

# AdaFuse: Adaptive Multimodal Fusion for Lung Cancer Risk Prediction via Reinforcement Learning

Chongyu Qu<sup>1</sup>

Zhengyi Lu<sup>1</sup>

Yuxiang Lai<sup>3</sup>

Thomas Z. Li<sup>1</sup>

Junchao Zhu<sup>1</sup>

Junlin Guo<sup>1</sup>

Juming Xiong<sup>1</sup>

Yanfan Zhu<sup>1</sup>

Yuechen Yang<sup>1</sup>

Allen J. Luna<sup>1,2</sup>

Kim L. Sandler<sup>2</sup>

Bennett A. Landman<sup>1,2</sup>

Yuankai Huo<sup>1</sup>

CHONGYU.QU@VANDERBILT.EDU

ZHENGYU.LU@VANDERBILT.EDU

YUXIANGLAI117@GMAIL.COM

THOMAS.Z.LI@VANDERBILT.EDU

JUNCHAO.ZHU@VANDERBILT.EDU

JUNLIN.GUO@VANDERBILT.EDU

JUMING.XIONG@VANDERBILT.EDU

YANFAN.ZHU@VANDERBILT.EDU

YUECHEN.YANG@VANDERBILT.EDU

ALLEN.J.LUNA@VANDERBILT.EDU

KIM.SANDLER@VUMC.ORG

BENNETT.LANDMAN@VANDERBILT.EDU

YUANKAI.HUO@VANDERBILT.EDU

<sup>1</sup> Vanderbilt University, Nashville, TN, USA, 37215

<sup>2</sup> Vanderbilt University Medical Center, Nashville, TN, USA, 37232

<sup>3</sup> Emory University Atlanta, GA, USA, 30322

**Editors:** Under Review for MIDL 2026

## Abstract

Multimodal fusion has emerged as a promising paradigm for disease diagnosis and prognosis, integrating complementary information from heterogeneous data sources such as medical images, clinical records, and radiology reports. However, existing fusion methods process all available modalities through the network, either treating them equally or learning to assign different contribution weights, leaving a fundamental question unaddressed: **for a given patient, should certain modalities be used at all?** We present AdaFuse, an adaptive multimodal fusion framework that leverages reinforcement learning (RL) to learn patient-specific modality selection and fusion strategies for lung cancer risk prediction. AdaFuse formulates multimodal fusion as a sequential decision process, where the policy network iteratively decides whether to incorporate an additional modality or proceed to prediction based on the information already acquired. This sequential formulation enables the model to condition each selection on previously observed modalities and terminate early when sufficient information is available, rather than committing to a fixed subset upfront. We evaluate AdaFuse on the National Lung Screening Trial (NLST) dataset. Experimental results demonstrate that AdaFuse (0.762 AUC) outperforms both the best single-modality baseline (0.732 AUC) and the best fixed fusion strategy (0.759). Our work demonstrates the potential of reinforcement learning for personalized multimodal fusion in medical imaging, representing a shift from uniform fusion strategies toward adaptive diagnostic pipelines that learn when to consult additional modalities and when existing information suffices for accurate prediction. Code is publicly available at: <https://github.com/hrlblab/adafuse>

**Keywords:** Reinforcement Learning, Multimodal Fusion, Risk Prediction

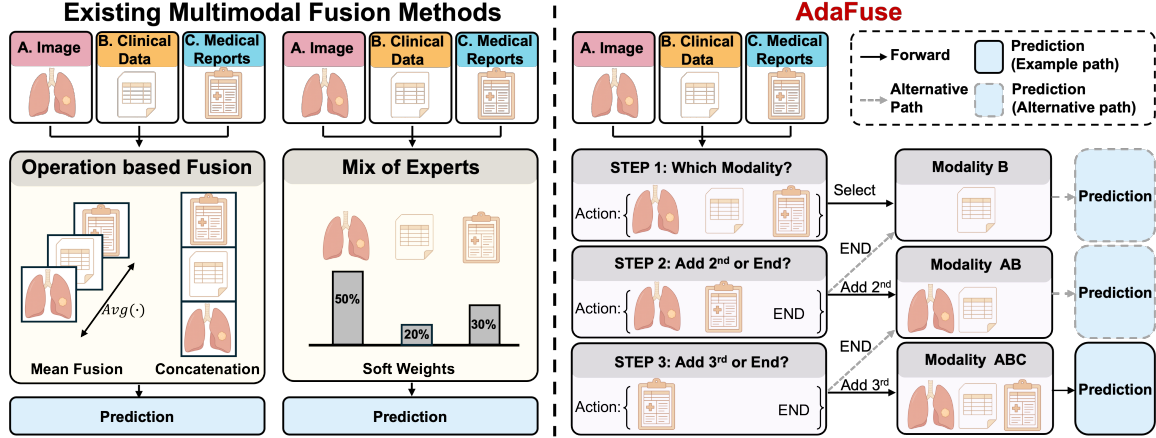


Figure 1: **Comparison of multimodal fusion paradigms.** Existing multimodal fusion methods (left two) process all modalities uniformly: operation-based fusion applies fixed combination rules, while Mixture-of-Experts learns soft weights but still requires all inputs. AdaFuse (right) makes sequential discrete decisions to select patient-specific modality subsets, with the flexibility to entirely exclude uninformative modalities from computation.

