A New Approach for 3D Shape Retrieval Based on an Explainable Boosting Classifier

Fatima Rafii Zakani[1], Mohcine Bouksim[2], Khadija Arhid[3], Taoufiq Gadi[4], Mohamed Aboulfatah[5]
University of Abdelmalek Essaâdi, ENSAH, Laboratory of Applied Sciences, SOVIA team Al Hoceïma, Morocco[1]
Mohammed First University, ENSAO, SmartICT Lab, Oujda, Morocco[2]
Cadi Ayyad University, ESTS, Lab. of Processes, Signals, Industrial Systems and Computer Science, Safi, Morocco[3]
University of Hassan I, FSTS, MISI Laboratory, Settat, Morocco[4, 5]

***Abstract*—In the past decade, the number of available 3D models has grown rapidly. This growth is mainly driven by the development of scanning devices and increasing demand for 3D meshes in many application fields. As a consequence, it becomes crucial to have consistent content-based retrieval systems of large 3D mesh repositories. However, many existing methods still struggle to achieve a good compromise between efficiency and accuracy, especially on large and heterogeneous databases. In this work, we propose a supervised framework for building compact but discriminative 3D shape descriptors. For each mesh, we start by computing three features - the dihedral angles between adjacent faces, the Shape Diameter Function (SDF), and the Shape Index - then we convert them into normalized histograms and concatenate them into a single feature vector, which is then used as input to an Explainable Boosting Classifier (EBC). After training, the classifier produces, for each mesh, a short probability vector that we use as its numerical descriptor. Experiments on the standard Princeton Benchmark database validate our approach, achieving a mean Average Precision (mAP) of 97.23%, outperforming the selected baseline methods under the adopted experimental protocol.**

***Keywords*—3D shape retrieval; content-based indexing; shape descriptor; explainable boosting classifier**

I. Introduction

As 3D model databases continue to grow rapidly in medicine, engineering, and multimedia, there is an urgent need for efficient retrieval systems that can support powerful discrimination capabilities. This necessity has motivated research in content-based indexing to become a significant area of study [1],[2],[3], Early attempts at addressing this challenge focused on extracting shape descriptors based on geometric measures [4], [5].

The expressive limitations of these handcrafted descriptors led to a paradigm shift toward deep learning. These new techniques achieved state-of-the-art performance by learning features directly from various data representations, such as multi-view projections and voxel-based grids [6], [7], [8]. This success, however, is counterbalanced by significant drawbacks: deep learning models typically require vast labeled datasets, substantial computational resources, and their 'black box' nature limits interpretability.

These challenges have led to a renewed interest in alternative approaches that blend compact, handcrafted geometric descriptors with interpretable learning algorithms. Among these, the Explainable Boosting Classifier (EBC) [9] offers a compelling compromise between predictive performance and full interpretability. By constructing an additive model from shallow decision trees, the EBC achieves high performance while generating a descriptor ideally suited for efficient 3D mesh indexing.

In this work, we introduce a new retrieval method that directly addresses the aforementioned challenges. We combine the Explainable Boosting Classifier (EBC) with a robust set of geometric features. Specifically, we employ three such features: dihedral angles, which reveal local angular relationships, and the Shape Diameter Function (SDF)[10], which provides information on object thickness, and the Shape Index [11], which describes local surface topology.

This study is structured as follows. Section II reviews related work. Section III details our proposed method. Section IV validates our approach through experimental analysis, and Section V concludes with a summary of our contribution and discusses future work.

II. Related Works

The field of Content-Based 3D Shape Retrieval has historically evolved along two axes: traditional methods based on geometric descriptors, and more recent approaches leveraging deep learning.

A. Traditional Methods

Initial research focused on extracting shape descriptors from geometric properties, leading to two main families.

1) Global descriptors: Represent an entire model in a single vector. Early works used statistical methods like shape distributions [4] or spectral approaches like spherical harmonics [12]. More recently, methods like DEA-based optimization [13] integrated multiple geometric properties to improve performance. However, the global nature of these descriptors limits their robustness to partial matching and articulation, and they often struggle to capture high-level semantics.

