

Unsupervised Change Detection for Space Habitats Using 3D Point Clouds

Jamie Santos^{*,¶}, Holly Dinkel^{†,¶}, Julia Di[‡], Paulo V.K. Borges[§],
Marina Moreira[¶], Oleg Alexandrov[¶], Brian Coltin[¶], and Trey Smith[¶]

This work presents an algorithm for scene change detection from point clouds to enable autonomous robotic caretaking in future space habitats. Autonomous robotic systems will help maintain future deep-space habitats, such as the Gateway space station, which will be uncrewed for extended periods. Existing scene analysis software used on the International Space Station (ISS) relies on manually-labeled images for detecting changes. In contrast, the algorithm presented in this work uses raw, unlabeled point clouds as inputs. The algorithm first applies modified Expectation-Maximization Gaussian Mixture Model (GMM) clustering to two input point clouds. It then performs change detection by comparing the GMMs using the Earth Mover’s Distance. The algorithm is validated quantitatively and qualitatively using a test dataset collected by an Astrobee robot in the NASA Ames Granite Lab comprising single frame depth images taken directly by Astrobee and full-scene reconstructed maps built with RGB-D and pose data from Astrobee. The runtimes of the approach are also analyzed in depth. The source code is publicly released to promote further development.

I. Nomenclature

D	=	the dimensions of a point cloud
P	=	the number of parameters specifying each distribution
τ	=	the log-likelihood convergence threshold
t_0, t	=	the times at which the state of the scene is captured
M, N	=	the number of points in a point cloud at t_0 and t
L	=	the maximum number of EM iterations
$\mathbf{S}_{M \times D}^{t_0}, \mathbf{S}_{N \times D}^t$	=	the sets of points representing a scene at times t_0 and t
$\varepsilon_{i,j}$	=	the Euclidean distance between two points
E	=	the Earth Mover’s Distance (EMD) between two distributions
$\mathbf{E}_{K^{*t_0} \times K^{*t}}$	=	the matrix representing the EMD between all cluster means in Θ^{t_0} and Θ^t
$\Theta_{K^{t_0} \times D}^{t_0}, \Theta_{K^t \times D}^t$	=	the Gaussian Mixture Models (GMMs) representing the scene at t_0 and t
K, K_{min}	=	the maximum and minimum, respectively, allowable number of distributions in a GMM
K^*	=	the optimal number of distributions in a GMM
Π	=	the GMM representing the areas within the map of where changes occurred

II. Introduction

As humanity ventures to establish a sustainable presence in deep space, the autonomous maintenance of space habitats will become paramount. Ensuring the safety, functionality, and longevity of these habitats—especially when they are uncrewed for long periods of time—requires advanced technologies for monitoring and detecting changes in environmental conditions. Previous work developed the first microgravity robotic assistants for space habitats [1, 2]. One of these robots is Astrobee [3], a free-flying robot currently operating on the International Space Station (ISS) that acts as a platform for experiments and research. One current research initiative for Astrobee is anomaly detection, using the robot as a mobile sensor platform. Anomaly detection with Astrobee could enable automatic detection of critical safety issues, including blocked vents or loose cargo, and identification of target areas for remapping to improve localization. These applications motivate an algorithm for autonomously detecting generalized changes over time within a map.

*Department of Physics, Chalmers University, Gothenburg, Sweden, 41296. Corresponding author. jamiesanto@gmail.com.

†Department of Aerospace Engineering, University of Illinois Urbana-Champaign, Urbana, IL, USA, 61801

‡Department of Mechanical Engineering, Stanford University, Stanford, CA, USA, 94305

§Robotics and Autonomous Systems Group, CSIRO, Brisbane, QLD, Australia

¶Intelligent Robotics Group, NASA Ames Research Center, Moffett Field, CA, USA, 94305

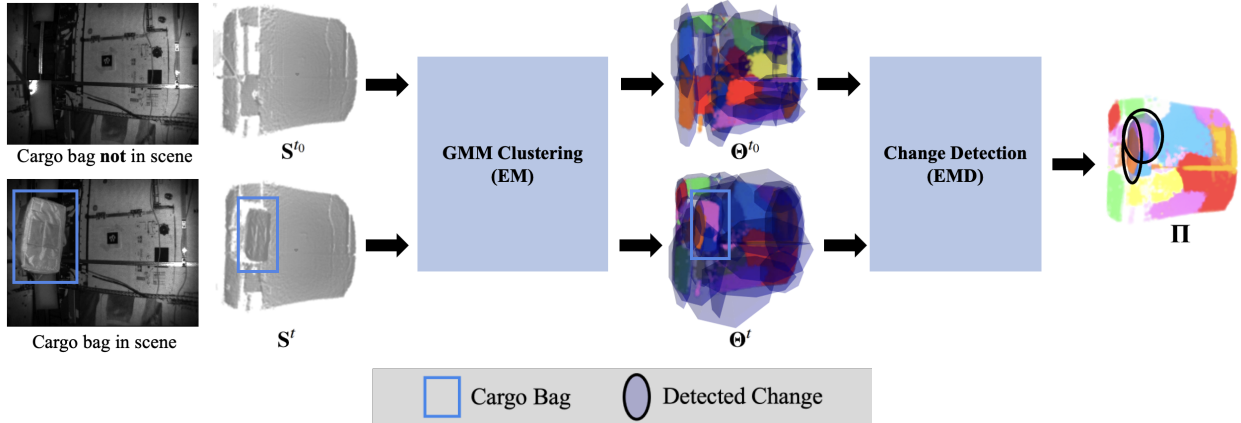


Fig. 1 The EM-EMD algorithm first clusters input point clouds S^{t_0} and S^t into Gaussian distributions using Expectation-Maximization (EM). Changes are detected using the Earth Mover’s Distance (EMD) to compare the distributions.

Current anomaly detection work for Astrobees is semantic- and image-based. In semantic-based change detection, changes can only be detected for items with known categories. Building detection models to categorize objects typically requires a large amount of manually-labeled training data. This strategy relies on up-to-date labeling and is limited to the number of object categories added manually [4, 5]. In image-based change detection, changes are detected through pairwise image comparison and inconsistency filtering against a 3D model given the camera pose [6]. While this method predicts change regions at near-real-time computation speeds, it is sensitive to the accuracy of the camera pose, the accuracy of the 3D model, and environment configurations such as lighting and reflectivity.

This work performs 3D scene change detection directly from point clouds [7]. An unsupervised clustering strategy using point cloud data from a depth sensor is a candidate method for generalizing existing research in anomaly detection [8]. The Expectation-Maximization (EM) algorithm [9, 10] consolidates point clouds into Gaussian Mixture Model (GMM) clusters with an unknown initial number of GMM clusters [11]. The Minimum Description Length (MDL) stopping criterion is used to determine when the optimal number of GMM clusters is achieved [12]. Finally, the Earth Mover’s Distance (EMD) [13] is computed between the GMM clusters at the initial and final times to select the clusters contributing the highest degree of change. This results in a set of GMM clusters in the scene indicating likely locations of change. The EM-EMD algorithm is summarized in Figure 1. This work contributes:

- 1) Detection of multiple appearing or disappearing changes within a map at one time using GMM clustering of 3D depth data.
- 2) Demonstration of change detection on real data collected using an Astrobees robot in the Granite Lab at NASA Ames Research Center. The Granite Lab is a space habitat analog, enabling testing detection and mapping algorithms under localization uncertainty, with 3-DOF microgravity simulation and imitated ISS visual features.
- 3) Evaluation of computational performance and accuracy impacts of two algorithm modifications, namely varying K , the initial number of GMM clusters, and using Principal Component Analysis as a pre-processing step to reduce the data dimensionality.
- 4) Open release of the source code at <https://github.com/nasa/isaac/tree/master/anomaly>. To the best of our knowledge, this is the first openly-available software that detects scene changes using GMM clustering on only 3D depth data.

