Development of Multi-robot 3D Thermal Modeling System for Built Environments

Jula Marzouk[1†], Kadin Sales[2†], Fanru Gao[2], Ci-Jyun Liang[2], and Jacob J. Lin[3]

[1]Department of Mechanical Engineering, Stony Brook University, USA
[2]Department of Civil Engineering, Stony Brook University, USA
[3]Department of Civil Engineering, National Taiwan University, Taiwan
[†] First and second authors contributed equally and are listed lexicographically

jula.marzouk@stonybrook.edu, kadin.sales@stonybrook.edu, fanru.gao@stonybrook.edu,
ci-jyun.liang@stonybrook.edu, jacoblin@ntu.edu.tw

Abstract -

Robotic thermal modeling deploys a single mobile robot to collect data and construct thermal point clouds, but such systems and processes are time-consuming, expensive, and intrusive to building operations. This paper introduces the preliminary development of a small multi-robot system that establishes the geometric and calibration foundations for three-dimensional (3D) thermal modeling of built environments. The system integrates three components: a 3D thermal modeling module that combines RGBD and thermal cameras; a point cloud generation module that constructs local and fused point clouds; and a multi-robot SLAM module that aligns point cloud data from multiple robots within a shared global reference frame. Experiments were conducted to evaluate intrinsic and cross-camera calibration of the thermal and RGBD cameras, as well as multi-point cloud reconstruction using two mobile robots. Results demonstrate that individual point clouds can be successfully aligned and merged into a consistent global point cloud, which indicates the feasibility of cooperative data collection using a group of small and low-cost robotic platforms. In summary, this work establishes a foundation for full multi-robot 3D thermal modeling. Future work will integrate thermal data into the global point cloud and evaluate in larger and more complex environments.

Keywords -

Multi-Robot System; 3D Thermal Modeling; Swarm SLAM; Camera Calibration

1 Introduction

Conventional building thermal inspections rely on 2D thermal images from handheld or stationary infrared cameras. While it is effective for small-scale qualitative assessment, 2D thermal analysis is viewpoint-dependent and lacks spatial context, which is difficult to locate thermal anomalies accurately, compare observations across viewpoints, or associate temperature measurements with specific building components [1]. In complex indoor environments, these limitations can lead to incomplete or ambiguous thermal analysis.

Three-dimensional (3D) thermal modeling addresses these limitations by embedding thermal measurements within 3D spatial representations, such as point clouds. 3D thermal modeling of built environments is crucial for energy efficiency and performance evaluation, occupant comfort analysis, and inspection. Accurate thermal models can help engineers and building managers identify the distribution of room temperatures, building pathologies, and structural delamination [1, 2], which can be used for building retrofitting and optimizing HVAC systems. Existing 3D thermal modeling methods rely on manual processes (e.g., stationary or handheld devices) that are time- and labor-intensive [3, 4]. Recent studies have applied a large mobile robot equipped with multiple sensors (e.g., cameras, depth cameras, laser scanners, and LiDAR) with sensor fusion methods to collect thermal and geometric data and construct 3D thermal models [2, 5, 6]. These robots navigate and localize themselves inside the built environment using Simultaneous Localization and Mapping (SLAM) [7]. However, deploying a single and large mobile robot to collect data is costly and time-consuming, and introduces safety hazards to occupants in built environments.

To address these challenges, we propose a multi-robot 3D thermal modeling system that uses a fleet of small, low-cost mobile robots to construct a 3D thermal point cloud. Multiple robots, or swarm robots, can operate in parallel to reduce data acquisition time, while the small-sized robot is less intrusive and poses lower safety risks to building occupants [8]. In this paper, we focus on the preliminary development of the three modules of the system: the 3D thermal camera module, the point cloud generation module, and the multi-robot SLAM module.

2 Related Work

2.1 Robot-based 3D Thermal Modeling

Robot-based 3D thermal modeling methods have demonstrated the feasibility of automatic acquisition of 3D thermal models across different built environments. Borrmann et al. [9] integrated terrestrial laser scanning and thermal imaging with 2D and 3D next-best-view planning algorithms and voxel-based 3D modeling method to generate thermal meshes. Adan et al. [5] developed a Mobile Platform for Autonomous Thermal Digitization (MoPATD), which is capable of omnidirectional 360°thermal point cloud reconstruction and automated multi-room navigation through a next-best-thermal-scan algorithm. Building on this, Adan et al. [2] presented a smaller and low-cost version (MoPATD2) that generates floor-scale omnidirectional thermal modeling with ROS-based navigation, autonomous docking, and a non-expert graphical user interface.

