

000 001 ICYMI²I: THE ILLUSION OF MULTIMODAL INFORMA- 002 TIVENESS UNDER MISSINGNESS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Multimodal learning is of continued interest in artificial intelligence-based applications, motivated by the potential information gain from combining different types of data. However, modalities observed in the source environment may differ from the modalities observed in the target environment due to multiple factors, including cost, hardware failure, or the perceived *informativeness* of a given modality. This shift in missingness between the source and target environment has not been carefully studied. Naïve estimation of the information gain associated with including an additional modality without accounting for missingness may result in improper estimates of that modality’s value in the target environment. We formalize the problem of missingness, demonstrate its ubiquity, and show that the subsequent distribution shift results in bias when the missingness process is not explicitly accounted for. To address this issue, we introduce ICYMI²I (In Case You Multimodal Missed It), a framework¹ for the evaluation of predictive performance and information gain under missingness through inverse probability weighting-based correction. We demonstrate the importance of the proposed adjustment to estimate information gain under missingness on synthetic, semi-synthetic, and real-world datasets.

029 1 INTRODUCTION

031 Multimodal learning is ubiquitous in machine learning as practitioners combine multiple data types to improve predictive performance in applications to healthcare (Perochon et al., 2023; Tu et al., 032 2024), robotics (Gao et al., 2024; Shah et al., 2023), and recommender systems (Chen et al., 2019). 033 However, factors such as privacy concerns (Jaiswal & Provost, 2020; Zhang et al., 2021), cost- 034 benefit tradeoffs of data-acquisition (Buck et al., 2010), and user preferences (Kossinets, 2006) 035 imprint multimodal data with missingness. Additionally, even if modality complete data is available 036 or curated at training, data noise (Cohen et al., 2004; Ma et al., 2023) and sensor failures (Inceoglu 037 et al., 2021; 2023) may result in missing modalities in the target environment.

039 **Although the missingness of modalities is a recurring challenge** in real-world settings, current multi- 040 modal machine learning methods often assume **that** modalities are fully observed, both in source and 041 target environments. When missingness is considered, the literature has focused on engineering ef- 042 forts (Le et al., 2025; Wu et al., 2024) such as data selection (Hosseini et al., 2022), imputation (Tran 043 et al., 2017; Cohen Kalafut et al., 2023; Malatesta et al., 2024), and architecture design (Chen et al., 044 2022; Zeng et al., 2022), which implicitly assume a stable missingness process between source and 045 target environments. When this assumption is violated, the missingness mechanism induces a dis- 046 tribution shift (Zhang et al., 2023; Liu et al., 2023b) that biases the estimated informativeness of 047 a given modality. Missingness is pervasive and impacts a broad range of application domains en- 048 countered in the multimodal literature: in breast cancer screening, biopsies are only performed if 049 there are abnormal findings in a mammogram; in autonomous vehicles, LiDAR sensor dropout can 050 occur due to weather and lighting conditions; and in online recommender systems, reviews are only 051 collected after certain consumer behaviors. Across these settings, ignoring the distribution shift be- 052 tween source and target due to missingness when quantifying modality informativeness may conflate 053 missingness with signal, leading to flawed data collection and modeling decisions.

¹Code available on Github: <https://anonymous.4open.science/r/ICYMI2I-BC18/>

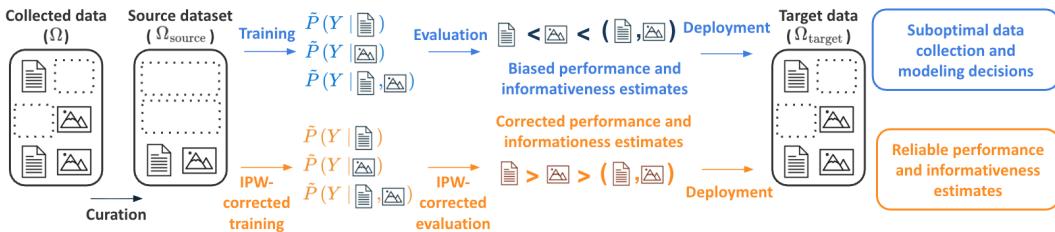


Figure 1. Overview of the proposed framework. Curation often discards missing data, resulting in a discrepancy between the collected Ω and source datasets Ω_{source} used for training. Current practice is denoted in blue: naïve training and evaluating on Ω_{source} leads to biased estimates of performance and informativeness on target data. The orange path illustrates the proposed ICYM²I: a double inverse probability weighting (IPW) mechanism that yields accurate performance and informativeness estimates under the target distribution.

In this work, we propose a framework to overcome the (mis)estimation of both inherent informativeness and predictive utility under missingness in multimodal learning. Our contributions, summarized in Figure 1, are as follows:

- **Framework for multimodal learning with missingness.** We formalize the impact of missingness as a distribution shift *intrinsic* to multimodal learning, where the *observed* source distribution differs from the target distribution due to missingness. We show that not accounting for missingness, a common practice, may bias the estimate of a modality’s predictive and information-theoretic utility.
- **ICYM²I.** Under the missingness-at-random (MAR) assumption, a much more realistic assumption than the common and often implicit assumption of missingness-completely-at-random (MCAR) made by state-of-the-art multimodal strategies, we propose ICYM²I (In Case You Multimodal Missed It), a double inverse-propensity weighting correction to overcome missingness-induced distribution shifts. Specifically, we demonstrate that ICYM²I improves correlation in predictive and information-theoretic utility of modalities.
- **Experiments on diverse data.** We demonstrate the broad applicability and utility of our methods in synthetic, semi-synthetic, and real-world benchmark datasets, including a case study in multimodal learning in health.

2 RELATED WORK

Multimodal benchmarks suppress missingness encountered in real-world environments. Prior work on multimodal models often assumes *fully observed modalities* (Ngiam et al., 2011; Zadeh et al., 2017; Hou et al., 2019). Missingness has largely been an overlooked problem (Le et al., 2025; Wu et al., 2024), to the extent that current benchmarks rarely contain samples with missing modalities. Curation often involves dropping incomplete or filtering samples based on data quality criteria, such as text length or file size (Sharma et al., 2018; Schuhmann et al., 2022) or imputing with automatic tools (Miech et al., 2019). This curation implicitly assumes that rejected samples follow the same distribution as the observed ones. This assumption may not hold. For example, in autonomous driving data, samples with sensor failure – often resulting from extreme weather or lighting conditions – may be filtered out. Models trained on complete data may, consequently, not generalize to these settings, creating real-world risk at deployment. When missingness is considered, previous works focused on robustness through imputation (Tran et al., 2017; Cohen Kalafut et al., 2023; Malitest et al., 2024), representation learning (Wu et al., 2024; Liu et al., 2023a), knowledge distillation (Li et al., 2024; Wang et al., 2020a), and model ensembling (Chen et al., 2022; Zeng et al., 2022) – all ignoring the potential shift resulting from the missingness process.

Multimodal missingness in the target environment. Prior work has explored missingness in the target environment (Lin & Hu, 2023; Zeng et al., 2022), e.g., when a captor fails at deployment (Ma et al., 2022). Broadly, two strategies have been proposed (Wu et al., 2024): (i) data preprocessing through cross-modal imputation (Cohen Kalafut et al., 2023; Malitest et al., 2024; Tran et al., 2017), where one replaces the missing modality (Ma et al., 2021; Zhou et al., 2022), as well as (ii) model training strategies such as architecture design (Lee et al., 2023; Ge et al., 2023), distillation-based methods (Li et al., 2024; Wang et al., 2020a), and ensembling (Chen et al., 2022; Zeng et al.,

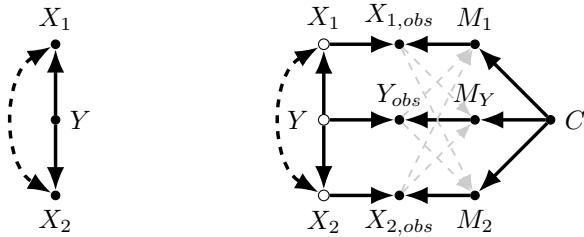
108 2022). Through the proposed formalization, our work distinguishes between different missingness
 109 assumptions, demonstrating that the previously studied framework is only one among various plau-
 110 sible mechanisms for which current strategies are not well-designed.

111 **Distribution shifts in multimodal learning.** Addressing multimodal shifts has been studied in
 112 vision-language models (Zhou et al., 2024; Verma et al., 2024) or using information-theoretic no-
 113 tions to understand multimodal behavior under distribution shifts (Oh et al., 2025). Augmentation
 114 and regularization strategies have been leveraged to address temporal shifts for conversation under-
 115 standing (Woo et al., 2023; Lian et al., 2023). Advances in learning, such as in-context learning,
 116 have been studied to characterize adaptation to multimodal distribution shifts (Zhou et al., 2024;
 117 Xue et al., 2024). However, existing strategies aim to improve robustness under *domain shifts only*,
 118 while ignoring the potential shift in missingness between source and target environments. *Instead,*
 119 *our work aims to correct estimates of performance and modality informativeness under missingness*
 120 *to inform modality collection at deployment.*

121 **Quantifying information-theoretic value of a modality.** Existing works often implicitly assume
 122 that additional modalities improve performance, ignoring the prohibitive cost, complexity, and po-
 123 tential noise added by these additional dimensions. When limited resources or constraints limit
 124 availability in the target environment, a central challenge is to quantify the information-theoretic
 125 value of a modality (Liang et al., 2024c). Liang et al. (2024a) proposed a method for recovering
 126 partial information decomposition measures of the redundancy, uniqueness, and synergy of the in-
 127 formation provided by the different modalities (Bertschinger et al., 2014; Williams & Beer, 2010).
 128 However, these decompositions fully ignore the impact of missingness.

129 **Correcting for missingness bias.** The lack of formalization of missingness in the multimodal
 130 literature has led to neglecting its potential impact. Ignoring this process risks biasing estimates of
 131 interest (Phelan et al., 2017) as the observed distribution differs from the underlying one practitioners
 132 aim to model. The statistical literature has introduced strategies such as matching (Stuart, 2010) and
 133 reweighting (Jethani et al., 2022) to correct for the missingness process. However, these works have
 134 overlooked the multimodal setting and the systematic shifts that may occur in this setting.

3 MULTIMODALITY AND MISSINGNESS



147 **Figure 2.** Directed Acyclic Graphs of the assumed data-generating processes. On the left is the com-
 148 monly assumed graph with no missingness. On the right is the proposed missingness formalism. X_1
 149 and X_2 are two modalities of interest, Y is the label of interest. The missingness process depends on
 150 C . Filled point nodes are observed variables, while unfilled nodes are unobserved. Gray edges indicate
 151 MAR missingness for a given modality.

152 Consider two modalities, $X_1 \in \mathcal{X}_1$ and $X_2 \in \mathcal{X}_2$ and the state of interest $Y \in \mathcal{Y}$. We denote the
 153 joint underlying distribution $\Omega = \mathcal{X}_1 \times \mathcal{X}_2 \times \mathcal{Y}$. Without loss of generality, we assume an anti-
 154 causal setting for the data-generating process, in which the modalities are dependent on the states
 155 Y , as shown in Figure 2 (left). We use the binary indicators of missingness M_1 , M_2 , and M_Y , which
 156 are equal to 1 if the associated variable is missing, 0 if observed, following the convention of Mohan
 157 & Pearl (2021). Observed variables are subscripted by ‘obs’, which corresponds to the underlying
 158 modality if observed, and unobserved otherwise (denoted by \emptyset). Formally, the observed variable
 159 Y_{obs} and observed modalities $X_{1,obs}$ and $X_{2,obs}$ can be defined as follows:

$$160 \quad 161 \quad Y_{obs} = \begin{cases} \emptyset & \text{if } M_Y = 1, \\ Y & \text{otherwise.} \end{cases}$$

162 In this setting, we denote the observed joint distribution $\Omega_{\text{obs}} := (M \cdot \mathcal{X}_1) \times (M \cdot \mathcal{X}_2) \times (M \cdot \mathcal{Y})$
 163 where $M = M_1 \cdot M_2 \cdot M_Y$. This complete modalities analysis has been the focus of multimodal
 164 learning.

165 **Missingness in multimodal learning.** We distinguish three mechanisms that cover potential **missing-
 166 modality in multimodal settings** (Rubin, 1976):
 167

- 168 • Missing Completely At Random (MCAR): A modality is missing completely at random if the
 169 missingness process is independent of any other variable.
- 170 • Missing At Random (MAR): The missingness mechanism depends on observed variables only.
- 171 • Missing Not At Random (MNAR): Missingness depends on unobserved variables.

173 In Figure 2 (right), we describe the missingness mechanisms as dependent on C , a set of covariates
 174 that determine the missingness mechanism. Note that C may include one of the modalities of interest,
 175 e.g., whether X_2 is observed may depend on the realization of X_1 . In general, the set C may
 176 differ for each modality depending on the data-generating process.

177 **Missingness-induced distribution shifts.** Missing modality X_i and/or missing label Y can induce
 178 distribution shifts between the source and the target distributions. For example, if a modality is
 179 observed in the target environment only if another one meets some criterion, then this distribution
 180 may not match the source distribution. Theoretically, we know that a non-MCAR missingness
 181 mechanism induces distribution shifts (Liu et al., 2023b; Zhang et al., 2023), i.e., the observed
 182 distribution differs from the underlying distribution. Critically, models trained and evaluated on the
 183 observed distribution are statistically biased estimates under any other missingness process.