2) Local descriptors: Local descriptors were developed to address the limitations of global methods against partial matching and articulation. For instance, Van Blokland et al. [14] introduced binary descriptors from local ray-intersections,

achieving robust partial matching on the SHREC'16 benchmark using a dissimilarity tree. Similarly, Arhid et al. [15] improved performance on articulated objects by segmenting the mesh and matching per-part DEA descriptors. However, local methods are computationally expensive due to the part-matching problem and sensitive to initial segmentation quality.

B. Machine Learning Approaches

The limitations of handcrafted descriptors on complex shapes led to the rise of learning-based techniques, which automatically learn discriminative features and consistently outperform traditional methods.

1) View-based methods: These methods were among the first to apply deep learning by converting 3D objects into sets of 2D projections for use with CNNs. MVCNN [7] pioneered this approach by aggregating features from multiple views. Later works like RotationNet [16] and GVCNN [17] enhanced this concept by improving the aggregation method and reducing sensitivity to object orientation. These techniques are very powerful; however, their success is quite dependent on the selection and number of representative views.

2) Voxel-based deep representations: Volumetric methods convert meshes into voxel grids to enable 3D convolutions. 3D ShapeNets [6] introduced a probabilistic volumetric representation, while VoxNet [8] proposed a more compact and efficient architecture. However, these methods suffer from high memory requirements and resolution limitations, restricting their ability to capture fine geometric details.

3) Point cloud-based deep architectures: Point-based neural networks were developed to operate directly on 3D point clouds, thus avoiding the problems of volumetric discretization. The foundational PointNet model [18] achieved permutation invariance by using symmetric functions to aggregate point features. This concept was extended by PointNet++ [19], which applied the architecture hierarchically to learn local features, and later by DGCNN [20], which used a dynamic graph to better encode neighborhood relations. While these approaches retain fine geometric detail, they are highly dependent on the initial point sampling and sensitive to density variations.

C. Mesh-Based Deep Learning Methods

To avoid the information loss and processing overhead associated with data conversion, mesh-based deep neural networks operate directly on the native polygonal structure. Pioneering works like MeshCNN [21] introduced convolutions on mesh edges, using edge-collapse pooling to preserve topology. MeshNet [22] took a different approach by learning from face attributes, while methods like SpiralNet++ [23] view the mesh as a graph and apply convolutions along spiral vertex sequences to capture local geometry. Although this family of models yields detailed descriptors, they generally incur a high computational cost, especially for meshes with complex geometry.

D. Hybrid Learning Methods

Between purely hand-crafted shape descriptors and deep neural networks, hybrid approaches have emerged that leverage explicit geometric features within supervised models. Bouksim et al. [24] proposed a neural network-based method before transitioning to a CatBoost-based method [25], [26]. In their approach, geometric features are extracted from 3D meshes and used as input to the classifier. Their findings demonstrate that this strategy significantly improves retrieval accuracy compared to using raw geometric descriptors, while maintaining a simple and efficient pipeline.

Our work addresses the gap between traditional efficiency and deep learning accuracy. We use a boosting classifier in a novel capacity: to transform simple geometric features into a compact descriptor. By leveraging the classifier's output probability vector as the final shape descriptor, our method achieves high retrieval accuracy while remaining computationally efficient, providing a practical alternative to resource-intensive deep learning architectures.

III. Proposed Method

In this study, we use a hybrid 3D shape retrieval method based on an Explainable Boosting Classifier (EBC) trained with geometric mesh properties. The main idea of the proposed approach is to use the class-probability vector obtained by our EBC model as a compact descriptor for each 3D object, which can be efficiently used for similarity-based retrieval. The following sections will elaborate on the theoretical foundation and implementation details.