## 1. Introduction

Lung cancer remains the leading cause of cancer-related mortality worldwide (Rivera et al., 2013), with early detection being critical for improving patient outcomes (Blandin Knight et al., 2017). Low-dose computed tomography (LDCT) screening has demonstrated significant potential in reducing lung cancer mortality (Bonney et al., 2022; Ostrowski et al., 2018), as evidenced by the National Lung Screening Trial (NLST) (Team, 2011). Beyond imaging, clinical variables such as age, smoking history, and family history provide complementary risk factors (Tammemägi et al., 2013), while radiology reports (Hans Vitzthum von Eckstaedt et al., 2020) capture expert observations and contextual information. The integration of these heterogeneous data sources through multimodal fusion (Cui et al., 2022; Liu et al., 2023) has emerged as a promising direction for more accurate and comprehensive risk prediction.

Multimodal fusion methods have evolved considerably over the past years. Early approaches relied on simple operations such as concatenation (Mobadersany et al., 2018; Yap et al., 2018) or mean pooling (Cheerla and Gevaert, 2019; Ghosal et al., 2021) to combine features from different modalities. Tensor-based methods employing Kronecker products (Chen et al., 2020; Wang et al., 2021) were later introduced to capture higher-order feature interactions, though at the cost of increased computational complexity. To address this, low-rank factorization techniques (Liu et al., 2018; Sahay et al., 2020) have been proposed to reduce the parameter overhead while preserving expressive power. Attention mechanisms (Zhu et al., 2020; Lu et al., 2023) have gained popularity for their ability to learn cross-modal interactions and dynamically weight modality contributions. More recently, Mixture-of-Experts (MoE) architectures (Cao et al., 2023; Han et al., 2024) route inputs

through specialized sub-networks, providing input-dependent feature processing. Despite their architectural differences, all these methods process every available modality through the network, either treating them equally or learning to assign different contribution weights (i.e., soft selection), as illustrated in Figure 1 (left). This leaves a fundamental question unaddressed: **for a given patient, should certain modalities be used at all?**

In clinical practice, the diagnostic value of each modality varies across individuals (Acosta et al., 2022; Huang et al., 2020). Some patients may benefit from multimodal integration, while for others, a single modality suffices or additional modalities introduce noise rather than complementary information. Existing methods cannot entirely exclude uninformative modalities from computation; instead, they uniformly process all inputs regardless of their utility for individual patients. This not only incurs unnecessary computational costs but also limits the model’s ability to truly personalize the diagnostic pipeline.

To address this limitation, we propose **AdaFuse**, an adaptive multimodal fusion framework that formulates modality selection as a sequential decision process. As shown in Figure 1 (right), at each step, a policy network decides whether to incorporate an additional modality or to proceed to prediction with the current selection. Unlike existing approaches that process all modalities and learn to weight their contributions, AdaFuse makes discrete decisions to select or skip each modality, providing the flexibility to use different modality combinations for different patients. This formulation naturally mirrors clinical practice, where physicians selectively order diagnostic tests based on patient-specific factors rather than exhaustively acquiring all available data (Winslow et al., 1988; Ball et al., 2015).

Our contributions are summarized as follows:

1. We propose AdaFuse, a reinforcement learning framework that formulates multimodal fusion as a sequential decision process, enabling patient-specific modality selection for lung cancer risk prediction.
2. We conduct comprehensive experiments on the NLST dataset with three modalities, comparing against single-modality baselines and various fusion strategies including concatenation, mean pooling, tensor fusion, and MoE.
3. We provide detailed analysis of learned fusion policies, offering insights into how the model adapts its modality selection across different patient subgroups and demonstrating the potential of adaptive fusion for personalized medical diagnosis.

## 2. Method

We present AdaFuse, an adaptive multimodal fusion framework that learns patient-specific modality selection strategies through reinforcement learning. The key insight is that different patients may benefit from different modality combinations: some patients may require comprehensive multimodal integration, while others may achieve accurate predictions with fewer modalities. Rather than applying a fixed fusion strategy uniformly, AdaFuse formulates modality selection as a sequential decision process, where a policy network learns to identify the optimal modality subset for each patient.

Figure 2 illustrates the overall architecture. Given a patient with three available modalities, AdaFuse first encodes each modality into a compact representation. A policy network