In addition to static indoor environments, complex urban environments and human-centered contexts are also studied with the use of robot-based thermal modeling. Lin et al. [10] developed a UAV-based multimodal data-fusion workflow and pose refinement method that combines RGB images with calibrated thermal images to produce 3D textured thermal models of building facades and rooftops. Cheng et al. [6] used a mobile robot equipped with RGBD and thermal cameras, real-time feature extraction, multi-view re-identification, and machine learning-based comfort estimation to dynamically map occupant distribution and thermal comfort. In summary, these works highlight the need for multi-robot deployment and cross-sensor calibration in robot-based 3D thermal modeling for built environments.

2.2 Multi-Robot SLAM

Multi-robot SLAM extends the single-robot SLAM by enabling a team of robots to estimate their trajectories and a shared map within a common reference frame [11]. In most systems, each robot executes its own SLAM pipeline locally, and communication between robots is used to align and fuse the individual maps. Compared with a single robot system, multi-robot SLAM offers broader spatial coverage, reduces exploration and data-collection time, and improves the robustness of state estimation through inter-robot loop closures. It also enhances overall system stability, which is particularly important in GPS-denied or hazardous environments such as subterranean mines, underground infrastructure, and active construction sites [12, 13].

Existing multi-robot SLAM systems can be divided into centralized and decentralized architectures. In a centralized architecture, the base station maintains a global map over all robots, performs loop-closure detection and outlier rejection, and optimizes the global map. The resulting map updates are then sent back to each robot. Chang et al. [14] introduced LAMP 2.0, a centralized multi-robot SLAM system that integrates a multi-robot front-end, 3D registration, and a robust pose-graph optimizer. This method enhances loop-closure reliability and achieves centimeter-level trajectory accuracy in real-world multi-robot deployments, while reducing the computational cost of global optimization by approximately an order of magnitude. Ebadi et al. [11] developed LAMP that incorporates an accurate LiDAR front-end, a vision front-end for artifact detection, and a pose graph optimization back-end equipped with an Incremental Consistent Measurement (ICM) module for outlier rejection. This system improves odometry accuracy by matching scans to sub-maps, reliably filters spurious loop closures, and demonstrates robust large-scale 3D mapping performance in subterranean environments. This design simplifies global consistency and enables the use of powerful computing resources and prior maps, but it depends on reliable communication and introduces a single point of failure [15].

In contrast, decentralized or distributed systems eliminate the need for a central server. Each robot maintains its own pose graph, detects inter-robot loop closures when encountering the others, and exchanges only selected measurements or compressed map data. Global consistency is achieved through peer-to-peer place recognition and distributed pose-graph optimization operating within explicit communication constraints. For example, Tian et al. [16] introduced Kimera-Multi, a SLAM system that performs distributed place recognition, robust inter-robot loop-closure verification, and distributed mesh reconstruction. This system enables scalable multi-robot mapping using peer-to-peer communication and demonstrates reliable trajectory estimation and consistent 3D mesh reconstruction across several real-world multi-robot experiments. In addition, Lajoie and Beltrame [17] proposed Swarm-SLAM, a sparse and decentralized collaborative SLAM framework that introduces a budgeted inter-robot loop-closure prioritization strategy based on algebraic connectivity maximization. Across multiple datasets and a three-robot field deployment, the spectral prioritization attains higher algebraic connectivity and lower absolute translation error with fewer loop closures than a greedy baseline. Decentralized multi-robot SLAM improves scalability and robustness by distributing computation and relying only on peer-to-peer communication. Its main challenge lies in maintaining global consistency under limited communication and decentralized data association [18].

Despite advances in robot 3D thermal modeling and multi-robot SLAM, there is still limited development between the two domains. Existing robot thermal modeling systems rely on single, large, and expensive platforms, limiting scalability, accessibility, and deployment feasibility in occupied buildings. Conversely, multi-robot SLAM has focused on geometric mapping and localization, without considering multimodal sensing requirements of thermal cameras, such as cross-sensor calibration.

