Post-clustering Prioritization Framework for Autonomous Decision Making in the Absence of Ground Truth via Hypothetical Probing

Wolfgang Fink[1] and Karm Al Hajhog[1]

[1] *Visual and Autonomous Exploration Systems Research Laboratory, College of Engineering, University of Arizona, Tucson, AZ 85721, USA*

wfink@arizona.edu
karm.hajhog@gmail.com

Abstract

A generic prioritization framework is introduced for addressing the problem of automated prioritization of region of interest or target selection. The framework is based on the assumption that clustering of preliminary data for pre-identified regions or targets of interest within an operational area has already occurred, i.e., post-classification, and that the clustering quality can be expressed as an energy/objective function. Region or target of interest prioritization then means to rank regions or targets of interest according to their probability of changing the energy/objective function value upon subsequent hypothetical probing as opposed to actually conducted reexamination, i.e., thorough follow-up or in-situ measurements. The mathematical formalism for calculating these probabilities to contribute to this change of the energy/objective function value is introduced and validated through numerical simulations. Moreover, these probabilities can also be understood as a confidence-check of the classification, i.e., the pre-clustering of the preliminary data. The operation of the prioritization framework is independent of the algorithm used to pre-cluster the preliminary data, and supports autonomous decision making. It is widely applicable across many scientific disciplines and areas, ranging from the microscopic to the macroscopic scale. Due to its ability to help maximize scientific return while optimizing resource utilization, it is particularly relevant in the context of resource-constrained autonomous robotic planetary exploration as it may extend the Remaining Useful Lifetime (RUL) – a key aspect in Prognostics and Health Management (PHM) – of space missions. On a more general, PHM-relevant level, the prioritization framework may provide an additional mechanism of identifying and correcting the maintenance status of system components to assist predictive maintenance or condition-based maintenance.

https://doi.org/10.36001/IJPHM.2024.v15i1.3206

1. Introduction

General scientific discovery, in some instances, can be thought of as driven by the reexamination of (pre-)clustered objects within an *operational area (OA)* given an initial clustering performed by a wide variety of supervised and unsupervised clustering approaches. Hereby, an OA can be a geologic field site on another planet, an agricultural area for precision agriculture, a medically imaged area (e.g., hyperspectral imaging), such as, but not limited to the fundus of an eye (e.g., Johnson et al., 2007), or a petri dish with bacterial or fungal cultures, etc.

Starting with a standoff/preliminary observation, a scientist (e.g., autonomous science craft or human scientist) seeks to better understand the OA through selected in-situ investigations, i.e., close-up measurements or investigations of a region or target of interest, that potentially reshape and alter the "knowledge" or understanding gained from the initial standoff observation. The question then arises which specific object(s) in the OA to reexamine. In an ideal situation the scientist or science craft will reexamine each identified object within that OA, but this may be prohibitive or unrealistic due to time and resource constraints, especially in the context of robotic planetary exploration, which is the motivation for this work. It is impractical and often even impossible, for example, to turn over every single rock on the surface of a planetary body, or to analyze each and every cell within a petri dish. Therefore, there is a need for a *prioritized* and *automated* selection of regions or targets of interest for in-situ or close-up/follow-up investigations to improve the overall understanding of the OA.

To that effect, this paper introduces a general purpose prioritization framework for identifying objects, i.e., regions or targets of interest, in a pre-clustered scene (the OA), which have the potential to increase the *"knowledge"* about the OA. Hereby "knowledge" is defined as a better *clustering quality*, i.e., tightness of the clustering as expressed by an energy/objective function (Section 2). The prioritization framework introduces the method of *hypothetical probing* whereby objects are hypothetically moved amongst the identified clusters by changing their respective cluster memberships and by adopting *hypothetical* rather than *actual* measurements to result in the calculation of a *probability* for that move to change the overall clustering quality. The hypothetical probing is not dependent on actual, additional measurements.

The proposed prioritization framework is a predictive method that operates in the complete absence of ground truth: It uses only preliminary or coarse data of an OA (e.g., gathered by standoff measurements) to predict probability-wise which previously identified objects have the potential to be game changers when subsequently examined in-situ or up-close.

In contrast to machine learning, confidence assessment, and predictive techniques that operate on noisy data or data with associated uncertainty (e.g., Hastie, Tibshirani, & Friedman, 2009), the proposed prioritization framework engages *after* these methods have been applied, i.e., the data at hand, represented as multi-dimensional feature vectors (Section 2; Fig. 2), are clustered according to their respective abilities. Hereby, the data are assumed to be only of as good a quality or resolution as preliminary or standoff measurement(s) can afford, e.g., from aboard an aerial platform above a planetary surface as opposed to an in-situ measurement, i.e., ground-truth. That means the quality/resolution of the data can only be increased through actual follow-up in-situ or close-up measurement(s) of ideally only a few well-selected regions or targets of interest. (Note, the term *resolution* hereby refers to both spatial and sensor/instrument resolution.) This is precisely where the proposed prioritization framework takes over by exploring via the *hypothetical probing* mechanism whether the clustering quality (assuming *soft clustering*) can be improved if:

a) The cluster membership of a feature vector of an identified object *hypothetically* changes or gets swapped but the preliminary measurement, i.e., the feature vector itself, remains unchanged, i.e., *no hypothetical* follow-up measurement is conducted;

b) The cluster membership of a feature vector of an identified object *hypothetically* changes or gets swapped *and* the feature vector is updated with a *hypothetical* follow-up measurement by assigning it the values of the centroid feature vector of the target cluster it is changing into because these are already known from the pre-clustering.

All possible configurations for the feature vector of a particular object that yield a better clustering quality this way compared to the pre-clustering one, contribute to the overall probability that this object is a worthwhile candidate for subsequent *actual* in-situ or close-up measurement(s). Figure 1 provides the motivation for and operational overview of the proposed prioritization framework that addresses this selection challenge.

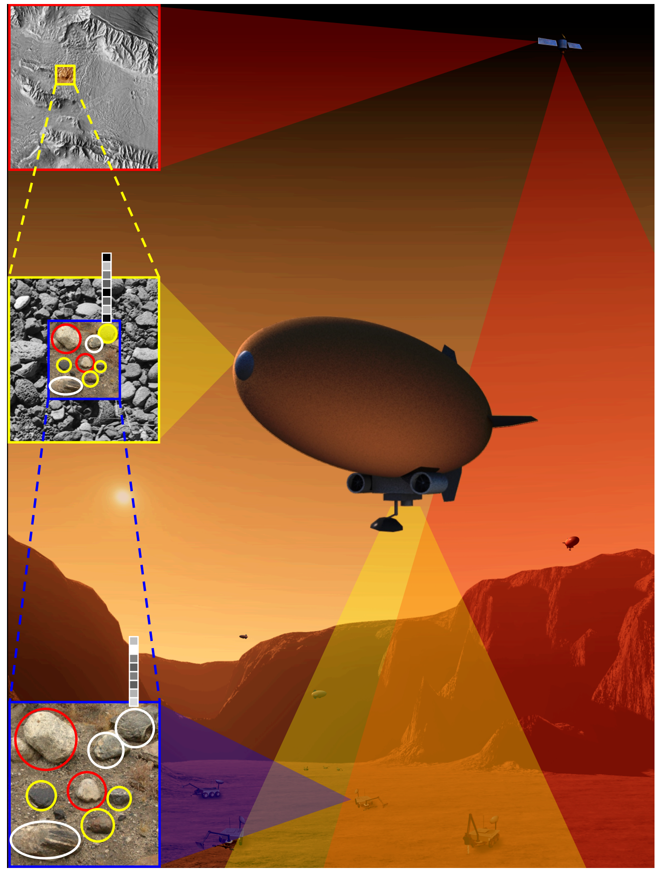

Figure 1: Motivation for and operational overview of the proposed prioritization framework, here in the context of multi-tiered and multi-agent autonomous robotic planetary exploration of Mars (e.g., Fink et al., 2005; Fink, Tarbell, & Jobling, 2008; Fink et al., 2022). The orbiter has a global view (at a lower spatial or sensor/instrument resolution but larger field of view) of the planetary surface and identifies an area for aerial platform deployment (here: blimp). The blimp obtains a higher resolution standoff view and assessment of an identified operational area and performs onboard preliminary clustering with established clustering techniques (red, yellow, and white circles). It subsequently engages the proposed prioritization framework that flags the upper right object as a candidate to be reexamined in-situ by one of the depicted ground-based rovers, which results in a correction of the cluster membership (yellow → white circle), i.e., a "knowledge" gain about the operational area. The two vectors indicate the preliminary blimp-view-based and the final ground-rover-based, i.e., ground-truth, feature vector (Section 2; Fig. 2) of the particular object that was flagged by the prioritization framework.

The generic prioritization framework presented here allows for an automated selection of objects with the highest potential/probability to alter the "knowledge" landscape without actually having to conduct reexaminations to make that selection, i.e., using preliminary data only. As such, it supports/augments autonomous decision making, e.g., as is required on autonomous science craft, with the ability to help maximize scientific return while optimizing resource utilization, especially in the context of usually resource-constrained (autonomous) robotic planetary exploration. The latter is a key aspect in the field of Prognostics and Health Management (PHM), as the prioritization framework proposed here has the potential to contribute to the extension of the *Remaining Useful Lifetime (RUL)* of space missions.

Potential cross-disciplinary application areas of such a generic prioritization framework range from autonomous planetary exploration, precision agriculture, microbiology, ophthalmology, to medical imaging, etc.

The paper presents and evaluates in the following four different prioritization scenarios using the method of hypothetical probing, derives the respective mathematical underpinnings (Sections 2-3), and assesses the effectiveness of each scenario in preliminary numerical simulations (Sections 4-6).

2. METHODS – *BACKGROUND*

Data about identified objects in an OA, i.e., regions or targets of interest, may be represented or described mathematically in form of multi-dimensional *feature vectors* gathered by standoff or preliminary examinations of an OA (Fig. 2). For example, in the context of planetary surface exploration in general and geologic field sites in particular, features of objects, such as rocks or rock formations, may comprise: *visual features*, such as color, albedo, texture, etc.; *geometric features*, such as moments, angularity, compactness, extent, circularity, eccentricity, size, etc.; and/or multi- or hyperspectral data – all in numerical form (Fig. 2; e.g., Fink, Brooks, & Tarbell, 2018; Fink et al., 2008).

