

Exosense: A Vision-Based Scene Understanding System For Exoskeletons

Jianeng Wang¹, Matias Mattamala¹, Christina Kassab¹, Guillaume Burger², Fabio Elnecave², Lintong Zhang¹, Marine Petriaux², and Maurice Fallon¹




Fig. 1: Exosense is a vision-based hardware and software system for scene understanding by exoskeletons. We developed a specialized multi-sensor unit (center), consisting of three global shutter wide-angle Alphasense cameras and Realsense D435i and D455 RGB-D units to provide 3D terrain and environment measurements. The hardware can be worn as a human-leg-mounted wearable device (left) or as an lower-limb attachment for an exoskeleton, such as the Wandercraft’s *Personal Exoskeleton* (right).

Abstract—Self-balancing exoskeletons are a key enabling technology for individuals with mobility impairments. While the current challenges focus on human-compliant hardware and control, unlocking their use for daily activities requires a scene perception system. In this work, we present *Exosense*, a vision-centric scene understanding system for self-balancing exoskeletons. We introduce a multi-sensor visual-inertial mapping device as well as a navigation stack for state estimation, terrain mapping, and long-term operation. We tested *Exosense* attached to both a human leg and Wandercraft’s *Personal Exoskeleton* in real-world indoor scenarios. This enabled us to test the system during typical periodic walking gaits, as well as future uses in multi-story environments. We demonstrate that *Exosense* can achieve an odometry drift of about 4 cm per meter traveled, and construct terrain maps under 1 cm average reconstruction error. It can also work in a visual localization mode in a previously mapped environment, providing a step towards long-term operation of exoskeletons.

Index Terms—Wearable Robotics, Prosthetics and Exoskeletons, RGB-D Perception, Mapping

I. INTRODUCTION

Recent advances in self-balancing exoskeletons, such as Wandercraft’s *Atalante* exoskeleton [1], are enabling individuals with lower-limb disabilities to walk independently without requiring additional support from crutches [2]. These

powered exoskeletons are being used in controlled clinical and therapeutic contexts [3]. However, the ultimate goal is to enable users to do everyday activities at home and outdoors.

Exoskeleton development has primarily focused on hardware and control challenges, aiming to design systems that can support and transport individuals while mimicking natural human walking. Many of these approaches employ control schemes with pre-defined gait trajectories [4]. This requires manual activation by an operator using a control panel, which increases the metabolic cost [5]. Integrating perception systems into the loop could reduce metabolic costs by partly automating low-level control tasks, such as switching gait modes, crossing doorways, or climbing stairs.

Vision sensing, because of its rich source of information [6], has been the primary sensor used to achieve this. Cameras have been used to estimate the semantic class of the terrain (namely stairs, ramps, and level ground walking) [7], [8], to determine basic geometric features such as ground inclination and step height [9], as well as to detect potential obstacles in the environment [10]. While these approaches are effective in providing the instantaneous information required for low-level decision making, they do not aim to integrate this information in long-term representations, which could be reused when revisiting environments.

In this work, we take initial steps towards developing long-term operation of self-balancing exoskeletons by presenting *Exosense*, a vision-centric scene understanding system. *Exosense* aims to generate home-scale, rich scene represen-

¹Oxford Robotics Institute, Dept. of Engineering Science, Uni. of Oxford, UK. {jianeng, matias, christina, lintong, mfallon}@robots.ox.ac.uk

²Wandercraft SAS, 5 Rue Pernelle, Paris, France. {guillaume.burger, fabio.elnecave, marine.petriaux}@wandercraft.health

tations from vision and geometry, capturing terrain structure, semantics, and traversability for localization and future navigation in previously visited environments. Our solution is primarily designed for self-balancing exoskeletons, such as the Wandercraft’s *Personal Exoskeleton* (Fig. 1), which aims to be the first self-balancing exoskeleton designed for domestic use. To achieve this, we introduce a versatile multi-sensor unit that can be attached to the exoskeleton’s leg or carried by a human as a leg-mounted wearable device. This placement allowed us to rigidly attach the Exosense to the exoskeleton, while also avoiding potential occlusions from the user’s body. Additionally, it enabled us to develop and test Exosense in realistic human walking scenarios and to seamlessly transfer the system to the exoskeleton given the similarities we observed in the gait dynamics (Fig. 3).

The key contributions of our work are:

- A versatile, leg-mounted multi-sensor unit that provides wide-angle vision and depth sensing for state estimation, terrain mapping, and localization.
- A scene understanding system that builds local maps embedded with the terrain geometry, room semantics, and traversability of indoor environments.
- A study of the performance of visual odometry for different camera configurations during typical walking patterns, where we obtained 4 cm drift per meter traveled for the selected odometry algorithm.
- Extensive experiments of the Exosense’s scene representation, particularly accuracy of the terrain reconstruction and traversable space estimates.
- A real-world demonstration of the Exosense integrated into the Wandercraft’s *Personal Exoskeleton* for indoor localization tasks, demonstrating the potential for future long-term operation in home environments.

