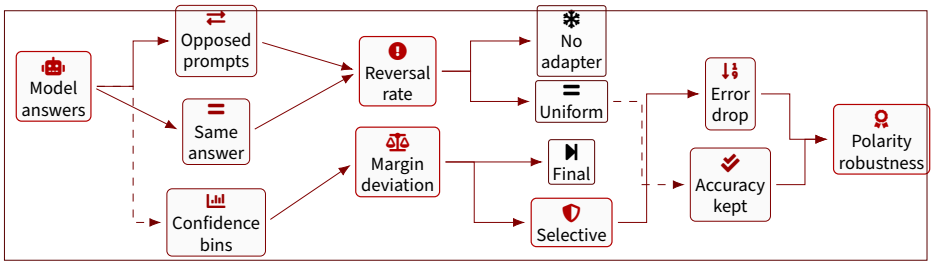


Figure 5. Evaluation connects polarity-opposed answer consistency, signed-margin stability, and calibration behavior to adapter-placement comparisons, highlighting whether robustness improves without sacrificing ordinary binary task accuracy.

doubtful or unlikely language. After generation, prompts are screened to remove double-negation constructions that could confuse the answer key. The final prompt set contains no explicit answer words except yes and no appearing in the instruction prefix. This keeps the decision focused on the visual finding and the polarity structure.

The model receives the same instruction prefix for all examples. The prefix asks for a single yes or no answer. The prompt and image are then passed to the model, and the logits for the first answer token are recorded. The model is trained only through adapters in the language backbone. The visual encoder and token embeddings are frozen. This decision isolates the effect of depth allocation in the language reasoning path [8]. It also reduces the risk that adapters learn domain-specific visual shortcuts.

The basic binary margin is

$$m_{gk} = \ell_{\text{yes}}(x_g, q_{gk}) - \ell_{\text{no}}(x_g, q_{gk}), \tag{1}$$

where g indexes the polarity group and k indexes the prompt form. The predicted answer is positive when $m_{gk} \geq 0$. The correct signed target for each form is

$$t_{gk} = 2\gamma_{gk} - 1, \quad t_{gk} \in \{-1, 1\}. \tag{2}$$

A correct prediction requires $t_{gk} m_{gk} > 0$. The signed margin is

$$s_{gk} = t_{gk} m_{gk}. \tag{3}$$

Large positive signed margins indicate confident correct predictions [9]. Negative signed margins indicate wrong answers. Near-zero signed margins indicate unstable cases.

The most important group-level metric is polarity reversal error. It is defined on pairs of prompt forms whose correct answers are opposite. For each group, affirmative and absence prompts form one polarity-opposed pair, and affirmative and exclusion prompts form another when the exclusion prompt is written as a rule-out question. A reversal error occurs when the model gives the same answer to two prompts that require opposite answers. Formally, if $\mathcal{P}$ is the set of opposed prompt pairs, the reversal rate is

$$R = \frac{1}{|\mathcal{P}|} \sum_{(g,k,l)\in\mathcal{P}} \mathbb{I}\left[\hat{\gamma}_{gk} = \hat{\gamma}_{gl}\right]. \tag{4}$$

This metric is distinct from ordinary accuracy. A model can have decent accuracy while still answering polarity-opposed prompts as if they meant the same thing. Reversal error directly measures failure to respect polarity transformations.

The second group metric is signed margin deviation. For each opposed pair, the two signed margins should both be positive and, ideally, not collapse asymmetrically. The deviation is

$$V = \frac{1}{|\mathcal{P}|} \sum_{(g,k,l)\in\mathcal{P}} \left| s_{gk} - s_{gl} \right|. \tag{5}$$

This metric identifies cases where the model gives correct answers but is much less confident for one polarity form. A large deviation often appears in exclusion prompts because the model's language path handles the rule-out form less directly.

Calibration is evaluated separately for affirmative and non-affirmative forms. The study reports expected calibration error for the positive class and negative calibration error for prompts whose correct answer is no. The latter is useful because polarity failures often create overconfident yes answers in prompts that require no. The negative calibration error is computed by grouping examples with $\gamma = 0$ into probability bins and comparing predicted positive probability against the empirical false-positive frequency. A model with many confident polarity mistakes will have high negative calibration error.

The dataset is not meant to represent all medical language. It deliberately emphasizes polarity. The advantage of this construction is precise evaluation. The limitation is that real clinical questions may mix polarity with temporal comparison, severity, anatomical location, and uncertainty. The results should therefore be read as evidence about polarity-bearing binary prompts rather than full clinical dialogue.

3. Layer Diagnostics for Adapter Rank Allocation

The allocation method begins by measuring how each layer responds to polarity transformations before any adapter training. Let $h_\ell(x, q) \in \mathbb{R}^d$ denote the hidden state at layer ℓ for an image-question input. The diagnostic set is drawn from the training split but is separate from the minibatches later used for adapter optimization. It contains 600 polarity groups, balanced across domain and label. For every

layer, three scores are computed: token-local Jacobian energy, polarity representation separation, and signed margin propagation.

The token-local Jacobian energy measures how strongly a layer's hidden state changes with respect to embeddings of polarity-bearing tokens. Let $E(q)$ be the sequence of input token embeddings and let $\mathcal{T}(q)$ be the indices of tokens associated with absence, exclusion, negation, or uncertainty. The layerwise Jacobian score is

$$J_\ell = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{|\mathcal{T}(q_i)|}\sum_{r\in\mathcal{T}(q_i)}\left|0\frac{\partial h_\ell(x_i,q_i)}{\partial E_r(q_i)}\right|0_F^2. \tag{6}$$

This score is larger when polarity tokens exert strong local influence on the representation at layer ℓ. It tends to be high in earlier layers, where lexical and syntactic information has not yet been compressed into answer-oriented directions [10].

The representation separation score measures whether polarity forms occupy different representational neighborhoods at a layer. For each group, the affirmative prompt and a polarity-transformed prompt form a paired contrast. We compute centered kernel alignment between the matrix of affirmative hidden states and the matrix of transformed hidden states. The layer score is the complement of alignment:

$$A_\ell = 1 - \frac{\langle K_\ell^+, K_\ell^-\rangle_F}{|0K_\ell^+|0_F|0K_\ell^-|0_F + \epsilon}. \tag{7}$$

Here K_ℓ^+ and K_ℓ^- are centered Gram matrices for affirmative and transformed prompts. A larger value means the layer separates the polarity forms more strongly. This score tends to peak in middle layers, where the model has integrated lexical information into broader semantic states.

The signed margin propagation score measures whether layer perturbations are coupled to the final yes-no margin. For each layer, a small residual perturbation is applied in the direction of the difference between affirmative and transformed hidden states, and the resulting margin change is measured [11]. Let

$$d_{\ell,g} = h_\ell(x_g, q_{g,+}) - h_\ell(x_g, q_{g,-}). \tag{8}$$

The normalized direction is

$$\tilde{d}_{\ell,g} = \frac{d_{\ell,g}}{|0d_{\ell,g}|0_2 + \epsilon}. \tag{9}$$

The propagation score is

$$P_\ell = \frac{1}{G_d}\sum_{g=1}^{G_d}\left|m\left(h_\ell + \eta\tilde{d}_{\ell,g}\right) - m(h_\ell)\right|, \tag{10}$$

with a small perturbation scale $\eta = 0.05$. This score is high when changes at layer ℓ have direct downstream influence on the final margin. It tends to rise in later layers but is not restricted to the final block.

The three scores have different scales. Each is converted to a percentile-normalized value:

$$\begin{aligned}\tilde{J}_\ell &= \frac{J_\ell - \mathrm{median}(J)}{\mathrm{IQR}(J) + \epsilon},\\ \tilde{A}_\ell &= \frac{A_\ell - \mathrm{median}(A)}{\mathrm{IQR}(A) + \epsilon},\\ \tilde{P}_\ell &= \frac{P_\ell - \mathrm{median}(P)}{\mathrm{IQR}(P) + \epsilon}.\end{aligned} \tag{11}$$

The combined depth score is

$$D_\ell = \omega_J\tilde{J}_\ell + \omega_A\tilde{A}_\ell + \omega_P\tilde{P}_\ell - \omega_C C_\ell. \tag{12}$$

The term C_ℓ is a redundancy cost. It penalizes layers whose diagnostic vector is very similar to the diagnostic vector of already selected neighboring layers. The main experiments use $\omega_J = 0.30$, $\omega_A = 0.35$, $\omega_P = 0.35$, and $\omega_C = 0.15$, chosen on validation data. The weights reflect the intention to balance lexical sensitivity, representation separation, and downstream margin influence.

