






Descriptor: *LYNRED Mobility Dataset Multimodal Detection Subset (LYNRED-MDS)*

LOÏC ARBEZ ¹, JESSY MATIAS ², XAVIER BRENIÈRE ²,
JOCELYN CHANUSSOT ¹ (FELLOW, IEEE), AND RONALD PHLYPO ³

¹University of Grenoble Alpes, INRIA, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

²Lynred, 38113 Veurey-Voroize, France

³University of Grenoble Alpes, CNRS, Grenoble INP, GIPSA-Lab, 38000 Grenoble, France

CORRESPONDING AUTHOR: Loïc Arbez (e-mail: loic.arbez@inria.fr).

This work was supported by the Fondation Grenoble INP: DeepRed Research Chair, under the patronage of Lynred.

ABSTRACT Current road safety systems primarily focus on minimizing post-collision damage. However, advances in algorithmic perception are shifting focus toward early collision prediction, especially in low-visibility conditions like nighttime or fog, where thermal infrared sensing outperforms both human vision and RGB imaging. While available RGB-infrared datasets such as FLIR ADAS and LLVIP are good benchmarks, they mostly consist of clear weather and overly simple scenarios. In this article, we introduce the LYNRED-MDS: Multimodal Detection Subset, a subset of the LYNRED Mobility Dataset, comprised of 4000 RGB-infrared image pairs captured under diverse weather, lighting, and road conditions around Grenoble, France. Our dataset spans varied driving contexts (urban, rural, mountainous, etc.) and a vehicle fleet compliant with Western European standards. Thermal cross-dataset evaluation using a YOLOv8n baseline suggests that our dataset offers strong generalization potential for pedestrian detection in driving scenarios. By covering critical edge cases, our dataset supports the development of more reliable and deployable vision systems for advanced driver-assistance systems.

IEEE SOCIETY/COUNCIL Signal Processing Society (SPS)

DATA TYPE/LOCATION Thermal images, Visible images, Multispectral images; Grenoble, Auvergne-Rhône-Alpes, FRANCE

DATA DOI/PID <https://www.lynred.com/lynred-mobility-dataset>

INDEX TERMS Advanced driver assistance systems (ADAS), dataset, infrared, object detection.

BACKGROUND

Recent advancements in technology have triggered a significant shift in regulatory frameworks designed to enhance the safety of vulnerable road users (VRUs). These frameworks increasingly focus on the integration of advanced technologies, particularly in the realm of vehicle safety systems. With technologies such as computer vision-based pedestrian and cyclist detection systems proving effective in reducing VRU-related collisions [1], [2], [3], many jurisdictions now require the inclusion of advanced driver assistance systems (ADAS) in new vehicles. For example, the European Union's General Safety Regulation now requires

features like automated emergency braking tailored for VRU scenarios [3], and organizations like the European New Car Assessment Programme have incorporated these systems into their vehicle safety assessment protocols [1]. In the United States, the National Highway Traffic Safety Administration similarly emphasizes the adoption of ADAS technologies, in particular for automatic emergency braking (AEB) systems [2]. In that regard, they require these AEB systems to avoid impact at 62 mph (about 100 km/h). This means emergency braking systems should be able to detect objects, in this case human beings, at about 50 meters assuming an AEB reaction time close to instantaneous.

In addition to vehicle-based systems, regulatory efforts have extended to urban infrastructure. Many cities have introduced mandates for adaptive traffic controls, smart crosswalks, and reduced speed limits in high-density VRU zones [3]. These efforts build on traditional safety measures, such as crashworthiness standards and pedestrian-friendly design and further reflect a broader shift toward leveraging cutting-edge technologies like computer vision and smart sensors to address VRU safety in a more comprehensive manner. This shift toward predictive safety has been facilitated by advancements in computer vision and sensor technologies, which have both reduced the cost of sensors and improved the robustness of detection systems over the last decades. The automotive industry has, since then, embraced these technologies.

The publication of public datasets sparked further research in the field. Notable among these are the Kitti [4] and Cityscapes [5] datasets, which introduced demanding tasks such as semantic segmentation and depth estimation, mostly leveraging light detection and ranging (LiDAR) data. Alongside those active sensor technologies, several thermal imaging datasets made their way into the pedestrian or VRU detection community, demonstrating the efficacy of passive IR imaging under various lighting conditions. Indeed, long-wave infrared (LWIR) or thermal sensors—unlike visible light or RGB sensors—do not rely on external light sources such as the sun or a car’s headlights, nor on emitted signals like LiDAR or ultrasound sensors. Indeed, infrared imaging uniquely relies on the thermal radiation of objects, i.e. the quantity of photons released by a given body as a function of its temperature [6]. This makes it an ideal candidate for pedestrian detection tasks. As a matter of fact, humans have an internal temperature of 36.6 °C [7] and hence tend to be easily distinguishable from their environment [8], [9].

The collection of visible data being relatively inexpensive, these IR datasets are often paired with visible RGB sensor data. Most modern vehicles are already equipped with a standard optical camera. As a result, RGB-T (visible-thermal) or RGB-IR (visible-infrared) datasets have been released. The first works on those paired visible and thermal images in object detection were through datasets such as KAIST multi-spectral [10] and CVC-14 [8]. The KAIST multispectral dataset has been particularly influential in the development of VRU detection tasks. Following KAIST, the FLIR thermal ADAS dataset [11] helped further establish the utility of thermal-visible pairs for driving scenarios [12]. It succeeded in raising awareness for VRU detection and paved the way for this new paradigm [12]. In addition to the on-vehicle datasets, more recent static datasets like LLVIP [13] have been published. Although they differ from the original focus on RGB-T object detection for on-vehicle scenarios, they provide valuable insights for research in RGB-T object detection. Additionally, some datasets have introduced new tasks for infrared modalities, such as depth estimation alongside object detection [14].