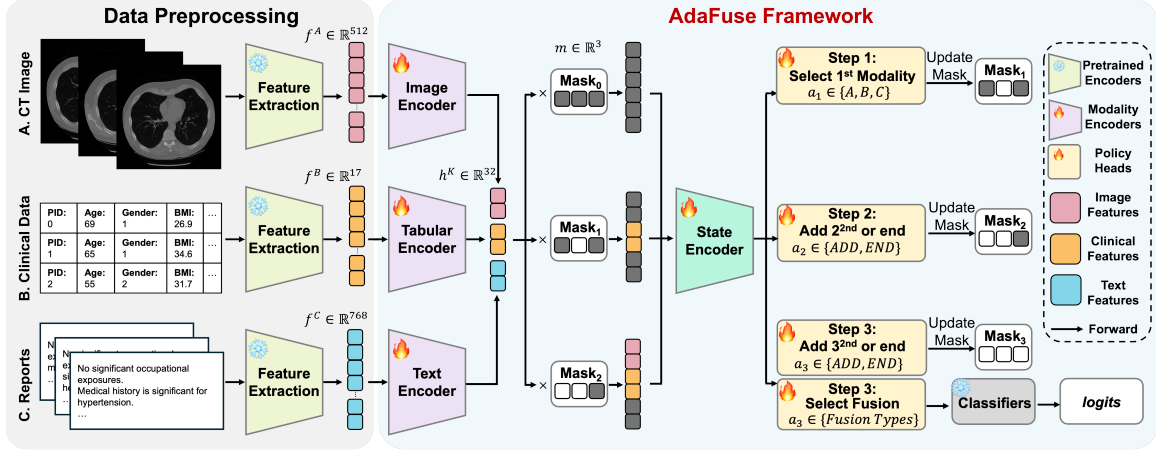


Figure 2: **Overview of the AdaFuse framework.** *Data preprocessing* extracts features from three modalities (CT images, clinical variables, and text reports) using pre-trained encoders; details are provided in Section 3.1. The **AdaFuse framework** consists of three components: (1) *Modality encoders* project each input feature to a shared 32-dimensional representation, with a binary mask  $m \in \{0, 1\}^3$  tracking selected modalities. (2) *State encoder* concatenates the masked features with the selection mask and maps them to a 64-dimensional state vector that captures the current selection status. (3) *Policy heads* make sequential decisions: Step 1 selects the primary modality from  $\{A, B, C\}$ ; Step 2 decides whether to add a second modality or proceed to prediction; Step 3 optionally incorporates the third modality and selects a fusion strategy from  $\{Concat, Mean, Tensor\}$ . After each selection, the corresponding mask entry is updated from 0 to 1, and the state encoder recomputes the state representation for the next decision. The selected modality combination is passed to the corresponding pretrained classifier among 15 fusion classifiers covering all valid modality-fusion combinations.

then makes sequential decisions: (1) selecting the primary modality, (2) deciding whether to add a second modality, and (3) optionally incorporating the third modality along with a fusion strategy. The selected modalities are fused and passed to a classifier for prediction. The entire framework is trained end-to-end using a combination of supervised learning and policy gradient optimization.

## 2.1. Preliminaries and Notations

We consider a multimodal learning setting with three modalities: CT images, clinical variables, and radiology text reports. For each patient, we denote the raw input features as  $f^A \in \mathbb{R}^{512}$  for CT image features,  $f^B \in \mathbb{R}^{17}$  for clinical variables, and  $f^C \in \mathbb{R}^{768}$  for text embeddings. Each modality is processed by a dedicated encoder  $E^A$ ,  $E^B$ ,  $E^C$  to produce encoded representations  $h^A = E^A(f^A)$ ,  $h^B = E^B(f^B)$ ,  $h^C = E^C(f^C)$ , where all encoded features share a common dimension  $d = 32$ . The modality selection process is tracked by

a binary mask  $\mathbf{m} = [m^A, m^B, m^C]^\top \in \{0, 1\}^3$ , where  $m^i = 1$  indicates that modality  $i$  has been selected.

## 2.2. Sequential Modality Selection

We formulate adaptive modality selection as a Markov Decision Process (MDP) (Bellman, 1957; Puterman, 2014), where the policy learns to sequentially construct the optimal modality subset for each patient.

**State Representation.** At each decision step  $t$ , the state  $s_t$  captures the information from currently selected modalities. We compute the state by concatenating the encoded features weighted by their selection status, along with the mask itself:

$$s_t = g_\theta ([h^A \odot m^A; h^B \odot m^B; h^C \odot m^C; \mathbf{m}]) \quad (1)$$

Here  $\odot$  denotes element-wise multiplication that zeros out unselected modalities,  $[\cdot; \cdot]$  denotes concatenation, and  $g_\theta$  is a two-layer MLP that maps the  $(3d + 3)$ -dimensional input to a 64-dimensional state vector. This formulation allows the state encoder to distinguish between “modality not yet selected” (zeroed features,  $m^i = 0$ ) and “modality selected but potentially uninformative” (non-zero features,  $m^i = 1$ ).

**Action Space.** The decision process unfolds over at most three steps, with the action space adapting based on previous selections:

- *Step 1:* The policy selects one modality from  $\{A, B, C\}$  as the primary modality.
- *Step 2:* The policy decides whether to stop with the current selection, or to add one of the two remaining modalities.
- *Step 3:* If two modalities have been selected, the policy decides whether to stop or incorporate the third modality, and selects a fusion strategy from {concatenation, mean, tensor}.

This sequential formulation naturally captures the hierarchical nature of modality selection: the primary modality anchors the decision, and subsequent choices refine the combination based on the accumulated information.

## 2.3. Policy Network Architecture

The policy network builds upon the modality encoders, augmented with a state encoder and step-specific decision heads.

**State Encoder.** The state encoder  $g_\theta$  is a two-layer MLP with ReLU activations that takes the concatenation of masked modality features and the selection mask as input, producing a 64-dimensional state vector for downstream policy decisions.

**Policy Heads.** Each decision step has dedicated heads that output action logits  $\mathbf{l}_t$ , from which we sample actions via  $\pi(a_t|s_t) = \text{softmax}(\mathbf{l}_t/\tau)$ . The temperature  $\tau$  is annealed from  $\tau_{\text{init}}$  to  $\tau_{\text{final}}$  over training to transition from exploration to exploitation; during inference, we use greedy decoding.