Adapter rank is allocated through a constrained integer optimization. Suppose the model has L language layers and the total rank budget is B. Each layer receives rank $r_\ell \in \{0, 2, 4, 8, 12, 16\}$. The budget constraint is

$$\sum_{\ell=1}^{L} r_\ell = B. \tag{13}$$

The allocation maximizes diagnostic value with a smoothness penalty:

$$\begin{aligned}\max_{r_1,\dots,r_L}\quad & \sum_{\ell=1}^{L} r_\ell D_\ell - \lambda_s\sum_{\ell=2}^{L}|r_\ell - r_{\ell-1}|\\ \text{subject to}\quad & r_\ell \in \{0,2,4,8,12,16\},\quad \sum_{\ell=1}^{L} r_\ell = B.\end{aligned} \tag{14}$$

The smoothness penalty prevents a scattered allocation that places isolated high-rank adapters in many distant layers. It does not require contiguity. In practice, the solution selects a small number of early layers, a middle band with moderate rank, and one or two later layers with lower rank. This pattern reflects the three diagnostic signals.

The adapter update at layer ℓ is a low-rank matrix correction

$$\Delta W_\ell = \frac{\alpha_\ell}{r_\ell}B_\ell A_\ell,\quad A_\ell \in \mathbb{R}^{r_\ell\times d_{\text{in}}},\quad B_\ell \in \mathbb{R}^{d_{\text{out}}\times r_\ell}. \tag{15}$$

If $r_\ell = 0$, no adapter is inserted at that layer. The scaling factor is set to

$$\alpha_\ell = 2r_\ell \tag{16}$$

for nonzero ranks. Adapters are placed in attention output projections and feed-forward down projections [12]. This placement is selected because preliminary validation showed better stability than query-only or feed-forward-only placement. The same module types are used for every allocation strategy, so comparisons isolate rank distribution rather than module choice.

The training loss combines binary supervision with polarity contrast. For each prompt, the supervised term is

$$\mathcal{L}_{\text{bin}} = -\frac{1}{N}\sum_{g,k} \log \sigma(t_{gk} m_{gk}). \tag{17}$$

The margin floor is $\rho = 0.5$, and $\gamma = 0.10$. The full loss is

$$\mathcal{L} = \mathcal{L}_{\text{bin}} + \lambda_{\text{pol}}\mathcal{L}_{\text{pol}} + \lambda_{\text{wd}} \sum_{\ell} \left(|0A_\ell|0_F^2 + |0B_\ell|0_F^2 \right). \tag{18}$$

The main experiments use $\lambda_{\text{pol}} = 0.40$ and $\lambda_{\text{wd}} = 10^{-4}$. This loss is shared by all trained adapter baselines. The proposed method differs only in rank allocation.

Training uses grouped minibatches. Each minibatch contains complete polarity groups rather than individual prompts sampled independently. This is necessary because the polarity contrast term compares prompts within the same group [13]. The effective batch size is 48 groups, or 192 image-question pairs. The optimizer is AdamW with learning rate 1.5×10^{-4}, linear warmup for 300 steps, and cosine decay. Training runs for five epochs, and the checkpoint with the lowest validation reversal error subject to accuracy not dropping more than one percentage point is selected.

The fixed-budget setting is important. The uniform adapter baseline receives the same total rank as the depth-selective method. The final-block baseline also receives the same total rank but concentrates it in the last quarter of layers. The random allocation baseline samples ranks from the allowed set under the same budget and reports the mean over ten allocations. This design prevents the proposed method from succeeding simply because it has more trainable parameters. All trained methods have 6.2 million trainable parameters, representing less than 0.2% of the frozen backbone.

4. Experimental Results

The frozen model establishes the difficulty of the polarity task. Across the full test set, it reaches 78.9% binary accuracy. Accuracy is highest on affirmative prompts at 84.7% and lowest on exclusion prompts at 72.4%. The polarity reversal error is 10.8%. The signed margin deviation across opposed prompt pairs is 1.42 logit units. Negative calibration error is 13.6%, indicating that the model is often overconfident when the correct answer is no under a polarity-transformed prompt. These baseline results show that ordinary accuracy understates the difficulty of the task. The model is not failing uniformly; it is especially fragile when polarity cues require the answer to change direction.

Uniform adapters improve overall accuracy and reduce polarity errors, but the gains are limited. Test accuracy rises to 81.6%, and reversal error falls to 7.3%. Signed margin deviation decreases to 0.97. Negative calibration error improves to 9.8%. The improvement confirms that the training loss is useful and that the task is learnable with a small number of parameters. However, uniform rank spends capacity on every layer, including layers with weak polarity diagnostics. The method is stable but not parameter-efficient.

Final-block adapters produce a different pattern. They improve affirmative accuracy and sharpen final margins, giving 81.1% overall accuracy. Reversal error falls to 8.0%, which is better than the frozen model but worse than uniform adapters. Exclusion prompts remain problematic, with 75.6% accuracy. The result suggests that late layers can adjust the answer surface but cannot fully repair polarity transformations that have already been compressed or misrouted earlier in the network. Final-layer control is useful but incomplete.

Random rank allocation is variable. Across ten allocations, mean accuracy is 80.9%, with a standard deviation of 0.6 percentage points. Mean reversal error is 7.8%, with a standard deviation of 0.9 percentage points. Some random allocations approach uniform performance, while others perform little better than final-block placement. The variability shows that layer placement matters under a fixed rank budget. It also provides a useful control against the possibility that any sparse allocation is sufficient.

Depth-selective allocation gives the strongest results. Test accuracy is 82.0%, reversal error is 4.9%, signed margin deviation is 0.58, and negative calibration error is 6.9%. Compared with the frozen model, reversal error decreases by 5.9 percentage points. Compared with uniform adapters, it decreases by 2.4 percentage points with the same parameter count. The accuracy difference between depth-selective and uniform adapters is 0.4 percentage points. A paired bootstrap over polarity groups gives a 95% interval of $[-3.5, -1.3]$ percentage points for the reversal-error difference between depth-selective and uniform adapters. This supports a genuine allocation effect.

The main test results are summarized as follows:

$$\begin{aligned} \text{Accuracy}_{\text{frozen}} &= 78.9\%, & \text{Reversal}_{\text{frozen}} &= 10.8\%, \\ \text{Accuracy}_{\text{uniform}} &= 81.6\%, & \text{Reversal}_{\text{uniform}} &= 7.3\%, \\ \text{Accuracy}_{\text{final}} &= 81.1\%, & \text{Reversal}_{\text{final}} &= 8.0\%, \\ \text{Accuracy}_{\text{depth}} &= 82.0\%, & \text{Reversal}_{\text{depth}} &= 4.9\%. \end{aligned} \tag{19}$$

Prompt-form analysis reveals where the gains arise. On affirmative prompts, the frozen model scores 84.7%, uniform adapters 86.1%, and depth-selective adapters 86.3%. The difference between trained methods is small. On negative prompts, the frozen model scores 79.8%, uniform adapters 82.4%, and depth-selective adapters 83.2%. On