3 Multi-Robot 3D Thermal Modeling System

3.1 System Overview

The proposed system consists of three modules: the 3D thermal camera module, the point cloud generation module, and the multi-robot SLAM module. Figure 1 illustrates the system architecture with two or more robots. Each robot's 3D thermal camera module is equipped with a thermal camera and an RGBD camera to construct a local thermal model (see Section 3.2). The depth camera from the RGBD camera can generate a local point cloud for each robot, and all local point clouds from different robots are fused to create a global point cloud (see Section 3.3). Finally, the pose of each robot can be estimated using the multi-robot SLAM module, which can be used to fuse the local thermal models to create a global thermal model (see Section 3.4). While decentralized multi-robot SLAM improves scalability and robustness, the implementation remains challenging. In this paper, we develop the point cloud generation module using centralized multi-robot SLAM as the preliminary work, and develop the multi-robot SLAM module using decentralized multi-robot SLAM to evaluate feasibility.

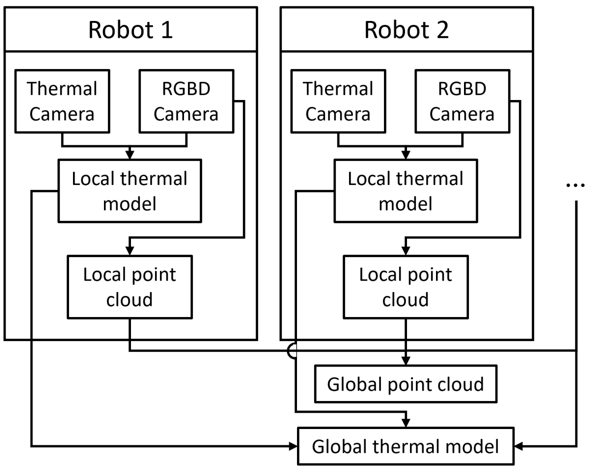

Figure 1. System architecture

3.2 3D Thermal Camera Module

The 3D thermal camera module consists of an RGBD camera and a thermal camera. Before data capture, the RGBD and thermal cameras need to be calibrated to ensure accurate data collection and comparison. In order to calibrate them, the RGBD camera is mounted on a camera stand. A camera mount was designed to connect the RGBD and thermal camera so that the relative distance between the cameras remains unchanged, which is crucial for the extrinsic and cross calibration. The mount was designed in Autodesk Fusion 360 and 3D printed in an effort to make it cost efficient. The 3D printed part, as well as the camera setup, are shown in Figure 2.

Figure 2. 3D printed part and camera setup

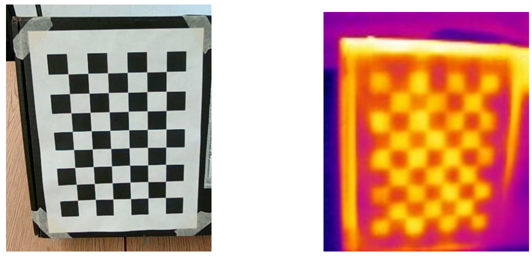

(a) RGB image (b) Thermal image

Figure 3. Chessboard used for camera calibration

The calibration process is divided into three major parts: intrinsic calibration of the RGB camera and thermal camera, and cross-calibration between both cameras [19]. For intrinsic calibration of the RGB camera, it was calibrated using a standard chessboard pattern printed on paper and taped to cardboard for stability, as shown in Figure 3(a). The camera detects the squares on the chessboard and computes the intrinsic matrix of the RGB camera. This matrix contains the focal lengths, principal points, and skew expressed in pixels. These values were collected using the camera calibration application in the MATLAB Computer Vision Toolbox, which was used for intrinsic calibration.

For intrinsic calibration of the thermal camera, it was calibrated using the same chessboard pattern but a different procedure. The chessboard was heated using a thermal lamp. Since the black squares of the pattern absorb more heat, a temperature gradient is observed in thermal camera images after heating (Figure 3(b)). As for the thermal camera, the intrinsic matrix was calculated using the same MATLAB camera calibration application.

For cross-calibration of the RGB and thermal cameras, the same chessboard was placed in the field of view of both cameras to allow simultaneous image capturing. Images were taken by setting up the checkerboard at different angles and positions. Using the stereo camera calibration application in the MATLAB Computer Vision Toolbox, the extrinsic parameters (translation and rotation) for the setup were obtained. For the depth camera, since it was already calibrated with the RGB camera, we will directly use the parameters provided by the manufacturer.

