

SURVEY

RGB-D and Thermal Sensor Fusion: A Systematic Literature Review

MARTIN BRENNER¹, NAPOLEON H. REYES¹, TEO SUSNJAK¹, AND ANDRE L. C. BARCZAK²

¹School of Mathematical and Computational Sciences, Massey University, Auckland 0632, New Zealand

²Centre for Data Analytics, Bond University, Gold Coast, QLD 4226, Australia

Corresponding author: Martin Brenner (martin.brenner.1@uni.massey.ac.nz)

ABSTRACT In the last decade, the computer vision field has seen significant progress in multimodal data fusion and learning, where multiple sensors, including depth, infrared, and visual, are used to capture the environment across diverse spectral ranges. Despite these advancements, there has been no systematic and comprehensive evaluation of fusing RGB-D and thermal modalities to date. While autonomous driving using LiDAR, radar, RGB, and other sensors has garnered substantial research interest, along with the fusion of RGB and depth modalities, the integration of thermal cameras and, specifically, the fusion of RGB-D and thermal data, has received comparatively less attention. This might be partly due to the limited number of publicly available datasets for such applications. This paper provides a comprehensive review of both, state-of-the-art and traditional methods used in fusing RGB-D and thermal camera data for various applications, such as site inspection, human tracking, fault detection, and others. The reviewed literature has been categorised into technical areas, such as 3D reconstruction, segmentation, object detection, available datasets, and other related topics. Following a brief introduction and an overview of the methodology, the study delves into calibration and registration techniques, then examines thermal visualisation and 3D reconstruction, before discussing the application of classic feature-based techniques and modern deep learning approaches. The paper concludes with a discourse on current limitations and potential future research directions. It is hoped that this survey will serve as a valuable reference for researchers looking to familiarise themselves with the latest advancements and contribute to the RGB-DT research field.

INDEX TERMS Multimodal, RGB-D, RGB-DT, RGB-T, sensor fusion, thermal.

I. INTRODUCTION

The extraction and analysis of features from RGB images have become a widely used processing technique in computer vision, finding its way into a diverse array of industrial, commercial, and everyday applications. However, this technique exhibits limitations, primarily from its confinement to the visible spectrum. As illustrated in Fig. 1, the visible imaging range is notably narrower compared to other spectra, which underscores the potential benefits of exploring alternative non-visible spectral regions to overcome these restrictions. The most significant constraint is that it only operates effectively under good lighting conditions and clear visibility. This has prompted researchers to explore using RGB-D and

thermal cameras for multi-spectral perception in recent years as shown in Fig. 2.

The increasing application of depth cameras can largely be attributed to the release of the Microsoft Kinect sensor in 2010. This sensor utilises an infrared (IR) structured light system, operating in the Near Infra Red (NIR) band, to capture depth information in addition to RGB colour data, and was the first depth camera to be widely available for the consumer market. Thermal cameras on the other hand capture temperature information and have been available for many years. Despite a drop in price, the cost of thermal cameras is still considerably high and the possible resolution of the sensors is low compared to RGB cameras due to the larger pixel pitch required for the Long-Wave Infrared (LWIR) band. With the introduction of the first microbolometric array camera in 1997, detector cooling in thermal cameras became unnecessary [1], as non-cooled thermal imagers now feature

The associate editor coordinating the review of this manuscript and approving it for publication was Mehul S. Raval¹.




FIGURE 1. Depiction of the electromagnetic spectrum, with a focus on the various infrared (IR) bands. The thermal IR range, which is radiation-based, is specifically indicated.




FIGURE 2. Number of publications with RGB-DT, RGB-T and RGB-D from 2012 to 2022. Source: Data from Google Scholar keyword search: (“RGB-DT” OR “rgb-dt”); (“rgb-t” OR “rgbt” OR “RGB-Thermal”); (“rgb-d” OR “rgbd” OR “RGB-Depth” OR “RGB+Depth”)[peer reviewed articles only].

electronic stabilisation. While non-cooled cameras offer advantages such as being lighter, faster, more affordable, and more reliable, cooled cameras still have the edge in terms of greater sensitivity [2].

In the context of sensor fusion, each sensor modality - RGB, Depth (D), and Thermal (T) - brings its own set of advantages and disadvantages.

RGB cameras, being ubiquitous and economical, offer high-resolution colour images that are readily interpreted by both human observers and computer vision algorithms. They excel in tasks such as object recognition, scene understanding, and texture analysis. However, their efficacy is heavily dependent on good lighting conditions.

Conversely, depth sensors, integral to RGB-D cameras, operate relatively independently of visible light, allowing them to function effectively under a variety of lighting conditions. Nevertheless, their range is typically limited, and they can be affected by factors such as sunlight interference (in the case of Time of Flight sensors) or low texture areas and lighting conditions (in the case of stereo vision). Despite these

limitations, depth sensors offer valuable 3D environmental information, proving advantageous for tasks such as object detection, localisation, and navigation.

Thermal infrared sensors, in contrast to the visible and depth modalities, can sense slight temperature differences between objects and their surroundings. This capability is not hindered by low-light conditions or complete darkness, as these sensors operate based on thermal radiation, independent of any light source. This unique capability makes thermal sensing a valuable modality for object detection under challenging conditions. However, thermal sensors typically offer lower resolution than RGB cameras and are more expensive.

The overarching aim of sensor fusion is to harmonise the strengths of each sensor modality to mitigate their individual limitations. However, achieving effective fusion requires careful calibration and alignment of the sensors, along with sophisticated algorithms to integrate the different types of data.

The field of surveillance has shown significant interest in the integration of RGB and Thermal (RGB-T) data. Similarly, the combination of LiDAR sensors or stereo depth cameras with RGB, polarised images, and radar is a well-explored area in autonomous vehicles and robotics. However, the fusion of RGB-D and thermal data has not been studied as extensively in comparison. Fusion of these three modalities has the potential to provide more robust and accurate perception in various applications, such as object recognition, tracking, and localisation for applications where no long-range detection is required or the detection of endotherms is beneficial.

LiDAR and RGB-D cameras are both used for capturing 3D data, but they have different characteristics. LiDAR produces a sparser 3D point cloud with decreasing resolution over distance, while RGB-D cameras produce a more densely packed depth map that is limited to a few metres

of distance. RGB-D sensors that rely on IR Time Of Flight (ToF) technology are not suitable for outdoor applications due to interference from sunlight, but devices based on stereo vision can overcome this issue.

In order to achieve an effective fusion of the different modalities, it is essential to calibrate each sensor and align them in the same coordinate system, which involves determining the intrinsic and extrinsic parameters using the pin-hole camera model. Aligning the modalities correctly is crucial for achieving precise data fusion. Although descriptor-based methods utilising feature point matching algorithms can accomplish registration, they are often not suitable for real-time applications involving moving cameras. This is due to their high computational complexity and the challenges in implementing them with thermal data, which displays distinct characteristics compared to visual data.

sensors, particularly within the realm of autonomous driving. These reviews often delve into the integration of these various sensor modalities and their specific challenges, yet they do not explore the specific tri-modal fusion of RGB, Depth (D), and Thermal (T) sensors. Our review uniquely situates itself at this intersection of sensor fusion, thereby distinguishing it from the broader landscape of sensor fusion literature.

The primary contributions of this review paper are designed to inform researchers working in this field by:

- Presenting a summary of various traditional and current methodologies being utilised.
- Identifying available datasets for furthering this research.
- Highlighting the current research trajectories and various application areas.

The ultimate goal is to provide a comprehensive resource that will ease the entry of interested researchers into this field while identifying trends for others.

As illustrated in Fig. 3, the paper’s structure begins with an introduction followed by a brief background to provide further context. Camera calibration and image registration are reviewed first since they are prerequisites for most approaches and fields of application. The discussion then shifts to modality fusion in general before examining the overlaying of thermal data onto visual data or 3D models for visual inspection or the extraction of thermal data from specific regions of interest. The use of one modality to support another in preprocessing is briefly addressed, followed by the exploration of RGB-DT applications in 3D reconstruction. Subsequently, the paper delves into manual descriptor-based methods and deep learning-based methods. Lastly, available datasets, limitations, and conclusions are presented.

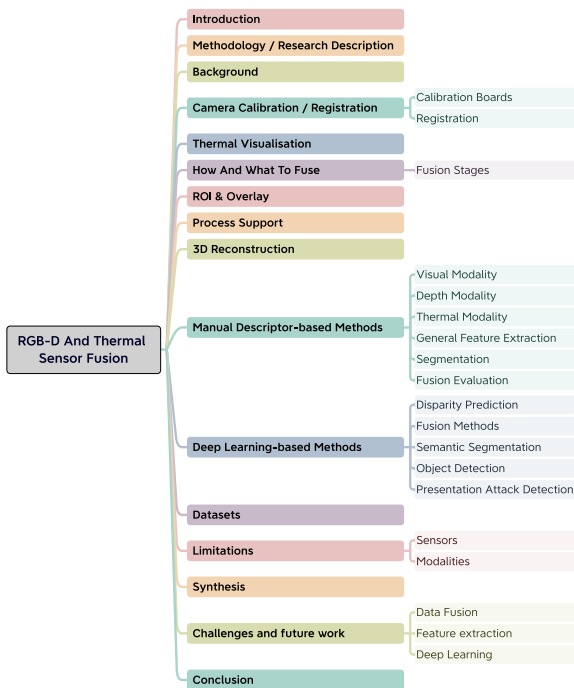


FIGURE 3. The overall structure of this paper.

Overall, the fusion of multiple modalities is an important area of research with many potential applications in various fields. With the advancements in deep learning, it is now possible to construct more advanced systems that can perform complex tasks using fused RGB-DT data.

A. CONTRIBUTION

The primary aim of this survey is to provide a comprehensive and all-encompassing overview of the use of thermal cameras in combination with RGB and depth data. At the time of writing, the authors were unaware of comparable surveys specifically focusing on these technologies. While there are numerous reviews on sensor fusion, they predominantly focus on the amalgamation of LiDAR, Radar, RGB and other

II. METHODOLOGY AND RESEARCH DESCRIPTION

The systematic literature review (SLR) for this study employed the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) methodology [3], which is a widely-used approach that involves a structured process for conducting a comprehensive literature search, applying eligibility criteria, extracting data, synthesising findings, and ensuring the search is reproducible with the same steps, keywords, and tags. The review began by defining the research topic of: RGB-D And Thermal Sensor Fusion. This was then followed by the definition of keywords and search tags used to search scientific databases via Google Scholar as listed in Table 1.

A comprehensive search resulted in the identification of 70 research papers related to the chosen topic. These papers were further refined by utilising exclusion criteria, such as language, repeated papers, and eliminating papers that were not relevant to the techniques under review, as depicted in the PRISMA flow diagram in Fig. 4. Following the implementation of the exclusion criteria, 31 papers were reviewed in detail. Additionally, 16 more relevant documents were added after analysing the references of the initially identified papers, bringing the total number of papers reviewed to 47.

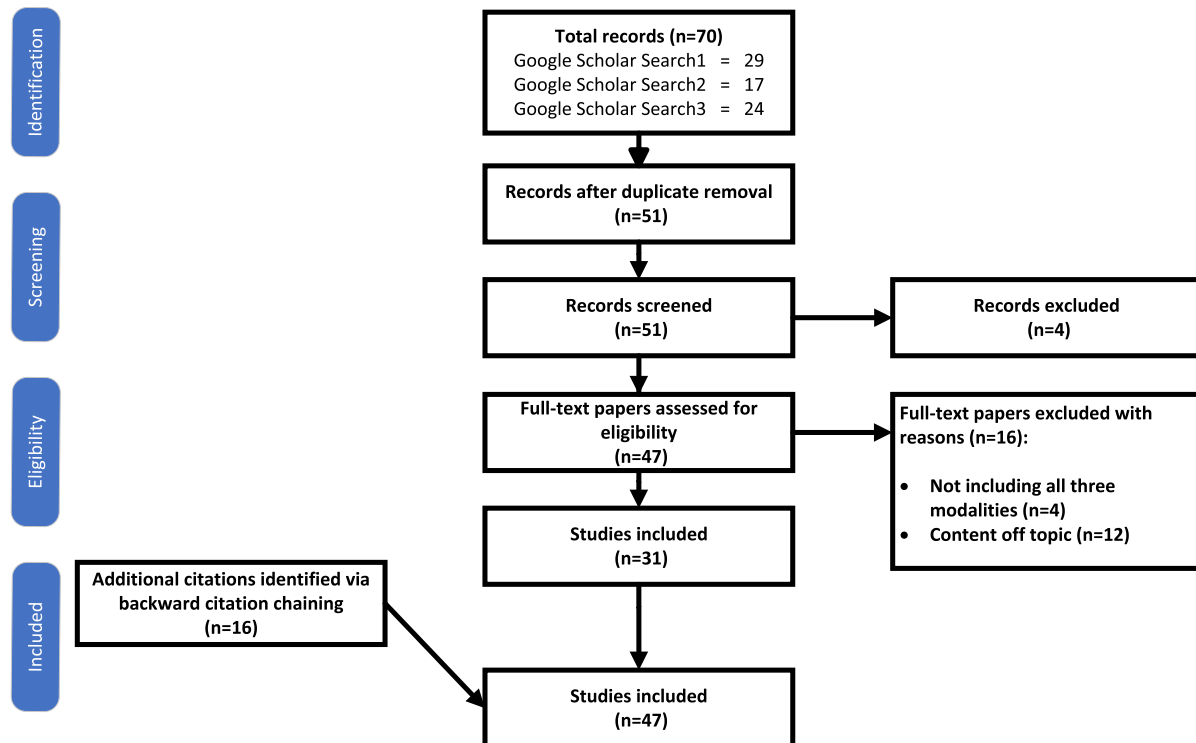


FIGURE 4. PRISMA flow diagram illustrating the search strategy and providing the phases of article identification and selection, which resulted in the identification of 47 papers that were deemed eligible for inclusion in the review. Prepared in accordance with Tricco AC, et al. PRISMA Extension for Scoping Reviews (PRISMA-ScR) [3].

TABLE 1. Literature searches on Google Scholar.

Date	Terms	Filter	Results
14.1.2023	allintitle: thermal ("fuse" OR "fusing" OR "object detection" OR "object-detection" OR detection) ("3D" OR "depth" OR "point cloud" OR "point clouds")	2018	29
29.1.2023	allintitle: thermal rgb-d	-	16 (-1 duplicate)
6.3.2023	allintitle: thermal rgb "rgb-d" OR depth OR "rgb-dt" OR "RGB-DT"	-	6(-18 duplicates)

Further studies that did not precisely match the three modalities, but were considered relevant to support the topic, were also included in this review. A list of the studies that have been included in the analysis can be found in Table 2, along with additional details such as the type of sensors used, their resolution, the frequency of data acquisition, the fusion method employed, and whether or not the system is capable of real-time processing.

A. REVIEW QUESTIONS

In this work, the aim is to answer the following review questions:

- What datasets are currently available for RGB-DT and what scenarios do they cover?
- What are the different methods for fusing the modalities?
- How are modalities weighted during fusion?
- What are the most suitable fusion and detection methods for real-time applications?

- What are the potential application areas for this technology?
- What are the limitations and future prospects?

III. BACKGROUND

The initial research papers that concentrated on fusing RGB, Depth, and Thermal data (RGB-DT) using RGB-D cameras emerged in 2011. Early works in this field investigated medical scans [2], while later in 2013, research expanded to include 3D thermal mapping of building interiors [45], sensor fusion for people tracking [46], and tri-modal person re-identification [47].

Some earlier works proposed systems using different technologies, such as a terrestrial laser scanner and thermal infrared camera [48], or a Structure from Motion (SfM) or MultiView Stereo (MVS) pipeline to generate a dense, coloured point cloud with optional thermal data overlay [49].

Over the last decade then, the fusion of multiple modalities has been increasingly researched as combining different

TABLE 2. Overview of reviewed studies.

Ref	Year	Method	Datasets	RGB(D)	Thermal	T Res	T Hz	Type	Reg	T Use	Fus.	RT	CPU / GPU
[4]	2023	Segmentation	none	RGB Stereo	HV DS-2TD2636	384x288	50	DL	Align	DL	E	Y	Unspec / GTX 2080Ti
[5]	2022	Detection general	own public	Kinect2	FLIR A655c	640x480	50	DL	Align	DL	MDL	N	Unspec / Unspec
[6]	2022	Image de-hazing	public	n/a	n/a			DL	Align	P/DL	L	N	Unspec / Unspec
[7]	2022	Face detection	own	Orbbec Astra	FLIR Lepton 3.5	160x120	8.7	DL	n/a	n/a	n/a	n/a	i7-6700HQ 2.6GHz / Unspec
[8]	2022	Lifeform detection	own	epc635	FLIR Lepton 3.5	160x120	8.7	F	Align	P	E	L	i7-8700 / GTX 1080Ti
[9]	2022	Human detection	own	Kinect1	Seek C. Pro	320x240	8	DL	Align	DL	E	L	Unspec / Unspec
[10]	2022	3D Reconstruction	none	Realsense D455	FLIR Boson 320	320x256	60	OL	Align	P/D	E	YG	Unspec / Unspec
[11]	2021	3D Reconstruction	none	Kinect1	Optris Pi400	382x288	80	SW	Feature	D	E	N	Unspec / Unspec
[12]	2021	Human detection	public	n/a	n/a			DL	Align	DL	EL	L	Unspec / 2x Titan Xp
[13]	2021	Visualization	none	RGB Stereo	Seak	320x240	15	P	Align	D	OL	n/a	Unspec / Unspec
[14]	2021	3D Reconstruction	own	DJI Zenmuse XT2	FLIR XT2	640x512	9	P	Feature	D	OL	N	Unspec / Unspec
[15]	2021	Segmentation	public	n/a	n/a			DL	Align	DL	MF	Y	Unspec / Unspec
[16]	2021	3D Reconstruction	own	Realsense d415	Optris Pi640	640x480	125	DL	Align	D	OL	Y	Unspec / Unspec
[17]	2021	3D Reconstruction	none	Realsense	FLIR A65	640x512	30	ICP	Align	D	OL	YG	Unspec / Unspec
[18]	2020	3D Reconstruction	none	Photogrammetry	FLIR Zenmuse XT	640x480	30	SW	Feature	D	OL	N	Unspec / Unspec
[19]	2020	3D Reconstruction	none	Photogrammetry	FLIR A65	640x512	30	F	Feature	D	OL	N	i7-10870H / RTX 2060
[20]	2020	Tracking	none	Kinect1	FLIR Lepton 2.5	80x60	9	SW	Align	F	E	N	Unspec / Unspec
[21]	2020	Human detection	public	n/a	n/a			DL	Align	DL	EML	L	Unspec / Unspec
[22]	2020	3D Reconstruction	none	Photogrammetry	FLIR E6	160x120	9	SW	Feature	D	OL	N	Unspec / Unspec
[23]	2020	Human detection	none	Kinect1	FLIR A320	320x240	9	F	Align	D	L OL	Y	Unspec / Unspec
[24]	2020	Face detection	own	n/a	n/a			DL	Align	DL	M	n/a	Unspec / Unspec
[25]	2020	Face detection	none	Kinect1	Optris Pi450	382x288	27	P	Align	F	OL	Y	Unspec / Unspec
[26]	2020	PAD	own	RealSense SR300	Seek C. Pro	320x240	15	DL	Align	DL	E M	L	Unspec / Unspec
[27]	2020	PAD	own	n/a	n/a			DL	Align	DL	E M	L	Unspec / Unspec
[28]	2019	Face detection	none	Kinect1	FLIR Lepton 2.5	80x60	9	P	Align	D	OL	Y	Unspec / Unspec
[29]	2019	3D Reconstruction	none	DAVID 3D	n/a			F	Feature	F	OL	n/a	Unspec / Unspec
[30]	2019	PAD	own	RealSense SR300	Seek C. Pro	320 × 240	15	DL	Align	DL	E M	L	i7-4800MQ / GTX 780M
[31]	2019	3D Reconstruction	own	Kinect2	FLIR Boson 640	640x512	30	ICP	Feature	D	OL	N	Unspec / Unspec
[32]	2018	Human detection	none	Realsense R200	FLIR Boson	n/a		P	Align	DL	L	L	i7-9700K / RTX 2060
[33]	2018	ROI Face detection	none	Asus Xtion	Optris Pi640	640x480	32	SW	Align	D	OL	N	Unspec / Titan RTX
[34]	2018	Human detection	own	Kinect2	Optris Pi640	640x480	32	DL	Align	F	L	N	Unspec / GTX 1080Ti
[35]	2018	3D Reconstruction	none	FLIR One	FLIR One	160x120	8.7	F	Align	D	OL	N	iPhone SE 1 (A9) / GT7600
[36]	2018	3D Reconstruction	none	Kinect2	Xenics Gobi 640	640x480	60	ICP	Align	D	OL	YG	Unspec / Unspec
[37]	2017	3D Reconstruction	own	RealSense SR300	FLIR One	160x120	8.7	P	Align	DL	E	N	Unspec / GTX 680M
[38]	2016	Tracking	none	Kinect2	FLIR A655c	640x480	50	F	Align	F	MF	n/a	Unspec / Unspec
[39]	2016	Human detection	own public	Kinect1	AXIS Q1922	640x480	30	F	Align	F	MF	n/a	i7-960 / GTX 400
[40]	2015	3D Reconstruction	none	Kinect1	Optris Pi160	160x120	120	ICP	Align	D	OL	N	i7 1.9GHz / GTX 1060
[41]	2015	Face detection	none	Kinect2	AXIS Q1921	384x288	30	F	Align	D	MF	L	Unspec / Unspec
[42]	2015	Visualization	none	ASUS Xtion Pro	Optris Pi450	382x288	80	F	Feature	D	OL	L	Unspec / Unspec
[43]	2014	3D Reconstruction	none	ASUS Xtion Pro	Optris Pi450	382x288	80	ICP	Feature	D	OL	YG	Unspec / Unspec
[44]	2014	3D Reconstruction	none	Kinect1	Jenoptik IR-TCM	640x480	60	OL	Align	D	OL	n/a	i7-6700K 4GHz / GTX 1080
[45]	2013	3D Reconstruction	none	Kinect1	TM Miricle 307K	640x480	240	OL	Align	D	OL	L	Unspec / Unspec
[46]	2013	Tracking	none	Kinect1/Hokuyo	Heimann HTPA	32x31	9.1	F	None	F	L	Y	i7-6700K 4GHz / GTX 1080
[47]	2013	Re-identification	own	Kinect1	AXIS Q1922	640x480	30	F	Align	F	MF	L	Unspec / Titan Xp
[48]	2012	3D Reconstruction	none	Riegl VZ-400 laser	Optris Pi160	160x120	120	P	Align	D	OL	n/a	Unspec / Unspec
[49]	2012	3D Reconstruction	none	RGB FLIR E60 (SiM)	FLIR E60	320x240	60	F	Feature	D	OL	n/a	Unspec / Unspec
[2]	2011	3D Reconstruction	none	Kinect1	TC384	384 × 288	50	OL	Align	D	OL	n/a	Unspec / Unspec