Despite the extensive literature on thermal datasets, most of the existing works focus on favorable detection scenarios, primarily—if not exclusively—using ideal weather conditions for both thermal and visible modalities. Extreme weather conditions and challenging real-world environments are rarely represented. Yet, the systems developed using these datasets are expected to perform well even under such conditions. To address this gap, we present a recently published RGB-T object detection dataset: the multimodal detection subset of the LYNRED Mobility Dataset (LYNRED MD). The latter offers images acquired during more diverse and realistic driving scenarios. A few snapshot pairs are presented in Fig. 1.

Related Datasets

In this section, we present the state-of-the-art datasets, their specifications, and assess their contribution to the literature. A summary can be found in Table I.

KAIST Multi-Spectral

Published in 2015 by the Korean Advanced Institute of Science and Technology, it includes over 95 k aligned color-thermal pairs of images and 103 128 annotations in its first version [10]. The original test split annotations were first criticized for their “problematic bounding boxes” and corrected by [15], while [16] subsequently extended this effort to the training split. Both corrected splits are now commonly distributed through [16]. Despite its relatively low resolution by today’s standards, this cleaned version remains widely used for model evaluation and comparison.

FLIR Thermal ADAS

Proposed by Teledyne FLIR, it is an Advanced Driver Assistance System dataset. Its second version is composed of 13 k pairs of RGB-T images, offered both in 16 and 8-bit for the infrared images. Initially annotated with 5 classes, it was extended to 16 classes with the release of its second version in 2022 [11]. It presents images acquired mostly in favorable weather conditions, where neither the visible nor the infrared sensors are significantly challenged. The dataset was criticized by Zhang et al. [17] for offering unaligned annotations for the RGB and infrared modalities. The misaligned images were eventually removed, leading to the FLIR aligned subset of the original FLIR ADAS dataset.

LLVIP

Although it consists only of static video surveillance images, the LLVIP dataset [13] is another recent dataset that caught the community’s interest thanks to its size and the quality of its images and labels. It makes a good benchmark for multispectral object detection models. Furthermore, the camera setup used offers optically aligned infrared and visible images, hence removing the geometrical alignment constraint that most other datasets of its kind suffer from. This explains



FIG. 1. Three samples of the multimodal detection subset of the LYNRED Mobility Dataset showing the diversity of the dataset.

TABLE I. Datasets Specifications for All the Presented RGB-T Datasets

Datasets	# Train Image Pairs (k)	# Classes	Camera Setup		Image Resolution (Width × Height)		Intermodal Alignment		Metadata		
			On Vehicle	Stationary	IR	RGB	Fully Aligned	Chessboard	Outside Temperature	Season	Occlusion
KAIST sanitized [18]	47.5	1	✓	—	640 × 480	640 × 480	✓	—	—	—	✓
FLIR ADAS [11]	13	16	✓	—	640 × 512	1800 × 160*	—	✓	—	—	✓
LLVIP [13]	15	1	—	✓	1080 × 720	1080 × 720	✓	—	—	—	—
M3FD [19]	4.2	6	✓	✓	1024 × 768*	1024 × 768*	✓	—	—	—	—
FLIR aligned [17]	5	3	✓	—	640 × 512	640 × 512	—	✓	—	—	—
LYNRED-MDS (ours)	4	9	✓	—	640 × 480	1280 × 960	—	✓	✓	✓	✓

Note: We use (*) to signal that image size is variable. In this case, only the most frequent size is displayed.

the growing attention this dataset has received since its release in 2021.

M3FD

Proposed by Liu et al. [19], it consists of both static and on-vehicle images captured in Dalian, China. This dataset offers 4.2 k RGB-T pairs for training and 300 for testing, and 6 annotated classes (People, Car, Bus, Motorcycle, Lamp, Truck) representing most road users. Even though it is on the smaller side of the spectrum for RGB-T datasets, the diversity of scenes offered, as well as the associated annotations on weather and lighting conditions, make it an interesting case study for real-world ADAS scenarios. In addition to the original annotations, [20] offers further labels for weather conditions that can be useful for evaluation purposes.

Honorable Mentions

Among existing datasets, InfraParis [14] is the most similar in scope to ours, as it provides annotations for both infrared and RGB images acquired around Paris, France. However, due to significant differences in the fields of view between

the two modalities, it is unsuitable for multispectral or multimodal object detection tasks. Another notable dataset is MFNet [21], which offers both bounding box annotations and semantic segmentation masks for a variety of road users in urban environments in Tokyo, Japan. Lastly, the SJTU Multispectral Object Detection (SMOD) Dataset [18] focuses on scenes with a high density of vulnerable road users, particularly pedestrians and cyclists, making it well-suited for safety-critical perception research.

COLLECTION METHODS AND DESIGN

Data Collection Setup Description

The multimodal detection subset of the LYNRED Mobility Dataset presented in this article, comprises 4000 RGB–infrared image pairs collected from 12 driving sequences recorded around Grenoble, France. Data acquisition was performed using two sensor pairs—each consisting of an infrared and an RGB camera—mounted on an aluminum frame fixed to the roof of a test vehicle.