III. Related Work

Various metrics have been suggested for identifying geometrical and appearance changes in the world. Typically, the objective is to compare the data from different times and identify spatial regions in the data that have a distance exceeding a specific threshold. The characterization of the environment is usually based on geometry, appearance, or semantics. Geometry and appearance approaches are arguably the most common, detecting changes based on inconsistency in maps from RGB or depth projections, without semantic labeling [6, 14–17]. The problem can be described as a 3D object discovery technique, where an object represents anything that can be moved. Concerning unsupervised object

discovery, some works focus on identifying parts of the input mesh as candidate objects [18]. The additions to the scene can be also be analyzed in a hybrid approach by combining global geometry with local semantics to highlight changes in the distribution of data [19].

The concept of semantic class segmentation can also be useful for change detection [20, 21]. The idea of using semantics closely relates to that behind semantic instance detection [4]. The key difference involves assigning individual voxels to classes within the scene, rather than to instances of objects. One algorithm splits voxelized data into segments that do not necessarily correlate to complete object classes, but to segments that optimize the change detection output [20]. With a voxel-based approach, it follows that the change detection algorithm may be able to attain a higher specificity than is afforded by the bounding box constraint of object detection. Nevertheless, the manual labeling of the data is still a major barrier for semantic class segmentation.

A number of works have also targeted changed detection from unmanned aerial vehicles (UAVs) or remote sensing [22, 23]. Targeting high-resolution UAV images, for example, the method outlined in [24] presents a pipeline of comparing point cloud data from two different periods. Initially, point clouds are generated from image-based 3D reconstruction. After acquiring the point clouds, separate maps are generated and a coarse-to-fine registration process is applied to overlay the two point clouds. Depth and grayscale difference maps are produced and random forest classification and component connectivity analysis techniques are applied to discover changed buildings. This algorithm was tuned and validated for a “2.5D” satellite image reconstruction, not to the 3D environment robots explore, which made this method unsuitable.

In the context of robots detecting changes online within the ISS, maps of modules reconstructed from images have historically been patchy in areas occluded by protruding objects. The quality of these maps precludes them from being a top candidate for change detection [4]. For fast 3D localization, a method of comparing updated point clouds to a base environment map has been proposed [25]. Point clouds produced from depth sensors are converted to Normal Distribution (ND) voxels using a normal distribution transform. The voxels are categorized for comparisons between the map and measured data. A similar method uses the Mahalanobis distance to compare local 3D point clouds to the nearest voxels in occupancy maps of the environment [26]. A clustering algorithm generates a list of change candidates, and outliers are removed with a random forest classifier. Classification scores and number of occurrences are used to map and report changes in real time. Another lightweight change detection algorithm detects changes between current image data and a previously constructed 3D model [6, 16]. Instead of relying on current environment maps online, an image of the current scene is back-projected onto the 3D model, and projected to a viewpoint of another recent image with an overlapping viewpoint. The projection from the model is compared to the other current image to identify differences. A short sequence of keyframes is used to remove ambiguities and to locate changes in the 3D map. This method suffers from robustness issues depending on the quality of the available 3D map, RGB camera localization, and RGB camera intrinsic calibration.

While aforementioned localization techniques demonstrate real-time change detection for robotic applications, they require high-quality and current 3D maps. This is difficult to generate in the ISS using current Astrobees and ISAAC software given the complexity and density of the environment. Additionally, methods which rely on labeling lack generalizability to new, unseen items entering space habitats. Change detection methods based on RGB images as input may lack robustness when lighting changes. A recent scene change detection algorithm addresses these issues by performing change detection directly from 3D point clouds [27]. This method summarizes the 3D point clouds as GMMs and extracts changes between the two input GMMs as the distributions contributing the greatest EMD. Because the method does not rely on RGB images, it is robust to changes in lighting. More recent work further improved the robustness of this algorithm by using a split-and-merge variation of EM to autonomously settle on the optimal number of distributions in the GMM to resolve algorithm initialization [8].

The EM-EMD algorithm in this work reliably extracts high-level representations of scene changes from 3D point clouds. It iterates between the E-step and a modified M-step in EM to delete distributions to autonomously select the final number of distributions in the GMM [28]. This algorithm is demonstrated to detect both object appearances and object disappearances. In contrast to many of the previous methods, this method can detect multiple scene changes of never-before-seen objects using only depth data, which makes it a candidate for general change detection on space habitats.

IV. Method

This section describes the new unsupervised method for change detection from 3D point clouds and data on which it was tested. First, the components of an algorithm that solves for Θ^f and Θ^{t_0} through Expectation-Maximization (EM) is

presented, where Θ^t is the GMM representing the map at time t (Section IV.A). Next, the procedure solving for the detected change regions, Π , by iteratively removing clusters using the Earth Mover's Distance (EMD) is presented (Section IV.B) [29]. As a pre-processing step, the algorithm applies the statistical outlier removal filter to the point clouds \mathbf{S}^0 and \mathbf{S}^t , and the voxel grid downsampling filter to reduce the number of point cloud points [30]. The data collected to test the algorithm were collected in a ground laboratory environment with an Astrobee robot (Section IV.C).

A. Gaussian Mixture Modeling with Expectation-Maximization

GMM clustering summarizes each point cloud with K initial Gaussian distributions. The K distributions are each initialized with a random mean, covariance, and weight, and the EM algorithm fits the data to the clusters most likely to produce them. Given K distributions and assuming each point has equal membership probability, the probability a sample point $\mathbf{s}_i^t \in \mathbf{S}^t = \{\mathbf{s}_i^t\}_{i=1,\dots,N}$ is drawn from a distribution $\theta_k^t \in \Theta^t = \{\alpha_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1,\dots,K}$ is given by

$$p(\mathbf{s}_i^t | \Theta^t) = \sum_{k=1}^K \alpha_k p(\mathbf{s}_i^t | \theta_k^t). \quad (1)$$

The prior model weights, α_k , require

$$\alpha_k \geq 0, \quad k = 1, \dots, K, \quad \text{and} \quad \sum_{k=1}^K \alpha_k = 1. \quad (2)$$

Using Bayes' theorem, the posterior probability $p(k | \mathbf{s}_i^t)$ is

$$p(k | \mathbf{s}_i^t) = \frac{p(\mathbf{s}_i^t | \theta_k^t)}{\sum_{k=1}^K p(\mathbf{s}_i^t | \theta_k^t)}. \quad (3)$$