"T Use": Thermal data application; Display(D), Post-processing or Process/Algorithm(P), Feature(F), Deep Learning(DL)

"Fus.": How the Thermal data was fused with the other modalities; Late(L), Middle(M), Early(E), Overlay/Align(OL), Feature(F)

"RT": Inference speed real-time (≥ 30 FPS); Yes(Y), Yes with GPU(YG), Likely but no data provided(L), No(N)

"Reg": Registration process: Image alignment (Align), Feature matching (Feature)

"Unspec": Unspecified

modalities, such as RGB-DT data, has been recognised to provide a richer and more comprehensive representation of the environment or scene. This has resulted

in achieving a more accurate and robust performance in a wide range of applications, including building mapping [45], person re-identification [47], 3D salient object

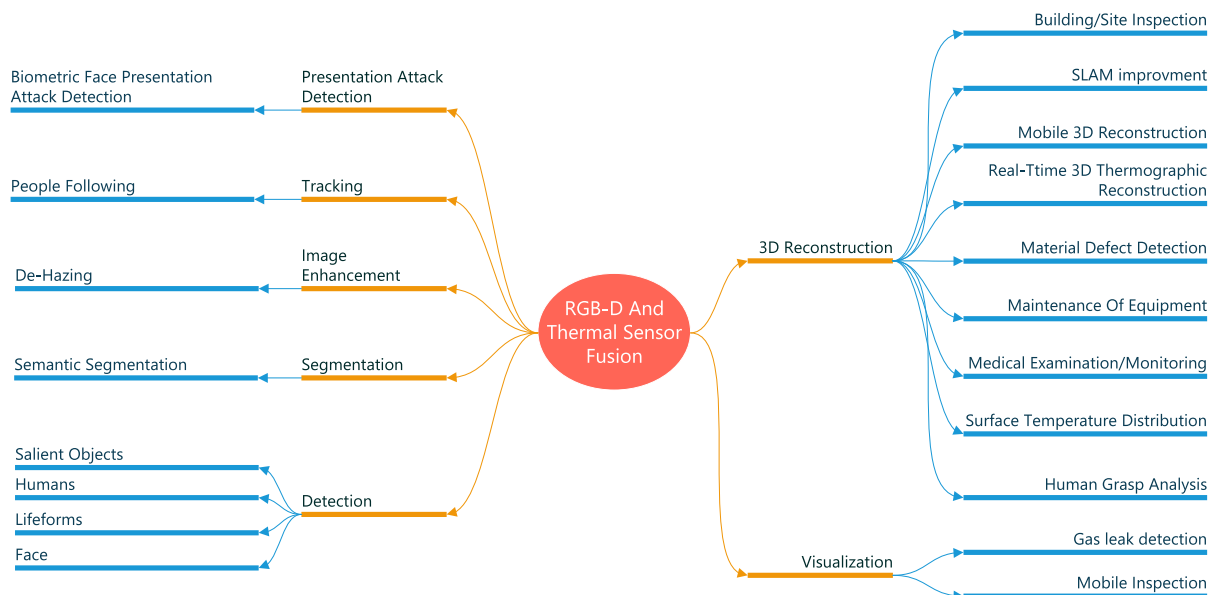


FIGURE 5. Overview of methods and areas of application of the reviewed documents.

detection [5], autonomous driving [6], [50], activity recognition [34], robotics [5], surveillance [12], 3D reconstruction [10], [43], defect detection [29], gas leak detection [42] and many others. For example, in robotics, the combination of RGB-DT data can enable robots to perceive and navigate through complex environments with greater accuracy and efficiency [5] or to interact with humans better by interpreting their emotions [24] and activities [34]. In surveillance, the fusion of RGB-DT data has been shown to improve the detection and recognition of objects, people, and activities in a monitored area [38] under difficult light conditions, and in autonomous driving, the fusion of multi-modal data can provide a more comprehensive understanding of the surrounding environment, enabling safer and more reliable driving [6], [50]. The temperature characteristics of maize under water stress, which serves as an example of multi-modal sensing in agriculture, has been investigated [11], which has the benefit of developing more efficient and sustainable agricultural practices. In the field of industrial maintenance, an Augmented Reality(AR) system, that visualises components and their temperature in real-time has been proposed [16], helping to identify faults and problems. It is also worth noting that in some cases, a single modality can indirectly improve the quality of another modality. This has been demonstrated in [6] by using the thermal data, and the extracted monodepth [51] data from the thermal data, in an algorithm used to dehaze the RGB image so that it can be used for object detection further down in the processing pipeline. Fig. 5 offers a summary of the methods and application areas covered in the reviewed documents. This figure highlights the wide range of approaches and techniques employed in the research papers, as well as the diverse fields where these methods have been implemented.

While early works employed traditional computer vision techniques, the field has evolved alongside advancements in deep learning. Although the fusion of modalities has been shown to outperform single-modality systems, only a limited number of researchers have tackled the topic of heterogeneous sensor fusion involving stereo vision or depth cameras with thermal cameras. This is despite the growing need to meet evolving requirements and develop more robust decision-making systems by integrating features from various sensors. The potential for improved performance in a range of applications highlights the importance of continuing to explore and develop these multi-modal fusion approaches. Figure 2 shows the evolving trends for RGB-D, RGB-T and RGB-DT research by depicting the number of studies published over the past 10 years.

IV. CAMERA CALIBRATION AND REGISTRATION

For successful multimodal environmental sensing using RGB, depth, and thermal data, it is crucial to acquire the data from these modalities in a properly aligned manner. This can pose a challenge since the sensors used for each modality may have varying fields of view (FOV), resolutions, and sensing capabilities. To facilitate data fusion, the system must be calibrated by determining the intrinsic (pin-hole camera model parameter matrix) and extrinsic (estimation of the relative sensor poses) parameters of each camera, which can then be used to align the data. This calibration, based on the pinhole camera model, has been simplified by using a stereo calibration process [52], which can be applied using these and similar modalities. This method has been implemented in numerous studies in different ways. Fig. 6 shows the pattern matching using stereo calibration.




FIGURE 6. Stereo calibration of RGB and Thermal cameras. The left image shows a heated bi-material checkerboard captured by an RGB camera, while the right image presents the same board as seen by a thermal camera. The overlaid lines illustrate the pattern recognition process of the stereo calibration.

A. CALIBRATION BOARDS

The most popular approach for the geometric calibration of thermal cameras used to be a printed chessboard heated by a flood lamp which was comparatively inaccurate and difficult to execute [53] as the temperature difference was fading quickly and the pattern was blurry. To address this a novel geometric mask with high thermal contrast that does not require a flood lamp has been proposed [53] as an alternative calibration pattern. This approach involves cutting a mask out of a thin material and holding it in front of a backdrop with a different level of thermal radiance. Building on this idea, various constructions have been developed in recent years, all based on the same principle.




FIGURE 7. RGB-DT calibration board made of a glass substrate and Alumina panel [5].

The multi-material calibration boards, which are essential for cross-calibrating thermal and visual modalities, with their distinct geometric patterns visible in all calibrated modalities, are used in the calibration process [52]. A checkerboard with 12×9 (30 mm for every square grid) with the pattern printed onto an alumina plate has been used [5] which is then mounted on a glass substrate as illustrated in Fig. 7. The board is heated from the back, while the white reflects the heat, the black conducts it to produce the pattern in the thermal modality. These boards are commercially available. The authors in [54] constructed a board where the

calibration pattern comprises a line-based grid with regularly sized square patterns. The pattern consists of thin copper lines milled onto a printed circuit board (PCB) with a width of 2 mm and a spacing of 40 mm, and it has six/seven intersections along the shorter/longer axis. Compared to conventional calibration patterns, the line-grid pattern is more robust in maintaining high contrast in thermal images due to the good conductivity of the copper lines, which ensures a uniform thermal distribution. Additionally, the proposed pattern has the same geometric relations as the conventional chessboard pattern, allowing for the use of existing algorithms for camera calibration.

However, calibration boards can be constructed simpler as demonstrated in [22] where an 11×11 checkerboard pattern made of cardboard paper and highly reflective metal squares was used. Alternatively, [39] constructed the calibration board using an A3-sized 10mm polystyrene foam board as a backdrop and a board of the same size with cut-out squares as the checkerboard. This is similar to [2] where a solid board was used that had rectangular holes cut out, as shown in Fig. 8(d), whereas [13] used fabric for the black pattern. In addition to using squares, circles can also be used, as demonstrated in [8] where the authors utilised a mask made of 3mm thin Depron[®] material with an asymmetric circle pattern, as shown in Fig. 8(a), while [40] proposed using 3D printed boards in their study, as shown in Fig. 8(g).




FIGURE 8. Calibration boards used in studies [2], [8], and [40] from top to bottom respectively: (a) RGB (b) Thermal (c) Depth (d) RGB (e) Thermal (f) Depth (g) Heating of lower plate (h) Thermal.

A distinct approach was adopted in [20] and [25], where resistors were placed onto the calibration board and heated up electrically, enabling a prolonged calibration process. Similarly, [23] and [32] employed incandescent light bulbs embedded at every other corner of the grid to emit heat.

In [55], a method was proposed for calibrating a UV camera with RGB-D and thermal cameras using a rectangular aluminium plate with evenly distributed circle holes. A heater strip is placed behind the plate to create sufficient contrast for the thermal camera. To ensure that all cameras can be calibrated together, a black box is used, which absorbs most of the light while allowing light to pass through the holes on the aluminium board. The white paper covering the board reflects visible and UV light, which can be detected by the RGB and UV cameras. Once the calibration tool's features are detected, the centre of each circle is marked, and OpenCV's [56] camera calibration function is used to obtain intrinsic and transformation matrices for each camera coordinate system. The proposed method allows for the accurate calibration of multiple cameras, including a UV camera, which can be beneficial in various applications. Reference [10] used this approach to calibrate RGB-D and thermal modalities.

The combination of thermal images and colour images typically involves the use of methods that require complex calculations. However, in [57], a 2-point approach was proposed that outperformed commonly used 8-point and 7-point approaches for equalising the epipolar geometries of different images. The study proposes a method for effectively combining images by determining two points on the epipolar plane. This technique was also employed in [9] for the calibration process. In order to find calibration points in both thermal and optical data, the method used in the study involves several operations. Firstly, the Canny Edge detection method is applied to the thermal image to determine the calibration points. Next, in the optical image, the Hough circle finding method is used to locate the circles containing the calibration points, and the centres of these circles are determined as calibration points. It is important to note that the calibration mechanism design consists of two black circle drawings on a white background with incandescent bulbs at the centre of these circles. This setup allows for the creation of distinguishable common points in both the thermal and optical data, which are essential for the calibration process. Following this, line segments are extracted and plotted on both thermal and optical data. The lengths of these segments are determined by the Euclidean distance, and the slopes of the lines between the points are calculated using the slope formula and are stored for the combining process. The rotation of the thermal image is based on the difference in the calculated slopes of the lines, followed by resizing the thermal image with respect to the line length ratio. The midpoints and distances between them are obtained from thermal and optical images to achieve precise alignment in the same plane. This allows the determination of the position of the thermal image relative to the optical image [9].

In some RGB-D sensors, like the Microsoft Kinect series, the depth stream originates from a time-of-flight camera that also generates an additional IR stream from amplitude information. Since both streams originate from the same sensor, it is referred to as the Depth/IR sensor. The IR stream can

be utilised for calibration purposes, eliminating the need for any 3D elements on a board and can provide supplemental data that may be beneficial in applications for object detection or tracking in low-light conditions. This IR stream senses the 850nm (NWIR) spectral band and does not contain any thermal data. It is important to clarify that this stream should not be mistaken for the stream from a thermal camera, which is based on wavelengths of roughly 8 – 14 μ m (LWIR). An overview of the spectral range is given in Fig. 1.

B. REGISTRATION

RGB-D cameras, including models like Microsoft Kinect (V1, V2, and Azure) and Intel RealSense (D415, D435, etc.), are engineered to simultaneously capture both visual and depth modalities. As a result, they inherently register and output both data types. To align the thermal data, the stereo calibration process can be used to register it against the visual data.

In earlier works, before calibration using geometric patterns was applied, researchers used the Hough Parameter Space to register modalities as demonstrated in [52]. This process involved detecting edges with the Canny edge detector, resulting in binary edge images. These images were then processed by the Hough transform, which extracted all linear image segments. The rotation and translation differences could be calculated using line correspondence analysis [58]. Nonetheless, considering the two modalities as a stereo pair and employing stereo calibration techniques simplifies this process. The algorithm [52] has since been conveniently integrated into various tools such as OpenCV [56], Matlab [59], and other tools and frameworks, facilitating the acquisition of the translation vector, rotation matrix, and distortion coefficients.

The calibration of multi-camera systems, each characterised by a unique field of view (FOV), can be a challenging task, particularly when it involves a variety of modalities and resolutions. RGB and RGB-D cameras typically offer higher resolutions and distinct FOVs compared to thermal cameras. For accurate sensor fusion, optimising the overlap between the RGB, depth, and thermal modalities is crucial. In RGB-D cameras, the RGB component is usually internally pre-adjusted to match the overlapping FOV of the depth data. In sensor fusion processes that are designed for subsequent analysis and necessitate overlapping data from all modalities for real-time processing, it is necessary to modify the RGB-D data through cropping or clipping to match the resolution and FOV of the thermal camera. This requires careful consideration of the FOV of each camera during system design to ensure maximum overlap. When aligning a lower-resolution image with a smaller FOV to a higher-resolution image with a larger FOV, a homography is typically used to transform the lower-resolution image to align with the corresponding part of the higher-resolution image. This approach, which allows for the incorporation of additional information from the

lower-resolution image while preserving the high-resolution data, is employed in the study [40]. In this study, the authors effectively align and fuse data from sensors with different FOVs and resolutions. They further address the challenges of occlusions and significant differences in the FOVs of the cameras, demonstrating the versatility and robustness of this approach in handling complex sensor fusion scenarios. In offline processing scenarios, more intricate techniques can be employed for alignment. For instance, the paper [14] utilises a combination of Scale-Invariant Feature Transform (SIFT) for keypoints computation and matching, Random Sample Consensus (RANSAC) for eliminating geometrically inconsistent matches, and Bundle Block Adjustment (BBA) for optimising camera parameters and producing an initial 3D structure of aligned images. Although these methods are computationally demanding, they offer superior accuracy and robustness, making them ideal for applications where precision is crucial.

While the stereo calibration approach is effective, it encounters a significant challenge from the different FOVs of the modalities, which can result in a parallax effect that varies at different depths. This phenomenon is due to the difference in viewing angles between the cameras, causing objects at various depths to appear at different positions in the different cameras. As a result, using a single homography, a transformation that maps points in one image to corresponding points in another image, only functions effectively on a specific plane. This variation in perspective leads to misalignment in the fused data.

One approach to overcoming this problem is presented in [39]. Firstly, a thermal-visible calibration device is used to establish the correspondences between the points extracted from the thermal and RGB modalities. Using a Microsoft Kinect camera, the depth sensor is already factory registered to the RGB camera; therefore, registration is focused only on the RGB to thermal data. Registration is performed using a weighted sum of multiple homographies. Multiple views of the calibration device scattered throughout the exploratory scene were used to generate homographies relating RGB and thermal modalities. Each homography is calculated using a RANSAC-based method, taking into account the approximate distance to the view of the calibration device represented by the homography. This strategy effectively compensates for parallax at different depths. The rationale behind the approach is that registration based on each homography is only accurate for points on the plane that are spanned by the particular view of the calibration device. Therefore, to register an arbitrary point in the scene, the 8 closest homographies are weighted and then summed up. It was observed that registration accuracy is primarily dependent on 3 factors: the distance in space to the nearest homography, the synchronisation of RGB and thermal cameras, as well as the accuracy of the depth estimate.

In photogrammetry-based 3D reconstruction, as demonstrated in [14], the registration process depends on the identification and matching of keypoints, which is followed by

Bundle Block Adjustment (BBA) [60]. Keypoint computation involves the detection and description of features using the SIFT [61] algorithm. Keypoints are unique locations in the image that correspond to the same real-world object across different images. The matching step entails finding matching keypoints across overlapping images. Subsequently, BBA is used to optimise the camera parameters, both internal and external, for each image, ensuring accurate calculation of ray paths inside and outside the camera for precise 3D reconstruction. These keypoint computations, matching, and BBA algorithms have been extensively studied and integrated into various software packages and frameworks for photogrammetric applications.

1) AUTOMATIC REGISTRATION

A different approach was taken by the authors in [11] by extracting edge images. To register, feature points were detected and matched. Common feature descriptors used for image registration include SIFT [61], SURF [62], and BRISK [63]. However, these methods often involve the use of a Gaussian filter, which can cause the loss of image details. To address this issue, [64] proposed a new feature descriptor called KAZE, which can detect image features in nonlinear scale spaces and obtain more feature points. The KAZE feature descriptor was utilised to register thermal and colour images of maize. The KAZE features and key points were detected from extracted edge images, and their descriptors were built. Feature points were then matched using the nearest neighbour distance ratio strategy, with outliers removed using the M-estimator Sample Consensus (MSC) algorithm, a variant of the RANSAC algorithm. This approach is akin to [29], which is elaborated in more detail in section X-F. The study proposed a feature-based registration method for aligning thermal and RGB-D images using the Shape Constrained SIFT Descriptor (SCSIFT).