Infrared data was acquired using two different infrared camera sensors, each contributing approximately half of the provided images. The first being a LYNRED PICO640Gen2

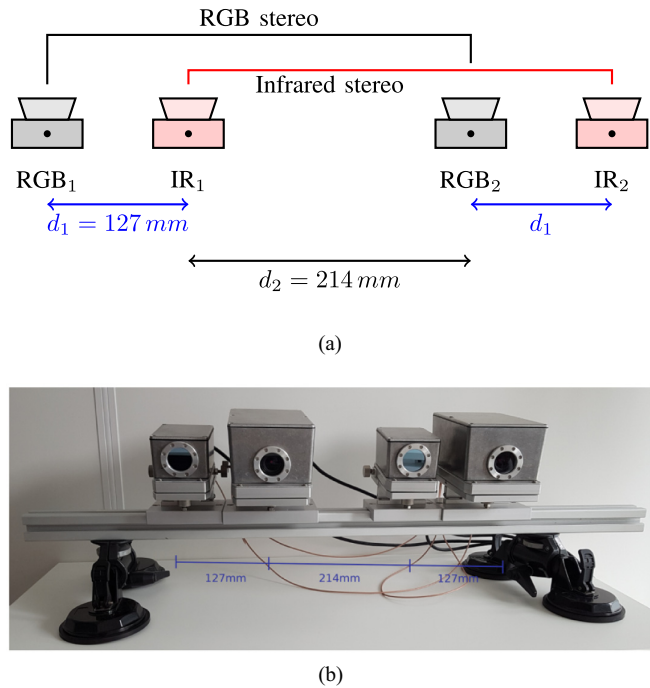


FIG. 2. Schematic representation of the double stereo setup used for the dataset acquisition. For the object detection problem, only the left pair is kept. In order to get an image aligned in the zone of interest, the sensors are angled such that object at infinity (in our case over 10 m) are aligned on all sensors. (a) Schematic of the setup. (b) Photograph of the actual setup on the roof of the car.

sensor, used for 2518 images while the remaining 1482 were acquired with a LYNRED ATTO640D-02. The ATTO sensor is equipped with optics providing a 30° horizontal field of view and an aperture of $f/1.2$. While the PICO sensor is mounted with a $42^\circ f/1.2$ lens. This setup ensures consistent framing across the different sensors while maintaining sensitivity under low-light and long-range conditions.

The RGB cameras integrate a Sony IMX273 sensor (1448×1086 pixels) with a lens allowing a horizontal field of view of 45° and an aperture of $f/1.4$. Given the difference in field of view between the RGB and IR camera sensors, RGB images are cropped from 1448×1086 pixels to 1280×960 pixels retaining only the rightmost and vertically centered region as it represents the region covered by their paired IR sensors.

The system’s geometry is described in Fig. 2. The main RGB camera (RGB₁ on Fig. 2) controls a trigger signal for the three other cameras to obtain synchronous acquisition from each sensor. Images are acquired at a standard frame rate of 30 images per second, and one image every 3 s is kept to maximize diversity in the training set. The sensors were initially oriented such that objects situated 10 m away from the vehicle and 1 m above the ground are approximately centered in all sensors’ fields of view. This configuration was chosen to ensure good alignment between modalities for relevant detection targets.

Dataset Content Description

This dataset is—to the best of our knowledge—the first large-scale European ADAS dataset specifically focusing on VRU detection in multimodal infrared imagery. LYNRED-MDS consists of a total of 4000 images, divided into 3168 image pairs for training and 832 for testing. The data was collected in the surroundings of Grenoble (France). It includes a diverse range of driving scenarios, from winter snowy ski resorts to sunny summer urban areas.

Images were acquired all year round, covering a wide range of ambient temperatures from 3.7°C to 30°C . This temperature variability presents a significant challenge for thermal or long-wave infrared (LWIR) sensors since they operate by capturing thermal radiation. When environmental temperatures approach that of the human body, the contrast between objects and their surroundings decreases, potentially affecting detection performance; see Fig. 3 for an example.

To prevent data leakage, the sequences were split into separate training and testing subsets while ensuring a representative distribution of time of day and seasonal conditions. The distribution of the dataset split across various parameters is illustrated in Fig. 4.

Additionally, nine different classes are annotated to reflect the diversity of road users. A comprehensive comparison of class labels across the existing datasets’ training sets is presented in Table II. Fewer classes were chosen to counteract class imbalance. Indeed, 6 out of the 16 classes of FLIR have under 100 samples in the train set and most of these are not present in the test set [11]. Furthermore, the occlusion status of human-class objects is also recorded, following the method used in other existing infrared object detection datasets [11], [13], [18]. The full annotation protocol is described in “Validation and Quality” section.

We provide both 16-bit-per-pixel and processed 8-bit-per-pixel infrared images, as tone mapping has been shown to be a relevant component of the embedded object detection pipeline [22]. The 16-bit images correspond to the minimally processed output of the infrared camera.

In Fig. 5, we plot the empirical cumulative distribution of bounding box height for each dataset; this allows us to pinpoint some important differences between the datasets concerning object size annotations. We can clearly see that our dataset is on par with most other RGB-T datasets, resembling M3FD the most with respect to its bounding box height distribution. The task shift between LLVIP and the others is easily noticeable, as objects in video surveillance clips are at a controlled distance from the camera, leading to a smaller variance and overall bigger boxes.