The likelihood for a given mixture model, $\mathcal{L}(\mathbf{S}^t | \Theta^t)$, is given by

$$\mathcal{L}(\mathbf{S}^t | \Theta^t) = \log \prod_{i=1}^N p(\mathbf{s}_i^t | \Theta^t) = \sum_{i=1}^N \log \sum_{k=1}^K \alpha_k p(\mathbf{s}_i^t | \theta_k^t) \quad (4)$$

and the probabilities in Eq. (1) and Eq. (3) are optimized by maximizing the log-likelihood function to solve for the set of distributions according to

$$\hat{\Theta} = \arg \max_{\Theta^t} \{\log \mathcal{L}(\mathbf{S}^t | \Theta^t)\}. \quad (5)$$

Because the maximum likelihood estimate in Eq. (5) cannot be computed analytically, EM is used to fit the data to these distributions iteratively until convergence. The modified EM algorithm interprets \mathbf{S}^t as incomplete data missing a set of $z_{i,k} \in \mathbf{Z}_{N \times K}$ binary labels indicating which distribution produced which sample point [28]. The EM algorithm estimates (the E-step) the conditional expectation of the log-likelihood at optimization iteration $l < L$ according to

$$Q(\Theta^t, \hat{\Theta}_l) \equiv \mathbb{E} [\log p(\mathbf{S}^t, \mathbf{Z} | \Theta^t) | \mathbf{S}^t, \hat{\Theta}_l] = \log p(\mathbf{S}^t, \mathbf{W} | \Theta^t), \quad (6)$$

where $w_{i,k} \in W_{N \times K} \equiv \mathbb{E} [\mathbf{Z} | \mathbf{S}^t, \hat{\Theta}_l]$ are the posterior model weights. These weights are the conditional expectation that each distribution produces each sample point. Bayes' theorem can be used to solve for $w_{i,k}$ at iteration l as

$$w_{i,k} \equiv \mathbb{E} [z_{i,k} | \mathbf{S}^t, \hat{\Theta}_l] = \frac{\hat{\alpha}_{k,l} p(\mathbf{s}_i^t | \hat{\theta}_k^l)}{\sum_{k=1}^K \hat{\alpha}_{k,l} p(\mathbf{s}_i^t | \hat{\theta}_k^l)}. \quad (7)$$

The EM algorithm then maximizes (the M-step) these conditional expectations to update model parameters according to

$$\hat{\Theta}_{l+1} = \arg \max_{\Theta^t} \{Q(\Theta^t, \hat{\Theta}_l)\}, \quad (8)$$

where the parameters of $\hat{\Theta}_{l+1}$ are updated as

$$\hat{\boldsymbol{\mu}}_{k,l+1} = \left(\sum_{i=1}^N w_{i,k} \right)^{-1} \sum_{i=1}^N \mathbf{s}_i^t w_{i,k} \quad (9)$$

$$\hat{\boldsymbol{\Sigma}}_{k,l+1} = \left(\sum_{i=1}^N w_{i,k} \right)^{-1} \sum_{i=1}^N (\mathbf{s}_i^t - \hat{\boldsymbol{\mu}}_{k,l}) (\mathbf{s}_i^t - \hat{\boldsymbol{\mu}}_{k,l+1})^\top w_{i,k}. \quad (10)$$

The EM algorithm is initialized with K , which must be carefully selected to guarantee adequate performance. By modifying the M-step of the EM algorithm to remove unused distributions, the algorithm can be more efficiently initialized with an overestimated K to avoid expensive tuning [28]. In the modified M-Step, distributions assigned too few sample points are removed. The estimated weights of the prior distributions, $\hat{\alpha}_k$, are driven to 0 if there are fewer than $\frac{P}{2}$ points represented by $\boldsymbol{\theta}_k$ according to*

$$\hat{\alpha}_{k,l} = \frac{\max \left\{ 0, \left(\sum_{i=1}^N w_{i,k} \right) - \frac{P}{2} \right\}}{\sum_{k=1}^K \max \left\{ 0, \left(\sum_{i=1}^N w_{i,k} \right) - \frac{P}{2} \right\}} \quad \text{for } k = 1, \dots, K \quad (11)$$

$$\hat{\boldsymbol{\Theta}}_{l+1} = \arg \max_{\boldsymbol{\Theta}^t} Q \left(\boldsymbol{\Theta}^t, \hat{\boldsymbol{\Theta}}_l \right), \quad \text{for } k : \hat{\alpha}_{k,l+1} > 0, \quad (12)$$

where P is the number of parameters specifying a distribution according to

$$P = D + \frac{D(D+1)}{2}, \quad (13)$$

and $D = 3$ is the dimensionality of the input point clouds. The EM loop iterates until the difference in values of the cost function between iterations is less than a threshold, τ , defined as

$$\mathcal{L}(\mathbf{S}^t | \boldsymbol{\Theta}_{l+1}^t) - \mathcal{L}(\mathbf{S}^t | \boldsymbol{\Theta}_l^t) < \tau. \quad (14)$$

The log-likelihood, Eq. (5), is the basis for the EM cost function. This work modifies this cost function to add a penalty term, $p(k)$, to penalize a large number of distributions in $\boldsymbol{\Theta}^t$ and reduce the overfitting introduced by initializing with a large K . The objective is transformed from maximization to minimization by negating the log-likelihood cost used to fit the data and adding the new penalty term to minimize the total number of distributions to become

$$\boldsymbol{\Theta}^* = \arg \min_{\boldsymbol{\Theta}^t} \{-\log \mathcal{L}(\mathbf{S}^t | \boldsymbol{\Theta}^t) + p(k)\}. \quad (15)$$

The Minimum Description Length (MDL) criterion measures the number of parameters in $\boldsymbol{\Theta}^t$ and is used as a stopping criterion to determine when the optimal number of GMM clusters, K^* , is achieved [8, 28]. The MDL criterion was used to generate the penalty term, $p(k)$, as

$$p(k) = \frac{P}{2} \sum_{k=1}^K \log \left(\frac{N\alpha_k}{12} \right) + \frac{k}{2} \log \left(\frac{N}{2} \right) + \frac{K(P+1)}{2}. \quad (16)$$

The overall cost function becomes

$$\mathcal{L}(\boldsymbol{\Theta}^t, \mathbf{S}^t) = \frac{P}{2} \sum_{k:\alpha_k>0} \log \left(\frac{N\alpha_k}{12} \right) + \frac{k_{nz}}{2} \log \left(\frac{N}{2} \right) + \frac{k_{nz}(P+1)}{2} - \log p(\mathbf{S}^t | \boldsymbol{\Theta}^t). \quad (17)$$

Optimization returns $\boldsymbol{\Theta}^t(w, \boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t)$. The EM algorithm is run on both \mathbf{S}^{t_0} and \mathbf{S}^t to produce $\boldsymbol{\Theta}^{t_0}$ and $\boldsymbol{\Theta}^t$. These two GMMs are used to detect changes.

*Consider a set $\mathbf{s}_i^t \in \mathbf{S}_{N \times D}^t$ represented by $\boldsymbol{\theta}_k^t \in \boldsymbol{\Theta}_{K \times D}^t$. Then $\forall (\mathbf{s}_i^t, \boldsymbol{\theta}_k^t)$ the probability each point belongs to each distribution $p(\mathbf{s}_i^t | \boldsymbol{\theta}_k^t) \rightarrow 1$ and $\sum_{i=1}^N w_{i,k} < \frac{P}{2}$.