II. RELATED WORK

We briefly review works on vision systems for wearable devices (Sec. II-A) and perception for self-balancing exoskeletons (Sec. II-B), which are relevant to Exosense.

A. Wearable Vision-based Systems

Vision-based systems provide richer information about the users and their surroundings than proprioceptive sensing, drawing increasing interest in wearable robotics research [11]. Integrating computer vision into upper-limb wearable robots has been commonly used in rehabilitation applications to assist object manipulation tasks including determining object dimensions [12] and detecting user intention [13]. These solutions are integrated as external egocentric cameras (e.g., mounted on glasses [14]) or directly attached to a robot [15]. In addition to manipulation tasks, vision-based wearables have enabled independent navigation for visually-impaired people [16]. In industrial settings, similar solutions have provided assistance to alleviate joint stress and to protect users from work-related musculoskeletal disorders [17].

Lower-limb exoskeletons mainly use vision to detect relevant ground features for reliable locomotion. Ramanathan *et*

al. [18] developed a vision-based perception system to enable exoskeletons to change the step size subject to detected obstacles. Follow-up work extended this approach to terrain recognition [19]. Tricomi *et al.* [20] used RGB images to identify types of terrain and adjust the walking controller of a hip exosuit. Karacan *et al.* [8] used a depth camera and an IMU mounted on a subject’s waist to classify objects (e.g., ramps, staircases, obstacles), and to predict staircase height. A similar setup has been deployed on a lower-limb prosthesis to recognize features in the environment [21].

In general, the aforementioned vision-based approaches focus on assisting short-horizon low-level decision making, without preserving historical data. With Exosense we instead aim to build a map representation that can be reused for future operations in indoor settings.

B. Perception for Self-balancing Exoskeletons

Existing self-balancing exoskeletons primarily rely on proprioceptive sensing for state estimation. Vigne *et al.* [22] incorporated multiple IMUs and robot joint encoder measurements to estimate position and velocity through a flexible kinematic model to account for deformations during exoskeleton walking. MOVIE [23] fused the same sensor measurements and takes a velocity-aided approach to estimate the robot’s orientation with respect to gravity. Elneceve *et al.* [24] built on top of MOVIE to estimate the 6 DoF pose and velocity of the exoskeleton’s body using an EKF.

Self-balancing exoskeletons also incorporate proprioceptive sensing into the control scheme. Tian *et al.* [25] used motor force sensors to estimate the center of mass deformation and physical parameters of the human operator during exoskeleton walking, which were then fed into a joint control framework to help the robot walk stably. Li *et al.* [26] used IMUs to estimate both robot and human center of mass; these estimates were integrated into a human-in-the-loop cooperative control scheme to adaptively adjust the motion controller to follow the user’s intention.

In contrast to previous work that is tightly tailored to the exoskeleton platform, we developed an integrated sensing unit that relies on exteroceptive sensing only, hence being independent of the particular platform, user gait, or control framework.

III. SYSTEM

The Exosense system overview is shown in Fig. 2. It is a scene understanding system which involves a leg-mounted vision-based sensing unit, and a navigation stack designed for highly dynamic walking motions. The system generates an environment representation encoding the terrain geometry, semantics, and traversability that can enable the continuous deployment of the exoskeleton for localization and navigation. The following sections describe the main components, from the hardware design to the navigation stack.

A. Multi-sensor Setup

The Exosense hardware consists of a lower-limb-mounted wearable device, shown in Fig. 1. The sensing device in-




Fig. 2: Exosense scene understanding system. The inputs are RGB images and 3D point clouds from the multi-sensor unit. Different modules provide odometry and sensor-rate terrain maps, which are integrated into a semantic pose graph and processed with a terrain reconstruction module. The final scene representation (bottom right) contains terrain geometry, semantics, and traversability as well as visual localization information to aid long-term operation.

cludes a Sevensense Alphasense unit with three hardware-synchronized wide-angle global shutter cameras plus an inertial sensor, and two Realsense RGB-D cameras (D435i and D455). The Alphasense unit is used mainly for state estimation (odometry estimation and localization) due to its wide field-of-view (FoV). The Realsense cameras provide 3D sensing for terrain mapping.

The multi-sensor unit can be attached to either a human thigh or an exoskeleton. We developed the sensing system to only require exteroceptive sensing so as to disregard challenges related to leg deformation and bending. On the exoskeleton, the device is rigidly attached to the thigh to avoid occlusions with the rest of the user’s body and to facilitate its future integration with the exoskeleton’s walking controller.