184 For instance, consider an autonomous vehicle setting where video and LiDAR represent two modalities
 185 of interest. If the LiDAR randomly dysfunctions, the missingness patterns are MCAR. However,
 186 as previously mentioned, LiDAR may malfunction under extreme weather conditions. If these
 187 conditions can be extracted from the video modality, one may assume MAR patterns. However, if the
 188 video cannot capture the variable explaining the LiDAR dysfunction – e.g., temperature – LiDAR
 189 would be MNAR, as dependent upon an unobserved variable. Under the last two scenarios, focusing
 190 solely on samples with LiDAR excludes all extreme condition settings.

191 A common and often implicit assumption in the multimodal literature is the absence of missing-
 192 ness, which corresponds to either a MCAR mechanism – $\Omega_{\text{obs}} = \Omega$ – or a stable missingness
 193 process **between the source and target environments**, i.e., the observed distributions are the same
 194 $\Omega_{\text{obs}}^{\text{source}} = \Omega_{\text{obs}}^{\text{target}}$. In other words, not adjusting for missingness assumes that the missingness process
 195 is uninformative or will remain the same in the target environment.

196 **When missingness is studied**, prior works focus on improving the robustness of multimodal models
 197 when performance may degrade due to a modality missing in the target distribution,

199 **i.e., the source distribution reflects the true distribution while the target may present missingness**
 200 $\Omega_{\text{obs}}^{\text{source}} \sim \Omega \neq \Omega_{\text{obs}}^{\text{target}}$.

201 Our work questions the applicability of these assumptions where modality collection is costly. While
 202 missingness may **result in a distributional shift** (Zhou et al., 2023), we emphasize that demonstrating
 203 the value of a modality in the source environment may lead to increased collection of this modality
 204 in the target environment, inducing a distribution shift akin to **Assumption A**.

205 **Assumption A** (Multimodal analysis informs data collection). *Demonstrated multimodal perfor-
 206 mance gain induces a shift in the missingness process in the target, i.e. $\Omega_{\text{obs}}^{\text{source}} \neq \Omega^{\text{target}} \sim \Omega$.*

208 We focus on settings where historical data **used to** train a model is marked by missingness. Under
 209 such settings, we aim to do the following: (i) identify which modalities are informative and may
 210 consequently be **collected** in the target environment, and (ii) train models that generalize to **this**
 211 **target environment where the modalities are fully observed**.

212 4 IS THIS MODALITY INFORMATIVE?

215 We aim to assess whether a partially missing modality would be informative if fully observed. To
 this end, we introduce ICYM²I (In Case You Multimodal Missed It), a framework for correcting

216 model performance trained on modality complete samples where all modalities and labels are ob-
 217 served (Ω_{obs}) to estimate the predictive utility of the partially missing modality if it were observed
 218 for the whole population (Ω). Additionally, we propose a correction to derive unbiased estimates of
 219 the information-theoretic utility of a modality, using Ω_{obs} . We rely on Partial Information Decom-
 220 position (PID) (Williams & Beer, 2010) bounds introduced by Bertschinger et al. (2014) for this task,
 221 which quantifies the information value of a target of interest captured by two input variables.

222 **Correction.** We propose an Inverse Probability Weighting (IPW) approach (Robins et al., 1994),
 223 which reweights samples based on their probability of being observed. Under Assumption B that
 224 relaxes the common MCAR assumption made in the multimodal literature, IPW recovers unbiased
 225 estimates of the true distribution, enabling learning and evaluation on the true distribution from
 226 observed samples. IPW-adjustment is critical for both training and evaluation of multimodal models
 227 under missingness. IPW-adjusted training results in a model trained to infer on the underlying
 228 distribution Ω , while correction of the evaluation allows for measuring performance on Ω , despite
 229 evaluating the model only on samples from the observed distribution Ω_{obs} .

230 **Assumption B** (MAR and Positivity). *The missingness mechanism is MAR, and $p_{\Omega}(M_1 = 0, M_2 =$
 231 $0, M_Y = 0 | C) > 0$.*

233 4.1 A MOTIVATING EXAMPLE

235 We consider the common multimodal example of learning bit-wise logic operators (Bertschinger
 236 et al., 2014; Harder et al., 2013; Liang et al., 2024a). We generate 10,000 points with two modal-
 237 ities drawn from Bernoulli distributions ($p = 0.5$). The output state Y is defined using the binary
 238 operators AND, OR, and XOR of input bits X_1 and X_2 . In this setting, we induce missingness M_2
 239 in X_2 and Y as a function of X_1 (MAR): $M_2 \sim \text{Bern}(0.6X_1 + 0.2)$, resulting in 50% missingness
 240 in X_2 . We investigate the impact of missingness on current strategies for evaluating the predictive
 241 and information-theoretic utility of a given modality.

242 **Estimating performance for informativeness.** A common practice to measure the predictive value
 243 of adding a modality is through modality ablation studies where practitioners train models on the
 244 subset of observed samples where all modalities are observed (Ω_{obs}). First, unimodal models $f(\mathbf{x}_i)$
 245 are trained to approximate $p_{\Omega_{\text{obs}}}(y | x_i)$, $\forall i \in \{1, 2\}$ and a multimodal model $f(\mathbf{x}_1, \mathbf{x}_2)$ to ap-
 246 proximate $p_{\Omega_{\text{obs}}}(y | x_1, x_2)$ on the same observed dataset. The performance is then compared in a
 247 hold-out set, which is also sampled from Ω_{obs} . The informativeness of a modality is attributed to
 248 the relative performance gain of the multimodal model compared to the unimodal model. However,
 249 multimodal models can perform worse than their unimodal counterparts due to data characteris-
 250 tics (Zhang et al., 2024) and learning dynamics (Wang et al., 2020b; Zhai et al., 2024). Thus, relying
 251 solely on performance as a proxy for informativeness, particularly under distribution shifts, can be
 252 misleading.

253 **Partial Information Decomposition** (Bertschinger et al., 2014). As an alternative to estimating
 254 performance, existing works have decomposed the informativeness associated with each modality
 255 (Liang et al., 2024a). Bertschinger et al. (2014) formalized this decomposition by analyzing the
 256 total (three-way) mutual information $I(Y : (X_1, X_2))$ (McGill, 1954; Te Sun, 1980), a measure of
 257 dependency between the target variable Y and the modalities (X_1, X_2) , decomposing it into shared
 258 information (information both X_1, X_2 share about Y), unique information 1 (information only X_1
 259 has about Y), unique information 2 (information only X_2 has about Y), and complementary infor-
 260 mation (information about Y that requires both X_1 and X_2) as follows:

$$261 I(Y : (X_1, X_2)) = \underbrace{SI(Y : X_1; X_2)}_{\text{shared information}} + \underbrace{UI(Y : X_1 \setminus X_2)}_{\text{unique information 1}} + \underbrace{UI(Y : X_2 \setminus X_1)}_{\text{unique information 2}} + \underbrace{CI(Y : X_1; X_2)}_{\text{complementary information}}$$

264 Bertschinger et al. (2014) specifies how to estimate these quantities. For instance, Bertschinger et al.
 265 (2014) show that the unique information between Y and X_1 can be estimated using the following:
 266

$$267 \widetilde{UI}(Y : X_1 \setminus X_2) = \min_{q \in \Delta_{\Omega}} [I_q(Y : (X_1, X_2)) - I_q(Y : X_2)],$$

270 where Δ_Ω is the set of joint distributions over (X_1, X_2, Y) such that, $q(X_i = x_i, Y = y) = p_\Omega(X_i = x_i, Y = y) \forall x_i \in \mathcal{X}_i, y \in \mathcal{Y}, i \in \{1, 2\}$, that is, the set of joint distributions that match 271 the true two-way data distributions. Notice that the objective function requires a minimization over 272 the three-way mutual information. Approximations for all other entities in the decomposition are 273 in Appendix B. Importantly, all approximations require minimizing the three-way mutual information, 274 and Bertschinger et al. (2014) demonstrates that the solution to any one objective specifies an 275 optimum for all decompositions.

276 Prior work that relies on this Partial Information Decomposition (PID) to attribute information- 277 theoretic value implicitly assumes that $\Omega_{\text{obs}}^{\text{source}} = \Omega^{\text{source}} = \Omega^{\text{target}} = \Omega$. Instead, we evidence the 278 limitations of these strategies performed on $\Omega_{\text{obs}}^{\text{source}}$ when the target decomposition is $\Omega^{\text{target}} = \Omega$, 279 i.e., the true data-generating mechanism.

280 **Table 1.** Impact of missingness on multimodality information for bitwise logic operators. Parentheses 281 denote standard deviation across batches.

		AUROC			Information Decomposition			
		X_1	X_2	$X_1 + X_2$	Unique 1	Unique 2	Shared	Complementary
AND	Oracle	0.83 (0.01)	0.84 (0.01)	1.00 (0.00)	0.05 (0.00)	0.03 (0.00)	0.26 (0.00)	0.47 (0.00)
	Observed	0.66 (0.01)	0.93 (0.01)	1.00 (0.00)	0.44 (0.00)	0.00 (0.00)	0.15 (0.00)	0.36 (0.00)
	ICYM ² I	0.83 (0.01)	0.85 (0.02)	1.00 (0.00)	0.03 (0.00)	0.03 (0.00)	0.27 (0.00)	0.45 (0.00)
OR	Oracle	0.84 (0.01)	0.83 (0.01)	1.00 (0.00)	0.04 (0.00)	0.05 (0.00)	0.27 (0.00)	0.46 (0.00)
	Observed	0.95 (0.01)	0.77 (0.01)	1.00 (0.00)	0.01 (0.00)	0.15 (0.00)	0.10 (0.00)	0.23 (0.00)
	ICYM ² I	0.85 (0.02)	0.82 (0.01)	1.00 (0.00)	0.03 (0.00)	0.02 (0.00)	0.27 (0.00)	0.50 (0.00)
XOR	Oracle	0.51 (0.02)	0.49 (0.01)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.99 (0.00)
	Observed	0.52 (0.02)	0.80 (0.02)	1.00 (0.00)	0.34 (0.00)	0.07 (0.00)	-0.07 (0.00)	0.62 (0.00)
	ICYM ² I	0.53 (0.03)	0.49 (0.03)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.01 (0.00)	0.96 (0.00)

295 As a motivating example, we analyze the impact of missingness on estimating PID in the case 296 of unidimensional modalities with a bitwise logic outcome (AND, OR, and XOR). Table 1 (left) 297 presents the discriminative performance associated with neural networks trained on each modality 298 and their combination under three scenarios: (i) access to all data (**Oracle**), (ii) focusing only on 299 datapoints with all covariates observed (**Observed**), and (iii) adequately accounting for missingness 300 (ICYM²I using IPW to adjust $\Omega_{\text{obs}} \mapsto \Omega$, by modeling the missingness mechanism), as proposed 301 in Section 4.2. Table 1 (right) presents PID, discussed in Section 4.3 under the same scenarios, 302 demonstrating how information decomposition is also biased due to missingness.

303 Specifically, relying on Ω_{obs} overestimates the performance of X_1 for OR but underestimates it for 304 AND. Similarly, biased decomposition results in overestimating the informativeness of X_1 (“Unique 305 1” compared to “Unique 2”) for OR. As X_1 informs the missingness process, it indirectly informs 306 the outcome of interest, despite the true underlying generative process being dependent on both. 307 The use of IPW can correct for such bias under positivity as long as the propensities for IPW can 308 be estimated (i.e., the MAR assumption). We study sensitivity to this assumption in Appendix D, 309 where we further evaluate the robustness of our method under MCAR and MNAR, demonstrating 310 robustness under MCAR.

311 We now formally describe two methods for reliably inferring the informativeness of modalities using 312 (i) unbiased estimation of unimodal versus multimodal model performance using supervised learning 313 (ICYM²I-learn), and (ii) high-dimensional autodifferentiable partial information decomposition 314 (ICYM²I-PID). In addition, we demonstrate the need for IPW-adjusted *evaluation* as a key element 315 to determine modality informativeness using supervised learning.

316 4.2 ICYM²I-LEARN: ESTIMATING PERFORMANCES FOR INFORMATIVENESS UNDER 317 MISSINGNESS

318 **Training.** Under the MAR assumption, i.e., the missingness is fully explained by observed covariates C ; that is, the probability of a data point being missing depends only on C , we propose to train 319 the model with a weighted loss using samples from Ω_{obs} . The proposed IPW-adjusted loss accounts 320 for the distributional shift ($\Omega_{\text{obs}} \mapsto \Omega$) by up-weighting under-observed points, as described in the 321 following lemma.

324 **Lemma 1** (IPW Training). *The loss function computed on the observed data $l_{\Omega_{obs}}(x_1, x_2, y)$ can be
325 reweighted to approximate the target loss $l_{\Omega}(x_1, x_2, y)$ as follows:*

$$327 \quad l_{\Omega}(x_1, x_2, y) = \frac{1}{1 - p(m_1, m_2, m_y \mid C)} l_{\Omega_{obs}}(x_1, x_2, y)$$

329 where $p(m_1, m_2, m_y \mid C)$ is the probability of missingness, given the covariates C .