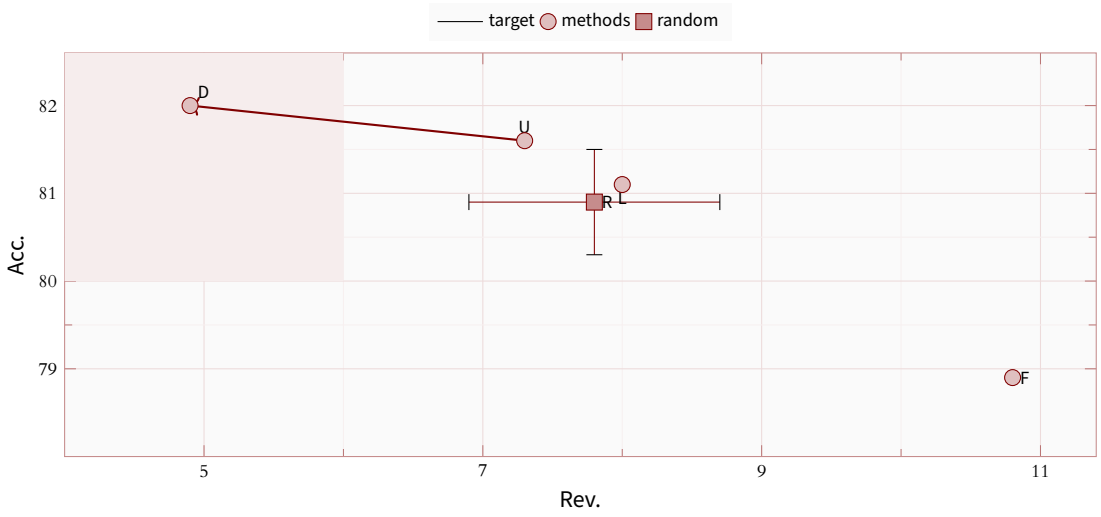


Figure 6. Accuracy–reversal operating map shows that depth-selective allocation moves the model toward the low-reversal, high-accuracy region while random allocation remains variable under the same budget.

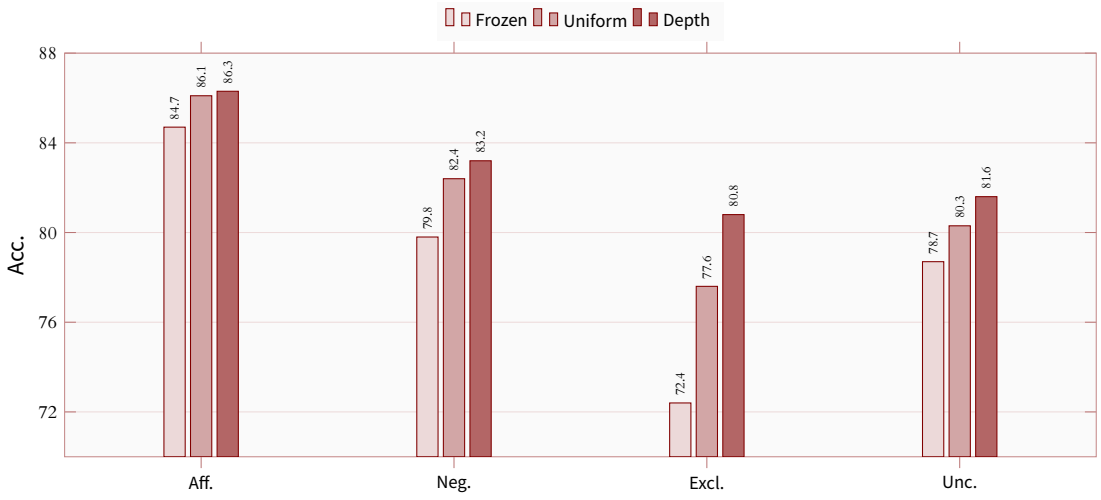


Figure 7. Prompt-form accuracy shows the largest gain under exclusion wording, where polarity reasoning is most fragile and uniform rank placement leaves a larger deficit.

exclusion prompts, the frozen model scores 72.4%, uniform adapters 77.6%, and depth-selective adapters 80.8%. On uncertainty prompts, the frozen model scores 78.7%, uniform adapters 80.3%, and depth-selective adapters 81.6%. The largest improvement therefore appears in exclusion prompts, the form most dependent on polarity reasoning.

Domain-level results show that the method is not confined to one image type. In radiography, depth-selective adapters reduce reversal error from 11.6% to 4.7%. In ultrasound, reversal error falls from 12.1% to 5.4%. In dermoscopy, it falls from 8.8% to 4.6%. The ultrasound task remains hardest, mainly because several findings have visually subtle negative cases. Dermoscopy has lower baseline reversal error because the prompt templates use shorter finding descriptions and fewer exclusion phrases. The method improves all domains, but the magnitude differs.

The paired McNemar test on group-level reversal outcomes shows significant improvement over uniform adapters with $p = 0.004$. The comparison with final-block adapters gives $p < 0.001$. The comparison with the best random allocation gives $p = 0.018$. Since many random allocations are tested, the best-random comparison is conservative in spirit but not a formal multiple-testing correction. It still suggests that the diagnostic allocation is not merely lucky. A mixed-effects logistic model with random intercepts for finding and domain estimates an odds ratio of 0.64 for reversal error under depth-selective allocation compared with uniform allocation.

Signed margin deviation provides a more continuous view. The frozen model has mean deviation 1.42 [14]. Uniform adapters reduce it to 0.97. Final-block adapters reduce it to 1.04. Depth-selective adapters reduce it to 0.58 [15]. The reduction is especially large in pairs involving exclusion prompts. The distribution also changes shape.

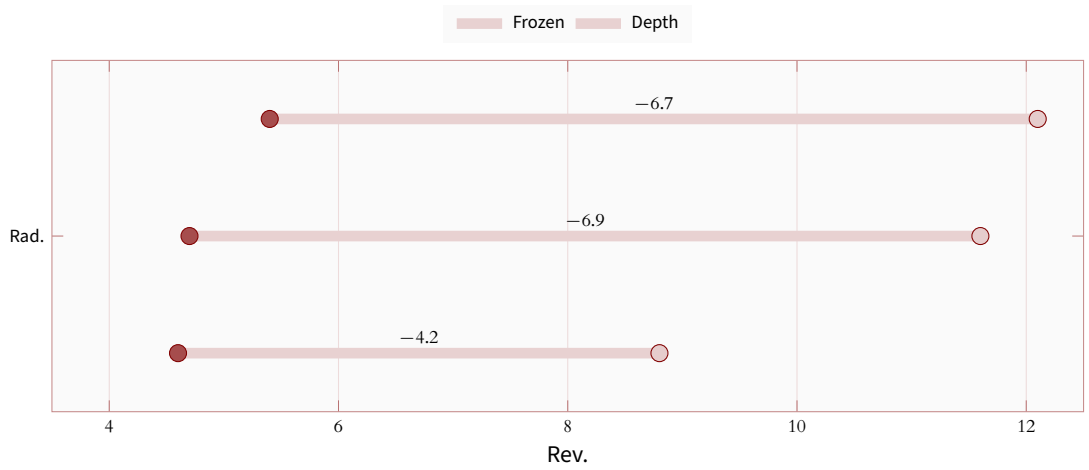


Figure 8. Domain-wise reversal error drops in radiography, ultrasound, and dermoscopy, with the hardest ultrasound setting retaining the largest absolute residual error.

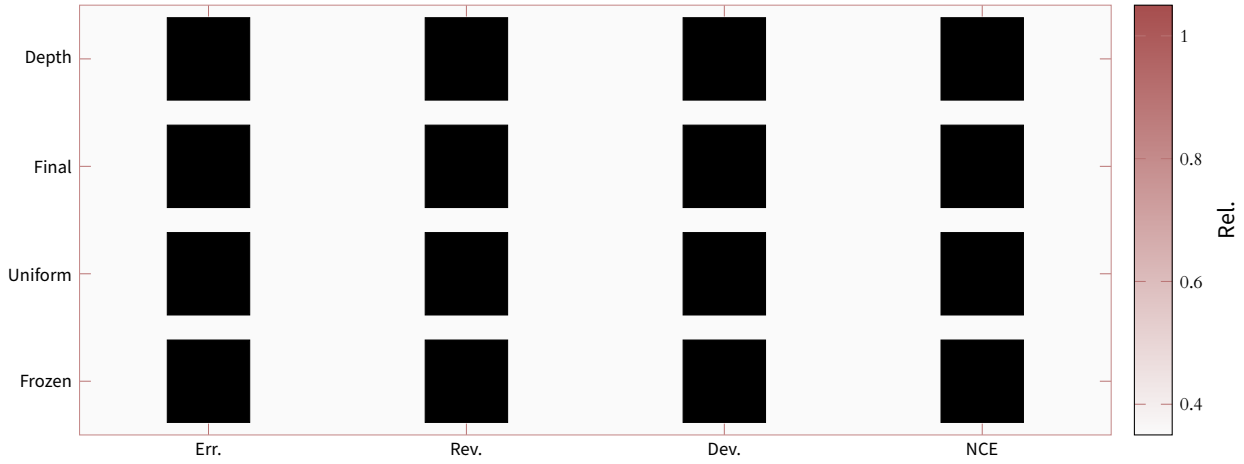


Figure 9. Normalized error matrix compares binary error, reversal error, signed-margin deviation, and negative calibration error; lower relative cells concentrate in the depth-selective row.