Because the device can also be attached to a human leg, we could develop the scene understanding system without needing permanent access to an exoskeleton—enabling us to test algorithms that go beyond the current capabilities of the exoskeleton, such as multi-floor navigation. This design decision is supported by the similarities we observed in the walking motions of the human-leg-mounted and exoskeleton-mounted sequences, as illustrated in Fig. 3.

B. Visual-Inertial Odometry

To estimate ego-motion, we use a visual-inertial odometry system to provide high-frequency state estimation and to handle high rotation rates and jerk during the exoskeleton locomotion. We considered OpenVINS [27], VILENS-MC [28] and ORB-SLAM [29]. We evaluated their performance in custom sequences recorded with our multi-sensor unit to assess their performance and reliability under walking




Fig. 3: Sample of linear acceleration and angular rotation rates measured by Exosense in exoskeleton (top) and human-leg-mounted (bottom) configurations. Both modes have a similar gait duration. The highlighted spikes (pink) occur during foot strikes.

patterns, which is discussed in Sec. IV-Exp B. For our experiments, we chose OpenVINS due to its better balance of estimation accuracy and computational cost.

C. Semantic Pose-graph SLAM

The odometry estimate serves as input to a visual SLAM system. We used LEXIS [30], as it provides a semantic pose graph representation that can be easily extended with other information sources, such as terrain maps. LEXIS constructs a pose graph representation with evenly spaced nodes. The nodes store corresponding RGB images, which are used for visual localization and loop closure detection. For semantics, LEXIS uses the CLIP visual-language model [31] to obtain visual embeddings, which are compared against text embeddings of a list of room classes (e.g., office, kitchen, corridor), providing potential room labels for each node in the graph.

The room labels associated to each node enable hierarchical place recognition by comparing the predicted room class of the current image to the keyframe labels in the graph. PnP [32] is used as a geometric verification step to propose loop closure candidates, which are jointly optimized in a factor graph to reduce the drift in the graph.

The output of LEXIS is a pose graph encoding odometry and loop closure connectivity with a pose-level room segmentation. In Exosense we extend this representation by adding terrain maps to each node, which can be refined by exploiting the room information. This enables us to produce an *elastic* globally-consistent terrain representation defined by the pose graph, as proposed in the Atlas framework [33].

D. Mapping and Reconstruction

To obtain the local terrain maps, we use the method by Fankhauser *et al.* [34], which integrates point clouds from both RGB-D cameras at 15 Hz, to generate a rolling local multi-layered 2.5D map [35] at 2 cm resolution around the Exosense’s sensing unit location. We used this method as a server, providing local terrain maps to be attached to LEXIS’ pose graph nodes on request.

While this representation enables a lightweight elastic terrain reconstruction, the individual submaps only represent a local region around the corresponding node, which might be suboptimal due to moving objects and partial visibility.

Hence, we propose to exploit the semantic room information already stored in the pose graph to refine the terrain estimates, creating single *room-based* terrain maps.

For submaps within the same room, we fuse the height values of overlapping map cells using the median:

$$h_i^{\text{merged}} = \text{median}(\{h_i\}), \quad (1)$$

where h_i^{merged} is the fused height value for the i^{th} cell of terrain map, $\{h_i\}$ is the set of all the valid overlapping height values in cell i . While other submap fusion strategies could be used, here we exploit the room information already provided by LEXIS, which provides a semantic guidance.

E. Terrain Traversability Analysis

Following the room-based fusion step, we wish to estimate the traversability of each room's terrain map. This traversability estimate is computed on a cell basis, characterizing which areas of each room should be accessible by the exoskeleton in a navigation setting. To obtain it, we perform a geometric analysis of the local terrain tailored to the exoskeleton's gait specifications.

Technically, the local terrain analysis module determines a traversability score for a cell i , $t_i \in [0, 1]$, which characterizes how difficult it would be for the exoskeleton to step on the cell (specifically, 0 for untraversable and 1 for traversable). For this, we assume that the exoskeleton has a nominal maximum stride length s^* (size of step forward) and step height h^* (maximum height it can step on). For each cell i with height h_i , we select the neighboring cells j within a radius s^* , denoted by $\mathcal{C}_{s^*}$, and compute the maximum height difference $h_i^{\max}$ in the neighborhood:

$$h_i^{\max} = \max(|h_j - h_i|), j \in \mathcal{C}_{s^*}. \quad (2)$$

We then define the traversability score of a cell as the percentage difference of the maximum height difference $h_i^{\max}$ with respect to the nominal step height h^* :

$$t_i = 1 - \min\left(\frac{h_i^{\max}}{h^*}, 1\right). \quad (3)$$

This traversability score then represents a conservative cell estimate of how safe it would be for the exoskeleton to step into any other cell given this nominal step height. Fig. 8 illustrates how this compares to a geometric approach based on surface normals on a staircase.