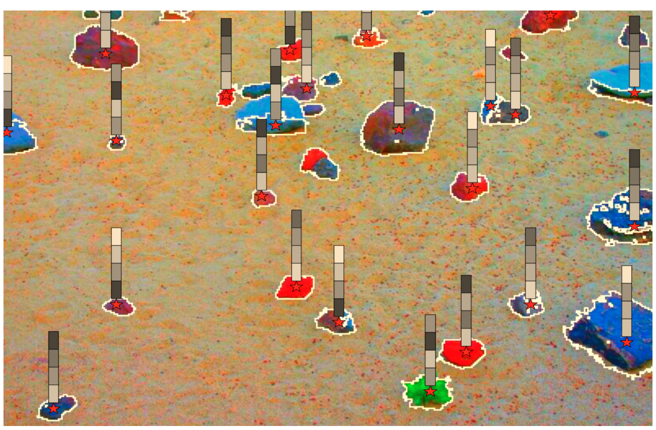

Figure 2: Operational area, i.e., regions or targets of interest, represented or described mathematically in form of multi-dimensional *feature vectors*, comprised of, e.g., visual, geometric, and (hyper)spectral features (e.g., Fink, Brooks, & Tarbell, 2018; Fink et al., 2008). [Source of slightly modified image: Fink et al., 2008]

Pre-identified objects described by such features/feature vectors can be prioritized for reexamination in multiple ways, such as, but not limited to: (a) by utilizing metrics such as a dot product or Hamming distance to compare the extracted feature data for all the identified objects within the OA; (b) by calculating the average and standard deviation for all the identified objects within the OA to identify feature vectors with a large deviation for subsequent high priority reexamination; or (c) by clustering.

Clustering algorithms, such as, but certainly not limited to, K-means (MacQueen, 1967; Duda, Hart, & Stork, 2000), Fuzzy C-means (Dunn, 1973; Bezdek, 1981; Sathishkumar, 2015), EM for mixture models (Bishop, 1995), hierarchical clustering (Williams, 2000), and the entire plethora of deep learning-based clustering approaches (e.g., Karim et al., 2021) depend on the presence of features being mathematically represented as feature vectors. Hereby each target is assigned a normalized membership value representing the membership confidence with respect to each occurring cluster (assuming *soft clustering*). This can be used to prioritize a target or sets of targets for close-up reexamination. For example, given an image containing a target that belongs to one of the pre-determined clusters with high confidence would be assigned a relatively low to medium reexamination priority, provided that sufficient images or measurements of this target type have been investigated. Clustering algorithms can also be used in conjunction with special-purpose models (Fink et al., 2001) to not only cluster the feature information and the spatial information of the encountered targets within an OA, but to organize these clusters into super-clusters, leading to a more global and comprehensive understanding of the OA. These special-purpose models are specifically tailored to the nature of the OA and the data types gathered within.

2.1. Clustering Quality: A Representation of "Knowledge"

In the following we derive an extensible prioritization framework, first introduced in part in Fink (2006), for objects, previously identified at a coarse/low resolution level in an OA, to be revisited more closely or in-situ for potential "knowledge gain" about the OA. It is assumed that preliminary feature or reconnaissance data about the regions or targets of interest have been gathered, and pre-clustering has occurred, using, for example, any of the clustering approaches listed above, etc. The quality of the data clustering can then be expressed in form of an *energy/objective function E* and can be formulated in more general mathematical terms as follows:

[Note: the next paragraphs until Section 3 are largely adopted from Fink (2006) and only slightly modified]

$$E(t) = \sum_{k=1}^{K} \sum_{i=1}^{N} M_{ki}(t)\left(\left\|\mathbf{c}_i(t) - \mathbf{cc}_k(t)\right\|^2 - \mu\right), \tag{1}$$

with $M_{ki}(t)$ denoting the membership value of an object $i \in \{1,\ldots,N\}$ with respect to cluster k at time t, with $0 \leq M_{ki}(t) \leq 1$ and the sum of all M_{ki} over all clusters $k \in \{1,\ldots,K\}$ being normalized to 1, $\mathbf{c}_i(t)$ the current feature vector of object i at time t, $\mathbf{cc}_k(t)$ the current cluster center vector (i.e., cluster centroid) in feature space (see definition below) of cluster k at time t, and μ a constant reward/penalty term that can be ignored without limitation of generality because of (2) below.

The value of the energy/objective function E is a measure for the "knowledge" about the objects within an OA and thereby the OA itself. Hence, a "knowledge gain" is defined as being synonymous with lowering the energy/objective function value, i.e., obtaining a tighter clustering. In this context, object prioritization then means to rank individual objects or sets of objects according to their probability to increase the "knowledge", i.e., to lower the energy/objective function E, upon close or in-situ reexamination. Using (1), a change in the energy/objective function E (i.e., ΔE) can be expressed as a difference between (a) the energy/objective function at a (future) time t, i.e., *hypothetical probing* (Section 3) of individual objects within the OA, and (b) the energy/objective function at a time $t-1$, i.e., pre-clustering of actually obtained/measured preliminary coarse or low resolution data of the identified objects within the OA:

$$\begin{aligned}\Delta E &= E(t) - E(t-1) \\ &= \sum_{k=1}^{K}\sum_{i=1}^{N} M_{ki}(t)\left(\left\|\mathbf{c}_i(t) - \mathbf{cc}_k(t)\right\|^2 - \mu\right) \\ &\quad - \sum_{k=1}^{K}\sum_{i=1}^{N} M_{ki}(t-1)\left(\left\|\mathbf{c}_i(t-1) - \mathbf{cc}_k(t-1)\right\|^2 - \mu\right) \\ &= \sum_{k,i}\begin{bmatrix} M_{ki}(t)\left\|\mathbf{c}_i(t) - \mathbf{cc}_k(t)\right\|^2 \\ -M_{ki}(t-1)\left\|\mathbf{c}_i(t-1) - \mathbf{cc}_k(t-1)\right\|^2 \end{bmatrix},\end{aligned} \tag{2}$$

with the cluster center vector $\mathbf{cc}_k$ calculated as follows:

$$cc_k(t) = \frac{\sum_{i=1}^{N}[M_{ki}(t) * c_i(t)]}{\sum_{i=1}^{N} M_{ki}(t)} \tag{3}.$$

It should be stressed that the notion of "time" merely denotes an instance at which the preliminary data were gathered and pre-clustered (i.e., $t-1$), and a subsequent instance where the hypothetical probing is performed (i.e., t). As such, t is an index describing steps of a process rather than actual time.

2.2. Calculation of Prioritization Probability

For a particular object $i^* \in \{1,\ldots,N\}$ and cluster $k^* \in \{1,\ldots,K\}$ the following indicator function can be defined:

$$\varphi_1\left(\Delta E\left(k^*, i^*\right)\right) := \begin{cases} 1 \text{ if } \Delta E\left(k^*, i^*\right) < 0 \\ 0 \text{ else} \end{cases}. \tag{4}$$

With the above definition (4), a probability $P(i^*)$ for a particular object i^* to lower the energy/objective function E upon changing its current cluster membership from $M_{ki^*}(t-1)$ to $M_{ki^*}(t)$ as well as its current feature vector $\mathbf{c}_{i^*}(t-1)$ to $\mathbf{c}_{i^*}(t)$ according to the prioritization scenarios derived below (Section 3) can be expressed as:

$$P\left(i^*\right) = \frac{\sum_k \varphi_1\left(\Delta E\left(k, i^*\right)\right)}{K}, \tag{5}$$

or, alternatively expressed as a *weighted* probability, which is the version adopted in the remainder of this manuscript:

$$P\left(i^*\right) = \frac{\sum_k \varphi_1\left(\Delta E\left(k, i^*\right)\right)\left|\Delta E\left(k, i^*\right)\right|}{\sum_k \left|\Delta E\left(k, i^*\right)\right|}. \tag{6}$$

As a side note, it is worthwhile pointing out that the computational effort for calculating the probabilities $P(i)$ for all N pre-identified objects in an OA is $O(N \cdot K)$, i.e., linear in N and K. This is particularly important in applications that are compute-constrained, such as multi-tier and multi-agent autonomous robotic planetary exploration (e.g., Fink et al., 2005; Fink, Tarbell, & Jobling, 2008; Fink et al., 2022).

3. MATHEMATICAL UNDERPINNING OF PRIORITIZATION FRAMEWORK – HYPOTHETICAL PROBING

The definition (2) above is the general expression for determining the change in clustering quality (ΔE) between two clustering states at time t and $t-1$. However, for certain cases of single object hypothetical probing, the mathematical definition of ΔE changes as object feature vectors and thus cluster center vectors change their values. In the context of *soft clustering*, we discuss in detail four cases (out of many more possible) in the following, for which we calculate ΔE to exemplify the workings of the prioritization framework.

3.1. Prioritization Case I

Case I considers the change in clustering quality (ΔE) when the (soft) membership of an object i^* in a selected target cluster k^* becomes 1, i.e., 100% membership, and zero for all other clusters (basically retroactive hard clustering): $M_{k^*i^*}(t)=1$, $M_{ki^*}(t)=0$ for all $k \neq k^*$, and without any changes (i.e., no hypothetical measurement) to the object-defining feature vector: $c_{i^*}(t)=c_{i^*}(t-1)$. After defining the object memberships of the target and source clusters, and keeping the memberships for all other objects i unchanged, we calculate ΔE as follows (note: all cluster center vectors $cc_k(t)$ change because of the membership change for i^* at t):

$$\Delta E = \sum_{k=1}^{K}\sum_{i=1}^{N}[\, M_{ki}(t)\|c_i(t) - cc_k(t)\|^2 - M_{ki}(t-1)\|c_i(t-1) - cc_k(t-1)\|^2]$$

3.2. Prioritization Case II

Case II considers the change in clustering quality (ΔE) when the membership of object i^* is switched between a target cluster k^* and a source cluster k^0: $M_{k^*i^*}(t)=M_{k^0i^*}(t-1)$ and $M_{k^0i^*}(t)=M_{k^*i^*}(t-1)$, but the object-defining feature vector c_{i^*} stays the same (i.e., no hypothetical measurement), i.e., $c_{i^*}(t)=c_{i^*}(t-1)$. In this case, ΔE is calculated in three parts (note: two cluster center vectors $cc_{k0}(t)$ *and* $cc_{k^*}(t)$ change because of the membership switch for i^*):

1. ΔE for all vectors except i^* in two clusters k^0 and k^* to capture the change in quality after the cluster centers $cc_{k0}(t)$ *and* $cc_{k^*}(t)$ are modified.
2. ΔE for the target vector i^* in the target cluster k^*.
3. ΔE for the target vector i^* in the source cluster k^0, since we excluded the calculation of ΔE for i^* in the first part.

