

Causal Mode Multiplexer: A Novel Framework for Unbiased Multispectral Pedestrian Detection

Taeheon Kim^{*1}, Sebin Shin^{*1}, Youngjoon Yu¹, Hak Gu Kim², Yong Man Ro^{†1}

¹Integrated Vision and Language Lab, KAIST, South Korea

²Immersive Reality and Intelligent Systems Lab, Chung-Ang University, South Korea
{eetaekim, ssbin0914, greatday, ymro}@kaist.ac.kr, hakgukim@cau.ac.kr

Abstract

RGBT multispectral pedestrian detection has emerged as a promising solution for safety-critical applications that require day/night operations. However, the modality bias problem remains unsolved as multispectral pedestrian detectors learn the statistical bias in datasets. Specifically, datasets in multispectral pedestrian detection mainly distribute between ROTO¹ (day) and RXTO (night) data; the majority of the pedestrian labels statistically co-occur with their thermal features. As a result, multispectral pedestrian detectors show poor generalization ability on examples beyond this statistical correlation, such as ROTX data. To address this problem, we propose a novel Causal Mode Multiplexer (CMM) framework that effectively learns the causalities between multispectral inputs and predictions. Moreover, we construct a new dataset (ROTX-MP) to evaluate modality bias in multispectral pedestrian detection. ROTX-MP mainly includes ROTX examples not presented in previous datasets. Extensive experiments demonstrate that our proposed CMM framework generalizes well on existing datasets (KAIST, CVC-14, FLIR) and the new ROTX-MP. Our code and dataset are available at: <https://github.com/ssbin0914/Causal-Mode-Multiplexer.git>.

1. Introduction

Multispectral pedestrian detection plays a critical role in many real-world applications that require both day/night operations, such as smart surveillance cameras (CCTVs), search and rescue (SAR) autopilots, and autonomous vehicles (AVs) [4, 6, 7, 9, 11, 14–16, 18, 28, 39, 41]. Despite its notable progress, an overlooked factor in multi-

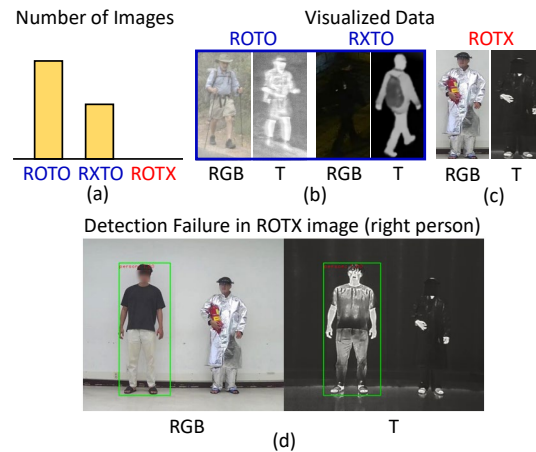


Figure 1. (a) ROTO, RXTO, and ROTX distribution in existing multispectral pedestrian datasets. (b) In the existing datasets, a high statistical correlation exists between pedestrian labels and their thermal features. (c) Example of a ROTX image. (d) Current multispectral pedestrian detection models fail prediction in ROTX image (the right person) due to learning statistical bias in datasets.

spectral pedestrian detection is the modality bias problem that occurs in multimodal models (e.g., Visual Question Answering) [3, 19, 23, 34, 40]. Multimodal models are known to often leverage spurious correlations between a certain modality and answers due to statistical biases in datasets [10, 20, 22, 36, 37]. For example, a Visual Question Answering (VQA) model may blindly answer “tennis” for the question “What sports?” just referring to the most co-occurring textual QA pairs (i.e., linguistic bias) in the train data [23, 29]. Models exploiting statistical bias in datasets often demonstrate poor generalization ability to out-of-distribution data, rarely providing proper multimodal evidence for prediction [8, 35].

As modality biases can significantly impact the reliability of multimodal models, it is essential to investigate statistical bias in multispectral pedestrian datasets [7, 9, 11, 13]. Analyzing their distribution, we find an ever-overlooked statistical bias. Illustrated in Fig.1 (a), most data are dis-

^{*}Equally contributed.

[†]Corresponding author.

¹R★T★ refers to the visibility (O/X) in each modality. Generally, ROTO refers to daytime images, and RXTO refers to nighttime images. ROTX refers to daytime images in obscured situations.

tributed between ROTO (daytime images) or RXTO (night-time images). In these data, pedestrian labels always statistically co-occur with their thermal features, as thermal sensors can generally capture robust pedestrian silhouettes all day/night. On the other hand, there are rarely ROTX images that can be captured in thermal obscured situations in the daytime e.g., firefighters wearing heat-insulation cloth as in Fig.1 (c). From this data distribution, we hypothesize that models may learn the statistical co-occurrence between the pedestrian label and the presence of thermal features. If so, there is a potential risk that multispectral pedestrian detectors may fail in scenes in which such statistical correlation does not hold, e.g., ROTX images in which thermal features are obscured.

In this paper, we test multispectral pedestrian detection models under ROTX data and validate our hypothesis. Since existing datasets [7, 9, 11, 13] do not include ROTX data, we first collect ROTX pedestrian images based on practical scenarios. This includes search and rescue (SAR) on firefighters and surveillance of pedestrians over a window. In this situation, pedestrians are visible in only RGB but obscured in thermal as the intermediate obstacle (i.e., heat-insulating cloth or window) blocks the thermal radiation emitted by the pedestrian from reaching the thermal sensors. Fig.1 (d) demonstrates that a multispectral pedestrian detection model fails prediction on ROTX, though the firefighter is obviously visible in RGB. As seen, multispectral pedestrian detectors make predictions on ROTX just based on the absence of the thermal feature, motivating us to formulate this phenomenon as a modality bias problem.