F. Localization within the Scene Representation

Exosense generates a scene representation that includes terrain geometry, semantics, and traversability. We extend it with images and depth maps obtained at each topological map node. This enables us to perform place recognition and metric localization in subsequent missions, enabling the reuse of previously built maps.

Our localization approach uses visual bags of words [36] to obtain place candidates, and PnP [32] to obtain a metric pose estimate from the RGB and depth images. We combine this with the odometry estimate to provide a continuous pose estimate between relocalizations, as well as providing an

estimate when the exoskeleton visits an unmapped area. We demonstrate this in Sec. IV-Exp F.

IV. EXPERIMENTS

We conducted several experiments to validate the Exosense's multi-sensor unit and navigation pipeline for indoor exoskeleton applications. We collected two datasets using the Exosense multi-sensor unit in two indoor environments:

- *Human* – Four sequences recorded with the Exosense device mounted on a human leg (Fig. 1, left): Sequences with 2-minute duration *H1*, *H2*, *H3* were captured using a Vicon motion capture system in a research lab, while *H4* was 10-minute long and recorded in a multi-floor office environment. The objective was to evaluate our design before testing on the exoskeleton, as well as testing *Exosense* mapping capabilities which go beyond the current exoskeleton locomotion capabilities.
- *Exo* – Two 7-minute sequences recorded with the Exosense device attached to the thigh of Wandercraft's *Personal Exoskeleton* (Fig. 1, right): *E1* was used to evaluate the terrain reconstruction, while *E2* was used to demonstrate the localization capabilities of Exosense. The exoskeleton was teleoperated while carrying a dummy during the recordings in a mixed office and lab environment with occasional passers-by. The objective of this dataset was to assess the mapped terrain quality in realistic conditions, as well as the potential of Exosense for indoor navigation.

All data was post-processed with a mid-range laptop, Intel i7 10750H @ 2.60Hz 12 core laptop, Nvidia GTX 1650Ti GPU. All the algorithms are CPU-based except for the CLIP feature extractor in LEXIS.

Our first two experiments, Sec. IV-Exp A and Sec. IV-Exp B, aim to assess our hardware and odometry estimators decisions prior to the deployment of Exosense on the exoskeleton—hence they are demonstrated with the human-leg mounted mode. We additionally show the potential of the full pipeline to operate in multi-story environments in Sec. IV-Exp C. The last three experiments assess the reconstruction accuracy, traversability quality, and demonstrate a localization use case in the *Exo* sequences.

Exp A. Study of Wide FoV Multi-Camera Systems [Human]

Our first experiment tested the suitability of our multi-camera setup. We achieved this by comparing the performance of a visual-inertial odometry system, for different numbers of cameras and different FoV in sequence *H1*. We used VILENS-MC [28], a fixed-lag optimization-based visual-inertial odometry algorithm designed in our group which works with multi-camera systems. The Exosense sensing unit was carried by a human as a wearable device on the person's thigh, while mimicking the exoskeleton walking pattern. Both wide and narrow field-of-view (FoV) images were recorded, and a Vicon motion capture system was used to provide ground truth poses. When evaluating the odometry performance, we report the Root Mean Square Error (RMSE) of the Relative Pose Error (RPE) as our main metric.

TABLE I: Exp A - Translation and rotation RPE at 1m, averaged over five runs under different camera configurations.

Relative Pose Error (RPE) – Translation [m] / Rotation [°]			
FoV (H × V)	No. Camera	Translation RMSE [m]	Rotation RMSE [°]
64° × 90°	2	Fail	Fail
64° × 90°	3	0.61	3.57
92.4° × 126°	2	0.34	2.38
92.4° × 126°	3	0.11	2.70

Tab. I reports the main results of this experiment. We observe that using a small FoV camera results in significant drift or even motion tracking failure due to the high accelerations and jerks present in the walking motion. Wider FoV cameras help to mitigate these effects, significantly reducing drift. Adding the lateral camera mitigates situations where no features are detected in front of the device. The lowest drift rates are achieved using both wide FoV cameras and the multi-camera setup with forward and lateral views. This configuration achieved reliable motion tracking, even in scenarios with significant viewpoint changes or occlusions under the jerky walking motion.

Exp B. Comparison of VI Odometry Algorithms [Human]

Next, we extended the evaluation of odometry estimators for Exosense to other open source algorithms, using the wide FoV and 3-camera configuration determined in the previous experiment. The objective was to assess the performance of other methods in these challenging walking conditions. We compared VILENS-MC to ORB-SLAM (optimization-based) [29] and OpenVINS (filtering-based) [27]. We must note that for ORB-SLAM we used the stereo-inertial configuration as it does not support multi-camera setups; we also disabled loop closure mechanisms for a fair comparison with the odometry systems. Further, OpenVINS processes each camera as a monocular input, while VILENS-MC treats the three cameras as a stereo pair and a monocular camera.