Therefore, the change in clustering quality (ΔE) is as follows:

$$\Delta E = \underset{k=k^0,k^*}{sum}\left(\sum_{i\neq i^*}^{N}\left[M_{ki}(t)\|c_i(t)-cc_k(t)\|^2 - M_{ki}(t-1)\|c_i(t-1)-cc_k(t-1)\|^2\right]\right)$$
$$+M_{k^0i^*}(t-1)\,\|c_{i^*}(t)-cc_{k^*}(t)\|^2 - M_{k^*i^*}(t-1)\|c_{i^*}(t-1)-cc_{k^*}(t-1)\|^2$$
$$+M_{k^*i^*}(t-1)\,\|c_{i^*}(t)-cc_{k^0}(t)\|^2 - M_{k^0i^*}(t-1)\|c_{i^*}(t-1)-cc_{k^0}(t-1)\|^2$$

and after further simplification:

$$\Delta E = \underset{k=k^0,k^*}{sum}\left(\sum_{i\neq i^*}^{N}\left[M_{ki}(t)\|c_i(t)-cc_k(t)\|^2 - M_{ki}(t-1)\|c_i(t-1)-cc_k(t-1)\|^2\right]\right)$$
$$+M_{k^0i^*}(t-1)\left[\|c_{i^*}(t)-cc_{k^*}(t)\|^2 - \|c_{i^*}(t-1)-cc_{k^0}(t-1)\|^2\right]$$
$$+M_{k^*i^*}(t-1)\left[\|c_{i^*}(t)-cc_{k^0}(t)\|^2 - \|c_{i^*}(t-1)-cc_{k^*}(t-1)\|^2\right]$$

3.3. Prioritization Case III

Case III considers the change in clustering quality (ΔE) when the membership of object i^* is switched between target cluster k^* and source cluster k^0: $M_{k^*i^*}(t)=M_{k^0i^*}(t-1)$ and $M_{k^0i^*}(t)=M_{k^*i^*}(t-1)$, and the object-defining feature vector $c_{i^*}(t)$ changes to the cluster center vector of the target cluster: $c_{i^*}(t)=cc_{k^*}(t-1)$, i.e., hypothetical measurement is applied. In this case ΔE is calculated in four parts:

1. ΔE for all vectors except i^* in all clusters to capture the change in quality after all cluster centers are modified.
2. ΔE for the target vector i^* in the target cluster k^*.
3. ΔE for the target vector i^* in the source cluster k^0, since we excluded the calculation of ΔE for i^* in the first part.
4. ΔE for the target vector i^* with respect to the remaining clusters.

Therefore, the change in clustering quality (ΔE) is as follows:

$$\Delta E = \sum_{k=1}^{K}\sum_{i\neq i^*}^{N}\left[M_{ki}(t)\|c_i(t)-cc_k(t)\|^2 - M_{ki}(t-1)\|c_i(t-1)-cc_k(t-1)\|^2\right]$$
$$+M_{k^0i^*}(t-1)\left[\|cc_{k^*}(t-1)-cc_{k^*}(t)\|^2 - \|c_{i^*}(t-1)-cc_{k^0}(t-1)\|^2\right]$$
$$+M_{k^*i^*}(t-1)\left[\|cc_{k^*}(t-1)-cc_{k^0}(t)\|^2 - \|c_{i^*}(t-1)-cc_{k^*}(t-1)\|^2\right]$$
$$+\sum_{k\neq k^0;k\neq k^*}^{K}\left[M_{ki^*}(t)\|cc_{k^*}(t-1)-cc_k(t)\|^2 - M_{ki^*}(t-1)\|c_{i^*}(t-1)-cc_k(t-1)\|^2\right]$$

3.4. Prioritization Case IV

Case IV considers the change in clustering quality (ΔE) when the membership of object i^* for the target cluster k^* becomes 1, i.e., 100% membership, and zero for all other clusters (basically retroactive hard clustering): $M_{k^*i^*}(t)=1$, $M_{ki^*}(t)=0$ for all $k \neq k^*$, and the object-defining feature vector $c_{i^*}(t)$ changes to the cluster center of the target cluster: $c_{i^*}(t)=cc_{k^*}(t-1)$, i.e., hypothetical measurement is applied. In this case, ΔE is calculated in four parts:

1. ΔE for all vectors except i^* in all clusters to capture the change in quality after all the cluster centers are modified.
2. ΔE for the target vector i^* in the target cluster k^*.
3. ΔE for the target vector i^* in the source cluster k^0, since we excluded the calculation of ΔE for i^* in the first part.
4. ΔE for the target vector i^* with respect to the remaining clusters.

Therefore, the change in clustering quality (ΔE) is as follows:

$$\Delta E = \sum_{k=1}^{K}\sum_{i\neq i^*}^{N}\left[M_{ki}(t)\|c_i(t)-cc_k(t)\|^2 - M_{ki}(t-1)\|c_i(t-1)-cc_k(t-1)\|^2\right]$$
$$+\|cc_{k^*}(t-1)-cc_{k^*}(t)\|^2 - M_{k^*i^*}(t-1)\|c_{i^*}-cc_{k^*}(t-1)\|^2$$
$$-M_{k_0i^*}(t-1)\left\|c_{i^*}(t-1)-cc_{k_0}(t-1)\right\|^2$$
$$-\sum_{k\neq k_0;k\neq k^*}^{K}\left[M_{ki^*}(t-1)\|c_{i^*}(t-1)-cc_k(t-1)\|^2\right]$$

4. NUMERICAL SIMULATION & VALIDATION OF PRIORITIZATION FRAMEWORK

4.1. Generation of ground truth

To test the devised prioritization framework, we created a sample data file to simulate Cases I-IV on *blurred*, i.e., noise added, versions of the data (Section 4.2), thus simulating the preliminary coarse/low resolution data resulting from a standoff observation of an OA. For visualization simplicity, the sample file had 15 two-dimensional (2D) data/feature vectors, partitioned in groups of 5 feature vectors that are close to each other to form a cluster, which results in 3 clusters overall. To create this data file, we used a Gaussian distribution on 3 distinct cluster center vectors to randomly produce vectors that are close-by. Figure 3 shows the resulting 2D data vectors of the sample file.

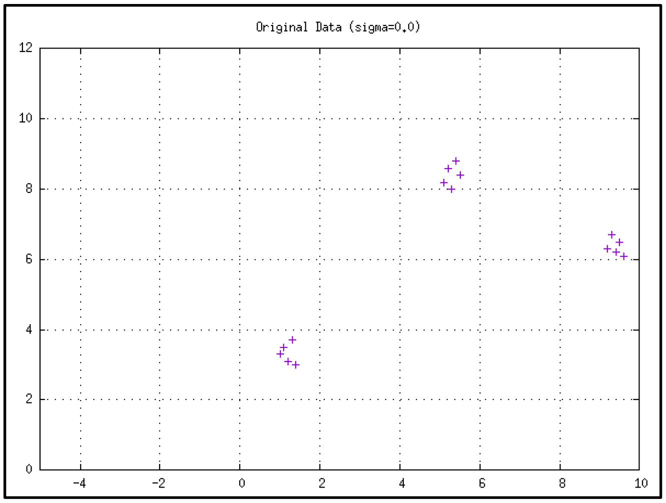

Figure 3: Plot of original, unclustered feature vectors with standard deviation s=0.0, i.e., no blurring.

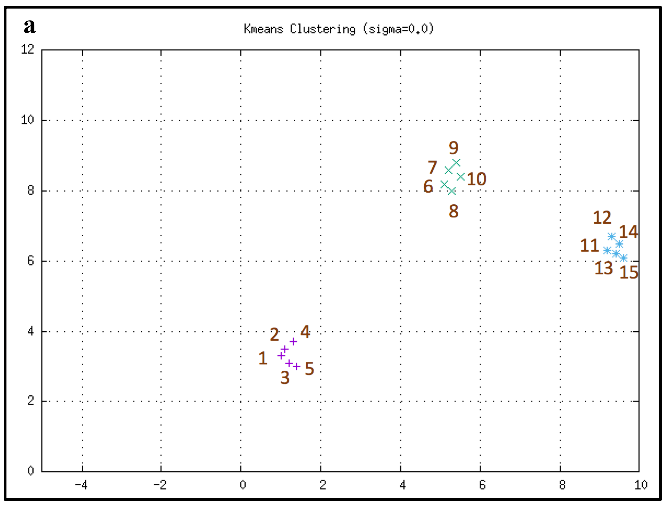

Figure 4a: K-means clustering results (here: 3 clusters) of original feature vectors with standard deviation s=0.0 (Fig. 3). The numbers indicate the respective feature vectors, and colors and symbols "+", "x", and "*" the respective clusters.

Subsequently, we clustered these data points into the predetermined number of clusters, i.e., 3 in this case. For simplicity, we used *K-means clustering* since the prioritization framework operates post-hoc on pre-clustered data and thus is independent of the clustering algorithm. The respective 5 vectors, which are relatively close to each other, were clustered together as expected (Fig. 4a). This resulting clustering is considered the *ground truth*, i.e., this clustering combination represents the goal configuration that is aimed for when applying the four cases to the *blurred* versions of the 15 feature vectors, representing the preliminary coarse/low resolution data resulting from the standoff observation of the objects in the OA.