To address this problem, we propose a novel Causal Mode Multiplexer (CMM) framework that performs unbiased inference from statistically biased multispectral pedestrian training data. We adopt this problem setting instead of explicitly changing the training priors (e.g., adding ROTX in the train set) to verify that multispectral pedestrian detectors can disentangle the learned multimodal knowledge and memorized priors by learning *causality*. Specifically, the CMM framework learns causality based on different cause-and-effects between ROTO, RXTO, and ROTX inputs and predictions. For ROTO data, we guide the model to learn the *total effect* [25] in the common mode learning scheme. Next, for ROTX and RXTO data, we utilize the tools of counterfactual intervention to eliminate the direct effect of thermal by subtracting it from the total effect. To this end, we modify the training objective from maximizing the posterior probability likelihood to maximizing the *total indirect effect* [25] in the differential mode learning scheme. Our design requires combining two different learning schemes; therefore, we propose a Causal Mode Multiplexing (CMM) Loss to optimize the interchange.

Moreover, we evaluate the modality bias in multispectral pedestrian detection with our new dataset: ROTX Multi-

spectral Pedestrian (ROTX-MP) dataset. Different from existing multispectral pedestrian datasets [7, 9, 11] which are constrained to ROTO (day) and RXTO (night) data, ROTX-MP mainly includes ROTX pedestrians comprised of 1000 test image pairs. Our experimental results demonstrate that our CMM framework generalizes well under ROTX-MP even with biased training and also performs robustly on existing datasets [7, 9, 11].

The main contributions of our paper are:

1. We propose a Causal Mode Multiplexer (CMM) framework that learns different causality between ROTO, RXTO, and ROTX inputs and pedestrian labels in multispectral pedestrian detection.
2. We propose a Causal Mode Multiplexing (CMM) Loss to optimize the interchange between learning different causal representations.
3. To evaluate modality bias in multispectral pedestrian detection, we conduct a new dataset: ROTX-MP.
4. Extensive experiments demonstrate that our CMM framework generalizes well under ROTX test data with even biased training data - ROTO, RXTO.

2. Preliminaries

Before introducing our method, we present the fundamental concepts of causal inference [2, 24–27, 31, 32].

2.1. Structural Causal Model (SCM)

Structural Causal Models reflect the cause-and-effect relationships (*links* \mathcal{E}) between the set of variables (*nodes* \mathcal{V}). The cause-and-effects (*cause* \rightarrow *effect*) are represented in an acyclic graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$. For random variables X and Y which the direct effect of X is on Y , the direct link could be formulated as $X \rightarrow Y$. If an indirect effect of X is on Y through the variable M , M is considered a mediator between X and Y ($X \rightarrow M \rightarrow Y$ in Fig.2 (a)). With structural causal models, the examination of causality links among variables can be achieved through variable intervention, which involves modifying the value of particular variables and subsequently observing the outcomes.

2.2. Counterfactual Intervention

Counterfactual intervention can break the direct link and eliminate the effect of particular variables. To do this, the effects on Y are compared from two different treatments to the cause variable X : factual and counterfactual. Take the SCM in Fig.2 (a) as an example i.e., $X \rightarrow M \rightarrow Y$. Suppose that $X = x$ represents the “treatment condition” and $X = x^*$ represents the “no-treatment condition” (lowercase letter indicates the observed value of the random variable). Then we can consider factual and counterfactual scenarios.

Counterfactual Notations. We denote the value that Y would obtain in the factual scenario if X is assigned x and M is assigned m as $Y_{x,m_x} = Y(X = x, M = m_x)$.




Figure 2. Structural Causal Models (SCMs) of (a) factual, (b) counterfactual, and (c) total indirect effect scenarios. The direct effect of $X \rightarrow Y$ can be eliminated due to counterfactual intervention.

Similarly, in the counterfactual scenario, Y would have the value $Y_{x,m_{x^*}} = Y(X = x, M = m_{x^*})$ when $X = x$, $M = m_{x^*} = M(X = x^*)$. X is set to a value x^* (usually zero value or mean features) and can break the direct link between M and its parent node X . We note that in the counterfactual scenario, X can be simultaneously assigned to different values x and x^* . So when the intervention is conducted on M , the variable X retains its original value of x as if x (observation on X) had existed as in Fig.2 (b).

Total Indirect Effect. Now we can estimate the total indirect effect (TIE) by comparing two hypothetical scenarios.

$$TIE = Y_{x,m_x} - Y_{x,m_{x^*}}. \quad (1)$$

Total indirect effect (TIE), as in Fig.2 (c) breaks the direct link of $X \rightarrow Y$. Furthermore, total indirect effect (TIE) can be decomposed into total effect (TE) and natural direct effect (NDE).

$$\begin{aligned} TIE &= TE - NDE \\ &= (Y_{x,m_x} - Y_{x^*,m_{x^*}}) - (Y_{x,m_{x^*}} - Y_{x^*,m_{x^*}}). \end{aligned} \quad (2)$$

When we examine the total effect (TE), we compare two hypothetical scenarios: one where $X = x$ and the other where $X = x^*$. In contrast, the natural direct effect (NDE) represents the effect of X on Y while keeping the mediator M blocked. It measures the change in Y as X transitions from x^* to x , with M assigned to the value of the no-treatment $X = x^*$, thereby nullifying M 's response to the treatment $X = x$. In the subsequent section, we will take a more in-depth look at the interpretations of these effects in the context of multispectral pedestrian detection.

3. Structural Causal Model of Multispectral Pedestrian Detection

Before pruning direct effects and performing counterfactual interventions, we investigate the causal links (causality) in multispectral pedestrian detection as below:

Link $X_R \rightarrow M \leftarrow X_T$ (Feature Fusion). RGB and thermal features (X_R and X_T) are fused to generate fusion features of multimodal knowledge M .