The frozen model has a long tail of deviations above 3.0 logits. Uniform adapters reduce but do not remove this tail. Depth-selective adapters cut the proportion of high-deviation pairs from 14.9% to 4.1%. These high-deviation pairs are clinically relevant because they indicate that the model's confidence depends heavily on wording.

Negative calibration error improves alongside reversal error. The frozen model has negative calibration error 13.6%. Uniform adapters reduce it to 9.8%. Final-block adapters reduce it to 10.4%. Depth-selective adapters reduce it to 6.9%. Calibration curves show that the largest correction occurs in bins where the model predicts yes with probability between 0.70 and 0.90 while the correct answer is no under a polarity-transformed prompt. The depth-selective method does not simply make probabilities less confident everywhere. Positive calibration error changes from 8.1% in the frozen model to 6.7% after depth-selective adaptation, a smaller but favorable shift.

The adapter rank pattern selected by the diagnostic rule is informative. Under the main budget, nonzero ranks appear in three regions. The early region receives ranks in layers 3, 4, and 6. The middle region receives ranks in layers 13, 15, 16, and 18. The late region receives ranks in layers 25 and 28. The highest individual rank is assigned to layer 15, where representation separation and propagation scores are both high. The early ranks are lower but consistently selected because Jacobian energy around polarity tokens is strong. The late ranks are modest and appear to support final margin correction.

Removing the early selected layers while redistributing their rank to selected middle layers increases reversal error from 4.9% to 6.1%. Removing the middle selected layers and redistributing rank to early and late layers increases reversal error to 7.0%. Removing the late selected layers increases reversal error to 5.7%. These ablations indicate that middle layers contribute most, early layers are also necessary, and late layers provide additional margin control. The full allocation works because it combines regions rather than because one region dominates completely.

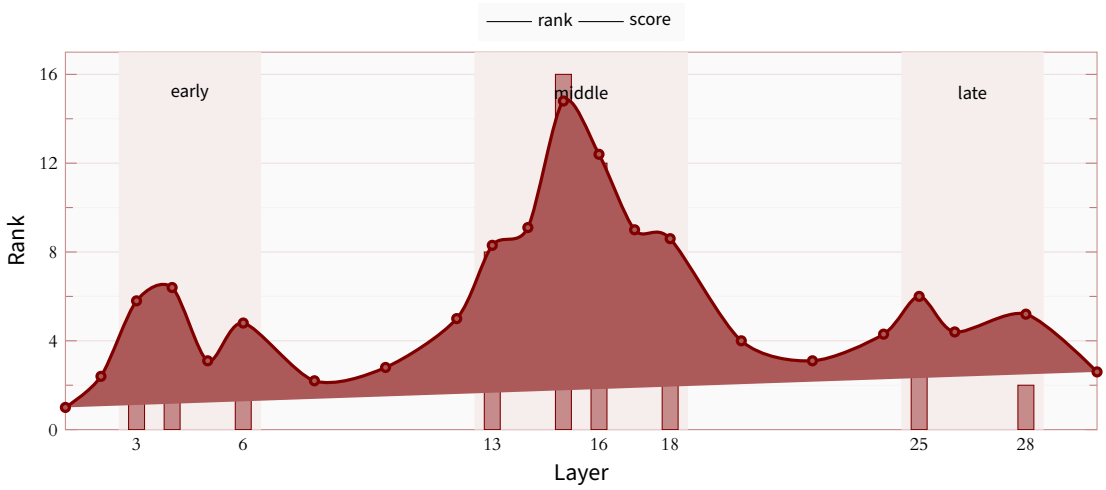


Figure 10. Selected adapter capacity spans token-sensitive early layers, representation-shaping middle layers, and margin-coupled late layers rather than forming a single final block.

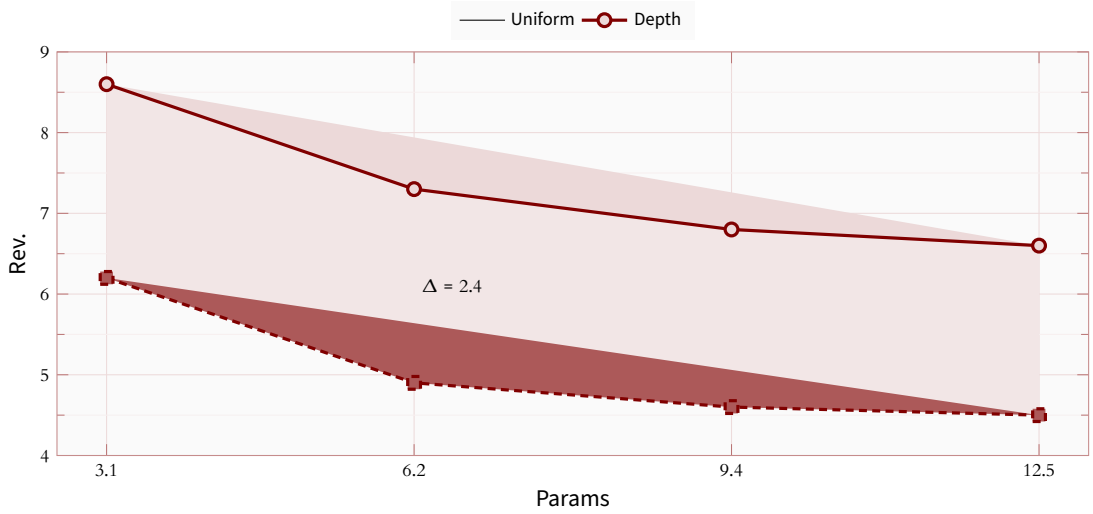


Figure 11. Rank-budget scaling shows diminishing returns for both methods, with the largest allocation advantage appearing at the moderate 6.2M-parameter setting.

The allocation ablation is

$$\begin{aligned} \text{Full allocation} &\Rightarrow R = 4.9\%, \quad V = 0.58, \\ \text{No early selected layers} &\Rightarrow R = 6.1\%, \quad V = 0.74, \\ \text{No middle selected layers} &\Rightarrow R = 7.0\%, \quad V = 0.89, \\ \text{No late selected layers} &\Rightarrow R = 5.7\%, \quad V = 0.66. \end{aligned} \tag{20}$$

Rank budget experiments show diminishing returns. With a 3.1 million parameter budget, uniform adapters reach 8.6% reversal error and depth-selective adapters reach 6.2%. With 6.2 million parameters, the values are 7.3% and 4.9%. With 9.4 million parameters, they are 6.8% and 4.6%. With 12.5 million parameters, they are 6.6% and 4.5%. The gap between allocation strategies is largest at moderate budgets. Once the budget becomes large, uniform placement improves because enough capacity reaches useful layers by chance. The practical value of depth selection is therefore strongest when parameter budgets are constrained.

The loss ablation shows that rank allocation and polarity contrast are complementary. If the binary loss is used without the polarity contrast term, depth-selective adapters reach 81.7% accuracy but 6.8% reversal error. Adding the contrast term reduces reversal error to 4.9%. Increasing the contrast weight too much lowers reversal error slightly on validation but reduces test accuracy, suggesting overemphasis on paired prompt agreement. The chosen weight balances ordinary correctness and polarity robustness [16].