4.2. Generation of standoff/preliminary observation of ground truth

To simulate the standoff/preliminary observation of the ground truth, we generated blurred versions of the original feature vectors by blurring each feature vector component-wise with a Gaussian distribution around a mean value of zero and a preselected standard deviation (sigma) to simulate the degree of standoff/preliminary observation quality, i.e., the degree of measurement resolution, respectively. In our case, we used standard deviations starting from zero (original data) up to sigma=3.0 in increments of 0.1 to study the performance of the prioritization framework (Cases I-IV) as a function of degree of blurriness, representing the standoff/preliminary observation quality or resolution (Figs. 4b-g).

Furthermore, similarly to the original unblurred ground truth feature vectors (sigma= 0.0), we pre-clustered these blurred feature vectors using K-means. As the feature vectors become increasingly blurred with increasing sigma, they move further away from their original feature vector locations, which ultimately results in cluster memberships different from the original, ground-truth K-means clustering, i.e., miss-classification. This step is essential for the simulation and validation of the prioritization framework as it compares the output of Cases I-IV to the feature vectors that are miss-classified compared to the "invisible" ground truth. The following graphs (Figs. 4b-g) show the initial K-means clustering for each of the blurred feature vector files; the feature vectors are numbered to keep track of each vector as it is increasingly blurred, i.e., increasing levels of noise added.

We tested these initially clustered data files with the prioritization Cases I-IV to check if the blurred vectors, which were wrongly clustered by K-means, show a higher respective probability to be chosen for reexamination (Section 2.2). The respective probability yields a percentage of how likely a vector is being "placed" in the "wrong" cluster, thereby becoming a candidate for in-situ, close-up/follow-up investigation. The results of such a close-up reexamination would stand a chance to yield the correct placement, e.g., through actual in-situ measurements (Fig. 1), thereby restoring the "invisible" ground truth, i.e., the "knowledge" about the operational area, which is the goal of the entire proposed prioritization framework.

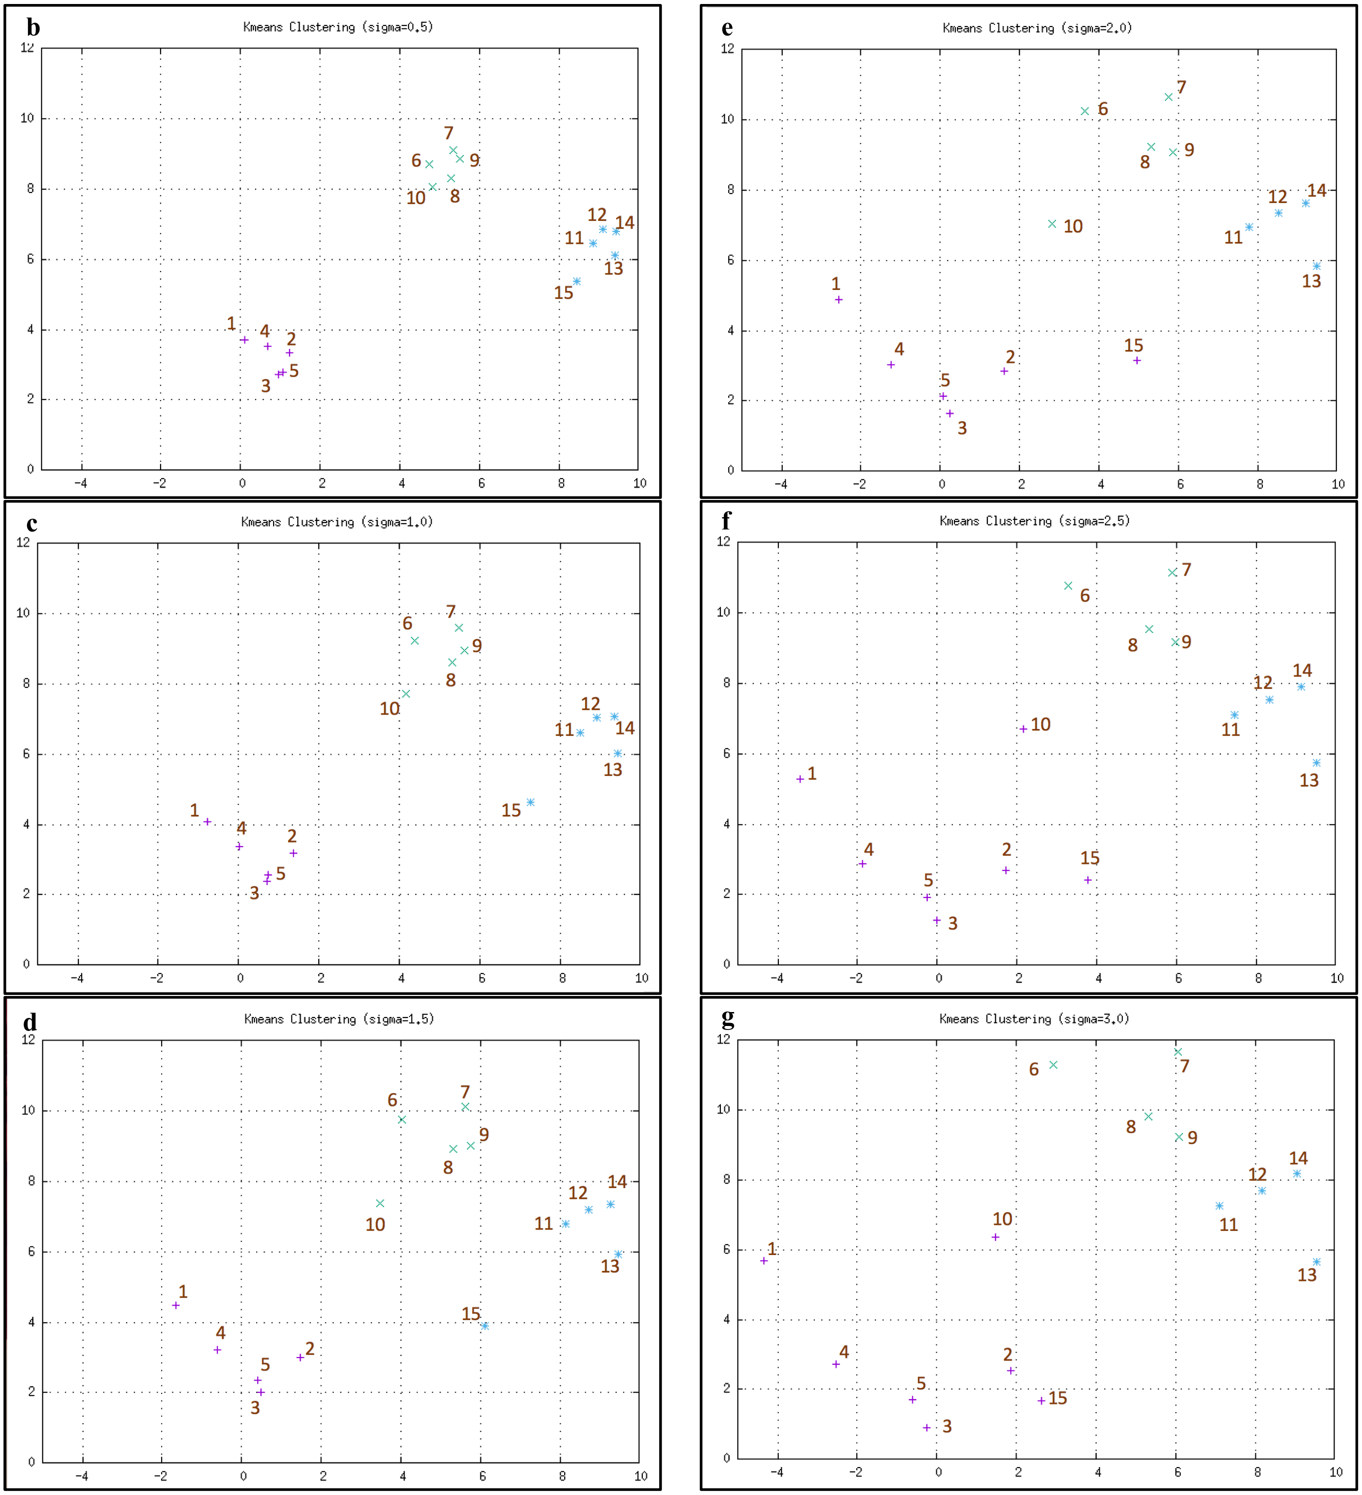

Figures 4b-g: K-means clustering results (here: 3 clusters) of blurred feature vectors with increasing sigma s: (b) s=0.5; (c) s=1.0; (d) s=1.5; (e) s=2.0; (f) s=2.5; and (g) s=3.0. The numbers indicate the respective feature vectors, and colors and symbols “+”, “x”, and “*” the respective clusters.