Link $M \rightarrow Y$ (Class Label Prediction with Multimodal




Figure 3. Structural Causal Models (SCMs) of multispectral pedestrian detection on (a) ROTO and (b) RXTO data. (b) The thermal direct effect is considered in RXTO scenarios as models rely heavily on the thermal features for making predictions in the nighttime.

Knowledge). Using the multimodal knowledge M , the head network outputs class prediction Y for each fused feature.

Beyond the above causal links, an additional direct link should be considered in practice. In RXTO data (nighttime), the predictions heavily rely on the thermal modality because RGB sensors degrade at night. Thus, a direct effect of the thermal feature X_T on the prediction Y can be formulated as a direct link below:

Link $X_T \rightarrow Y$ (Thermal Direct Effect). As machines memorize priors, models will learn the skewed preference toward thermal in RXTO data. This direct effect causes detection failures when inferring ROTX data. We formulate this cause-and-effect as the thermal direct link.

From the above discussions, we represent structural causal models separately on (a) ROTO (daytime) data and (b) RXTO data (nighttime) in Fig.3 (a)-(b). Our goal is to prune the direct link $X_T \rightarrow Y$ when training RXTO.

4. Proposed Method

4.1. Causal Mode Multiplexer

Based on the different causal graphs in ROTO and RXTO data, our key idea is to interchangeably learn causality from two different learning schemes. First, the 1) *common mode* scheme is designed to learn the total effect in ROTO data (daytime data). Second, the 2) *differential mode* learning scheme learns the total indirect effect in RXTO data (night-time images) to prune the thermal direct effect.

1. Common Mode - ROTO Train/Test Graph. In the common mode, we learn the total effect. Based on the causal graph in Fig.3 (a), we intentionally add the links $X_R \rightarrow Y$ and $X_T \rightarrow Y$ to estimate the uni-modal direct effect of X_R on Y and X_T on Y . We add these direct links for two reasons. 1) To estimate the input types as we will describe in Section 4.2. 2) To provide an estimate of the direct effect of $X_T \rightarrow Y$ in the differential mode. To implement this direct link, we train a uni-modal neural model denoted as $H_{\theta_{X_R}}(\cdot)$ and $H_{\theta_{X_T}}(\cdot)$. We denote the prediction scores obtained from direct links as $Y_{x_R}^{CM}$ and $Y_{x_T}^{CM}$ which can be

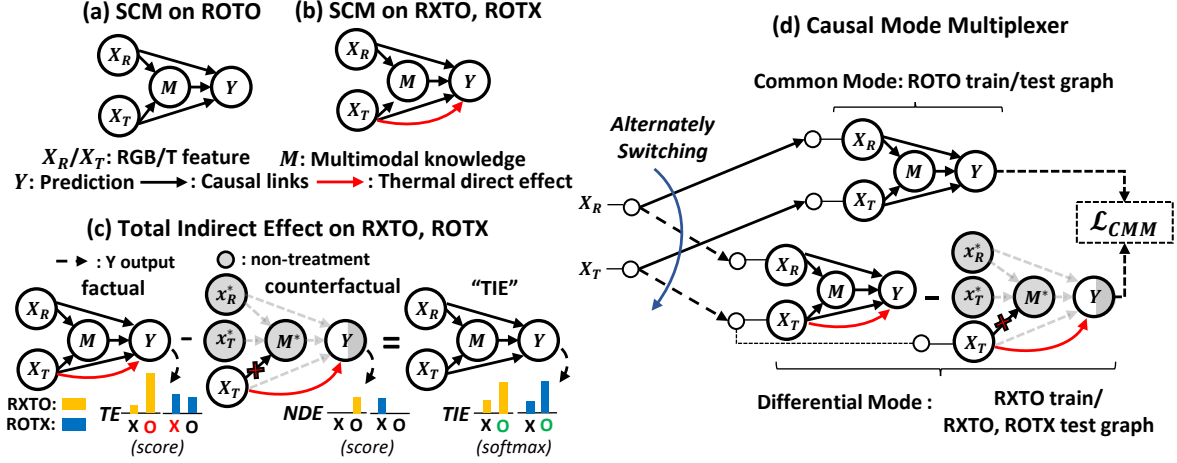


Figure 4. Our Structural Causal Model (SCM) formulations on (a) ROTO (day) and (b) RXTO (night) data. (a) We add direct links ($X_R \rightarrow Y$, $X_T \rightarrow Y$) to the conventional causal graph to determine modes and measure the thermal direct link. (b) The thermal direct link (red) hinders the model from learning causality. (c) implies the total indirect effect on RXTO and ROTX for which the thermal direct link is pruned. (d) We propose a Causal Mode Multiplexer framework that learns causality both on ROTO and RXTO data and thus generalizes well on all ROTO, RXTO, and ROTX.

expressed as below:

$$Y_m^{CM} = H_{\theta_M}(m), Y_{x_R}^{CM} = H_{\theta_{X_R}}(x_R), Y_{x_T}^{CM} = H_{\theta_{X_T}}(x_T), \quad (3)$$

where the neural network $H_{\theta_M}(\cdot)$ refers to the head network, which takes fusion features $M = m$ and outputs the class scores Y_m^{CM} (direct link $M \rightarrow Y$). The structural causal model of *common mode* can be illustrated as in Fig.4 (a).

In order to calculate prediction scores out of Y_m^{CM} , $Y_{x_R}^{CM}$, and $Y_{x_T}^{CM}$, we design a fusion function $\mathcal{F}(\cdot)$ using the nonlinear Log-Harmonic (LH):

$$Y_{m,x_R,x_T}^{CM} = Y_{LH}^{CM} = \log(\sigma(Y_m^{CM}) \times \sigma(Y_{x_R}^{CM}) \times \sigma(Y_{x_T}^{CM})), \quad (4)$$

where Y_{m,x_R,x_T}^{CM} denotes the final prediction score, and σ denotes the sigmoid function.