Cross-domain transfer is tested by training adapters on two domains and evaluating on the held-out third domain. When radiography is held out, depth-selective adapters trained on ultrasound and dermoscopy reduce radiography reversal error from 11.6% to 7.2%. When ultrasound is held out, reversal error falls from 12.1% to 8.1%. When dermoscopy is held out, reversal error falls from 8.8% to 6.0%. These gains are smaller than in-domain training but

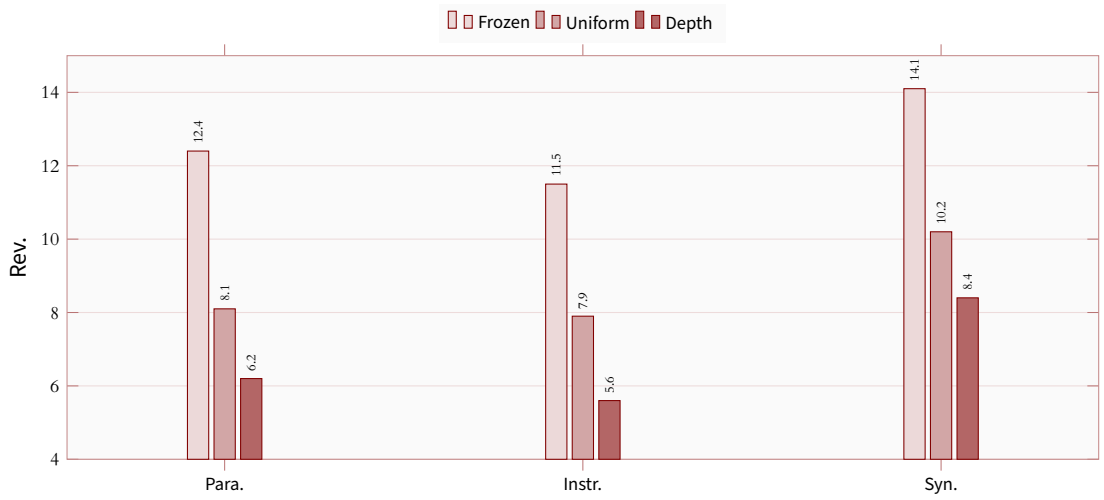


Figure 12. Robustness tests retain the depth-selective advantage under paraphrased prompts, altered instructions, and rare clinical synonym substitution, although synonym transfer remains weaker.

remain meaningful. The result suggests that some polarity adaptation transfers across medical image domains, though domain-specific examples improve performance.

A further test trains the depth allocation diagnostics on one domain but trains adapter weights on all domains. When diagnostics are computed only from radiography, the resulting allocation achieves 5.5% overall reversal error. Ultrasound-only diagnostics give 5.7%. Dermoscopy-only diagnostics give 6.0%. All-domain diagnostics give 4.9%. This indicates that the diagnostic layer pattern has shared structure but benefits from multi-domain measurement. Dermoscopy-only diagnostics underselect early layers, which may explain its weaker transfer.

The method also affects abstention quality when an optional uncertainty threshold is applied. If predictions with $|m| < 0.4$ are abstained, the frozen model abstains on 9.7% of examples and has non-abstained reversal error 8.9%. Depth-selective adapters abstain on 8.4% and have non-abstained reversal error 3.8%. At matched 10.0% abstention, depth-selective adapters retain higher accuracy and lower reversal error than uniform adapters. This suggests that the method reduces genuine polarity mistakes rather than merely pushing errors into the uncertain region [17].

5. Layer Behavior After Adaptation

The diagnostic scores are computed before training, but adaptation changes the model's layer behavior. To examine this, hidden states are remeasured after adapter training. The analysis focuses on the test set and compares uniform and depth-selective adapters. The first question is whether the selected layers become more or less sensitive to polarity tokens. The second is whether polarity representation separation becomes better organized. The third is whether answer margins respond more symmetrically to opposed prompt forms.

After depth-selective adaptation, early selected layers show reduced excessive Jacobian concentration on individual polarity tokens. This does not mean the model ignores polarity. Rather, the gradient becomes distributed across a slightly broader set of surrounding tokens. For example, in exclusion prompts, sensitivity shifts from the single rule-out token toward the full predicate phrase. This change is associated with fewer cases where the model treats rule-out language as a weak synonym of absence or presence. Uniform adapters show a similar but smaller effect.

Middle selected layers show the clearest representational change. Before adaptation, affirmative and exclusion prompts for the same image often occupy representational neighborhoods that are separated mainly by surface form. After depth-selective adaptation, the separation aligns more strongly with correct answer polarity. This is measured by comparing two linear probes. A prompt-form probe predicts whether a hidden state comes from affirmative, negative, exclusion, or uncertainty wording. A target-answer probe predicts the correct yes-no answer. In middle selected layers, prompt-form probe accuracy decreases from 71.8% to 63.2%, while target-answer probe accuracy increases from 76.4% to 81.5%. This suggests that representations become less organized by superficial prompt form and more organized by the answer implied by image and polarity together.

Late selected layers show improved signed margin symmetry. For opposed prompt pairs, the final margins should move in opposite raw directions but similar signed directions. Before adaptation, many pairs have one strong signed margin and one weak or negative signed margin. After depth-selective adaptation, the signed margins become more balanced. The mean product of signed margins in opposed pairs increases from 1.18 in the frozen model to 2.07 under uniform adapters and 2.84 under depth-selective adapters. A higher product indicates that both prompts in the pair are confidently correct rather than one prompt dominating.

The signed margin product is

$$B_{\text{pair}} = \frac{1}{|\mathcal{P}|} \sum_{(g,k,l)\in\mathcal{P}} s_{gk} s_{gl}. \tag{21}$$

The test values are [18]

$$B_{\text{frozen}} = 1.18, \quad B_{\text{uniform}} = 2.07, \quad B_{\text{depth}} = 2.84. \tag{22}$$

The improvement is largest for exclusion pairs, where the frozen value is only 0.72 and the depth-selective value is 2.43. This indicates that the method does not merely flip answers at the threshold but strengthens paired correctness [19].

Centered kernel alignment between prompt forms also changes. For affirmative and exclusion prompts, baseline alignment at selected middle layers is 0.61. Uniform adapters raise it to 0.66, while depth-selective adapters raise it to 0.72 after correcting for answer polarity. This means the representations for clinically linked prompts become more similar once the correct answer relation is accounted for [20]. For affirmative and absence prompts, raw alignment remains lower because the correct answers are opposite, but signed-answer-conditioned alignment improves. This supports the idea that adaptation reorganizes representations around polarity-adjusted task meaning.

The method does not erase all prompt-form information. A prompt-form probe still performs above chance after adaptation. This is expected and probably desirable. The model should know whether a question asks about presence or exclusion because that information determines the answer. The goal is not prompt-form invariance [21]. The goal is polarity competence. A model that made all prompt forms identical would fail absence and exclusion questions. Depth-selective adaptation reduces unhelpful surface separation while retaining answer-relevant polarity structure.

Error cases after adaptation reveal remaining weaknesses. Some errors involve nested uncertainty, such as prompts asking whether it is unlikely that a finding cannot be excluded. These were mostly removed from the main corpus, but a small challenge set includes them. The depth-selective method improves performance on this set from 59.3% to 66.8%, but errors remain common. Another weakness is visual ambiguity. When the image evidence is weak, the model may choose the same answer across opposed prompts because the visual signal dominates the linguistic operation. A third weakness is synonym variation. Some rare clinical aliases are handled less well because the diagnostic set contains limited examples.

The study includes a causality-inspired perturbation analysis, though not a full mechanistic proof. At selected layers, the hidden-state difference between affirmative and exclusion prompts is scaled by coefficients from -1.0 to 1.0, and the output margin is measured. In the frozen model, the margin response is nonlinear and often asymmetric. Under depth-selective adaptation, the response becomes more monotonic for correctly classified groups. The mean monotonicity score increases from 0.68 to 0.82 [22]. Uniform adapters reach 0.75. This suggests that the adapted model has a smoother mapping from polarity representation displacement to answer margin.

The monotonicity score is computed as

$$M_\ell = 1 - \frac{1}{T-1} \sum_{a=1}^{T-1} \mathbb{I}\left[\Delta m_\ell(\alpha_{a+1}) < \Delta m_\ell(\alpha_a)\right], \tag{23}$$

where α_a indexes increasing perturbation coefficients after aligning the expected sign. A score of 1 indicates fully monotonic response. This analysis is descriptive but useful. It shows that adaptation changes not only the final predictions but also the local geometry of polarity effects.