A similar auto registration approach was taken in [42], Edge-Based Mutual Information (EMI). However, they encountered issues when utilising the thermal images because of the Automatic Gain Control (AGC) employed in the thermal video stream. This AGC results in a variable colour range, as depicted in Fig. 19, which shows an example from the VDT-2048 dataset. Their proposed method combines mutual information (MI), edge detection, and image separation to achieve image registration with the following steps:

Image filtering: The input images are first filtered using a Gaussian filter to reduce noise. This is done with a 9×9 kernel size and a standard deviation (σ) of 1.85.

Edge detection: A Canny edge detector is applied to both filtered images to generate edge images.

Region separation: After obtaining the edge images, region separation is performed. The primary goal of this step is to constrain the mutual information (MI) optimisation functions to focus on grey values that are in the vicinity of edges. This approach helps to ensure that the MI optimisation

process is more accurate and robust, as it considers only the most relevant information in the image.

When software tools are used for the 3D reconstruction, as in [18] and [22], the registration algorithms are applied by those software packages and are mainly based on a combination of feature detection, feature matching, and bundle block adjustment:

Feature detection: Identifies keypoints or features in each image. These features are typically distinct and easily recognisable patterns, such as corners, edges, or textures. The software employs SIFT (Scale-Invariant Feature Transform) or similar algorithms to extract features from the images.

Feature matching: After detecting features in each image, match corresponding features across multiple overlapping images. The software uses a matching algorithm, such as approximate nearest neighbour matching, to find the best matches between the features detected in different images.

Bundle block adjustment: Once the matching features have been identified, employ a bundle block adjustment technique to optimise the camera positions and orientations, as well as the 3D coordinates of the keypoints. This process involves minimising the reprojection error, which measures the discrepancy between the observed image coordinates and the projected coordinates of the keypoints in 3D space. Bundle block adjustment refines the initial estimates of camera parameters and 3D points to improve the overall accuracy of the reconstructed scene.

By combining these techniques, the software registers the images, ultimately creating a consistent and accurate 3D representation of the surveyed area.

V. THERMAL DATA VISUALISATION

The Automatic Gain Control (AGC) technique is a histogram-based processing method that transforms raw data formats into 8-bit image data. However, this processing results in data compression, leading to a significant loss of information. In the case of 16-bit data, with a possible value range of 0 to 65,535, the resulting image is represented with values in the 0 to 255 interval, further decreasing detail. To address this issue, AGC algorithms are designed to enhance image contrast and brightness, thereby emphasising the contextual details of the scene [65].

Most LWIR cameras produce a grayscale or colour-range image stream with 8-bit per pixel. They typically use an AGC algorithm to generate the 8-bit image with high contrast. The 8-bit data represents gain-controlled values that depend on the temperature of objects in the scene and are more appropriate for human vision. However, the 8-bit representation results in a lower thermal resolution and the algorithm causes colour changes based on minimum and maximum measurements.

VI. HOW AND WHAT TO FUSE

In multimodal sensor fusion, deciding how and what to fuse depends on the specific application, the data modalities involved, and the desired outcome. The fusion of features or decisions can be achieved in many ways, such as

concatenating feature vectors, averaging or weighted averages of data or decisions, weighted voting schemes to combine decisions, or applying machine learning techniques such as neural networks, decision trees, or support vector machines. However, the fast and massive data collection capabilities of the sensors and the representation of the obtained large data in the memory, possibly with different data types, are one of the challenges of real-time sensor fusion [9].

Alongside the fusion of different modalities, another important aspect to consider is the methodology of sensor fusion implementation. Two primary approaches dominate this field: model-based and data-driven techniques. Model-based methods, as explored in study [66], utilise pre-established models to interpret and integrate sensor data. These methods often exhibit robustness and interpretability, but their effectiveness can be constrained by the accuracy of the models they employ. In contrast, data-driven techniques, as outlined in the research [67], learn to merge sensor data directly from the data itself, typically employing machine learning techniques. These methods can potentially achieve superior performance, but they may require substantial data quantities and may be less interpretable.

A. FUSION STAGES

Features of multiple modalities can be fused at different points in a process, and these fusion points are generally categorised into three levels: Data level, Feature level, and Decision level. These levels can also be referred to as low, mid, and high or early, middle, and late fusion. Each level of fusion has its advantages and disadvantages so it is essential to consider the specific context when selecting the fusion point. The three levels can be categorised as:

- 1) Data level(early) fusion: At the data level, the fusion of different modalities involves combining raw data from all modalities to create an integrated dataset, often by concatenating or averaging. This approach is useful when the raw data from different modalities are directly comparable and compatible. For RGB-DT data, often multi-channel images are created by blending and combining the data, primarily for deep learning purposes [9], [12], [21].
- 2) Feature level(middle) fusion: In this approach, features are extracted separately from each modality and then combined before being fed into a classifier or a learning algorithm. Feature-level fusion can involve concatenating the feature vectors or using other methods to merge the extracted features. This method often results in a more compact and informative representation of the data, as the features from each modality are combined after being extracted, retaining information specific to each modality. In manually crafted feature-based approaches, this is a common approach while in deep learning, this method usually enhances accuracy but has higher computational requirements. There are many variations of middle fusion depending on the

processing pipeline. The authors of [38] propose an algorithm optimised for human tracking based on an enhanced Bhattacharyya coefficient, and in [39] features are fused for body segmentation using stacked learning and Random Forest while in [41] features are combined by applying landmark-based energy filters for pain level recognition.

- 3) **Decision level(late) fusion:** At this stage, each modality is processed separately, with features extracted and then classified or analysed independently. The results or decisions from each modality are then combined to produce a final decision or output. This approach is suitable when the modalities are diverse and difficult to compare directly, or when separate classifiers have been optimised for each modality. Decision-level fusion can involve using majority voting [19], [47], weighted voting [8], or other decision-fusion techniques like a Support Vector Machine(SVM) [34].

The performance of a fusion method is highly dependent on the sensing modalities, data, and network architectures being used. This rough categorisation also holds true when applying Deep Neural Networks(DNN), which is discussed in more detail in section XI and in section X for the feature-based approach. The fusion of the modalities however is not limited to a single stage but can be applied at multiple stages in a processing pipeline. Besides the listed fusion methods, it is also worth mentioning that the direct fusion of multiple modalities is not the only way to enhance the quality of data. A single modality can also indirectly enhance the quality of another modality. In [6] for example, the authors used the thermal data, together with the monodepth [51] data extracted from it, to improve the quality of RGB images by applying a dehazing algorithm.

B. FUSION METHODOLOGIES

Sensor data fusion methodologies can be broadly categorised into two main approaches: model-based and data-driven.

- **Model-Based Approaches** These methods rely on pre-defined models to interpret and combine sensor data. They are often robust and interpretable but may be limited by the accuracy of the models they use. Some common techniques under this category include:
 - **Kalman Filters** These are utilised in linear systems characterised by Gaussian noise, offering optimal performance in terms of minimising the mean squared error. As demonstrated in the study by the authors of [46], a Kalman filter, when combined with a probabilistic model of a leg shape, can ensure robust tracking in scenarios such as person-following.
 - **Particle Filters** These are used for non-linear and non-Gaussian systems. They are more flexible than Kalman filters but require more computational resources. In the context of person tracking, the authors of [38] employed a simple particle filter approach, which estimates the target's probability

distribution using a set of weighted particles while study [20] presents an adaptive human tracking method using. The method incorporates adaptive weighting based on velocity and head position, allowing it to handle fast motion, partial occultation, and scale variation. The fusion of depth and thermal data enhances the robustness and accuracy of the tracking process, as demonstrated in various challenging scenarios.

- **Bayesian Networks** These models are utilised in probabilistic modelling to represent the probabilistic connections between a group of variables. They are particularly beneficial when the relationships between the sensors are either known or can be learned. In the domain of Presentation Attack Detection (PAD), the authors of study [68] employed Bayesian Networks to differentiate between a genuine face and a fraudulent attack. Their approach involved designing an attack detector module based on Bayesian principles, with the decision boundary set at a log-likelihood ratio of attack to bona fide equal to 0. This design choice ensures that the classifier operates independently and maximises the confidence score in its classification.
- **Data-Driven Approaches** These methods learn to combine sensor data directly from the data itself, often using machine learning techniques. They can achieve higher performance but require large amounts of data and can be less interpretable. Some common techniques under this category include:
 - **Support Vector Machines (SVMs)** These are powerful supervised learning models that perform well in high-dimensional spaces and can be customised with different Kernel functions for the decision function. However, their effectiveness can be surpassed by more complex models such as CNNs in certain contexts, as shown in [69] for fall detection systems. In the context of activity recognition and emotion classification, SVMs have demonstrated promising results when combined with various types of features. For example, the authors of [34] utilised an SVM model trained with depth and skeleton features in conjunction with thermal sensor data to enhance activity recognition accuracy. Similarly, [24] used an SVM model with both gait Power Spectral Density (PSD) and thermal features, achieving an offline testing accuracy of 70% in emotion classification while the authors of [39] compared their human body segmentation, based on Random Forest, with one using HOG + SVM. Despite the HOG + SVM approach being trained on larger, varied datasets, the study's proposed method significantly outperformed it. Further, the authors of [27], in the context of Presentation Attach Detection (PAD), noted that the

SVM baseline generally performed worse than the other approaches, suggesting that the local, pixel-wise classification approach may not be as effective as the more holistic view provided by CNN models in their evaluation.

- **Decision Trees** These flowchart-like structures are used for decision-making, where each internal node signifies a test on an attribute, each branch represents the outcome of the test, and each leaf node holds a class label. They are valued for their simplicity and interpretability. For instance, the authors of [67] utilised a decision tree-based algorithm in a novel data association approach. This method used polar rays to find correspondences between trifocal camera objects and fused hypothesis, or super-sensor objects. The decision tree gradually eliminated unwanted associations by considering object characteristics such as area, visible façade, dimension ratio, and relative position in different coordinate systems.
- **Random Forest** This ensemble learning method constructs multiple decision trees during training and outputs the class that is the mode of the classes of the individual trees. A practical application of this technique is demonstrated in [19], where a Random Forest was used to predict the conditional probabilities of different class labels based on point descriptors. The Random Forest, an ensemble of decision trees, was trained on randomly sampled subsets of training data. This approach resulted in decorrelated trees that enhanced the generalisation and robustness of the classification. The final point label was determined by majority voting across all decision trees in the Random Forest.
- **Neural Networks and Deep Learning Models** Neural networks excel at discerning complex patterns within high-dimensional data, such as images. Deep learning, a subset of neural networks, utilises multiple hidden layers to automatically learn and extract features from raw data, proving highly effective for RGB-D and Thermal sensor fusion tasks. A more comprehensive discussion on this topic can be found in section XI.

It's important to note that these categories are not mutually exclusive, and sensor fusion systems may use a combination of these approaches. The choice of methodology, similar to the selection of fusion methods, depends on the specific requirements of the task, the available data, and the computational resources.

VII. ROI & OVERLAY

In some applications, thermal data serves as supplementary information for analysis purposes, such as site [14], [22] and building [18], [19], [43] inspections, medical examinations [23], or human thermal comfort assessments [28]. Large-area inspections for sites and buildings are generally

not performed in real-time or with RGB-D sensors. Instead, photogrammetry [60] is employed, either with custom-built processing pipelines as in [14] or established tools like Pix4Dmapper, 3DF Zephyr, Context Capture, PhotoScan, and others as in [18] and [22], to generate point clouds offline. By aligning thermal images, the point clouds are enriched with thermal data for offline analysis. In contrast, [35] used mobile devices and proposed image-based modelling (IBM), a passive mapping technique that uses image datasets with multiple fields of view (FOV) to reconstruct 3D models. This study employed a low-cost thermal camera and two smartphones to capture visible and thermal images. The work established that the proposed method is cost-effective and achieves a temperature precision of 2°C in the 3D thermal models, albeit at a slower pace. Since these approaches are not the primary focus of this study, they are not pursued further but are mentioned for completion as they also represent a type of fusion of these modalities. However, the modalities are not used to enhance a process but merely for post-analysis.

In [4], the authors utilised stereo vision and trained a neural network for disparity estimation to generate depth data. They also applied semantic segmentation, further discussed in section XI-D, using depth and RGB data to define the ROI for extracting thermal data and producing a 3D reconstruction for post-processing. Meanwhile, [23] identified a region of interest (ROI) in the RGB modality also by segmentation but did this by classic methodologies not involving neural networks and extracting the thermal data by applying the ROI to the aligned thermal modality. In certain applications, such as those previously mentioned, the actual temperature values are relevant. However, in other studies like [13], the focus was on using the visual information derived from the thermal image rather than the actual temperature values. In these cases, transformations like stretching the brightness histogram values are applied to enhance the contrast, and additional denoising techniques are used to improve the image quality.

Unlike the previously discussed studies, the authors of [33] configured a system in which ROIs are identified in the RGB modality based on the facial landmark points detected using the CLM Face Tracker [70], and their coordinates are converted to thermal frame coordinates. Key regions of interest include the facial area, ocular and periocular areas, and nose area, and evaluated parameters include position, orientation, green colour component, depth (distance), and temperature. The average values of each variable are computed for each region of interest, and the relative positions and temperatures are computed with respect to the average values computed for the entire face. Finally, each computed value is logged to an individual stamped CSV file for post-experimental processing and analysis.

Similarly, in [25], face detection and extraction of landmark points from RGB images are accomplished by using the Dlib [71] machine learning toolkit based on histogram-oriented-gradient (HOG) features. The authors assumed that the target person does not move significantly between two

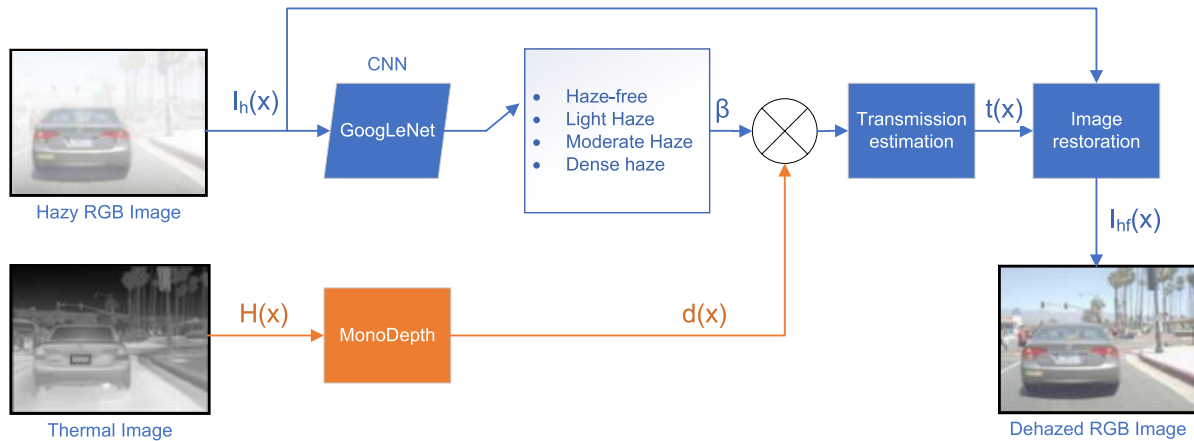


FIGURE 9. Depiction of the image dehazing process using GoogLeNet, a CNN-based classification model, to learn about weather conditions and select an appropriate atmospheric scattering coefficient based on the level of haze. The model performs depth estimation between the object and the camera using Monodepth and the thermal image. With the selected atmospheric scattering coefficient and depth information, a transmission map is estimated and a haze-free image is produced [6].

consecutive frames and limited consecutive detection to the previously identified area to increase the processing speed. The facial ROIs in the thermal image are located using calibrated landmark points. The forehead centre is computed as the middle of the two eyebrow corner points, and the average temperature in the forehead area is taken as the body temperature. The mean temperatures in the nose and cheek areas are used for the measurement of the respective respiration and heartbeat rates through harmonic analysis. The dominant frequency in the temperature signal's spectrum is identified by Fast Fourier Transform (FFT), and then multiplied by 60 to obtain the respiration or heartbeat rate in cycles per minute.

For the purpose of thermal comfort of humans, the authors in [28] used algorithms implemented in OpenCV [56] for facial tracking, but unlike in [25], there was no guarantee that a face faces the camera why the thermal images used for facial skin temperature measurements contain various types of noise, such as false detection of background as faces and interference from high-temperature objects in the environment, which are represented as sudden spikes in measurements. To remove such noise, the median filter was applied before data analysis. Unlike previous studies, that segmented the frontal face into several regions and collected skin temperature from each region, this study used global skin temperature features, including the highest, lowest, first quartile, third quartile, and average temperature measurements of all pixels in the detected facial region. These features provide an overall description of the distribution of skin temperature over a detected face, including both frontal and profile faces.

The authors in [8] adopted a different approach for processing aligned modalities. They applied background substitution and evaluated the size of connected pixel areas from the delta image to determine whether a living being was detected or not. This study fused these regions of thermal and depth data at different levels to determine the optimal result. The study did not find any significant differences in the results based on

the different fusion methods used. The evaluation resulted in an accuracy of 90.1%. However, since the authors used their own data, no comparison with other methods was possible.

Numerous other studies [24], [33], [34], [41] have employed various detection methods to identify ROIs for extracting feature data to be used in decision systems or deep learning algorithms. For example, the average face temperature or the nostril area can be tracked to predict human behaviour. Further details on studies that extract data based on ROIs but process them further are presented in sections X and XI.

VIII. PROCESS SUPPORT

As briefly mentioned in the Fusion Stages VI-A section, there is also an indirect way of using a modality to improve the quality of the data of another modality. In [6] the authors proposed a dehazing network with RGB and thermal depth (DN-RTD). To effectively remove haze, the DN-RTD dehazing network is designed to estimate β , the atmospheric scattering coefficient for the current atmospheric conditions, and $d(x)$, the depth between the camera and the object, using both RGB and thermal images. This network is shown in Fig. 9.

In essence, the dehazing algorithm utilises GoogLeNet, a CNN-based classification model, to categorise captured hazy images $I_h(x)$ into four haze levels: haze-free, light haze, moderate haze, and dense hazy. The model then selects β that corresponds to the classified weather condition. Additionally, the algorithm estimates depth information $d(x)$ from a thermal image $H(x)$ using Monodepth, rather than an RGB image. The transmission map $t(x)$, which expresses the level of atmospheric light transmission, is derived from an Equation using the estimated β and $d(x)$. Finally, the clear image $I_{hf}(x)$ is extracted through the image restoration process. The authors then used two You Only Look Once (YOLO) [72] detectors for both, the thermal and dehazed RGB image, and fused using late fusion. However, the

dehazing process takes 659.1ms to compute why it is not suitable for real-time applications yet.

IX. 3D RECONSTRUCTION

3D thermal mapping reconstruction is a crucial application area for RGB-DT images. Based on the type of 3D reconstruction equipment used, 3D thermal mapping reconstruction methods can be categorised into five groups: RGB-D (ToF or Stereo Vision), Laser Scanning, binocular stereo-structured light encoding, Photogrammetry and Structure from Motion.

The first depth camera employed to aid in 3D thermal mapping reconstruction was the Kinect v1, which has been used in various studies [40], [44], [45]. The authors in [43] developed a handheld 3D thermal mapping system using the Xtion Pro camera, and more recently, Kinect v2 and Intel RealSense [10], [37] have emerged as the most commonly used cameras for 3D thermal mapping reconstruction [17], [31], [36].

The most commonly used technique for large-scale 3D geometrical reconstruction is however Structure from Motion (SfM) [73] which was utilised in [18], [35], [48], and [49]. SfM-based 3D reconstruction approaches typically extract and track robust visual features (e.g. SIFT or SURF) on 2D images captured from different viewpoints and only work well under good illumination conditions (e.g. during daytime). Feature extraction and matching, which involves the detection of SIFT features, SURF features, ORB features, and AKAZE features, is a crucial part of the SFM algorithm. However, it only produces sparse 3D point clouds, and the generated 3D models lack absolute scale information, which is not ideal for thermal diagnosis applications. To overcome these limitations, RGB-D-based 3D modelling approaches nowadays utilise depth sensors to acquire depth data of 3D objects/scenes from different viewpoints and apply 3D point cloud registration techniques, such as the iterative closest point algorithm, to align the current view with the global model [74], [75], [76], [77]. Besides the better quality, it is also worth noting that binocular stereo-structured light, as used by [78] and [79], or time-of-flight depth sensors, can acquire 3D geometrical information in darkness.