3.3 Point Cloud Generation Module

The point cloud generation module is responsible for constructing a 3D geometric representation of the environment, as well as the thermal point cloud, using spatial and thermal data obtained by each mobile robot. Each robot is equipped with a depth camera (the RGBD camera in the 3D thermal camera module) that captures range information as it navigates the indoor space. This depth data is transformed into a local point cloud expressed in the robot's coordinate frame [12]. The thermal and RGB data will also be transformed to robot's coordinate frame and merged with the local point cloud to generate the 3D thermal point cloud using the camera matrices from Section 3.2. While this thermal-to-point-cloud merging pipeline is designed, the experimental validation and deployment remain ongoing work.

The point clouds captured from each robot are sent to a central frame, where SLAM pose estimates are used to align them in a shared global frame [16, 20]. Once complete, the separate point clouds are merged to produce a combined point cloud representation of the environment. This centralized fusion approach enables parallel data collection across multiple robots, reducing overall mapping time compared to single-robot systems [8]. The detailed multi-robot SLAM algorithm is discussed in the next section.

3.4 Multi-Robot SLAM Module

The multi-robot SLAM module enables several robots to cooperatively build a consistent global map while each robot maintains its own local estimation process. Each robot incrementally constructs a pose graph following a keyframe-based strategy. A new keyframe is created when the robot's motion exceeds a predefined threshold. The relative constraint between consecutive keyframes is computed from the relative transform of the odometry estimates of the two keyframe timestamps, which captures the accumulated motion over the interval [21]. The associated point cloud is first filtered using an adaptive downsampling procedure to eliminate noise and maintain uniform density. Then it is stored as a local scan associated with the corresponding keyframe.

Incremental pose-graph updates are transmitted by each robot as new nodes or constraints are generated. These updates include keyframes, their estimated poses, and relative constraints. Local scans are transmitted as separate messages and are associated with keyframes through their shared key identifiers. Each unique key identifier ensures that data from different robots can be merged without conflict. A global process maintains a unified pose graph and integrates updates from all robots into a shared structure. As new information arrives, the system updates existing nodes as needed and inserts new constraints from the robots' local estimation processes. Once loop-closure constraints are incorporated, the global pose graph is optimized using a robust back-end based on Kimera-RPGO [20]. The optimizer performs a pairwise consistency assessment to evaluate loop-closure compatibility and employs a Graduated Non-Convexity (GNC) strategy [22] to adaptively adjust constraint weights during the optimization process. This approach mitigates the influence of inconsistent measurements while maintaining the geometric consistency of the pose graph.

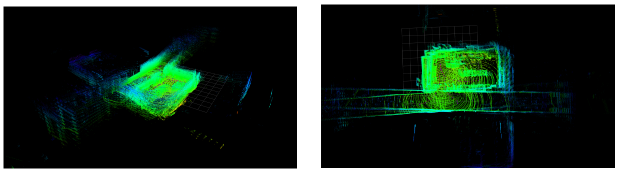

(a) ISO view (b) Top view

Figure 4. Multi-robot SLAM point cloud result

Based on the optimized pose estimates, a globally consistent map is reconstructed by expressing local scans from all robots in a shared reference frame. Subsequently, local scans are incorporated once the corresponding pose estimates become available. As a result, observations from all robots are consistently aligned and aggregated into a unified point cloud representation of the environment. By decoupling local pose graph construction, global data fusion, and robust back-end optimization into independent components, this module exhibits scalability and supports different sensing and odometry configurations. Figure 4 shows the result of the multi-robot SLAM point cloud.

4 Evaluation

4.1 Experiments

We evaluate the 3D thermal camera module and the point cloud generation module with two experiments. The first experiment is the camera calibration experiment for the 3D thermal camera module. The second experiment is the point cloud experiment for the point cloud generation module. Table 1 summarizes the hardware used in the experiment and implementation. Although thermal data acquisition was implemented at the sensor level, experimental results on thermal texture mapping and global 3D thermal reconstruction are outside the scope of this preliminary study.