We used sequences *H2* and *H3* to test the performance of these systems in new conditions not considered in *H1*. In *H2*, the operator walked at a slow pace in a loop that included a small staircase. The sequence had a peak linear acceleration of 37.6 m/s² and rotation rate of 4.4 rad/s. In *H3* we included periods of abrupt rotation change and occasional occlusions of the front stereo cameras. This sequence had peak acceleration and rotation rates of 55.1 m/s² and 5.0 rad/s respectively. The algorithms were run five times on each sequence. RPE at 1 m and 5 m are presented in Tab. II.

Additionally, to provide further insights on the computational budget required by each method, we logged the CPU usage of the algorithms. These results are presented in Fig. 4.

Our odometry evaluation results show that the three tested VIO algorithms are robust and reliable even during walking patterns. This follows our design decision to use wide FoV cameras as discussed in Sec. IV-Exp A. ORB-SLAM experiences higher drift rates at times as it lacks multi-camera support, which we also noted in a previous paper [28].




Fig. 4: Exp B – CPU usage over time for the evaluated odometry algorithms over five runs. The darker lines show the mean, while the shaded areas are the 95% confidence interval. OpenVINS is significantly more lightweight—using about half the computation.

TABLE II: Exp B - Translation and rotation RPE (at 1 m and 5 m) for each odometry algorithm, averaged over five runs.

Relative Pose Error (RPE) – Translation [m] / Rotation [°]					
Dist	Method	Seq. H2		Seq. H3	
		Translation RMSE	Rotation RMSE	Translation RMSE	Rotation RMSE
1m	ORB-SLAM	0.08	3.14	0.19	2.49
	OpenVINS	0.06	3.45	0.15	2.55
	VILENS-MC	0.10	3.62	0.14	2.55
5m	ORB-SLAM	0.27	3.62	0.58	3.56
	OpenVINS	0.20	3.40	0.40	4.10
	VILENS-MC	0.26	3.20	0.38	4.25

While OpenVINS and VILENS-MC achieve comparable odometry accuracy (20 cm RPE at 5 m, i.e. 4 cm drift per meter traveled), OpenVINS uses less computation (i.e., under 10 % CPU usage). As a result, we chose OpenVINS as the odometry source for Exosense for our next experiments. In future we envisage these algorithms being run on low power hardware such as an ASIC or FPGA chip.

Exp C. Multi-story Mapping with Exosense [Human]

The last experiment in the human-leg mounted mode tested the full Exosense navigation pipeline for sequence *H4*, featuring a multi-story building. This sequence shows the potential of Exosense to build multi-floor representations that are beyond the current locomotion capabilities of self-balancing exoskeletons on multi-step staircases.

Fig. 5 demonstrates the terrain reconstruction of a multi-story office environment carrying the Exosense unit on a human leg. The system handles elevation maps in the multi-floor scenario and merges them by semantic room labels, which yields a smooth terrain map. We demonstrate the detailed view of multi-floor mapping in the attached multi-media material and present the quantitative reconstruction evaluation in Sec. IV-Exp D.

Exp D. Evaluation of Terrain Reconstruction Quality [Both]

We evaluated the Exosense terrain reconstruction quality in both the exoskeleton and human-leg mounted modes. For this experiment we focused on staircases present in




Fig. 5: Exp C – Multi-story mapping of sequence *H4* in the *Home* dataset. Exosense generated a globally consistent multi-floor terrain map. Each room is a single individual submap colored by its type.




Fig. 6: Exp D – Qualitative mapping result after submap merging of the *Exo* sequence, colored by the elevation. Staircases and part of the ground areas are shown in detail both before and after applying submap merging (bottom). The median-based merging method removed outliers in the terrain submap while preserving the sharp features and edges of terrain geometry.

sequences *H4* (*Human*) and *E1* (*Exo*), using millimeter-accurate reconstructions from tripod-based laser scanners, Leica RTC360 and Faro Focus 3D-X130, respectively. We ran Exosense with the same setup used in Sec. IV-Exp C.

Fig. 6 shows the reconstruction of the *Exo* sequence colored by elevation. We also show maps before and after applying the submap fusion strategy introduced in Sec. III-D. We observed improvements in the terrain flatness, and better crispness at the edges of the staircase steps. Furthermore, we note that outliers from odometry drift and dynamic objects, though not explicitly modeled in the individual submaps, are mitigated by our fusion strategy, resulting in a consistent reconstruction of the terrain.

Further, we performed a quantitative evaluation against the laser scans, by extracting key areas from the elevation map that the exoskeleton could traverse (e.g., staircases). We cropped these regions and converted the terrain maps into meshes to preserve the geometry of the terrain, and then sampled 10000 points per square meter from meshes to compute point-to-point distances to the ground truth scans.