Recently, the authors in [78] introduced a fast and reliable 3D thermographic reconstruction method using stereo vision. The system features adjustable measurement fields and distances, based on the chosen optics for the cameras and projector. It can reach frame rates of up to 12.5 kHz for VIS cameras and 1 kHz for the LWIR camera at full resolution. By lowering the resolution, even higher frame rates can be attained.

Meanwhile, researchers in [48] utilised terrestrial laser scanners (TLS) to acquire dense 3D point clouds, and temperature information obtained by an infrared camera is mapped onto 3D surfaces. To improve the mobility of 3D thermal imaging systems, a multi-sensor system consisting of a thermal camera and a depth sensor was built to generate 3D

models with both visual and temperature information to be used for building energy efficiency monitoring [45].

Another method proposed was a thermal-guided 3D point cloud registration method (T-ICP) that improves the robustness and accuracy of 3D thermal reconstruction by integrating complementary information captured by thermal and depth sensors [36], but the method requires high computing resources to calculate several feature points. A set of experiments were performed to analyse how the key factors, such as sensing distance, specularity of the target, and scanning speed, affect the performance of high-fidelity 3D thermographic reconstruction. The authors in [10] implemented a similar idea but the localisation method combines the ORB-SLAM2 with the thermal direct method, and the entire system runs on the Robot Operation System (ROS).

Based on the Thermal-guided Iterative Closest Point (T-ICP) algorithm presented in [36], the authors of [17] developed a multi-sensor system that consists of a thermal camera, an RGB-D camera, and a digital projector. This method utilises an effective coarse-to-fine approach to enhance the robustness of pose estimation, allowing it to handle significant camera motion during large-scale thermal scanning processes. This system enables multimodal data acquisition, real-time 3D thermographic reconstruction, and spatial augmented reality through projection.

A new dataset consisting of objects and their corresponding thermal imprints resulting from grasping was proposed in [31]. To generate a coherent contact map of an object, the object is placed on a turntable which rotates as RGB-D and thermal images are captured from multiple viewpoints. The thermal images are texture-mapped onto the object's 3D mesh using a data processing technique. The steps involved in this process include extracting corresponding turntable angle and RGB, depth, and thermal images at nine locations where the turntable pauses, converting the depth maps to point clouds, estimating the turntable plane and segmenting the object using white colour segmentation, estimating the full 6D pose of the object in the nine segmented point clouds using the Iterative Closest Point (ICP) algorithm implemented in PCL, obtaining a least squares estimate of the 3D circle described by the moving object using the object origins in the nine views, and interpolating the object poses for views that are unsuitable for the ICP step. Finally, the 3D mesh along with the nine pose estimates and thermal images are input into a colourmap optimisation algorithm, which is implemented in Open3D [80], to minimise the photometric texture projection error and generate a mesh that is coherently textured with contact maps. Examples of the resulting contact maps are shown in Fig. 10.

X. MANUAL DESCRIPTOR-BASED METHODS

In contrast to deep learning methods where the feature extraction is done by the Neural Network(NN), like by convolutional layers in Convolutional Neural Networks (CNN), handcrafted descriptors are manually designed features extracted from the input data. These include histogram of




FIGURE 10. Examples from ContactDB, constructed from multiple 2D thermal images of hand-object contact resulting from human grasps [31].

oriented gradients (HOG), Histogram of Optical Flow (HOF), scale-invariant feature transform (SIFT), local binary patterns (LBP), histograms of thermal intensities, oriented gradients (HIOG), and others. These techniques are considered to be the traditional methods that have been mostly replaced by CNNs, and more recently by transformer networks, in modern detection pipelines [6]. The authors in [14] noted that the SIFT algorithm is robust and accurate for matching features in RGB images, but it only computes low-level features and cannot recognise high-level representations.

A. VISUAL MODALITY (V)

The authors in [39] employed a combination of HOG, HOF, and HIOG to extract features. HOGs and HOFs are used to extract features from the RGB and depth data, while HIOG is used for thermal data. These features are then combined into a single feature vector and fed into a Random Forest classifier for body segmentation. The classic implementation of HOG was used for the RGB data but with a lower descriptor dimension than the original by not overlapping the HOG blocks. No gamma correction was used for the gradient computations and the Sobel kernel was applied. This means that for each pixel, the gradient orientation is determined by looking at the dominant colour channel (red, green, or blue) of that pixel, and then quantising it into a histogram over each HOG-cell [39].

HOF is a feature extraction method used to obtain motion information from an image. It works by computing dense optical flow and describing the distribution of the resultant vectors. The optical flow vectors are computed using the luminaries information of image pairs with the Gunnar Farneback's [81] algorithm. In [39], the authors used the implementation in OpenCV [56]. The resulting motion vectors are then masked and quantised to produce weighted votes for local motion based on their magnitude, taking into account only those motion vectors that fall inside the colour grids. The votes are locally accumulated into a v-bin histogram over each grid cell according to the signed (0° – 360°) vector orientations. Unlike HOG, HOF uses signed optical flow as the orientation information provides more discriminative power [39].

Similarly [38] also used histogram-based descriptors but to process the colour modality, the RGB image is converted to a normalised colour space denoted as rgb , where $r = R/(R+G+B)$, $g = G/(R+G+B)$, and $b = B/(R+G+B)$. The colour

normalisation approach is used to eliminate the illumination information in order to achieve robustness against lighting variations. Due to the fact that two components are adequate for describing the normalised colour space, with $r+g+b = 1$, a 2D histogram H_C is computed using the pair (r, g) .

B. DEPTH MODALITY (D)

For depth, the authors in [39] used Histogram of Oriented Depth Normals (HON) to describe points in a point cloud. The depth modality contains a depth-dense map that represents a planar image of pixels measuring depth values in millimetres. The intrinsic parameters of the depth sensor can be used to obtain the actual coordinates from this depth representation, which can be seen as a 3D point cloud structure. This new representation allows measuring actual Euclidean distances that reflect the real world. After converting the depth modality, the surface normals for each point in the point cloud are computed, and their distribution of angles is summarised in an α -bin histogram. Then a histogram describing the distribution of the normal vectors' orientations is built. A normal vector is expressed in spherical coordinates using three parameters: the radius, the inclination θ , and the azimuth φ . In this case, the radius is a constant value, so this parameter can be omitted.

For θ and φ the calculation of the cartesian-to-spherical coordinate transformation is:

$$\theta = \arctan\left(\frac{n_z}{n_y}\right), \quad \varphi = \arccos\frac{\sqrt{n_y^2 + n_z^2}}{n_x} \quad (1)$$

Thus, a 3D normal vector can be represented by a pair of angles (θ, φ) , and the depth description comprises two histograms for δ_θ -bin and δ_φ -bin, which are L1-normalised and combined. These histograms describe the angular distributions of the surface normals on the body.

Similarly, [38] used an approach where a 3D normal vector is computed for each data point by fitting a 3D plane to a pre-defined local neighbourhood. Using the corresponding polar angle θ and azimuthal angle φ information, a 2D histogram H_D is computed.

The authors in [46] used a Leg Detection method proposed in an earlier work [82] which utilises a probabilistic leg pattern. The leg model is implemented as a sequence of maximum, minimum, maximum, minimum, and maximum values based on the laser readings, as in [83]. Various measures are defined, such as the distance between the legs and the distance between the legs and background based on these five points. Besides the laser, the depth data of an RGB-D sensor is used to detect a particular emergency vest of a person. After detecting the corners, the Lucas-Kanade method is used to calculate the optical flow. The optical flow is computed for each corner, and in each frame, the centroid of the corners is then extracted, providing an estimation of the target's position.

C. THERMAL MODALITY (T)

Reference [39] used the Histogram of Thermal Intensities and Oriented Gradients (HIOG) descriptor derived from the thermal cue. This descriptor is a concatenation of two histograms. The first histogram provides a summary of thermal intensities, which are distributed over the range [0, 255]. The second histogram represents the orientations of thermal gradients. These gradients are calculated by convolving a first derivative kernel in both directions and then binned into a histogram, with their magnitude serving as a weighting factor. The two histograms are L1-normalised and concatenated. For the intensities, α_i bins, and for the gradient orientations α_g bins are used.

In a similar way, but solely relying on summarising the distribution of thermal intensities, the authors in [38] proposed a method to generate a one-dimensional histogram for the thermal modality by directly utilising the intensity values of the thermal image.

Meanwhile, [46] proposed a method to generate a 32-dimensional vector from a thermal image, where each element of the vector corresponds to the estimated probability of a person being present in a particular column of the image. This approach was chosen as the used thermal sensor had a resolution of 32×31 pixels. The computation of the vector involves three steps: firstly, a likelihood of a pixel corresponding to a person is assigned based on the assumption that the temperature of a person follows a normal distribution with mean and standard deviation values of 36 and 2, respectively, which are determined from several thermal images of people. Secondly, the likelihood matrix is smoothed by convolving it with a Gaussian kernel of a width of five pixels. Finally, the maximum value in each column of the likelihood matrix is used to determine the corresponding element of the output vector. The computation is based on established techniques such as the Lucas-Kanade optical flow method and Gaussian smoothing.

D. GENERAL FEATURE EXTRACTION

The authors in [41] studied the detection of pain levels in faces and used the same feature extraction for all three modalities as a descriptor that considers both, spatial and temporal domains. This is needed to capture the spatiotemporal phenomena of changes due to pain in a facial expression. To achieve this, a steerable separable spatiotemporal filter has been selected, which utilises the second derivative of a Gaussian filter and their corresponding Hilbert transforms to measure the orientation and level of energy in the 3D space of x , y , and t . The filter provides information on the spatial texture of the face through its spatial responses and the dynamic of the features such as velocity through its temporal responses. The filter is applied independently to all three modalities, and for each pixel, the energy is calculated and normalised to improve comparability in different facial expressions. Finally, to improve localisation, the normalised energy is weighted using histograms of directions, and

pixel-based energies are combined into region-based energies. For each pixel, the energy is calculated by:

$$E(x, y, t, \theta, \gamma) = [G_2(\theta, \gamma) * I(x, y, t)]^2 \quad (2)$$

The convolution operator $*$ is used to denote the operation in which (x, y, t) represents the pixel value located at position (x, y) of the t th frame (temporal domain) in the aligned video sequence I . $E(x, y, t, \theta, \gamma)$ represents the energy released by this pixel in the direction θ and scale γ . To ensure that the obtained energy measure is comparable across different facial expressions, normalisation is performed using:

$$\hat{E}(x, y, t, \theta, \gamma) = \frac{E(x, y, t, \theta, \gamma)}{\sum E(x, y, t, \theta_i, \gamma) + \epsilon}, \quad (3)$$

After considering all directions θ_i , where i considers all directions and ϵ is a small bias to prevent numerical instability when the overall estimated energy is too small, the normalised energy is weighted to improve localisation using the method proposed in [84]:

$$\dot{E}(x, y, t, \theta, \gamma) = \hat{E}(x, y, t, \theta, \gamma) \cdot z(x, y, t, \theta) \quad (4)$$

where:

$$z(x, y, t, \theta) = \begin{cases} 1 & \sum_{\gamma_i} \hat{E}(x, y, t, \theta, \gamma_i) > Z_\theta \\ 0 & \text{Otherwise} \end{cases} \quad (5)$$

The resulting weighted normalised energy obtained in equation 4 assigns a value to each pixel based on the level of energy released by that pixel, corresponding to the chosen directions of $\theta = 0, 90, 180$, and 270 . To combine these pixel-based energies into region-based energies, the authors follow study [85], by using their histograms of directions:

$$H_{R_i}(t, \theta_i, \gamma) = \sum_{R_i} \dot{E}(x, y, t, \theta_i, \gamma), \quad (6)$$

The histogram H_{R_i} represents the directions of the i -th region of the face, where $i = 1, 2$ or 3 , and is used to combine regional histograms that are directly related to each other during the pain process. This is necessary because the muscles return to their original locations after being moved due to pain. In accordance with [85], two directions of up-down (UD) and left-right (LR) are used to combine these histograms. The directional histograms are obtained for each modality of RGB, depth, and thermal, and are subsequently used separately to determine the level of pain.

E. SEGMENTATION

This section explores various methods of basic segmentation using multiple modalities from the reviewed studies.

In [23] the authors isolate the abdominal region of newborns. The regions of interest are extracted using depth information, followed by the refinement of the human body area using the colour information to remove the background and isolate the individual. First, a dynamic depth threshold is applied to separate the body from the flat bedding surface. The distance threshold is automatically determined based on the histogram of the depth map and the first significant

observed cluster according to the imaging conditions, which involve imaging the subject from above. The second step involves utilising a skin colour model that is encoded in the YCbCr space to improve the segmentation of exposed body regions from other objects in the field of view, such as probes, tubes, or clothes. The method includes multiple steps using Canny edge detection and polygonal approximation algorithms. Then, an additional refinement step is introduced in the form of a skeleton recognition method based on the depth image. This method utilises depth data to recognise various skeleton points and describe different parts of the human body.

The authors in [37] proposed a multimodal egocentric SLAM (Simultaneous Localisation and Mapping) system based on ORB-SLAM [86] which faces a significant challenge in segmenting the input frame into left-hand, right-hand, object in interaction, and static environment classes. This segmentation is crucial for two reasons: first, removing dynamic points from the input frame is essential for successful SLAM operation, and second, these labels provide the necessary structure for proper scene understanding. The semantic segmentation algorithm the authors proposed is based on priors for the hands, including their colour model, temperature, and shape. Hand location is also a prior for the object in interaction. The segmentation is performed in two steps, first segmenting the left and right hands and then the object in interaction. The right and left hands are distinguished using the prior that the right hand is at the right side of the image frame and the left hand is at the left side. CRF-based image segmentation is used to segment the hands, defining an energy minimisation problem:

$$\min_{\alpha_i^t} \sum_i U(\alpha_i^t, \mathbf{y}_i^t) + \sum_i \sum_{j \in \mathcal{N}(i)} V(\mathbf{y}_i^t, \mathbf{y}_j^t) 1[\alpha_i^t \neq \alpha_j^t] \quad (7)$$

In this equation, α_i^t represents a binary value of 1 if pixel i is classified as part of the hand at time t , and 0 otherwise. The neighbouring set of i is represented by $\mathcal{N}(i)$, and the indicator function is represented by $1(\cdot)$. The concatenated vector of $\mathbf{z}, \mathbf{c}, d, \tau$ is represented by \mathbf{y} . The unary energy function, $U(\alpha_i^t, \mathbf{y}_i^t)$, expresses the likelihood of pixel i being part of the hand, and it is a weighted combination of the probabilities of temperature (T), colour (C), hand-detector outputs (S), and history over time (H).

$$\begin{aligned} U(\alpha_i^t, \mathbf{y}_i^t) = & w_T U^T(\alpha_i^t, \mathbf{y}_i^t) + w_C U^C(\alpha_i^t, \mathbf{y}_i^t) \\ & + w_S U^S(\alpha_i^t, \mathbf{y}_i^t) \\ & + w_H \sum_i U(\alpha_i^{t-1}, \mathbf{y}_i^{t-1}) e^{-\Delta(\mathbf{y}_i^t, \mathbf{y}_i^{t-1})} \end{aligned} \quad (8)$$

where $\Delta(\cdot, \cdot)$ calculates the geodesic distance over RGB-thermal space between two voxels. $V(\cdot, \cdot)$ is a binary consistency term that is defined over neighbouring pixels and takes

the following form:

$$V(\mathbf{y}_i^t, \mathbf{y}_j^t) = \exp\left(-\frac{|\mathbf{y}_i^t - \mathbf{y}_j^t|_2}{\gamma}\right) \quad (9)$$

where $\gamma = \frac{1}{N} \sum_i \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} |\mathbf{y}_i^t - \mathbf{y}_j^t|_2$, and N is the total number of pixels.

They defined each component of the unary energy as:

$$\begin{aligned} U^T(\alpha_i^t, \mathbf{y}_i^t) &= \tau_i^t 1[\alpha_i^t = 1] + (1 - \tau_i^t) 1[\alpha_i^t = 0] \\ U^C(\alpha_i^t, \mathbf{y}_i^t) &= p(\mathbf{c}_i^t | \alpha_i^t) \\ U^S(\alpha_i^t, \mathbf{z}_i^t) &= \sum_{k \in \mathcal{H}} p_k e^{-\Delta(\mathbf{y}_i^t, \mathbf{y}_k)} 1[\alpha_i^t = 1] \end{aligned} \quad (10)$$

where the RGB-colour model $p(\mathbf{c}_i^t | \alpha_i^t)$ is represented using a Gaussian Mixture Model (GMM) with five components and is learned separately for the hand and static scene from training data. \mathcal{H} is a collection of hand detections, where each detection is represented by a centroid \mathbf{c}_k and a detection likelihood $p_k \cdot \mathbf{y}_k$ which includes colour, position, depth, and temperature of the centroid of the detected hand.

All components of this energy function can be computed using bi-linear filters in log-linear time and minimised using the min-cut/max-flow framework as explained in [87]. The authors used the open-source code released by the authors of [87] and the original paper provides further details.

After segmenting the hands, the process continues to segment the remaining part of the image into static and dynamic object components. The same energy minimisation framework is used, with an additional motion prior and the removal of the colour prior. The motion prior accounts for the disparity between the motion of the object in interaction and the camera motion, and is defined as:

$$U^M(\alpha_i^t, \mathbf{y}_i^t) = \rho\left(\left|\mathbf{z}_i^t - \mathbf{z}_{\pi(\mathbf{R}^t \mathbf{x}_i^{t-1})}\right|\right) \quad (11)$$

In the above equation, ρ denotes the Huber function, π is the pinhole projection, \mathbf{R} and \mathbf{t} are the estimated camera pose, and \mathbf{X}_i represents the 3D position of the i^{th} point in homogeneous coordinates. Here, α_i is a binary variable, which is equal to 1 if the i^{th} pixel belongs to the object in interaction and 0 otherwise. The tradeoff parameters $\omega_T, \omega_S, \omega_H, \omega_M$ were learned by cross validation.

F. FUSION & EVALUATION

In this section, we examine the various approaches adopted for fusing the extracted features and the evaluation processes of the reviewed studies.

While [39] applied a background removal algorithm based on V and D for segmentation purposes to define the Ground Truth (GT), the authors in [20] used background removal based only on D for their body segmentation in a fixed camera set-up. The paper proposes a modality fusion method for head tracking using D and T information for fall detection. This approach uses a particle filter to estimate the head position based on both, the D and T data. A Silhouette is constructed

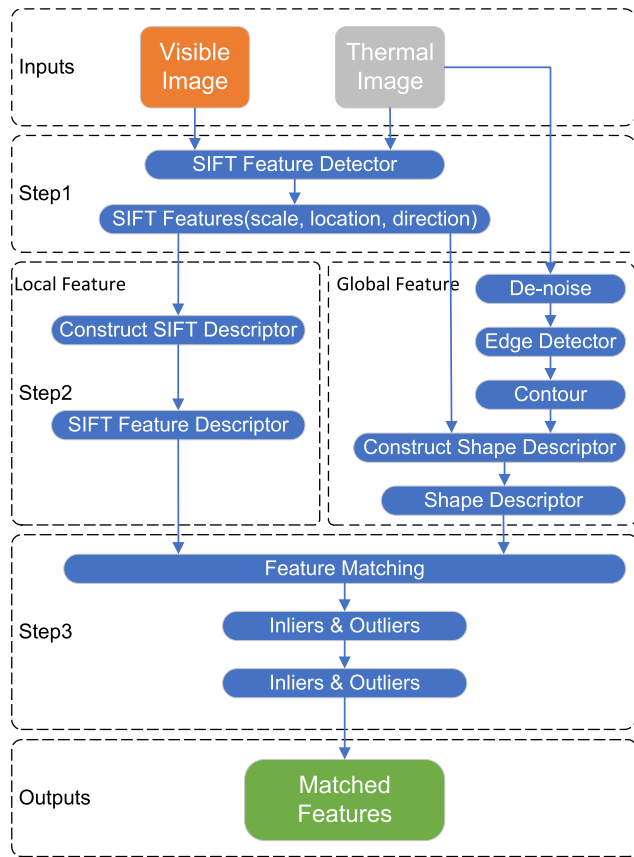


FIGURE 11. Structure of proposed feature matching algorithm [29].

based on D and basic body shape assumptions, while the thermal data is used to distinguish the head from the background. The fusion is performed by combining and weighting D and T based on their reliability. The authors evaluated 4 different models and concluded that the D and T data was improving the results. However, the method was limited to 8 FPS and the authors further concluded to use Deep Learning models for future refinements of this application. In [39], the authors evaluated uni-modal classification and a multi-modal fusion based on a Random Forest classifier to achieve a human body segmentation with the extracted features discussed in section X-B, X-A, and X-C.