Total Effect (TE). We measure the total effect by comparing Y_{m,x_R,x_T}^{CM} with its no-treatment condition. Following the definition from the preliminaries, the total effect of the common mode causal graph can be written as:

$$TE = Y_{m,x_R,x_T}^{CM} - Y_{m^*,x_R^*,x_T^*}^{CM} \quad (5)$$

where $Y_{m^*,x_R^*,x_T^*}^{CM}$ is the no-treatment condition. The implementation details of the no-treatment condition are described in the *supplementary material*.

2. Differential Mode - RXTO Train/RXTO, ROTX Test Graph. In the *differential mode*, we learn the total indirect effect. Specifically, we intentionally assign node values and prune the direct link $X_T \rightarrow Y$. To this end, we introduce the natural direct effect.

Natural Direct Effect (NDE). We estimate the natural

direct effect (NDE), the “counterfactual scenario”, which refers explicitly to the thermal direct effect i.e., $X_T \rightarrow Y$ in Fig.4 (b). This direct effect can be estimated by blocking the effects of X_R and M , using the no-treatment condition definitions in the preliminary section. X_T is set to x_T , and M would attain the value m^* when X_T had been x_T^* and X_R is x_R^* . The natural direct effect (NDE) can be obtained by comparing the counterfactual graph to the no-treatment conditions:

$$NDE = Y_{m^*,x_R^*,x_T}^d - Y_{m^*,x_R^*,x_T^*}^d, \quad (6)$$

where m^* , x_R^* and x_T^* represent the no-treatment condition.

Total Indirect Effect (TIE). The elimination of the thermal direct effect can be achieved by subtracting the natural direct effect (NDE) from the total effect (TE).

$$TIE = TE - NDE = Y_{m,x_R,x_T}^{CM} - Y_{m^*,x_R^*,x_T}^d. \quad (7)$$

Fig.4 (c) shows an illustration of the total indirect effect of the differential mode. Note that the no-treatment conditions $Y_{m^*,x_R^*,x_T^*}^d$ and $Y_{m^*,x_R^*,x_T^*}^{CM}$ are the same. For inference, we opt for the label with the highest TIE, in contrast to the conventional strategies that primarily rely on posterior probability i.e., $p(y|x_R, x_T)$.

4.2. Causal Mode Multiplexing (CMM) Loss

Determination of Causal Modes. To assign different learning schemes based on the input type, we leverage the uni-modal prediction scores $Y_{x_R}^{CM}$ and $Y_{x_T}^{CM}$ based on the following rationale: For ROTO inputs, both the $Y_{x_R}^{CM}$

and $Y_{x_T}^{CM}$ value differences between pedestrians and backgrounds will have the same sign. In this case, we determine the input as ROTO and assign the common mode learning scheme. In the other case, we determine the input as RXTO or ROTX, assign the differential mode learning scheme, and prune the thermal direct link. Using these properties, we design a binary translation function to calculate the mode number.

We denote the mode number as K_{mode} , for which we want the value 1 for the common mode and -1 for the differential mode. The details of our function design are as follows: Let $\pi_R = [\pi_R^b, \pi_R^f]$ and $\pi_T = [\pi_T^b, \pi_T^f]$ the approximate one-hot representation of the prediction labels derived from the $argmax$ of $Y_{x_R}^{CM}$ and $Y_{x_T}^{CM}$. Here, index 0 refers to the prediction label of the background (no pedestrian), and index 1 refers to the prediction label of the pedestrian. Since the $argmax$ operation is non-differentiable on the gradient descent method, we adopt the Gumbel-softmax [12] estimation. Then we can write π_R and π_T as:

$$\pi_R = \text{softmax} [g_R + \log(Y_{x_R}^{CM})/\tau], \quad (8)$$

$$\pi_T = \text{softmax} [g_T + \log(Y_{x_T}^{CM})/\tau], \quad (9)$$

where we set the Gumbel noises g_R and g_T to zero since we do not need random sampling variations for our purpose. From the formula of Eq.(8) and Eq.(9), we can obtain values $\{\pi_R^b = 0, \pi_R^f = 1, \pi_T^b = 0, \pi_T^f = 1\}$ for ROTO, $\{\pi_R^b = 1, \pi_R^f = 0, \pi_T^b = 0, \pi_T^f = 1\}$ for RXTO, and $\{\pi_R^b = 0, \pi_R^f = 1, \pi_T^b = 1, \pi_T^f = 0\}$ for ROTX. From these values, we can determine the causal mode number K_{mode} according to the above rationale.

$$K_{mode} = \Delta\pi_R \times \Delta\pi_T = (\pi_R^f - \pi_R^b) \times (\pi_T^f - \pi_T^b). \quad (10)$$

This K_{mode} value will obtain 1 for the common mode and -1 for the differential mode. Then we can design a “switchable total indirect effect” (sTIE) in which the cause-effect is differently calculated between TE and TIE according to the causal mode number K_{mode} .

$$\begin{aligned} sTIE &= TE - ReLU(-K_{mode}) \times NDE \\ &= Y_{m,x_R,x_T}^{CM} - ReLU(-K_{mode}) \times Y_{m^*,x_R^*,x_T}^d \\ &= \begin{cases} Y_{m,x_R,x_T}^{CM} & \text{if } K_{mode} = 1 \\ Y_{m,x_R,x_T}^{CM} - Y_{m^*,x_R^*,x_T}^d & \text{if } K_{mode} = -1, \end{cases} \end{aligned} \quad (11)$$

in which sTIE is calculated as the total effect in Eq.(5) for ROTO inputs and as the total indirect effect in Eq.(7) for RXTO and ROTX inputs. We summarize the relationship between cause-effect (CE) and K_{mode} assignment with respect to data type in Table 1.

Causal Mode Multiplexing Loss. We formulate the sTIE

Table 1. Cause-effect (CE) and K_{mode} assignment with respect to data type.