330 **Evaluation.** Existing works suffer from an analogous bias in model evaluation, by relying on a
331 hold-out set from the observed distribution (Ω_{obs}). To estimate a given metric on the true underlying
332 distribution, one must correct this metric using a similar correction as previously described. Li et al.
333 describes how to correct for both AUC and Brier score using IPW.

334 **Corollary 1** (ICYM²I-learn). *Consider a model f trained and evaluated on data drawn from Ω_{obs} .
335 To correct the model and estimate its performance on Ω , one must correct both its training and
336 evaluation following the previous corrections.*

338 4.3 ICYM²I-PID: PARTIAL INFORMATION DECOMPOSITION FOR MULTIMODAL 339 INFORMATIVENESS

341 Under missingness, we have samples from Ω_{obs} instead of Ω . Estimating PID measures in this
342 setting requires adjusting for the $\Omega_{obs} \mapsto \Omega$ shift. Our approach introduces a correction to ensure
343 that we optimize an unbiased estimate of the three-way mutual information using samples from Ω_{obs} .

344 **Lemma 2** (Corrected mutual information).

$$346 \quad I_{\Omega}(Y : (X_1, X_2)) = \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y|x_1, x_2)}} \left[\frac{1 - p(m_1, m_2)}{1 - p(m_1, m_2|x_1, x_2, y)} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right]$$

349 See Appendix A for the complete proof. Effectively, our approach corrects the optimization pro-
350 posed by Liang et al. (2024a) to account for the distribution shift induced by missing modalities. In
351 other words, ICYM²I-PID solves:

$$352 \quad \min_{q \in \Delta_{\Omega}^{\text{ICYM}^2\text{I}}} I_q(Y : (X_1, X_2))$$

355 where

$$356 \quad \Delta_{\Omega}^{\text{ICYM}^2\text{I}} = \{q \propto \exp(f_1(x_1) \cdot f_2(x_2)) : q(x_i, y) = \text{IPW}_{\Omega}(p_{\phi}(y, x_i)) \forall x_i \in \mathcal{X}, y \in \mathcal{Y}, i \in \{1, 2\}\}$$

357 where, $\text{IPW}_q(p)$ is an IPW correction for the shift $p \mapsto q$ using samples from p and p_{ϕ} is a re-
358 parametrization of Ω using neural networks and learned using samples from Ω_{obs}

359 via the weighted loss introduced in Lemma 1. Importantly, the proposed correction is agnostic to
360 the parametrisation of q . The resulting PID estimation framework consists of the following steps
361 (detailed in Appendix C):

- 363 **1. Model the missingness mechanism.** Train a model to estimate the probability of missing-
364 ness given C to obtain importance weights for correcting the distribution shift.
- 365 **2. Train corrected unimodal and multimodal models.** Train each model ($f_1(x_1)$, $f_2(x_2)$,
366 and $f(x_1, x_2)$) using the IPW-corrected loss introduced in Lemma 1. These models can be
367 trained using flexible inductive biases depending on the application of interest.
- 368 **3. Solve PID optimization.** Estimate $q \in \Delta_{\Omega}^{\text{ICYM}^2\text{I}}$ ² minimizing $I_q(Y : (X_1, X_2))$, where
369 q is parameterized by the product of two unimodal networks. To enforce the marginal
370 constraints $q(x_i, y)$, we apply a modified Sinkhorn–Knopp procedure (Knight, 2008) using
371 IPW-corrected unimodal distributions². Note that the parametrization of q chosen for this
372 method follows Liang et al. (2023) due to the flexibility of implementation. However, our
373 framework is agnostic to the choice of this parametrization. The choice of this parametriza-
374 tion for fusion is application-driven, as long as calibrated probabilistic scores are learned.

375 2

376 ²Prior work typically matches $q(y \mid x_i)$ to $\Omega_{obs}(y \mid x_i)$, which yields biased PID estimates under missing-
377 ness.

378 4. **Estimate the PID components.** Given q , compute PID quantities using the bounds
 379 of Bertschinger et al. (2014), corrected via the IPW-correction introduced in Appendix A.
 380

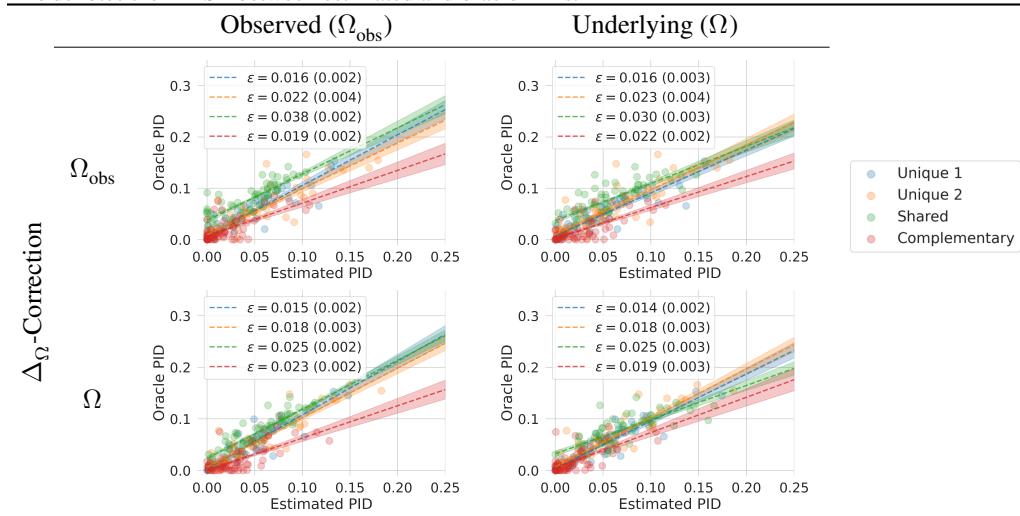
381 **5 EXPERIMENTS**

383 To better understand the connection between performance, information decomposition and miss-
 384 ingness, we propose a simulation (detailed in Appendix E), two semi-synthetic studies that reflect
 385 real-world missingness mechanisms (see Appendix F), and a real-world case-study.
 386

387 **5.1 SIMULATION AND SEMI-SYNTHETIC EXPERIMENTS**

389 In Table 2, each point represents the estimated PID value (Unique 1, Unique 2, Shared, and Com-
 390plementary) for one simulation under the training and evaluation IPW-corrections and the oracle
 391 performance, i.e., a model trained and tested on Ω . Specifically, columns reflect evaluation correc-
 392tion, while rows reflect training correction. These results underline the importance of correcting
 393 both training and evaluation, as proposed in ICYM²I, to best align with the performance one would
 394 obtain on Ω , as shown by the smallest Root Mean Squared Error (RMSE) observed when both
 395 corrections are applied. This observation shows that the proposed ICYM²I best recovers the true
 396 informativeness of each modality, despite relying on Ω_{obs} . Appendix E echoes the same observation
 397 when evaluating model performance.

398 **Table 2.** Comparison between estimated PID using training and PID corrections, and oracle PID on Ω .
 399 ϵ denotes the RMSE between estimated and oracle PIDs.



400 The semi-synthetic experiments examine the effect of enforcing increasing missingness on the
 401 performance and information decomposition of UR-FUNNY (Hasan et al., 2019) and hateful
 402 memes (Kiela et al., 2020), two foundational real-world datasets used in the multimodal litera-
 403 ture for affective computing and content moderation. Table 3 summarizes the effect of enforcing
 404 70% missingness on estimating multimodality informativeness across these datasets, demon-
 405 strating the generalizability of our proposed strategy across real-world datasets. Appendix F further
 406 illustrates the robustness of the methodology under different levels of missingness in these datasets and
 407 explores MNAR patterns.

408 **5.2 CHEST RADIOGRAPHS ARE UNINFORMATIVE OVER ELECTROCARDIOGRAMS FOR
 409 STRUCTURAL HEART DISEASE DETECTION.**

410 While our core contribution is methodological, this section illustrates how ignoring missingness can
 411 lead to biased estimates of the informativeness of a given modality in a real-world setting where
 412 modalities are commonly missing. Specifically, we study structural heart disease (SHD), a set
 413 of conditions that affect the heart's physiology, which is typically diagnosed using transthoracic

432
433
434
435
436
437
438
439
440
441
442
443**Table 3.** Impact of 70% missingness on multimodality information for UR-FUNNY (Hasan et al., 2019) and Hateful Memes (Kiela et al., 2020). Parentheses denote standard deviation across batches.

		AUROC			Information Decomposition			
		Text	Image/Video	Image + Text	Unique Text	Unique Image	Shared	Complementary
UR-FU.	Oracle	0.68 (0.01)	0.60 (0.02)	0.69 (0.02)	0.10 (0.00)	0.02 (0.00)	0.00 (0.00)	0.00 (0.00)
	Observed	0.61 (0.03)	0.54 (0.04)	0.63 (0.03)	0.05 (0.00)	0.00 (0.00)	0.03 (0.00)	0.00 (0.00)
	ICYM ² I	0.66 (0.03)	0.57 (0.04)	0.62 (0.04)	0.07 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Memes	Oracle	0.71 (0.01)	0.57 (0.01)	0.72 (0.01)	0.09 (0.01)	0.00 (0.00)	0.04 (0.00)	0.05 (0.01)
	Observed	0.68 (0.02)	0.61 (0.02)	0.71 (0.02)	0.13 (0.00)	0.04 (0.00)	0.01 (0.00)	0.00 (0.00)
	ICYM ² I	0.67 (0.02)	0.61 (0.02)	0.71 (0.02)	0.10 (0.00)	0.01 (0.00)	0.02 (0.03)	0.03 (0.01)

echocardiograms (TTEs) (Writing Committee Members et al., 2021). However, TTEs are often underutilized in the United States due to diagnostic stewardship and competing financial incentives (Papolos et al., 2016). Prior work using unimodal models with common modalities in electrocardiograms (ECGs) (Elias et al., 2022; Ulloa-Cerna et al., 2022) and chest radiographs (CXR) (Bhave et al., 2024) has demonstrated that non-TTE modalities can detect structural heart disease labels. However, CXRs are not systematically collected in conjunction with ECGs, leading to systematic missingness patterns. We, therefore, evaluate ICYM²I on this clinical task to evaluate the informativeness of CXRs in diagnosing SHD, despite its missingness.

Dataset. Our study population consists of a retrospective study gathering 98,397 adult patients who received an ECG and a TTE within one year of each other. The population has 20.56% SHD prevalence. In this cohort, 12,587 members (12.79%) have recorded CXRs. For subjects with multiple echocardiograms, we select the first TTE to model opportunistic screening with non-TTE modalities. All data were collected from an academic urban medical system between 2008 and 2022. Data are split temporally, where subjects with TTEs collected on or after 2018 ($n = 40,734$) are allocated to the test set. All data were de-identified, retrospective, and collected for clinical purposes from an academic hospital system, with approval from the Institutional Review Board. Appendix G contains further details regarding preprocessing, [embedding generation](#), and the [ICYM²I implementation](#).

Results. Table 4 presents the performance of each uni- and multimodal model, along with the associated information decomposition. While both the observed and corrected analyses demonstrate the importance of ECG in modeling SHD, the corrected results raise questions about the information gain associated with CXR. Naive decomposition suggests the unique information in CXRs at about 5% of the total information. However, ICYM²I reduces this unique contribution to 1.8% while increasing estimates of shared information between ECG and CXRs for SHD detection. In contrast to domain knowledge, where ECGs capture electrophysiology while CXRs capture structure and anatomy, two distinct aspects of cardiac health, the corrected complementary and shared results, and low unique information of CXRs suggest that CXRs are not independently useful for SHD diagnosis. Note that our results indicate that the multimodal model performs slightly worse than the unimodal ECG model, reflecting the potential overfitting risk associated with a large number of features.

Table 4. Informativeness of ECG and CXR modalities on model-based structural heart disease detection. Parentheses denote standard deviation across batches ($n = 1024$).

