2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

AutoExplorers: Autoencoder-Based Strategies for High-Entropy Exploration in Unknown Environments for Mobile Robots

Lennart Puck[1] Maximilian Schik[1] Tristan Schnell[1] Timothée Buettner[1] Arne Roennau[1] Rüdiger Dillmann[1]

***Abstract*— Deciding where to go next is a challenging task for humans. However, for robots in unknown environments, this becomes even more demanding. In planetary explorations, the robots are continuously challenged with the task of exploring novel areas, yet so far, humans decide for the robots where to go. Even then, prioritizing the next target based on previous knowledge is complex. In our proposed work, the robot utilizes data about its surroundings from drone or satellite images. Alternatively, a volumetric representation can be reduced to form a suitable input. From the input, tiles are selected and embedded by different autoencoder variants. The robot can select the most promising next exploration goal through the distance in the embedding to the previous samples. In this work, a variational autoencoder, a Wasserstein autoencoder, and a spherical autoencoder are evaluated against each other. The latter two variants yield a high information gain when evaluated on satellite data from the Netherlands. Additionally, the framework was employed on data from an analog mission in the Tabernas desert. Through the framework, the robots get an understanding of which goals yield the most information gain and, therefore, can quickly improve their knowledge about their surroundings.**

I. INTRODUCTION

Exploration is a fundamental aspect of any robotic mission into unknown environments. Without detailed information about the surroundings, the robot cannot safely plan paths or traverse the environment. Furthermore, and especially in planetary exploration missions, each exploration step comes with a cost in terms of resources. This might be energy consumption or the wear of the robot's components, which might lead to fatal failures. The success and efficiency of exploration missions could be defined by the amount of novel information gained throughout the mission. As such, selecting the right exploration goals has a crucial impact on the overall mission.

There are different ways of representing the surroundings. The simplest is a 2D representation of free and occupied spaces as proposed by Moravec and Elfes in [1]. More recent approaches target 2.5D and 3D spaces. The Gridmap proposed by Fankhauser et al. [2] is a multi-layered 2D map. It provides additional functionality, such as traversability estimation depending on the elevation map, slope, step size, and terrain roughness. In the 3D case, the most well-known approach is the OctoMap by Hornung et al. in [3]. The approach utilizes an Octree to discretize the environment.

[1] Department of Interactive Diagnosis and Service Systems (IDS), FZI Research Center for Information Technology, Haid-und-Neu-Straße 10–14, 76131 Karlsruhe, Germany.

The research leading to these results has received funding from the DLR Space Administration under grant agreement No. 50RA2026 by the German Bundestag.

Fig. 1. ANYmal C [7] in the Spanish Tabernas desert during an analog mission. The world is depicted in the VDBMap with additional information. The possible exploration goals are marked with blue flags. The robot needs to prioritize the goals according to the possible information gain.

Recent developments utilize a different underlying data structure. With the OpenVDB Framework by Museth [4], an efficient and real-time capable data structure was proposed. Both the Spatio Temporal Voxel Grids by Macsenski et al. [5] and the VDBMapping framework by Grosse Besselmann et al. [6] apply OpenVDB to generate highly efficient volumetric maps. In our proposed work, we utilize the VDBMapping Framework [6] with an extension for different layers, which include ground properties and RGB information.

Novelty Detection or Anomaly Detection can be based on different methods. Scholkopf et al. [8] proposed a Support Vector Machine based method. They fit hyperplanes, so they separate between outliers and inliers. For more complex scenarios, deep learning architectures are used [9]. For example, the variational autoencoder (VAE) [10] can be used for novelty detection. An and Cho [11] apply a VAE for anomaly detection by utilizing the reconstruction error.

In the context of robotics, anomaly detection is used for both internal fault detection and detection of changes in the environment. Hornung et al. [12] propose a model-free system based on SVMs and dimensionality reduction. They apply both positive and negative examples in their data. Schnell et al. [13] only use the positive data to train a generative adversarial network for fault detection.

For novelty detection in outdoor inspection tasks, Contreras-Cruz et al. [14] propose a method based on evolving connectionist systems (ECoS). These are trained on the normal data and then can find novelties in new scenes. The ECoS has the advantage of being trained by one-pass learning and incrementally adapting to new data. Özbildge et al. [15] use a recurrent neural network to build dynamic models of the sensor information. Through a prediction error in the predictions, novelties can be detected and then

dynamically included in the model. Puck et al. [16] applied an ensemble of small autoencoders to find novelties in the environment. Through the modular approach, the ensemble can be reconfigured in different contexts or when new data is available.

Exploration tasks in mobile robotics deal with the challenge of maximizing the knowledge about an unknown environment [17]. The current known map is used to achieve this, and a pose, the so-called Next-Best-View (NBV), is generated, maximizing the knowledge gained. The two primary approaches are frontier-based [18] and sampling-based [19] exploration.