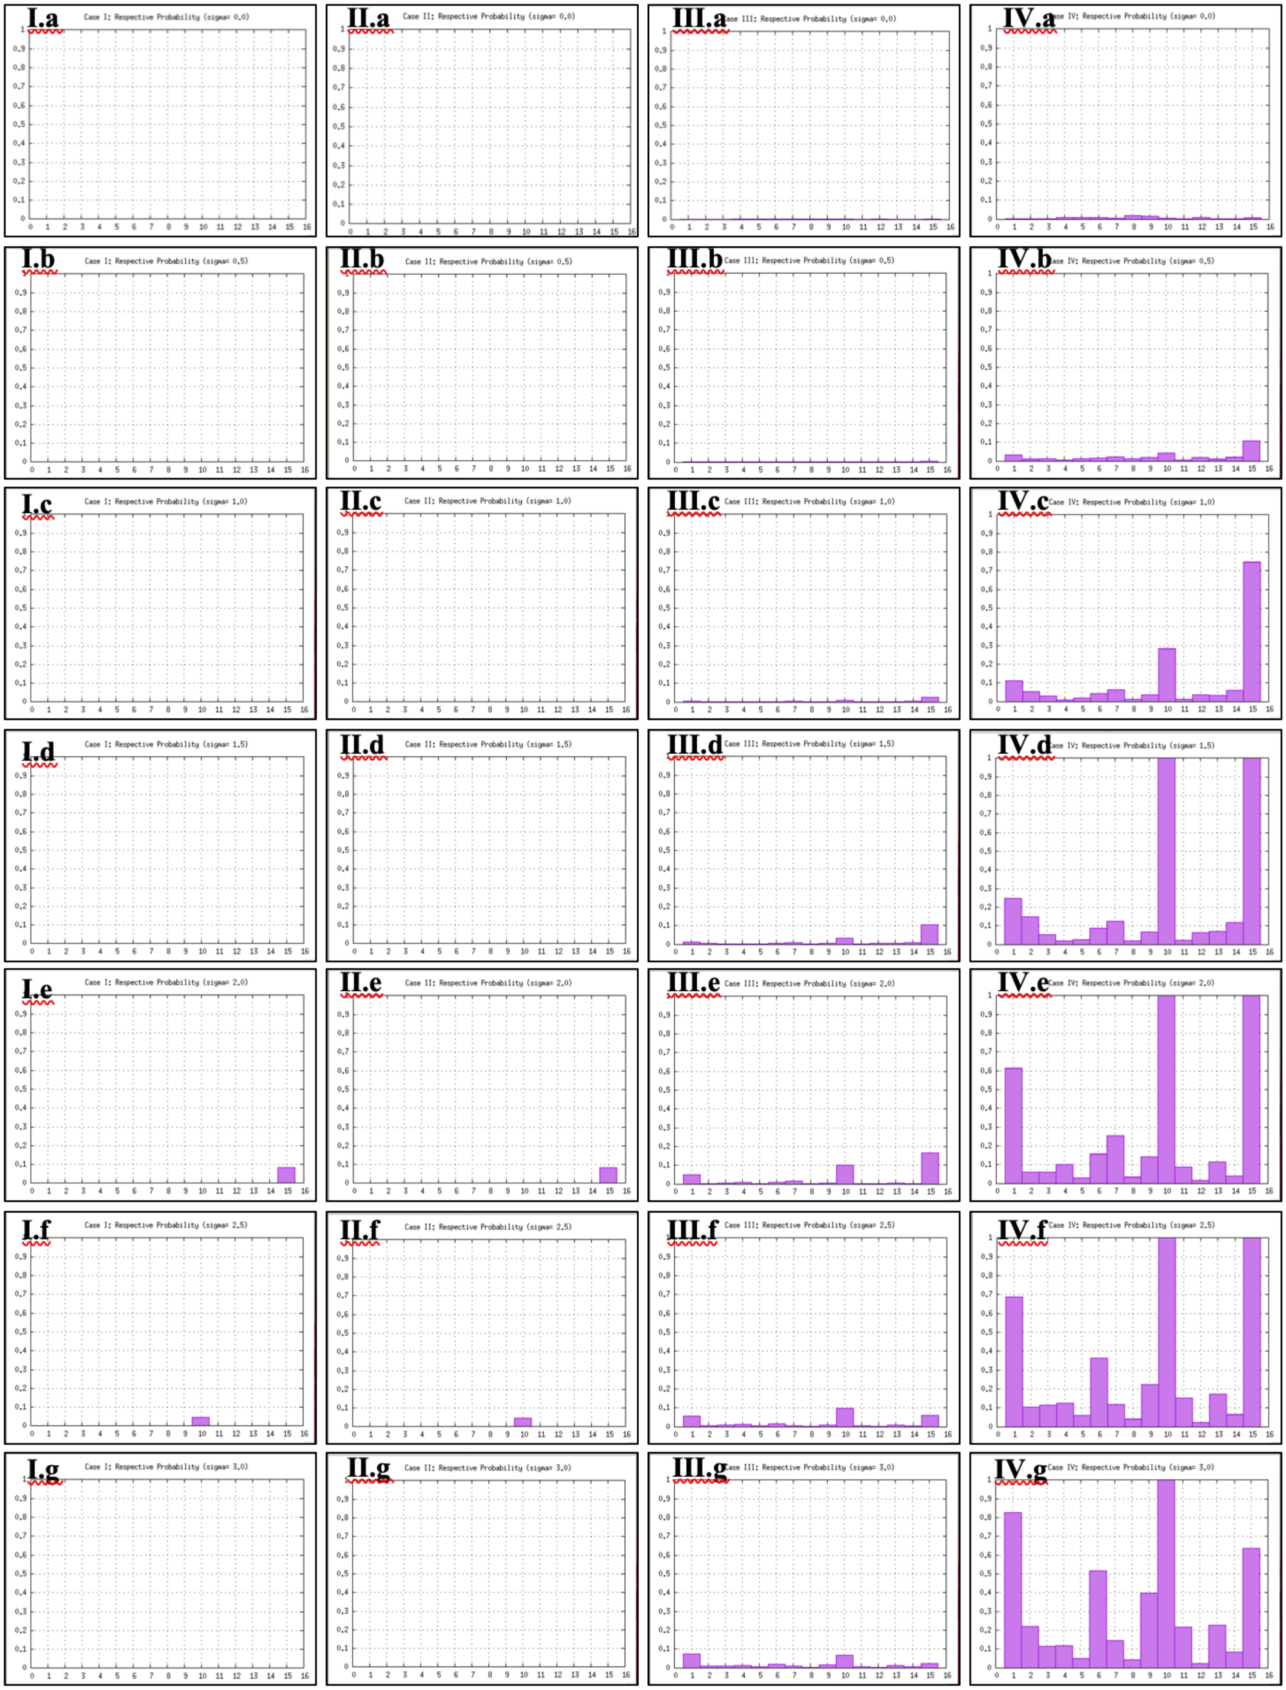

Figure 5: Probability values (y-axes) obtained with K-means clustering, i.e., *hard clustering*, for Cases I-IV and sigma s=0.0 to 3.0 for each feature vector 1 to 15 (x-axes). Column I corresponds to Case I; II to Case II; III to Case III; and IV to Case IV. Row a corresponds to s=0.0; b to s=0.5; c to s=1.0; d to s=1.5; e to s=2.0; f to s=2.5; and g to s=3.0.