	AUROC			Information Decomposition			
	ECG	CXR	ECG + CXR	Unique ECG	Unique CXR	Shared	Complementary
Observed	0.83 (0.01)	0.72 (0.02)	0.82 (0.01)	0.11 (0.00)	0.01 (0.00)	0.10 (0.00)	0.00 (0.00)
ICYM ² I	0.82 (0.01)	0.73 (0.02)	0.83 (0.01)	0.07 (0.00)	0.01 (0.00)	0.48 (0.00)	0.01 (0.00)

6 DISCUSSION

This work formalizes the issue of [partially observed modalities in multimodal settings](#). We emphasize that existing works commonly overlook missingness by discarding samples with [any missing modality](#) at the curation stage, or implicitly assume that the missingness mechanism remains stable when a model is deployed in the target environment. Our work formalizes this problem and

486 demonstrates its ubiquity in the multimodality literature. Most critically, prior work ignores that any
 487 perceived informativeness of a modality may result in increased rates of data collection, inducing
 488 different missingness patterns at deployment. Our work, therefore, introduces ICYM²I, a correction
 489 to estimate the information gain associated with a *partially observed modality*. Our results demon-
 490 strate the methodology’s capacity to correct for biases introduced by missingness across synthetic,
 491 semi-synthetic, and real-world multimodal datasets. Finally, we **demonstrate** the practical utility of
 492 this methodology in a healthcare dataset, **showing the divergent** conclusions that one would reach
 493 if ignoring missingness. Our work highlights the critical importance of missingness in multimodal
 494 research and urges practitioners to pay particular attention to this issue by systematically *collecting*
 495 data with incomplete modalities and carefully *modeling* and *accounting* for missingness to enhance
 496 robustness.
 497

Limitations. The key assumption in our work is that **missingness is MAR**. No theoretical guarantees exist under MNAR patterns. While distinguishing these assumptions is empirically untestable, practitioners should ensure that this assumption is appropriate for their data. Importantly, MAR is less restrictive than the implicit MCAR assumption made in the multimodal literature, and does not require unrealistic distributional assumptions that one must assume to tackle MNAR patterns. Additionally, our work is based on Partial Information Decomposition (PID Bertschinger et al. (2014)), which focuses on two input modalities. In practice, practitioners could consider a one-vs-all approach to inform modality informativeness using our method. However, extending the decomposition to more than two modalities remains an open challenge (Griffith & Koch, 2014; Kolchinsky, 2022) where notions of mutual information itself are not well outlined beyond three-way mutual information. As in prior work on PID-based measures of information gain on high-dimensional data (Liang et al., 2023; 2024a), the quality of the representations used may impact the measures returned by ICYM²I. We ensure that our probabilistic estimates are calibrated to mitigate any such challenges.
 500

Ethics statement. Our work demonstrates the impact of missingness on performance estimates in multimodal learning. We demonstrate the utility of our method in a crucial healthcare use case. However, the methodology remains a proof of concept that would require additional testing to be deployed in a real-world context. Our study is approved by the [Anonymized] Institutional Review Board. We have extracted all data in HIPAA-compliant servers, and our experiments are also conducted on HIPAA-compliant compute despite being deidentified for extra caution. While beyond the scope of this work, modality completeness is not uniform across demographic subgroups and can manifest in data collection policies, such as differential access to care based on insurance status. Our method could provide important insights into the utility of multimodal predictions in such settings. Finally, our proposed method relies on a notion of instance for which all modalities can be observed. Extending this method when there is no notion of an instance, i.e., unaligned modalities, could be considered but requires different inductive biases to model the underlying unimodal and multimodal probabilities.
 510

Reproducibility statement. Theoretical proofs are provided in Appendix A. All code for applying the proposed ICYM²I and reproducing all synthetic and semi-synthetic results presented in this work is publicly available on Github³. A summary of the computational resources required to reproduce our results is given in Appendix H.
 520

528 REFERENCES

530 Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S.
 531 Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew
 532 Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath
 533 Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris
 534 Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker,
 535 Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wat-
 536 tenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learn-
 537 ing on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software
 538 available from tensorflow.org.
 539

³ <https://anonymous.4open.science/r/ICYM2I-BC18/>

540 Tadas Baltrušaitis, Peter Robinson, and Louis-Philippe Morency. Openface: an open source facial
 541 behavior analysis toolkit. In *2016 IEEE winter conference on applications of computer vision*
 542 (*WACV*), pp. 1–10. IEEE, 2016.

543

544 Nils Bertschinger, Johannes Rauh, Eckehard Olbrich, Jürgen Jost, and Nihat Ay. Quantifying unique
 545 information. *Entropy*, 16(4):2161–2183, 2014.

546

547 Shreyas Bhave, Victor Rodriguez, Timothy Poterucha, Simukayi Mutasa, Dwight Aberle, Kathleen
 548 M Capaccione, Yibo Chen, Belinda Dsouza, Shifali Dumeer, Jonathan Goldstein, et al. Deep
 549 learning to detect left ventricular structural abnormalities in chest x-rays. *European Heart Journal*,
 550 pp. ehad782, 2024.

551

552 Andreas K Buck, Ken Herrmann, Tom Stargardt, Tobias Dechow, Bernd Joachim Krause, and Jonas
 553 Schreyögg. Economic evaluation of pet and pet/ct in oncology: evidence and methodologic ap-
 554 proaches. *Journal of nuclear medicine technology*, 38(1):6–17, 2010.

555

556 Lei Chen and Chungmin Lee. Predicting audience’s laughter during presentations using convo-
 557 lutional neural network. In *Proceedings of the 12th Workshop on Innovative Use of NLP for*
Building Educational Applications, pp. 86–90, 2017.

558

559 Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler,
 560 Huan Zhao, and Binqiang Zhao. Pog: personalized outfit generation for fashion recom-
 561 mendation at alibaba ifashion. In *Proceedings of the 25th ACM SIGKDD international conference on*
knowledge discovery & data mining, pp. 2662–2670, 2019.

562

563 Yi-Ting Chen, Jinghao Shi, Zelin Ye, Christoph Mertz, Deva Ramanan, and Shu Kong. Multimodal
 564 object detection via probabilistic ensembling. In *European Conference on Computer Vision*, pp.
 565 139–158. Springer, 2022.

566

567 Ira Cohen, Fabio Gagliardi Cozman, Nicu Sebe, Marcelo Cesar Cirelo, and Thomas S Huang.
 568 Semisupervised learning of classifiers: Theory, algorithms, and their application to human-
 569 computer interaction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(12):
 570 1553–1566, 2004.

571

572 Noah Cohen Kalafut, Xiang Huang, and Daifeng Wang. Joint variational autoencoders for multi-
 573 modal imputation and embedding. *Nature machine intelligence*, 5(6):631–642, 2023.

574

575 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Ad-
 576 vances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf.

577

578

579 Gilles Degottex, John Kane, Thomas Drugman, Tuomo Raitio, and Stefan Scherer. Covarep—a
 580 collaborative voice analysis repository for speech technologies. In *2014 ieee international con-
 581 ference on acoustics, speech and signal processing (icassp)*, pp. 960–964. IEEE, 2014.

582

583 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
 584 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
 585 pp. 248–255. Ieee, 2009.

586

587 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
 588 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of*
589 the North American chapter of the association for computational linguistics: human language
590 technologies, volume 1 (long and short papers), pp. 4171–4186, 2019.

591

592 Pierre Elias, Timothy J Poterucha, Vijay Rajaram, Luca Matos Moller, Victor Rodriguez, Shreyas
 593 Bhave, Rebecca T Hahn, Geoffrey Tison, Sean A Abreau, Joshua Barrios, et al. Deep learning
 electrocardiographic analysis for detection of left-sided valvular heart disease. *Journal of the*
American College of Cardiology, 80(6):613–626, 2022.

594 Jensen Gao, Bidipta Sarkar, Fei Xia, Ted Xiao, Jiajun Wu, Brian Ichter, Anirudha Majumdar, and
 595 Dorsa Sadigh. Physically grounded vision-language models for robotic manipulation. In *2024*
 596 *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12462–12469. IEEE,
 597 2024.

598 Chongjian Ge, Junsong Chen, Enze Xie, Zhongdao Wang, Lanqing Hong, Huchuan Lu, Zhenguo
 599 Li, and Ping Luo. Metabev: Solving sensor failures for 3d detection and map segmentation.
 600 In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8721–8731,
 601 2023.

602 Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G
 603 Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank,
 604 physiotoolkit, and physionet: components of a new research resource for complex physiologic
 605 signals. *circulation*, 101(23):e215–e220, 2000.

606 Virgil Griffith and Christof Koch. Quantifying synergistic mutual information. In *Guided self-
 607 organization: inception*, pp. 159–190. Springer, 2014.

608 Malte Harder, Christoph Salge, and Daniel Polani. Bivariate measure of redundant information.
 609 *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 87(1):012130, 2013.

610 Md Kamrul Hasan, Wasifur Rahman, AmirAli Bagher Zadeh, Jianyuan Zhong, Md Iftekhar Tanveer,
 611 Louis-Philippe Morency, and Mohammed (Ehsan) Hoque. UR-FUNNY: A multimodal language
 612 dataset for understanding humor. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.),
 613 *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and*
 614 *the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp.
 615 2046–2056, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:
 616 10.18653/v1/D19-1211.

617 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
 618 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.
 619 770–778, 2016.

620 Seyedmajid Hosseini, Raju Gottumukkala, Satya Katragadda, Ravi Teja Bhupatiraju, Ziad Ashkar,
 621 Christoph W Borst, and Kenneth Cochran. A multimodal sensor dataset for continuous stress
 622 detection of nurses in a hospital. *Scientific Data*, 9(1):255, 2022.

623 Ming Hou, Jiajia Tang, Jianhai Zhang, Wanzeng Kong, and Qibin Zhao. Deep multimodal multi-
 624 linear fusion with high-order polynomial pooling. *Advances in Neural Information Processing
 625 Systems*, 32, 2019.

626 Arda Inceoglu, Eren Erdal Aksoy, Abdullah Cihan Ak, and Sanem Sariel. Fino-net: A deep mul-
 627 timodal sensor fusion framework for manipulation failure detection. In *2021 IEEE/RSJ interna-
 628 tional conference on intelligent robots and systems (IROS)*, pp. 6841–6847. IEEE, 2021.

629 Arda Inceoglu, Eren Erdal Aksoy, and Sanem Sariel. Multimodal detection and classification of
 630 robot manipulation failures. *IEEE Robotics and Automation Letters*, 9(2):1396–1403, 2023.

631 Mimansa Jaiswal and Emily Mower Provost. Privacy enhanced multimodal neural representations
 632 for emotion recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol-
 633 ume 34, pp. 7985–7993, 2020.

634 Neil Jethani, Aahlad Puli, Hao Zhang, Leonid Garber, Lior Jankelson, Yindalon Aphinyanaphongs,
 635 and Rajesh Ranganath. New-onset diabetes assessment using artificial intelligence-enhanced elec-
 636 trocardiography. *arXiv preprint arXiv:2205.02900*, 2022.

637 Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng,
 638 Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. Mimic-iv, a freely accessible
 639 electronic health record dataset. *Scientific data*, 10(1):1, 2023.

640 Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ring-
 641 shia, and Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal
 642 memes. *Advances in neural information processing systems*, 33:2611–2624, 2020.

648 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
 649 *arXiv:1412.6980*, 2014.
 650

651 Philip A Knight. The sinkhorn–knopp algorithm: convergence and applications. *SIAM Journal on*
 652 *Matrix Analysis and Applications*, 30(1):261–275, 2008.

653 Artemy Kolchinsky. A novel approach to the partial information decomposition. *Entropy*, 24(3):
 654 403, 2022.

655 Gueorgi Kossinets. Effects of missing data in social networks. *Social networks*, 28(3):247–268,
 656 2006.

657 Lien P Le, Thu Nguyen, Michael A Riegler, Pal Halvorsen, and Binh T Nguyen. Multimodal missing
 658 data in healthcare: A comprehensive review and future directions. *Computer Science Review*, 56:
 659 100720, 2025.

660 Kwanhyung Lee, Soojeong Lee, Sangchul Hahn, Heejung Hyun, Edward Choi, Byungeun Ahn, and
 661 Joohyung Lee. Learning missing modal electronic health records with unified multi-modal data
 662 embedding and modality-aware attention. *arXiv preprint arXiv:2305.02504*, 2023.

663 Mingcheng Li, Dingkang Yang, Xiao Zhao, Shuaibing Wang, Yan Wang, Kun Yang, Mingyang Sun,
 664 Dongliang Kou, Ziyun Qian, and Lihua Zhang. Correlation-decoupled knowledge distillation for
 665 multimodal sentiment analysis with incomplete modalities. In *Proceedings of the IEEE/CVF*
 666 *Conference on Computer Vision and Pattern Recognition*, pp. 12458–12468, 2024.

667 Pin Li, Jeremy MG Taylor, Daniel E Spratt, R Jeffery Karnes, and Matthew J Schipper. Evaluation
 668 of predictive model performance of an existing model in the presence of missing data. *Statistics*
 669 *in medicine*, 40(15):3477–3498, 2021.

670 Zheng Lian, Lan Chen, Licai Sun, Bin Liu, and Jianhua Tao. Gcnet: Graph completion network
 671 for incomplete multimodal learning in conversation. *IEEE Transactions on pattern analysis and*
 672 *machine intelligence*, 45:8419–8432, 2023.

673 Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Chen, Peter Wu,
 674 Michelle A Lee, Yuke Zhu, et al. Multibench: Multiscale benchmarks for multimodal represen-
 675 tation learning. *Advances in neural information processing systems*, 2021(DB1):1, 2021.

676 Paul Pu Liang, Yun Cheng, Ruslan Salakhutdinov, and Louis-Philippe Morency. Multimodal fu-
 677 sion interactions: A study of human and automatic quantification. In *Proceedings of the 25th*
 678 *International Conference on Multimodal Interaction*, pp. 425–435, 2023.

679 Paul Pu Liang, Yun Cheng, Xiang Fan, Chun Kai Ling, Suzanne Nie, Richard Chen, Zihao Deng,
 680 Nicholas Allen, Randy Auerbach, Faisal Mahmood, et al. Quantifying & modeling multimodal
 681 interactions: An information decomposition framework. *Advances in Neural Information Pro-
 682 cessing Systems*, 36, 2024a.

683 Paul Pu Liang, Akshay Goindani, Talha Chafekar, Leena Mathur, Haofei Yu, Ruslan Salakhutdi-
 684 nov, and Louis-Philippe Morency. Hemm: Holistic evaluation of multimodal foundation models.
 685 *Advances in Neural Information Processing Systems*, 37:42899–42940, 2024b.