**Fusion Classifiers.** We maintain 15 fusion classifiers corresponding to all valid modality-fusion combinations (3 single-modality, 9 dual-modality for 3 pairs  $\times$  3 fusion types, and 3 triple-modality for 3 fusion types), and invoke the appropriate classifier based on the policy’s selection.

## 2.4. Learning Objective

**Reward Design.** After the policy completes its sequential decisions and produces a prediction  $\hat{p}$  for a patient with label  $y$ , we compute a reward signal based on two components:

$$r = r_{\text{BCE}} + \lambda_{\text{auc}} \cdot r_{\text{auc}} \quad (2)$$

The first term  $r_{\text{BCE}} = y \log \hat{p} + (1 - y) \log(1 - \hat{p})$  is the negative binary cross-entropy, which provides a continuous signal based on prediction confidence. The second term  $r_{\text{auc}}$  is a mini-batch AUC reward: for each sample, we compute how well its prediction ranks relative to samples of the opposite class, normalized to  $[-1, 1]$ . Unlike a fixed-threshold indicator (e.g.,  $\hat{p} > 0.5$ ), this formulation is well-suited for imbalanced settings like lung cancer prediction where positive probabilities are typically low.

**Policy Gradient Optimization.** We adopt REINFORCE (Williams, 1992; Hu et al., 2023) over more complex algorithms (e.g., PPO (Schulman et al., 2017), GRPO (Shao et al., 2024)) due to the simplicity of our decision process: with at most three steps and a small discrete action space, the variance reduction from the batch-mean baseline and the auxiliary supervised loss provide sufficient stability. For a trajectory  $\boldsymbol{\tau} = (a_1, a_2, \dots)$  of sequential actions with log-probability  $\log \pi_{\theta}(\boldsymbol{\tau}) = \sum_t \log \pi_{\theta}(a_t | s_t)$ , the policy gradient loss is  $\mathcal{L}_{\text{PG}} = -\log \pi_{\theta}(\boldsymbol{\tau}) \cdot (r - \bar{r})$ , where  $\bar{r}$  is the mean reward over the mini-batch. The total training objective combines this with an entropy regularization term to encourage exploration:

$$\mathcal{L} = \mathcal{L}_{\text{PG}} - \lambda_{\text{ent}} \sum_t H(\pi_{\theta}(\cdot | s_t)) + \lambda_{\text{sup}} \mathcal{L}_{\text{BCE}} \quad (3)$$

where the final term is a supervised cross-entropy loss that directly optimizes the classifier, providing stable gradients independent of the stochastic policy. To accelerate convergence, we initialize the encoders and classifiers from pre-trained baseline models and use separate learning rates for different components.

## 3. Experiments

We conduct comprehensive experiments to evaluate AdaFuse on the lung cancer risk prediction task. Section 3.1 describes the dataset and feature extraction pipeline. Section 3.2 introduces the baseline fusion strategies. Section 3.3 compares AdaFuse against baseline methods. Section 3.4 presents ablation studies on training configurations.

### 3.1. Dataset

We evaluate on the National Lung Screening Trial (NLST) dataset (Team, 2011), a large-scale multi-center study of low-dose CT screening for lung cancer. To ensure fair evaluation, we use the held-out test set from Ardila et al. (Ardila et al., 2019) that was not seen during

Sybil model (Mikhael et al., 2023) training, since our CT image features are extracted using the pre-trained Sybil encoder. The dataset contains 1,847 patients for training and 462 patients for testing, with lung cancer prevalence of 6.44% and 6.06% respectively.

**Feature Extraction.** For each patient, we extract three modalities: (1) *CT image features* ( $f^A \in \mathbb{R}^{512}$ ): extracted from the Sybil model, a state-of-the-art lung cancer risk prediction network trained on NLST; (2) *Clinical variables* ( $f^B \in \mathbb{R}^{17}$ ): risk factors from the PLCO<sub>m2012</sub> model (Tammemägi et al., 2013) including age, smoking history, BMI, and family history; (3) *Text embeddings* ( $f^C \in \mathbb{R}^{768}$ ): we generate synthetic radiology reports from structured clinical variables covering occupational exposures (e.g., asbestos, chemical work, coal mining), medical history (e.g., diabetes, heart disease, hypertension), and secondhand smoke exposure, then extract embeddings using CORE (Yang et al., 2020), a BERT-based model (Lee et al., 2020) pre-trained on chest radiograph reports.

### 3.2. Baseline Fusion Strategies

We compare AdaFuse against 15 fixed fusion baselines covering all valid modality combinations, plus a MoE baseline.

**Single-Modality Baselines.** We train three single-modality models ( $A$ ,  $B$ ,  $C$ ) using only CT, clinical, or text features respectively.

**Operation-Based Fusion.** For multi-modality combinations, we evaluate three fusion operations: (1) *Concatenation*: features are concatenated along the channel dimension; (2) *Mean*: features are averaged element-wise after projection to a common dimension; (3) *Tensor*: features are fused via Kronecker product following the Tensor Fusion Network formulation (Zadeh et al., 2017). This yields 9 dual-modality models ( $AB$ ,  $AC$ ,  $BC \times 3$  fusion types) and 3 triple-modality models ( $ABC \times 3$  fusion types).