Layerwise adapter norm analysis shows that allocated rank is not always fully used. Some low-rank matrices develop small effective singular values. The effective rank is computed as the number of singular values explaining 95.0% of adapter update energy. Early selected layers have effective ranks near their assigned ranks. Some late layers have lower effective rank than assigned, suggesting that smaller late ranks may be sufficient. Middle layers use most of their assigned capacity. This matches the ablation finding that removing middle layers is most damaging.

The effective rank is

$$r_{\text{eff},\ell} = \min\left\{ r : \frac{\sum_{i=1}^{r} \sigma_{\ell i}^2}{\sum_{i=1}^{r_\ell} \sigma_{\ell i}^2} \geq 0.95 \right\}. \tag{24}$$

Across selected layers, the mean ratio $r_{\text{eff},\ell}/r_\ell$ is 0.84 for early layers, 0.88 for middle layers, and 0.61 for late layers. This suggests that future allocation rules could assign late layers more conservatively or use adaptive rank pruning after training.

6. Robustness Analysis

Several robustness experiments test whether the main result depends on specific design choices. The first varies the prompt grammar. A paraphrastic prompt set is created by rewriting the four prompt forms with different lexical choices while preserving polarity. The depth-selective model trained on the original grammar is evaluated on this rewritten test set without additional training. Reversal error rises from 4.9% to 6.2%, but it remains below uniform adapters at 8.1% and the frozen model at 12.4%. The increase indicates some template dependence, but the retained advantage suggests that the method learns more than memorized wording.

The second robustness test changes the answer instruction. Instead of asking for only yes or no, the prompt asks for a concise binary answer. This minor instruction change affects the frozen model, raising reversal error from 10.8% to 11.5%. Uniform adapters reach 7.9%, and depth-selective adapters reach 5.6%. The method remains beneficial, but the result shows that instruction wording can still matter

[23]. Adapter training under multiple instruction prefixes may improve robustness further.

The third test evaluates label imbalance. The main corpus is balanced, but real datasets may not be. A skewed test set is created with 70.0% negative labels in each domain. The frozen model has 80.4% accuracy but 12.7% reversal error. Depth-selective adapters reach 83.1% accuracy and 5.8% reversal error. Negative calibration error improves from 15.2% to 7.6%. The method performs well under imbalance, partly because the polarity contrast term is group-based rather than prevalence-based.

The fourth test evaluates reduced training data. With 25.0% of the training groups, depth-selective adapters reach 7.4% reversal error. With 50.0%, they reach 5.9%. With the full training set, they reach 4.9%. Uniform adapters are more affected by data reduction, reaching 10.1%, 8.6%, and 7.3% respectively. Depth-selective allocation appears more data-efficient because it trains fewer layers with greater diagnostic relevance. However, the diagnostic set itself must still be large enough to estimate layer scores reliably.

The fifth test examines diagnostic noise. Layer scores are estimated using 100, 200, 400, and 600 diagnostic groups. With 100 groups, selected layers vary noticeably across random samples, and reversal error is 6.5%. With 200 groups, reversal error is 5.7%. With 400 groups, it is 5.1%. With 600 groups, it is 4.9%. The diagnostic stage therefore benefits from several hundred groups, but it does not require the full training set. This matters because layer diagnostics can be computationally expensive.

The sixth test compares alternative diagnostic combinations. Using only Jacobian energy gives 6.4% reversal error. Using only representation separation gives 5.8%. Using only margin propagation gives 6.0%. Combining all three gives 4.9%. Pairwise combinations fall between 5.2% and 5.6%. Representation separation is the strongest single score, but it does not fully replace the others. The combined rule works because polarity processing involves token sensitivity, hidden-state organization, and final margin influence.

The seventh test changes the rank smoothness penalty. With no smoothness penalty, the allocation becomes scattered and reversal error is 5.4%. With the main penalty, reversal error is 4.9%. With double the penalty, the allocation becomes too contiguous and reversal error is 5.3%. A moderate penalty is best [24]. The result suggests that nearby layers share useful behavior, but strict contiguity is not optimal.

The eighth test evaluates whether visual encoder adaptation would change conclusions. A small adapter is added to the visual projection layer while keeping the same total parameter budget by reducing language adapter rank. Accuracy improves slightly on affirmative prompts, but reversal error worsens to 5.6%. The likely reason is that polarity errors are primarily language-conditional in this task; spending capacity on visual adaptation leaves less capacity for polarity processing. This does not mean visual adaptation is unimportant for other tasks. It means that under a fixed budget for this polarity problem, language-layer allocation is more useful.

The ninth test evaluates out-of-template clinical synonyms. A synonym challenge set replaces common target terms with less frequent clinical phrases. In radiography, airspace opacity replaces consolidation in some prompts. In ultrasound, peritoneal fluid replaces free fluid. In dermoscopy, atypical vessels replaces atypical vascular pattern. The frozen model has 14.1% reversal error, uniform adapters 10.2%, and depth-selective adapters 8.4%. The improvement is smaller than in the main test. This indicates that polarity robustness and synonym robustness are related but not identical. Additional lexical coverage would be needed for rare phrase variants [25].

The tenth test evaluates random seed stability. Across five training seeds, depth-selective reversal error ranges from 4.6% to 5.3%, with standard deviation 0.3 percentage points. Uniform adapters range from 7.0% to 7.7%. Final-block adapters range from 7.5% to 8.4%. Random allocation has much wider variation because the allocation itself changes [26]. The proposed method is therefore stable once diagnostics are fixed.

The robustness results support the central finding while defining its limits. Depth-selective allocation remains stronger than baselines under rewritten prompts, instruction changes, imbalance, smaller training sets, and seed variation. It is less effective under rare synonym substitution and nested uncertainty. It depends on reasonably stable layer diagnostics. It works best when the main failure is polarity processing rather than missing visual recognition or unfamiliar terminology.

7. Error Analysis and Practical Implications

The remaining reversal errors after depth-selective adaptation are not randomly distributed. The largest cluster involves exclusion prompts with visually subtle positive findings. For example, if a small effusion is present and the prompt asks whether effusion can be excluded, the correct answer is no. The model sometimes answers yes, apparently treating uncertainty in the image as support for exclusion. This is a clinically meaningful confusion because image uncertainty and linguistic exclusion are different [27]. The model needs to understand that inability to see a finding confidently does not always mean the finding can be ruled out.

A second cluster involves negative prompts with high-salience visual distractors. In dermoscopy, an image may contain an atypical feature related to the target but not exactly matching it. The model sometimes answers yes to absence prompts because the visual abnormality dominates the language polarity. The adapter reduces this error type but does not eliminate it. This suggests that the language route can be improved while still being overridden by strong visual associations.

A third cluster involves uncertainty wording. Prompts asking whether a finding is unlikely can be difficult when the image is borderline. The correct answer key treats un-

likely as a binary polarity form, but clinical interpretation may be less discrete. If a finding is subtle, both yes and no may seem defensible depending on threshold. The model's errors in this cluster may partly reflect label simplification rather than pure language failure. This is one reason the paper avoids making clinical safety claims from binary metrics alone.

A fourth cluster involves domain-specific phrasing. Ultrasound prompts contain anatomical and procedural terms that are less common in radiography-style VQA data. Even when the model recognizes the image pattern, it may handle phrases such as appendiceal noncompressibility or intrauterine gestational sac differently across polarity forms. Depth-selective allocation improves general polarity handling, but it cannot fully compensate for uneven domain terminology. Domain-specific prompt augmentation may be needed.

The practical implication is that adapter allocation can be treated as a design variable rather than an afterthought. When a team has a fixed parameter budget, placing the same rank everywhere is not always the best use of capacity. Measuring where the target linguistic operation affects the model can guide placement. This is particularly relevant for medical settings where the target operation is known in advance. If the goal is polarity robustness, diagnostics should focus on polarity. If the goal is temporal comparison, the diagnostic set should test before-after language. If the goal is anatomical grounding, layer scores should measure location-sensitive representations.