Frontiers are borders of free cells with unknown cells clustered, and the cluster's centroid is selected as the NBV [18]. The approach has been adapted for 3D space [20] and found applications in high-speed flight exploration with UAV [21].

An evaluation metric is introduced for sampling-based approaches, which tries to approximate the information gain. Borugault et al. [19] evaluate the samples by the increase the location brings to the confidence in the robot's location. Bircher et al. [22] use a Rapid-exploring Random Tree (RRT), which evaluates the number of unmapped voxels seen in each node. The node with the highest score is then picked as the NBV.

Recent advances use reinforcement learning to select the NBV, for example, the exploration planners by Zhu et al. [23] and Tai et al. [24]. A mapless reinforcement learning approach based on an Intrinsic Curiosity Model was proposed by Zhelo et al. in [25].

Schmid et al. [26] train a VAE on a set of high utility viewpoints combined with an embedding of a 2D occupancy grid. A second model was trained on the viewpoints to approximate the information gain. This leads to a fast sampling of viewpoints with a higher information gain.

Finally, Girdhar et al. [27] propose the use of *realtime online spatiotemporal topic modeling framework* (ROST) for exploration. They extract a set of fixed features and index these to represent their encoding. The distribution of those features is then modeled by a topic distribution where topic perplexity is used as the metric by which exploration goals are chosen. Additionally, the authors apply their research to satellite images [28] and underwater exploration [29]. The approach yields a higher information gain measured as Shannon entropy than a coverage-orientated approach.

In this work, an exploration planner is developed, which selects goals based on their novelty score. The embeddings of different autoencoder variants are used to achieve this. The applied autoencoders are variational, spherical, and Wasserstein autoencoders. Through the distance to the previous embeddings, new samples can be analyzed according to their value in the current mission. We compare our approach against the ROST framework from Girdhar et al. [27] and against a random environment sampling.

The paper is structured as follows: Section I gave the intro and summarised the current state of the art. Section II details the concept of the exploration goal selection and the different used variants of autoencoders. In Section III, the evaluation of two datasets of the Netherlands and an analog mission in the Tabernas desert are shown. Finally, in Section IV, a conclusion and the next steps are detailed.

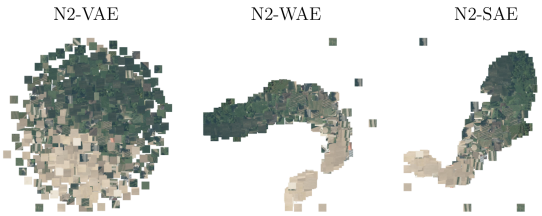

Fig. 2. Visualization of the latent space of the different autoencoders. Each autoencoder is trained with the N2 dataset around, and the same dataset is then embedded and visualized. The visualization is done with t-SNE [32]. The VAE differentiates more between brighter and darker tiles, whereas the other two have a more defined latent space. The more differentiated the latent space is, the better a new sample can be rated. A novel context would be embedded in a completely different part of the latent space and, therefore, has a high distance and will be chosen as the next exploration target.

II. EXPLORATION GOAL SELECTION

The proposed exploration goal generation is embedded into our autonomous robot stack and should be able to replace the currently applied frontier detection. The framework is trained and evaluated on satellite data from the Netherlands [30] and data from an analog mission in the Spanish Tabernas desert as part of the intelliRISK2 project [31].

The more risks are associated with a mission, the more critical it is to select informative goals for an asset. Exploration methods such as frontier detection only consider geometric structures. However, the gained information should be maximized when each action comes with a risk. Additional knowledge, such as physical properties or color, is mostly not considered. This work generates a structured latent space through different autoencoder variants for multidimensional inputs to measure the information gain. This helps in selecting exploration goals for the robot and can extend frontier detections or sampling-based methods. The goal selection will be triggered once the robot needs to select a new exploration goal; thereby, it is executed on demand.

A. Structured Latent Space

A structured latent space needs to fulfill specific requirements. Each dimension should represent a basic visual concept such as orientation, size, or color [33]. Close samples from the input space should produce spatially close embeddings in the latent space. Thereby, with the so-called *continuity*, the distance between two embeddings can be used to determine how similar two points from the input space are. Through a well-structured latent space, samples with a higher information gain should be further away from the previous embedded samples. Furthermore, changing factors such as sunlight or wind should not affect the embedding. The embeddings of the input images should only contain information about the surroundings.

Regular autoencoders lack the regularization to produce such a latent space. Therefore, three different models have been applied and evaluated, which add the regularizations for a structured latent space. In this work, a variational autoencoder (VAE) [10], a Wasserstein autoencoder (WAE) [34], and a spherical autoencoder (SAE) [35] are applied. All three models share the same architecture based on a stack of five convolutions. Each block consists of a batchnorm, convolutional layer, leakyReLU, and max-pooling layer. Additionally to their regularization term, all models were trained to reduce reconstruction errors.