Finally, in the multimodal detection subset of LYNRED-MDS objects as small as ten pixels in height are annotated. This is done to evaluate the efficiency of infrared sensors for pedestrian detection in the context of regulatory standards cited in the introduction [2]. To assess the capability of detecting pedestrians at a distance of 50 m, we estimate the

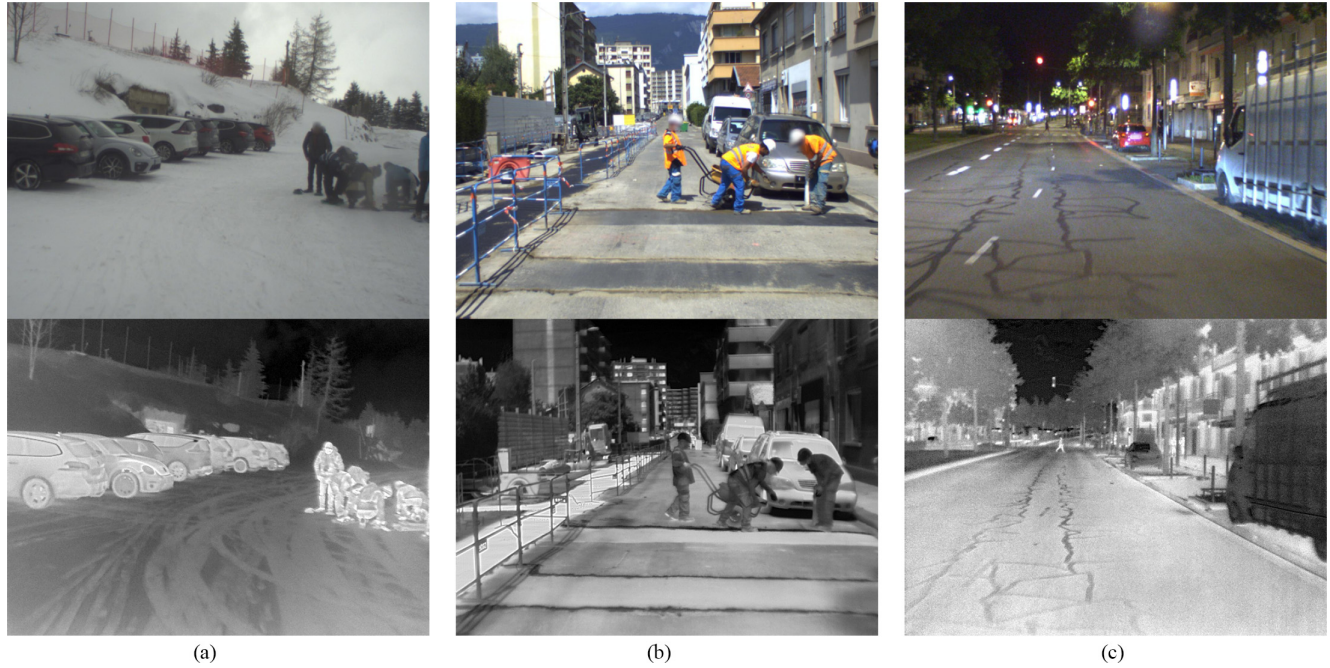


FIG. 3. Images from a broad range of seasonal and heat scenarios are present in the LYNRED Mobility Dataset, offering challenging conditions for developing infrared-based AEB. (a) External temperature: 5.4 °C. (b) External temperature: 30 °C. (c) External temperature: 20.6 °C.

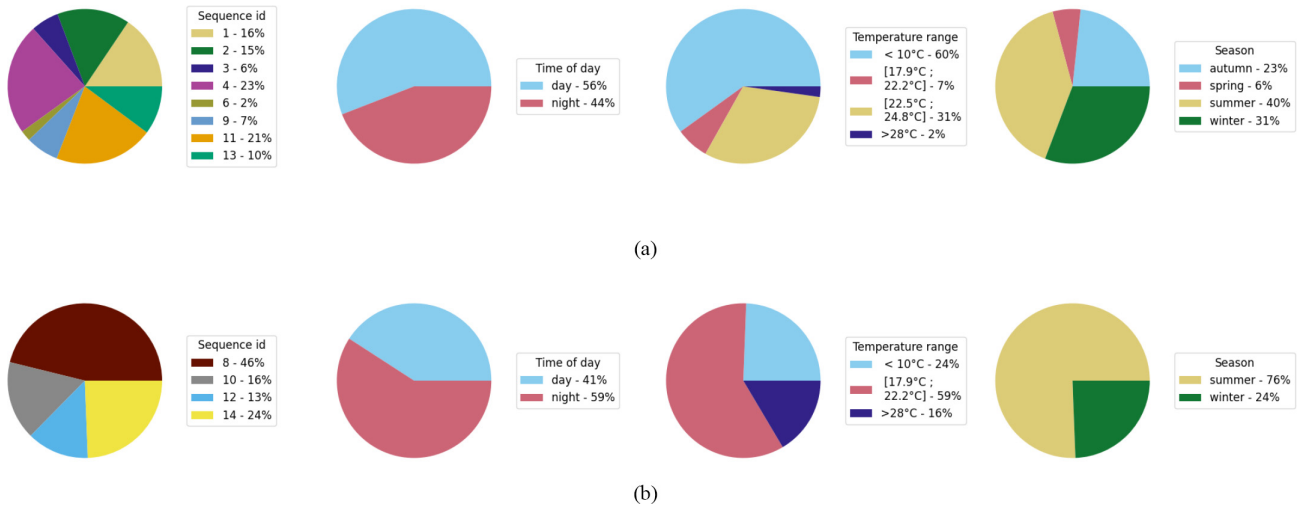


FIG. 4. Decomposition of the different data splits. (a) Train dataset. (b) Test dataset.