686 Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations & trends in multimodal
 687 machine learning: Principles, challenges, and open questions. *ACM Computing Surveys*, 56(10):
 688 1–42, 2024c.

689 Ronghao Lin and Haifeng Hu. Missmodal: Increasing robustness to missing modality in multimodal
 690 sentiment analysis. *Transactions of the Association for Computational Linguistics*, 11:1686–
 691 1702, 2023.

692 Hong Liu, Dong Wei, Donghuan Lu, Jinghan Sun, Liansheng Wang, and Yefeng Zheng. M3ae:
 693 Multimodal representation learning for brain tumor segmentation with missing modalities. In
 694 *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 1657–1665, 2023a.

702 Jiashuo Liu, Tianyu Wang, Peng Cui, and Hongseok Namkoong. On the need for a language de-
 703 scribing distribution shifts: Illustrations on tabular datasets. *Advances in Neural Information*
 704 *Processing Systems*, 36:51371–51408, 2023b.

705 Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. Smil: Multimodal
 706 learning with severely missing modality. In *Proceedings of the AAAI Conference on Artificial*
 707 *Intelligence*, volume 35, pp. 2302–2310, 2021.

708 Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. Are multimodal transform-
 709 ers robust to missing modality? In *Proceedings of the IEEE/CVF conference on computer vision*
 710 *and pattern recognition*, pp. 18177–18186, 2022.

711 Yingbo Ma, Mehmet Celepkolu, Kristy Elizabeth Boyer, Collin F Lynch, Eric Wiebe, and Maya
 712 Israel. How noisy is too noisy? the impact of data noise on multimodal recognition of confusion
 713 and conflict during collaborative learning. In *Proceedings of the 25th International Conference*
 714 *on Multimodal Interaction*, pp. 326–335, 2023.

715 Daniele Malitesta, Emanuele Rossi, Claudio Pomo, Tommaso Di Noia, and Fragkiskos D Malliaros.
 716 Do we really need to drop items with missing modalities in multimodal recommendation? In *Pro-
 717 ceedings of the 33rd ACM International Conference on Information and Knowledge Management*,
 718 pp. 3943–3948, 2024.

719 William McGill. Multivariate information transmission. *Transactions of the IRE Professional Group*
 720 *on Information Theory*, 4(4):93–111, 1954.

721 Kaden McKeen, Laura Oliva, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. Ecg-fm:
 722 An open electrocardiogram foundation model. *arXiv preprint arXiv:2408.05178*, 2024.

723 Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef
 724 Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated
 725 video clips. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.
 726 2630–2640, 2019.

727 Karthika Mohan and Judea Pearl. Graphical models for processing missing data. *Journal of the*
 728 *American Statistical Association*, 116(534):1023–1037, 2021.

729 Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, Andrew Y Ng, et al. Multi-
 730 modal deep learning. In *ICML*, volume 11, pp. 689–696, 2011.

731 Changdae Oh, Zhen Fang, Shawn Im, Xuefeng Du, and Yixuan Li. Understanding multimodal llms
 732 under distribution shifts: An information-theoretic approach. *arXiv preprint arXiv:2502.00577*,
 733 2025.

734 Alexander Papolos, Jagat Narula, Chirag Bavishi, Farooq A Chaudhry, and Partho P Sengupta. Us
 735 hospital use of echocardiography: insights from the nationwide inpatient sample. *Journal of the*
 736 *American College of Cardiology*, 67(5):502–511, 2016.

737 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 738 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-
 739 performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

740 Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word
 741 representation. In *Proceedings of the 2014 conference on empirical methods in natural language*
 742 *processing (EMNLP)*, pp. 1532–1543, 2014.

743 Sam Perochon, J Matias Di Martino, Kimberly LH Carpenter, Scott Compton, Naomi Davis, Brian
 744 Eichner, Steven Espinosa, Lauren Franz, Pradeep Raj Krishnappa Babu, Guillermo Sapiro, et al.
 745 Early detection of autism using digital behavioral phenotyping. *Nature Medicine*, 29(10):2489–
 746 2497, 2023.

747 Matthew Phelan, Nrupen A Bhavsar, and Benjamin A Goldstein. Illustrating informed presence
 748 bias in electronic health records data: how patient interactions with a health system can impact
 749 inference. *EGEMS*, 5(1), 2017.

756 Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer,
 757 Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equaliza-
 758 tion and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.

759

760 M. Reyna, N. Sadr, A. Gu, E. A. Perez Alday, C. Liu, S. Seyed, A. Shah, and G. Clifford. Will
 761 two do? varying dimensions in electrocardiography: The physionet/computing in cardiology
 762 challenge 2021 (version 1.0.3). *PhysioNet*, 2022. doi: 10.13026/34va-7q14.

763

764 Matthew A Reyna, Nadi Sadr, Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux,
 765 Ali Bahrami Rad, Andoni Elola, Salman Seyed, Sardar Ansari, et al. Will two do? varying
 766 dimensions in electrocardiography: the physionet/computing in cardiology challenge 2021. In
 767 *2021 computing in cardiology (CinC)*, volume 48, pp. 1–4. IEEE, 2021.

768

769 James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when
 770 some regressors are not always observed. *Journal of the American statistical Association*, 89
 (427):846–866, 1994.

771

772 Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

773

774 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi
 775 Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An
 776 open large-scale dataset for training next generation image-text models. *Advances in neural in-
 777 formation processing systems*, 35:25278–25294, 2022.

778

779 Dhruv Shah, Błażej Osiński, Sergey Levine, et al. Lm-nav: Robotic navigation with large pre-
 780 trained models of language, vision, and action. In *Conference on robot learning*, pp. 492–504.
 781 PMLR, 2023.

782

783 Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned,
 784 hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th
 785 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
 786 2556–2565, 2018.

787

788 Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matri-
 789 ces. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.

790

791 Elizabeth A Stuart. Matching methods for causal inference: A review and a look forward. *Statistical
 792 science: a review journal of the Institute of Mathematical Statistics*, 25(1), 2010.

793

794 Han Te Sun. Multiple mutual informations and multiple interactions in frequency data. *Inf. Control*,
 795 46:26–45, 1980.

796

797 Luan Tran, Xiaoming Liu, Jiayu Zhou, and Rong Jin. Missing modalities imputation via cascaded
 798 residual autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern
 799 recognition*, pp. 1405–1414, 2017.

800

801 Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang,
 802 Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, et al. Towards generalist biomedical ai.
 803 *NEJM AI*, 1(3):A1oa2300138, 2024.

804

805 Alvaro E Ulloa-Cerna, Linyuan Jing, John M Pfeifer, Sushravya Raghunath, Jeffrey A Ruhl,
 806 Daniel B Rocha, Joseph B Leader, Noah Zimmerman, Greg Lee, Steven R Steinhubl, et al. Re-
 807 chommend: an ecg-based machine learning approach for identifying patients at increased risk of
 808 undiagnosed structural heart disease detectable by echocardiography. *Circulation*, 146(1):36–47,
 809 2022.

810

811 Aayush Atul Verma, Amir Saeidi, Shamanthak Hegde, Ajay Therala, Fenil Denish Bardoliya, Na-
 812 garaju Machavarapu, Shri Ajay Kumar Ravindhiran, Srijan Malyala, Agneet Chatterjee, Yezhou
 813 Yang, et al. Evaluating multimodal large language models across distribution shifts and augmenta-
 814 tions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 815 pp. 5314–5324, 2024.

810 Qi Wang, Liang Zhan, Paul Thompson, and Jiayu Zhou. Multimodal learning with incomplete
 811 modalities by knowledge distillation. In *Proceedings of the 26th ACM SIGKDD International*
 812 *Conference on Knowledge Discovery & Data Mining*, pp. 1828–1838, 2020a.

813

814 Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks
 815 hard? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 816 pp. 12695–12705, 2020b.

817

818 Paul L Williams and Randall D Beer. Nonnegative decomposition of multivariate information. *arXiv*
 819 *preprint arXiv:1004.2515*, 2010.

820

821 Sangmin Woo, Sumin Lee, Yeonju Park, Muhammad Adi Nugroho, and Changick Kim. Towards
 822 good practices for missing modality robust action recognition. In *Proceedings of the AAAI Con-*
 823 *ference on Artificial Intelligence*, volume 37, pp. 2776–2784, 2023.

824

825 Catherine M Writing Committee Members, Otto, Rick A Nishimura, Robert O Bonow, Blase A
 826 Carabello, John P Erwin III, Federico Gentile, Hani Jneid, Eric V Krieger, Michael Mack, et al.
 827 2020 acc/aha guideline for the management of patients with valvular heart disease: a report of the
 828 american college of cardiology/american heart association joint committee on clinical practice
 829 guidelines. *Journal of the American College of Cardiology*, 77(4):e25–e197, 2021.

830

831 Renjie Wu, Hu Wang, Hsiang-Ting Chen, and Gustavo Carneiro. Deep multimodal learning with
 832 missing modality: A survey. *arXiv preprint arXiv:2409.07825*, 2024.

833

834 Aaron D Wyner. A definition of conditional mutual information for arbitrary ensembles. *Information*
 835 and *Control*, 38(1):51–59, 1978.

836

837 Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung
 838 Weng, Atilla Kiraly, Sahar Kazemzadeh, Zakkai Melamed, et al. Elixir: Towards a general purpose
 839 x-ray artificial intelligence system through alignment of large language models and radiology
 840 vision encoders. *arXiv preprint arXiv:2308.01317*, 2023.

841

842 Yihao Xue, Siddharth Joshi, Dang Nguyen, and Baharan Mirzasoleiman. Understanding the robust-
 843 ness of multi-modal contrastive learning to distribution shift. In *The Twelfth International Confer-*
 844 *ence on Learning Representations*, 2024. URL <https://openreview.net/forum?id=rt14XnJYBh>.

845

846 Jiahong Yuan, Mark Liberman, et al. Speaker identification on the scotus corpus. *Journal of the*
 847 *Acoustical Society of America*, 123(5):3878, 2008.

848

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

851

852 Jiandian Zeng, Jiantao Zhou, and Tianyi Liu. Mitigating inconsistencies in multimodal sentiment
 853 analysis under uncertain missing modalities. In *Proceedings of the 2022 conference on empirical*
 854 *methods in natural language processing*, pp. 2924–2934, 2022.

855

856 Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Yi Ma. Investigating
 857 the catastrophic forgetting in multimodal large language model fine-tuning. In *Conference on*
 858 *Parsimony and Learning*, pp. 202–227. PMLR, 2024.

859

860 Haoran Zhang, Harvineet Singh, Marzyeh Ghassemi, and Shalmali Joshi. "Why did the model
 861 fail?": Attributing model performance changes to distribution shifts. In Andreas Krause, Emma
 862 Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Pro-*
 863 *ceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings*
 864 *of Machine Learning Research*, pp. 41550–41578. PMLR, 23–29 Jul 2023.

865

866 Peng-Fei Zhang, Yang Li, Zi Huang, and Hongzhi Yin. Privacy protection in deep multi-modal
 867 retrieval. In *Proceedings of the 44th International ACM SIGIR Conference on Research and*
 868 *Development in Information Retrieval*, pp. 634–643, 2021.

864 Yuhui Zhang, Alyssa Unell, Xiaohan Wang, Dhruba Ghosh, Yuchang Su, Ludwig Schmidt, and
865 Serena Yeung-Levy. Why are visually-grounded language models bad at image classification?
866 In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL
867 <https://openreview.net/forum?id=MwmmBg1VYg>.

868

869 Guanglin Zhou, Zhongyi Han, Shiming Chen, Biwei Huang, Liming Zhu, Salman Khan, Xin Gao,
870 and Lina Yao. Adapting large multimodal models to distribution shifts: The role of in-context
871 learning. *arXiv preprint arXiv:2405.12217*, 2024.

872

873 Helen Zhou, Sivaraman Balakrishnan, and Zachary Lipton. Domain adaptation under missingness
874 shift. In *International Conference on Artificial Intelligence and Statistics*, pp. 9577–9606. PMLR,
875 2023.

876

877 Tongxue Zhou, Pierre Vera, Stéphane Canu, and Su Ruan. Missing data imputation via conditional
878 generator and correlation learning for multimodal brain tumor segmentation. *Pattern Recognition
Letters*, 158:125–132, 2022.