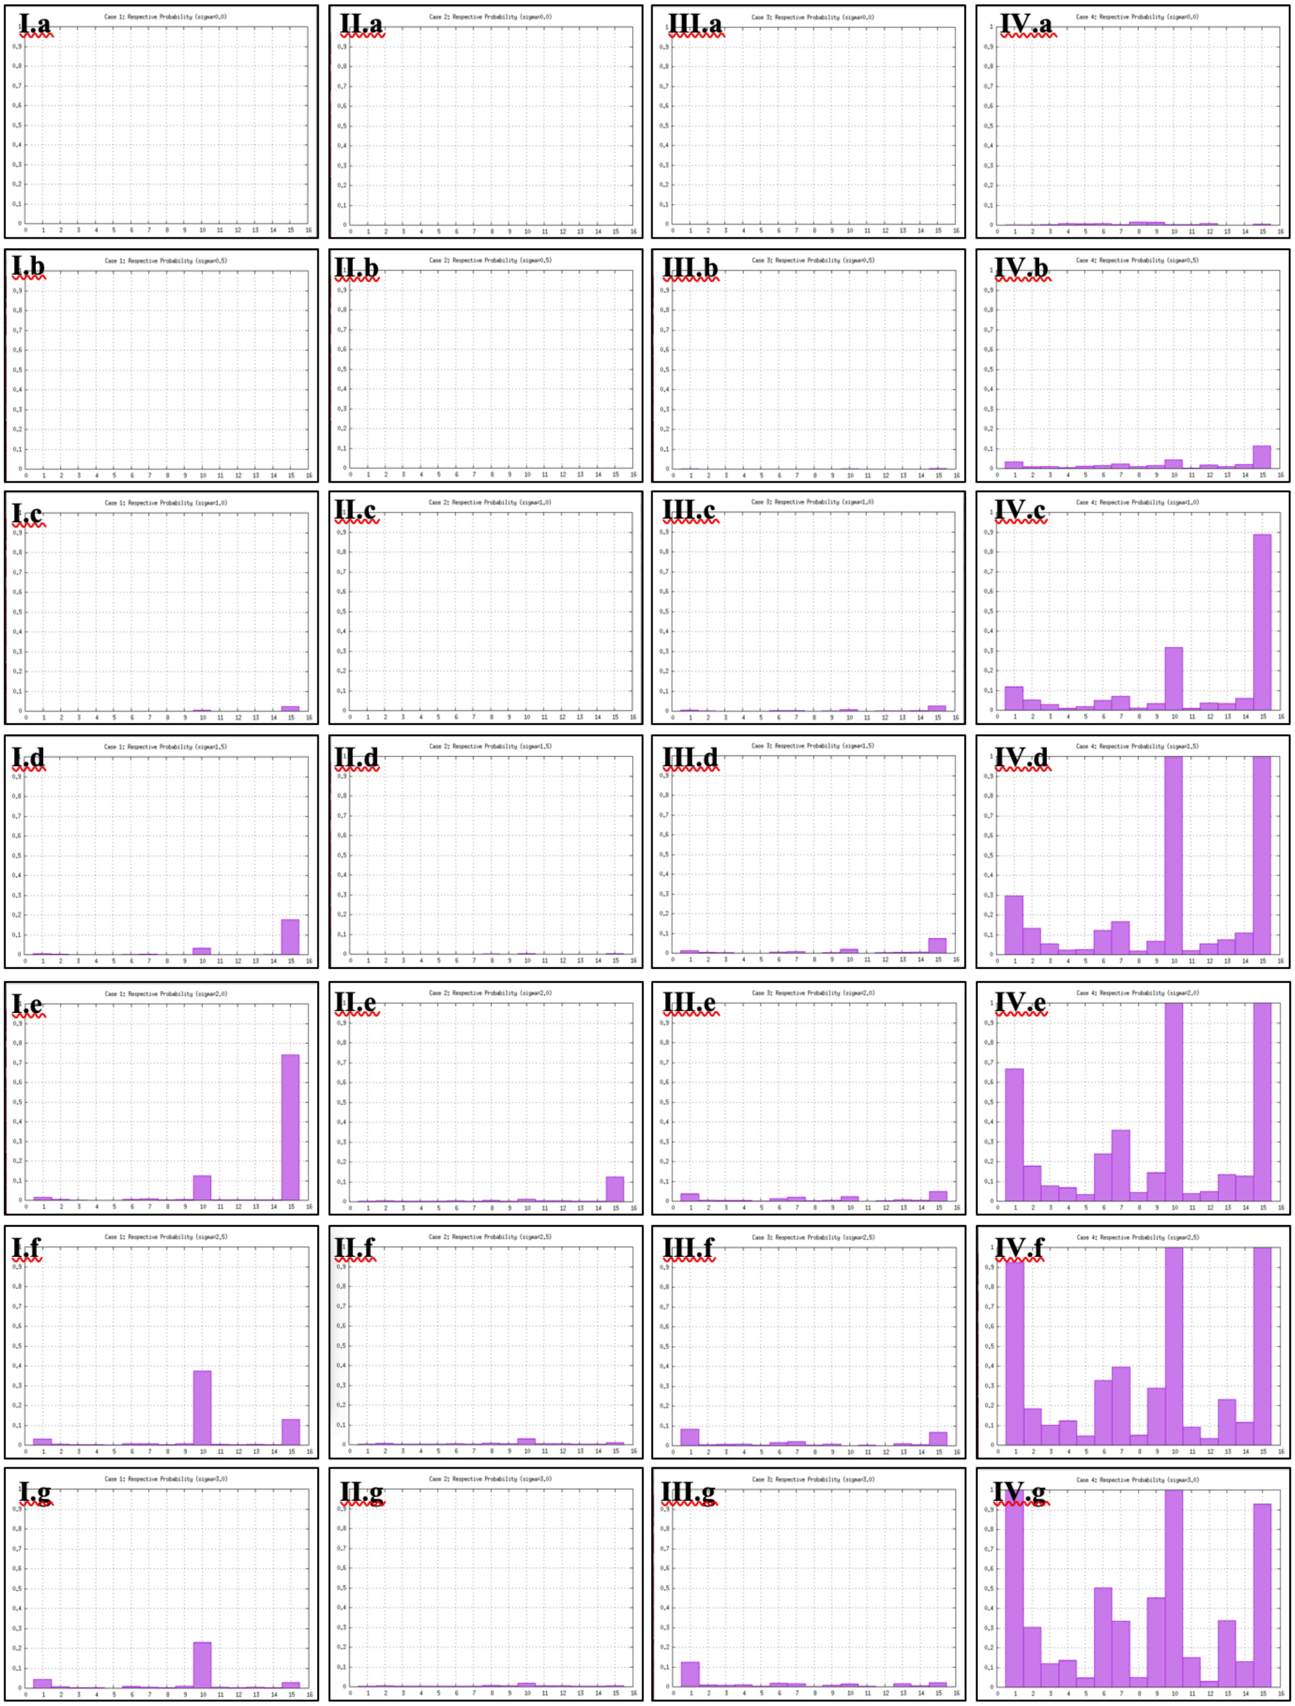

Figure 6: Probability values (y-axes) obtained with Fuzzy C-Means clustering (with fuzzification parameter m=1.25), i.e., *soft clustering*, for Cases I-IV and sigma s=0.0 to 3.0 for each feature vector 1 to 15 (x-axes). Column I corresponds to Case I; II to Case II; III to Case III; and IV to Case IV. Row a corresponds to s=0.0; b to s=0.5; c to s=1.0; d to s=1.5; e to s=2.0; f to s=2.5; and g to s=3.0.

5. SIMULATION RESULTS

The probability outputs of the respective prioritization Cases I-IV for each of the 15 feature vectors as a function of sigma s=0.0, 0.5, 1.0, 1.5, 2.0, 2.5, and 3.0, i.e., for various degrees of blurriness/resolution, are shown for *hard clustering* via K-means (Fig. 5) and for *soft clustering* via Fuzzy C-Means (Fig. 6), the latter being implemented according to Sathishkumar (2015) with fuzzification parameter m=1.25.

In the case of *hard clustering*, each feature vector is a member of exactly one cluster (membership value 1.0), respectively, and none of the others (membership value 0), both for time *t* (i.e., pre-clustering) and time *t-1* (i.e., hypothetical probing).

In the case of *soft clustering*, all feature vectors are members of all clusters, with their respective membership values adding up to 1.0 for time *t* (i.e., pre-clustering) and most of them for time *t-1* (i.e., hypothetical probing) except for the feature vector describing object i^* in Case I and IV (see Sections 3.1 and 3.4 above) which are governed by hard clustering.

We also ran the prioritization Cases I-IV for extreme sigma values of s=10.0 and s=100.0: as expected this leads to a general breakdown of the prioritization framework, i.e., random probability value assignments that yield no useful, i.e., actionable, information for each of the prioritization Cases I-IV, because the applied degree of blurriness is so high that the new cluster memberships become the natural "new homes" of the feature vectors as opposed to their original ground truth ones. In these cases, there is no chance of recovering/revealing the ground truth, at least in the context of the proposed prioritization framework.

6. DISCUSSION

In the following the respective ability of the prioritization Cases I-IV to recover/reveal ground truth is discussed to identify candidate feature vectors for targeted close-up/follow-up examination as a function of degree of blurriness: for *hard clustering* via K-means and for *soft clustering* via Fuzzy C-Means.

6.1. Prioritization using Hard Clustering (here: via K-Means)

As a reminder for the following discussion, the cluster membership values for all feature vectors are exactly 1.0 for one particular cluster, respectively, and zero for all other clusters.

As for **Case I**:
The respective probability values for revisiting any of the 15 feature vectors are zero for all sigma values up to s=2.0, thereby indicating no need for revisiting any of the feature vectors up close as the preliminary clustering would reflect ground truth (Fig. 5, column #1). At s=2.0, only feature vector 15 has a respective probability value that is greater than zero, indicating the need for revisiting it up close, which is commensurate with the fact that feature vector 15 has moved into a position that is halfway between its original, ground truth cluster, and the wrongly assigned target cluster. Similarly, at s=2.5 only feature vector 10 has a respective probability value that is greater than zero, indicating the need for revisiting it up close, because feature vector 10 has now moved halfway between its original, ground truth cluster, and the wrongly assigned target cluster. At that sigma value, feature vector 15 is now considered "native" in its wrongly assigned target cluster due to its proximity in feature space. The values of respective probabilities go back to zero for feature vectors 10 and 15 as sigma increases beyond 2.5 because the feature vectors are so close to the new, wrongly assigned target clusters that they are now considered "native" in them. At that point, a potential return to ground truth via Case I prioritization is no longer indicated by means of in-situ or close-up/follow-up examination.

As for **Case II**:
The reason why the results of Case II (Fig. 5, column #2) are the exact same as for Case I (Fig. 5, column #1) is because the simulation was run on an example with hard clustering (i.e., K-means). Therefore, when memberships of zeros and ones are switched (Case II) it yields the exact same results as giving a full membership of 1.0 to a feature vector in one cluster and zero memberships in all the other clusters.

As for **Case III**:
As opposed to Cases I and II, Case III assigns probabilities to all feature vectors simultaneously (i.e., not in sequence) starting with sigma as low as 0.0. Far more responsive than Cases I and II, Case III exhibits non-zero probability values as soon as the respective feature vectors move further away from their original clusters in feature space towards the ultimately wrongly assigned target clusters even while they are still being clustered correctly (Fig. 5, column #3). Moreover, it should be noted that while the probabilities for the wrongly classified feature vectors increase up until sigma=2.0, they do not drop to zero for larger sigma values, e.g., as is the case for sigma=3.0 in Cases I and II. This means that feature vectors that are now "firmly" embedded, i.e., considered "native," in their wrongly assigned target clusters, can still be retroactively identified and flagged for targeted in-situ or close-up/follow-up examinations to potentially return to ground truth.

As for **Case IV**:
Similarly to Case III but much more amplified (i.e., all the way to probability saturation at 100%), Case IV assigns probabilities to all feature vectors simultaneously (i.e., not in sequence) starting with sigma as low as 0.0. As soon as the respective feature vectors move further away from their original clusters in feature space towards an ultimately wrong target cluster assignment due to increased blurring, Case IV assigns rapidly increasing probabilities in the correct order of

occurrence, even while the feature vectors are still being clustered correctly according to ground truth (Fig. 5, column #4). The prioritization Case IV results in the highest respective probability values. For example, feature vectors 10 and 15 have a respective probability of 100% (or 1.0) when sigma is between 1.5 and 2.7 with only a slight decrease over all feature vector probabilities beyond sigma = 2.7. Again, this means that feature vectors which are now "firmly" embedded, i.e., considered "native," in their wrongly assigned target clusters, can still be retroactively identified and flagged for targeted in-situ or close-up/follow-up examinations to potentially return to ground truth.

6.2. *Prioritization using Soft Clustering (here via Fuzzy C-Means with fuzzification parameter m=1.25)*

In the following discussion, all feature vectors are members of all clusters, with their respective membership values adding up to 1.0.

Cases II-IV (Fig. 6, columns 2-4) behave qualitatively and for the most part also quantitatively just like the prioritization counterparts using hard clustering (Fig. 5, columns 2-4). The only significant difference is between Cases I and II, which are behaving differently now when using soft-clustering (Sections 3.1 and 3.2 above): in particular Case I becomes now far more pronounced and responsive probability wise, because it forces a retroactive "hard clustering" for the particular feature vector under consideration, while all other feature vectors remain soft-clustered (Fig. 6, column 1). The probabilities for reexamination are pointing to the correct feature vectors though, i.e., 15 and 10.

7. CONCLUSION & OUTLOOK

Given the above considerations and the results, Case IV appears to be the most qualitatively and quantitatively sensitive and responsive prioritization scheme, both for hard (Fig. 5) and soft clustering (Fig. 6). Since non-zero probabilities for all feature vectors are generated starting with sigma=0.0, as a caveat, Case IV may potentially lead to a higher false-positive rate than Case III. However, since the premise for this entire study is that only preliminary data at a coarser, more uncertain or lower resolution are available, it is generally impossible to know what constitutes ground truth in an operational area unless that operational area had already been fully characterized before, in which case (re)examination would be a moot point.