4.1.1 Camera Calibration Experiment

For the camera calibration experiment, the cameras were first stationed in fixed positions so their parameters would not change throughout the experiment. Figure 2 shows that the camera module is mounted on a tripod at a fixed position. The RGB camera was used to capture images of the chessboard paper from multiple positions. Ten images

Table 1. Experiment setup and implementation

Module	Component	Hardware	Price
3D Thermal Camera	RGBD Camera	RealSense Depth Camera D455	$400
	Thermal Camera	FLIR Lepton 3.5	$280
Point Cloud Generation	Depth Camera	Astra Pro Plus depth camera	$150
	Robot	JetAuto Pro Robot Car	$1,350

(a) Front and back viewpoints

(b) Side robot viewpoints

Figure 5. Multi-robot point cloud data collection setup

(a) Mean reprojection error

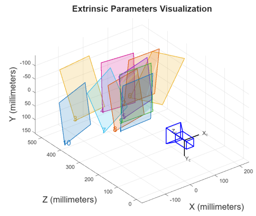

(b) chessboard locations

Figure 6. RGB camera calibration results

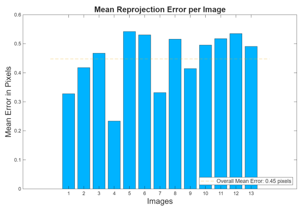

(a) Mean reprojection error

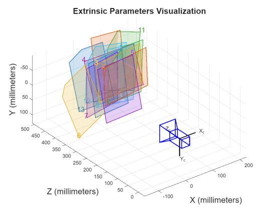

(b) chessboard locations

Figure 7. Thermal camera calibration results

were captured and imported into the MATLAB Computer Vision Toolbox to be used for calibration. The same procedure was repeated for the thermal camera except the chessboard was heated under the thermal lamp for a few minutes before capturing images to allow for the temperature gradient to occur. 13 images were captured and used for the thermal camera calibration. The intrinsic and extrinsic parameters were then computed for each camera. To cross-calibrate the RGB and thermal cameras, images were captured near simultaneously of the chessboard angled in different positions using both cameras. Nine pairs of images were captured and imported into the same program to be used for calibration.

4.1.2 Point Cloud Experiment

The point cloud experiment was designed to evaluate the ability of the multi-robot system to collect and combine spatial data into a single three-dimensional representation. The experiment was conducted in a controlled indoor environment using two JetAuto Pro mobile robots, as shown in Figure 5. A separate ground robot, different from the JetAuto Pro platforms, was set stationary in the lab and served as the target object for 3D mapping. This object was selected to provide a consistent geometric object for the point cloud generation between the two objects. The two JetAuto Pro robots were positioned approximately 2.5 meters from the stationary robot and set at eight distinct angles distributed around the target. Each JetAuto Pro collected spatial measurements from multiple perspectives, generating local point clouds that together provided coverage of the object from all sides.

Each local point cloud was sent to a central controller and aligned in a shared global reference frame using SLAM-based pose estimates [12]. Once received and aligned, the individual point clouds were merged to produce a combined point cloud representation of the stationary robot.

4.2 Camera Calibration Results

The following RGB camera intrinsic matrix (eq 1) is computed from the data gathered (in pixels) after calibrating using ten images. The overall mean reprojection error is 0.35 pixels, indicating good calibration accuracy. Figure 6 shows the RGB camera calibration results.

$$\begin{bmatrix} 674.7999 & 0 & 639.7857 \\ 0 & 678.9357 & 387.8659 \\ 0 & 0 & 1 \end{bmatrix} \tag{1}$$

The thermal camera intrinsic matrix (eq 2) is computed from the data gathered (in pixels) after calibrating using 13 images. The overall mean reprojection error is 0.45 pixels, which also indicates good calibration accuracy. Figure 7 shows the thermal camera calibration results.

$$\begin{bmatrix} 162.2003 & 0 & 73.7213 \\ 0 & 164.4104 & 58.9480 \\ 0 & 0 & 1 \end{bmatrix} \tag{2}$$

From the cross-calibration, the extrinsic parameters (rotation and translation vectors) of the RGB camera relative to the thermal camera can be obtained (eq 3 and eq 4). The overall mean reprojection error is 10.46 pixels. Although an estimate of the relative pose was obtained, the error indicates limited calibration accuracy. The observed high error in the RGB and thermal cameras' cross-calibration can be attributed to several factors. First, since the thermal camera and RGB camera applications were not connected, images had to be captured separately and quickly on each camera. A slight time offset (a few seconds) between RGB and thermal camera image captures could result in misaligned data due to subtle hand movements. Second, the chessboard pattern was reheated multiple times, which may have caused uneven or inconsistent temperature distribution across the squares. Finally, thermal cameras typically have lower resolution and higher noise levels compared to the RGB camera. In our implementation, the thermal camera resolution is 160×120 and the RGB camera resolution is 1280×800.