B. Change Detection Using the Earth Mover's Distance

Once the \mathbf{S}^{t_0} and \mathbf{S}^t data are represented by Θ^{t_0} and Θ^t , the scene differences can be efficiently calculated using the EMD [8]. The EMD is a metric which measures the distance between two distributions. The EMD is the work needed to move an object from a position at time t_0 to a position at time t . In this context, the objects are the distributions $\theta_k^t \in \Theta^t$, and the sizes of the objects are their weights, w_k^t . The EMD is the minimum amount of work necessary to make the two distributions Θ^{t_0} and Θ^t equal. The minimum work is computed as the optimal flow \mathbf{F} , the smallest total weight moved. This optimal transport problem is solved through linear programming with the formulation

$$\begin{aligned}
 & \underset{\mathbf{F}}{\text{minimize}} \quad \text{Work}(\Theta^{t_0}, \Theta^t, \mathbf{F}) = \sum_{j=1}^M \sum_{i=1}^N \varepsilon_{ij} f_{ij} \\
 & \text{subject to} \\
 & f_{ji} \geq 0, 1 \leq j \leq M, 1 \leq i \leq N, \\
 & \sum_{j=1}^M f_{ji} \leq w_j^{t_0}, 1 \leq j \leq M, \\
 & \sum_{i=1}^N f_{ji} \leq w_i^t, 1 \leq i \leq N, \\
 & \sum_{j=1}^M \sum_{i=1}^N f_{ji} = \min \left(\sum_{j=1}^M w_j^{t_0}, \sum_{i=1}^N w_i^t \right)
 \end{aligned} \tag{18}$$

The EMD, $E(\Theta^{t_0}, \Theta^t)$, results from normalizing the work by the total flow according to

$$E(\Theta^{t_0}, \Theta^t) = \frac{\text{Work}(\Theta^{t_0}, \Theta^t, \mathbf{F})}{\sum_j^M \sum_i^N f_{ij}}. \tag{19}$$

The EMD is used in a greedy selection algorithm to extract the clusters from Θ^t which contribute the highest amount of change [8]. The distribution whose removal best decreases the EMD between Θ^{t_0} and Θ^t is extracted. Extracted clusters are transferred to the final model, Π , which stores detected changes. This process iterates until the EMD between the two distributions stops decreasing. The final model, Π , is used to identify the points in \mathbf{S}^t where change occurred. This process is summarized in Algorithm 1.

Algorithm 1 EM + EMD Algorithm for Change Detection

```

 $\Theta^{t_0}, \Theta^t \leftarrow \text{EM}(\mathbf{S}^{t_0}, \mathbf{S}^t)$ 
 $\Pi \leftarrow \mathbf{0}$ 
 $E_{old} \leftarrow E(\Theta^{t_0}, \Theta^t)$ 
 $\Theta^* \leftarrow$  random initialization
while  $E_{old} > E_{lowest}$  do
  for  $\theta_k^t$  in  $\Theta^t$  do
     $k^* \leftarrow 0$ 
     $\Theta_{temp} \leftarrow \Theta^t - \theta_k^t$ 
     $E_{new} \leftarrow E(\Theta^{t_0}, \Theta_{temp})$ 
    if  $E_{new} < E_{lowest}$  then
       $E_{lowest} \leftarrow E_{new}$ 
       $k^* \leftarrow k$ 
       $\Theta^* \leftarrow \Theta_{temp}$ 
    end if
  end for
   $\Theta^t \leftarrow \Theta^*$ 
   $\Pi \leftarrow \Pi + \theta_{k^*}^t$ 
end while
return  $\Pi$ 

```




Fig. 2 The Granite Lab at NASA Ames Research Center simulates the visual environment of the ISS and mimics 3-DOF microgravity by mounting Astrobee on a near-frictionless air bearing carriage. This system is used for testing Astrobee perception, motion planning, and control.

C. Data Collection with Astrobee

The data used to verify the EM-EMD algorithm were collected with an Astrobee robot in the Granite Lab at NASA Ames Research Center, shown in Figure 2. The Granite Lab is a facility that imitates the visual features of the ISS. The facility mimics 3-DOF microgravity by placing Astrobee on a near-frictionless air bearing, allowing it float freely on a 2×2 m granite monolith. In this environment, Astrobee uses its own onboard sensors and actuators to detect landmarks and complete maneuvers, replicating its behaviors on the ISS. A collocated ground station computer allows for real-time tracking of Astrobee’s position and downlinking sensor data. The Navigation Camera (NavCam) on Astrobee collects Bayer images at 1280×960 resolution at 5 Hz. It is a fixed-focus RGB camera with a wide field of view [3, 31]. The Hazard Camera (HazCam) on Astrobee is a PMD Pico Flexx 2 time-of-flight sensor that collects depth images at 224×172 resolution at 5 Hz. To create the full-scene 3D point cloud maps of the environment used throughout this work, the NASA Ames Stereo Pipeline package uses Astrobee NavCam, HazCam, and localization data [2, 32, 33]. It registers image data from the NavCam with depth information from the HazCam using Theia structure-from-motion [34], and then fuses the depth point clouds into a mesh [35]. This work also uses single-frame point clouds, which were taken directly from the Astrobee HazCam raw data stream.

Astrobee conducted surveys of five different scenes where up to three large objects were placed in, moved, or removed from the scene. These objects include another Astrobee robot, a crate, and a cargo bag, as shown in Figure 3. The first scene contained none of these objects and served as the basis for comparison for change detection from both single frame point cloud and reconstructed map data. The second scene contained an added cargo bag used to test change detection on single frame point cloud data. The third scene contained an added Astrobee also used to test change detection on single frame point cloud data. The data from the second and third scenes is shown in Figure 5. The fourth and fifth scenes were both processed into reconstructed maps. The fourth scene contained an added Astrobee and an added cargo bag. The fifth scene contained an added Astrobee, an added cargo bag, and an added crate. The data from the fourth and fifth scenes are shown in Figure 4. The two single frame point clouds containing change objects are also referred to as one-object scenes and the reconstructed maps containing two and three changed objects are referred to as two- and three-object scenes in Section V.E.

V. Results

The performance of the EM-EMD algorithm is demonstrated on both full-scene reconstructed maps (Section V.A) and single-frame point cloud data collected by the HazCam on Astrobee (Section V.B). The impact on accuracy of two algorithm variations, namely varying K and pre-processing the input point clouds with PCA, is discussed (Section V.C), and runtime is presented (Section V.D). Finally, performance of the algorithm on both single frame and full-scene reconstructed maps is evaluated (Section V.E). For all experiments, $K_{min} = 1$, $K = 25$, $\tau = 10^{-5}$, and $L = 100$ unless otherwise stated. The notions of True Positive (TP), False Positive (FP), and False Negative (FN) detection are used




Fig. 3 Three large objects were placed in, moved, or removed from each scene. These images are captured with the NavCam on Astrobee for Scene 5. (Left) An Astrobee was placed near the edge of the scene. (Middle) A crate was placed in one corner of the scene. (Right) A cargo bag was placed in one corner of the scene.

throughout the qualitative and quantitative discussion of results.