The proposed prioritization framework in general and Case IV in particular offer an opportunity to identify potential candidate objects in an operational area, characterized by feature vectors, for in-situ or close-up/follow-up examination to potentially contribute to an increase of "knowledge", i.e., ground truth, without *any* additional measurements, i.e., via *hypothetical probing* only. This represents a unique opportunity, especially for autonomous robotic planetary exploration scenarios (e.g., Fink et al., 2005; Fink, Tarbell, & Jobling, 2008; Fink et al., 2022), which might otherwise be missed in the absence of such a prioritization framework. As such, the prioritization framework has the potential to help maximize scientific return while optimizing resource utilization. Especially in the context of resource-constrained autonomous science craft for space exploration, it has the potential to contribute to the extension of the *Remaining Useful Lifetime (RUL)* of a space mission or system – a key aspect in Prognostics and Health Management (PHM).

As shown in Fink (2006), the above introduced prioritization Cases I-IV of the overarching prioritization framework can readily be expanded to the prioritization of *n-tuples* of objects, i.e., pairs, triplets, etc., as well as to the prioritization of instrument/sensor usage for in-situ, follow-up measurements of the prioritized objects. The latter is particularly critical for autonomous robotic planetary exploration missions, where often additional constraints have to be taken into account, such as risk vs. benefit considerations, time, power consumption, instrument/sensor resolution limits, spectral ranges in case of a spectrometer, accessibility of objects in case of a rover/lander arm, etc.

It is also conceivable and potentially beneficial to combine multiple cases of the prioritization framework, e.g., to form a committee or ensemble. For example, Cases I-III are more selective and exhibit the "come" and "go" of probabilities for target vectors to be reexamined (Figs. 5 and 6), whereas Case IV is more sensitive and flags target vectors earlier.

In addition to the RUL considerations detailed above, on a more general, PHM-relevant level, the proposed prioritization framework may assist *predictive maintenance* or *condition-based maintenance* by providing an additional mechanism of identifying system components whose actual maintenance status was initially wrongly classified, i.e., through pre-clustering of preliminary, low(er) resolution and/or noisy data.

7.1. *Comparison to other post-clustering quality improvement methodologies*

There are several *extrinsic* and *intrinsic* methods to measure the clustering quality post-clustering, i.e., to assess how good the clustering is (Han, Kamber, & Pei, 2012b). *Extrinsic* methods can be applied when ground truth is available to compare your clustering against. These methods comprise, but are not limited to (see Han, Kamber, & Pei, 2012b for details): *cluster homogeneity, cluster completeness, "rag bag" or "miscellaneous" category, and small cluster preservation*. It is important to note, that these methods are *not* applicable in autonomous robotic planetary exploration as ground truth not only is generally not known but the determination of it is the actual exploration goal. *Intrinsic* methods, in contrast, can be applied in the absence of ground truth. In general, these methods assess the clustering quality by examining cluster separation and compactness (Han, Kamber, & Pei, 2012b). In addition, similarity metrics between feature vectors in the data set, such as *silhouette coefficients*, can be employed (Han, Kamber, & Pei, 2012b). However, while intrinsic methods seem in principle applicable in autonomous robotic planetary exploration, the inherent challenge is that the measured or observed data originate from standoff/preliminary observation, and, as such, mostly eliminate the *similarity* argument, since the preliminary or coarse data of an OA may not be indicative of or similar to the actual ground truth at all.

The probabilities for prioritizing regions or targets of interest for in-situ, close-up investigation, introduced here, can also be understood as *quantitative confidence measures* of the classification, i.e., the pre-clustering of the preliminary data. Post-classification techniques, similar in purpose to the prioritization framework, i.e., to assess and quantify the reliability of classification decisions in the presence of inherent ambiguity of non-discriminative features, inadequate number of training samples, and in particular the degree of noise in measurements or observations have been proposed. For example, Banerjee et al. (2017) have devised a framework that allows for the incorporation of major sources of classification errors into a single quantitative measure, i.e., a confidence metric, to determine the reliability of automated signal classification (ASC) in non-destructive evaluation (NDE) applications, i.e., post-classification. To that end, the authors use *"bootstrapping and weighting Bayes posterior probability with estimated noise distribution"* to embed the effect of noise in NDE measurements into the confidence metric (Banerjee et al., 2017). In stark contrast, the prioritization framework proposed here operates only on a *single* observation or measurement of the object feature vectors in the OA (as is often the case in planetary exploration), i.e., it does not depend on the existence of multiple measurements that would allow bootstrapping or the calculation/estimation of any kind of distribution.

Furthermore, there is a relatively scarce body of work on the specific problem of post-clustering improvement (by not changing the number of clusters, see below), particularly when considering the concept of *hypothetical probing*, i.e., altering the feature vectors to be clustered post-hoc, which is at the core of the prioritization framework introduced here. Most post-clustering improvement approaches only operate on the clustered data at hand. For example, Borlea et al. (2022) introduce *"a way of improving the resulted clusters generated by the K-means algorithm by post processing the resulted clusters with supervised learning algorithm. The proposed approach is focused on improving the quality of the resulting clusters and not on reducing the processing time."* More specifically, Borlea et al. pre-cluster a data set with K-means, and subsequently use a (data point, cluster centroid)-based distance measure as a *split criterion* to divide the pre-clustered data set into a training data set and a misclassified data set. They then train a supervised machine learning algorithm on the training data set in order to re-classify the data points in the misclassified data set, ultimately arriving at the final, improved clustering (Fig. 2 in Borlea et al. (2022)). The fundamental difference to our approach is that Borlea et al. never modify the data points (feature vectors) themselves, which is especially the case in prioritization cases III and IV because of hypothetical probing. Moreover, in none of our cases splitting of the pre-clustered data set and training of a supervised machine learning algorithm on a subset are required. In addition, as mentioned in Section 2.2 above, the computational effort of our prioritization framework is $O(N{\cdot}K)$, i.e., linear in the number of data points N and the number of clusters K, hence computationally inexpensive.

Another example for post-clustering improvement are *partitioning techniques* or *optimization techniques* that allow reassignment of data points (feature vectors) during the clustering process to correct sub-optimal or incorrect initial clustering at a later stage (Marzo et al., 2006). As Marzo et al. (2006) state: *"The majority of these techniques can be formulated as attempts to partition the set of objects so as to optimize some predefined criterion. These techniques employ three distinct procedures: (1) initiating clusters; (2) allocating objects to initial clusters; and (3) relocating the objects in alternative clusters. The differences between the partition techniques lie primarily in the methods for initiating clusters (procedure 1) and in the relocation techniques (procedure 3) [Everitt, 1980]."* K-means (MacQueen, 1967) itself is such a technique. However, as mentioned before, the key difference to our approach is that the data points (feature vectors) to be clustered are not modified in the process.

It is worth pointing out that finding/determining the natural, "right", or "appropriate" number of clusters is a constantly encountered challenge in clustering in general, especially in the absence of ground truth and/or deeper insights into the data-generating process(es) as is the case with both in autonomous robotic planetary surface exploration, as it determines the proper *granularity* of the cluster analysis by finding a good balance between *compressibility* and *accuracy* (Han, Kamber, & Pei, 2012b). The prioritization framework introduced here does not address this challenge as the focus is on *post*-clustering, i.e., the choice of number of clusters has already occurred by other means, e.g., using the elbow method or cross-validation (Han, Kamber, & Pei, 2012a).

Given the above, we believe our approach to post-clustering improvement is novel, conceptually straightforward, and computationally inexpensive, and therefore especially suitable for compute-constrained autonomous robotic planetary surface exploration.

Robotic planetary surface exploration, akin to geological field studies, is largely investigative in nature (Gilbert, 1886; Baker, 2014). It is mainly focused on and driven by the making of discoveries that are not easily accomplished by the mere testing of preconceived hypotheses. Rather, new hypotheses are generated when encountering anomalies, i.e., those features or phenomena that do not fit or agree with preconceived hypotheses because of their identification through the *investigative approach* (Baker, 2017). These new hypotheses are formulated in agreement with (and are dependent on) the investigator's experience with analogues to the anomalies or phenomena encountered/discovered (Gilbert, 1896; Baker, 2014). They turn into *"working hypotheses"* (Chamberlin, 1890) that are subsequently assessed and judged by their resulting productivity or fruitfulness in guiding the investigation/exploration along to further productive inquiry (Baker, 2017). It is in this context that the prioritization framework introduced here, while attempting to improve upon the clustering of the feature vectors obtained from standoff or preliminary observations, effectuates in fact the process of generating working hypotheses during autonomous robotic planetary surface exploration, i.e., in the absence of human explorers on the one hand, and the general lack of high performance computing power on the other hand.

7.2. Potential next steps

Again, it is crucial to reiterate that the operation (not performance, see below) of the prioritization framework introduced here is completely independent of the algorithm used to cluster the preliminary data, because it operates *post-hoc* on the pre-clustered data. As such, we used purposely basic clustering algorithms, i.e., K-means and Fuzzy C-Means, as well as low-dimensional feature vectors in order to visualize the inner workings of the prioritization framework. As next steps, it will be important to investigate how the performance of the prioritization framework is affected by, but not limited to: (1) increasing feature space/vector dimension n; (2) pairwise Euclidean distance in n-dimensional feature space between the initial cluster centroids resulting from the pre-clustering, i.e., dependence on the initial clustering algorithm used; and (3) noise-scaling behavior, such as, $O(n^x)$ with x>0, to model standoff observation/measurement(s) quality or resolution. As such, follow-up numerical simulations with higher-dimensional feature vectors and in the presence of various noise models will certainly have to be conducted to more fully explore and validate the capabilities, performance, and limits of the prioritization framework introduced here, but are beyond the scope of this paper which focuses on its theoretical underpinning.

