

A Step towards Explainable Artificial Neural Networks in Image Processing by Dataset Assessment

Nina Felicitas Heide*, Alexander Albrecht*,
and Michael Heizmann^{+,*}

*Fraunhofer IOSB, Fraunhofer Center for Machine Learning,
Fraunhoferstraße 1, Karlsruhe, Germany
{*nina.heide, alexander.albrecht*}@iosb.fraunhofer.de
⁺ Karlsruhe Institute of Technology,
Institute of Industrial Information Technology,
Kaiserstraße 12, Karlsruhe, Germany
michael.heizmann@kit.edu

Abstract We propose a methodology for generalized exploratory data analysis focusing on artificial neural network (ANN) methods. Our method is denoted *IC-ACC* due to the combined assessment of information content (*IC*) and accuracy (*ACC*) and aims at answering a frequently posed question in ANN research: “What is good data?” As the dataset has the primary influence on the development of the model, *IC-ACC* provides a step towards explainable ANN methods in the pre-modeling stage by a better insight in the dataset. With this insight, detrimental data can be eliminated before a negative influence on the ANN performance occurs. *IC-ACC* constitutes a guideline to generate efficient and accurate data for a specific, data-driven ANN method. Moreover, we show that training an ANN for the semantic segmentation of 3D data from unstructured environments with *IC-ACC*-assessed and -customized training data contributes to a more efficient training. The *IC-ACC* method is demonstrated on application examples for the visual perception of robotic platforms.

Keywords Artificial neural networks, image processing, pre-modeling explainability, robot vision systems

1 Introduction

In the development of classic methods – without the use of artificial intelligence (AI) – the scientist defines the behavior of a method by domain knowledge. In contrast to this, AI methods can be separated into data-driven and model-driven methods [1]. Initial design considerations in ANNs are specified with expert knowledge. Apart from this, the input data constitutes the major impact on the performance of a data-driven ANN approach [2]. To develop powerful AI methods, it is advisable to examine the input data in the pre-modeling stage.

We classify the data depending on the target application of the ANN. The 2D imaging domain can be divided into segmentation, depth estimation, object detection and tracking, and classification. 3D imaging splits into segmentation, object detection and tracking, shape classification, and registration [3,4].

In explainable AI research, the stages of explainability are subdivided into pre-modeling explainability, explainable modeling, and post-modeling explainability [2]. Pre-modeling explainability includes exploratory data analysis, dataset description standardization, dataset summarization, and explainable feature engineering. So far, most methods for exploratory data analysis examine and summarize the main characteristics and focuses on statistic parameters such as the Google Facets toolkit which maps the characteristics into numeric and categorical features.

Dataset assessment methods can be subdivided into adversarial testing methods, testing methods based on model and data coverage, and testing based on metrics [5]. In coverage testing, a high quality of a dataset is derived from a high percentage of activated neurons [5]. Testing based on metrics mainly focuses on the prediction accuracy of the ANN method and dismisses the in-depth analysis of the underlying dataset. Most research focuses on only one target application such as classification in [5] and [6]. [5] proposes an in-depth method to test the coverage of deep neural network models by examining the dataset quality with statistical measures such as centroid positioning. [6] detects class structure ambiguities in classification and proposes a reorganization strategy in case of decreasing

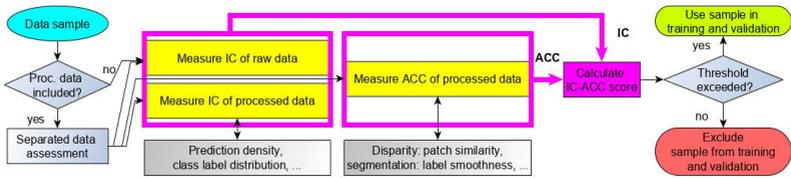


Figure 2.1: IC-ACC method: separated assessment of raw and processed data, the IC-ACC score decides whether to include the sample in the final dataset.

accuracy. Usually, the processed reference data is assumed as sufficiently accurate ground truth without verification.