A. Reconstructed Maps

Detecting changes on full-scene reconstructed maps enables anomaly identification across large regions and allows for correcting localization uncertainty. Figure 4 shows two different scene changes using reconstructed maps as input. In the first scene, an Astrobee robot and a cargo bag were placed along the same wall. The Astrobee docking station is also fixed along this wall and present in every scene. Despite the presence of the dock, which may be falsely detected as a scene change along with the variable scene objects, the Astrobee robot and cargo bag are both correctly identified scene changes. Figure 4 also shows a scene with three added objects. In this scene, a crate, cargo bag, and Astrobee were placed along separate walls, with the Astrobee placed on its docking station. The location of the docking station at the scene edge places the Astrobee next to significant background noise, shown in Figure 3. The Astrobee was not directly detected (FN), however the region behind it was. The crate and cargo bag were both correctly identified.

B. Raw Point Cloud Data

The performance of the EM-EMD algorithm is demonstrated on raw point cloud data acquired by an Astrobee robot for a mapping survey of the Granite Lab. The robot traverses the lab, capturing still depth images with its HazCam at specific locations. The EM-EMD change detection is robust to robot localization uncertainty as shown in Figure 5. For each of the four input point clouds shown in Figure 5, the EM algorithm converged to $K^* \in [20, 23]$ with $K = 25$. In both scenarios, the object was correctly identified as a change in the scene (TP), and no objects were falsely identified as scene changes (FP).

C. Algorithm Variations

Two variations were applied to the algorithm to study their effect on performance: varying the initial number of GMM clusters, K , and PCA as a pre-processing step. Figure 6 demonstrates the impact of increasing K . It shows that when K is too small, objects are clustered together with the surrounding walls and surfaces, reducing recall. This is evident with $K = 10$, where the cargo bag was clustered along with the wall and not detected as a change (FN). As K increases, the resolution of the model increases. Objects are represented by an increasing number of distributions and the total detected change area approximately decreases. For $K = 50$, the objects begin to fragment. This could be overcome by inferring distributions with centers close in space as a single scene change. However, increasing K increases runtime significantly. For the reconstructed maps and single-frame point clouds used in this work, $K = 25$ struck the best balance between accurately detecting appearances, accurately detecting disappearances, and runtime.

Data dimensionality reduction on the input point clouds was performed using PCA. The change objects protrude from planar walls in the scene, so points belonging to the scene and points belonging to the change objects are distinct. Applying PCA to this data removes the z -dimension of the point clouds. Changes are accurately identified when using PCA for pre-processing as shown in Figure 7, however pre-processing with PCA led to a 10% speedup in total change detection runtime for this test case. These results show pre-processing with PCA to reduce the complexity of the input point clouds could reduce runtime while maintaining performance. However, in complex environments such as the ISS,




Fig. 4 Scene change detection with the EM-EMD algorithm identifies changes on reconstructed maps of the environment. (Left) Two objects, an Astrobee and cargo bag, are added along one wall of the scene, and this region along the wall is identified as a change region. (Right) Three objects, an Astrobee, cargo bag, and crate, are each placed along separate walls in the scene. The cargo bag and crate are correctly detected areas of change (TP), while significant noise at the map boundary led to a shift in the change region. Astrobee was not correctly identified as a scene change (FN).

Table 1 Change Detection Runtime [s]

K	Data Loading	PCA	GMM Clustering (EM)	Change Detection (EMD)	Total
15	0.044	3.727	2043.113	0.677	2047.561
20	0.037	2.580	2804.327	1.491	2808.535
25	0.037	1.979	4496.449	3.070	4501.535
30	0.043	2.920	3307.551	2.708	3313.222
40	0.023	1.981	5209.999	5.254	5217.257
50	0.026	2.169	7714.528	9.415	7726.138

removing one spatial dimension may have more pronounced impacts on change detection performance and may require further study.

D. Runtime

A runtime analysis was performed to verify the standalone performance of the change detection algorithm. Computational timing data were collected for two-object change detection from reconstructed maps (see Figure 4) on a computer with an 8-core, 3.2 GHz Apple M1 CPU with 8 GB RAM. Table 1 reports runtimes for four change detection computational processes based on the number of Gaussian clusters, K , used to compute changes. Filtering was not applied to the input point clouds and therefore all of the tunable parameters were solely in the GMM clustering portion of the algorithm and left as their default values. Data are loaded as .ply point clouds. The EM portion of the algorithm accounts for most of the runtime, increasing approximately linearly with K . Runtime could be improved by increasing K_{min} to a more reasonable lower-bound, downsampling the input point clouds, or decreasing J .

E. Performance Metrics

The standard performance metrics of accuracy, precision, recall, and $F1$ score are used to compare the performance of the EM-EMD algorithm on the single frame point clouds with reconstructed maps. Accuracy measures the ratio of predicted TP and TN to all true and false detections, precision measures the ratio of predicted TP to predicted FP detections, recall measures the ratio of predicted TP to predicted FN detections, and the $F1$ score reflects the ratio of TP to FP and FN detections. The TPs were counted as the number of distributions in Π overlapping with a real change in