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918 A PROOFS
919920 This section provides the proofs for Lemma 1 and Lemma 2.
921922 **Lemma 1.** *The separable loss function computed on the observed data $l_{\Omega_{obs}}(x_1, x_2, y)$ can be
923 reweighted to approximate the target loss $l_{\Omega}(x_1, x_2, y)$ as follows:*

924
$$l_{\Omega}(x_1, x_2, y) = \frac{1}{1 - p_{\Omega}(m_1, m_2, m_y | C)} l_{\Omega_{obs}}(x_1, x_2, y)$$

925

926 where $p_{\Omega}(m_1, m_2, m_y | C)$ is the probability of missingness, given the covariates C .
927928 *Proof.* The proof is analogous to that of Lemma 2, which we show in detail, for any separable loss
929 function $l(x_1, x_2, y)$. \square
930931 **Lemma 2.**
932

933
$$I_{\Omega}(Y : (X_1, X_2)) = \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y | x_1, x_2)}} \left[\frac{1 - p(m_1, m_2)}{1 - p(m_1, m_2 | x_1, x_2, y)} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right]$$

934

935 *Proof.* Let $m = (m_1, m_2)$,
936

937
$$\begin{aligned} I_{\Omega}(Y : (X_1, X_2)) \\ 938 &= \mathbb{E}_{\Omega} \left[\log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right] \\ 939 &= \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y | x_1, x_2)}} \left[\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega_{obs}}(x_1, x_2, y)} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right] \\ 940 &= \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y | x_1, x_2)}} \left[\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2, y | m = 0)} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right] \\ 941 &= \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y | x_1, x_2)}} \left[\frac{p_{\Omega}(x_1, x_2, y)}{\frac{p(m=0 | x_1, x_2, y)p_{\Omega}(x_1, x_2, y)}{p(m=0)}} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right] \\ 942 &= \mathbb{E}_{\substack{x_1, x_2 \sim p_{\Omega_{obs}}(x_1, x_2) \\ y \sim p_{\Omega}(y | x_1, x_2)}} \left[\frac{1 - p(m = 1)}{1 - p(m = 1 | x_1, x_2, y)} \log \left(\frac{p_{\Omega}(x_1, x_2, y)}{p_{\Omega}(x_1, x_2)p_{\Omega}(y)} \right) \right] \end{aligned}$$

943

944 That is, to estimate the mutual information under the true data distribution, we adjust for the shift in
945 $p_{\Omega_{obs}}(x_1, x_2) \mapsto p_{\Omega}(x_1, x_2)$ and sample y from the IPW-adjusted (parametrized approximations) of
946 $p_{\Omega}(y | x_1, x_2)$.
947948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

972 B PARTIAL INFORMATION DECOMPOSITION (PID)
973

974 Partial information decomposition (PID [Williams & Beer \(2010\)](#)) consists in decomposing the total
975 mutual information ([McGill, 1954](#); [Te Sun, 1980](#)) between a target variable and two input vari-
976 ables into information about the target variable that both input variables share (“Shared” infor-
977 mation), only one input variable has (“Unique” information) and emerges from the interactions of both
978 (“Complementary” information). [Bertschinger et al. \(2014\)](#) introduces bounds for these, reiterated
979 below. In this Appendix, we express these bounds as entropy. First, Table 5 summarizes the nota-
980 tions used in the literature and those used in our work.

981
982 **Table 5.** Quantities and associated variables. Note that the four information measures are approxima-
983 tions.

984 Quantity	Bertschinger	ICYM ² I
985 Input Variable 1	Y	X_1
986 Input Variable 2	Z	X_2
987 Target Variable	X	Y
988 Redundant / Shared Information	$\widetilde{SI}(X : Y : Z)$	$\widetilde{SI}(Y : X_1; X_2)$
989 Unique Information (Input Variable 1)	$\widetilde{UI}(X : Y \setminus Z)$	$\widetilde{UI}(Y : X_1 \setminus X_2)$
990 Unique Information (Input Variable 2)	$\widetilde{UI}(X : Z \setminus Y)$	$\widetilde{UI}(Y : X_2 \setminus X_1)$
991 Synergistic / Complementary Information	$\widetilde{CI}(X : Y; Z)$	$\widetilde{CI}(Y : X_1; X_2)$

992 PID decomposition of the three-way mutual information $I(Y : (X_1, X_2))$ results in the quantities
993 of interest as follows:

$$994 I(Y : (X_1, X_2)) = \underbrace{SI(Y : X_1; X_2)}_{\text{Shared}} + \underbrace{UI(Y : X_1 \setminus X_2)}_{\text{Unique 1}} + \underbrace{UI(Y : X_2 \setminus X_1)}_{\text{Unique 2}} + \underbrace{CI(Y : X_1; X_2)}_{\text{Complementary}}$$

995 Let Δ be the space of all distributions over (X_1, X_2, Y) and let Ω denote the true data distribution
996 (without missingness) and define $\Delta_\Omega := \{q \in \Delta : q(X_i = x_i, Y = y) = p_\Omega(X_i = x_i, Y = y) \forall x_i \in \mathcal{X}_i, y \in \mathcal{Y}, i \in \{1, 2\}\}$. That is, Δ_Ω is the set of all distributions over (X_1, X_2, Y) such
997 that the two-way joints between X_i and Y match the true data-generating distribution. Equipped
998 with this set, [Bertschinger et al. \(2014\)](#) provides the following bounds \widetilde{SI} , \widetilde{UI} , and \widetilde{CI} on the
999 analogous quantities:

$$\begin{aligned} 1000 \widetilde{SI}(Y : X_1; X_2) &= \max_{q \in \Delta_\Omega} CoI_q(Y; X_1; X_2) \\ 1001 &= \max_{q \in \Delta_\Omega} [I_q(Y : X_1) - I_q(Y : X_1 | X_2)] \\ 1002 &= \max_{q \in \Delta_\Omega} [I_q(Y : (X_1, X_2)) - I_q(Y : X_2 | X_1) - I_q(Y : X_1 | X_2)] \\ 1003 &= \max_{q \in \Delta_\Omega} [I_q(Y : (X_1, X_2)) - [I_q(Y : X_2 | X_1) + I_q(Y : X_1 | X_2)]] \\ 1004 \widetilde{UI}(Y : X_1 \setminus X_2) &= \min_{q \in \Delta_\Omega} I_q(Y : X_1 | X_2) \\ 1005 &= \min_{q \in \Delta_\Omega} [I_q(Y : (X_1, X_2)) - I_q(Y : X_2)] \\ 1006 \widetilde{UI}(Y : X_2 \setminus X_1) &= \min_{q \in \Delta_\Omega} I_q(Y : X_2 | X_1) \\ 1007 &= \min_{q \in \Delta_\Omega} [I_q(Y : (X_1, X_2)) - I_q(Y : X_1)] \\ 1008 \widetilde{CI}(Y : X_1; X_2) &= I_\Omega(Y : (X_1, X_2)) - \min_{q \in \Delta_\Omega} I_q(Y : (X_1, X_2)) \end{aligned}$$

1009 In this context, [Bertschinger et al. \(2014\)](#) demonstrates that solving the optimization for $q \in \Delta_\Omega$
1010 that satisfies one of the four conditions above is sufficient to obtain all the quantities of interest.

Importantly, these bounds are tight if there exists a $q_0 \in \Delta_\Omega$ such that $\widetilde{CI}_{q_0}(Y : X_1; X_2) = 0$. Bertschinger et al. (2014) further shows that under common (but unverifiable) assumptions on the data-generating process, the inequalities are tight for all $q \in \Delta_\Omega$. This results in a compelling argument for relying on these entities, as it suggests that it is not possible to decide whether complementary information exists when only the marginals (Y, X_1) and (Y, X_2) are known.

Formulating PID quantities in terms of entropy. For stability, we propose to formalize the previous bound in terms of entropy, $H(\cdot)$, defined for general distributions of X and Y as follows:

$$\begin{aligned} H(X) &:= - \sum_{x \in \mathcal{X}} p(x) \log p(x) \\ H(Y, X) &:= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(y, x) \log (p(y, x)) \\ H(Y|X) &:= - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(y, x) \log \left(\frac{p(y, x)}{p(x)} \right) \\ &= H(Y, X) - H(X) \end{aligned}$$

Using these notations, the mutual information $I(\cdot)$ can be defined as:

$$\begin{aligned} I(Y : X) &:= H(X) - H(X|Y) \\ &= H(X) + H(Y) - H(Y, X) \\ I(Y : X_2|X_1) &:= H(Y, X_1) + H(X_1, X_2) - H(Y, X_1, X_2) - H(X_1) \end{aligned}$$

The previous quantities of interest can then be derived as:

$$\begin{aligned} I(Y : (X_1, X_2)) &:= I(Y : X_1) + I(Y : X_2|X_1) \\ &= \underbrace{H(Y) + H(X_1) - H(Y, X_1)}_{I(Y:X_1)} + \underbrace{H(Y, X_1) + H(X_1, X_2) - H(Y, X_1, X_2) - H(X_1)}_{I(Y:X_2|X_1)} \\ &= H(Y) + H(X_1, X_2) - H(Y, X_1, X_2) \end{aligned}$$

where the first equation comes from the chain rule of mutual information (Wyner, 1978).

Similarly, we can get the expression for co-information $CoI(Y; X_1; X_2)$:

$$\begin{aligned} CoI(Y; X_1; X_2) &= I(Y : X_1) + I(Y : X_2) - I(Y : (X_1, X_2)) \\ &= \underbrace{[H(Y) - H(Y|X_1)]}_{I(Y:X_1)} + \underbrace{[H(Y) - H(Y|X_2)]}_{I(Y:X_2)} \\ &\quad - \underbrace{[H(Y) + H(X_1, X_2) - H(Y, X_1, X_2)]}_{I(Y:(X_1,X_2))} \\ &= [H(Y) - [H(Y, X_1) - H(X_1)]] + \underbrace{[H(Y) - [H(Y, X_2) - H(X_2)]]}_{I(Y:(X_1,X_2))} \\ &\quad - \underbrace{[H(Y) + H(X_1, X_2) - H(Y, X_1, X_2)]}_{I(Y:(X_1,X_2))} \\ &= H(Y) + H(X_1) + H(X_2) \\ &\quad - [H(X_1, X_2) + H(Y, X_1) + H(Y, X_2)] \\ &\quad + H(Y, X_1, X_2) \end{aligned}$$

1080 The PID bounds can then be expressed in terms of entropy:
 1081

1082

$$\begin{aligned}
 1083 \widetilde{SI}(Y : X_1; X_2) &= \max_{q \in \Delta_\Omega} CoI_q(Y; X_1; X_2) \\
 1084 &= \max_{q \in \Delta_\Omega} [H_q(Y) + H_q(X_1) + H_q(X_2) \\
 1085 &\quad - [H_q(X_1, X_2) + H_q(Y, X_1) + H_q(Y, X_2)] + H_q(Y, X_1, X_2)] \\
 1086 \widetilde{UI}(Y : X_1 \setminus X_2) &= \min_{q \in \Delta_\Omega} I_q(Y : X_1 | X_2) \\
 1087 &= \min_{q \in \Delta_\Omega} [H_q(Y, X_2) + H_q(X_1, X_2) - H_q(Y, X_1, X_2) - H_q(X_2)] \\
 1088 \widetilde{UI}(Y : X_2 \setminus X_1) &= \min_{q \in \Delta_\Omega} I_q(Y : X_2 | X_1) \\
 1089 &= \min_{q \in \Delta_\Omega} [H_q(Y, X_1) + H_q(X_1, X_2) - H_q(Y, X_1, X_2) - H_q(X_1)] \\
 1090 \widetilde{CI}(Y : X_1; X_2) &= I_\Omega(Y : (X_1, X_2)) - \min_{q \in \Delta_\Omega} I_q(Y : (X_1, X_2)) \\
 1091 &= [H_\Omega(Y) + H_\Omega(X_1, X_2) - H_\Omega(Y, X_1, X_2)] \\
 1092 &\quad - \min_{q \in \Delta_\Omega} [H_q(Y) + H_q(X_1, X_2) - H_q(Y, X_1, X_2)]
 \end{aligned}$$

1093

1094 where $H_q(\cdot)$ and $H_\Omega(\cdot)$ is the entropy of a variable under probability distributions q and Ω , respectively.
 1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134 C ICYM²I
11351136 C.1 ICYM²I-LEARN
11371138 Table 6 summarizes the proposed approach to estimate performance using Ω_{obs} . Critically, one must
1139 correct both training and evaluation to obtain the performance on Ω .
11401141 **Table 6.** ICYM²I: Inverse probability weighting-adjusted multimodal training and evaluation under
1142 missingness shift.

		Evaluation distribution	
		Observed Ω_{obs}	Underlying Ω
Training		Current practice	IPW-adjusted evaluation alone
		IPW-adjusted training alone	ICYM ² I (IPW-adjusted training and evaluation)

1156 C.2 ICYM²I-PID
11571158 To obtain the PID decomposition, one must solve one of the bounds introduced in Appendix B. Due
1159 to the equivalence demonstrated by Bertschinger et al. (2014), one can focus on solving:
1160

1161
$$\min_{q \in \Delta_{\Omega}} I_q(Y : (X_1, X_2))$$

1162

1163 To this end, Liang et al. (2024a) proposes to minimize this quantity using the set
1164

1165
$$\Delta_{\Omega} \approx \{q \propto \exp(f_1(x_1) \cdot f_2(x_2)) : q(x_i, y) = \Omega_{\text{obs}\phi}(x_i, y) \forall x_i \in \mathcal{X}_i, y \in \mathcal{Y}, i \in \{1, 2\}\}$$

1166 Our proposed approach uses the following projection set to operationalize PID-bound estimation
1167 while accounting for the distribution shift associated with missing modalities:
1168

1169
$$\begin{aligned} & \Delta_{\Omega}^{\text{ICYM}^2\text{I}} \\ & \approx \{q \propto \exp(f_1(x_1) \cdot f_2(x_2)) : q(X_i = x_i, Y = y) = p_{\Omega\phi}(y, x_i) \forall x_i \in \mathcal{X}_i, y \in \mathcal{Y}, i \in \{1, 2\}\} \\ & = \{q \propto \exp(f_1(x_1) \cdot f_2(x_2)) : q(X_i = x_i, Y = y) = \text{IPW}_{p_{\Omega\phi}}(p_{\Omega_{\text{obs}\phi}}(y, x_i)) \forall x_i \in \mathcal{X}_i, y \in \mathcal{Y}, i \in \{1, 2\}\} \end{aligned}$$

1173

1174 where $\text{IPW}_q(p)$ denote IPW reweighting to correct $p \mapsto q$ using samples from p .
11751176 We summarize our proposed auto-differentiable PID estimation algorithm with IPW-based
1177 correction in Algorithm 1, based on the following steps:
1178

- **Estimate IPW correction to adjust for the distribution shift.** To this end, we recommend training a flexible, calibrated model that controls for all covariates explaining the missingness process.
- **Train corrected unimodal and multimodal models.** Training must account for the distribution shift by weighting the loss using Lemma 1. These models are used to estimate Ω .
- **PID optimization.** One must estimate $q \in \Delta_{\Omega}^{\text{ICYM}^2\text{I}}$ that minimizes $I_q(Y : (X_1, X_2))$. To ensure $q \in \Delta_{\Omega}^{\text{ICYM}^2\text{I}}$, we use a modified SINKHORN-KNOPP algorithm that matches q 's marginals to the estimated probabilities of the corrected unimodal neural networks (obtained in the previous step).