A. From Classification to Similarity

Our main hypothesis is that the class-probability vector from a well-trained classifier can serve as a compact descriptor. Indeed, by learning to separate classes, the classifier builds an internal model of the key geometric and semantic features of the data. We argue that two semantically similar shapes, even from different classes (e.g., a bird and an airplane), will be processed similarly by this model, leading to comparable probability distributions. Consequently, the distance between these probability vectors becomes a measure of high-level semantic similarity, capturing a likeness that goes beyond simple class categories. Measuring the similarity between these probability vectors therefore requires a metric that captures their directional alignment in the probability space. As demonstrated empirically in Section IV-B(5), the Cosine distance is the most effective choice for this purpose, since the discriminative information in EBC signatures is primarily encoded in the orientation of the probability vector rather than its magnitude.

B. Boosting and Explainable Boosting Classifier (EBC)

The theory behind our approach is based on ensemble learning, and more specifically, boosting. Boosting is an iterative technique that combines several weak classifiers (whose performance is slightly better than random choice) to form a single strong and highly accurate classifier. The principle is that each new weak model is trained to correct the prediction errors made by previous models, gradually improving the overall performance of the system.

The Explainable Boosting Classifier (EBC), also known as the Explainable Boosting Machine (EBM) [9], belongs to the Generalized Additive Models (GAM) family. These models represent predictions as a sum of functions, where each function

captures the contribution of an individual feature or a pair of features. The structure of an EBM is defined as:

$$g(E[y]) = \beta_0 + f_1(x_1) + f_2(x_2) + .. + f_k(x_k) \quad (1)$$

where, $\boldsymbol{g}$ is a link function, $\boldsymbol{E[y]}$ is the expected prediction, and each $\boldsymbol{f_j}$ is a shape function that models the impact of the feature $\boldsymbol{x_j}$

The Explainable Boosting Classifier (EBC) employs gradient boosting algorithms to learn feature functions $\boldsymbol{f_j}$, which captures the value of each feature to the model, with shallow decision trees. Its primary strength is that it is interpretable: EBC, in contrast to neural networks, being a black box model, is a glass box model, i.e., the contribution of each feature could be isolated and visualized. Through this transparency, we can know not only the general behavior of the model, but also can explain the individual predictions. With the combination of good precision and easy interpretability, EBC was chosen for this work to enable content-based indexing and retrieval by providing a competitive performance.

C. Geometric Features Used

To train our classifier, we combined three histograms of features extracted from the 3D object. These features were chosen for their simplicity of calculation, their invariance to rigid transformations, and their discriminating power proven in the literature.

Dihedral angle [27]: A widely used feature that measures the angle between two adjacent faces. It effectively captures local geometry and the sharpness of creases and is calculated from the faces' normal vectors. The following equation can express the dihedral angle:

$$\theta = arcos\left(\frac{\overrightarrow{n_1}.\overrightarrow{n_2}}{||\overrightarrow{n_1}||\,||\overrightarrow{n_2}||}\right) \quad (2)$$

where:

- $\boldsymbol{\theta}$ is the dihedral angle,
- $\overrightarrow{\boldsymbol{n_1}}$ and $\overrightarrow{\boldsymbol{n_2}}$ are the normal vectors of the two adjacent faces,
- $\cdot$ represents the dot product,
- $||\,\overrightarrow{n_1}\,||$ and $||\,\overrightarrow{n_2}\,||$ are the norms of the normal vectors.

Shape Index [11] proposed by Koenderink, the Shape Index describes the local surface topology based on its principal curvatures. It provides a continuous scale for categorizing local shape (e.g., convex, concave), making it an effective feature for segmentation and indexing. The Shape Index $\boldsymbol{S}$ is calculated using the principal curvatures $\boldsymbol{k_1}$ and $\boldsymbol{k_2}$ at a point on the surface, as expressed in the following equation:

$$S = \frac{2}{\pi}\arctan\left(\frac{k_2+k_1}{k_2-k_1}\right) \quad (3)$$

where,

- $\boldsymbol{k_1}$ and $\boldsymbol{k_2}$ are the main curvatures with $\boldsymbol{k_2 \geq k_1}$, note that this index is not defined for flat areas ($\boldsymbol{k_1 = k_2}$).