Type	$\Delta\pi_R$	$\Delta\pi_T$	K_{mode}	CE
ROTO	1	1	1	TE
RXTO	-1	1	-1	TIE
ROTX	1	-1	-1	TIE

as a loss function. Before introducing our Causal Mode Multiplexing Loss, we first revisit the loss function of the conventional model. Given a triplet (x_T, x_R, y) where y is the ground-truth class label of RGB/T ROI feature pair: x_R/x_T , the ROI classification branches of the conventional multispectral pedestrian detection model [14] are optimized by:

$$\mathcal{L}_{cls} = \mathcal{L}_Y(TE, y) \quad (\text{Conventional}), \quad (12)$$

where \mathcal{L}_Y denotes the cross-entropy loss. This conventional loss guides the model to learn the total effect on both ROTO and RXTO, which provokes the thermal direct effect. Different from them, Causal Mode Multiplexing (CMM) Loss learns the causality in both ROTO and RXTO training data based on two causal modes:

$$\mathcal{L}_{CMM} = \mathcal{L}_Y(sTIE, y), \quad (13)$$

which sTIE refers to the formula in Eq.(11). The overall classification branch loss can be written as:

$$\mathcal{L}_{cls} = \mathcal{L}_{CMM} + \mathcal{L}_Y(Y_{x_R}^{CM}, y) + \mathcal{L}_Y(Y_{x_T}^{CM}, y), \quad (14)$$

where $\mathcal{L}_Y(Y_{x_R}^{CM}, y)$ and $\mathcal{L}_Y(Y_{x_T}^{CM}, y)$ are over $Y_{x_R}^{CM}$ and $Y_{x_T}^{CM}$.

4.3. Implementation

Training. The final training loss is the combination of \mathcal{L}_{cls} , bounding box regression loss \mathcal{L}_{bbox} , and \mathcal{L}_{model} , which include the RPN and the uncertainty module. We follow the implementation details of the paper [14].

$$\mathcal{L}_{total} = \sum_{(x_R, x_T, y) \in D} \mathcal{L}_{cls} + \mathcal{L}_{bbox} + \mathcal{L}_{model}. \quad (15)$$

Inference. We use the switchable total indirect effect (sTIE) for inference.

$$sTIE = Y_{m,x_R,x_T}^{CM} - ReLU(-K_{mode}) \times Y_{m^*,x_R^*,x_T}^d. \quad (16)$$

5. New Dataset: ROTX-MP

To evaluate modality bias in multispectral pedestrian detectors, we propose to collect a new dataset: The ROTX Multispectral Pedestrian (ROTX-MP) dataset. ROTX-MP consists of 1000 ROTX test images collected from two practical scenarios related to the applications of multispectral pedestrian detection. The details of ROTX-MP are described below.

5.1. Comparison to Existing Datasets

Our ROTX-MP dataset mainly contains ROTX data, compared to existing datasets that consist of ROTO and RXTO data. We compare the data distributions of datasets using the statistics criteria as the following.

5.1.1 Statistics Criteria

For the distribution comparison, we count the number of images for each data type: ROTO, RXTO, and ROTX. We explain our criteria for counting ROTO, RXTO, and ROTX (or ROTO/ROTX) data in datasets.

ROTO: We count the number of daytime images for each dataset for ROTO.

RXTO: We count the nighttime images for each dataset for RXTO.

ROTX: ROTX data refer to daytime images where there is an intermediate medium that obscures the thermal radiation of the pedestrian. In this case, pedestrians are only imaged by RGB sensors. We count images that contain these scenarios. There can be images where ROTO and ROTX scenarios both occur. We note them separately, as ROTO/ROTX.

5.1.2 Dataset Distribution

From the above criteria, we obtain the distribution of each ROTO, RXTO, and ROTX (or ROTO/ROTX) in different datasets: 1) KAIST [11], 2) CVC-14 [9], 3) FLIR [7], and 4) our ROTX-MP. We count the images of all train and test data for KAIST, CVC-14, and FLIR, as well as the images of the test data of ROTX-MP. ROTX-MP contains only test data. The distribution is plotted in Fig.5 and Fig.6. It can be seen from Fig.5 that KAIST, CVC-14, and FLIR mainly contain ROTO and RXTO data and no ROTX scenes (including ROTO/ROTX). On the other hand, the ROTX-MP dataset (1000 images) primarily consists of ROTX data (737 images) and ROTO/ROTX (263 images) data. Moreover, we compare the overall size of the dataset. Since ROTX-MP contains only test data, we compare it with the test sets of KAIST, CVC-14, and FLIR. ROTX-MP contains 1000 test data, comparable to KAIST test-2252 images (ROTO: 1455, RXTO: 797), CVC-14 test-1417 images (ROTO: 690, RXTO: 727), and FLIR test-1013 images (ROTO: 702, RXTO: 311).

5.2. Data Collection Process

ROTX-MP is a dataset of RGBT multispectral image pairs captured from an RGBT camera. The RGBT camera we used for data collection was the FLIR Duo Pro R (FPA 640×512) camera manufactured by FLIR Systems, Inc. This product supports simultaneous RGB and thermal ($NETD < 50mK$, $\lambda = 7.5 \sim 13.5\mu m$) imaging in a




Figure 5. ROTO, RXTO, and ROTX distribution of popular multi-spectral pedestrian datasets: (a) KAIST [11], (b) CVC-14 [9], and (c) FLIR [7]. Images from all train/test sets are counted.




Figure 6. ROTO/ROTX and ROTX distribution of ROTX-MP test set.

dual pip mode. We invited 8 volunteers to participate in our data collection. All of the volunteers agreed to the data collection and its future release to the public. On streets and indoors, we mounted the camera device on a portable tripod and captured video of volunteers from different shooting distances ranging from 0 to 15 meters. Each video involved 2–8 volunteers. From the recorded videos, we meticulously hand-picked high-quality image pairs. RGBT pairs were spatially aligned and synchronized, and finalized of 1000 RGBT pairs. The faces of volunteers are blurred for privacy protection. Specifically, we collected ROTX data based on the below practical scenarios.