**Mixture-of-Experts (MoE).** We implement an MoE baseline where a gating network learns to softly weight the contributions of each modality. Unlike AdaFuse which makes discrete selection decisions, MoE processes all modalities and learns continuous weights.

### 3.3. Comparison with Baselines

Table 1 presents the test AUC of all methods. AdaFuse achieves the best performance (0.762 AUC), outperforming both the best fixed fusion baseline  $ABC$ -tensor (0.759) and the MoE baseline that processes all three modalities with learned soft weights (0.742). Several observations emerge from these results.

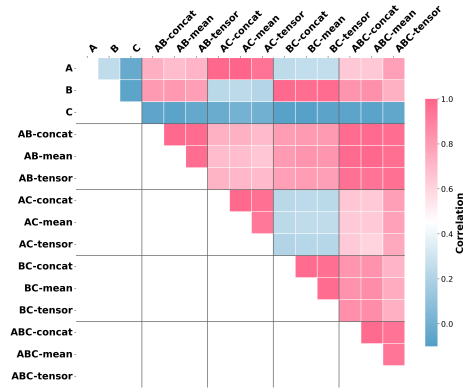


Figure 3: **Prediction correlation across baseline models.** Models containing CT features exhibit high mutual correlation, while text-only predictions show near-zero correlation with others.

Table 1: **Test AUC comparison on NLST.** We compare AdaFuse against single-modality baselines ( $A$ : CT image,  $B$ : clinical variables,  $C$ : text reports), operation-based fusion baselines with three fusion strategies (concatenation, mean pooling, tensor fusion), and a MoE baseline that learns soft modality weights. Background colors indicate modality count: single-modality, dual-modality, and triple-modality. The best AUC within each category and the overall best are shown in **bold**.

Method	AUC	Method	AUC
$A$ (CT)	<b>0.732</b>	$AC$ -tensor	0.739
$B$ (Clinical)	0.662	$BC$ -concat	0.661
$C$ (Text)	0.576	$BC$ -mean	0.678
$AB$ -concat	<b>0.758</b>	$BC$ -tensor	0.685
$AB$ -mean	0.755	$ABC$ -concat	0.735
$AB$ -tensor	0.735	$ABC$ -mean	0.748
$AC$ -concat	0.733	$ABC$ -tensor	<b>0.759</b>
$AC$ -mean	0.745	MoE	0.742
		<b>AdaFuse (Ours)</b>	<b>0.762</b>

**CT features dominate prediction.** The single-modality CT model ( $A$ , 0.732) already achieves competitive performance, outperforming many multi-modality fusion methods. This aligns with the fact that CT image features are extracted from Sybil, a model specifically trained for lung cancer risk prediction.

**Text reports provide limited information.** The text-only model ( $C$ , 0.576) performs near random, and combinations involving text without CT ( $BC$ ) consistently underperform. This is expected since our text reports are synthetically generated from structured clinical variables, which limits their informativeness compared to real radiology reports.

**Naive fusion can hurt performance.** Adding modalities does not guarantee improvement. For instance,  $ABC$ -concat (0.735) underperforms  $AB$ -concat (0.758), suggesting that indiscriminately incorporating all modalities can introduce noise rather than complementary information.

**AdaFuse learns to filter uninformative modalities.** Figure 3 shows the prediction correlation matrix across baseline models and the modality selection distribution of AdaFuse. We observe high correlation among models that include CT features, while text-only predictions show near-zero correlation with others. AdaFuse learns to leverage this structure: it predominantly selects CT-based combinations while adaptively incorporating clinical variables when beneficial, effectively filtering out the less informative text modality for most patients. Detailed analysis of the learned policy behavior is provided in Appendix Figure 5.

### 3.4. Ablation Studies

We conduct ablation studies to analyze the contribution of each component in AdaFuse. Section 3.4.1 investigates training configurations, and Section 3.4.2 examines the learning objective design.

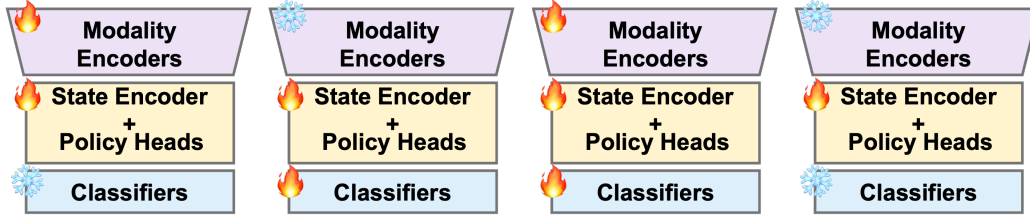


Figure 4: **Ablation study on training configurations.** From left to right: (1) freezing classifiers while training modality encoders; (2) freezing modality encoders while training classifiers; (3) training both components; (4) freezing both components. The flame icon indicates which components receive gradients during RL training. Quantitative results are provided in Table 2.

#### 3.4.1. TRAINING CONFIGURATION

Figure 4 illustrates the four training configurations, and Table 2 summarizes the results for freezing versus training encoders and classifiers during RL training.

Table 2: **Ablation study on training configurations.**  $\Delta$  denotes relative change compared to the best configuration.