Shape Diameter Function (SDF) [10]: A scalar function that captures the local thickness of the mesh. For each surface point, SDF measures the distance to the opposite side of the mesh along the inverted normal vector.

For each of the three geometric properties (SDF, Dihedral Angle, and Shape Index), a histogram is computed using 128 bins. To ensure a consistent representation, the value range for each histogram is determined globally from the statistics of the entire training dataset.

Finally, these three histograms are concatenated to form a single feature vector of dimension 384 (3 * 128), which serves as the input for our EBC classifier.

D. Descriptor Generation and Comparison

The process of creating and comparing descriptors involves three main steps:

1) Training and optimization: Before feature extraction, all 3D meshes underwent a standardized preprocessing pipeline consisting of size normalization and centering, whereby each model's center of mass was translated to the origin. An Explainable Boosting Classifier (EBC) from the InterpretML library was then trained on the concatenated histogram of geometric features. To identify the optimal hyperparameter configuration, a systematic grid search was conducted over the following ranges: n_estimators $\in \{10, 50, 100\}$, learning_rate $\in \{0.01, 0.05, 0.1\}$, and max_bins $\in \{128, 256, 512\}$. Model selection was performed using stratified 10-fold cross-validation, ensuring that each fold maintains a proportional representation of all 19 semantic classes. This stratification prevents class imbalance across folds and guarantees that no class is absent from either the training or validation partition during any fold, thus ensuring reproducible and unbiased hyperparameter selection. The optimal configuration identified by this process is reported in Table I.

TABLE I. OPTIMAL EBC HYPERPARAMETERS SELECTED VIA GRID SEARCH WITH STRATIFIED 10-FOLD CROSS-VALIDATION. THESE SETTINGS WERE RETAINED FOR ALL EXPERIMENTS REPORTED IN THIS STUDY

Hyperparameter	Value	Description
n_estimators	50	Number of boosting rounds (trees).
learning_rate	0.1	Step size shrinkage to prevent overfitting.
max_bins	256	Maximum number of bins for feature discretization.
random_state	42	Ensures reproducibility of the training process

2) Descriptor generation and refinement: The trained classifier is applied to generate a descriptor for each 3D model. For a given object, the feature vector is passed through the classifier to produce its probability vector. To ensure consistency across all models in the dataset, we enforce a fixed class ordering using the classifier's classes_ attribute, ensuring that all descriptors share identical dimensional structure and class correspondence. A key advantage of the Explainable Boosting Classifier is that it inherently produces well-calibrated probability estimates due to its Generalized Additive Model (GAM) structure and gradient boosting mechanism. Unlike

many ensemble methods that require post-hoc calibration (e.g., Platt scaling or Isotonic Regression), the EBC's additive formulation given in Eq. (1) and round-wise optimization naturally yield reliable probability outputs. The resulting probability vector serves as a compact and discriminative descriptor, encoding the shape's similarity distribution across all learned classes.

3) Similarity measurement: Finally, after generating the refined descriptors, we compute the similarity between two 3D models. While descriptors represent probability distributions, our analysis (Table III) reveals that the semantic classes are highly separated in the probability space. Among the distance metrics evaluated (Jensen-Shannon divergence [28], a symmetrized variant of the Kullback-Leibler divergence [29], Euclidean, Cosine Similarity), Cosine Similarity achieves the best performance (97.23% mAP), demonstrating that the angular direction of the probability vector is highly discriminative. This metric is particularly effective for our high-dimensional probability descriptors, as it focuses on the direction rather than the magnitude of the vectors, making it robust to scale variations and computationally efficient.

IV. Experimental Results

We conduct experiments on the Princeton Shape Benchmark (PSB), which contains 390 3D mesh models distributed across 19 semantic classes. This database serves as a standard evaluation testbed for 3D shape retrieval research.