Reference [38] proposed a simple person-tracking algorithm that combines the parameters of the three modalities in a way that gives less weight to modalities where camouflaging occurs. The tracker's ability to resist significant radial motions was demonstrated using the Jaccard index. The target model was modelled using a single histogram for each data source and a histogram of 3D normals was used as the depth descriptor. This did not significantly improve the tracker's accuracy compared to the same approach without depth descriptors but the authors argued it could improve its robustness in more complicated sequences.

The People tracking system proposed in [46] used four modalities and is based on a laser sensor, a thermal sensor and

an RGB-D camera in a mobile setup. These sensors supply input to three detection units: Leg detection, Vest Detection, and Thermal Detection, which have been discussed previously. Once their individual likelihoods are calculated, the final likelihood is calculated using coefficients to weigh the three likelihoods. The inclusion of these coefficients enables assigning more importance to one of the information sources if desired, and the authors determined these values during the evaluation process.

The researchers in [29] propose a feature-based registration method to register thermal and RGB-D images using the Shape Constrained SIFT Descriptor (SCSIFT). The registration process involves three steps: feature detection, feature description, and feature matching. In the first step, SIFT detector is applied to extract SIFT features from both visible and thermal images. In the second step, the proposed SCSIFT descriptor is constructed by combining the traditional SIFT descriptor with the shape descriptor extracted from the thermal image. In the third step, feature matching is performed by calculating the Euclidean distance between each shape descriptor vector and each SIFT descriptor vector, followed by normalisation and RANSAC to eliminate outlier matches. A detailed explanation of each this step can be found in the original study, Fig. 11 depicts the proposed algorithm.

Shape Feature Description: Global descriptors to support local descriptors are added in multi-modality image feature matching. As the thermal image is noisy, anisotropic diffusion is applied for effective smoothing before edge extraction. The canny edge operator is used to extract edges, but contour-based methods alone are insufficient for correct feature matching. A circular template is generated around each feature point, with evenly plotted bins for edge point fitting. To describe the global position of the feature point and construct the shape descriptor, a spiral of Theodorus is applied to build the weighting function. The weighting of close region pixels is enhanced while the weighting of far region pixels is suppressed. The proposed SCSIFT descriptor is constructed from the entire image and adds a global shape constraint to the traditional SIFT descriptor which uses only local neighbourhood information.

Feature Matching Scenario Based on SCSIFT Descriptor: Normalisation is necessary before implementing RANSAC because the global shape descriptor vectors and local SIFT descriptor vectors are the statistical analysis of different information. For each feature i in the source image f_S^i with the descriptor denoted as d_S^i , the Euclidean distance of the global descriptor $d_{S(G)}^i$ and local descriptor $d_{S(L)}^i$ to all global and local features descriptors $D_{\text{ref}(G)}^{\text{all}}$ and $D_{\text{ref}(L)}^{\text{all}}$ in the reference image are calculated respectively, denoted as set E_G^i and set E_L^i shown in Eq. 12 and 13.

$$E_G^i = \sqrt{\left(D_{\text{ref}(G)}^{\text{all}} - d_{S(G)}^i\right)^2} \quad (12)$$

$$E_L^i = \sqrt{\left(D_{\text{ref}(L)}^{\text{all}} - d_{S(L)}^i\right)^2} \quad (13)$$

The ratio of the maximum value of sets E_G^i and E_L^i in Eq. 14 represents the scaling factor S^i . As the process of calculating Euclidean distances for descriptors is already an integral part of feature matching, this normalisation step does not add to the computational complexity.

$$S^i = \frac{\max(E_L^i)}{\max(E_G^i)} \quad (14)$$

And the normalisation process is done by the Eq. 15

$$E^i = E_G^i \cup \frac{E_L^i}{S^i} \quad (15)$$

The unified distance set E^i is used to determine the most likely match for f_S^i to all features in the reference image, with the minimum value in E^i indicating the best match. Based on the maximum global and local distances, an appropriate scaling value is calculated for each feature to improve the matching accuracy. RANSAC is then employed to eliminate any outlier matches and refine the image transformation.

XI. DEEP LEARNING-BASED METHODS

Deep learning-based approaches for multi-modal sensor fusion have gained increasing attention due to their ability to learn complex relationships between the different modalities and effectively fuse the information from multiple sensors. However, despite that images from multiple modalities can be beneficial in highlighting salient regions and providing more comprehensive information, they can also introduce interference between the different modalities [5]. The application of Deep Neural Networks (DNN) can be categorised into semantic segmentation and object detection. In contrast to semantic segmentation, where multi-modal features are fused at various stages within the Fully Convolutional Network (FCN), object detection involves a wider range of network architectures and fusion variants. This diversity allows for greater flexibility and adaptability in addressing specific challenges related to object detection tasks [50].

Convolutional Neural Networks (CNN) have long been the dominant architecture for image processing tasks. However, recent developments in applying transformer networks [88] to Computer Vision (CV), known as Vision Transformers (ViT) [89], have demonstrated high performance in segmentation, recognition, and detection tasks. These advances indicate the growing potential for transformer-based approaches in the field of CV.

Within the realm of multimodal object detection and segmentation, a considerable amount of research is directed towards autonomous driving applications, where the fusion of LiDAR point cloud data and RGB camera data is crucial. However, this paper focuses on RGB-D sensor data, which incorporates pre-aligned depth data or can be aligned using stereo calibration techniques. As a result, the methods for aligning point clouds with RGB data will not be discussed in this paper. However, once the RGB-D data is transformed into a point cloud, the succeeding methods for object recognition and detection can still be utilised. Although, papers using

RGB stereo vision applying disparity prediction, briefly discussed in Section XI-A, are included as this method generates similar data as RGB-D sensors.

A. DISPARITY PREDICTION

To obtain depth data from a stereo image the disparity can either be computed, like in OpenCV which implements the block matching algorithm for calculating disparity with stereo calibration, or by training Neural Networks like AANet (Atrous Adaptive Network) [90]. AANet is a deep learning approach for stereo matching that can provide more accurate results than traditional methods such as block matching or simple disparity calculation. It can handle occlusions, textureless regions, and large baselines better than traditional methods. AANet is also more robust to lighting changes and can handle different camera configurations. Additionally, it can learn from large amounts of data, making it more adaptable to a wide range of scenarios. Overall, AANet provides a more flexible and accurate solution for stereo matching compared to traditional methods. Similar techniques can be applied to monocular vision to produce depth data as applied in [6] using monodepth [51].

Reference [4] utilised AANet for disparity prediction of chicken images for feather damage analysis. The network extracts the down-sampled feature pyramid and constructs multi-scale 3D cost volumes [91]. The cost volumes are then aggregated with six stacked Adaptive Aggregation Modules (AA Modules). Each AA Module consists of three Intra-Scale Aggregation (ISA) and a Cross-Scale Aggregation (CSA). The multi-scale disparity predictions are regressed by the soft argmin mechanism and hierarchically up-sampled and refined to the original resolution. The pre-trained AANet model for the Scene Flow dataset was used for direct inference on the dataset. The dataset was augmented by random colour augmentations and vertical flipping. The initial learning rate of the pre-trained AANet model was set to 0.001 and decreased by half at 400th, 600th, 800th and 900th epochs. Adam was used to optimise the parameters of the network to minimise the average loss of the model on the training data. The disparity range was from 0 to 192 pixels.

B. FUSION METHODS (DL)

1) DATA LEVEL FUSION (DL)

In study [50], two primary advantages of early fusion were identified. Firstly, the network learns the joint features of multiple modalities at an early stage, allowing it to fully utilise the information present in the raw data. Secondly, early fusion has lower computational demands and requires less memory, as it processes multiple sensing modalities together. However, these benefits come at the cost of reduced model flexibility. For instance, when an input is replaced with a new sensing modality or the input channels are extended, the early fused network must be completely retrained which was also noted by authors in [12]. This study created a two-channel image out of the two modalities for training a Faster R-CNN

architecture with a ResNet-50 backbone. On the same dataset, IPHD, the authors in [21] created a three-channel image by concatenating two duplicated thermal images and one depth image similar to a three-channel RGB image to make use of the ImageNet pre-trained weights for initialisation. They achieved similar results on the AP_{50} metric compared to [12]. In [9] the authors fused the data into a 3-channel image, combining the grey channel of each modality per channel and applying different weights to the thermal and optical channels by applying the `addWeighted` method of OpenCV [56]. This process involved a manual search to determine the optimal weighting scheme for the given data. To perform these unifications, the intensity values of the pixels from both images are multiplied by the desired weights and then added to compute the pixel intensity values of the resulting image. In this method, the thermal image pixel values are added to the optical image pixel values using weights ranging from 0.1 to 0.9 with a step size of 0.1.

The way a colour scheme is applied to translate the thermal to visual data plays a significant role, which is discussed in section XIII-B and XI-E. Further, the method is sensitive to spatial-temporal data misalignment among sensors, which can be caused by calibration errors, different sampling rates, or sensor defects. This sensitivity further highlights the limitations of early fusion in certain scenarios [50].

2) FEATURE LEVEL FUSION (DL)

Middle fusion can be seen as a compromise between early and late fusion. By combining feature representations from different sensing modalities at intermediate layers, this approach allows the network to learn cross-modal information with varying feature representations and depths. Authors in [50] argued that middle fusion is quite flexible, but that finding the “optimal” way to fuse intermediate layers for a specific network architecture can be challenging. This difficulty arises from the intricate interactions between features and the potentially vast array of possible fusion configurations. Nevertheless, merging feature representations from different sensing modalities at intermediate layers enables the network to learn cross-modalities with varying feature representations and depths. This fusion can occur at a specific layer only once or can be hierarchically fused, as depicted in Fig. 12, such as through deep fusion or ‘short-cut fusion’. Based on [50], this figure further illustrates the intricate nature of middle fusion and the variety of approaches that can be taken to combine information from different sensing modalities.

Shortcut Fusion, as discussed in detail in the paper [92], is a technique employed in deep neural networks that involves creating additional pathways within the network. This allows early layers to directly contribute to later layers, aiming to combine the advantages of both early and late fusion. By utilising low-level feature fusion and high-level decision fusion, shortcut fusion has the potential to enhance accuracy. It preserves detailed information from earlier stages and incorporates it into the final decision stages. However, it’s

worth noting that this method may increase the complexity of the network, potentially requiring additional computational resources.

Deep Fusion, on the other hand, operates hierarchically at multiple levels within the network. This is beneficial in capturing intricate interactions between different modalities at an intermediate stage. The authors of [93] emphasise the importance of deep feature alignment for multi-modal object detection and how Deep Fusion improves detection accuracy and robustness against input corruptions and out-of-distribution data. Deep fusion allows for a more comprehensive understanding of the data, as it integrates information from various stages of processing. Furthermore, deep middle fusion is often favoured over late fusion due to its superior feature integration capabilities. By combining features at a deeper level, it can lead to a more robust and reliable model, thereby enhancing the overall performance.

The complexity of middle fusion is further illustrated by the variety of techniques that can be implemented, as identified by the authors in [94]. These techniques include:

- **Additive Fusion:** Individual networks or branches process each sensing modality up to a designated intermediate layer. The feature maps from these intermediate layers are either added element-wise or concatenated. The resulting feature maps are further processed by the network to produce the final output.
- **Multiplicative Fusion:** Separate networks or branches handle each sensing modality up to a specific intermediate layer. The feature maps from these intermediate layers are multiplied element-wise. The combined feature maps undergo further processing within the network to generate the final output.
- **Skip Connections:** Separate networks or branches process each sensing modality. Feature representations from intermediate layers are combined via skip connections. The combined feature maps are further processed by the network to produce the final output.

Drawing on the insights from the authors of [50] and [94], it can be stated that Shortcut Fusion and Deep Fusion serve as overarching frameworks for integrating features from diverse modalities. Within these frameworks, specific techniques such as Additive Fusion, Multiplicative Fusion, and Skip Connections can be employed. Each of these methods brings its own unique strengths and potential challenges to the table. Additive Fusion, Multiplicative Fusion, and Skip Connections are distinct techniques that can be utilised to realise these fusion strategies. The choice between them depends on the specific requirements of the task at hand.

In [34], CNNs were used for feature extraction on the visual input of two modalities. In addition, hand-crafted feature extraction was applied to the thermal data for fusing these features in a Support Vector Machine (SVM) model. In their study, the authors asserted that face temperature variation contains significant differences between different actions, and thus can enhance the accuracy of activity recognition. They utilised synchronised thermal images to extract the face

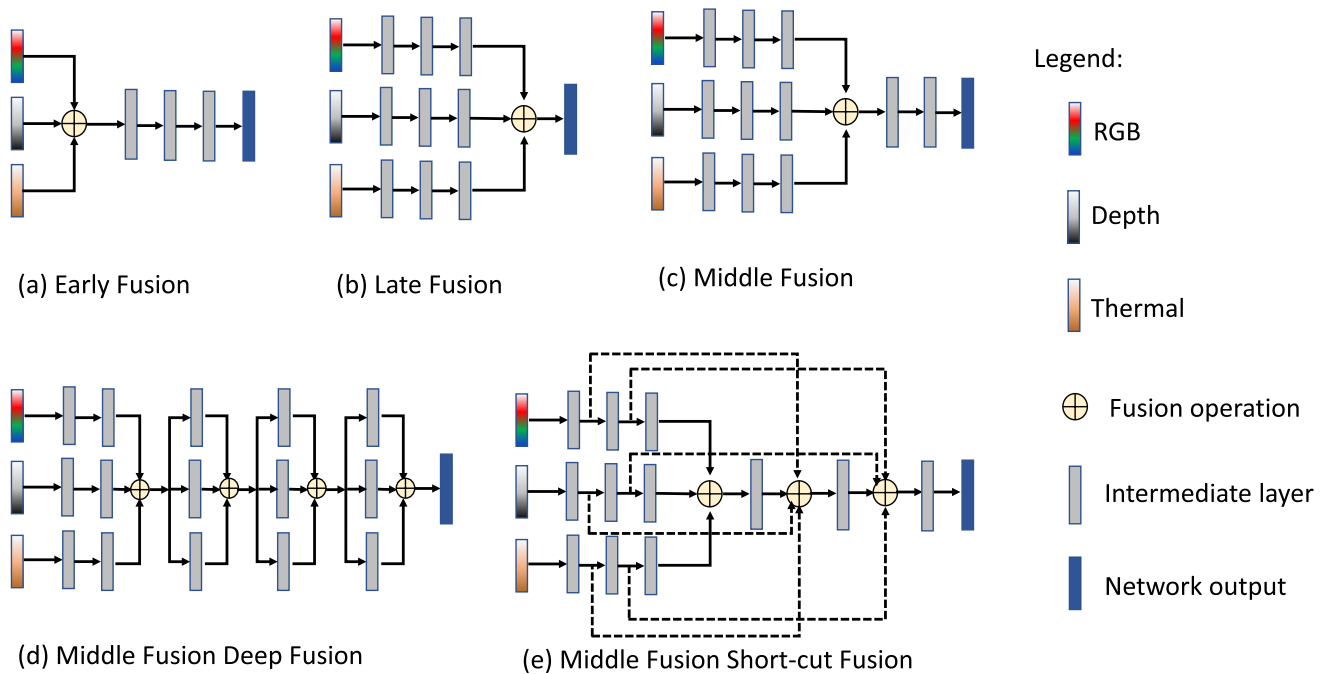


FIGURE 12. A depiction of architectures of different fusion schemes: early fusion, late fusion, and various middle fusion techniques employed in deep learning environments based on [50].

temperature variation of participants while they performed the actions. To achieve this, they manually selected the face region in the first frame and tracked it across frames using a KCF tracker [95]. Outliers in face temperature were removed when the movement was sudden or when the person was partially out of the frame. They divided the temperature values into 25 intervals and computed the average temperature for each interval. Then, they calculated the difference between every two consecutive intervals ($t_i - t_{i-1}$), resulting in 24 features that were added to the SVM model.

The authors in [24] extracted regions of interest (ROI) on the face using the Dlib [71] library to obtain the mean and variance of the temperatures in the ROIs as thermal features. The gait data of lower limbs were combined with these features for emotion detection, as lower limbs have more repeatable movements than the upper body. Joint angles and angular velocities were chosen as the features to characterize gait, including eight gait features based on the angle and velocity of the knees and hip. Convolutional Neural Networks (CNN), Hidden Markov Models (HMM), Support Vector Machines (SVM), and Random Forest (RF) models were employed to train and test the gait and thermal data. CNN and HMM models were trained with time series, whereas SVM and RF models were trained with static features such as the Power Spectral Density (PSD) of time series and the average temperatures of thermal image time series.

3) DECISION LEVEL FUSION (DL)

Late fusion on the other hand offers high flexibility and modularity. When a new sensing modality is introduced,

only the network associated with that modality needs to be trained, leaving the other networks unaffected. However, late fusion comes with drawbacks, such as increased computation costs and memory requirements. Additionally, it discards rich intermediate features that could be highly beneficial if fused, potentially limiting the performance of the overall system [50]. Late fusion in DL is commonly realised by the application of different versions of the non-maximum suppression algorithm (NMS) which works by first selecting the bounding box with the highest object detection score. Then, it compares the remaining bounding boxes and removes the ones that have a high degree of overlap as applied in [6]. Reference [12] investigated the NMS method further and compared the Dual-NMS with the simple method and concluded that the Dual-NMS had a better performance. The Dual-NMS involves sorting two lists of detection boxes based on their confidence scores and collecting pairs from them. Similar to the basic NMS method, the boxes with the highest scores from each list are selected one by one and compared with the boxes from the other list. If a sufficient intersection over union (IOU) is found, the detection box is paired with the candidate with the highest score from the other list. The paired boxes are then merged into a single result, and the final detection box coordinates are updated through weighted averaging of the coordinates of the components [12].

A depiction of various fusion types is shown in Fig. 12.

C. FUSION

The authors in [50] noted that they did not find definitive evidence that one fusion method would be superior to others

based on their review of various methods using different stages. However, [12] and [21] conducted a comparison of the performance between the early middle and late fusion techniques and concluded that early fusion yielded better results in their use cases. These studies focused on the fusion of depth and thermal data for human detection, and it was discovered in [12] that using only depth data did not produce satisfactory detection outcomes. While the late fusion approach was slightly superior to using only depth data, it was inferior to using only the thermal data in comparison. However, [21] further argued that early fusion outperforms the other fusion methods in both the final detection results and computational complexity. Unlike late fusion, which only merges the detected boxes, early fusion enables cooperation between the depth and thermal information during feature extraction, allowing the model to extract and combine useful information from both modalities. It should be noted that while intermediate fusion also merges feature maps, the merging in early fusion is accomplished by a deep backbone network, leading to more effective cooperation between the depth and thermal information. In their study, they also demonstrate that the use of a Receptive Enhancement Module (REM) improves AP by 0.4, 0.9, and 2.3 at IOU thresholds of 0.25, 0.5, and 0.75, respectively. These findings suggest that the REM module enhances the accuracy of bounding box localisation but [12] achieved slightly better scores without it, which could be due to the slightly different fusion and training approach. In [21], the authors used a ResNet-50 backbone that was initialised with pre-trained parameters from ImageNet. However, the REM module and box prediction module were trained from scratch. Reference [12] used the same backbone but models that were pre-trained using the Common Objects in Context dataset (COCO).

Besides constructing custom networks, some studies, such as [32], have utilised deep learning-based algorithms like OpenPose [96] to detect the pose of human occupants in a vehicle. Based on the derived bounding boxes this study applied late fusion. However, because OpenPose can only be applied to visual and thermal data, the authors did not utilise depth data. Other studies like [34] use CNNs for the feature extraction on the visual input of two modalities and apply a hand-crafted feature extraction for the thermal data to fuse these in a Support Vector Machine (SVM) model. However, they dropped the thermal feature in their experiments due to too much noise in the data.