1188
 1189 • **Estimate mutual information.** Equipped with q , one can estimate the PID decomposition
 1190 using the adjustment presented in Appendix A.

Algorithm 1 ICYM²I-PID

1191 **Require:** $X_1, X_2, Y \sim p_{\Omega_{\text{obs}}}$
 1192 1: **# Step 1: Adjust** $\Omega_{\text{obs}} \mapsto \Omega$.
 1193 2: Estimate missingness mechanisms $p_{\Omega_{\phi}}(M_1, M_2, M_Y \mid C)$ for IPW.
 1194 3: **# Step 2: Train corrected unimodal and multimodal models.**
 1195 4: Training each model with weighting IPW-loss: $f(y \mid x_i) \approx p_{\Omega}(y \mid x_i)$, $\forall i \in \{1, 2\}$, and
 1196 $f(y \mid x_1, x_2) \approx p_{\Omega}(y \mid x_1, x_2)$.
 1197 5: **# Step 3: PID optimization.**
 1198 6: Initialize parameterizations θ for q : $f_i(y \mid x_i)$, $\forall i \in \{1, 2\}$.
 1199 7: $q_{\theta}(y \mid x_1, x_2) \leftarrow \exp(f_1(y \mid x_1)f_2(y \mid x_2)^T)$
 1200 8: **while** not converged **do**
 1201 9: **for** samples in batch **do**
 1202 10: **# Ensure** $q \in \Delta_{\Omega}^{\text{ICYM}^2\text{I}}$ by projection.
 1203 11: $q_{\theta}(y \mid x_1, x_2) \leftarrow \text{SINKHORN-KNOPP}(q_{\theta}(y \mid x_1, x_2), \{p_{\Omega}(y, x_i)\}_{i=1}^2)$.
 1204 12: Estimate the loss $I_q(Y : (X_1, X_2))$ as a batch sample mean.
 1205 13: $\theta \leftarrow \theta - \nabla_{\theta} I_q(Y : (X_1, X_2))$.
 1206 14: **end for**
 1207 15: **end while**
 1208 **# Step 4: Estimate mutual information under** p_{Ω} .
 1209 16: Estimate $I_{\Omega}(Y : (X_1, X_2))$, and $I_{\Omega}(Y : X_i)$, $\forall i \in \{1, 2\}$ using adjustment in Appendix A.
 1210 17: $PID(\Omega) \leftarrow (\widetilde{CI}(Y : X_1; X_2), \widetilde{SI}(Y : X_1; X_2), \widetilde{UI}(Y : X_1 \setminus X_2), \widetilde{UI}(Y : X_2 \setminus X_1))$
 1211 18: **return** $PID(\Omega)$

1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

1242 The traditional SINKHORN-KNOPP algorithm updates a matrix to enforce its marginals to be unit
 1243 vectors. In our work, we adapt the algorithm to enforce the marginals to match p_Ω -marginals, en-
 1244 suring that $q_\theta(\cdot) \in \Delta_\Omega$. To ensure proper gradient propagation and reduce memory use, we use the
 1245 unrolled SINKHORN-KNOPP (Sinkhorn & Knopp, 1967; Cuturi, 2013) algorithm. In the follow-
 1246 ing, we use subscripts q_{x_1, x_2} to denote $q_\theta(y, x_1, x_2)$ and p_{x_i} to denote $p_\phi(y, x_i)$. The algorithm is
 1247 detailed below:

1248

1249 **Algorithm 2** Unrolled SINKHORN-KNOPP update1250 **Require:** $q_{x_1, x_2}, p_{x_1}, p_{x_2}$, tolerance atol 1251 1: $q_{x_1} \leftarrow \sum_{x_2} q_{x_1, x_2}$ 1252 2: $q_{x_2} \leftarrow \sum_{x_1} q_{x_1, x_2}$ 1253 3: **while** do1254 4: **# Avoid update if both exit conditions have been met.**1255 5: **if** $\left| \frac{q_{x_1} - p_{x_1}}{p_{x_1}} \right| \leq \text{atol}$ **and** $\left| \frac{q_{x_2} - p_{x_2}}{p_{x_2}} \right| \leq \text{atol}$ **then**1256 6: **return** q_{x_1, x_2} 1257 7: **end if**1258 8: **# Update marginal.**1259 9: $q_{x_1, x_2} \leftarrow \frac{q_{x_1, x_2}}{q_{x_2}} \cdot p_{x_2}$ 1260 10: $q_{x_1} \leftarrow \sum_{x_2} q_{x_1, x_2}$ 1261 11: **# If the other marginal still matches, done.**1262 12: **if** $\left| \frac{q_{x_1} - p_{x_1}}{p_{x_1}} \right| \leq \text{atol}$ **then**1263 13: **return** q_{x_1, x_2} 1264 14: **end if**1265 15: **# Repeat for the other marginal.**1266 16: $q_{x_1, x_2} \leftarrow \frac{q_{x_1, x_2}}{q_{x_1}} \cdot p_{x_1}$ 1267 17: $q_{x_2} \leftarrow \sum_{x_1} q_{x_1, x_2}$ 1268 18: **if** $\left| \frac{q_{x_2} - p_{x_2}}{p_{x_2}} \right| \leq \text{atol}$ **then**1269 19: **return** q_{x_1, x_2} 20: **end if**21: **end while**

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296 **D BIT-WISE LOGITS**
 1297

1298 In this section, we perform a sensitivity analysis of the logit setting presented in Section 4.1 under
 1299 two additional missingness patterns: MCAR (Missing Completely at Random) and MNAR (Missing
 1300 Not at Random). In this setting, we fit a logistic regression to estimate the probability of missingness
 1301 on the observed modality, which is then used to estimate the IPW. The performance estimates and
 1302 PID for these two missingness processes are illustrated in Tables 7 and 8.

1303 Since MCAR does not result in a distribution shift, one expects the same performance estimates
 1304 for both the full and observed populations. Furthermore, in this setting, the IPW correction corre-
 1305 sponds to a constant value, as any point has the same probability of observing both modalities. This
 1306 correction also results in no change in performance estimates.

1307 On the contrary, MNAR patterns do not guarantee similar behavior. Particularly, this missingness
 1308 process may result in a distribution shift that cannot be assessed or accounted for without assump-
 1309 tions about the data distribution, as one does not observe the covariates that impact the missingness
 1310 process. The results demonstrate that both the observed and corrected strategies result in biased
 1311 estimates.

1313 **Table 7.** Impact of missingness on multimodality information for bitwise logic operators under MCAR.
 1314 Parentheses denote standard deviation across batches.

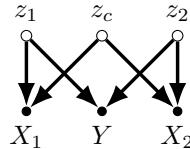
		AUROC			Information Decomposition			
		X_1	X_2	$X_1 + X_2$	Unique 1	Unique 2	Shared	Complementary
AND	Oracle	0.83 (0.01)	0.84 (0.01)	1.00 (0.00)	0.05 (0.00)	0.03 (0.00)	0.26 (0.00)	0.47 (0.00)
	Observed	0.83 (0.01)	0.83 (0.01)	1.00 (0.00)	0.05 (0.00)	0.03 (0.00)	0.23 (0.00)	0.52 (0.00)
	ICYM ² I	0.83 (0.01)	0.85 (0.01)	1.00 (0.00)	0.03 (0.00)	0.06 (0.00)	0.27 (0.00)	0.44 (0.00)
OR	Oracle	0.84 (0.01)	0.83 (0.01)	1.00 (0.00)	0.04 (0.00)	0.05 (0.00)	0.27 (0.00)	0.46 (0.00)
	Observed	0.84 (0.01)	0.84 (0.01)	1.00 (0.00)	0.06 (0.00)	0.03 (0.00)	0.25 (0.00)	0.51 (0.00)
	ICYM ² I	0.85 (0.01)	0.83 (0.01)	1.00 (0.00)	0.06 (0.00)	0.02 (0.00)	0.25 (0.00)	0.51 (0.00)
XOR	Oracle	0.51 (0.02)	0.49 (0.01)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.99 (0.00)
	Observed	0.51 (0.02)	0.50 (0.02)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.95 (0.00)
	ICYM ² I	0.51 (0.02)	0.51 (0.02)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.95 (0.00)

1327 **Table 8.** Impact of missingness on multimodality information for bitwise logic operators under MNAR.
 1328 Parentheses denote standard deviation across batches.

		AUROC			Information Decomposition			
		X_1	X_2	$X_1 + X_2$	Unique 1	Unique 2	Shared	Complementary
AND	Oracle	0.83 (0.01)	0.84 (0.01)	1.00 (0.00)	0.05 (0.00)	0.03 (0.00)	0.26 (0.00)	0.47 (0.00)
	Observed	0.93 (0.01)	0.67 (0.01)	1.00 (0.00)	0.45 (0.00)	0.00 (0.00)	0.17 (0.00)	0.33 (0.00)
	ICYM ² I	0.93 (0.01)	0.67 (0.01)	1.00 (0.00)	0.45 (0.00)	0.00 (0.00)	0.17 (0.00)	0.33 (0.00)
OR	Oracle	0.84 (0.01)	0.83 (0.01)	1.00 (0.00)	0.04 (0.00)	0.05 (0.00)	0.27 (0.00)	0.46 (0.00)
	Observed	0.78 (0.01)	0.95 (0.01)	1.00 (0.00)	0.00 (0.00)	0.17 (0.00)	0.11 (0.00)	0.23 (0.00)
	ICYM ² I	0.78 (0.01)	0.95 (0.01)	1.00 (0.00)	0.00 (0.00)	0.17 (0.00)	0.11 (0.00)	0.23 (0.00)
XOR	Oracle	0.51 (0.02)	0.49 (0.01)	1.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.99 (0.00)
	Observed	0.80 (0.02)	0.52 (0.02)	1.00 (0.00)	0.35 (0.00)	0.07 (0.00)	0.00 (0.00)	0.61 (0.00)
	ICYM ² I	0.80 (0.02)	0.52 (0.02)	1.00 (0.00)	0.35 (0.00)	0.07 (0.00)	0.00 (0.00)	0.61 (0.00)

1350 E SYNTHETIC DATA RESULTS
1351

1352 **Data generation.** Our work builds on the example introduced in (Liang et al., 2024a), in which
1353 we enforce additional missingness. Three latent variables (z_1 , z_2 , and z_c) are drawn from multi-
1354 dimensional clustered data; the observed covariates are a concatenation of z_c and one of the other
1355 latent variables, as illustrated in Figure 3. Then, the outcome Y is generated as $Y = \sigma(p_1\mathbb{E}(z_1) +$
1356 $p_2\mathbb{E}(z_2) + (1 - p_1 - p_2)\mathbb{E}(z_c))$, with the proportion $p_i \in [0, 1]$ such that $p_1 + p_2 \leq 1$. We simulate
1357 datasets with varying values of p_1 and p_2 . Then, we enforce a 50% MAR missingness pattern in
1358 X_2 by modeling the probability of missingness. We do this by clustering X_1 into 100 groups using
1359 Kmeans. Then, the probability of missingness is generated using a random forest that regresses X_1
1360 to predict $c_j \cdot Y$. **Empirical setting.** Data were split into three: 80% for training, 10% for validation,
1361