1. Pedestrians Over a Glass Window (737 ROTX images). Multispectral pedestrian detection is attractive for all-day/night smart surveillance. However, pedestrians observed through a window are visible in RGB but obscured in thermal because thermal radiation cannot penetrate through glass. Failure in these scenes constrains multispectral pedestrian detection within the window.

2. Pedestrians Wearing Heat-insulation Clothes (263 ROTO+ROTX images). Heat-insulation clothing such as fire protection gear or low-emissivity clothing provides a way of thermal invisibility or *stealth*. Wearing fire protection gear can make firefighters or evacuees undetectable from autopilots equipped with multispectral cameras in search and rescue (SAR) situations. Also, criminals (e.g., bank robbers) can wear low-emissivity clothing to evade multispectral pedestrian detection on security cameras. We collected 150 images for firefighter scenes and 113 images for criminal scenes. In both scenarios, we simultaneously contain ROTO scenes for comparison to the obscured pedestrians. Thus, we note the ROTO+ROTX data type for this scenario.




Figure 7. **Visualized examples of multispectral pedestrian detection on ROTO/ROTX data.** (a) ROTO and ROTX scenarios. (b) Conventional models [14] fail to detect pedestrians in ROTX scenes. (c) Our CMM framework estimates the total effect for ROTO and the total indirect effect for ROTX. As a result, it produces correct debiased results for ROTX samples, even from biased training data. We provide more examples in the supplementary material.

6. Experimental Setup

6.1. Dataset

We conduct experiments on multiple multispectral pedestrian datasets: KAIST [11], CVC-14 [9], FLIR [7], and our collected dataset: ROTX-MP. KAIST contains 95,328 RGBT data with 640×512 image resolutions. We use the annotation labels provided by the original authors [11] and organize the dataset by 7601 training images and 2252 test images following the baselines [14, 41] we compare. The 2252 (‘All’) test images are composed of 1455 (‘Day’) images and 797 (‘Night’) images. CVC-14 [9] contains 3618 training data and 1417 test data with grayscale RGBT images of 640×471 resolution. 1417 test data (‘All’) include 690 daytime images (‘Day’) and 727 nighttime images (‘Night’). Also, we evaluate under the Teledyne FLIR Free ADAS Thermal Dataset v2.0.0 [7] (FLIR for short). Following previous works [17, 28, 38], we use 4129 train image pairs and 1013 test image pairs organized by Zhang et al. [38] for fair comparison. As evaluations in previous work do not compare day and night performance separately, we note (‘All’) for the FLIR test-set. ROTX-MP is composed of 1000 test images. Similar to FLIR, we note (‘All’) for ROTX-MP.

6.2. Baseline Model

We implement CMM based on the Uncertainty-Guided model [14]. Our method is compared with four competitive multimodal pedestrian detection architectures recently proposed: Kim et al. [14], Cross-modality Fusion Transformer (CFT) [28], MBNet [41], and Halfway Fusion [21]+Faster RCNN [30] version (HFF). More details are described in the supplementary material.

7. Experimental Results

7.1. Evaluation on Existing Datasets

We first report the experimental results on the KAIST [11], CVC-14 [9], and FLIR [7] data. These datasets are

Table 2. Detection performance of different multispectral pedestrian detection models on the KAIST, CVC-14, and FLIR datasets.

Dataset	KAIST			CVC-14			FLIR
	MR(\downarrow)			MR(\downarrow)			AP(\uparrow)
Model	Day	Night	All	Day	Night	All	All
HFF	14.24	10.26	13.04	50.48	36.05	44.31	75.85
CFT	13.99	7.05	11.60	18.81	25.25	21.83	84.10
MBNet	9.64	8.26	9.04	24.70	13.50	21.10	-
Kim et al.	10.11	5.05	8.67	23.87	11.08	18.70	84.67
CMM (Ours)	9.60	5.93	8.54	27.81	7.71	17.13	87.80

comprised of day (ROTO) and night (RXTO) images. We train and test on these datasets to evaluate whether CMM over-corrects the modality bias. Our motivation for designing the CMM framework is to learn causality based on different input types (ROTO, and RXTO) in the training set; therefore, CMM’s intentional behavior is to generalize on existing data as well.

For the evaluation metrics, we follow previous multispectral pedestrian detection works for measuring performance. Specifically, we calculate the Log Average Miss-Rate (MR \downarrow) for the ‘All’, ‘Day’, and ‘Night’ test images of the KAIST and CVC-14 test sets. From our definition, ‘Day’ can be interpreted as ROTO data, ‘Night’ as RXTO data, and ‘All’ as their combination. To evaluate performance on the FLIR dataset, we use the Average Precision (AP \uparrow) to conduct a fair comparison with other methods [28]. A low Miss Rate and a high Average Precision value indicate high detection performance.

Results are reported in Table 2. The CMM framework demonstrates the lowest miss rates for ‘Day’ and ‘All’ in the KAIST set, ‘Night’ and ‘All’ in the CVC-14 set, and the highest AP in the FLIR set. Since ‘All’ represents the full dataset, it can be concluded that CMM achieves the best performance across all KAIST, CVC-14, and FLIR datasets. It is worth noting that CMM performs better than its baseline model, [14], with 0.13 MR, 1.57 MR, and 3.13 AP improvements in KAIST, CVC-14, and FLIR, respectively. Such results meet the original intention of our CMM design, and we extend the evaluation on ROTX-MP.

Table 3. **Ablation study on debiasing strategy.** “A+B” denotes the debiasing strategy that train the model with “A” cause-effect and test with “B” cause-effect. The baseline, TE+TIE, TIE+TIE, and sTIE (ours) are compared.