The VAE uses the Kullback Leibler Divergence to encourage the encoder to learn a distribution $P(x|Z)$ close to the standard normal distribution $\mathcal{N}(0,1)$ [10]. It creates blurrier reconstructions compared to a normal autoencoder. However, its regularization improves the structure of the latent space by packing embeddings close to a standard normal. It produces a more continuous latent space than an autoencoder. Due to the conflicting objectives of having a low-resolution error and well-structured latent space, it is hard to train [34]. The learned distributions tend to overlap, which can lead to samples producing erroneous reconstructions.

A model with similar advantages, however easier to train, is the Wasserstein autoencoder (WAE)[34]. The encoder of the WAE learns to map the input data to a latent space distribution that approximates a target distribution. For this, the target distribution is often a Gaussian distribution or uniform distribution. The Wasserstein distance between input and output is minimized to train the WAE. This distance measures the cost of moving the probability mass from one distribution to another. The encoder and decoder are trained using an adversarial loss function, where the critic network is used to estimate the Wasserstein distance between input and output distribution. While the encoder and decoder networks are trained to minimize the estimated Wasserstein distance, the critic network is trained to maximize it.

Additionally, a spherical autoencoder (SAE), which regularizes the latent space by forcing the embeddings onto the surface of a hypersphere, is utilized. Initially, the SAE was proposed to improve the performance of the VAE in text processing tasks due to latent collapse [35]. The original implementation uses a von Mises-Fisher distribution [35]. In this work, the regularization term is simplified with $L_{reg} = ||f_{enc}(x)|| - 1$.

B. Scoring

The different autoencoder variants are used to generate the embeddings of the input. Each embedding needs to be scored using the previously collected samples to select the exploration points. This work evaluates two approaches for setting a novelty score for the samples. The first approach uses the Probability Density Function (PDF) of a probability distribution in the latent space, while the second calculates the distance in the latent space.

To use the novelty of the sample, the PDF of the latent distribution $P(Z)$ can be used. However, since this relies on the pre-trained model and can not be updated except through retraining with new samples, an approximation for $P(Z)$ is generated using the samples acquired to this point. A method that can be used for this is Gaussian Mixture Models (GMMs). The GMMs fit multiple normal distributions to a dataset. In the used implementation, the GMM is refitted every five data points. When the new samples were analyzed using the GMM, the sample with the lowest probability of belonging to the distribution was selected as the next exploration point.

The second approach directly uses the embeddings. Due to the structure of the latent space, it is assumed that the distance of embeddings in the latent space is proportional to the distance in the input space. This means that the minimum distance between the previously acquired samples and the new sample can be defined as a novelty score

$$s(x) = \min_{\hat{z} \in \hat{Z}} (\hat{z} - f_{enc}(x)), \tag{1}$$

where $\hat{Z}$ is the set of all previously collected embeddings. Then, the sample with the largest distance is used as the next exploration point. The idea is that this maximizes the coverage of the latent space, which should result in a high entropy over the sampled path.

Entropy is used to calculate the score of the resulting paths. Since the input is not discrete, the differential distribution is used [36]

$$h(X) = -\mathbb{E}[\log p(X)] = -\int_X p(x) \log p(x) dx. \tag{2}$$

In contrast to the discrete entropy, the differential entropy is not bound to $[0,1]$ and can also get negative. For a multivariate normal distribution $X \sim \mathcal{N}(\mu, \Sigma)$ this can be rewritten to

$$h(X) = -\mathbb{E}[\log p(X)] = \frac{n}{2} \log 2\pi + \frac{1}{2} \log |\Sigma| + \frac{1}{2} n. \tag{3}$$

The entropy of the embeddings is computed to measure the overall information gained along a path. Each sample is embedded, and the differential entropy for multivariant normal distribution from Equation 3 is used. When x describes the sampled measurements, and Z represents a random variable that models the current environment in the latent space, the distribution can be written as $p(x|Z)$. The aim is to maximize the knowledge of the environment by maximizing the information about Z. We assume the sampled path from the measurements x is $P(\hat{Z})$; this means $\hat{Z}$ is sampled from Z. Therefore, we can infer that the entropy over Z equals the combined entropy, therefore $h(Z) = h(\hat{Z}, Z)$. Then, the mutual information for the differential entropy [37] can be simplified as in Eq. 4.

$$I(\hat{Z}; Z) = h(\hat{Z}) + h(Z) - h(\hat{Z}, Z) \tag{4}$$

$$= h(\hat{Z}) \tag{5}$$

This means a high entropy for $P(\hat{Z})$ would mean a high information gain for $P(Z)$.