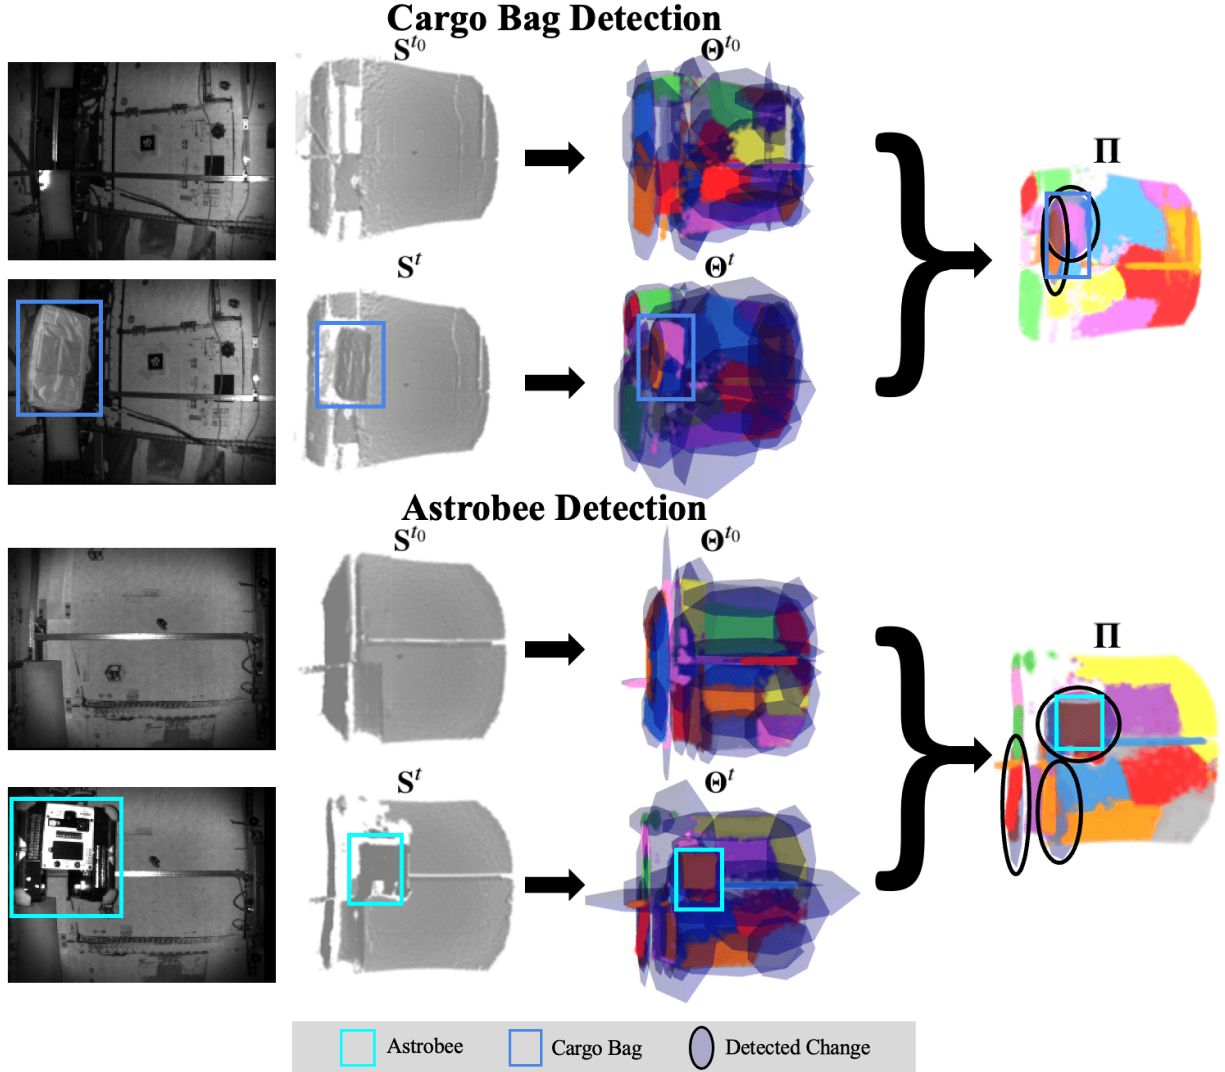


Fig. 5 A cargo bag and Astrobee robot are detected by comparing two single-frame depth images taken at the same location at different times. (Top) Cargo bag detection (Bottom) Astrobee detection (1st Stage) Images and input point clouds (2nd Stage) GMM clustering using EM (3rd Stage) Change detection using EMD

the scene. The FPs were counted as the number of distributions identified as changes in Π which were not real changes in the scene. The FNs were counted as the number of distributions not marked as changes in Π which were real changes in the scene. The scores of EM-EMD change detection on appearances in the single frame point clouds (one-object scene) and reconstructed maps (two- and three-object scenes) are reported in Table 2. These metrics are defined as

$$\text{Accuracy} = \frac{TP + TN}{TP + FN + TN + FP} \quad (20)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (21)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (22)$$

$$\text{F1} = \frac{2TP}{2TP + FP + FN}. \quad (23)$$

The EM-EMD algorithm achieved the highest recall on single frame point cloud input data. This could be because

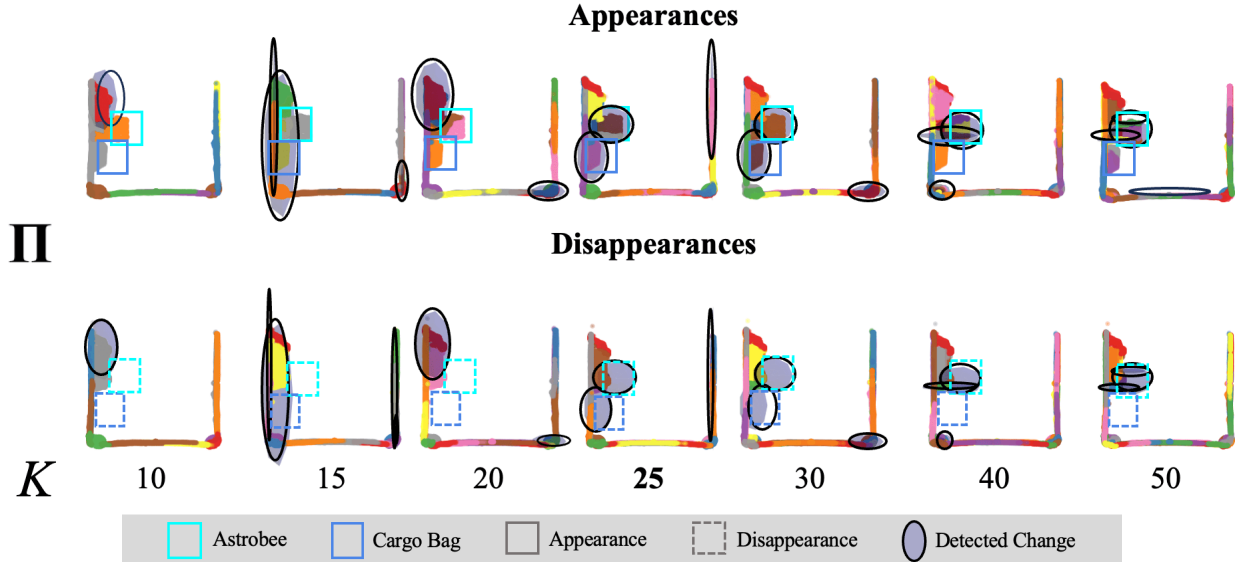


Fig. 6 When varying the initial number of GMM clusters, K , a wider area of the map is generally detected as a change region when fewer clusters are used while a smaller area of the map is generally detected as a change region when K is larger. The parameter $K = 25$ provided the best balance for this scene.

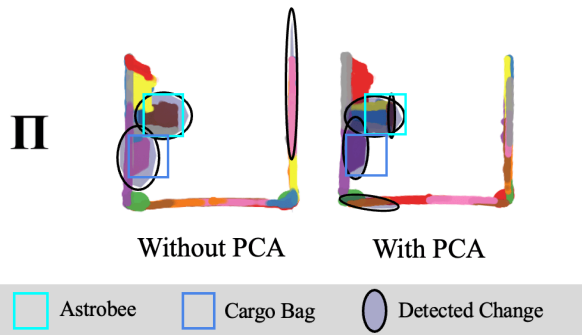


Fig. 7 Applying PCA as a pre-processing step reduces the dimensionality of the point clouds input to the EM-EMD algorithm. The EM-EMD algorithm still accurately detects changes on the PCA-processed data. (Left) Change detection result without pre-processing with PCA. (Right) Change detection result with pre-processing with PCA.

the single frame point clouds have the largest object-to-search-boundary ratio, simplifying change detection. The EM-EMD algorithm achieved the highest precision and $F1$ score on the reconstructed map data for the three-object scene. Objects were more evenly spaced in this scene than in the two-object scene, and it was not subject to the alignment error of the single frame point cloud data. However, the $F1$ score for the single frame point cloud data was significantly higher than the $F1$ score for the reconstructed map of the two-object scene because of high recall for this test case. The $F1$ scores provide the most complete understanding of expected behavior across the input data type and number of scene object test cases presented in this work.