Train	KAIST				Train	CVC-14				Train	FLIR	
Test	ROTX-MP (AP \uparrow)	KAIST (MR \downarrow)			Test	ROTX-MP (AP \uparrow)	CVC-14 (MR \downarrow)			Test	ROTX-MP (AP \uparrow)	FLIR (AP \uparrow)
Model	All	Day	Night	All	Model	All	Day	Night	All	Model	All	All
Baseline	21.69	10.11	5.05	8.67	Baseline	13.36	23.87	11.08	18.70	Baseline	12.23	84.67
TE+TIE	57.05	26.89	27.66	27.27	TE+TIE	27.97	32.85	12.11	22.33	TE+TIE	19.33	64.27
TIE+TIE	56.45	12.53	9.24	12.10	TIE+TIE	27.75	33.53	9.27	21.63	TIE+TIE	12.02	79.39
sTIE (Ours)	70.44	9.60	5.93	8.54	sTIE (Ours)	34.96	27.81	7.71	17.13	sTIE (Ours)	57.09	87.80

7.2. Evaluation on ROTX-MP

The CMM framework is introduced to enhance the generalizability of models when there is a substantial distribution difference in the training and test splits. Therefore, we train models on conventional datasets [7, 9, 11] and test on ROTX-MP to evaluate the generalizability on ROTX data of each model. Here, we do not directly compare with data augmentation methods [1, 5], which explicitly generate training samples, e.g., counterfactual generation. Generalizing on ROTX data by modifying training priors violates the original purpose of the CMM model, i.e., to evaluate whether multispectral pedestrian detection models infer from memorized priors in training data.

For each evaluation on the ROTX-MP test set (1000 test images), we calculate the Average Precision (AP \uparrow) since it is a standard metric to analyze the accuracy of a detection model. Evaluations are reported in Table 4. Other models fail on ROTX-MP test data when trained on conventional datasets. Models trained on the KAIST [11] dataset: HFF [21], CFT [28], MBNET [41], Kim et al. [14] achieve 36.95, 3.64, 18.88, 21.69 AP on ROTX-MP. Compared to these methods, CMM shows superior generalizability on ROTX-MP test data, achieving 70.44 AP, outperforming other baselines with at least 33.49 AP. Similar results are obtained from models trained on CVC-14 and FLIR when tested on ROTX-MP. Fig. 7 shows qualitative results.

8. Ablation Study

Furthermore, we conduct ablation studies to validate the design choice of the switchable total indirect effect (sTIE) in Eq.(11) for our debiasing strategy. Known debiasing strategies [23, 33] are generally based on training the model via total effect (TE) and inferring with total indirect effect (TIE). For the comparison study, we conducted an experiment on (2) TE training + TIE inference of the (1) baseline model [14]. Moreover, we also evaluate the (3) TIE training + TIE inference strategy to provide an extensive comparison. The results are shown in Table 3. Firstly, the (1) baseline model performs well on existing datasets but generalizes poorly on ROTX-MP. Secondly, the (2) TE+TIE model demonstrates enhanced model generalizability for

Table 4. Detection performance on the ROTX-MP test set. Models are trained from the KAIST, CVC-14, and FLIR datasets.

Train	KAIST	CVC-14	FLIR
Test	ROTX-MP	ROTX-MP	ROTX-MP
Metric	AP(\uparrow)	AP(\uparrow)	AP(\uparrow)
Model	All	All	All
HFF	36.95	8.80	13.21
CFT	3.64	8.58	5.28
MBNet	18.88	-	-
Kim et al.	21.69	13.36	12.23
CMM (Ours)	70.44	34.96	57.09

ROTX data compared to the baseline, but degrades on existing data. Moreover, the (3) TIE+TIE model demonstrates a moderate level of performance, yet degrades on existing datasets compared to the baseline. Compared to them, the (4) CMM framework shows superior performance on both ROTX-MP and conventional test sets. From this ablation study, we verify the effectiveness of our design choice of the switchable total indirect effect (sTIE) in Eq.(11).

9. Conclusion

In this paper, we address the modality bias problem in multispectral pedestrian detection with our innovative solution: The Causal Mode Multiplexer (CMM) framework. Using the tools of counterfactual intervention, CMM enables the model to interchangeably learn between two distinct causal graphs depending on the input data type. We propose the Causal Mode Multiplexing (CMM) Loss to optimize the interchange between two causal graphs. Additionally, we introduce the ROTX-MP dataset to evaluate modality bias in multispectral pedestrian detection. Experimental results on KAIST, CVC-14, FLIR, and our ROTX-MP dataset demonstrate that CMM effectively learns multimodal reasoning and performs well on ROTX test data with training ROTO and RXTO data.

10. Acknowledgements

This work was partly supported by two funds: IITP grant funded by the Korea government (MSIT) (No.2022-0-00984) and Center for Applied Research in Artificial Intelligence (CARAI) grant funded by DAPA and ADD (UD230017TD).