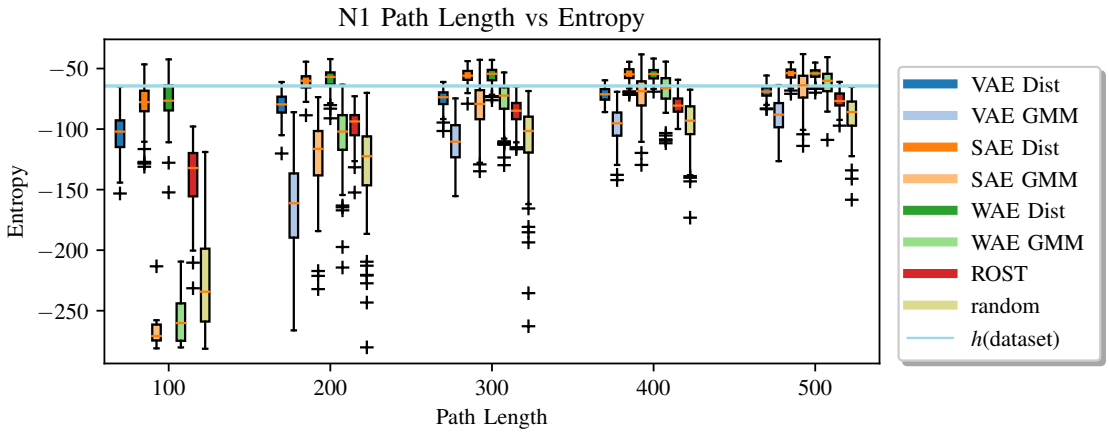

Fig. 3. The multivariant differential entropy over the path length for the different variants. For all variants, the entropy increases with longer paths. In comparison, the SAE and WAE with the latent distance perform the best. The GMM scoring does not work as well as the latent distance for all approaches. The SAE and WAE with latent distance even have a higher entropy than the overall dataset. Compared to the baselines of random tile selection and the ROST approach [27]m, the WAE and SAE perform significantly better. The VAE is comparable to the ROST framework from a performance point.

III. RESULTS AND EXPERIMENTS

For the training and the evaluation, two areas of the AHN2 dataset [38] as hosted by geotiles.nl [30] are used. The AHN2 dataset contains LIDAR measurements fused with aerial photography and satellite images. Therefore, it contains details on the height and color of each point. The two areas randomly chosen from the dataset are around Achtereind (N1) and the Kennemerduinen National Park (N2). The first consists of various urban areas with open fields, whereas the second is a more diverse landscape. Additionally, a dataset from an analog mission in the Spanish Tabernas desert is used. The dataset consists of a VDBMap with additional information [6] fused with drone images. A point cloud from a top view (2.5D) can be extracted from each dataset and used as input. The implemented autoencoders work on both point cloud and image data. Since using point cloud data in training and inference is more time-consuming, top-view images are generated from the top-view for the evaluation. Each image consists of roughly 3,000 tiles, which can be explored.

The VAE, WAE, and SAE all consist of a stack of convolutional and pooling layers. The decoder has the reverse structure as the encoder. The latent space for all of them has 64 dimensions. A rate scheduler and early stopping were utilized during the training, with the validation loss as the observed metric. In the experiments on the different datasets, each autoencoder variant was evaluated with the latent distance scoring and the GMM scoring. The results are compared against a slightly modified version of the ROST planner by Girdhar [27]. The parameters of the ROST planner are as described in their work. However, duplicated visited areas are not regarded in the entropy scoring since they would decrease the result for the planner.

For the evaluation, the top-view images were transformed into tiles of 64 x 64px. A random starting point was chosen for each round. The strategies then iteratively selected the next goal from all unexplored cells adjacent to the current position. This can be changed to all previously explored cells. However, this restriction was implemented to be comparable to the ROST framework. This was repeated up to a path length of 500 tiles and repeated 100 times for each strategy. The same series of starting points was used for each strategy to ensure fairness.

TABLE I
MULTIVARIANT DIFFERENTIAL ENTROPY OF THE THREE DATASETS FOR THE DIFFERENT APPROACHES

	N1	N2	Spain
VAE Dist	-69.31 ± 5.13	-67.70 ± 4.57	-42.84 ± 8.44
VAE GMM	-89.04 ± 13.13	-78.14 ± 9.66	-52.18 ± 11.42
SAE Dist	$\mathbf{-54.30 \pm 4.84}$	-59.62 ± 4.82	-37.27 ± 7.38
SAE GMM	-65.66 ± 14.22	-64.44 ± 7.49	-43.75 ± 10.05
WAE Dist	-54.59 ± 4.93	$\mathbf{-58.84 \pm 4.75}$	$\mathbf{-36.55 \pm 6.29}$
WAE GMM	-61.89 ± 11.23	-62.96 ± 6.80	-43.78 ± 9.81
ROST	-76.42 ± 7.16	-71.86 ± 5.58	-46.15 ± 6.41
random	-89.31 ± 16.76	-76.68 ± 9.66	-53.32 ± 11.58