The method also suggests a validation workflow. Before adaptation, construct paired prompt groups for the linguistic operation of interest. Measure layer diagnostics on a sample. Allocate adapter rank under a fixed budget. Train with a loss that includes ordinary supervision and a task-specific contrast. Evaluate not only accuracy but also operation-specific errors. This workflow is more involved than ordinary adapter training, but it provides more information about why capacity is placed where it is. It also creates an audit trail for model modification [28].

The study's results caution against relying only on final-block adaptation. Final layers are attractive because they are close to the output and cheaper to adjust. They do help, but they leave some polarity errors unresolved. The language operation begins before the answer surface. If earlier layers encode the prompt in a way that weakens or distorts polarity, late adapters may only patch the final decision. A small amount of earlier capacity can make the later correction easier. The best allocation in this study therefore spans multiple regions.

Uniform adaptation has the opposite issue [29]. It is robust and simple but may waste rank. In large budgets, this waste matters less because enough capacity reaches useful layers. In small budgets, it matters more. Medical models are often adapted under memory, compute, or governance constraints. A method that improves the use of a fixed budget can be valuable even if it does not change the underlying training data. The rank budget experiments show this clearly: the proposed method's advantage is largest when capacity is limited.

The method also has interpretive value. The selected layers and their diagnostic scores reveal how the model handles a linguistic operation. The result is not a complete explanation, but it gives a structured view of where token sensitivity, representation separation, and margin coupling occur. In this study, polarity is distributed across early, middle, and late layers. That pattern argues against a simple story in which one layer is the polarity layer. It also explains why an allocation spanning several regions works better than a single contiguous final block.

For deployment, polarity reversal error is an intuitive metric. Clinicians and auditors can understand that a model should not answer the same way to a presence question and a rule-out question when the correct answers differ. Accuracy alone can hide this failure because one of the prompts may still be correct. Reporting reversal error alongside accuracy provides a clearer picture of linguistic reliability. Signed margin deviation adds another layer by identifying cases where answers are correct but confidence is highly uneven across polarity forms.

The method does not remove the need for human-designed prompt standards. A model can be made more robust, but clinical systems should still avoid needlessly confusing language. Prompt templates should minimize double negation and ambiguous exclusion phrasing when possible. Robustness training is a safeguard, not an invitation to use poor prompts. The best practical approach is likely both: clearer prompts and models that are less brittle when prompt variation occurs.

8. Limitations and Future Directions

The first limitation is the controlled nature of the prompt grammar. The study uses carefully constructed polarity forms to ensure correct answer keys. Real clinical language is less tidy. Users may ask compound questions, include multiple findings, or mix uncertainty with comparison and negation. The results therefore show that depth-selective allocation helps under controlled polarity transformations, not that it handles every natural prompt. Future studies should use real user logs with expert annotation of polarity intent.

The second limitation is the binary answer space [30]. Polarity errors are easiest to define when answers are yes or no. Many medical vision-language tasks require free-text explanation, localization, measurement, or differential diagnosis. A model might correctly answer yes but provide an explanation with inconsistent polarity. Extending the method to structured or free-text outputs would require new metrics. One possibility is to score multiple clinically relevant answer propositions and apply polarity contrast at the proposition level.

The third limitation is that the study freezes the visual encoder. This is appropriate for isolating language-layer allocation, but some errors involve visual ambiguity or

domain-specific image features. Joint visual-language adaptation might improve those cases. However, visual adaptation also increases validation burden and may introduce new failure modes. Future work should examine whether depth-selective language adapters combine well with small visual adapters when the task includes both polarity and visual recognition gaps.

The fourth limitation is that the diagnostic scores require additional computation. Jacobian energy, representation alignment, and perturbation-based propagation are more expensive than ordinary validation metrics. The cost is manageable for the present diagnostic set, but it may be substantial for very large models or many candidate tasks. Approximate diagnostics could reduce cost. For example, gradient norms could be estimated on smaller token subsets, and propagation scores could use linear probes rather than perturbation passes.

The fifth limitation is that the allocation rule uses fixed weights for the three diagnostics. These weights are selected on validation data and may not transfer to every model. A more principled approach could learn the weights through bilevel optimization, where validation reversal error directly supervises allocation. Such an approach might improve performance but would also be more complex and more expensive. The current rule is intentionally simpler and easier to inspect.

The sixth limitation is that rank choices are discrete and manually bounded. The allowed ranks are practical, but they may not be optimal. Continuous relaxation of rank allocation could produce finer distributions, followed by rounding. Another direction is adaptive rank pruning after training. The effective-rank analysis suggests that some late-layer rank is underused. A train-prune-retrain cycle might reduce parameters without losing robustness.

The seventh limitation is that clinical labels may be imperfect. Binary labels simplify complex image findings. Some uncertainty prompts are especially sensitive to label thresholds. If the label says a finding is present but the visual evidence is minimal, the correct answer to an uncertainty question may not be straightforward. Human expert review of the hardest polarity cases would improve the validity of future evaluations. The present study treats labels as fixed because it focuses on model behavior under controlled prompt transformations.

The eighth limitation is that the method addresses one linguistic operation at a time. Polarity is important, but medical language also contains temporality, laterality, severity, quantification, and causal relations. A model adapted for polarity might not improve on these other operations. A future system could define multiple diagnostic sets and allocate adapter capacity for several operations jointly. The challenge would be avoiding conflict between operations competing for the same layers.

Future work should also evaluate whether depth-selective allocation is useful outside medicine. Polarity reasoning appears in legal, financial, and safety-critical technical domains. The general principle of measuring layer sensitivity to a target linguistic operation and allocating adapter capacity accordingly may transfer. However, medical images add a multimodal constraint that may change the allocation pattern. Cross-domain studies would clarify which parts of the approach are general and which are specific to medical vision-language inference.

Another direction is active data selection. The diagnostic scores identify layers, but the training data are still sampled broadly. The model might benefit from emphasizing groups whose polarity transformations produce large hidden-state changes or margin conflicts. An active curriculum could select examples that most directly exercise the selected layers. This may reduce the amount of training data needed for robust adaptation.

Finally, future work should connect polarity robustness with abstention and explanation. A model that remains uncertain under difficult polarity transformations should perhaps abstain rather than force a yes or no. The signed margin measures developed here could support abstention rules. Explanations could also be checked for polarity agreement with the final answer. A model that answers no to a rule-out question but explains as if the finding is present would still be unsafe. Robustness needs to apply to both answer and rationale.

9. Conclusion

This study evaluated depth-selective adapter allocation for improving polarity robustness in medical vision-language inference. The problem was framed around binary clinical prompts whose wording changes the required answer direction. Instead of increasing the adapter budget, the method redistributed a fixed low-rank budget across layers using diagnostic scores for token-local sensitivity, representation separation, and signed margin propagation.

The empirical results show that allocation matters. Under the same parameter budget and training loss, depth-selective adapters reduced polarity reversal error more than uniform adapters, final-block adapters, and random rank allocation. The method was especially helpful for exclusion-framed prompts, where ordinary accuracy was weakest and polarity mistakes were most common. It also reduced signed margin deviation and improved negative calibration error while preserving overall task accuracy.

The layer analyses indicate that polarity processing is distributed. Early layers respond strongly to polarity-bearing tokens, middle layers reorganize prompt representations, and later layers connect those representations to the answer margin. The best allocation spans these regions rather than concentrating all capacity in one place. This finding supports a practical view of adapter design: the placement of trainable capacity should reflect the linguistic operation being adapted.

The study has limits. The prompt grammar is controlled, the tasks are binary, and the labels simplify clinical uncertainty. The method requires diagnostic computation and hidden-state access [31]. It should therefore be treated as a

targeted adaptation strategy rather than a complete solution for medical language reliability. It improves a specific class of polarity failures but does not replace domain validation, careful prompt design, or expert review.

A useful next step is to apply the same allocation principle to other clinical language operations, including temporal comparison, laterality, and severity. Each operation may produce a different layer profile and require a different rank distribution. The broader lesson is that adapter capacity is not only a quantity but also a location. For medical vision-language systems, where small wording changes can reverse meaning, that location can materially affect robustness.

References

[1] K. Curyto, L. M. Dockler, and K. V. Haitsma, "Preference consistency: Veteran and non-veteran nursing home resident self-reported preferences for everyday living.," *Journal of gerontological nursing*, vol. 46, pp. 15–25, 3 2020.