References

- [1] Vedika Agarwal, Rakshith Shetty, and Mario Fritz. Towards causal vqa: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9690–9698, 2020. **8**
- [2] Elias Bareinboim and Judea Pearl. Controlling selection bias in causal inference. In *Artificial Intelligence and Statistics*, pages 100–108. PMLR, 2012. **2**
- [3] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32, 2019. **1**
- [4] Jiale Cao, Yanwei Pang, Jin Xie, Fahad Shahbaz Khan, and Ling Shao. From handcrafted to deep features for pedestrian detection: a survey. *IEEE transactions on pattern analysis and machine intelligence*, 2021. **1**
- [5] Chun-Hao Chang, George Alexandru Adam, and Anna Goldenberg. Towards robust classification model by counterfactual and invariant data generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15212–15221, 2021. **8**
- [6] Yukyung Choi, Namil Kim, Soonmin Hwang, Kibaek Park, Jae Shin Yoon, Kyoungwan An, and In So Kweon. Kaist multi-spectral day/night data set for autonomous and assisted driving. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):934–948, 2018. **1**
- [7] Inc FLIR Systems. Free teledyne flir thermal dataset for algorithm training. <https://www.flir.com/oem/adas/adas-dataset-form/>, 2021. Accessed: 2022-08-05. **1, 2, 6, 7, 8**
- [8] Itai Gat, Idan Schwartz, Alexander Schwing, and Tamir Hazan. Removing bias in multi-modal classifiers: Regularization by maximizing functional entropies. *Advances in Neural Information Processing Systems*, 33:3197–3208, 2020. **1**
- [9] Alejandro González, Zhijie Fang, Yainuvis Socarras, Joan Serrat, David Vázquez, Jiaolong Xu, and Antonio M López. Pedestrian detection at day/night time with visible and fir cameras: A comparison. *Sensors*, 16(6):820, 2016. **1, 2, 6, 7, 8**
- [10] Jianqiang Huang, Yu Qin, Jiaxin Qi, Qianru Sun, and Hanwang Zhang. Deconfounded visual grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 998–1006, 2022. **1**
- [11] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1037–1045, 2015. **1, 2, 6, 7, 8**
- [12] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. **5**
- [13] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. Llvip: A visible-infrared paired dataset for low-light vision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3496–3504, 2021. **1, 2**
- [14] Jung Uk Kim, Sungjune Park, and Yong Man Ro. Uncertainty-guided cross-modal learning for robust multi-spectral pedestrian detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(3):1510–1523, 2022. **1, 5, 7, 8**
- [15] Taeheon Kim, Hong Joo Lee, and Yong Man Ro. Map: Multispectral adversarial patch to attack person detection. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4853–4857. IEEE, 2022.
- [16] Taeheon Kim, Youngjoon Yu, and Yong Man Ro. Defending physical adversarial attack on object detection via adversarial patch-feature energy. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1905–1913, 2022. **1**
- [17] Taeheon Kim, Youngjoon Yu, and Yong Man Ro. Multispectral invisible coating: Laminated visible-thermal physical attack against multispectral object detectors using transparent low-e films. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1151–1159, 2023. **7**
- [18] Taeheon Kim, Sangyun Chung, Damin Yeom, Youngjoon Yu, Hak Gu Kim, and Yong Man Ro. Mscotdet: Language-driven multi-modal fusion for improved multispectral pedestrian detection. *arXiv preprint arXiv:2403.15209*, 2024. **1**
- [19] Gouthaman Kv and Anurag Mittal. Reducing language biases in visual question answering with visually-grounded question encoder. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 18–34. Springer, 2020. **1**
- [20] Bing Liu, Dong Wang, Xu Yang, Yong Zhou, Rui Yao, Zhiwen Shao, and Jiaqi Zhao. Show, deconfound and tell: Image captioning with causal inference. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18041–18050, 2022. **1**
- [21] Jingjing Liu, Shaoting Zhang, Shu Wang, and Dimitris N Metaxas. Multispectral deep neural networks for pedestrian detection. *arXiv preprint arXiv:1611.02644*, 2016. **7, 8**
- [22] Guoshun Nan, Rui Qiao, Yao Xiao, Jun Liu, Sicong Leng, Hao Zhang, and Wei Lu. Interventional video grounding with dual contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2765–2775, 2021. **1**
- [23] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12700–12710, 2021. **1, 8**
- [24] Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995. **2**
- [25] Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018. **2**
- [26] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- [27] Judea Pearl et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 19(2):3, 2000. **2**

- [28] Fang Qingyun, Han Dapeng, and Wang Zhaokui. Cross-modality fusion transformer for multispectral object detection, 2022. [1](#), [7](#), [8](#)
- [29] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. Overcoming language priors in visual question answering with adversarial regularization. *Advances in Neural Information Processing Systems*, 31, 2018. [1](#)
- [30] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [7](#)
- [31] James M Robins. Semantics of causal dag models and the identification of direct and indirect effects. *Oxford Statistical Science Series*, pages 70–82, 2003. [2](#)
- [32] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005. [2](#)
- [33] Bing Tian, Yixin Cao, Yong Zhang, and Chunxiao Xing. Debiasing nlu models via causal intervention and counterfactual reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11376–11384, 2022. [8](#)
- [34] Zhiquan Wen, Guanghui Xu, Mingkui Tan, Qingyao Wu, and Qi Wu. Debaised visual question answering from feature and sample perspectives. *Advances in Neural Information Processing Systems*, 34:3784–3796, 2021. [1](#)
- [35] Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. In *International Conference on Machine Learning*, pages 24043–24055. PMLR, 2022. [1](#)
- [36] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9847–9857, 2021. [1](#)
- [37] Chuanqi Zang, Hanqing Wang, Mingtao Pei, and Wei Liang. Discovering the real association: Multimodal causal reasoning in video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19027–19036, 2023. [1](#)
- [38] Heng Zhang, Elisa Fromont, Sébastien Lefevre, and Bruno Avignon. Multispectral fusion for object detection with cyclic fuse-and-refine blocks. In *2020 IEEE International conference on image processing (ICIP)*, pages 276–280. IEEE, 2020. [7](#)
- [39] Lu Zhang, Xiangyu Zhu, Xiangyu Chen, Xu Yang, Zhen Lei, and Zhiyong Liu. Weakly aligned cross-modal learning for multispectral pedestrian detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5127–5137, 2019. [1](#)
- [40] Sixian Zhang, Xinhang Song, Weijie Li, Yubing Bai, Xinyao Yu, and Shuqiang Jiang. Layout-based causal inference for object navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10792–10802, 2023. [1](#)
- [41] Kailai Zhou, Linsen Chen, and Xun Cao. Improving multispectral pedestrian detection by addressing modality imbalance problems. In *European Conference on Computer Vision*, pages 787–803. Springer, 2020. [1](#), [7](#), [8](#)