A. Experimental Protocol

To ensure a rigorous evaluation of our method, we adopted an experimental protocol designed to prevent information leakage and ensure comparability with state-of-the-art methods.

1) Method overview: The retrieval system operates in two stages. First, an Explainable Boosting Classifier is trained on geometric feature histograms. The classifier learns to map shapes to a 19-dimensional probability space that captures semantic similarity. These vectors are then used as descriptors, and the similarity between models is computed using Cosine distance to rank the results.

2) Evaluation setup: We follow the standard PSB query-by-example protocol, where each of the 390 models is used as a query against the full dataset (excluding itself), and results are ranked using Cosine similarity. Performance is evaluated with standard metrics (NN, FT, ST, DCG, mAP). Although PSB is a widely used benchmark, it remains a moderate-scale dataset; it was selected to ensure a fair comparison with baseline methods evaluated under the same conditions. Evaluating the approach on larger datasets such as SHREC or ModelNet40 is left for future work. This protocol ensures a rigorous and consistent evaluation framework widely adopted in 3D shape retrieval research.

3) Comparison methods: We evaluate our approach against four well-known methods: two learning-based methods (an ANN-based descriptor and a CatBoost-based method) and two classical methods (Panorama and LightField). All comparisons use the same query set and evaluation metrics to ensure fairness.

4) Evaluation metrics: The performance of our method is evaluated globally across all queries using a comprehensive set of standard retrieval metrics, including the Precision-Recall Curve, Nearest Neighbor (NN), First Tier (FT), Second Tier (ST), and Discounted Cumulative Gain (DCG). These metrics are calculated as defined by the standard PSB evaluation scripts, ensuring our results are directly comparable to the literature.

B. Performance Evaluation

1) Performance analysis via precision-recall curve: We initially assess the general quality of our method through the Precision- Recall (P-R) curve analysis, as shown in Fig. 1. This curve measures a descriptor's ability to retrieve all relevant models (Recall) without returning incorrect results (Precision).

As can be seen by inspection of the graph, our EBC method (in red) is clearly superior. Its curve remains at a precision of 1.0 across almost the entire recall scale, an exceptional result that demonstrates a near-perfect discriminative power. In practice, this means that our descriptor can retrieve the vast majority of similar shapes before making a single classification error.

Our approach outperforms CatBoost, the closest competitor, and significantly surpasses view-based methods (LightField, Panorama) as well as ANN, whose performance typically degrades at higher recall levels. Overall, the P-R analysis confirms the strong effectiveness of our descriptor for high-precision 3D shape retrieval.

2) Per-class performance analysis: We analyse here the robustness of our descriptor across the different classes of the database. For this test, we measure the precision within the top K=10 and K=20 results to evaluate the performance midway and across the entire class. Fig. 2 (for K=10) and Fig. 3 (for K=20) present the results.

At K=10 Fig. 2, our EBC method (in red) exhibits a remarkably stable and high performance, close to 1.0 across almost all classes. CatBoost (in green) follows a similar trend, confirming the robustness of boosting approaches. In contrast, the other methods already show signs of weakness: ANN and Panorama suffer from notable performance drops on certain categories, and Lightfield exhibits by far the greatest variability.

The more demanding analysis at K=20 Fig. 3 is even more telling. Our EBC method maintains exceptional consistency and precision, proving it does not just succeed on the easy cases. This level of stability is not matched by the competing methods, whose inter-class variability becomes more pronounced and whose weaknesses observed at K=10 are here clearly exacerbated.

In conclusion, this per-class evaluation confirms the superior robustness of our descriptor. It is not only more performant on average but also more reliable and consistent across a wide spectrum of 3D shapes.