D. SEMANTIC SEGMENTATION

Image semantic segmentation is a crucial task in computer vision, serving as an ideal perception solution for transforming image inputs into semantically meaningful regions and enabling pixel-wise dense scene understanding. Networks that rely solely on RGB information may face limitations in segmentation performance in complex environments or under challenging conditions. To enhance input information and improve performance, researchers have extensively explored

multimodal sensor data fusion, which integrates additional data sources to provide a more comprehensive understanding of the scene. Various approaches have been proposed, such as FuseNet [94], which incorporates depth information, and HeatNet [97], which leverages thermal data for improved performance at night. Polarisation information has also been integrated into models, as seen in EAFNet [98]. Event data has been utilised in dense-to-sparse fusion to capture dynamic context information and improve segmentation performance, as in ISSAFE [99]. Furthermore, there are specialised methods for RGB-D [100], [101], [102], RGB-T [103], [104], [105], [106] and RGB-P [15] semantic segmentation [15].

The authors in [107] argue that recently, vision transformers [89] have gained attention as they handle inputs as sequences and can acquire long-range correlations, providing a unified framework for diverse multi-modal tasks. But that multi-modal data often contain noisy measurements in different sensing modalities, such as low-quality distance estimation regions caused by limited effective depth ranges [108] and that compared to existing multi-modal fusion modules based on Convolutional Neural Networks (CNNs), it is not yet clear whether vision transformers can lead to significant improvements in RGB-X, where X stands for a different modality than RGB, semantic segmentation. Importantly, while some previous works like [108] and [109] use a simple global multi-modal interaction strategy, it may not generalise well across different sensing data combinations [106]. This is why the authors in [107] hypothesise that for RGB-X semantic segmentation with various supplements and uncertainties, comprehensive cross-modal interactions should be provided to fully exploit the potential of cross-modal complementary features why they propose CMX, a method designed to enhance semantic segmentation by incorporating diverse and complementary information from multiple modalities. CMX is a transformer-based cross-modal fusion framework that uses two streams to extract features from RGB images and the X-modality and includes a Cross-Modal Feature Rectification Module (CM-FRM) in each feature extraction stage to calibrate the feature of the current modality by combining the feature from the other modality. A Feature Fusion Module (FFM) is then used to mix the rectified feature pairs for the final semantic prediction. FFM includes a cross-attention mechanism, enabling the exchange of long-range contexts, and enhancing bi-modal features globally. Using a SegFormer-B2 backbone to visualise the segmentation results demonstrated that CMX improves the semantic segmentation of RGB-D data and identifies objects correctly, which are misclassified by the RGB-only model. For RGB-T segmentation, CMX provides clearer boundary distinctions between persons and unlabeled backgrounds in low illumination conditions. For RGB-P, CMX accurately segments specular glass areas, cars with polarisation cues, and pedestrians. For RGB-Event, CMX enhances the segmentation of moving objects. For RGB-LiDAR, CMX correctly segments the scene as compared to the RGB-only method. The results show that CMX

is a suitable approach for multi-modal sensing combinations, providing robust semantic scene understanding [107]. The proposed CMX framework achieves state-of-the-art performances in different benchmarks but is limited to two simultaneous modalities at the time of writing.

Similar to [107], the authors in [15] put the focus on developing a generalisable multimodal perception system for various image modalities with an attention-based fusion architecture for outdoor scene understanding called NLFNet. This network is designed to effectively address the challenges of object segmentation in various complex scenarios. The NLF (Non-Local Fusion) module, a key component of the network, is capable of adaptively extracting and fusing complementary information from different modal input images. It also leverages dependence information along with long-range contextual and positional priors to enhance the accuracy of semantic segmentation and applies a weighting mechanism based on a sigmoid activation function for fusing the modalities. By addressing these challenges, NLFNet aims to improve the performance of outdoor scene understanding across a range of conditions and input modalities. The network architecture is inspired by efficient networks such as SwiftNet [110] and RFNet [111]. NLFNet uses an encoder-decoder structure and adopts a ResNet-18 [112] backbone for each of its two independent branches. The encoder extracts latent features from RGB and other modal images, which are then merged using fusion operations. The Spatial Pyramid Pool (SPP) module [111], [113] is employed to expand effective receptive fields and generate feature maps with more global contextual information.

NLFNet incorporates efficient upsampling modules from SwiftNet [110] and merges RGB branch information through skip connections, improving segmentation accuracy. The Non-Local Fusion (NLF) module, inspired by Non-Local block [114] and NANet [101], integrates complementary information from RGB and other branches for the multi-level fusion of feature maps. The NLF module consists of two sub-modules: the Spatial Dependency Module (SDM) and the Channel Dependency Module (CDM).

The SDM establishes long-range contextual dependency between RGB and other modal branches in space, using global average pooling and convolutions to expand receptive fields. The CDM concatenates outputs from the SDM module along the channel dimension, obtaining a merged feature map, and performs global average pooling to obtain a squeezed feature map. It then adaptively transforms these embeddings into dependency weights via a sigmoid activation layer. This process establishes non-local contextual dependencies between different modalities and extracts nonlinear interactions between cross-modal channels.

The authors of the study demonstrate the effectiveness and generalisation ability of NLFNet across various multimodal sensor combinations. By conducting experiments with different sensor data, such as RGB-Depth, RGB-Polarisation, and RGB-Thermal images, they showcase the ability of NLFNet

to handle diverse modalities and effectively fuse the complementary information. The results indicate that NLFNet is capable of providing accurate semantic segmentation in various challenging scenarios, proving its potential as a robust solution for outdoor scene understanding. But like CMX [107], the solution is bi-modal only.

In [4] applied the Residual Encoder-Decoder Network (RedNet) [102], which is a high-performing semantic segmentation network proposed in [102] that improves segmentation results by incorporating depth information into RGB signals. RedNet utilises an encoder-decoder network structure [115] with residual blocks as building modules, as well as a pyramid supervision training scheme to optimise the network. The encoder structure includes two convolutional branches, one for RGB and one for depth, that have the same configuration except for the feature channel number of the convolution kernel, and feature fusion is achieved through element-wise summation. During training, the dataset was augmented and stochastic gradient descent (SGD) was used to optimise the network parameters with an initial learning rate of 0.002. The model is capable of segmenting target mask images from the background through inference. The authors augmented the training dataset by applying random scale and crop, followed by random hue, brightness, and saturation adjustment, which increased the dataset from 600 to 60,000 groups. Their model converged after approximately 100 epochs of training.

E. OBJECT DETECTION

There are two main types of object detection algorithms that utilise convolutional neural networks (CNNs): two-stage detectors and single-stage detectors. The R-CNN family is a popular example of two-stage detectors, which typically use region-based methods. One such version is Faster R-CNN, used in [12], which introduced the region proposal network (RPN). The RPN can predict both the bounding box and the score at each position simultaneously, leading to a significant decrease in prediction time. An example of a popular single-stage network is YOLO [72]. Most studies, [6], [9], [37], utilise these networks with slight adjustments and perform early data level fusion for RGB-DT or just RGB-D or RGB-T object detection.

Only a single study, [5], focused on RGB-DT data fusion for Salient Object Detection (SOD) and implemented feature-level fusion with a CNN. In this work, the VGG16 classification network is used as a backbone for feature extraction. The tri-modal images are encoded separately using a three-stream encoding network, which extracts five-level features with varying resolutions. The authors proposed a hierarchical weighted suppress interference (HWSI) method to achieve an effective fusion of cross-modal information while also suppressing interference. The approach taken can be classified as a middle fusion with skip connections. This method involves assigning weights to each modality based




FIGURE 13. The overall architecture of the deep multimodal detection strategy. Object detection using YOLO from clear RGB images with rich colour information and thermal images with clear object bounding lines. The model detects the object with the highest probability through late fusion [6].

on their importance for the given task and then using these weights to selectively suppress the interference introduced by each modality. By hierarchical weighting and selectively suppressing interference, the HWSI method can effectively fuse the cross-modal information but comes with a high computational cost which makes it less suitable for real-time applications. The HWSI method is composed of three distinct modules: the dual-modal attention fusion module (DMAFM), the triple-modal interactive weighting module (TMIWM), and the global attention-weighted fusion module (GAWFM). Each module is specifically designed to employ cross-modal information weighting to emphasize the salient regions and suppress interference effectively. The feature extraction is achieved by applying atrous convolutions with different dilation rates which can improve the performance of the network in tasks such as image segmentation and object detection. This approach, however, relies on the visual representation of the thermal and depth modality and some of the limitations that affect them are discussed in Section XIII-B. The dataset created by the authors of [5] limits the thermal and depth representation to 256 values and a dynamic colour AGC algorithm is applied to the thermal data, as shown in Fig. 19 and no gain control is applied to the depth data as shown in Fig. 21. This can reduce the performance of object detection.

In contrast to the previous study, the authors of [6] utilised late fusion with two separate YOLO [72] models, following the thermal data dehazing process discussed in the Process Support VIII section. The overall system architecture is illustrated in Figure 13. After completing the dehazing process, the resulting $I_{hf}(x)$ and $H(x)$ are fed into two YOLO models, denoted as $YOLO^R$ and $YOLO^T$, which use the image of their respective modality. Non-maximum suppression (NMS) is

then employed to achieve late fusion. Their proposed model also allows for an RGB image with improved quality, as some of the haze removal can be performed using the haze level estimates. Moreover, by using late fusion and thermal images, the proposed model can process the rich colour and clear boundary information from both the RGB and thermal images simultaneously.

Reference [9] investigated using a pre-trained YOLOv4 network on the COCO dataset, as well as training YOLOv4 on their own dataset. They limited the scope to human detection only but investigated different ways of fusing the images in an early fusion as discussed here XI-B1. As their proposed data level fusion comes at a lower computational cost (average of 20ms. per frame) and the single-stage object detectors are fast, this system could be applied to real-time problems such as surveillance similar to study [12]. However, neither study published data related to the inference speed, and while [12] used a public dataset that was temporally aligned, [9] did not discuss temporal sensor alignment and used a thermal camera that was limited to 8 FPS.

In [37], the authors employed all three modalities, RGB, depth, and thermal data, to detect hands using a YOLO-based object detection algorithm [72] for real-time performance. Their analysis of the results suggested that using 2D bounding box detection with all three modalities led to higher accuracy compared to state-of-the-art model-based RGB-D hand pose detection algorithms. Additionally, the authors found that RGB and thermal data were the most crucial modalities for this task. To train the YOLO detector, pre-trained features on ImageNet [116] were used, and the annotated bounding boxes in the dataset were employed for training. Since pre-training is only available for RGB images, knowledge




FIGURE 14. The CNN architecture is trained with two losses, the proposed One Class Contrastive Loss (OCCL) and Binary Cross Entropy (BCE), and then used as a fixed feature extractor with frozen weights. The one-class Gaussian Mixture Model (GMM) is trained using the embeddings obtained from the bona fide class alone [26].




FIGURE 15. Preprocessed images resulting from a rigid mask attack [26].

distillation [117] was used to transfer pre-trained features to the thermal and depth modalities.

Overall, deep learning-based approaches for multi-modal sensor fusion have shown promising results in various applications, and their continued development is expected to significantly advance the capabilities of multi-modal sensing systems in the future.

F. PRESENTATION ATTACK DETECTION (PAD)

Biometrics provides a secure and convenient method for access control. Among various biometric modalities, face biometrics is one of the most preferred due to its non-intrusive nature. Despite the high performance of systems in identifying individuals in many challenging datasets, they are still vulnerable to presentation attacks (PA). It was identified that PAD in visual spectra alone is insufficient for security-critical applications. This area has seen a lot of attention in recent years and PAD systems are another area where fused RGB-DT data was applied. A multi-channel PAD framework called the Multi-Channel Convolutional Neural Network (MCCNN) was proposed in [30]. The MCCNN architecture is an extended version of the LightCNN model [118] adapted specifically for multi-channel PAD tasks and was then also applied in [26]. The main idea behind the MCCNN architecture is to leverage the joint representation from multiple channels for PAD tasks, using a pre-trained face recognition network. In this approach, a pre-trained LightCNN face

recognition model is extended to accept multiple channels, and the embeddings from all channels are concatenated. Two fully connected layers are added on top of this joint representation layer for the PAD task. The first fully connected layer has ten nodes, and the second layer has only one output node. The higher-level features in the LightCNN part are shared among all modalities. The advantage of this architecture is that only lower layer features, known as Domain Specific Units (DSUs) [26], and higher-level fully connected layers are adapted in the training phase. This approach has two main advantages: first, a smaller number of parameters since the high-level features are shared across modalities, and second, adapting only DSUs and the final fully connected layers reduces possible over-fitting since PAD databases are typically small in size. An optimal set of layers to be adapted was obtained empirically and was used in the baseline MCCNN and the proposed approach. Fig. 15 shows a set of preprocessed images and Fig. 14 the Schematic diagram of the proposed system found in [26].

In [26], the authors proposed a framework that utilises a one-class classifier along with a novel loss function, which encourages the CNN to learn a compact yet discriminative representation for face images.

As part of study [30], a publicly available dataset called The Wide Multi-Channel Presentation Attack (WMCA) database, was released. Reference [27] made a similar dataset with higher quality and more modalities available which was called HQ-WCMA.

Three studies [26], [27], [30] were co-written by some of the same authors who further developed their ideas in [27]. The data was collected using a custom-made sensor suite that enabled the recording of both genuine faces and presentation attacks across five different image modalities, including RGB, NIR, SWIR, thermal, and depth. Moreover, four banks of six LED modules were employed for illumination, providing coverage in 10 different wavelengths ranging from

TABLE 3. Publicly available Bi-Modal datasets used in studies.

Study	Name	Type	Modalities	Link
[5]	VDT-2048	General	RGB D T	VDT-Dataset
[39]	VAP Trimodal People Seg.	People	RGB D T	VAP TPS-Dataset
[119]	KAIST	Driving	RGB D* T	KAIST-Dataset
[26], [30]	WMCA	Faces/Masks	RGB D T	WMCA-Dataset
[27]	HQ-WMCA [120]	Faces/Masks	RGB D T	HQ-WMCA-Dataset
[7]	TriModal Face Detection Dataset	Faces	RGB D T	TMFD-Dataset

Includes LiDAR and stereo RGB images.

TABLE 4. Publicly available Bi-Modal datasets used in studies.

Study	Name	Type	Modalities	Link
[12], [21]	IPHD [121]	People	D T	IPHD-Dataset
[107]	NYU Depth V2 [122]	Indoor	RGB D	NYU Depth V2-Dataset
[107]	SUN-RGBD [123]	Indoor	RGB D	SUN-RGBD-Dataset
[107]	Stanford2D3D [124]	Indoor	RGB D	Stanford2D3D-Dataset
[107]	ScanNetV2 [125]	Indoor	RGB D	ScanNetV2-Dataset
[107]	Cityscapes [126]	Driving	RGB D	Cityscapes-Dataset
[107]	MFNet [92]	Driving	RGB T	MFNet-Dataset
[107]	EventScape [127]	Driving	RGB E	EventScape-Dataset
[107]	KITTI-360 [128]	Driving	RGB L	KITTI-360-Dataset

735nm to 1650nm, encompassing the NIR and SWIR spectra. Sequential switching of these infrared emitters, synchronised with camera exposure periods, yielded multi-spectral reflectivity data across the sample. These wavelengths were chosen to provide the best possible multi-spectral coverage given market availability, resulting in 14 different modalities in each recording, including four NIR and seven SWIR wavelengths. The cameras were co-registered using a calibration procedure, enabling the captured data to be aligned in each modality. Experimental results showed that the investigated CNN models with SWIR outperformed baselines when a wide variety of attacks was considered, with almost perfect detection of all impersonation attacks while maintaining a low BPCER. However, the generalisation ability of the models using SWIR data was assessed on a cross-database experiment, revealing a noticeable difference on bona fide attempts, suggesting future research directions.

The proposed database 3 and code for studies [26]¹ and [27]² for reproducing the experiments are freely available for research purposes.

XII. DATASETS

The majority of studies examined in this paper faced a scarcity of publicly available datasets, leading them to develop their own. Although some studies claimed to make their datasets public, like the TriModal Face Detection dataset(TMFD) [7], it could not be found during the writing of this survey. However, subsequent to the preprint release of our paper, we were contacted by the authors of the TMFD, who have now made their dataset publicly available. This dataset is a comprehensive resource that encompasses a wide range of variations, including different numbers of people in the scene, various backgrounds and distances. The dataset is categorised

into three separate groups based on the complexity and difficulty level of face detection. Other studies, including [5] and [39], created and made their datasets available and a list of all available tri-modal datasets can be found in Table 3. Public bi-modal datasets used by some reviewed papers are listed in Table 4. As this paper is centred on RGB-DT tri-modal fusion, only datasets used by the studies included in this review are presented.

XIII. LIMITATIONS

A. SENSORS

One of the limitations of using thermal cameras in conjunction with RGB-D cameras is the potential mismatch in their respective field of view (FOV) and focal length, which can restrict the effective distance between objects or subjects being monitored. This discrepancy can result in inconsistencies in the size, position, and orientation of objects in the captured images, which can impact the accuracy of object detection and tracking algorithms. Another thermal sensor limitation is the need for Non-Uniformity Correction (NUC). NUC compensates for inconsistencies in the sensor's response to temperature changes, which can lead to inaccuracies in temperature measurements. This correction is required periodically and involves a mechanical shutter operation that blocks the imaging sensor with a material of uniform temperature for a short time up to a second.

The authors in [119] identified limitations with the temporal alignment in capturing images simultaneously using multiple devices. Despite the use of a signal generator to match the shutter times between devices, there can still be drift due to differences in exposure times. This can lead to an asynchronous phenomenon, especially during excessive movement of a vehicle or object.

It is important to note that temporal alignment is a critical factor in multi-camera systems, as it ensures that the images

¹https://gitlab.idiap.ch/bob/bob.paper.oneclass_mccnn_2019

²https://gitlab.idiap.ch/bob/bob.paper.pad_mccnns_swirldiff

captured by different cameras are synchronised and can be properly used in applications such as 3D reconstruction or object detection.

B. MODALITIES

In [5], the authors identified that visual perception systems that rely solely on RGB cameras face challenges such as:

- 1) The objects to be recognised in indoor environments are often small, numerous, dense, and vulnerable to background interference.
- 2) In low-light conditions, the ability to detect objects is greatly reduced, as illustrated in Fig. 16.

To overcome the above problems, thermal and depth modalities can be introduced but despite that those sensors improve the detection of salient objects, they also introduce interference challenges and have their own individual challenges as can be seen in Fig. 17 and Fig. 18.




FIGURE 16. RGB modality challenges. (a) Similar appearance. (b) Small salient object. (c) Low illumination [5].

Fig. 17(a) illustrates that the background of the depth image without any salient objects is very cluttered, which can distract the detection of salient object detection algorithms. Also, the depth information of a salient object can be incomplete when there is no distance difference between it and the surrounding objects, or when the difference is minimal. Furthermore, depth sensing can still be challenging for detecting some small objects.




FIGURE 17. Difficult challenges of depth images. (a) Background messy. (b) Depth information is incomplete. (c) Small salient objects [5].

Thermal sensors also present several challenges that need to be addressed, including thermal crossover, thermal radiation dispersion, and heat reflection. Thermal crossover occurs when the temperature of a salient object is the same as that of a portion of the background, as illustrated in 18(a), greatly increasing the difficulty of object detection. Fig. 18(b) demonstrates an example of thermal radiation dispersion, where a portion of a salient object appears more salient than the rest of the object, causing interference to detection. Additionally, some objects exhibit heat reflection phenomena, as shown in 18(c), which is another important interference that needs to be addressed.

Accurate assignment of thermal values, as identified in [45], requires careful attention due to the nature of thermal-infrared sensors. As regular NUCs are required, real-time systems must be able to cope with thermal data interruption. If the correctly measured temperature is of importance, it should be considered that the thermal value can also be affected by the incident angle between the sensor and radiation emitted from the surface. Minimising this angle is considered best practice. The authors of [45] suggested three possible strategies to mitigate this:

- Perform Non-Uniformity Corrections (NUCs) more frequently, approximately every minute.
- Disregard frames obtained from the camera while a NUC is in progress.
- Assign temperatures only to rays with an incident angle of less than 30 degrees.