$$\text{Rotation vector} = \begin{bmatrix} -2.1085 & -1.8912 & -0.1084 \end{bmatrix} \tag{3}$$

Translation vector (in mm) =

$$\begin{bmatrix} -209.2981 & -2215.5433 & 534.9036 \end{bmatrix} \tag{4}$$

This level of cross-calibration reprojection error has important implications for thermal-point-cloud data alignment. In particular, inaccuracies in the estimated RGB–thermal extrinsic parameters can lead to spatial misalignment when projecting thermal data onto 3D point clouds. Such misalignment will result in blurred or shifted thermal patterns in a 3D thermal point cloud. In the multi-robot mapping context, these local misalignments may further accumulate during global point cloud fusion. Therefore, future work will focus on increasing calibration accuracy. Errors could be reduced by automating image capture on both camera apps to ensure that the images being used for cross-calibration are exactly the same and aligned. Automation also allows more images to be captured faster, therefore reducing the time needed for calibration. In addition, using an asymmetric chessboard, where one side is even and the other is odd, could increase accuracy as cameras may detect the orientation of symmetric boards incorrectly.

4.3 Point Clouds Results

Figure 5 shows the experimental setup in which two JetAuto Pro robots captured point cloud data from multiple viewpoints around a stationary ground robot positioned approximately 2.5 meters away. The viewpoints were distributed to provide coverage of the object from different angles under controlled laboratory conditions.

Figure 8 presents the resulting combined point cloud generated after transmitting the individual point clouds to a central controller and merging them using MATLAB. Figure 8(a) provides a top-down view of the reconstructed scene, while Figure 8(b) shows a side view, illustrating the vertical structure of the stationary robot and the alignment achieved across multiple viewpoints. The merged point cloud demonstrates that spatial data collected from multiple mobile robots can be successfully aligned within a shared global reference frame. Compared to individual captures, the combined representation offers improved geometric representation of the stationary object by integrating observations from multiple angles.

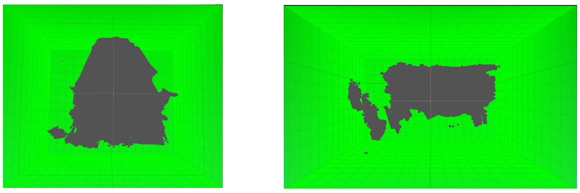

(a) Top view (b) Side view

Figure 8. Combined point cloud reconstructed from multi-view data collected by two mobile robots

This experiment demonstrates that the proposed framework can support spatial data collection from multiple small robots and successfully merge the resulting point clouds into a combined three-dimensional representation [6]. While the experiment focused on reconstructing a single object under controlled laboratory conditions, it provides a foundational success of the system's multi-robot communication, localization, and mapping capabilities. Future work will extend this approach to larger environments and incorporate thermal sensing for full 3D thermal mapping.

5 Discussion

The preliminary results of this study demonstrate the success of a multi-robot framework for 3D mapping in indoor environments and thermal modeling calibration. The system enabled small robots to communicate with a central controller and transmit data to generate a combined point cloud representation of the environment. These experiments, conducted in a controlled setting, used multi-robot SLAM, the core communication and mapping components of our work, consistent with prior work in multi-robot localization and mapping [12, 14, 16].

A central motive for this study is the opportunity to reduce costs and avoid causing disruption to the ongoing environment of a building. Previously, thermal inspections have relied on expensive single large robots (e.g., ~$26,000 based on the specification in [2]) that required slow control through an area [2, 5]. In comparison, our approach utilizes multiple small, low-cost robots (each robot costs ~$2,200, as listed in Table 1) that can operate together, reducing data collection time and lowering overall costs. The use of smaller robots also reduces physical risk and makes the system better suited to environments with a heavy presence. This design choice builds on previous findings on safety perception and trust in human-robot interaction in working environments [7, 8].

Although this study primarily focused on distributed spatial mapping and thermal camera module calibration, the main objective of the research is to integrate thermal sensing capabilities to enable 3D thermal representations of indoor environments. As shown in Section 4, thermal data collection was successful, but integration with the point cloud remains ongoing work. In addition, the benchmarking with existing robotic 3D thermal modeling is not included in this paper. For example, the thermal camera used in the proposal has lower resolution compared with other systems [2]. A more detailed benchmarking of the 3D thermal models will be conducted in future work.