Compared to the extensive research on ANN methods, data analysis for ANN methods is greatly underrepresented. With *IC-ACC*, we contribute a generalized, step-by-step approach in exploratory data analysis for pre-modeling explainability. We focus on the amount and diversity of the information inside the data which is required for proper ANN training as well as on the accuracy of its reference data for supervised learning approaches. Hence, we target the step prior to dataset assessment as proposed in [5] which analyzes biases or correlation among the variables. Only classical methods are considered in *IC-ACC* as AI-based data analysis would require additional assessment.

2 Exploratory Dataset Analysis with *IC-ACC*

Contrasting most works on exploratory data analysis, *IC-ACC* focuses on a generalized assessment of training data for ANN methods in the domain of image processing. The workflow of the proposed *IC-ACC* method is illustrated in Fig. 2.1. *IC-ACC* does not only provide a statistical measure for the diversity of the information such as the number of categories. But, we combine this with information content measures for images and point clouds, which can optimize the ANN performance, as well as with a first step towards an accuracy analysis of datasets.

2D and 3D data divides into raw and processed data: is the output of a sensor system after applying the intrinsic calibration and

is assumed to be error-free, processed data designates the reference data obtained in processing of the raw data which is usually utilized as training data. As the faultlessness of in data processing cannot be guaranteed, processed data can be subject to errors. With supervised training, an ANN learns to interpret raw data and requires processed data as a reference for the training loss. For unsupervised training, raw data is sufficient.

2.1 Information Content (IC) of Raw and Processed Data

To assess the information content of a message, Shannon [7] proposed the entropy measure $H = -(\sum_{N_I} p(i) \cdot \log_2(p(i)))$, with $i \in I$ a single symbol of all available symbols I , $p(i)$ the probability of the symbol to occur in the message, and N_I for $\#I$. In *IC-ACC*, the Shannon entropy is transferred on 2D and 3D data to assess the respective *IC*.

The *IC* of raw 2D data (IC_{r2D}) is contained in the intensity values of the captured spectral channels inside a pixel structure. The *IC* of a defined 2D pixel grid can be measured by its Shannon entropy. In 8-bit images, I contains all possible intensities $I \in \{0, 1, \dots, 255\}$.

Raw 3D data provides geometric information in 3D space with the position of each measurement point. Following [8] and [9], we regard point density and geometric structure as most conclusive criteria for the *IC* (IC_{r3D}). For point density, the density related to the distance from the sensor origin is chosen as the most promising representation [9]. 3D data is transformed from the Cartesian coordinates into homogenized coordinates ϕ , r , and z : $\phi = \arcsin(y/\sqrt{x^2 + y^2})$, $r = \sqrt{x^2 + y^2}$, and $z = z$. This yields a more uniform point distribution for active, rotating sensors [9]. A normalization in r and the binning of the homogenized points by their values of ϕ illustrates the point distribution, and thus the density. If notably different areas are included in each cloud, it is advisable to set a high number of bins. As a higher number of points naturally denotes a higher *IC*, also the total number of points as well as the number of points inside each previously defined bin can be applied to compare samples. We calculate the empirical mean to represent the relative distribution of the density values: $\mu = 1/N \sum_{i=1}^N x_i$, with N the number of bins, and x_i the relative density inside bin i . A uniform point distribution –

and thus a high μ – illustrates a proper representation of all cloud sectors and thus a high IC .