The results in the AHN2 dataset around Achtereind (N1) are displayed in Fig. 3. The entropy is given over the path length. Each autoencoder is trained on the same dataset from Achtereind. The best-performing models are the SAE and WAE with the latent distance metric, with a slight edge for the WAE with latent distance. With a short path of 100, the results are already close to the entropy of the overall dataset. Both strategies have a better entropy than the overall dataset of around 3000 tiles once the paths are longer than 200 tiles. Once the path reaches the maximum of tiles, the entropy decreases towards the entropy of the dataset. That shows that tiles with similar information are not explored through a good selection.

In comparison, the VAE with latent distance performs not as well. The results for the three autoencoders, with the GMMs scoring, show that the entropy drops. It takes

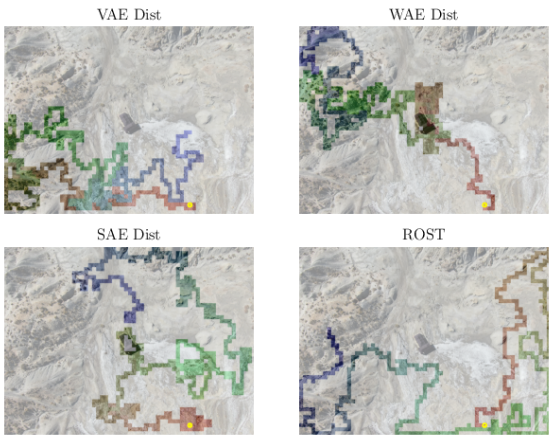

Fig. 4. A randomly selected path for the three strategies with the latent distance scoring and the ROST approach. Each strategy was pre-trained on the AHN2 dataset around Achtereind. The strategies cover different areas of the top view image of the Tabernas desert. The color range shows the exploration order, starting from red over green towards blue. According to the results, the WAE and SAE produce paths with higher information gain than the VAE and ROST.

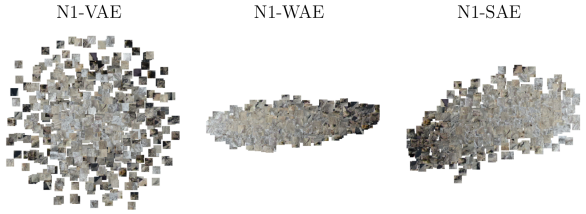

Fig. 5. Visualization of the latent space of the different autoencoders of the embeddings of the Tabernas desert. Each autoencoder is pre-trained on the N1 data from the Netherlands. The visualization is done with t-SNE [32]. The VAE differentiates more between brighter and darker tiles, whereas the other two have a more defined latent space. The clearly defined latent space of the WAE and SAE helps in scoring new tiles and selecting an appropriate next exploration goal.

a longer path to reach the dataset entropy. The path planner from Gridhar et al. [27] outperforms the GMM models with shorter paths. However, it saturates quickly and does not reach the dataset entropy even with a path of 500. Finally, the random sampling baseline is worse than most models. However, it is better than the VAE with GMM scoring.

Similar results can be observed in the AHN2 dataset around the Kennemerduinen National Park and the analog mission in the Tabernas desert. The results of the different areas are displayed in Table I. The autoencoders are all trained on the data from Achtereind and put into different contexts. The results for each case are embedded into the latent space of the WAE, and then the multivariant differential entropy is calculated. A different autoencoder could be chosen for the embeddings; then, the entropy would change. However, the relative performance stays similar between the models. The SAE and WAE with the latent distance perform the best again. When viewing the results on the image directly, it appears that the SAE and the WAE focus on high-complexity areas, whereas the VAE targets brighter tiles. This could be the result of the VAE only learning the difference between bright and dark tiles, as can be seen in Fig.2 (visualized with TSNE [32]). In contrast, the SAE and WAE learn better representation in their latent space of different areas.

Next to the results from the available datasets, an analog mission was conducted in the Spanish Tabernas desert. The location is where the first ExoFit trials for the Rosalind Franklin Rover were conducted in 2018 [39]. In the desert, two trials were conducted. First, a drone image was used to sample paths from an initial starting position. In the second, a small extended VDBMap was generated and then used to sample the following exploration goals for the robot. In the trials, the different strategies are trained on the Achtereind area of the AHN2 dataset [38].