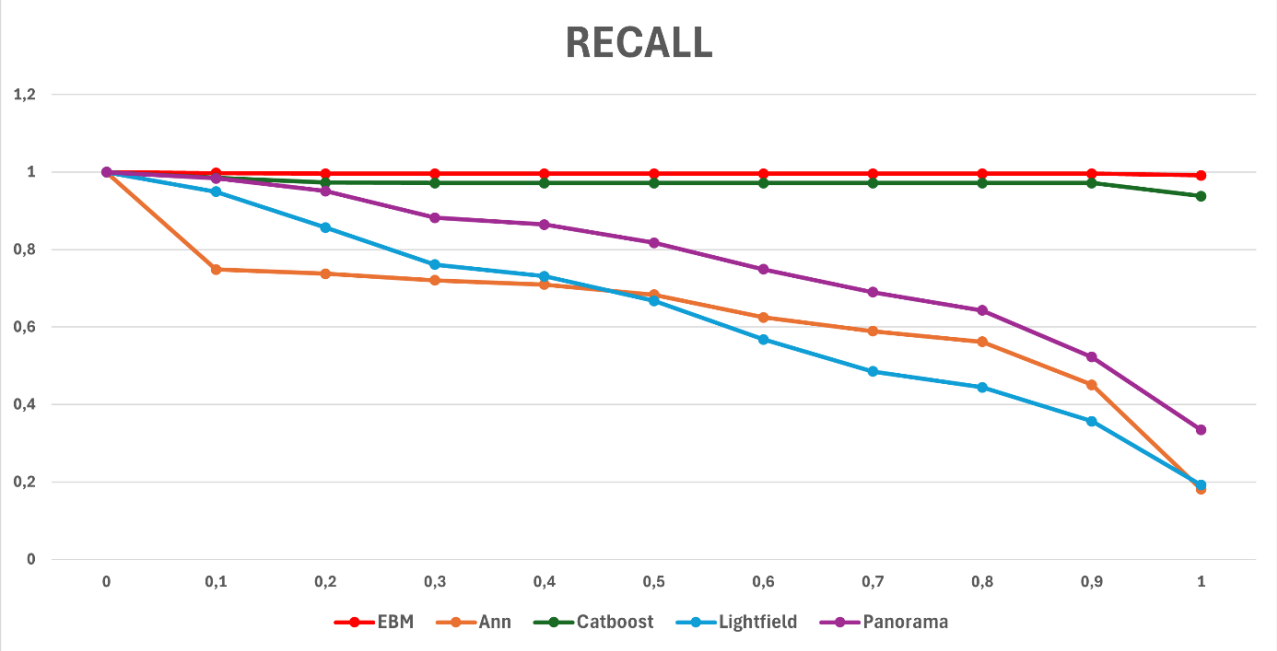

Fig. 1. Precision-Recall curves for the proposed EBC method and four baseline descriptors on the PSB. The EBC (red) maintains near-perfect precision across the full recall range, demonstrating superior discriminative power over all competing methods.