XIV. SYNTHESIS

This section is intended to synthesise the key findings from our extensive review of the literature on the fusion of RGB-DT sensor modalities. It encapsulates the current state of the art, summarising the significant advancements and methodologies in this field across various applications. Before we transition into discussing the challenges, future work, and conclusion, this synthesis serves as a succinct recapitulation of the key points, including some insightful observations.

The traditional approach for the geometric calibration of thermal cameras, using a printed chessboard and a flood lamp, was inaccurate and difficult to execute. Geometric masks with high thermal contrast were introduced as an alternative calibration pattern, and multi-material calibration boards made of two materials with different emissivities have been developed for cross-calibrating thermal and visual modalities which have proven to be reliable and accurate. The registration of the modalities is applied based on the processing requirements, with offline approaches utilising computationally intensive feature-point matching algorithms. One widely used technique, especially for large-scale 3D reconstruction, is the Bundle Block Adjustment (BBA) with current advancements and improvements to the basic approach, such as using more advanced optimisation algorithms (e.g. Levenberg-Marquardt [129], Gauss-Newton [130]) [10]. In real-time 3D reconstruction, thermal data was added to the back-end of SLAM systems to enhance the robustness under unstable illumination environments and research in this area focuses on improving real-time performance by further reducing computational time and offering better model quality. In contrast, real-time processing for semantic segmentation and object detection requires performing geometric image rectification and alignment as a preprocessing step to ensure correct feature extraction. While studies for multi-modal semantic segmentation of RGB, depth and thermal data based on recent transformer networks were found [15], [107], those works only processed two




FIGURE 18. Taken from the VDT-2048 dataset, demonstrating the three identified thermal challenges: (a) Thermal crossover, (b) Thermal radiation dispersion, (c) Heat reflection [5].

modalities at a time and future research is aimed at processing more modalities simultaneously. However, the two studies demonstrate how to fuse bi-modal in real-time successfully and [15] used an adaptive weighting of modalities with a sigmoid activation layer to limit interference. The majority of the reviewed papers on multi-modal object detection utilised early fusion to generate a fused 8-bit three-channel image. Common object detectors such as YOLO were employed for the detection task, and some studies incorporated late fusion with the bounding boxes acquired from individual streams. A single study, [5], was identified that fused all three modalities (RGB, depth, and thermal) in a neural network using VGG16 as the feature extraction backbone. The study incorporated an interference suppression module to weigh the feature information across different modalities, mitigating interference from a single modality and compensating for potential information gaps in certain modalities. However, the computational requirements of this approach rendered it unsuitable for real-time processing. This study additionally created a publicly available, generic tri-modal dataset. Apart from this dataset, there are only three other RGB-DT datasets, which are specialised in presentation attack detection (PAD) and human detection applications. Another area that attracted a lot of attention is PAD where researchers focused on the problem of generalisation of the system. The challenge was addressed in [26] by building upon the Multi-Channel CNN (MCCNN) originally proposed in [30]. The authors developed a one-class classifier framework that employs learned features and a new loss function. This innovative loss function compels the CNN to acquire a concise and discriminative representation of face images, which enhances the overall performance even when used with the RGB modality alone. The authors showcased that their CNN method surpasses existing state-of-the-art feature-based techniques, while future research will focus on addressing the issue of potential attackers attempting to impersonate others.

XV. CHALLENGES AND FUTURE WORK

A. DATA FUSION

When it comes to fusing different modalities, one of the main challenges is sensor calibration and registration,

especially when the sensors have different fields of view (FOV) that can cause parallax. To address this, techniques such as geometric calibration and image registration can be used to align the data from different sensors and reduce the effects of parallax. However, these techniques can be complex and time-consuming and still not produce perfect results, which can affect the accuracy of the fused data. Aside from employing software solutions to rectify parallax, another alternative involves using a beam-splitter, which enables two cameras to view the scene from the same point and utilise similar lenses to minimise parallax effects. Eliminating misalignment between modalities is vital in early fusion, which is why sensor calibration and registration continue to be significant research topics in this domain.




FIGURE 19. The image on the left shows a hot bottle on a table and the image on the right shows a cold bottle on the same table. Taken from the VDT-2048 dataset demonstrating the AGC colour shift of the same object (table) due to the application of a dynamic colour range based on the global minimum and maximum temperature in the frame.

B. THERMAL DATA

The mapping of colours for the display of thermal data is a crucial element for systems that utilise thermal data in a visual form. However, there appears to be a shortage of discussion on this topic in the reviewed literature. Most cameras apply automatic gain control (AGC) which is based on the lowest and highest temperature at any given time, causing the colours to shift. Besides that, grayscale or colour images limited to 256 values are being used but it is not specified over what range of temperatures it is used for e.g. when monitoring a range of -20 to 120°C , the resolution would be 0.55°C . However, thermal cameras have a sensitivity expressed as




FIGURE 20. Taken from the VDT-2048 dataset, the image on the left shows an overlay of Visual(V) and Thermal(T) with some objects in T outlined in yellow and the same objects outlined in red in V. A parallax and distortion between the modalities can be observed. The images on the top right and on the bottom right are T and V respectively.

Noise Equivalent Delta Temperature (NEdT), which can range from 0.020°C up to 0.075°C . It is crucial to consider both aspects when analysing thermal images as limiting the data to 8-bit discards detail. For example, the VDT-2048 dataset uses 256-colour thermal images with a dynamic thermal-to-colour range association, as shown in Fig. 19. The adaptive AGC algorithm may be appropriate for some applications; however, it could lead to difficulties if the intensity or colour information is essential for feature extraction or used for network training. Furthermore, the dataset exhibits some distortion and parallax between the visual and thermal modalities, as illustrated in Fig. 20. Notably, the KAIST [119] driving dataset features raw 14-bit thermal data; conversely, the pedestrian dataset only includes 8-bit data. It is important to mention that while a 14-bit sensor can represent values up to 16,383, in environments with ambient temperatures around 20°C , the raw data captured falls within a narrow band of the full range. As a result, compression and contrast enhancement is crucial for encoding thermal images. However, it is essential to recognise that enhancement operations in thermal images can artificially distort the data, causing the loss of the physical correlation between the radiant flux from infrared radiation and pixel intensity [65]. Besides the data preprocessing, authors in [5] identified thermal crossover, thermal radiation dispersion, and heat reflection as challenges when processing thermal data. It is believed that a thorough preprocessing of thermal data and addressing the identified challenges in the field of multi-modal fusion constitutes a relevant future research direction.

C. DEPTH DATA

Similar to thermal data, depth data is initially captured with a 16-bit resolution, but it is later converted to an 8-bit format when used as a depth map unlike point clouds, which are usually generated from the raw values. This conversion from 16-bit to 8-bit can lead to a loss of depth resolution and information due to the reduction in detail. Although this process can result in significant information loss, no studies in the reviewed literature have addressed this issue. It is essential to

conduct thorough preprocessing of this modality when using it in the form of an 8-bit depth map. While the dynamically applied AGC algorithm in thermal images could cause issues, applying no processing at all will result in a loss of details. Figure 21 shows image 18 taken from the VDT-2048 dataset, in the original image on the left, it can be noted that visually almost nothing can be recognised as the observed depth is limited to a narrow band in the 16-bit data that was converted to 8-bit. On the right the same image with adjusted tonal balance by redistributing its brightness values. This is done by mapping the brightest and darkest pixel values in the image to white and black, respectively, and redistributing all the intermediate values evenly across the entire range.



FIGURE 21. A comparative display of the original (left) and equalised (right) version of image 18 from the VDT-2048 dataset. The stark contrast between the two images accentuates the pivotal role of pre-processing in enhancing feature visibility, a critical step for the effective application of convolutional neural networks.

D. DATASETS

The research potential in the field of tri-modal RGB-DT object detection is currently limited due to the lack of publicly available datasets. Apart from the VDT-2048 dataset mentioned earlier, there are no other tri-modal datasets suitable for general object detection, highlighting the need for more comprehensive datasets to advance research in this area.

E. DEEP LEARNING

Integrating multiple modalities can improve object detection or segmentation accuracy, but it also increases computational requirements and makes real-time processing difficult. In the literature, studies show that early and late fusion of data has a relatively small impact on real-time performance; however, the potential for more sophisticated enhancements remains to be a challenge. More complex processing techniques applied in middle fusion, as demonstrated in [5] and [6], can result in frame rates dropping below 5 FPS. Even with the advances in Deep Neural Networks (DNNs) and hardware technology, achieving real-time object detection or segmentation with RGB-DT data still remains a challenge. While transformer-based architectures, as demonstrated in [107], and CNN-based architectures with non-local blocks, as demonstrated in [15], have shown promising results, they were limited to two concurrent modalities at the time of writing. Thus, further research is required to develop efficient and accurate fusion algorithms that can utilise three modalities and meet the requirements of real-time processing.

F. PAD

Currently, research in the field of PAD (Presentation Attack Detection) algorithms is centred on devising new methods capable of accurately detecting both known and unknown attacks. A major challenge for existing PAD algorithms is generalisation, as they often exhibit bias towards the training data. In [27], the authors recognised the SWIR (Short-Wave Infrared) spectrum, typically defined as light in the $0.9 - 1.7\mu\text{m}$ wavelength range but can also range from $0.7 - 2.5\mu\text{m}$, as complementary and valuable. However, InGaAs (Indium Gallium Arsenide) sensor-based cameras for this spectrum are costly and reserved for specific applications. Meanwhile, in [26] the researchers tackled the generalisation issue by building upon the Multi-Channel CNN (MCCNN) initially proposed in [30]. Their CNN approach with an innovative loss function outperformed all other methods. Future research will concentrate on addressing the challenge of potential attackers attempting to impersonate others.

XVI. CONCLUSION

This paper presents a comprehensive overview of the fusion between RGB-D and thermal modalities, exploring their applications, and the techniques employed. Over the past decade, there has been a surge of interest in fusing these modalities, demonstrating their considerable potential across diverse fields, including robotics, surveillance, medical imaging, and maintenance systems. Combining these modalities has proven to enhance the accuracy, robustness, and reliability of computer vision systems, contributing to the overall effectiveness of the technology. To systematically summarise the findings, a search strategy based on the PRISMA framework was used. The literature review has revealed several approaches for integrating RGB-D and thermal data, including feature-level fusion, decision-level fusion, and data-level fusion. Furthermore, the use of deep learning techniques has emerged as a popular approach for effectively combining RGB-D and thermal data, surpassing traditional feature-based approaches. Overall, the reviewed literature suggests that the fusion of RGB-D and thermal modalities holds great potential for enhancing the performance of computer vision systems in diverse applications and even creating new ones.

It was observed that researchers have primarily focused on the higher-level architecture of neural networks when conducting sensor fusion with deep learning while overlooking the importance of preprocessing steps. While Visual Transformers have shown promising results in sensor fusion, no existing tri-modal RGB-DT fusion has been developed thus far. Therefore, further research is necessary to develop advanced fusion techniques that can enhance the accuracy and reliability of the results while operating in real-time, thereby unlocking the full potential of this approach and making it applicable for various practical applications. In conclusion, this study aims to serve as a supplementary resource for researchers in the field of RGB-DT sensor fusion, providing a robust foundation and guidance for ongoing investigation and advancements.

REFERENCES

- [1] W. Minkina and S. Dudzik, *Infrared Thermography: Errors and Uncertainties*. Hoboken, NJ, USA: Wiley, 2009.
- [2] K. Skala, T. Lipić, I. Sović, L. Gjenero, and I. Grubić, "4D thermal imaging system for medical applications," *Periodicum Biologorum*, vol. 113, no. 4, pp. 407–416, 2011.
- [3] M. J. Page et al., "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *Systematic Rev.*, vol. 10, no. 1, p. 89, Dec. 2021, doi: [10.1186/s13643-021-01626-4](https://doi.org/10.1186/s13643-021-01626-4).
- [4] X. Zhang, Y. Zhang, J. Geng, J. Pan, X. Huang, and X. Rao, "Feather damage monitoring system using RGB-depth-thermal model for chickens," *Animals*, vol. 13, no. 1, p. 126, Dec. 2022.
- [5] K. Song, J. Wang, Y. Bao, L. Huang, and Y. Yan, "A novel visible-depth-thermal image dataset of salient object detection for robotic visual perception," *IEEE/ASME Trans. Mechatronics*, vol. 28, no. 3, pp. 1558–1569, Jun. 2023.
- [6] S. Yoon and J. Cho, "Deep multimodal detection in reduced visibility using thermal depth estimation for autonomous driving," *Sensors*, vol. 22, no. 14, p. 5084, Jul. 2022.
- [7] W. Mucha and M. Kampel, "Depth and thermal images in face detection—A detailed comparison between image modalities," in *Proc. 5th Int. Conf. Mach. Vis. Appl. (ICMVA)*, Feb. 2022, pp. 16–21.
- [8] M. Oppliger, J. Gutknecht, R. Gubler, M. Ludwig, and T. Loeliger, "Sensor fusion of 3D Time-of-Flight and thermal infrared camera for presence detection of living beings," in *Proc. IEEE Sensors*, Oct. 2022, pp. 1–4.
- [9] A. Ozcan and O. Cetin, "A novel fusion method with thermal and RGB-D sensor data for human detection," *IEEE Access*, vol. 10, pp. 66831–66843, 2022.
- [10] T. Zhang, L. Hu, Y. Sun, L. Li, and D. Navarro-Alarcon, "Computing thermal point clouds by fusing RGB-D and infrared images: From dense object reconstruction to environment mapping," in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Dec. 2022, pp. 1707–1714.
- [11] R. Qiu, Y. Miao, M. Zhang, and H. Li, "Detection of the 3D temperature characteristics of maize under water stress using thermal and RGB-D cameras," *Comput. Electron. Agricult.*, vol. 191, Dec. 2021, Art. no. 106551.
- [12] W. Gutfeter and A. Pacut, "Fusion of depth and thermal imaging for people detection," *J. Telecommunications Inf. Technol.*, vol. 4, no. 2021, pp. 53–60, Dec. 2021.
- [13] E. Semenishchev, V. Voronin, S. Agaian, M. Zhdanova, and A. Zelensky, "Algorithm for fusing data obtained by thermal, 3D, and the visible range cameras," *Proc. SPIE*, vol. 11732, Apr. 2021, Art. no. 117320E.
- [14] Y. Zefri, I. Sebari, H. Hajji, and G. Aniba, "In-depth investigation of applied digital photogrammetry to imagery-based RGB and thermal infrared aerial inspection of large-scale photovoltaic installations," *Remote Sens. Appl., Soc. Environ.*, vol. 23, Aug. 2021, Art. no. 100576.
- [15] R. Yan, K. Yang, and K. Wang, "NLFNet: Non-local fusion towards generalized multimodal semantic segmentation across RGB-depth, polarization, and thermal images," in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Dec. 2021, pp. 1129–1135.
- [16] M. Ortega, E. Ivorra, A. Juan, P. Venegas, J. Martínez, and M. Alcañiz, "MANTRA: An effective system based on augmented reality and infrared thermography for industrial maintenance," *Appl. Sci.*, vol. 11, no. 1, p. 385, Jan. 2021.
- [17] Y. Cao, Y. Dong, F. Wang, J. Yang, Y. Cao, and X. Li, "Multi-sensor spatial augmented reality for visualizing the invisible thermal information of 3D objects," *Opt. Lasers Eng.*, vol. 145, Oct. 2021, Art. no. 106634.
- [18] H. Zheng, X. Zhong, J. Yan, L. Zhao, and X. Wang, "A thermal performance detection method for building envelope based on 3D model generated by UAV thermal imagery," *Energies*, vol. 13, no. 24, p. 6677, Dec. 2020.
- [19] M. Jarzabek-Rychard, D. Lin, and H.-G. Maas, "Supervised detection of Façade openings in 3D point clouds with thermal attributes," *Remote Sens.*, vol. 12, no. 3, p. 543, Feb. 2020.
- [20] I. Halima, J.-M. Laferté, G. Cormier, A.-J. Fougères, and J.-L. Dillenseger, "Depth and thermal information fusion for head tracking using particle filter in a fall detection context," *Integr. Comput.-Aided Eng.*, vol. 27, no. 2, pp. 195–208, Feb. 2020.
- [21] Z. Zhao, J. Zhang, and S. Shan, "Noise robust hard example mining for human detection with efficient depth-thermal fusion," in *Proc. 15th IEEE Int. Conf. Autom. Face Gesture Recognit. (FG)*, Nov. 2020, pp. 809–813.