We present the results in Tab. III. The mapping results under both mounting modes showed similar accuracy, indicating the design choices made using the human-leg mounted mode can be successfully transferred to the exoskeleton-mounted mode. For the exoskeleton-mounted mode, Exosense produced an average error under 1 cm with 90th percentile error under 2 cm, indicating a sensible reconstruction quality for the use with the exoskeleton. Fig. 7 additionally shows the error distribution for staircases present in the sequences. The large errors mainly appear on the vertical

TABLE III: Exp D - Terrain Reconstruction Quality. We measured the point-to-point distance between our terrain reconstruction and the ground truth point cloud scan. We present the mean, max and 90th percentile error to quantify the mapping quality.

Point-to-point Distance of Staircase Reconstruction				
dataset	res. [cm]	mean [cm]	max [cm]	90% [cm]
<i>Human</i>	2	1.36	8.45	2.84
<i>Exo</i>	2	0.80	6.64	1.85




Fig. 7: Exp D – Maps of two areas of interest from the Exosense system in (a) exoskeleton mounted and (b) human-leg mounted modes. The mapping results are compared with ground truth laser scans and colored by the point-to-point distance. Errors in the near-vertical surfaces of the terrain map should be ignored.

areas, which is expected from an elevation-based terrain reconstruction method and the chosen evaluation procedure.

Exp E. Evaluation of Terrain Traversability [*Exo*]

Next, we assessed the traversability estimation result of the Exosense system. In Fig. 8, we present an example of per-cell traversability obtained from the mapping result in sequence *E1*: On the left, we show a baseline based on surface normals [37], while the right is our proposed step height-based method. We set our method with a maximum stride length of 20 cm and nominal step height of 20 cm. We observed that it correctly assigned high traversability scores to the riser and treads of the staircase, which reflects the effective traversable areas of a walking system. In contrast, methods based on surface normals correctly determine walls to be untraversable, but may assign low traversability scores to the risers—which is undesired for navigation tasks.

To quantitatively show the benefits of our approach, we stored the per-cell traversability predictions and estimated elevation, and hand-labeled the traversable and untraversable areas. We then binarized the traversability scores of the normals-based method and ours with different thresholds. Then we evaluated the quality of the predictions as a classification problem, where the traversable areas were positive and the untraversable areas negative. We computed precision, recall, and F-score values over different thresholds to evaluate the traversability estimation accuracy. Our method obtained F-scores above 0.9 for different threshold values, with an optimal threshold of 0.5 (F-score value 0.93). In contrast, the normals-based method reported lower F-scores (below 0.87) for all threshold values and was more sensitive to changes in the optimal traversability threshold.




Fig. 8: Exp E – Comparison between traversability analysis obtained based on (a) terrain normals, and (b) our method based on the exoskeleton’s step height. We observe that for the same traversability range, our method assigns a higher traversability score to the complete staircase compared to the normals-based method.

Exp F. Localization Demonstration [Exo]

Finally, we demonstrated the ability of the Exosense to visually relocalize within a prior map. To achieve this, we recorded an initial sequence using the leg mounted configuration, and then used the localization mode in the Exo-mounted configuration (sequence E2).

Fig. 9 shows the prior map built with Exosense in human-leg-mounted mode. The path followed by the exoskeleton in a subsequent experiment is shown in green. The exoskeleton was teleoperated to walk around 80m, with our system obtaining about 60 visual relocalizations (about one every 0.8m on average, shown as orange triangles). Our localization system was able to correct the estimate when the exoskeleton returns back along the corridor (a) and when passing through the narrow doorway (c). We observed that the odometry estimate was generally reliable when moving through unmapped regions (b), and the system was able to relocalize when returning to previously visited areas.

V. CONCLUSIONS

We introduced Exosense, a vision-based scene understanding system for self-balancing exoskeletons. Our system consists of a multi-sensor unit and a navigation stack, designed to be independent of the exoskeleton hardware and also usable as a wearable device. We investigated the hardware and dynamics of the problem, concluding that a visual-inertial unit with wide-angle cameras overcomes most of the challenges of the walking motion. We further introduced a mapping pipeline able to capture accurate terrain structure, semantics and traversability, as well as demonstrated how Exosense can relocalize in previously visited places. This provides input on the advantages of exteroceptive sensing for the eventual deployment of exoskeletons in indoor environments. In future work, we aim to extend the applicability of Exosense for long-term, multi-session localization and mapping applications under environment changes, explore its usage in outdoor scenarios, and further integrate the system into the exoskeleton’s navigation and control stack additionally exploiting its proprioceptive sensing.

ACKNOWLEDGMENT

This work is supported by a Royal Society University Research Fellowship (Fallon, Kassab), Horizon Europe project DigiForest 101070405 (Wang), and EPSRC C2C Grant

EP/Z531212/1 (Mattamala). We thank Wayne Tubby and Matthew Graham for hardware design support.