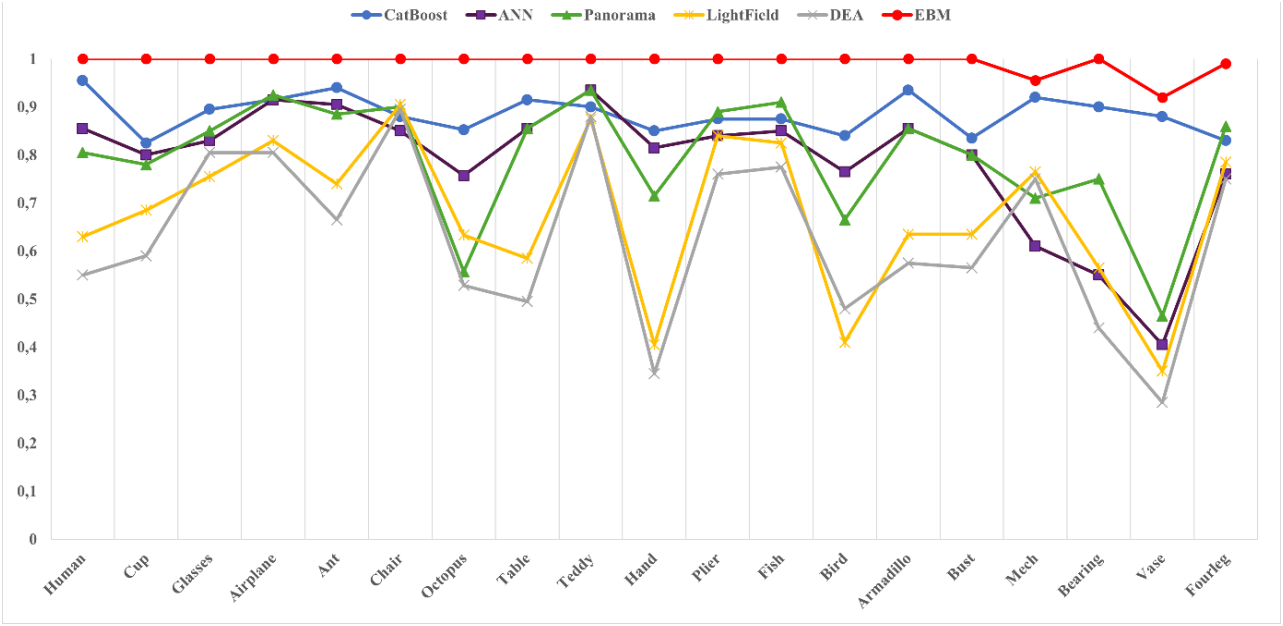

Fig. 2. Per-class retrieval precision at K=10 across the 19 PSB semantic classes. The EBC (red) achieves consistently high and stable scores, while competing methods exhibit notable variability across categories.

TABLE II. QUANTITATIVE COMPARISON OF THE PROPOSED EBC DESCRIPTOR AGAINST FOUR BASELINE METHODS ACROSS STANDARD PSB RETRIEVAL METRICS. BEST SCORES ARE HIGHLIGHTED IN BOLD. THE EBC ACHIEVES THE HIGHEST PERFORMANCE ON ALL METRICS

Descriptors/Metrics	NN	NN+1	1st Tier	2nd Tier	DCG	F-Measure
EBC	**0.99**	**0.99**	**0.99**	**0.53**	**0.99**	**0.46**
CatBoost	0.98	0.97	0.96	0.52	0.98	0.45
Panorama	0.97	0.92	0.73	0.43	0.92	0.42
LightField	0.91	0.84	0.57	0.36	0.86	0.38
ANN	0.95	0.93	0.80	0.45	0.86	0.38