The structure of the point cloud can be described with the surface variation $s = \lambda_3/(\lambda_1+\lambda_2+\lambda_3)$, with λ the eigenvalues when decomposing the covariance matrix of a point set. s is calculated for each point and indicates the structured or unstructured character of a point cloud. To combine the values s of a point set, we calculate the Gaussian mean of s , denoted \bar{s} . In general, structured environments contain controlled, clearly separable topological objects as well as a high number of smooth surfaces. Unstructured environments are dominated by natural elements such as grassland, trees, bushes, or rocks [9, 10]. Hence, a higher \bar{s} indicates less structured elements. The future application environment defines if a high or a low IC measure is achieved: for more complex, unstructured environments, a high \bar{s} indicates a high IC , while in structured environments a high IC is synonymous to a clear structure and thus a low \bar{s} . This selection is justified in the subsequent proof of concept.

The IC of processed data depends on the prediction density and diversity of the information added during the processing step in relation to the number of 2D pixels or 3D points. For 2D data (IC_{p2D}), the prediction density is related to the number of pixels. An example for 2D prediction density is provided in stereo depth estimation: a high prediction density indicates a high percentage of valid depth estimates and thus a decent quality of the reference data [11]. For 3D data (IC_{p3D}), the number of points inside a point cloud is used accordingly. The diversity of the information can be measured using the Shannon entropy. Here, I contains all possible values of the added information such as class labels in segmentation tasks. The relative frequency of these values determines $p(i)$.

2.2 Accuracy (ACC) of Processed Data

ACC provides a measure of confidence and error characteristics for the data used as reference in supervised training. To overcome the common lack of a verified, error-free ground truth to compare against, we assess ACC with indirect measures. In contrast to IC , the evaluation of ACC has to be adapted to the type of the processed information to some extent. Two groups can be distinguished:

data that is used to train an ANN for similarity matching (ACC_{2Ds} , ACC_{3Ds}) and data for interpretation (ACC_{2Di} , ACC_{3Di}). Similarity includes depth estimation in 2D and the registration in 3D, whereas segmentation, object detection and tracking, and classification aim at the interpretation of imaging data.

In ACC_{2Ds} and ACC_{3Ds} , source and target to be matched have to be examined for similarity. The target remains in its original representation, the source is transformed by applying the reference data to be assessed. Following [11], the similarity of 2D samples for depth estimation is measured using the structural similarity index measure (SSIM) and the normalized root mean squared error (NRMSE). For 3D registration, the processed data typically consists of transformations. A high similarity between source and target cloud, after applying the reference transformation to the source, indicates a high ACC. Hence, difference measures such as L_1 norm, L_2 norm, or NRMSE are applicable. NRMSE generates a scale-invariant difference measure, which can be problematic in case of an undetected, different scaling of the input data. Hence, the L_2 norm is applied to assess the similarity [9].

For ANN methods in interpretation (ACC_{tDi}) the processed data typically consists of labeling information. One obvious strategy is to check a small number of random samples manually and to deduce a qualitative statement. This is a time-consuming, but often a straightforward strategy for experts. Currently, human annotators generate labeling data with the assistance of labeling tools and errors tend to occur in border regions or transitions between objects. As objects are rarely represented by a small number of points or pixels, a high number of different labels in a small area or space can indicate noisy and inaccurate data. As a first step towards a verifiable and quantitative ACC measure, pixel- and point-wise labels can be examined for smoothness, and thus for the existence of outliers. To identify outliers automatically, a nearest neighbor search can be applied similar to [8], requiring a minimum number of neighbors with identical labels. Also, a qualitative visual assessment of label smoothness is possible. A scoring from 0 to 10 allows a detailed rating for experts with domain knowledge [12]. Here, 10 stands for the highest ACC possible.

Table 1: Elements of exploratory data analysis with *IC-ACC* with proposed measures.