6 Conclusion

This paper presents the preliminary development and evaluation of a multi-robot system for 3D thermal modeling of indoor built environments. The proposed framework integrates a 3D thermal camera module, a point cloud generation module, and a multi-robot SLAM module, which demonstrates the feasibility of using multiple small and low-cost mobile robots to cooperatively capture and reconstruct spatial and thermal data. A controlled indoor experiment was conducted to validate the ability to generate consistent merged point clouds from multiple robots and to calibrate the thermal and RGBD cameras. These results established a foundation for subsequent thermal data integration.

Although the current implementation remains at an early development stage, the results indicate significant potential to improve the efficiency, safety, and cost-effectiveness of indoor thermal inspection compared with existing single-robot or manual approaches. Future work will focus on integrating thermal and spatial data, refining cross-camera calibration, and improving coverage through multi-robot task allocation and path planning [23]. Additional experiments in more complex built environments will be conducted to evaluate robustness, scalability, and performance under real-world operational settings.

Overall, this study provides a strong foundation for a scalable and minimally intrusive approach to indoor mapping and thermal modeling. With further development, the proposed multi-robot thermal modeling system has the potential to support energy audits, fault detection, and performance monitoring in occupied buildings, enabling building owners and facility managers to identify thermal inefficiencies and implement energy-saving strategies without disrupting normal building operations.

References

[1] Sandra Pozzer, Gabriel Ramos, Parham Nooralishahi, Ehsan Rezazadeh Azar, Ahmed El Refai, Fernando López, Clemente Ibarra-Castanedo, and Xavier Maldague. Integration of thermographic inspection data with bim for enhanced concrete infrastructure assessment. *Automation in Construction*, 171:105965, 2025. doi:10.1016/j.autcon.2025.105965.

[2] Antonio Adan, Alejandro Lopez-Rey, and Amanda Ramon. Robot for thermal monitoring of buildings. *Automation in Construction*, 154:105009, 2023. doi:10.1016/j.autcon.2023.105009.

[3] Chao Wang, Yong K Cho, and Mengmeng Gai. As-is 3D thermal modeling for existing building envelopes using a hybrid LIDAR system. *Journal of Computing in Civil Engineering*, 27(6):645–656, 2013. doi:10.1061/(ASCE)CP.1943-5487.0000273.

[4] Jingwei Zhu, Yusheng Xu, Zhen Ye, Ludwig Hoegner, and Uwe Stilla. Fusion of urban 3d point clouds with thermal attributes using mls data and tir image sequences. *Infrared Physics & Technology*, 113: 103622, 2021. doi:10.1016/j.infrared.2020.103622.

[5] Antonio Adan, Samuel A Prieto, Blanca Quintana, Tomas Prado, and Juan Garcia. An autonomous thermal scanning system with which to obtain 3D thermal models of buildings. In *Advances in Informatics and Computing in Civil and Construction Engineering: Proceedings of the 35th CIB W78 2018 Conference: IT in Design, Construction, and Management*, pages 489–496. Springer, 2018. doi:10.1007/978-3-030-00220-6_58.

[6] Chenxi Cheng, Xiangtian Deng, Xiaoyong Zhao, Yuhan Xiong, and Yi Zhang. Multi-occupant dynamic thermal comfort monitoring robot system. *Building and Environment*, 234:110137, 2023. doi:10.1016/j.buildenv.2023.110137.

[7] Ci-Jyun Liang, Thai-Hoa Le, Youngjib Ham, Bharadwaj RK Mantha, Marvin H Cheng, and Jacob J Lin. Ethics of artificial intelligence and robotics in the architecture, engineering, and construction industry. *Automation in Construction*, 162:105369, 2024. doi:10.1016/j.autcon.2024.105369.

[8] Justin M Haney and Ci-Jyun Liang. A literature review on safety perception and trust during human–robot interaction with autonomous mobile robots that apply to industrial environments. *IISE transactions on occupational er-*

gonomics and human factors, 12(1-2):6–27, 2024. doi:10.1080/24725838.2023.2283537.

[9] Dorit Borrmann, Andreas Nuchter, Marija Dakulovic, Ivan Maurovic, Ivan Petrovic, Dinko Osmankovic, and Jasmin Velagic. A mobile robot based system for fully automated thermal 3D mapping. *Advanced Engineering Informatics*, 28 (4):425–440, 2014. doi:10.1016/j.aei.2014.06.002.