For the first evaluation, a drone image is taken from the desert. The desert contains different ground types, including salty, sandy, and rocky areas with some vegetation. This detailed top view image is then tiled into sections of 64x64 pixels. For the tests, 100 paths were generated, each from a different starting location. The results of the differential entropy(see Equation 3) are shown in Fig.6 and Table I. Similar to the experiments on the available datasets, the SAE and WAE with latent distance perform the best. The differential entropy for the WAE is at -36.55, slightly higher than the SAE's at -37.27. The GMM scoring models perform worse and have a higher standard deviation from the different runs. The ROST [27] approach is better than the baseline of the random goal selection and just below the different approaches with the GMM scoring, beating the VAE with GMM scoring.

One sampled path can be seen in Fig.4. The path segments are colored, beginning from red, going over green, and moving towards blue. Each method samples different areas of the image starting in the same region. In the selected paths, the VAE explores further toward the top, whereas both the ROST and WAE approaches sample more toward the right. The SAE variant explores first towards the right and then towards the left. It has to be noted that this is one randomly selected path, so the different starting locations will influence the area that will be explored. The resulting differential entropy of the methods over all paths shows that the SAE and WAE exploit the most different information from the available tiles.

The resulting embeddings of the different tiles of the Tabernas data are visualized in Fig.5. As before, the VAE produces a differentiation between brighter and darker tiles. Whereas both the SAE and WAE show a more defined latent space, where the salty areas, the sand, and the more rocky areas are all in their own defined area of the latent space,

The final experiment generated a sample VDBMap of the area, including RGB information. This data was then embedded into a WAE with a latent distance approach. Afterward, the VDBMap was projected into a 2.5D representation by projecting the voxels along the z-axis. Then, to ensure random sampling, the projection was sampled using an RRT approach, and each node of the RRT tree was scored using the embeddings. The RRT is not yet used for exploration

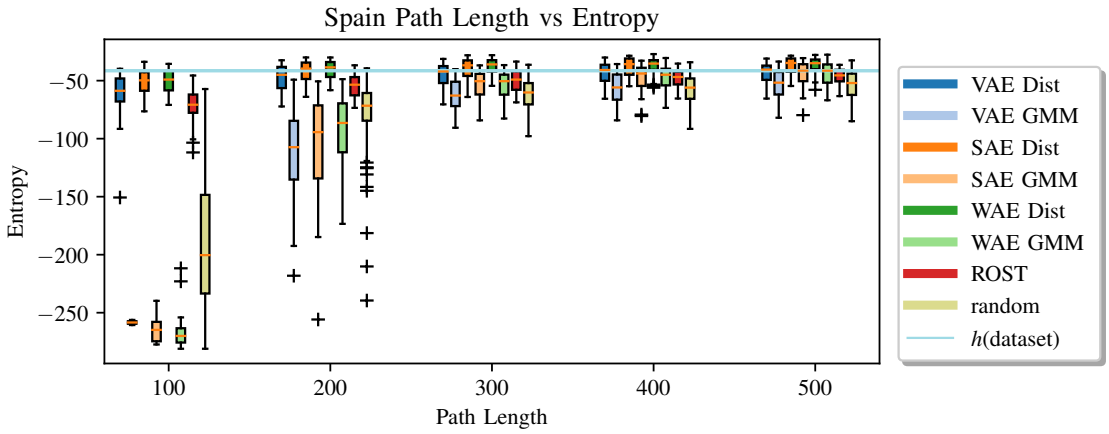

Fig. 6. The multivariant differential entropy over the path length for the different variants in the data from the analog mission in the Tabernas desert. For all variants, the entropy increases with longer paths. In comparison, the SAE and WAE with the latent distance perform the best. The GMM scoring does not work well with the latent distance for all approaches. The SAE and WAE with latent distance have a higher entropy than the overall dataset. Compared to the baselines of random tile selection and the ROST approach [27]m, the WAE and SAE perform significantly better. The VAE is comparable to the ROST framework from a performance point.

or path planning. An area around the node was taken as a tile and embedded (see Fig. 7). In yellow, areas with low information are presented. This might be due to occlusion from the map. So, if the tile does not contain enough information, it is marked. This clearly shows the frontier of the area, which could then be explored. Additionally, nodes with enough information are colored blue to red, where red depicts nodes that yield a high information gain. Especially man-made structures have a high information gain.

The results of the different experiments show that the approach with the WAE and SAE with the latent distance scoring outperforms the VAE encoder with both scoring methods and the GMM scoring for the WAE and SAE. The proposed methods select a subset from the area and can even increase the multivariant differential entropy compared to the overall dataset. This is due to the selection of more diverging samples compared to the complete dataset with similar entries. Compared to the ROST method by Girdhar et al. [27], better results could be achieved by embedding into the WAE and SAE with latent distance scoring. By combining existing methods, such as an RRT exploration, the nodes can be scored using the presented approach, and a value or priority can be assigned to each node. This shows that our proposed approach yields a higher information gain for the exploration goal than previous work and could aid mobile robots in their exploration missions.