<i>IC-ACC</i> element	Measure
IC_{r2D} : raw 2D	Shannon entropy H
IC_{r3D} : raw 3D	Surface variation \bar{s} , relative density μ
IC_{p2D} : processed 2D	Prediction density (pixels), diversity H
IC_{p3D} : processed 3D	Prediction density (points), diversity H
ACC_{2Ds} : similarity	NRMSE, SSIM
ACC_{2Di} : interpretation	Qualitative visual assessment, label smoothness
ACC_{3Ds} : similarity	L_2 norm (MSE)
ACC_{3Di} : interpretation	Qualitative visual assessment, label smoothness

2.3 Deriving the *IC-ACC* Score

To derive a holistic *IC-ACC* score, each *IC* and *ACC* measure is normalized individually. Tab. 1 provides an overview of all *IC-ACC* elements. The *IC-ACC* score is calculated using $IC-ACC = 1/3 \cdot (IC_{rtD} + IC_{ptD} + ACC_{ptD})$, with $t \in 2, 3$. If more than one measure is included in an *IC-ACC* element, the average of both measures is considered. For normalization, the respective values are mapped to $[0, 1]$ using the maximum value of the respective measure. Depending on data availability, we recommend to distinguish weak, medium, and strong data inclusion thresholds for the *IC-ACC* score. In reference to $\mathcal{N}_{0,1}$, we include samples achieving more than 68.27 % of the possible maximum *IC-ACC* score of 1.0, hence it is *IC-ACC* score > 0.6827 for a weak threshold. The medium threshold is set to 0.8664 ($\mu \pm 1.5\sigma$), the strong threshold is 0.9545 ($\mu \pm 2\sigma$). Regarding one dataset in-depth, bad samples can be detected by applying the threshold on all elements of the dataset. To compare different datasets, the *IC-ACC* scores, prior to and after applying the inclusion requirements, can be compared.

2.4 Proof of Concept: *IC-ACC* for 2D and 3D Data

For the IC_{r2D} , the Shannon entropy H of an image is calculated. Fig. 2.2 shows the IC_{r2D} and ACC_{2Ds} assessment of image patches used to train a CNN for stereo matching on the KITTI 2012 dataset

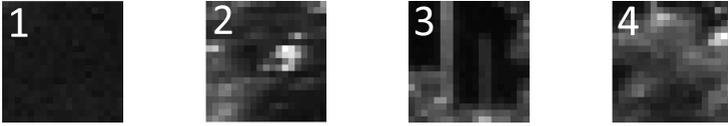


Figure 2.2: IC_{r2D} and ACC_{2Ds} for 19×19 pixel patches to train a depth estimation CNN: patch 1 has a low IC with $H_1 = 2.5$, whereas $H_2 = 6.24$ and $H_3 = 5.75$ indicate a high IC ; for pair 3-4 $SSIM = -0.04$ is sufficient, but $NRMSE = 1.15$ is too high for similarity [11] and the reference disparity is rated as inaccurate.

[13] using the reference disparity for ACC_{2Ds} as proposed in [11]. The prediction density concept for IC_{p2D} is illustrated on disparity maps in [11].

We demonstrate the IC analysis for raw 3D data on sequence (seq.) 00–10 of the SemanticKITTI dataset [4] using the raw 3D data of the KITTI Vision Odometry Benchmark [13] captured in urban and sub-urban areas. As terrain in urban and suburban areas mostly includes cultivated and rather structured terrain, only two of the 28 classes predominantly represent unstructured elements: vegetation and trunk.

In case of clearly separable sectors in point clouds, a subdivision into sectors improves the in-depth IC analysis. The clouds are divided into four sectors of 90° which are axisymmetric to the axes of the Velodyne LiDAR frame: front, right, back, and left. For seq. 02–04, \bar{s} of the left and right sectors is notably higher than \bar{s} of front and back: $\bar{s}_{left,02-04} = 0.0460$, $\bar{s}_{right,02-04} = 0.0649$, $\bar{s}_{front,02-04} = 0.0211$, and $\bar{s}_{back,02-04} = 0.0228$. Targeting on unstructured environments, this shows a higher IC for the left and right sectors and justifies the separation into structured and unstructured sectors for SemanticKITTI. Fig. 2.3 shows the point-wise estimates of s in scene 245 of seq. 04. Tab. 2 shows the \bar{s} and the class distribution for seq. 01 and 09 with the highest and seq. 06 with the lowest \bar{s} in seq. 00-10 of SemanticKITTI. Seq. 06 consequently has the lowest IC for unstructured target environments. As in seq. 06 only 9.77 % of the labels represent unstructured classes, compared to 23.91 % and 29.96 % in the seq. 01 and 09, this justifies the surface variation metric as a measure for the structured or unstructured character of a point cloud.