Table 2 only shows data for detected object appearances, however these scores are similar for detecting object disappearances. This is because the algorithm treats object disappearances as object appearances by reversing the order of the t and t_0 input point clouds. These data represent the case where the algorithm is either predicting appearances when only appearances exist in the scene or predicting disappearances when only disappearances exist in the scene. No test cases which mix predicting both appearances and disappearances were evaluated in this work.

Table 2 Performance by Number of Appearance Objects

Metric	One-Object Scene	Two-Object Scene	Three-Object Scene
Accuracy	0.953	0.929	0.952
Precision	0.600	0.617	0.933
Recall	1.000	0.667	0.750
F1	0.727	0.650	0.829

VI. Conclusions

This work presented a change detection algorithm for detecting multiple object appearances or disappearances in a scene from point clouds. The algorithm does not depend on manually-labeled data and is able to detect changes where new, previously unseen objects may be introduced. The system was demonstrated and validated in experiments on data collected by an Astrobee robot at the Granite Lab facility at NASA Ames Research Center as a ground environment simulating robotic caretaking of space habitats. The runtime of each step of the algorithm was also analyzed.

A. Limitations

One current limitation of the EM-EMD algorithm is that it cannot detect both object appearances and object disappearances between t and t_0 . This could be addressed by adding a semantic layer to identify objects by category within each scene before applying change detection. The EM-EMD algorithm also predicts FP detections around scene edges which may be noisy or misaligned between time steps. Single frame point clouds were extracted by manually finding the times at which the Astrobee was looking at the same position in each scene, rather than by automatically finding the two closest robot body poses, introducing human error and misalignment between \mathbf{S}^{t_0} and \mathbf{S}^t . When Astrobee scanned the scene to create the reconstructed maps, it scanned beyond the edge of the lab so manual cropping is necessary to set the search boundary. When Θ^{t_0} contains more distributions than Θ^t , whether due to a real scene change (TP) or due to noise or spatial misalignment (FP), a change will be detected.

B. Future Work

One area of future work could analyze the trade-off between runtime and accuracy for the EM-EMD algorithm. Future work could also build semantic understanding into the change detection architecture to enable detecting object appearances and object disappearances within the same scene. Similarly, the impacts of data pre-processing methods such as point cloud downsampling and filtering could be analyzed to improve change detection accuracy. Improving the quality of map reconstruction from ISS data is ongoing. When ISS reconstructed maps become available, the EM-EMD algorithm could be applied to this data. The ISS reconstructed maps are expected to present more complex scenes with both appearing and disappearing objects, a wide range of object sizes and scales, more point cloud noise, and more scene changes between time steps.

Finally, future robotic assistants will increasingly support deep space exploration through consistent caretaking operations. NASA is developing an intermittently-crewed lunar Gateway as infrastructure for venturing to distant locations. Future work could advance anomaly detection in these facilities, enabling advanced autonomous caretaking for Astrobee and future robotic assistants.

Acknowledgements

Part of this work was submitted by Jamie Santos in partial fulfillment of requirements for the Complex Adaptive Systems master’s degree at Chalmers University of Technology. The NASA Space Technology Graduate Research Opportunity 80NSSC21K1292 and the P.E.O. Scholar Award supported Holly Dinkel’s involvement. Julia Di was previously supported by the NASA Space Technology Graduate Research Opportunity 80NSSC18K1197. The CSIRO Data61 and the Winston Churchill Fellowship supported Paulo V.K. Borges. The NASA Game Changing Development (GCD) Program supported Jamie Santos, Marina Moreira, Oleg Alexandrov, Brian Coltin, and Trey Smith. The authors thank Ryan Soussan and the Astrobee Facilities team for their assistance.


References

- [1] Otero, A., Chen, A., Miller, D., and Hilstad, M., “SPHERES: Development of an ISS Laboratory for Formation Flight and Docking Research,” *Proceedings, IEEE Aerospace Conference*, Vol. 1, IEEE, Big Sky, MT, USA, 2002, pp. 1–59–1–73. <https://doi.org/10.1109/AERO.2002.1036828>.
- [2] Smith, T., Bualat, M., Akanni, A., Alexandrov, O., Barron, L., Benton, J., Coltin, B., Fong, T., Garcia, J., Hamilton, K., Hill, L., Moreira, M., Morris, R., Ortega, N., Pea, J., Rogers, J., Savchenko, M., Sharif, K., and Soussan, R., “ISAAC: An Integrated System for Autonomous and Adaptive Caretaking,” *ISS R&D Conference*, 2021, pp. 1–21.
- [3] Smith, T., Barlow, J., Bualat, M., Fong, T., Provencher, C., Sanchez, H., and Smith, E., “Astrobee: A New Platform for Free-Flying Robotics on the International Space Station,” *Int. Sympos. Artif. Intell. Robot. Autom. in Space (i-SAIRAS)*, 2016, pp. 1–8.
- [4] Miller, I. D., Soussan, R., Coltin, B., Smith, T., and Kumar, V., “Robust Semantic Mapping and Localization on a Free-Flying Robot in Microgravity,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022, pp. 4121–4127.
- [5] Coltin, B., Fusco, J., Moratto, Z., Alexandrov, O., and Nakamura, R., “Localization From Visual Landmarks on a Free-Flying Robot,” *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2016, pp. 4377–4382.
- [6] Dinkel, H., Di, J., Santos, J., Albee, K., Borges, P., Moreira, M., Alexandrov, O., Coltin, B., and Smith, T., “Multi-Agent 3D Map Reconstruction and Change Detection in Microgravity with Free-Flying Robots,” *IAF Int. Astronaut. Cong. (IAC)*, 2023, pp. 1–11.
- [7] Santos, J., “Detecting Changes on the ISS Autonomously with 3D Point Clouds: An Unsupervised Learning Approach Using GMM Clustering,” Master’s thesis, Chalmers University, Gothenburg, Sweden, 2023.
- [8] Manso, L. J., Núñez, P., da Silva, S., and Drews-Jr, P., “A Novel Robust Scene Change Detection Algorithm for Autonomous Robots Using Mixtures of Gaussians,” *Int. J. Adv. Robot. Syst.*, Vol. 11, 2014, p. 18.
- [9] Dempster, A. P., Laird, N. M., and Rubin, D. B., “Maximum Likelihood from Incomplete Data Via the *EM* Algorithm,” *J. R. Stat. Soc., Ser. B, Methodol.*, Vol. 39, 1977, pp. 1–22.
- [10] Bishop, C. M., and Nasrabadi, N. M., *Pattern Recognition and Machine Learning*, Vol. 4, Springer, 2006.
- [11] Li, Y., and Li, L., “A Novel Split and Merge EM Algorithm for Gaussian Mixture Model,” *Int. Conf. Nat. Comput.*, Vol. 6, 2009, pp. 479–483.
- [12] Kyrgyzov, I. O., Kyrgyzov, O. O., Maître, H., and Campedel, M., “Kernel MDL to Determine the Number of Clusters,” *Int. Conf. Mach. Learn. and Data Min. Pattern Recognit. (MLDM)*, 2007, pp. 203–217.
- [13] Pele, O., and Werman, M., “Fast and Robust Earth Mover’s Distances,” *IEEE Int. Conf. Comput. Vis. (ICCV)*, 2009, pp. 460–467.
- [14] Taneja, A., Ballan, L., and Pollefeys, M., “City-Scale Change Detection in Cadastral 3D Models Using Images,” *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2013, pp. 113–120.
- [15] Xiao, W., Vallet, B., and Paparoditis, N., “Change Detection in 3D Point Clouds Acquired by a Mobile Mapping System,” *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, Vol. 1, No. 2, 2013, pp. 331–336.
- [16] Palazzolo, E., and Stachniss, C., “Fast Image-Based Geometric Change Detection Given a 3D Model,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2018, pp. 6308–6315.
- [17] Palazzolo, E., and Stachniss, C., “Change Detection in 3D Models Based on Camera Images,” *9th Workshop on Planning, Perception and Navigation for Intelligent Vehicles at the IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2017, pp. 1–6.
- [18] Karpathy, A., Miller, S., and Li, F.-F., “Object Discovery in 3D Scenes Via Shape Analysis,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2013, pp. 2088–2095.
- [19] Langer, E., Patten, T., and Vincze, M., “Robust and Efficient Object Change Detection by Combining Global Semantic Information and Local Geometric Verification,” *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2020, pp. 8453–8460.
- [20] Neuman, B., Sofman, B., Stentz, A., and Bagnell, J. A., “Segmentation-Based Online Change Detection for Mobile Robots,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2011, pp. 5427–5434.