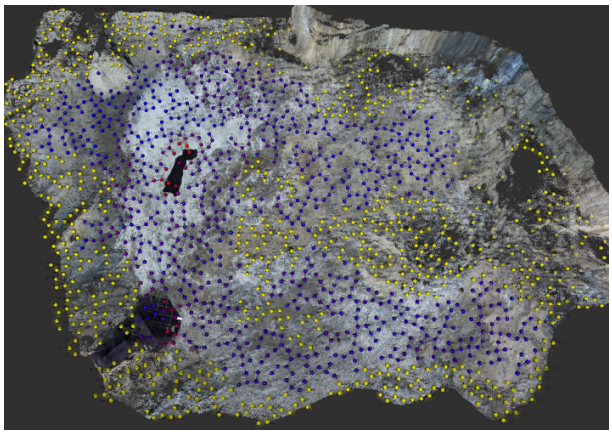

Fig. 7. Random sampling of a complete VDBMap using an RRT, each node is scored using the embeddings of the WAE with latent distance. The RRT is only for random sampling and is not yet used for any exploration or path planning. Nodes with insufficient measurements in the surrounding area are marked in yellow, showing the frontier. Nodes with high information gain are highlighted in red, while a low information gain is colored in blue. The man-made structures especially pose novel information for the exploration strategy.

IV. CONCLUSIONS AND FUTURE WORKS

In this work, three different autoencoder variants were evaluated for embedding and selecting exploration goals. The variational autoencoder (VAE), the Wasserstein autoencoder (WAE), and the spherical autoencoder (SAE) were evaluated on two datasets from the Netherlands and in an analog mission by using the multivariant differential entropy as measurement. Additionally, the approaches were compared against ROST by Girdhar et al. [27] and random sampling.

The work shows that using the embeddings from the SAE and WAE results in a clearly defined latent space. Using the latent distance as a metric to score the next best exploration goal was selected. This approach allows novel areas to be quickly explored, and the selected path even has a higher entropy than the complete dataset by reducing redundant information. The approach with the SAE and WAE outperforms the previously existing ROST method. By incorporating the scoring methods in existing exploration algorithms, such as an RRT, the nodes can be prioritized, and the next best view can be selected. The approach could be included in different exploration techniques such as frontier detection [18] to evaluate the frontiers better.

REFERENCES

[1] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *1985 IEEE International Conference on Robotics and Automation Proceedings*, vol. 2, Mar. 1985, pp. 116–121.

[2] P. Fankhauser and M. Hutter, "A universal grid map library: Implementation and use case for rough terrain navigation," *Robot Operating System (ROS) The Complete Reference (Volume 1)*, pp. 99–120, 2016.

[3] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: an efficient probabilistic 3D mapping framework based on octrees," *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, Apr. 2013.

[4] K. Museth, "VDB: High-resolution sparse volumes with dynamic topology," *ACM Transactions on Graphics*, vol. 32, no. 3, pp. 27:1–27:22, July 2013.

[5] S. Macenski, D. Tsai, and M. Feinberg, "Spatio-temporal voxel layer: A view on robot perception for the dynamic world," *International Journal of Advanced Robotic Systems*, vol. 17, no. 2, p. 1729881420910530, 2020, publisher: SAGE Publications Sage UK: London, England.

[6] M. Grosse Besselmann, L. Puck, L. Steffen, A. Roennau, and R. Dillmann, "VDB-Mapping: A High Resolution and Real-Time Capable 3D Mapping Framework for Versatile Mobile Robots," Aug. 2021, pp. 448–454.

[7] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, *et al.*, "Anymal-a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 38–44.

[8] B. Schölkopf, R. C. Williamson, A. Smola, J. Shawe-Taylor, and J. Platt, "Support Vector Method for Novelty Detection," in *Advances in Neural Information Processing Systems*, vol. 12. MIT Press, 1999.

[9] G. Pang, C. Shen, L. Cao, and A. v. d. Hengel, "Deep Learning for Anomaly Detection: A Review," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–38, Mar. 2022, arXiv:2007.02500 [cs, stat].

[10] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," May 2014, arXiv:1312.6114 [cs, stat].

[11] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.

[12] R. Hornung, H. Urbanek, J. Klodmann, C. Osendorfer, and P. van der Smagt, "Model-free robot anomaly detection," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sept. 2014, pp. 3676–3683, iSSN: 2153-0866.

[13] T. Schnell, K. Bott, L. Puck, T. Buettner, A. Roennau, and R. Dillmann, "Robigan: A bidirectional wasserstein gan approach for online robot fault diagnosis via internal anomaly detection," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4332–4337.