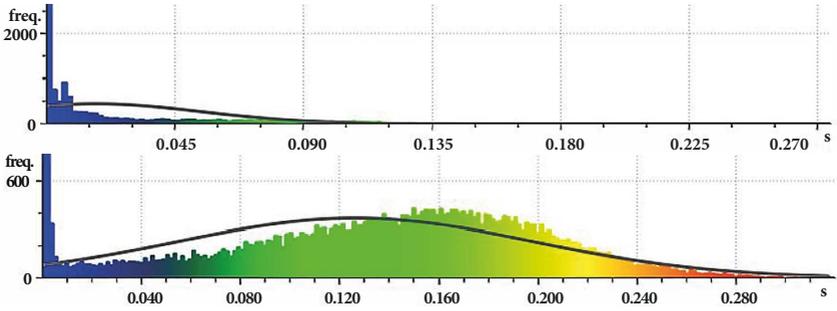


Figure 2.3: IC_{r3D} : histogram of surface variation s in the front sector (above) and right sector (below) of scene 245, seq. 04. The normal estimation radius is 0.40 m [8]. The low $\bar{s} = 0.0181$ of the front sector highlights its structured nature, while the high $\bar{s} = 0.1277$ of the right sector shows its unstructured character. Frequency for cut-off bin 0 is 13515 (front) and 3783 (right).

To demonstrate $IC-ACC$ in the comparison of two 3D clouds, we select scene 245 of seq. 04, denoted as (245,04), with a medium \bar{s} and scene 778 of seq. 09 (779,09) with a high \bar{s} . Tab. 3 illustrates the $IC-ACC$ results. The point density in IC_{r3D} is calculated using $N = 12$ bins, mapping 30° in one bin. ACC_{3Di} is demonstrated in qualitative manner. The renowned SemanticKITTI dataset comes with a high labeling accuracy and label smoothness which can both be verified manually. For IC_{p3D} , $N_I = 28$ is set in H with 28 classes in SemanticKITTI. In 245,04, the most frequent classes are road and vegetation with 35.46 % and 20.36 %. In 778,09 it is 26.89 % vegetation and 17.87 % building. The normalization values are derived from the maximums such as $\max(\bar{s}) = \bar{s}_{778,09}$ for s . We get $IC-ACC_{245,04} = 1/3 \cdot ((0.574 + 1.0)/2 + (0.828 + 1)/2 + 1.0) = 0.90$ and $IC-ACC_{778,09} = 1.0$. Applying a weak or medium threshold, both samples exceed the requirement with $86.64 \% < 90.0 \%$. For a strong threshold only 778,09 would be included in the final dataset.

2.5 Semantic Segmentation Efficiency with $IC-ACC$

The benefits of data analysis in training an ANN for the semantic segmentation of 3D point clouds is shown with SqueezeSeg [14, 15]. We follow the implementation of [15], but train with the full and

Table 2: \bar{s} and class distribution of points (in %) for lowest and highest \bar{s} in seq. 00-10.