[10] Dong Lin, Na Yang, Qi Miao, Xiaojie Cui, and Dinggen Xu. True 3d thermal inspection of buildings using multimodal uav images. *Journal of Building Engineering*, 100:111806, 2025. doi:10.1016/j.jobe.2025.111806.

[11] Kamak Ebadi, Yun Chang, Matteo Palieri, Alex Stephens, Alex Hatteland, Eric Heiden, Abhishek Thakur, Nobuhiro Funabiki, Benjamin Morrell, Sally Wood, et al. LAMP: Large-scale autonomous mapping and positioning for exploration of perceptually-degraded subterranean environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 80–86, 2020. doi:10.1109/ICRA40945.2020.9197082.

[12] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2017. doi:10.1109/TRO.2016.2624754.

[13] Ricardo Hogan, Wei-Yi Hsu, Jacob J Lin, and Ci-Jyun Liang. BIM-driven unmanned ground vehicle swarm system for geometric and visual data collection. *Journal of Computing in Civil Engineering*, 39 (5):04025063, 2025. doi:10.1061/JCCEE5.CPENG-6619.

[14] Yun Chang, Kamak Ebadi, Christopher E Denniston, Muhammad Fadhil Ginting, Antoni Rosinol, Andrzej Reinke, Matteo Palieri, Jingnan Shi, Arghya Chatterjee, Benjamin Morrell, Ali-akbar Agha-mohammadi, and Luca Carlone. LAMP 2.0: A robust multi-robot slam system for operation in challenging large-scale underground environments. *IEEE Robotics and Automation Letters*, 7(4):9175–9182, 2022. doi:10.1109/LRA.2022.3191204.

[15] Josep Aulinas, Yvan Petillot, Joaquim Salvi, and Xavier Lladó. The SLAM problem: a survey. In *Proceedings of the 11th International Conference of the Catalan Association for Artificial Intelligence*, pages 363–371, 2008. doi:10.3233/978-1-58603-925-7-363.

[16] Yulun Tian, Yun Chang, Fernando Herrera Arias, Carlos Nieto-Granda, Jonathan P How, and Luca Carlone. Kimera-multi: Robust, distributed, dense metric-semantic slam for multi-robot systems. *IEEE Transactions on Robotics*, 38(4), 2022. doi:10.1109/TRO.2021.3137751.

[17] Pierre-Yves Lajoie and Giovanni Beltrame. Swarm-SLAM: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems. *IEEE Robotics and Automation Letters*, 9(1):475–482, 2023. doi:10.1109/LRA.2023.3333742.

[18] Miquel Kegeleirs, Giorgio Grisetti, and Mauro Birattari. Swarm SLAM: Challenges and perspectives. *Frontiers in Robotics and AI*, 8:618268, 2021. doi:10.3389/frobt.2021.618268.

[19] Akshay Iyer, Subacz Subacz, and Sai Vineeth K V. Finding a best-fit approach for 3D thermal modeling of indoor environments, 2025. URL `https://github.com/akshay-iyer/3D-Thermal-Modeling-`. Accessed: 2026-01-12.

[20] Antoni Rosinol, Marcus Abate, Yun Chang, and Luca Carlone. Kimera: an open-source library for real-time metric-semantic localization and mapping. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1689–1696. IEEE, 2020. doi:10.1109/ICRA40945.2020.9196885.

[21] Weifeng Chen, Xiyang Wang, Shanping Gao, Guangtao Shang, Chengjun Zhou, Zhenxiong Li, Chonghui Xu, and Kai Hu. Overview of multi-robot collaborative slam from the perspective of data fusion. *Machines*, 11(6):653, 2023. doi:10.3390/machines11060653.

[22] Heng Yang, Pasquale Antonante, Vasileios Tzoumas, and Luca Carlone. Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robotics and Automation Letters*, 5(2):1127–1134, 2020. doi:10.1109/LRA.2020.2965893.

[23] Zhaofeng Hu and Ci-Jyun Liang. LLM-PCPP: A novel robot prioritized coverage path planning approach for complex environments. In *Proceedings of the International Symposium on Automation and Robotics in Construction (ISARC)*, volume 42, pages 250–257. IAARC, 2025. doi:10.22260/ISARC2025/0034.