ACKNOWLEDGEMENT

This research was supported in part by the Edward & Maria Keonjian Endowment in the College of Engineering at the University of Arizona.

DECLARATION OF INTEREST STATEMENT

The authors report no financial interest in the generic prioritization framework presented here.

REFERENCES

Baker, V.R. (2014) *Terrestrial analogs, planetary geology, and the nature of geological reasoning*; Planetary and Space Science, 95, 5-10.

Baker, V.R. (2017) *Debates – Hypothesis testing in hydrology: pursuing certainty versus pursuing uberty*; Water Resources Research, 53, 1770-1778.

Banerjee, P., Udpa, L., Udpa, S., & Benson, J. (2017) *Confidence metric for signal classification in Non Destructive Evaluation*; NDT & E International, 91, 88-96. https://doi.org/10.1016/j.ndteint.2017.06.003.

Bezdek, J.C. (1981) *Pattern Recognition with Fuzzy Objective Function Algorithms*; ISBN 0-306-40671-3.

Bishop, C.M. (1995) *Neural Networks for Pattern Recognition*; Clarendon Press, Oxford, 1995.

Borlea, I.-D., Precup, R.-E., & Borlea, A.-B. (2022) *Improvement of K-means Cluster Quality by Post Processing Improvement of K-means Cluster Quality by Post Processing Resulted Clusters*; Procedia Computer Science, 199, 63–70. https://doi.org/10.1016/j.procs.2022.01.009.

Chamberlin, T.C. (1890) *The method of multiple working hypotheses*; Science 15:92–96. (reprinted in Science 148: 754–759 [1965]).

Duda, R.O., Hart, P.E., & Stork, D.G. (2000) *Pattern Classification and Scene Analysis*; John Wiley & Sons, 2nd edition, 2000.

Dunn, J.C. (1973) *A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters*; Journal of Cybernetics, 3 (3): 32–57; https://doi.org/10.1080/01969727308546046.

Everitt, B.S. (1980) *Cluster Analysis*; John Wiley, Hoboken, N.J.

Fink, W. (2006) *Generic Prioritization Framework for Target Selection and Instrument Usage for Reconnaissance Mission Autonomy*; IEEE WCCI Conf. Proc., 11116-11119; https://doi.org/10.1109/IJCNN.2006.247316.

Fink, W., Brooks, A.J.-W., & Tarbell, M.A. (2018) *Anomaly Detection and Target Prioritization in Planetary Imagery via the Automated Global Feature Analyzer™ (AGFA™): Progress Towards a Driver for Autonomous C⁴ISR Missions*; SPIE DCS Conf. Proc., 106391Z; https://doi.org/10.1117/12.2303795.

Fink, W., Castano, R., Davies, A., & Mjolsness, E. (2001) *Clustering Algorithm for Mutually Constraining Heterogeneous Features*; JPL Technical Report; JPL Open Repository, V1; https://hdl.handle.net/2014/13062 01-1569.pdf [fileName] [last accessed: 07/23/2023].

Fink, W., Datta, A., Dohm, J.M., Tarbell, M.A., Jobling, F.M., Furfaro, R., Kargel, J.S., Schulze-Makuch, D., & Baker, V.R. (2008) *Automated Global Feature Analyzer (AGFA) – A Driver for Tier-Scalable Reconnaissance*; IEEE Aerospace Conf. Proc., paper #1273; https://doi.org/10.1109/AERO.2008.4526422.

Fink, W., Dohm, J.M., Tarbell, M.A., Hare, T.M., & Baker, V.R. (2005) *Next-Generation Robotic Planetary Reconnaissance Missions: A Paradigm Shift*; Planetary and Space Science, 53, 1419-1426; https://doi.org/10.1016/j.pss.2005.07.013.

Fink, W., Dohm, J.M., Tarbell, M.A., Hare, T.M., & Baker, V.R. (2022) *20 Years of Tier-Scalable Reconnaissance: Adoption of a Game-changing Mission Paradigm*; IEEE Aerospace Conf. Proc., paper #2886; https://doi.org/10.1109/AERO53065.2022.9843542.

Fink, W., Tarbell, M.A., & Jobling, F.M. (2008) *Tier-Scalable Reconnaissance – A Paradigm Shift in Autonomous Remote Planetary Exploration of Mars and Beyond*; Chapter 1 in "Planet Mars Research Focus", Ed. L. A. Costas. Nova Science Publishers, Hauppauge, NY. ISBN: 1-60021-826-1.

Gilbert, G.K. (1886) *The inculcation of scientific method by example*; American Journal of Science, 31, 284–299.

Gilbert, G.K. (1896) *The origin of hypotheses*; Science, 3, 1-12.

Han, J., Kamber, M., & Pei, J. (2012a) *8 - Classification: Basic Concepts and Methods*; Chapter 8 in "The Morgan Kaufmann Series in Data Management Systems, Data Mining: Concepts and Techniques" (Third Edition), Editor(s): Jiawei Han, Micheline Kamber, Jian Pei, Morgan Kaufmann, 2012, Pages 327-391, ISBN 9780123814791, https://doi.org/10.1016/B978-0-12-381479-1.00008-3.

Han, J., Kamber, M., & Pei, J. (2012b) *10 - Cluster Analysis: Basic Concepts and Methods*; Chapter 10 in "The Morgan Kaufmann Series in Data Management Systems, Data Mining: Concepts and Techniques" (Third Edition), Editor(s): Jiawei Han, Micheline Kamber, Jian Pei, Morgan Kaufmann, 2012, Pages 443-495, ISBN 9780123814791, https://doi.org/10.1016/B978-0-12-381479-1.00010-1.

Hastie, T., Tibshirani, R., & Friedman, J. (2009) *The elements of statistical learning: data mining, inference, and prediction*; Second Edition. Publisher: Springer New York, NY. eBook ISBN 978-0-387-84858-7; https://doi.org/10.1007/978-0-387-84858-7.

Johnson, W.R., Wilson, D.W., Fink, W., Humayun, M.S., & Bearman, G.H. (2007) *Snapshot hyperspectral imaging in ophthalmology*; Journal for Biomedical Optics, 12(1):014036; https://doi.org/10.1117/1.2434950.

Karim, Md.R., Beyan, O., Zappa, A., Costa, I.G., Rebholz-Schuhmann, D., Cochez, M., & Decker, S. (2021) *Deep learning-based clustering approaches for bioinformatics*; Briefings in Bioinformatics, 22(1), 2021, 393–415; https://doi.org/10.1093/bib/bbz170.

MacQueen, J. (1967) *Some methods for classification and analysis of multivariate observations*; in 5th Berkeley Symposium on Mathematical Statistics and Probability, vol. 1, edited by L. M. Le Cam and J. Neyman, pp. 281–296, Univ. of Calif. Press, Berkeley.

Marzo, G.A., Roush, T.L., Blanco, A., Fonti, S., & Orofino, V. (2006) *Cluster analysis of planetary remote sensing spectral data*; J. Geophys. Res., 111, E03002; https://doi:10.1029/2005JE002532.

Sathishkumar, E. (2015) *Fuzzy C Means Manual Work; Slideshare;* https://www.slideshare.net/ENSathishkumar/fuzzy-c-means-manual-work [last accessed: 07/23/2023].

Williams, C.K.I. (2000) *An MCMC Approach to Hierarchical Mixture Modelling*; Advances in Neural Information Processing Systems 12, Solla, S.A., Leen, T.K., Mueller, K.-R., eds., MIT Press, 2000.

BIOGRAPHIES

Prof. Dr. Wolfgang Fink is the inaugural Edward & Maria Keonjian Endowed Chair with joint appointments in the Departments of Electrical and Computer Engineering, Biomedical Engineering, Systems and Industrial Engineering, Aerospace and Mechanical Engineering, and Ophthalmology & Vision Science at the University of Arizona. He was a Visiting Associate in Physics at the California Institute of Technology (2001–2016), a Visiting Research Associate Professor of both Ophthalmology and Neurological Surgery at the University of Southern California (2005–2014), and a Senior Researcher at NASA's Jet Propulsion Laboratory (2001–2009). Dr. Fink is the founder and director of the Visual and Autonomous Exploration Systems Research Laboratory both at Caltech (http://autonomy.caltech.edu) and at the University of Arizona (https://www.vaesrl.com/). He obtained a B.S. and M.S. degree in Physics and Physical Chemistry from the University of Göttingen, Germany in 1990 and 1993, respectively, and a Ph.D. *summa cum laude* in Theoretical Physics from the University of Tübingen, Germany in 1997. Pursuing a trans-disciplinary systems engineering approach in "smart service systems" in general, he has focused his research and technology development efforts on autonomous systems (robotic space exploration), biomedical engineering for healthcare (especially ophthalmology and vision care), human-computer interfaces (artificial vision implants), smart platforms for mobile- and tele-health, and computeroptimized design. Dr. Fink is a NAI Fellow, ARVO Fellow, SPIE Fellow, PHMS Fellow, AIMBE Fellow, UArizona da Vinci Fellow and ACABI Fellow, Senior Member IEEE, and the current Vice President of the Prognostics and Health Management (PHM) Society. He has over 267 publications (including journal, book, and conference contributions), 6 NASA Patent Awards and 5 Techbrief Awards, as well as 30 U.S. and foreign patents awarded to date in areas, such as: autonomous systems; biomedical devices and tests with emphasis on ophthalmology and ocular healthcare; neural stimulation; signal processing; data fusion, analysis, and anomaly detection; nano-ultra-capacitors for implantable devices; MEMS fabrication; network-on-chip design; and multi-dimensional optimization. In 2021, Dr. Fink co-won the $200,000 DOE/NREL-funded E-ROBOT Prize, and in 2023 he was honored with the SPIE Aden and Marjorie Meinel Technology Achievement Award for his work in optics and imaging.

Karm Al Hajhog graduated with a B.S. degree in Biomedical Engineering from the University of Arizona in 2020. She joined the Visual and Autonomous Exploration Systems Research Laboratory from 2018-2020, working under direction of Prof. Dr. Wolfgang Fink on data analysis for ophthalmic examination devices as well as further developing and simulating the generic prioritization framework introduced in this paper.