Encoder	Classifier	Test AUC	$\Delta$
Train	Freeze	<b>0.762</b>	–
Freeze	Train	0.722	-5.25%
Train	Train	0.691	-9.32%
Freeze	Freeze	0.674	-11.55%

**Freezing classifiers yields the best performance.** As shown in Table 2, training classifiers with policy gradients degrades AUC from 0.762 to 0.722 (−5.25%), and jointly training both components further drops to 0.691 (−9.32%). This is because unfrozen classifiers only receive gradients when selected, causing undertrained combinations to produce unreliable rewards that further discourage their selection. Freezing pretrained classifiers provides stable reward signals and allows the policy to focus on selection without shifting decision boundaries.

**Training encoders is essential for policy learning.** The freeze-both configuration yields the worst performance (0.674, −11.55%), indicating that encoder adaptation is necessary. Unlike classifiers that provide fixed decision boundaries, encoders must learn representations that help the policy distinguish when each modality combination is beneficial for a given patient.

#### 3.4.2. LEARNING OBJECTIVE

We ablate the reward function (Eq. 2) and loss function (Eq. 3) to understand the contribution of each component. Table 3 presents the results.

Table 3: **Ablation study on learning objective.** We separately vary the loss composition (top) and reward design (bottom) while keeping the other fixed.  $\mathcal{L}_{\text{PG}}$ : policy gradient loss,  $\mathcal{H}$ : entropy regularization,  $\mathcal{L}_{\text{sup}}$ : supervised cross-entropy loss.  $\Delta$  denotes relative change compared to the best configuration.

	Configuration	Test AUC	$\Delta$
<b>Best</b>	$\mathcal{L}_{\text{PG}} + 0.1\mathcal{H} + 0.3\mathcal{L}_{\text{sup}}, r = 0.7r_{\text{BCE}} + 0.3r_{\text{AUC}}$	<b>0.762</b>	–
Same Reward	$\mathcal{L}_{\text{PG}} + 0.1\mathcal{H}$ (no supervision)	0.696	-8.7%
	$\mathcal{L}_{\text{PG}} + 0.3\mathcal{L}_{\text{sup}}$ (no entropy)	0.610	-20.0%
	$\mathcal{L}_{\text{PG}}$ (policy gradient only)	0.689	-9.6%
	$\mathcal{L}_{\text{PG}} + 1.0\mathcal{L}_{\text{sup}}$ (over-weighted supervision)	0.641	-15.9%
Same Loss	$r = r_{\text{AUC}}$ (AUC reward only)	0.647	-15.1%
	$r = r_{\text{BCE}}$ (BCE reward only)	0.637	-16.4%

**Both entropy and supervision are necessary for stable training.** Under the same reward, removing entropy regularization causes the largest performance drop from 0.762 to 0.610 (−20%), as the policy converges prematurely to a narrow set of modality combinations before exploring alternatives. Removing supervision degrades AUC to 0.696, since the supervised loss provides stable gradients independent of stochastic action sampling. However, over-weighting supervision ( $\lambda_{\text{sup}} = 1.0$ ) drops AUC to 0.641 by diminishing the influence of the reward signal.

**Mixed reward outperforms single-objective alternatives.** Under the same loss, using either BCE or AUC reward alone degrades AUC to 0.637 and 0.647, respectively. BCE provides dense per-sample feedback but can be misleading in imbalanced settings. AUC rewards correct ranking and is essential when positive prevalence is only 6%, but provides sparser feedback. The combination leverages the stability of BCE and the ranking awareness of AUC.

## 4. Conclusion

We presented AdaFuse, an adaptive multimodal fusion framework that addresses a fundamental question in multimodal learning: for a given patient, should certain modalities be used at all? By formulating modality selection as a sequential decision process, AdaFuse learns patient-specific fusion strategies with the flexibility to entirely exclude uninformative modalities rather than processing all inputs uniformly.

Our experiments on NLST reveal that the learned policy predominantly selects CT-based combinations while adaptively incorporating clinical variables, effectively filtering out the less informative text modality for most patients. This adaptive behavior outperforms fixed fusion strategies while demonstrating that uniform fusion is not always optimal. Ablation studies further show that freezing pretrained classifiers is essential for stable policy learning by providing reliable reward signals.

**Potential Limitations.** Our evaluation is limited to three modalities on a single dataset, where text is synthetically generated. Future work could extend AdaFuse to real clinical text and more modalities, and investigate when multimodal integration provides the most benefit across patient subgroups.

## Acknowledgements

This research was supported by the National Institutes of Health (NIH) through grants F30CA275020, 2U01CA152662, R01CA253923 (Landman & Maldonado), R01CA275015 (Maldonado & Lenburg), U01CA152662 (Grogan), U01CA196405 (Maldonado), and P30CA068485-29S1, as well as the National Science Foundation (NSF) through CAREER 1452485 and grant 2040462. Additional support was provided by the Vanderbilt Institute for Surgery and Engineering through T32EB021937-07, the Vanderbilt Institute for Clinical and Translational Research via UL1TR002243-06, the Pierre Massion Directorship in Pulmonary Medicine, and the American College of Radiology Fund for Collaborative Research in Imaging (FCRI) Grant. This manuscript was polished using AI-assisted editing (ChatGPT) with a “rephrase” prompt; no scientific content was generated or altered.