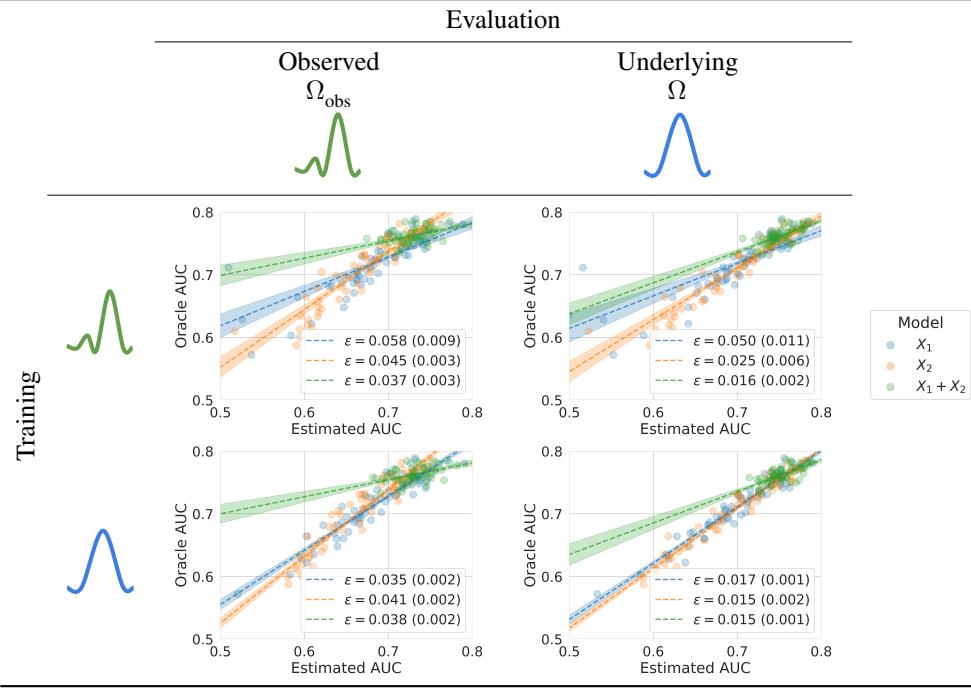


1362 **Figure 3.** Data generating processes for synthetic experiments. z_i denote latent vectors, while all other
1363 variables are observed. Filled point nodes are observed variables, while unfilled nodes are unobserved.
1364

1365 and the rest for testing. We consider neural networks with 2 hidden layers with 32 nodes, trained
1366 using an Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.001 over 100 epochs. Our
1367 evaluation relies on discriminative performance measured through AUROC.
1368

1369 **Estimating predictive performance under Ω_{obs} .** Table 9 presents the estimated performance ob-
1370 tained under different corrections. These results underline the importance of correcting both training
1371 and evaluation, as proposed in ICYM²I, to best align with the performance one would obtain on
1372 Ω , as shown by the smallest Root Mean Squared Error (RMSE) observed when both corrections are
1373 applied. Note that in this setting, we rely on the true IPW correction that one would obtain with a
1374 properly specified model, as the MAR setting is met.
1375

1376 **Table 9.** Comparison between estimated AUC performance under the different training and evaluation
1377 corrections and oracle performance on Ω . ϵ denotes the RMSE between estimated and oracle PIDs.
1378



1404 **F SEMI-SYNTHETIC DATA RESULTS**
 1405

1406 **F.1 UR-FUNNY**
 1407

1408 We illustrate the impact of missingness on estimating the informativeness of different modalities on
 1409 real-world data with UR-FUNNY (Hasan et al., 2019), a multimodal dataset for humor detection
 1410 from human speech used in affective computing. The dataset comprises text, audio, and visual
 1411 modalities from 10 - 20 second videos sourced from TED talks, and the task is to detect whether
 1412 a punchline would trigger a laugh. Labels were generated using the markup “(Laughter)” (Chen &
 1413 Lee, 2017) from the transcript.

1414 **Dataset.** The processed dataset from MultiBench (Liang et al., 2021) is a modality-complete dataset
 1415 with 10,166 samples of paired audio, text, and vision embeddings. Audio embeddings were generated
 1416 with COVAREP (Degottex et al., 2014), text with Glove (Pennington et al., 2014), and visual
 1417 features through the Facet (Yuan et al., 2008) library and OpenFace (Baltrušaitis et al., 2016), and
 1418 aligned using the Penn Phonetics Lab Forced Aligner (P2FA) (Yuan et al., 2008).

1419 **F.1.1 MAR**
 1420

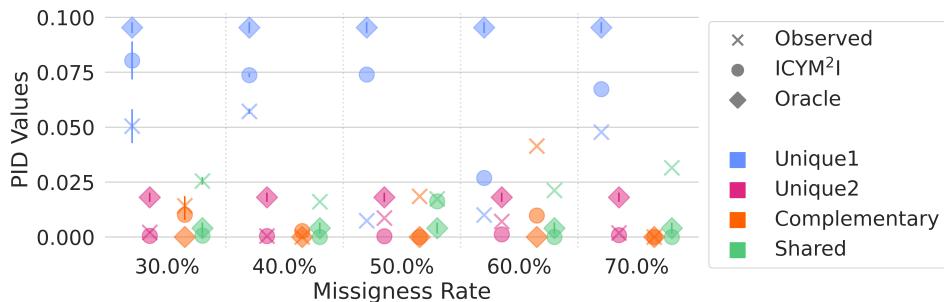
1421 **Enforcing missingness.** To explore the impact of missingness on informativeness, we simulate a
 1422 MAR missingness pattern on the audio and visual features given the textual modality. We vary the
 1423 missingness from 30% to 70%, using the same mechanism as described for synthetic data. This
 1424 semi-synthetic setting enables the evaluation of the proposed correction as the missingness mech-
 1425 anism is known. Note that the original dataset does not contain missing values, as the source data
 1426 (TED Talks) have transcripts, and data labeling was generated based on these transcripts. How-
 1427 ever, settings with systematic transcripts are rare and may reflect a shift from the audio and textual
 1428 modalities observed online for which such a match may not exist.

1429 **Results.** Following the same empirical setting as in the synthetic experiment for each missingness
 1430 rate, we measure the impact of missingness on PID decomposition. Figure 4 displays the PID values
 1431 obtained under three strategies:

1432

- 1433 • Observed: All quantities are estimated using Ω_{obs} .
- 1434 • ICYM²I: All quantities are estimated using Ω_{obs} but corrected for the distribution shift through
 1435 IPW.
- 1436 • Oracle: All quantities are estimated on Ω .

1437 This figure shows that the proposed strategy is consistently closer to the Oracle’s PID values. This
 1438 demonstrates that under Assumption A, the proposed correction yields better estimates of each
 1439 modality’s informativeness – specifically, the audio-visual modality (Unique 1) carries more in-
 1440 formation.



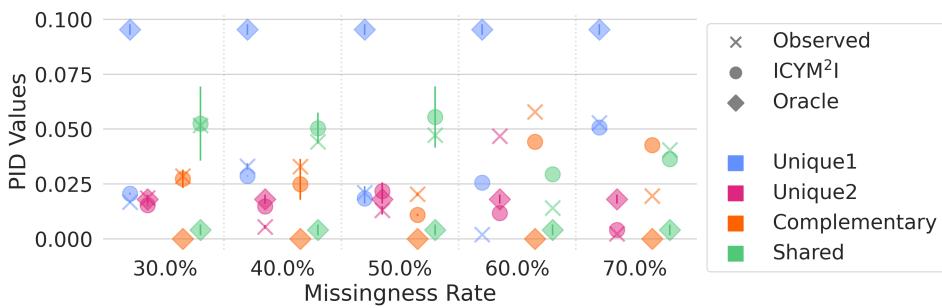
1452 **Figure 4:** Comparison between estimated PID values under increasing missingness in UR-FUNNY.
 1453

1454 **F.1.2 MNAR**
 1455

1456 **Enforcing missingness.** A central assumption of our method is that missingness is MAR. We pro-
 1457 pose to analyze the impact of violations of this assumption, specifically the presence of MNAR

1458 patterns, on the quality of estimates obtained using our correction. To this end, we simulate audio
 1459 and visual missingness as a function of the modality itself. Similarly to the previous analysis, we
 1460 vary missingness from 30% to 70%. To estimate propensity in this setting, we rely on a logistic
 1461 regression model based on the fully observed modality.

1462 **Results.** Figure 5 shows the PID values obtained under the three previously described strategies.
 1463 Critically, the proposed correction leads to performance similar to the model without correction,
 1464 as the missingness probabilities cannot be estimated from the observed modality. This example
 1465 illustrates the importance of assessing the plausibility of Assumption B in real-world settings, as no
 1466 theoretical guarantees hold in such settings.
 1467



1478 **Figure 5:** Comparison between estimated PID values under increasing missingness in UR-FUNNY.
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

1512
1513

F.2 HATEFUL MEMES

1514
1515
1516

We run experiments using the dataset from the Hateful Memes Challenge (Kiela et al., 2020), which investigates text-image multimodal reasoning in the context of hate speech detection in online memes. The dataset comprises text-image pairs with an associated label indicating hate speech.

1517
1518
1519
1520
1521
1522
1523
1524

Dataset. We utilize the [Kaggle version](#) of the Facebook Hateful Memes dataset, as referenced in the [Holistic Evaluation of Multimodal Foundation Models \(HEMM\)](#) (Liang et al., 2024b) repository. Our analysis focuses on the 9,000 samples with associated labels. For each sample, embeddings were extracted for both modalities using a ResNet-50 (He et al., 2016) for images and a BERT-base-uncased (Devlin et al., 2019) model for text. The proposed ResNet-50 was pre-trained on ImageNet (Deng et al., 2009) with the final layer replaced to extract 2048-dimensional feature vectors, and BERT-base-uncased (Devlin et al., 2019) extracts embeddings of dimension 784 from the penultimate layer.

1525
1526
1527
1528
1529
1530

Enforcing missingness. Similarly to the previous experiment, we vary the missingness from 30% to 70% by enforcing the same MAR missingness mechanism on the text modality, given the image modality, as we assume not all memes may contain text. Note that memes in the dataset were created by combining text from collected online memes with images sourced from stock images on Getty Images. Consequently, the dataset did not contain missing modality, but may not match the true distribution of memes one would observe online.

1531
1532
1533
1534
1535

Results. As above, we measure the impact of increasing percentages of missingness on PID estimates. While the missingness mechanism results in a limited distribution shift, and therefore small differences in estimates between the corrected and observed strategies, the difference at 70% missingness shows the superiority of the proposed methodology in recovering the Unique contributions.

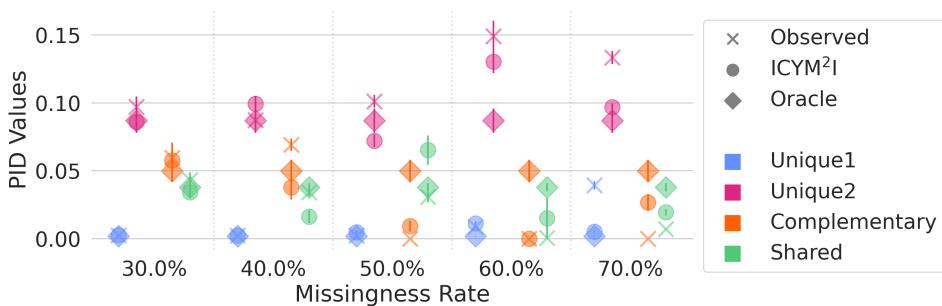
1536
1537
1538
1539
1540
1541
1542
1543
1544
15451546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

Figure 6. Comparison between estimated PID values under increasing missingness in Hateful Memes.

1566
1567

G STRUCTURAL HEART DISEASE DATA PROCESSING

1568
1569

G.1 EMBEDDING GENERATION

1570
1571

We generate embeddings using modality-specific foundation models—ECG embeddings are generated using ECG-FM (McKeen et al., 2024) and CXR embeddings with ELIXR-C (Xu et al., 2023).

1572
1573

All electrocardiograms were 10-second, standard 12-lead ECG signals collected at abstracted to 250 Hz, which we resampled to 500 Hz, and standard normalized by channel to match the inputs for ECG-FM (McKeen et al., 2024). We used the version of ECG-FM with weights pretrained on MIMIC-IV (Johnson et al., 2023; Goldberger et al., 2000) and PhysioNet 2021 (Reyna et al., 2021; 2022). We averaged the output feature embeddings along the temporal dimension and flattened them to produce vectors of length 768.

1579
1580
1581
1582
1583

The chest radiographs used in our study were all postero-anterior (PA) view CXRs. We extracted pixel values from the DICOM files as grayscale images, center-cropped each image along the shorter dimension, applied contrast-limited adaptive histogram equalization (CLAHE) (Pizer et al., 1987) with a clip limit of 0.2, and resized each image to 1284×1284 pixels. All outputted embeddings were flattened to 4098-dimensional vectors.

1584

G.2 IPW CORRECTION

1585

To address missingness in the observed CXRs, we apply the proposed propensity-based correction. The propensity scores are obtained from a logistic regression model using the ECG embedding, along with sex and age as predictors, serving as proxies for the socio-medical factors that influence whether a CXR is collected. Controlling for these covariates aims to render the MAR assumption more plausible. In practice, all relevant covariates, even outside of modalities being modeled, can be used for the correction.

1592

1593

H COMPUTE INFRASTRUCTURE

1594

All experiments were performed on a server with an AMD EPYC 7313 CPU, 256 GB of memory, and two NVIDIA RTX A6000 GPUs, as well as a server with an Intel Xeon E5-2640 CPU, 128 GB of memory, and a NVIDIA GTX Titan X GPU. Our software stack includes Python 3.12, PyTorch 2.2.1 (Paszke et al., 2019), and standard Python scientific libraries. Chest radiograph embeddings used Tensorflow 2.19 (Abadi et al., 2015) and Tensorflow-Text 2.19 based on the requirements for the ELIXR models (Xu et al., 2023). Electrocardiogram embeddings were generated using an environment with Python 3.9 and fairseq-signals 1.0 to match the requirements for fairseq-signals and ECG-FM (McKeen et al., 2024). Generating embeddings for our structural heart disease data took approximately 10 hours on our server with a Titan X GPU. All synthetic experiments require 12 hours of compute time using one NVIDIA RTX A6000 GPU.

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619