TABLE II. Number of Instances Per Class Across the Different Reference RGB-T Datasets

Datasets	Number of Instances Per Class															
	Person	Bicycle	Car	Motorcycle	Bus	Train	Truck	Light	Hydrant	Sign	Animal	Skateboard	Stroller	Scooter	Other Vehicle	
KAIST [10]	103 128	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
FLIR ADAS [11]	50 478	7237	73 623	1116	2245	5	829	16 198	1095	20 770	12*	29	15	15	1373	
LLVIP [13]	34 137	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
M3FD [19]	7443	-	11 737	344	449	-	654	1542	-	-	-	-	-	-	-	
KAIST sanitized [18]	24 247	-	-	-	-	-	-	-	-	-	-	-	-	-	-	
FLIR aligned [17]	8987	2566	20608	-	-	-	-	-	-	-	-	-	-	-	-	
LYNRED-MDS (ours)	9152	1755	15 924	310	244	304	514	-	-	-	80	-	-	-	91	

Note: The number of annotations displayed corresponds to the infrared train split. We notice that even though FLIR offers far more classes than others, the class imbalance makes most of those useless for training. The Starred * value represents blended classes (four dogs and eight deers).

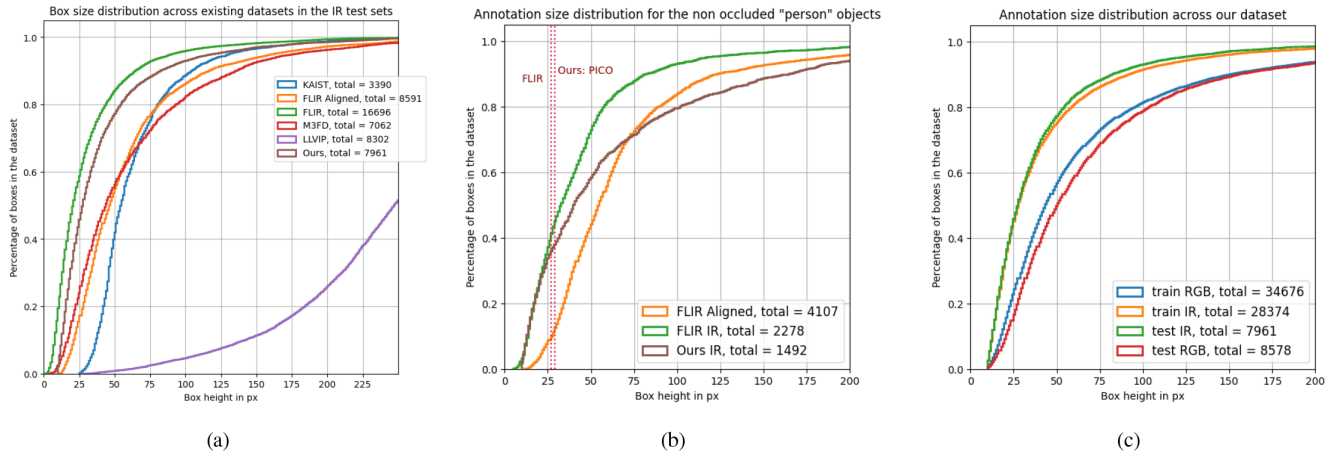


FIG. 5. Cumulative histograms of bounding boxes height for all presented datasets (a), comparison between FLIR ADAS, FLIR Aligned, and LYNRED-MDS IR splits (b), and LYNRED-MDS’s RGB and IR modalities (c). The gap between IR and RGB annotation sizes for curve (c) can be explained by the difference between the IR and RGB sensors specifications.

expected pixel height h_{object} of an unobstructed human in our dataset by $h_{\text{object}} = (f \times H_{\text{object}} \times h_{\text{image}}) / (D_{\text{object}} \times H_{\text{sensor}})$ derived from [23] and where H_x and h_x refer to physical height (mm) or height in pixels, respectively.

For an average adult pedestrian height provided by NHTSA as $H_{\text{object}} = 1.750$ m [24] and in order to meet regulatory requirements with our sensor setup specifications, we compute a necessity to detect objects with a minimum height of approximately 28.8 pixels on the PICO sensor (focal length $f = 14$ mm, image height $h_{\text{image}} = 480$ pixels, sensor height $H_{\text{sensor}} = 8.16$ mm, object distance $d_{\text{object}} = 50$ m), respectively 40.8 pixels for the ATTO sensor (sensor height $H_{\text{sensor}} = 5.76$ mm).

Applying the same calculation to FLIR’s dataset yields a comparable value of 26.8 pixels. We also performed the computation for a child-sized pedestrian (approximately 100 cm tall), which establishes more challenging detection threshold of 16.5 and 23.3 pixels for the PICO and the ATTO sensors, respectively. Meanwhile the value would be 11.9 pixels in FLIR’s dataset.

In Fig. 5(b), we illustrate the pixel height thresholds for all unobstructed pedestrian annotations in the FLIR ADAS validation set, as well as the FLIR Aligned and LYNRED-MDS test datasets. Notably, in both datasets, more than 35% of annotated humans are 29 pixels tall or smaller. Per NHTSA regulations [2], a vehicle traveling at 100 km/h must detect pedestrians at a minimum of 50 m to allow sufficient braking time, making this the most demanding AEB scenario—successful detections at 50 m implying successful detections at shorter distances and lower speeds. These datasets thus provide a challenging benchmark for AEB-relevant small pedestrian detection, especially in critical situations where detecting small objects remains a significant challenge for object detection algorithms [25], [26], [27], though a comprehensive AEB evaluation would additionally require consideration of motion patterns and spatiotemporal continuity.

Images were anonymized using the blurring algorithm from [28] targeting license plates and faces to make the dataset compliant with the regulations of the CNIL (*Commission Nationale de l’Informatique et des Libertés*), the French data protection authority.