\bar{s}	Vegetation	Trunk	Mainly unstructured	Terrain
$\bar{s}_{01} = 0.051$	23.87	0.04		23.91 13.83
$\bar{s}_{06} = 0.027$	9.31	0.46		9.77 26.10
$s_{09} = 0.051$	29.29	0.67		29.96 8.88

with a reduced version (seq. 02–04 training, seq. 08 validation) of SemanticKITTI [4]. Intersection over Union (IoU) is used to measure the segmentation performance. Training is conducted for 150 epochs. Representative classes are grouped into structured (car, road, parking, sidewalk, building, fence, pole, traffic sign) and unstructured, nature classes (vegetation, terrain, trunk). The full dataset achieves a mean IoU of $\bar{\text{IoU}} = 0.173$ in training and 0.210 in validation, the reduced dataset reaches 0.166 and 0.175. For the nature classes, it is $\bar{\text{IoU}} = 0.335$ on the reduced dataset, the structured classes reach $\bar{\text{IoU}} = 0.266$. This is remarkable as the nature classes are attributed to less than 30 % of the points present in the training data. It underlines the statement that unstructured data has a higher *IC* making it favorable in training and inference due to an increased unambiguity and a higher *IC*.

To rate the training efficiency, the customized segmentation efficiency metric $\eta_{\text{IoU}} = \text{IoU}/(D_T \cdot N_T)$ is proposed. The clouds are subdivided into structured (front, back) and unstructured (left, right)

Table 3: *IC-ACC* assessment of scene 245, seq. 04 and scene 778, seq. 09.

Measure	Raw measures		Normalization	
	245, 04	778, 09	245, 04	778, 09
$IC_{\text{r3D}}: \bar{s}$	0.027	0.047	$\frac{0.027}{0.047} = 0.574$	$\frac{0.047}{0.047} = 1.0$
$IC_{\text{r3D}}: \mu$	0.083	0.083	$\frac{0.083}{0.083} = 1.0$	$\frac{0.083}{0.083} = 1.0$
$IC_{\text{p3D}}: H$	2.405	2.903	$\frac{2.405}{2.903} = 0.828$	$\frac{2.903}{2.903} = 1.0$
$IC_{\text{p3D}}: \text{pred. density}$	100 %	100 %	$\frac{100\%}{100\%} = 1.0$	$\frac{100\%}{100\%} = 1.0$
$ACC_{\text{3D}}: \text{qual.}$	10, 10	10, 10	$\frac{10}{10} = 1.0$	$\frac{10}{10} = 1.0$
<i>IC-ACC</i> score	–	–	0.90	1.0

sectors. D_T represents the amount of data inside each tensor and is calculated for each scene with $D_T = hwc$. It is $h = 64$ and $w = 512$, height and width of the spherical projection, and $c = 5$ the number of features per point. The separation into sectors yields four projections with $4 \cdot D_T$ data points per cloud. N_T denotes the number of randomly selected scenes from the reduced dataset that are used in training. Hence, η_{IoU} measures the IoU in relation to the amount of training data. We test $N_T \in \{350, 700, 1050, 1750, 2800\}$. Overfitting is evaluated using the front and right sectors and can be prevented with $N_T \geq 1050$. For reference, it is $\eta_{IoU} = 3.1 \cdot 10^{-8}$ using all four sectors with $N_T = 2800$. Also with $N_T = 2800$, it is $\eta_{IoU} = 6.0 \cdot 10^{-8}$ using the right sectors only, $\eta_{IoU} = 3.9 \cdot 10^{-8}$ combining the two unstructured sectors, and $\eta_{IoU} = 5.2 \cdot 10^{-8}$ using the unstructured left and the structured front sectors. This shows that a notably higher η_{IoU} is achieved with a similar amount of training data, but with different structure. For seq. 02–04 of SemanticKITTI, the \overline{IoU} can be raised by more than 30 % combining data with different surface variations instead of data with a similar structure. Hence, the composition of IC-efficient datasets can improve the performance of ANN methods or reduce the amount of labeled training data required to achieve comparable results.