- [21] Shi, W., Zhang, M., Zhang, R., Chen, S., and Zhan, Z., “Change Detection Based on Artificial Intelligence: State-of-the-Art and Challenges,” *Remote Sens.*, Vol. 12, No. 10, 2020, p. 1688.
- [22] Asokan, A., and Anitha, J., “Change Detection Techniques for Remote Sensing Applications: A Survey,” *Earth Sci. Inform.*, Vol. 12, 2019, pp. 143–160.
- [23] Khelifi, L., and Mignotte, M., “Deep Learning for Change Detection in Remote Sensing Images: Comprehensive Review and Meta-Analysis,” *IEEE Access*, Vol. 8, 2020, pp. 126385–126400.
- [24] Chen, B., Chen, Z., Deng, L., Duan, Y., and Zhou, J., “Building Change Detection with RGB-D Map Generated from UAV Images,” *Neurocomputing*, Vol. 208, 2016, pp. 350–364.
- [25] Katsura, U., Matsumoto, K., Kawamura, A., Ishigami, T., Okada, T., and Kurazume, R., “Spatial Change Detection Using Voxel Classification by Normal Distributions Transform,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2019, pp. 2953–2959.
- [26] Wellhausen, L., Dubé, R., Gawel, A., Siegart, R., and Cadena, C., “Reliable Real-Time Change Detection and Mapping for 3D Lidars,” *IEEE Int. Symp. Saf. Secur. Rescue Robot. (SSRR)*, 2017, pp. 81–87.
- [27] Núñez, P., Drews, P., Rocha, R., Campos, M., and Dias, J., “Novelty detection and 3D shape retrieval based on Gaussian Mixture Models for autonomous surveillance robotics,” *IROS*, 2009, pp. 4724–4730.
- [28] Figueiredo, M., and Jain, A., “Unsupervised Learning of Finite Mixture Models,” *IEEE Trans. Pattern Anal. Mach. Intell.*, Vol. 24, No. 3, 2002, pp. 381–396.
- [29] Rubner, Y., Tomasi, C., and Guibas, L. J., “The Earth Mover’s Distance as a Metric for Image Retrieval,” *Int. J. Comput. Vis.*, Vol. 40, No. 2, 2000, p. 99.
- [30] Rusu, R. B., and Cousins, S., “3D is Here: Point Cloud Library (PCL),” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2011, pp. 1–4.
- [31] NASA, “Astrobee Robot Software.” [Online] Available: <https://github.com/nasa/astrobee>, 2023.
- [32] NASA, “ISAAC (Integrated System for Autonomous and Adaptive Caretaking),” [Online] Available: <https://github.com/nasa/isaac>, 2023.
- [33] Soussan, R., Kumar, V., Coltin, B., and Smith, T., “AstroLoc: An Efficient and Robust Localizer for a Free-flying Robot,” *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022, pp. 4106–4112.
- [34] Sweeney, C., “Theia Multiview Geometry Library: Tutorial & Reference.” [Online] Available: <http://theia-sfm.org>, 2023.
- [35] Beyer, R. A., Alexandrov, O., and McMichael, S., “The Ames Stereo Pipeline: NASA’s Open Source Software for Deriving and Processing Terrain Data,” *Earth Space Sci.*, Vol. 5, 2018.

Authors





Jamie Santos  completed an M.S. in Complex Adaptive Systems at Chalmers University in Gothenburg, Sweden. She researches change detection to enable NASA's free-flying Astrobees robots to detect anomalies on the International Space Station with the NASA Ames Research Center Intelligent Robotics Group.





Holly Dinkel   is pursuing a Ph.D. in aerospace engineering at the University of Illinois Urbana-Champaign where she researches robotic caretaking as a NASA Space Technology Graduate Research Fellow with the NASA Ames Research Center Intelligent Robotics Group and the NASA Johnson Space Center Dexterous Robotics Laboratory.




Julia Di   is pursuing a Ph.D. in mechanical engineering at Stanford University where she researches tactile sensing and perception. She was a NASA Space Technology Graduate Research Fellow with the NASA Ames Research Center Intelligent Robotics Group and NASA Jet Propulsion Laboratory Mobility and Robotic Systems Section.




Paulo Borges   is a Principal Research Scientist in the Robotics and Autonomous Systems Group at CSIRO in Brisbane, Australia. He completed a Ph.D. in Electronic Engineering from Queen Mary University of London, London, United Kingdom. His research interests include robotic automation for the manufacturing, energy, and agriculture industries, bridging industry, innovation, and research.






Marina Moreira  is a research engineer in the NASA Ames Research Center Intelligent Robotics Group. She completed a M.S. in aerospace engineering from Instituto Superior Técnico, Lisbon, Portugal. She develops and maintains open-source software for the ISAAC problem and is interested in problems related to robotic systems and control.





Oleg Alexandrov  is a Research Scientist in the NASA Ames Research Center Intelligent Robotics Group. He received a Ph.D. in applied mathematics from the University of Minnesota. His research focuses on mapping using satellite and robot images.



Brian Coltin   is a Computer Scientist in the NASA Ames Research Center Intelligent Robotics Group where he works on the Astrobees robot, the VIPER rover, and flood mapping. He earned a Ph.D. in Robotics from Carnegie Mellon University. His research interests include planning, scheduling, multi-robot coordination, localization, and computer vision.



Trey Smith   is a Computer Scientist in the NASA Ames Research Center Intelligent Robotics Group where he works on the Astrobees robot and the VIPER rover. He earned a Ph.D. in Robotics from Carnegie Mellon University. His research interests include planning, scheduling, multi-robot coordination, mapping, and computer vision.