VALIDATION AND QUALITY

Annotation Protocol and Guidelines

Annotations were performed by a professional labeling company and subsequently verified by an in-house expert in infrared imaging. Annotators labeled RGB and infrared images separately and independently, annotating objects exclusively in the modality where they are visible. Bounding boxes were required to be as tight as possible, with a minimum object height of ten pixels. Every annotation carries a mandatory occlusion tag: None (fully visible), Partial (more than 30% visible), or Heavy (less than 30% visible). Objects were only annotated if identifiable by a human annotator. Nine object classes were defined: person, car (smaller than a Mercedes Sprinter), truck (bigger than a Mercedes Sprinter), bus, bicycle, motorcycle, train (including tramways), animal, and construction machine, with a special tag used for atypical instances. Riders were annotated separately from their vehicles and tagged accordingly. Bicycles and motorcycles, were annotated as such even if not ridden.

The dataset repository contains two annotation versions `metadata` and `metadata_small_objects`. The `metadata_small_objects` folder contains the original annotations from a first expert phase, which include objects as small as four pixels in height. However, since this finer labelling was only applied to a subset of the data, it introduces inconsistencies that can lead to unreliable evaluation and are only kept as legacy. The `metadata` folder contains the unified annotation set, in which all bounding boxes below ten pixels in height are excluded regardless of their phase of origin, ensuring full consistency with the protocol described

TABLE III. Mean Average Precision (mAP@50) for Human Detection Across (a) IR and (b) RGB Datasets

Test Dataset	Train Dataset					
	ALL	LLVIP	FLIR	FLIR Aligned	M3FD	Ours
ALL	0.88	0.61	0.54	0.54	0.59	0.73
LLVIP	0.96	0.96	0.41	0.44	0.64	0.77
FLIR ADAS	0.61	0.11	0.80	0.55	0.43	0.62
FLIR Aligned	0.85	0.22	0.79	0.81	0.48	0.76
M3FD	0.85	0.29	0.67	0.54	0.86	0.61
ours	0.65	0.18	0.49	0.45	0.34	0.60

(a) IR

Test Dataset	Train Dataset					
	ALL	LLVIP	FLIR	FLIR Aligned	M3FD	Ours
ALL	0.77	0.54	0.36	0.31	0.55	0.45
LLVIP	0.97	0.54	0.17	0.11	0.36	0.34
FLIR ADAS	0.44	0.14	0.70	0.39	0.39	0.52
FLIR Aligned	0.69	0.24	0.58	0.64	0.51	0.47
M3FD	0.74	0.23	0.42	0.24	0.77	0.38
ours	0.49	0.07	0.36	0.19	0.22	0.51

(b) RGB

Note: Only the “person” class is retained for comparison purposes. Each column uses a same training dataset and each row uses a same evaluation dataset. The “ALL” dataset combines LLVIP, FLIR Aligned, M3FD, and ours. All experiments were conducted using YOLOv8n trained for 200 epochs. For each row, the best result is highlighted in bold and dark gray; the second best is highlighted in light gray.

above. All experiments and benchmarks reported in this study are based exclusively on the metadata annotations.

Experimental Results: Unimodal Cross-Dataset Evaluation

To benchmark our dataset, we evaluated the finetuning performances of YOLOv8n [29], pretrained on the COCO dataset [30]. We selected this model for its ease of use and suitability for deployment. The model is fine-tuned for 200 epochs with a batch size of 16, using the “auto” optimizer setting provided by Ultralytics.

Rather than focusing solely on in-domain performance, our work emphasizes the model’s generalization capabilities. To assess this, we adopted mAP@50 (mean Average Precision at an Intersection-over-Union threshold of 0.5) as our evaluation metric. This threshold is commonly used in object detection benchmarks and offers a reasonable balance between detection quality and tolerance for localization error, particularly important when detecting small or low-contrasted objects in thermal imagery.

Our generalization test protocol involves training the model on one dataset and evaluating it on the others. We focus exclusively on the “person” or “human” class, as it is the most relevant for VRU detection tasks. The datasets used in our experiments are FLIR ADAS, FLIR Aligned, M3FD, LLVIP, and ours. Additionally, we include a combined dataset, referred to as ALL, which merges LLVIP, FLIR Aligned, M3FD, and ours. This serves as a topline reference. The comparative results across datasets are presented in Table III.

It highlights several key findings. First, the model trained on the “ALL” dataset achieves the best performance on its constituent datasets: LLVIP, FLIR Aligned, M3FD, and ours. In contrast, performance on FLIR ADAS—excluded from the “ALL” set—is notably lower. This, as well as FLIR Aligned’s performance on FLIR ADAS’s test set, suggests that FLIR Aligned may not fully capture the data distribution of the broader FLIR ADAS dataset.

Second, we observe a modality gap: IR generally outperforms RGB across datasets in object detection. This trend is especially pronounced in datasets such as LLVIP, which focuses on nighttime or low-light conditions where RGB sensors are less effective. Interestingly, despite LLVIP’s

size being far larger than most other datasets, its models perform poorly in cross-dataset validation. This emphasizes the relevance of task-specific datasets focusing on driving scenarios.

These preliminary results, suggest that LYNRED-MDS offers competitive generalization potential among the compared datasets. The model trained on our IR images achieves the second-best cross-dataset mAP on three out of five target datasets, and the LYNRED-MDS test split consistently ranks among the more challenging evaluation sets for models trained on other driving datasets. While these observations are encouraging, broader conclusions about generalization would require validation across additional architectures and metrics. Nevertheless, LYNRED-MDS appears to be a promising and challenging benchmark for infrared-based pedestrian detection in ADAS contexts.