2.6 Guidelines for Data Generation

Naturally, guidelines for future data generation can be derived from the proposed IC-ACC method. We recommend to ensure that the captured data achieves a high IC and a high ACC. With this, the central point in the generation of data is fulfilled: it does neither contain too similar or too little, nor erroneous information. To ensure this, test samples can be assessed prior to capturing the final dataset. For 3D data, the target application of the ANN method defines the desired surface variation as previously stated. Capturing 3D for applications in unstructured environments such as in off-road robotics, a high \bar{s} is required, whereas a dataset for indoor scenes rather requires a low \bar{s} measure. Furthermore, we recommend to apply the ACC measures on the test data samples to verify a high ACC for the full, subsequently generated dataset.

3 Conclusion and Future Work

We present *IC-ACC*, a generalized methodology for exploratory data analysis for ANN methods in image processing. The *IC* examination can be applied to filter detrimental data and facilitates the composition of efficient datasets. The proposed *ACC* measures present a first step towards confidence and error assessment for supervised learning data. We demonstrate *IC-ACC* on ANN methods for robotic perception. Applying *IC-ACC* in the semantic segmentation of 3D data from unstructured environments shows an increased performance when using properly assessed and customized training data. Furthermore, *IC-ACC* presents a guideline for an efficient and less error-prone data generation. As data-driven AI methods can learn erroneous behaviors from erroneous training data, the analysis of the input data is an important step towards reliable and explainable AI methods.

Future works include evaluating *IC-ACC*-efficient data for ANN methods in image processing as well as extending *IC-ACC* to other domains.

Acknowledgement: The described research has been conducted within the competence center "ROBDEKON – Robotic Systems for Decontamination in Hazardous Environments", which is funded by the Federal Ministry of Education and Research (BMBF) within the scope of the German Federal Government's "Research for Civil Security" program.

References

1. R. Ashri, "Building AI software: Data-driven vs model-driven AI and why we need an AI-specific software development paradigm," URL: <https://hackernoon.com/building-ai-software-data-driven-vs-model-driven-ai-and-why-we-need-an-ai-specific-software-640f74aaf78f>, 2018.
2. B. Khaleghi, "The How of Explainable AI: Pre-modeling Explainability," URL: <https://towardsdatascience.com/the-how-of-explainable-ai-pre-modelling-explainability-699150495fe4>, 2019.
3. Y. Guo et al., "Deep Learning for 3D Point Clouds: A Survey," *arXiv preprint arXiv:1912.12033*, 2019.

4. J. Behley et al., "SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences," *ICCV*, 2019.
5. S. Mani et al., "Coverage Testing of Deep Learning Models using Dataset Characterization," *arXiv preprint arXiv:1911.07309*, 2019.
6. J.-D. Wang and H.-C. Liu, "An approach to evaluate the fitness of one class structure via dynamic centroids," *Expert Systems with Applications*, vol. 38, no. 11, pp. 13 764–13 772, 2011.
7. C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
8. N. Heide, T. Emter, and J. Petereit, "Calibration of multiple 3D lidar sensors to a common vehicle frame," *ISR*, 2018.
9. N. F. Heide et al., "UCSR: Registration and Fusion of Cross-Source 2D and 3D Sensor Data in Unstructured Environments," *Fusion*, 2020.
10. J. Petereit et al., "ROBDEKON: Robotic Systems for Decontamination in Hazardous Environments," *SSRR*, 2019.
11. N. F. Heide, S. Gamer, and M. Heizmann, "UEM-CNN: Enhanced Stereo Matching for Unstructured Environments with Dataset Filtering and Novel Error Metrics," *ISR*, 2020, in press.
12. N. F. Heide, A. Albrecht, and M. Heizmann, "SET: Stereo Evaluation Toolbox for Combined Performance Assessment of Camera Systems, 3D Reconstruction and Visual SLAM," *ICICSP*, 2019.
13. A. Geiger et al., "Vision meets robotics: The KITTI dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
14. B. Wu et al., "SqueezeSeg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3D lidar point cloud," *ICRA*, 2018.
15. A. Milioto et al., "RangeNet++: Fast and Accurate LiDAR Semantic Segmentation," *IROS*, 2019.