- [22] F. Javadnejad, D. T. Gillins, C. E. Parrish, and R. K. Slocum, "A photogrammetric approach to fusing natural colour and thermal infrared UAS imagery in 3D point cloud generation," *Int. J. Remote Sens.*, vol. 41, no. 1, pp. 211–237, Jan. 2020.
- [23] Y. Shi, P. Payer, M. Frize, and E. Bariciak, "Thermal and RGB-D imaging for necrotizing enterocolitis detection," in *Proc. IEEE Int. Symp. Med. Meas. Appl. (MeMeA)*, Jun. 2020, pp. 1–6.
- [24] C. Yu and A. Tapus, "Multimodal emotion recognition with thermal and RGB-D cameras for human-robot interaction," in *Proc. Companion ACM/IEEE Int. Conf. Human-Robot Interact.*, Mar. 2020, pp. 532–534.
- [25] F. Zhao, S. Cosar, N. Bellotto, and S. Yue, "ROI detection and tracking for physiological monitoring based on calibration between RGB-D and thermal cameras," Liverpool John Moores Univ., 2018. [Online]. Available: <https://bpb-eu-w2.wpmucdn.com/blogs.lincoln.ac.uk/dist/2/1110/files/2020/06/Zhao2018.pdf>
- [26] A. George and S. Marcel, "Learning one class representations for face presentation attack detection using multi-channel convolutional neural networks," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 361–375, 2021.
- [27] G. Heusch, A. George, D. Geissbühler, Z. Mostaani, and S. Marcel, "Deep models and shortwave infrared information to detect face presentation attacks," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 2, no. 4, pp. 399–409, Oct. 2020. [Online]. Available: http://publications.idiap.ch/downloads/papers/2020/Heusch_TBIOM_2020.pdf
- [28] D. Li, C. Menassa, and V. R. Kamat, "Thermal and RGB-D sensor fusion for non-intrusive human thermal comfort assessment," in *Proc. CIB World Building Congr.*, Hong Kong, Jun. 2019, pp. 1–11.
- [29] X. Chen, G. Tian, J. Wu, C. Tang, and K. Li, "Feature-based registration for 3D eddy current pulsed thermography," *IEEE Sensors J.*, vol. 19, no. 16, pp. 6998–7004, Aug. 2019.
- [30] A. George, Z. Mostaani, D. Geissenbuhler, O. Nikisins, A. Anjos, and S. Marcel, "Biometric face presentation attack detection with multi-channel convolutional neural network," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 42–55, 2020.
- [31] S. Brahmabhatt, C. Ham, C. C. Kemp, and J. Hays, "ContactDB: Analyzing and predicting grasp contact via thermal imaging," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8701–8711.
- [32] C. Song and S.-H. Kim, "Robust vehicle occupant detection based on RGB-depth-thermal camera," *J. Korea Robot. Soc.*, vol. 13, no. 1, pp. 31–37, Mar. 2018.
- [33] D. Jacob and A. Tapus, "First attempts in deception detection in HRI by using thermal and RGB-D cameras," in *Proc. 27th IEEE Int. Symp. Robot Human Interact. Commun. (RO-MAN)*, Aug. 2018, pp. 652–658.
- [34] M. Sorostinean and A. Tapus, "Activity recognition based on RGB-D and thermal sensors for socially assistive robots," in *Proc. 15th Int. Conf. Control, Autom., Robot. Vis. (ICARCV)*, Nov. 2018, pp. 1298–1304.
- [35] M.-D. Yang, T.-C. Su, and H.-Y. Lin, "Fusion of infrared thermal image and visible image for 3D thermal model reconstruction using smartphone sensors," *Sensors*, vol. 18, no. 7, p. 2003, Jun. 2018.
- [36] Y. Cao, B. Xu, Z. Ye, J. Yang, Y. Cao, C.-L. Tisse, and X. Li, "Depth and thermal sensor fusion to enhance 3D thermographic reconstruction," *Opt. Exp.*, vol. 26, no. 7, pp. 8179–8193, 2018.
- [37] R. Luo, O. Sener, and S. Savarese, "Scene semantic reconstruction from egocentric RGB-D-thermal videos," in *Proc. Int. Conf. 3D Vis. (3DV)*, Oct. 2017, pp. 593–602.
- [38] I. R. Spremolla, M. Antunes, D. Aouada, and B. E. Ottersten, "RGB-D and thermal sensor fusion-application in person tracking," in *Proc. 11th Joint Conf. Comput. Vis., Imag. Comput. Graph. Theory Appl. (VISIGRAPP)*, vol. 3, 2016, pp. 612–619.
- [39] C. Palmero, A. Clapés, C. Bahnsen, A. Møgelmoose, T. B. Moeslund, and S. Escalera, "Multi-modal RGB-depth-thermal human body segmentation," *Int. J. Comput. Vis.*, vol. 118, no. 2, pp. 217–239, Jun. 2016.
- [40] W. Nakagawa, K. Matsumoto, F. de Sorbier, M. Sugimoto, H. Saito, S. Senda, T. Shibata, and A. Iketani, "Visualization of temperature change using RGB-D camera and thermal camera," in *Proc. Eur. Conf. Comput. Vis. (ECCV) Workshops*. Zurich, Switzerland: Springer, Sep. 2014, pp. 386–400.
- [41] R. Irani, K. Nasrollahi, M. O. Simon, C. A. Corneanu, S. Escalera, C. Bahnsen, D. H. Lundtoft, T. B. Moeslund, T. L. Pedersen, M.-L. Klitgaard, and L. Petrini, "Spatiotemporal analysis of RGB-D-T facial images for multimodal pain level recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2015, pp. 88–95.
- [42] J. Rangel, J. Garzón, J. Sofrony, and A. Kroll, "Gas leak inspection using thermal, visual and depth images and a depth-enhanced gas detection strategy," *Revista de Ingeniería*, vol. 42, no. 1, pp. 8–15, 2015.
- [43] S. Vidas, P. Moghadam, and S. Sridharan, "Real-time mobile 3D temperature mapping," *IEEE Sensors J.*, vol. 15, no. 2, pp. 1145–1152, Feb. 2015.
- [44] J. Rangel, S. Soldan, and A. Kroll, "3D thermal imaging: Fusion of thermography and depth cameras," in *Proc. Int. Conf. Quant. Infr. Thermograp.*, vol. 3, 2014, pp. 1–10.
- [45] S. Vidas, P. Moghadam, and M. Bosse, "3D thermal mapping of building interiors using an RGB-D and thermal camera," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2013, pp. 2311–2318.
- [46] L. Susperregi, J. M. Martínez-Otzeta, A. Ansuategui, A. Iburguren, and B. Sierra, "RGB-D, laser and thermal sensor fusion for people following in a mobile robot," *Int. J. Adv. Robotic Syst.*, vol. 10, no. 6, p. 271, Jun. 2013.
- [47] A. Møgelmoose, C. Bahnsen, T. B. Moeslund, A. Clapes, and S. Escalera, "Tri-modal person re-identification with RGB, depth and thermal features," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2013, pp. 301–307.
- [48] D. Borrmann, A. Nüchter, M. Dakulović, I. Maurović, I. Petrović, D. Osmanković, and J. Velagić, "The project ThermalMapper—Thermal 3D mapping of indoor environments for saving energy," *IFAC Proc. Volumes*, vol. 45, no. 22, pp. 31–38, 2012.
- [49] Y. Ham and M. Golparvar-Fard, "Rapid 3D energy performance modeling of existing buildings using thermal and digital imagery," in *Proc. Construct. Res. Congr., Construct. Challenges Flat World*, May 2012, pp. 991–1000.
- [50] D. Feng, C. Haase-Schütz, L. Rosenbaum, H. Hertlein, C. Gläser, F. Timm, W. Wiesbeck, and K. Dietmayer, "Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 3, pp. 1341–1360, Mar. 2021.
- [51] C. Godard, O. M. Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 6602–6611.
- [52] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
- [53] S. Vidas, R. Lakemond, S. Denman, C. Fookes, S. Sridharan, and T. Wark, "A mask-based approach for the geometric calibration of thermal-infrared cameras," *IEEE Trans. Instrum. Meas.*, vol. 61, no. 6, pp. 1625–1635, Jun. 2012.
- [54] N. Kim, Y. Choi, S. Hwang, K. Park, J. S. Yoon, and I. S. Kweon, "Geometrical calibration of multispectral calibration," in *Proc. 12th Int. Conf. Ubiquitous Robots Ambient Intell. (URAI)*, Oct. 2015, pp. 384–385.
- [55] T. Zhang, L. Hu, L. Li, and D. Navarro-Alarcon, "Towards a multispectral RGB-IR-UV-D vision system—Seeing the invisible in 3D," in *Proc. IEEE Int. Conf. Robot. Biomimetics (ROBIO)*, Dec. 2021, pp. 1723–1728.
- [56] G. Bradski and A. Kaehler. (2000). *OpenCV Library*. Accessed: Mar. 22, 2023. [Online]. Available: <https://opencv.org/>
- [57] G. Ben-Artzi, T. Halperin, M. Werman, and S. Peleg, "Epipolar geometry based on line similarity," in *Proc. 23rd Int. Conf. Pattern Recognit. (ICPR)*, Dec. 2016, pp. 1864–1869.
- [58] R. Istenic, D. Heric, S. Ribaric, and D. Zazula, "Thermal and visual image registration in Hough parameter space," in *Proc. 14th Int. Workshop Syst., Signals Image Process. 6th EURASIP Conf. Focused Speech Image Process., Multimedia Commun. Services*, Jun. 2007, pp. 106–109.
- [59] T. M. Inc. (2022). *MATLAB*. Accessed: Mar. 22, 2023. [Online]. Available: <https://www.mathworks.com/products/matlab.html>
- [60] K. Kraus, *Photogrammetry: Geometry from Images and Laser Scans*, 2nd ed. Berlin, Germany: Walter de Gruyter, 2007.
- [61] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004.
- [62] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," in *Proc. 9th Eur. Conf. Comput. Vis.* (Lecture Notes in Computer Science), vol. 3951, May 2006, pp. 404–417.
- [63] S. Leutenegger, M. Chli, and R. Y. Siegwart, "BRISK: Binary robust invariant scalable keypoints," in *Proc. Int. Conf. Comput. Vis.*, Nov. 2011, pp. 2548–2555.
- [64] P. F. Alcantarilla, A. Bartoli, and A. J. Davison, "KAZE features," in *Proc. 12th Eur. Conf. Comput. Vis. (ECCV)*. Florence, Italy: Springer, Oct. 2012, pp. 214–227.

- [65] M. J. Sousa, A. Moutinho, and M. Almeida, "Thermal infrared sensing for near real-time data-driven fire detection and monitoring systems," *Sensors*, vol. 20, no. 23, p. 6803, Nov. 2020.
- [66] J. Nie, J. Yan, H. Yin, L. Ren, and Q. Meng, "A multimodality fusion deep neural network and safety test strategy for intelligent vehicles," *IEEE Trans. Intell. Vehicles*, vol. 6, no. 2, pp. 310–322, Jun. 2021.
- [67] M. P. Muresan, I. Giosan, and S. Nedeveschi, "Stabilization and validation of 3D object position using multimodal sensor fusion and semantic segmentation," *Sensors*, vol. 20, no. 4, p. 1110, Feb. 2020.
- [68] M. Kowalski and K. Mierzejewski, "Detection of 3D face masks with thermal infrared imaging and deep learning techniques," *Photon. Lett. Poland*, vol. 13, no. 2, pp. 22–24, 2021.
- [69] D.-E. Kim, B. Jeon, and D.-S. Kwon, "3D convolutional neural networks based fall detection with thermal camera," *J. Korea Robot. Soc.*, vol. 13, no. 1, pp. 45–54, Mar. 2018.
- [70] T. Baltusaitis, P. Robinson, and L.-P. Morency, "Constrained local neural fields for robust facial landmark detection in the wild," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Dec. 2013, pp. 354–361.
- [71] D. E. King, "Dlib-ml: A machine learning toolkit," *J. Mach. Learn. Res.*, vol. 10, pp. 1755–1758, Jul. 2009.
- [72] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 779–788.
- [73] S. Ullman, "The interpretation of structure from motion," *Proc. Roy. Soc. London B, Biol. Sci.*, vol. 203, no. 1153, pp. 405–426, Jan. 1979.
- [74] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-time dense surface mapping and tracking," in *Proc. 10th IEEE Int. Symp. Mixed Augmented Reality*, Oct. 2011, pp. 127–136.
- [75] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, "BundleFusion: Real-time globally consistent 3D reconstruction using on-the-fly surface reintegration," *ACM Trans. Graph.*, vol. 36, no. 4, p. 1, Aug. 2017.
- [76] T. Whelan, M. Kaess, H. Johannsson, M. Fallon, J. J. Leonard, and J. McDonald, "Real-time large-scale dense RGB-D SLAM with volumetric fusion," *Int. J. Robot. Res.*, vol. 34, nos. 4–5, pp. 598–626, Apr. 2015.
- [77] M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison, "CodeSLAM—Learning a compact, optimisable representation for dense visual SLAM," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2560–2568.
- [78] M. Landmann, S. Heist, P. Dietrich, P. Lutzke, I. Gebhart, J. Templin, P. Kühmstedt, A. Tünnermann, and G. Notni, "High-speed 3D thermography," *Opt. Lasers Eng.*, vol. 121, pp. 448–455, Oct. 2019.
- [79] J. Sun, H. Ma, and D. Zeng, "Three-dimensional infrared imaging method based on binocular stereo vision," *Opt. Eng.*, vol. 54, no. 10, Oct. 2015, Art. no. 103111.
- [80] The Open3D Development Team. (2023). *Open3D: A Modern Library for 3D Data Processing*. Version 0.13.0. [Online]. Available: <http://www.open3d.org/>
- [81] G. Farneback, "Two-frame motion estimation based on polynomial expansion," in *Proc. 13th Scand. Conf. (SCIA)*. Halmstad, Sweden: Springer, Jun. 2003, pp. 363–370.
- [82] J. M. Martínez-Otzeta, A. Ibarguren, A. Ansuategi, and L. Susperregi, "Laser based people following behaviour in an emergency environment," in *Proc. 2nd Int. Conf. Intell. Robot. Appl. (ICIRA)*. Singapore: Springer, Dec. 2009, pp. 33–42.
- [83] N. Bellotto and H. Hu, "A bank of unscented Kalman filters for multimodal human perception with mobile service robots," *Int. J. Social Robot.*, vol. 2, no. 2, pp. 121–136, Jun. 2010.
- [84] K. J. Cannons and R. P. Wildes, "The applicability of spatiotemporal oriented energy features to region tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 4, pp. 784–796, Apr. 2014.
- [85] R. Irani, K. Nasrollahi, and T. B. Moeslund, "Pain recognition using spatiotemporal oriented energy of facial muscles," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2015, pp. 80–87.
- [86] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [87] O. Sener, K. Ugur, and A. A. Alatan, "Efficient MRF energy propagation for video segmentation via bilateral filters," *IEEE Trans. Multimedia*, vol. 16, no. 5, pp. 1292–1302, Aug. 2014.
- [88] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–11.
- [89] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16 × 16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2021, pp. 1–21.
- [90] H. Xu and J. Zhang, "AANet: Adaptive aggregation network for efficient stereo matching," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2020, pp. 1956–1965.
- [91] X. Gu, Z. Fan, S. Zhu, Z. Dai, F. Tan, and P. Tan, "Cascade cost volume for high-resolution multi-view stereo and stereo matching," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 2492–2501.
- [92] Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, "MFNet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2017, pp. 5108–5115.
- [93] Y. Li, A. W. Yu, T. Meng, B. Caine, J. Ngiam, D. Peng, J. Shen, Y. Lu, D. Zhou, Q. V. Le, A. Yuille, and M. Tan, "DeepFusion: LiDAR-camera deep fusion for multi-modal 3D object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 17161–17170.
- [94] C. Hazirbas, L. Ma, C. Domokos, and D. Cremers, "FuseNet: Incorporating depth into semantic segmentation via fusion-based CNN architecture," in *Proc. 13th Asian Conf. Comput. Vis. (ACCV)*. Taipei, Taiwan: Springer, Nov. 2017, pp. 213–228.
- [95] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 3, pp. 583–596, Mar. 2015.
- [96] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2D pose estimation using part affinity fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1302–1310.
- [97] J. Vertens, J. Zürn, and W. Burgard, "HeatNet: Bridging the day-night domain gap in semantic segmentation with thermal images," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2020, pp. 8461–8468.
- [98] K. Xiang, K. Yang, and K. Wang, "Polarization-driven semantic segmentation via efficient attention-bridged fusion," *Opt. Exp.*, vol. 29, 4, pp. 4802–4820, 2021.
- [99] J. Zhang, K. Yang, and R. Stiefelhagen, "ISSAFE: Improving semantic segmentation in accidents by fusing event-based data," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 1132–1139.
- [100] H. Zhou, L. Qi, Z. Wan, H. Huang, and X. Yang, "RGB-D co-attention network for semantic segmentation," in *Proc. Asian Conf. Comput. Vis. (ACCV)*, 2020, pp. 519–536.
- [101] G. Zhang, J.-H. Xue, P. Xie, S. Yang, and G. Wang, "Non-local aggregation for RGB-D semantic segmentation," *IEEE Signal Process. Lett.*, vol. 28, pp. 658–662, 2021.
- [102] J. Jiang, L. Zheng, F. Luo, and Z. Zhang, "RedNet: Residual encoder-decoder network for indoor RGB-D semantic segmentation," 2018, *arXiv:1806.01054*.
- [103] S. S. Shivakumar, N. Rodrigues, A. Zhou, I. D. Miller, V. Kumar, and C. J. Taylor, "PST900: RGB-thermal calibration, dataset and segmentation network," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 9441–9447.
- [104] Y.-H. Kim, U. Shin, J. Park, and I. S. Kweon, "MS-UDA: Multi-spectral unsupervised domain adaptation for thermal image semantic segmentation," *IEEE Robot. Autom. Lett.*, vol. 6, no. 4, pp. 6497–6504, Oct. 2021.
- [105] C. Li, W. Xia, Y. Yan, B. Luo, and J. Tang, "Segmenting objects in day and night: Edge-conditioned CNN for thermal image semantic segmentation," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 7, pp. 3069–3082, Jul. 2021.
- [106] Q. Zhang, S. Zhao, Y. Luo, D. Zhang, N. Huang, and J. Han, "ABMDRNet: Adaptive-weighted bi-directional modality difference reduction network for RGB-T semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2633–2642.
- [107] J. Zhang, H. Liu, K. Yang, X. Hu, R. Liu, and R. Stiefelhagen, "CMX: Cross-modal fusion for RGB-X semantic segmentation with transformers," 2022, *arXiv:2203.04838*.

- [108] X. Hu, K. Yang, L. Fei, and K. Wang, "ACNet: Attention based network to exploit complementary features for RGBD semantic segmentation," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2019, pp. 1440–1444.
- [109] X. Chen, K.-Y. Lin, J. Wang, W. Wu, C. Qian, H. Li, and G. Zeng, "Bi-directional cross-modality feature propagation with separation-and-aggregation gate for RGB-D semantic segmentation," in *Proc. 16th Eur. Conf. Comput. Vis. (ECCV)*. Glasgow, U.K.: Springer, Aug. 2020, pp. 561–577.
- [110] M. Oršič, I. Krešo, P. Bevandic, and S. Šegvic, "In defense of pre-trained ImageNet architectures for real-time semantic segmentation of road-driving images," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 12599–12608.
- [111] L. Sun, K. Yang, X. Hu, W. Hu, and K. Wang, "Real-time fusion network for RGB-D semantic segmentation incorporating unexpected obstacle detection for road-driving images," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 5558–5565, Oct. 2020.
- [112] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [113] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6230–6239.
- [114] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7794–7803.
- [115] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. 18th Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*. Munich, Germany: Springer, Oct. 2015, pp. 234–241.
- [116] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [117] S. Gupta, J. Hoffman, and J. Malik, "Cross modal distillation for super-resolution transfer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2827–2836.
- [118] X. Wu, R. He, Z. Sun, and T. Tan, "A light CNN for deep face representation with noisy labels," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 11, pp. 2884–2896, Nov. 2018.
- [119] Y. Choi, N. Kim, S. Hwang, K. Park, J. S. Yoon, K. An, and I. S. Kweon, "KAIST multi-spectral day/night data set for autonomous and assisted driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 3, pp. 934–948, Mar. 2018.
- [120] Z. Mostaani, A. George, G. Heusch, D. Geissenbuhler, and S. Marcel, "The high-quality wide multi-channel attack (HQ-WMCA) database," Idiap, Martigny, Switzerland, Tech. Rep. Idiap-RR-22-2020, Sep. 2020.
- [121] A. Clapés, J. C. J. Junior, C. Morral, and S. Escalera, "ChaLearn LAP 2020 challenge on identity-preserved human detection: Dataset and results," in *Proc. 15th IEEE Int. Conf. Autom. Face Gesture Recognit. (FG)*, Nov. 2020, pp. 801–808.
- [122] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from RGBD images," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, vol. 7576, 2012, pp. 746–760.
- [123] S. Song, S. P. Lichtenberg, and J. Xiao, "SUN RGB-D: A RGB-D scene understanding benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 567–576.
- [124] I. Armeni, S. Sax, A. R. Zamir, and S. Savarese, "Joint 2D-3D-semantic data for indoor scene understanding," 2017, *arXiv:1702.01105*.
- [125] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3D reconstructions of indoor scenes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2432–2443.
- [126] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.
- [127] D. Gehrig, M. Rüegg, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, "Combining events and frames using recurrent asynchronous multimodal networks for monocular depth prediction," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2822–2829, Apr. 2021.
- [128] Y. Liao, J. Xie, and A. Geiger, "KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2D and 3D," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 3, pp. 3292–3310, Mar. 2023.
- [129] J. J. Moré, "The Levenberg–Marquardt algorithm: Implementation and theory," in *Numerical Analysis*. Berlin, Germany: Springer, 2006, pp. 105–116.
- [130] F. Dan Foresee and M. T. Hagan, "Gauss–Newton approximation to Bayesian learning," in *Proc. Int. Conf. Neural Netw. (ICNN)*, Jun. 1997, pp. 1930–1935.

MARTIN BRENNER received the B.Sc. degree in information technology and computing from The Open University, Milton Keynes, U.K., in 2008, and the M.Sc. degree in computer science from the Massey University, Auckland, New Zealand, in 2021, where he is currently pursuing the Ph.D. degree.

NAPOLEON H. REYES received the B.Sc. degree in physics and the M.Sc. and Ph.D. degrees (Hons.) in computer science from De La Salle University, Manila, Philippines, in 1993, 1999, and 2004, respectively. He is currently a Teacher with Massey University, Auckland, New Zealand. His research interests include evolving meta-level reasoning, intelligent systems, and deep learning.

TEO SUSNJAK received the Ph.D. degree in computer science with a focus on machine learning. He teaches data science topics at both bachelor's and master's levels with Massey University. He has broad industry experience, as a machine learning analyst. He continues to apply his research to solving real-world problems using data-driven approaches.

ANDRE L. C. BARCZAK received the B.Eng. and M.S. degrees in mechanical engineering from Unicamp, Brazil, in 1987 and 1996, respectively, and the Ph.D. degree in computer science from Massey University, New Zealand, in 2008. From 1989 to 1995, he was a System Engineer with IBM. He was a Research Assistant with Unicamp, from 1995 to 1998. In 1998, he emigrated to New Zealand and worked as a system engineer, until 2002. From 2002 to 2022, he was an Academic Staff of computer science with Massey University. Since 2023, he has been with Bond University, Australia. He has published more than 70 technical articles in journals and conferences. His research interests include computer vision, machine learning, big data, and optimization algorithms.

• • •