RECORDS AND STORAGE

Data is stored on LYNRED’s website (<https://www.lynred.com/lynred-mobility-dataset>), alongside two other datasets. The dataset is organized in six folders as described in Fig. 6. The folders tagged “_aligned” correspond to images that have been geometrically aligned through the chessboard method described earlier. The dataset annotation data is presented using the standard COCO JSON format [30] to which we add a few additional fields. These fields are described in Table IV.

INSIGHTS AND NOTES

Multispectral Object Detection

Although not used in this study, the paired images can be used for multispectral object detection. In that sense, it is comparable to other commonly used datasets such as FLIR Aligned and M3FD. To further facilitate the use of the dataset in multispectral object detection, it is planned to release an aligned version that corrects parallax for close objects.

Tone Mapping

The release of 16-bit “RAW” infrared images allows their use for tone-mapping-free object detection, further reducing computational load in an AEB scenario, as well as enabling

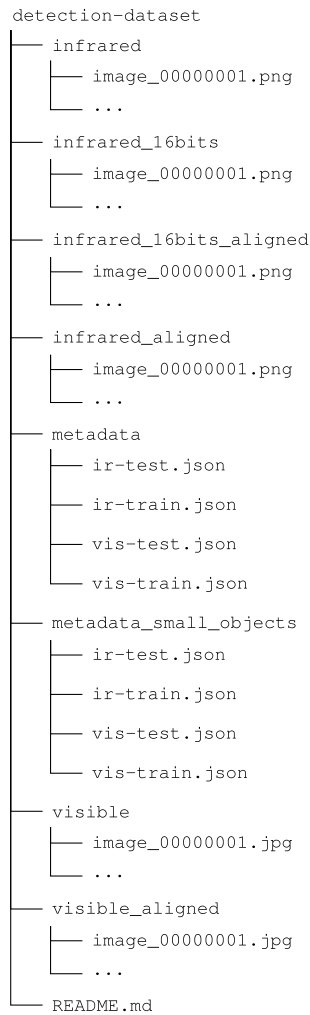


FIG. 6. Overview of the structure of the dataset folder.

TABLE IV. Description of JSON Fields for Images and Annotations

Field	Description/Use
Images	
id	Unique identifier of the image
width, height	Image resolution in pixels
file_name	File name of the thermal image (PNG)
visible_image	Corresponding RGB image file (JPG)
season	Season of acquisition (e.g., summer, winter)
time_of_day	Time of acquisition (e.g., day, night)
tamb	Ambient temperature at acquisition (°C)
sequence_id	Sequence identifier (groups images from the same sequence)
author	Source/creator of the data (LYNRED or Neovision)
Annotations	
id	Unique identifier of the annotation
image_id	Reference to the corresponding image (via images.id)
category_id	Object class (links to categories.id)
bbox	Bounding box coordinates [x, y, width, height] in pixels
area	Area covered by the bounding box
iscrowd	COCO convention flag (0: normal object, 1: crowd/ambiguous)

the development of optimal tone mapping algorithms for object detection tasks.

Regional Adaptation

While our dataset reflects Western European driving conditions and vehicle standards, we acknowledge that adaptation to other regional contexts—such as different vehicle typologies, road infrastructure, or climate profiles—may require additional data collection or domain adaptation techniques. We encourage the community to build upon this dataset and extend it to other geographical contexts using the protocol provided in this article.

SOURCE CODE AND SCRIPTS

The code used to obtain Figs. 4, 5, and Table III can be found on this GitHub repository: https://github.com/arbezlo/LYNRED_Mobility_Dataset.

ACKNOWLEDGMENT

All the authors contributed to the manuscript.

L.A. was the main author of the manuscript and was responsible for the data split and the experiments. J.M. and X.B. handled dataset collection and supervised labeling.

The authors have declared no conflicts of interest.

REFERENCES

- [1] M. Beedham, “How risk tolerance shapes ADAS regulation: The U.S. and European approaches” 2025. Accessed: Jan. 30, 2025. [Online]. Available: <https://www.tomtom.com/newsroom/explainers-and-insights/risk-tolerance-and-adas-regulation-eu-and-usa/>
- [2] National Highway Traffic Safety Administration, part of the U.S. Department of Transportation, “NPRM: Automatic emergency braking systems,” 2025. Accessed: May 13, 2025. [Online]. Available: <https://www.nhtsa.gov/document/nprm-automatic-emergency-braking-systems>
- [3] “REGULATION (EU) 2019/2144 of the EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 November 2019 on Type-Approval Requirements for Motor Vehicles and Their Trailers, and Systems, Components and Separate Technical Units Intended for Such Vehicles, as Regards Their General Safety and the Protection of Vehicle Occupants and Vulnerable Road Users, Amending Regulation (EU) 2018/858 of the European Parliament and of the Council and Repealing Regulations (EC) no 78/2009, (EC) no 79/2009 and (EC) no 661/2009 of the European Parliament and of the Council and Commission Regulations (EC) no 631/2009, (EU) no 406/2010, (EU) no 672/2010, (EU) no 1003/2010, (EU) no 1005/2010, (EU) no 1008/2010, (EU) no 1009/2010, (EU) no 19/2011, (EU) no 109/2011, (EU) no 458/2011, (EU) no 65/2012, (EU) no 130/2012, (EU) no 347/2012, (EU) no 351/2012, (EU) no 1230/2012 and (EU) 2015/166” 2024. Accessed: Jan. 24, 2024. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2019/2144/2024-07-07>
- [4] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? The KITTI vision benchmark suite,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2012.
- [5] M. Cordts et al., “The Cityscapes dataset for semantic urban scene understanding,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016.
- [6] C. Corsi, “New frontiers for infrared,” *Opto-Electron. Rev.*, vol. 23, no. 1, pp. 3–25, 2015.
- [7] C. Ley, F. Heath, T. Hastie, Z. Gao, M. Protsiv, and J. Parsonnet, “Defining usual oral temperature ranges in outpatients using an unsupervised learning algorithm,” *JAMA Internal Med.*, vol. 183, no. 10, pp. 1128–1135, 2023.
- [8] A. González et al., “Pedestrian detection at day/night time with visible and FIR cameras: A comparison,” *Sensors*, vol. 16, no. 6, 2016, Art. no. 820.
- [9] R. Donà et al., “Thermal cameras and their safety implications for pedestrian protection: A mixed empirical and simulation-based