[2] C. K. McMullen, J. L. Schneider, A. Firemark, J. V. Davis, and M. Spofford, "Cultivating engaged leadership through a learning collaborative: Lessons from primary care renewal in oregon safety net clinics," *Annals of family medicine*, vol. 11, pp. S34–40, 5 2013.

[3] K. C. Kuban, T. M. O'Shea, E. N. Allred, H. Tager-Flusberg, D. J. Goldstein, and A. Leviton, "Positive screening on the modified checklist for autism in toddlers (m-chat) in extremely low gestational age newborns.," *The Journal of pediatrics*, vol. 154, pp. 535–540, 1 2009.

[4] C. Bricker-Anthony, J. Hines-Beard, and T. S. Rex, "Molecular changes and vision loss in a mouse model of closed-globe blast trauma.," *Investigative ophthalmology & visual science*, vol. 55, pp. 4853–4862, 7 2014.

[5] B. Sadanandan and V. Behzadan, "Mechanistically guided lora improves paraphrase consistency in medical vision-language models," *arXiv preprint arXiv:2603.00148*, 2026.

[6] R. Wang, L. V. Tang, and Y. Hu, "Genetic factors, risk prediction and ai application of thrombotic diseases.," *Experimental hematology & oncology*, vol. 13, pp. 89–89, 8 2024.

[7] C. M. Bann, J. E. Newman, B. Poindexter, K. Okoniewski, S. DeMauro, S. A. Lorch, D. Wilson-Costello, N. Ambalavanan, M. Peralta-Carcelen, C. Limperopoulos, K. Kapse, J. M. Davis, M. Walsh, and S. Merhar, "Outcomes of babies with opioid exposure (oboe): protocol of a prospective longitudinal cohort study.," *Pediatric research*, vol. 93, pp. 1772–1779, 8 2022.

[8] M.-C. Fung, "Beyond the hype: navigating the real-world applications of artificial intelligence (ai) in healthcare.," *Medical review (2021)*, vol. 6, pp. 14–34, 1 2026.

[9] K. Johari, N. Riccardi, S. Malyutina, M. Modi, and R. H. Desai, "Hd-tdcs over motor cortex facilitates figurative and literal action sentence processing.," *Neuropsychologia*, vol. 159, pp. 107955–107955, 7 2021.

[10] A. Jeyabose, B. Robinson, O. Leggio, M. H. Kazemi, and V. L. Boerwinkle, "Multimodal contrastive learning on rs-fmri to quantify whole-brain network recovery after hypothalamic hamartoma surgery.," *Biomedical engineering online*, vol. 24, pp. 125–125, 10 2025.

[11] R. R. Krishna, A. Priyadarshini, A. V. Jha, B. Appasani, A. Srinivasulu, and N. Bizon, "State-of-the-art review on iot threats and attacks: Taxonomy, challenges and solutions," *Sustainability*, vol. 13, pp. 9463–, 8 2021.

[12] Y. Li, J. Li, M. Li, E. Yu, D. Rhee, M. Amith, L. Tang, L. S. Savas, L. Cui, and C. Tao, "Vaxbot-hpv: a gpt-based chatbot for answering hpv vaccine-related questions.," *JAMIA open*, vol. 8, pp. ooaf005–ooaf005, 12 2024.

[13] H. Miao, C. Li, and J. Wang, "A future of smarter digital health empowered by generative pretrained transformer.," *Journal of medical Internet research*, vol. 25, pp. e49963–e49963, 9 2023.

[14] Y. Lin, Z. Luo, Z. Ye, N. Zhong, L. Zhao, L. Zhang, X. Li, Z. Chen, and Y. Chen, "Applications, challenges, and prospects of generative artificial intelligence empowering medical education: Scoping review.," *JMIR medical education*, vol. 11, pp. e71125–e71125, 10 2025.

[15] Z. Bi, J. Li, Q. Liu, and Z. Fang, "Deep learning-based optical coherence tomography and retinal images for detection of diabetic retinopathy: a systematic and meta analysis.," *Frontiers in endocrinology*, vol. 16, pp. 1485311–1485311, 3 2025.

[16] D. Guo and K.-K. R. Choo, "Applications of federated large language model for adverse drug reactions prediction: Scoping review.," *Journal of medical Internet research*, vol. 27, pp. e68291–e68291, 9 2025.

[17] V. S. S. S. Bhuvaneswari and M. Thangamuthu, "Towards intelligent safety: A systematic review on assault detection and technologies.," *Sensors (Basel, Switzerland)*, vol. 25, pp. 3985–3985, 6 2025.

[18] R. P. Mucelli, "Staging renal cancer," *Insights into imaging*, vol. 4, pp. 1–144, 3 2013.

[19] Y. Belinkov and J. Glass, "Analysis methods in neural language processing: A survey," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 49–72, 4 2019.

[20] V. Shariff, C. Paritala, and K. M. Ankala, "Optimizing non small cell lung cancer detection with convolutional neural networks and differential augmentation.," *Scientific reports*, vol. 15, pp. 15640–15640, 5 2025.

[21] Y. Belinkov and J. Glass, "Analysis methods in neural language processing: A survey," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 49–72, 4 2019.

[22] M. M. Rahman, M. E. Masry, S. C. Gnyawali, Y. Xue, G. Gordillo, and J. P. Wachs, "Ai-driven integrated system for burn depth prediction with electronic medical records: Algorithm development and validation.," *JMIR medical informatics*, vol. 13, pp. e68366–e68366, 8 2025.

[23] D. C. Nguyen, M. Ding, P. N. Pathirana, and A. Seneviratne, "Blockchain and ai-based solutions to combat coronavirus (covid-19)-like epidemics: A survey," *IEEE access : practical innovations, open solutions*, vol. 9, pp. 95730–95753, 6 2021.

[24] B.-D. D. R. Banda, P. A. Okungu, N. Griffin-Shirley, M. K. Meeks, and O. Landa-Vialard, "Teaching orientation and mobility skills to students with autism and vision impairment in public schools: A data-based study," *Vision Rehabilitation International*, vol. 7, pp. 34–43, 1 2017.

[25] M. P. Moeller, G. L. Carr, L. Seaver, A. Stredler-Brown, and D. Holzinger, "Best practices in family-centered early intervention for children who are deaf or hard of hearing: an international consensus statement.," *Journal of deaf studies and deaf education*, vol. 18, pp. 429–445, 9 2013.

[26] C. L. Reighard, M. R. Pillai, S. Shroff, G. L. Spaeth, S. G. Schilling, S. S. Wizov, J. D. Stein, A. L. Robin, V. Raja, and J. R. Ehrlich, "Glaucoma-associated visual task performance and vision-related quality of life in south india," *Ophthalmology. Glaucoma*, vol. 2, pp. 357–363, 6 2019.

[27] M. Abbasian, I. Azimi, A. M. Rahmani, and R. Jain, "Conversational health agents: a personalized large language model-powered agent framework.," *JAMIA open*, vol. 8, pp. ooaf067–ooaf067, 7 2025.

[28] S. S. Mohammadi and Q. D. Nguyen, "A user-friendly approach for the diagnosis of diabetic retinopathy using chatgpt and automated machine learning.," *Ophthalmology science*, vol. 4, pp. 100495–100495, 2 2024.

[29] B. M. Yakubu, S. M. Ali, M. I. Khan, and P. Bhattarakosol, "Patcen: A blockchain-based patient-centric mechanism for the granular access control of infectious disease-related test records.," *PloS one*, vol. 19, pp. e0310407–e0310407, 9 2024.

[30] M. R. Greco, M. A. Lopez, M. L. Beltran-Quintero, E. T. Bengur, M. D. Poe, and M. L. Escolar, "Infantile krabbe disease (0-12months), progression, and recommended endpoints for clinical trials.," *Annals of clinical and translational neurology*, vol. 11, pp. 3064–3080, 11 2024.

[31] M.-R. Ki, D. H. Kim, M. A. A. Abdelhamid, and S. P. Pack, "Cancer and aging biomarkers: Classification, early detection technologies and emerging research trends.," *Biosensors*, vol. 15, pp. 737–737, 11 2025.