REFERENCES

- [1] T. Gurriet, S. Finet, G. Boeris, A. Duburcq, A. Hereid, O. Harib, M. Masselin, J. Grizzle, and A. D. Ames, “Towards restoring locomotion for paraplegics: Realizing dynamically stable walking on exoskeletons,” in *IEEE Int. Conf. Robot. Autom.*, 2018.
- [2] V. Huynh, G. Burger, Q. V. Dang, R. Pelgé, G. Boéris, J. W. Grizzle, A. D. Ames, and M. Masselin, “Versatile dynamic motion generation framework: Demonstration with a crutch-less exoskeleton on real-life obstacles at the cybathlon 2020 with a complete paraplegic person,” *Front. Robot. AI*, vol. 8, 2021.
- [3] D. Tian, W. Li, J. Li, F. Li, Z. Chen, Y. He, J. Sun, and X. Wu, “Self-balancing exoskeleton robots designed to facilitate multiple rehabilitation training movements,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 32, pp. 293–303, 2024.
- [4] T. Yan, M. Cempini, C. M. Oddo, and N. Vitiello, “Review of assistive strategies in powered lower-limb orthoses and exoskeletons,” *Robot. Auton. Syst.*, vol. 64, pp. 120–136, 2015.
- [5] M. Kim, H. Jeong, P. Kantharaju, D. Yoo, M. Jacobson, D. Shin, C. Han, and J. Patton, “Visual guidance can help with the use of a robotic exoskeleton during human walking,” *Sci. Rep.*, vol. 12, 03 2022.
- [6] A. H. Al-dabbagh and R. Ronsse, “A review of terrain detection systems for applications in locomotion assistance,” *Robot. Auton. Syst.*, vol. 133, p. 103628, 2020.
- [7] A. G. Kurbis, B. Laschowski, and A. Mihailidis, “Stair recognition for robotic exoskeleton control using computer vision and deep learning,” in *Int. Conf. Rehabil. Robot.*, 2022.
- [8] K. Karacan, J. T. Meyer, H. I. Bozma, R. Gassert, and E. Samur, “An environment recognition and parameterization system for shared-control of a powered lower-limb exoskeleton,” in *IEEE RAS EMBS Int. Conf. Biomed. Robot. Biomechanics*, 2020.
- [9] A. H. A. Al-Dabbagh and R. Ronsse, “Depth vision-based terrain detection algorithm during human locomotion,” *IEEE Trans. Med. Robot. Bionics.*, vol. 4, no. 4, pp. 1010–1021, 2022.
- [10] D.-X. Liu, J. Xu, C. Chen, X. Long, D. Tao, and X. Wu, “Vision-assisted autonomous lower-limb exoskeleton robot,” *IEEE Trans. Syst. Man Cybern. Syst.*, vol. 51, no. 6, pp. 3759–3770, 2021.
- [11] L. Gionfrida, D. Kim, D. Scaramuzza, D. Farina, and R. D. Howe, “Wearable robots for the real world need vision,” *Sci. Robot.*, vol. 9, no. 90, p. ead8812, 2024.
- [12] C. Hu, D. Kim, S. Luo, and L. Gionfrida, “Pointgrasp: Point cloud-based grasping for tendon-driven wearable robotic applications,” in *IEEE Int. Conf. Robot. Autom.*, 2024.
- [13] E. Rho, H. Lee, Y. Lee, K.-D. Lee, J. Mun, M. Kim, D. Kim, H.-S. Park, and S. Jo, “Multiple hand posture rehabilitation system using vision-based intention detection and soft-robotic glove,” *IEEE Trans. Ind. Inform.*, vol. 20, no. 4, pp. 6499–6509, 2024.
- [14] D. Kim, B. B. Kang, K. B. Kim, H. Choi, J. Ha, K.-J. Cho, and S. Jo, “Eyes are faster than hands: A soft wearable robot learns user intention from the egocentric view,” *Sci. Robot.*, vol. 4, no. 26, p. eaav2949, 2019.
- [15] J. Kuhn, J. Ringwald, M. Schappler, L. Johannsmeier, and S. Haddadin, “Towards semi-autonomous and soft-robotics enabled upper-limb exoprosthetics: First concepts and robot-based emulation prototype,” in *IEEE Int. Conf. Robot. Autom.*, 2019.
- [16] H.-C. Wang, R. K. Katzschnmann, S. Teng, B. Araki, L. Giarré, and D. Rus, “Enabling independent navigation for visually impaired people through a wearable vision-based feedback system,” in *IEEE Int. Conf. Robot. Autom.*, 2017.
- [17] F. Missiroli, P. Mazzoni, N. Lotti, E. Tricomi, F. Braghin, L. Roveda, and L. Masia, “Integrating computer vision in exosuits for adaptive support and reduced muscle strain in industrial environments,” *IEEE Robot. Autom. Lett.*, vol. 9, no. 1, pp. 859–866, 2024.
- [18] M. Ramanathan, L. Luo, J. K. Er, M. J. Foo, C. H. Chiam, L. Li, W. Y. Yau, and W. T. Ang, “Visual environment perception for obstacle detection and crossing of lower-limb exoskeletons,” in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022.
- [19] M. Ramanathan, L. Luo, M. J. Foo, C. H. Chiam, W.-Y. Yau, and W. T. Ang, “Heuristic vision based terrain recognition for lower limb exoskeletons,” in *International Convention on Rehabilitation Engineering and Assistive Technology*, 2024, p. 25–28.