- characterization,” *Transp. Res. Rec.*, vol. 2679, no. 3, 2024, Art. no. 03611981241278346.
- [10] S. Hwang, J. Park, N. Kim, Y. Choi, and I. So Kweon, “Multispectral pedestrian detection: Benchmark dataset and baseline,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1037–1045.
- [11] Teledyne FLIR OEM, “FREE teledyne FLIR thermal dataset for algorithm training” 2025. Accessed: May 12, 2025. [Online]. Available: <https://www.flir.ca/oem/adas/adas-dataset-form/>
- [12] A. Wilson, K. A. Gupta, B. H. Koduru, A. Kumar, A. Jha, and L. R. Cenkeramaddi, “Recent advances in thermal imaging and its applications using machine learning: A review,” *IEEE Sensors J.*, vol. 23, no. 4, pp. 3395–3407, Apr. 2023.
- [13] X. Jia, C. Zhu, M. Li, W. Tang, and W. Zhou, “LLVIP: A visible-infrared paired dataset for low-light vision,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 3496–3504.
- [14] G. Franchi, M. Hariat, X. Yu, N. Belkhir, A. Manzanera, and D. Filliat, “InfraParis: A multi-modal and multi-task autonomous driving dataset,” in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vision*, 2024, pp. 2973–2983.
- [15] L. Jingjing, “Multispectral deep neural networks for pedestrian detection,” 2026. Accessed: Mar. 17, 2026. [Online]. Available: <https://web.archive.org/web/20170911034415/http://paul.rutgers.edu:80/~jl1322/multispectral.htm>
- [16] C. Li, D. Song, R. Tong, and M. Tang, “Multispectral Pedestrian Detection via Simultaneous Detection and Segmentation,” 2018, *arXiv:1808.04818*.
- [17] H. Zhang, E. Fromont, S. Lefevre, and B. Avignon, “Multispectral fusion for object detection with cyclic fuse-and-refine blocks,” in *Proc. IEEE Int. Conf. Image Process. (ICIP)*. Piscataway, NJ, USA: IEEE Press, 2020, pp. 276–280.
- [18] Z. Chen, Y. Qian, X. Yang, C. Wang, and M. Yang, “AMFD: Distillation via adaptive multimodal fusion for multispectral pedestrian detection,” 2024, *arXiv:2405.12944*.
- [19] J. Liu et al., “Target-aware dual adversarial learning and a multi-scenario multi-modality benchmark to fuse infrared and visible for object detection,” in *Proc. IEEE/CVF Conf. Comput. Vision Pattern Recognit.*, 2022, pp. 5802–5811.
- [20] S. A. Deevi, C. Lee, L. Gan, S. Nagesh, G. Pandey, and S.-J. Chung, “RGB-x object detection via scene-specific fusion modules,” in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vision (WACV)*, Jan. 2024, pp. 7366–7375.
- [21] K. Takumi, K. Watanabe, Q. Ha, A. Tejero-De-Pablos, Y. Ushiku, and T. Harada, “Multispectral object detection for autonomous vehicles,” in *Proc. Thematic Workshops ACM Multimedia*, 2017, pp. 35–43.
- [22] C. Karam, J. Matias, X. Breniere, and J. Chanussot, “Optimizing the image correction pipeline for pedestrian detection in the thermal-infrared domain,” 2024, *arXiv:2407.04484*.
- [23] R. Szeliski, *Computer Vision: Algorithms and Applications*. Switzerland: Springer Nature, 2022.
- [24] National Highway Traffic Safety Administration, part of the U.S. Department of Transportation, “Federal Motor Vehicle Safety Standards; Pedestrian Head Protection” 2025. Accessed: May 13, 2025. [Online]. Available: <https://www.nhtsa.gov/document/nprm-pedestrian-head-protection-standard>
- [25] M. M. Gündoğan, T. Aksoy, A. Temizel, and U. Halici, “IR reasoner: Real-time infrared object detection by visual reasoning,” in *Proc. IEEE/CVF Conf. Comput. Vision Pattern Recognit.*, 2023, pp. 422–430.
- [26] S. Li, Y. Li, Y. Li, M. Li, and X. Xu, “YOLO-FIRI: Improved YOLOv5 for infrared image object detection,” *IEEE Access*, vol. 9, pp. 141861–141875, 2021.
- [27] W. Wei, Y. Cheng, J. He, and X. Zhu, “A review of small object detection based on deep learning,” *Neural Comput. Appl.*, vol. 36, no. 12, pp. 6283–6303, 2024.
- [28] Understand.ai, “Anonymizer,” 2022. Accessed: Jul. 9, 2022. [Online]. Available: <https://github.com/understand-ai/anonymizer>
- [29] G. Jocher, J. Qiu, and A. Chaurasia, “Ultralytics YOLO,” Jan. 2023. Accessed: Feb. 21, 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [30] T.-Y. Lin et al., “Microsoft COCO: Common objects in context,” in *Proc. Eur. Conf. Comput. Vision (ECCV)*. Zurich, Switzerland: Springer, 2014, pp. 740–755.