## References

- Julián N Acosta, Guido J Falcone, Pranav Rajpurkar, and Eric J Topol. Multimodal biomedical ai. *Nature medicine*, 28(9):1773–1784, 2022.
- Diego Ardila, Atilla P Kiraly, Sujeeth Bharadwaj, Bokyung Choi, Joshua J Reicher, Lily Peng, Daniel Tse, Mozziyar Etemadi, Wenxing Ye, Greg Corrado, et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nature medicine*, 25(6):954–961, 2019.
- John R Ball, Bryan T Miller, and Erin P Balogh. Improving diagnosis in health care. 2015.
- Richard Bellman. A markovian decision process. *Journal of mathematics and mechanics*, pages 679–684, 1957.
- Sean Blandin Knight, Phil A Crosbie, Haval Balata, Jakub Chudziak, Tracy Hussell, and Caroline Dive. Progress and prospects of early detection in lung cancer. *Open biology*, 7(9):170070, 2017.
- Asha Bonney, Reem Malouf, Corynne Marchal, David Manners, Kwun M Fong, Henry M Marshall, Louis B Irving, and Renee Manser. Impact of low-dose computed tomography (ldct) screening on lung cancer-related mortality. *Cochrane Database of Systematic Reviews*, (8), 2022.
- Bing Cao, Yiming Sun, Pengfei Zhu, and Qinghua Hu. Multi-modal gated mixture of local-to-global experts for dynamic image fusion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 23555–23564, 2023.
- Anika Cheerla and Olivier Gevaert. Deep learning with multimodal representation for pancancer prognosis prediction. *Bioinformatics*, 35(14):i446–i454, 2019.
- Richard J Chen, Ming Y Lu, Jingwen Wang, Drew FK Williamson, Scott J Rodig, Neal I Lindeman, and Faisal Mahmood. Pathomic fusion: an integrated framework for fusing histopathology and genomic features for cancer diagnosis and prognosis. *IEEE Transactions on Medical Imaging*, 41(4):757–770, 2020.

- Can Cui, Han Liu, Quan Liu, Ruining Deng, Zuhayr Asad, Yaohong Wang, Shilin Zhao, Haichun Yang, Bennett A Landman, and Yuankai Huo. Survival prediction of brain cancer with incomplete radiology, pathology, genomic, and demographic data. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 626–635. Springer, 2022.
- Sayan Ghosal, Qiang Chen, Giulio Pergola, Aaron L Goldman, William Ulrich, Karen F Berman, Giuseppe Blasi, Leonardo Fazio, Antonio Rampino, Alessandro Bertolino, et al. G-mind: an end-to-end multimodal imaging-genetics framework for biomarker identification and disease classification. In *Medical imaging 2021: image processing*, volume 11596, pages 63–72. SPIE, 2021.
- Xing Han, Huy Nguyen, Carl Harris, Nhat Ho, and Suchi Saria. Fusemoe: Mixture-of-experts transformers for fleximodal fusion. *Advances in Neural Information Processing Systems*, 37:67850–67900, 2024.
- V Hans Vitzthum von Eckstaedt, Andrea B Kitts, Christina Swanson, Michael Hanley, and Arun Krishnaraj. Patient-centered radiology reporting for lung cancer screening. *Journal of Thoracic Imaging*, 35(2):85–90, 2020.
- Mingzhe Hu, Jiahan Zhang, Luke Matkovic, Tian Liu, and Xiaofeng Yang. Reinforcement learning in medical image analysis: Concepts, applications, challenges, and future directions. *Journal of Applied Clinical Medical Physics*, 24(2):e13898, 2023.
- Shih-Cheng Huang, Anuj Pareek, Saeed Seyyedi, Imon Banerjee, and Matthew P Lungren. Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ digital medicine*, 3(1):136, 2020.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
- Quan Liu, Jiawen Yao, Lisha Yao, Xin Chen, Jingren Zhou, Le Lu, Ling Zhang, Zaiyi Liu, and Yuankai Huo. M 2 fusion: Bayesian-based multimodal multi-level fusion on colorectal cancer microsatellite instability prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 125–134. Springer, 2023.
- Zhun Liu, Ying Shen, Varun Bharadhwaj Lakshminarasimhan, Paul Pu Liang, AmirAli Bagher Zadeh, and Louis-Philippe Morency. Efficient low-rank multimodal fusion with modality-specific factors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2247–2256, 2018.
- Siyu Lu, Mingzhe Liu, Lirong Yin, Zhengtong Yin, Xuan Liu, and Wenfeng Zheng. The multi-modal fusion in visual question answering: a review of attention mechanisms. *PeerJ Computer Science*, 9:e1400, 2023.
- Peter G Mikhael, Jeremy Wohlwend, Adam Yala, Ludvig Karstens, Justin Xiang, Angelo K Takigami, Patrick P Bourgouin, PuiYee Chan, Sofiane Mrah, Wael Amayri, et al. Sybil:

- a validated deep learning model to predict future lung cancer risk from a single low-dose chest computed tomography. *Journal of Clinical Oncology*, 41(12):2191–2200, 2023.
- Pooya Mobadersany, Safoora Yousefi, Mohamed Amgad, David A Gutman, Jill S Barnholtz-Sloan, José E Velázquez Vega, Daniel J Brat, and Lee AD Cooper. Predicting cancer outcomes from histology and genomics using convolutional networks. *Proceedings of the National Academy of Sciences*, 115(13):E2970–E2979, 2018.
- Marcin Ostrowski, Tomasz Marjański, and Witold Rzyman. Low-dose computed tomography screening reduces lung cancer mortality. *Advances in medical sciences*, 63(2):230–236, 2018.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- M Patricia Rivera, Atul C Mehta, and Momen M Wahidi. Establishing the diagnosis of lung cancer: Diagnosis and management of lung cancer: American college of chest physicians evidence-based clinical practice guidelines. *Chest*, 143(5):e142S–e165S, 2013.
- Saurav Sahay, Eda Okur, Shachi H Kumar, and Lama Nachman. Low rank fusion based transformers for multimodal sequences. *arXiv preprint arXiv:2007.02038*, 2020.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Martin C Tammemägi, Hormuzd A Katki, William G Hocking, Timothy R Church, Neil Caporaso, Paul A Kvale, Anil K Chaturvedi, Gerard A Silvestri, Tom L Riley, John Commins, et al. Selection criteria for lung-cancer screening. *New England Journal of Medicine*, 368(8):728–736, 2013.
- National Lung Screening Trial Research Team. Reduced lung-cancer mortality with low-dose computed tomographic screening. *New England Journal of Medicine*, 365(5):395–409, 2011.
- Zhiqin Wang, Ruiqing Li, Minghui Wang, and Ao Li. Gpdbn: deep bilinear network integrating both genomic data and pathological images for breast cancer prognosis prediction. *Bioinformatics*, 37(18):2963–2970, 2021.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- Constance Monroe Winslow, Jacqueline B Kosecoff, Mark Chassin, David E Kanouse, and Robert H Brook. The appropriateness of performing coronary artery bypass surgery. *Jama*, 260(4):505–509, 1988.

Cheng Yang, Shengnan Wang, Yuechuan Li, Chao Yang, Ming Yan, Jingqiao Zhang, and Fangquan Lin. Core: An efficient coarse-refined training framework for bert. *arXiv preprint arXiv:2011.13633*, 2020.

Jordan Yap, William Yolland, and Philipp Tschandl. Multimodal skin lesion classification using deep learning. *Experimental dermatology*, 27(11):1261–1267, 2018.

Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor fusion network for multimodal sentiment analysis. *arXiv preprint arXiv:1707.07250*, 2017.

Hu Zhu, Ze Wang, Yu Shi, Yingying Hua, Guoxia Xu, and Lizhen Deng. Multimodal fusion method based on self-attention mechanism. *Wireless Communications and Mobile Computing*, 2020(1):8843186, 2020.

# Appendix

## Appendix A. Analysis of Learned Policy Behavior

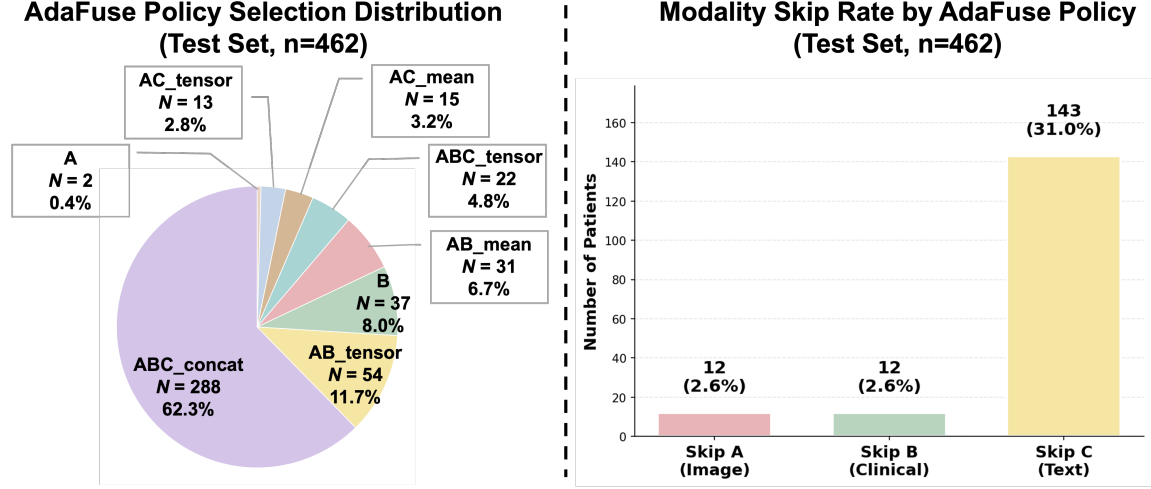


Figure 5: **AdaFuse policy selection distribution and modality skip rates on the test set.** Left: Distribution of modality-fusion combinations selected by the learned policy, where  $N$  denotes the number of patients. The policy most frequently selects ABC-concat ( $N = 288$ , 62.3%), followed by AB-tensor ( $N = 54$ , 11.7%) and clinical-only ( $N = 37$ , 8.0%). Right: Frequency of skipping each modality. The text modality (C) is skipped for 143 patients (31.0%), while CT (A) and clinical variables (B) are each skipped for only 12 patients (2.6%). This confirms that the policy learns to filter out the less informative text modality while consistently relying on imaging and clinical data.