Fig. 9: Exp F – Exoskeleton-mounted localization demonstration within a prior map. The green line indicates the path estimated by *Exosense* using our navigation system in localization mode; the triangles denote a subset of areas where a relocalization fix was achieved. Left images show part of the testing area: (a) while crossing a long corridor, (b) in an unmapped region, and (c) passing through a narrow doorway. The top images show visual matches between the prior map (human-leg-mounted, left), and live exoskeleton stream (right).

- [20] E. Tricomi, M. Mossini, F. Missiroli, N. Lotti, X. Zhang, M. Xiloyannis, L. Roveda, and L. Masia, "Environment-based assistance modulation for a hip exosuit via computer vision," *IEEE Robot. Autom. Lett.*, vol. 8, no. 5, pp. 2550–2557, 2023.
- [21] K. Zhang, C. Xiong, W. Zhang, H. Liu, D. Lai, Y. Rong, and C. Fu, "Environmental features recognition for lower limb prostheses toward predictive walking," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 27, no. 3, pp. 465–476, 2019.
- [22] M. Vigne, A. El Khoury, F. Di Meglio, and N. Petit, "State estimation for a legged robot with multiple flexibilities using imu: A kinematic approach," *IEEE Robot. Autom. Lett.*, vol. 5, no. 1, pp. 195–202, 2020.
- [23] M. Vigne, A. E. Khoury, M. Pétriaux, F. D. Meglio, and N. Petit, "Movie: A velocity-aided imu attitude estimator for observing and controlling multiple deformations on legged robots," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 3969–3976, 2022.
- [24] F. E. Xavier, G. Burger, M. Pétriaux, J.-E. Deschaud, and F. Goulette, "Multi-imu proprioceptive state estimator for humanoid robots," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023.
- [25] D. Tian, Z. Chen, F. Li, M. Yang, J. Li, W. Li, Y. He, L. Zhang, and X. Wu, "Dual-loop control framework of a self-balancing lower-limb exoskeleton for assisted walking," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–11, 2024.
- [26] Z. Li, T. Zhang, P. Huang, and G. Li, "Human-in-the-loop cooperative control of a walking exoskeleton for following time-variable human intention," *IEEE Trans. Cybern.*, vol. 54, no. 4, pp. 2142–2154, 2024.
- [27] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVins: A research platform for visual-inertial estimation," in *IEEE Int. Conf. Robot. Autom.*, 2020.
- [28] L. Zhang, D. Wisth, M. Camurri, and M. Fallon, "Balancing the budget: Feature selection and tracking for multi-camera visual-inertial odometry," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 1182–1189, 2022.
- [29] C. Campos, R. Elvira, J. J. Gomez, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM," *IEEE Trans. Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [30] C. Kassab, M. Mattamala, L. Zhang, and M. Fallon, "Language-EXTended Indoor SLAM (LEXIS): A Versatile System for Real-time Visual Scene Understanding," in *IEEE Int. Conf. Robot. Autom.*, 2024.
- [31] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *Intl. Conf. on Machine Learning*, 2021.
- [32] M. A. Fischler and R. C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Commun. ACM*, vol. 24, no. 6, p. 381–395, 1981.
- [33] M. Bosse, P. M. Newman, J. J. Leonard, and S. J. Teller, "Simultaneous localization and map building in large-scale cyclic environments using the Atlas framework," *Intl. J. of Robot. Res.*, vol. 23, no. 12, pp. 1113–1139, 2004.
- [34] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [35] P. Fankhauser and M. Hutter, "A Universal Grid Map Library: Implementation and Use Case for Rough Terrain Navigation," in *ROS – The Complete Reference (Volume 1)*. Springer, 2016, ch. 5.
- [36] D. Gálvez-López and J. D. Tardós, "Bags of binary words for fast place recognition in image sequences," *IEEE Trans. Robotics*, vol. 28, no. 5, pp. 1188–1197, October 2012.
- [37] M. Wermelinger, P. Fankhauser, R. Diethelm, P. Krüsi, R. Siegwart, and M. Hutter, "Navigation planning for legged robots in challenging terrain," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016.