[14] M. A. Contreras-Cruz, J. P. Ramirez-Paredes, U. H. Hernandez-Belmonte, and V. Ayala-Ramirez, "Vision-Based Novelty Detection Using Deep Features and Evolved Novelty Filters for Specific Robotic Exploration and Inspection Tasks," *Sensors*, vol. 19, no. 13, p. 2965, Jan. 2019, number: 13 Publisher: Multidisciplinary Digital Publishing Institute.

[15] E. Özbilge, "On-line expectation-based novelty detection for mobile robots," *Robotics and Autonomous Systems*, vol. 81, pp. 33–47, July 2016.

[16] L. Puck, M. Schik, T. Schnell, T. Buettner, A. Roennau, and R. Dillmann, "Ensemble based anomaly detection for legged robots to explore unknown environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 511–517.

[17] I. Lluvia, E. Lazkano, and A. Ansuategi, "Active Mapping and Robot Exploration: A Survey," *Sensors*, vol. 21, p. 2445, Apr. 2021.

[18] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, July 1997, pp. 146–151.

[19] F. Bourgault, A. Makarenko, S. Williams, B. Grocholsky, and H. Durrant-Whyte, "Information based adaptive robotic exploration," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, Sept. 2002, pp. 540–545 vol.1.

[20] C. Dornhege and A. Kleiner, "A frontier-void-based approach for autonomous exploration in 3D," *Advanced Robotics*, vol. 27, no. 6, pp. 459–468, Apr. 2013.

[21] T. Cieslewski, E. Kaufmann, and D. Scaramuzza, "Rapid exploration with multi-rotors: A frontier selection method for high speed flight," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2017, pp. 2135–2142, iSSN: 2153-0866.

[22] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding Horizon "Next-Best-View" Planner for 3D Exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1462–1468.

[23] D. Zhu, T. Li, D. Ho, C. Wang, and M. Q.-H. Meng, "Deep Reinforcement Learning Supervised Autonomous Exploration in Office Environments," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 7548–7555, iSSN: 2577-087X.

[24] L. Tai and M. Liu, "Towards cognitive exploration through deep reinforcement learning for mobile robots," *arXiv preprint arXiv:1610.01733*, 2016.

[25] O. Zhelo, J. Zhang, L. Tai, M. Liu, and W. Burgard, "Curiosity-driven Exploration for Mapless Navigation with Deep Reinforcement Learning," May 2018, arXiv:1804.00456 [cs].

[26] L. Schmid, C. Ni, Y. Zhong, R. Siegwart, and O. Andersson, "Fast and Compute-efficient Sampling-based Local Exploration Planning via Distribution Learning," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7810–7817, July 2022, arXiv:2202.13715 [cs].

[27] Y. Girdhar, P. Giguère, and G. Dudek, "Autonomous adaptive exploration using realtime online spatiotemporal topic modeling," *The International Journal of Robotics Research*, vol. 33, no. 4, pp. 645–657, Apr. 2014, publisher: SAGE Publications Ltd STM.

[28] Y. Girdhar, D. Whitney, and G. Dudek, "Curiosity Based Exploration for Learning Terrain Models," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 578–584, arXiv:1310.6767 [cs].

[29] Y. Girdhar and G. Dudek, "Exploring Underwater Environments with Curiosity," in *2014 Canadian Conference on Computer and Robot Vision*, May 2014, pp. 104–110.

[30] A. v. N. (fwrite.org), "Geotiles," Oct 2020. [Online]. Available: https://geotiles.nl/

[31] L. Puck, T. Schnell, T. Buettner, A. Roennau, and R. Dillmann, "Intellirisk 2: Risk-Aware Exploration Through Active Learning, Anomaly Detection and Episodic Memories," 2022.

[32] L. van der Maaten and G. Hinton, "Viualizing data using t-SNE," *Journal of Machine Learning Research*, vol. 9, pp. 2579–2605, Nov. 2008.

[33] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework," July 2022.

[34] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf, "Wasserstein Auto-Encoders," Dec. 2019, arXiv:1711.01558 [cs, stat].

[35] J. Xu and G. Durrett, "Spherical Latent Spaces for Stable Variational Autoencoders," Oct. 2018, arXiv:1808.10805 [cs].

[36] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Nov. 2012, google-Books-ID: VWq5GG6ycxMC.

[37] I. Kontoyiannis and M. Madiman, "Sumset and inverse sumset inequalities for differential entropy and mutual information," *IEEE transactions on information theory*, vol. 60, no. 8, pp. 4503–4514, 2014.

[38] Ahn, "Open data," Apr 2022. [Online]. Available: https://www.ahn.nl/open-data

[39] M. Balme, S. Schwenzer, B. Dobke, M. Azkarate, D. Juan, and D. Ludovic, "The exofit rover field trial-simulating exomars rosalind franklin rover operations," Copernicus Meetings, Tech. Rep., 2020.