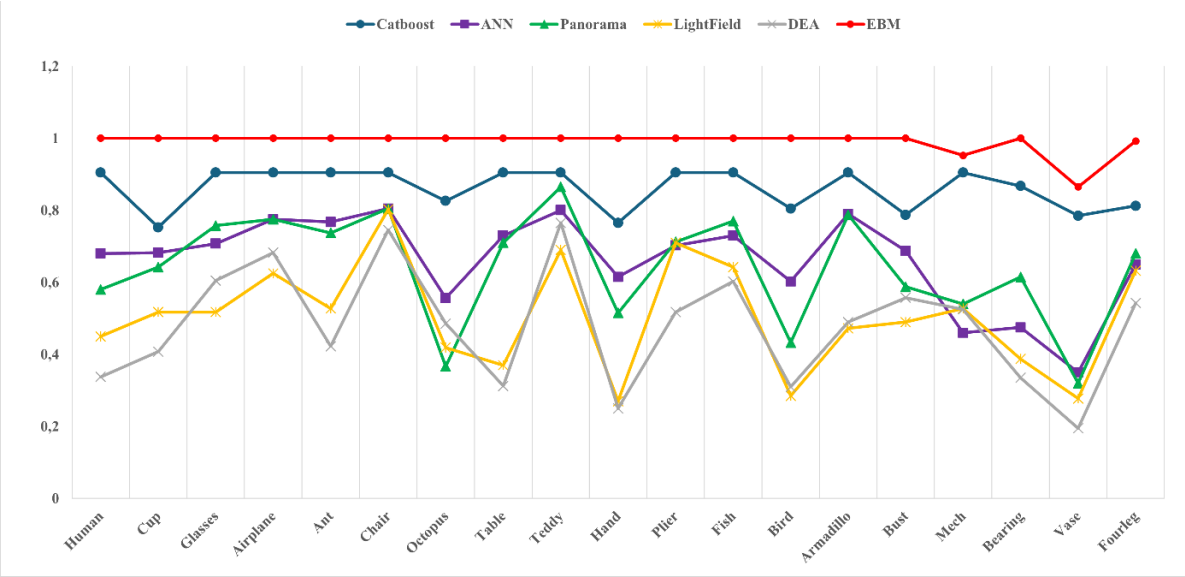

Fig. 3. Per-class retrieval precision at K=20 across the 19 PSB semantic classes. The performance gap between the EBC (red) and competing methods widens compared to K=10, further confirming the robustness and consistency of the proposed descriptor.

3) Quantitative evaluation with standard metrics: To supplement the graphical analyses, we now perform a quantitative comparison of our EBC descriptor against competing methods using a set of standard metrics. Table II summarizes the results for Nearest Neighbor (NN), First Tier (FT), Second Tier (ST), F-Measure, and Discounted Cumulative Gain (DCG).

The data in the table unequivocally confirms the superiority of our EBC method. It achieves the highest score across all metrics, which quantifies the dominance observed in the previous visual analyses (P-R curve and per-class stability). The nearly perfect Nearest Neighbor (NN) score demonstrates the exceptional accuracy of our descriptor in identifying the single best match for a query. Furthermore, the excellent DCG score proves that the overall quality of the results ranking is also superior, with relevant models being consistently placed at the top of the list.

The duel with CatBoost is once again won by our approach on every measure, although the margin is tighter, confirming the competitiveness of boosting algorithms. The other methods (ANN, Panorama, Lightfield) systematically rank lower, validating the performance trends already observed. In conclusion, this set of quantitative tests provides definitive, numerical proof of our descriptor's effectiveness, establishing it as a new state-of-the-art for 3D shape retrieval.

4) Qualitative evaluation via dissimilarity matrix: To assess the quality of our descriptor, we have computed a dissimilarity matrix for each pair of objects in the database based on their distances. A method that performs well should produce, through visualization, "diagonal blocks" of low dissimilarity (high intra-class similarity) and "off-diagonal" areas of high dissimilarity (low inter-class similarity).

Our matrix is shown in Fig. 4 displays a deep red color along the diagonal (indicating low dissimilarity), corresponding to a quantitatively very low dissimilarity score in the [0, 0.2] range, confirming that the methods groups the objects from the same category with a high degree of cohesion; whereas the off-diagonal region shows a uniformly red color (indicating a systematic high dissimilarity score, in the [0.8, 1.0] range) demonstrating the good discrimination capacity of our method in distinguishing the various categories.

5) Impact of distance metric choice: To select the most appropriate distance metric for comparing EBC probability descriptors, we conducted a systematic evaluation of three candidates: Cosine distance, Jensen-Shannon (JS) divergence, and Euclidean (L2) distance. As shown in Table III, Cosine distance achieves the highest retrieval mAP (97.23%), outperforming Jensen-Shannon (96.94%) and Euclidean (96.61%). This result is theoretically consistent with the nature of our descriptors: since the EBC produces normalized probability vectors, the discriminative information is primarily encoded in their directional orientation rather than their magnitude — a property that the Cosine distance is specifically designed to capture. Accordingly, Cosine distance is adopted as the reference metric for all results reported in this study, including the Abstract, Table II, and Table IV.

TABLE III. RETRIEVAL PERFORMANCE (MAP) OF THE PROPOSED METHOD UNDER THREE DISTANCE METRICS

Distance Metric	mAP (%)	Interpretation
Cosine Distance	97.23%	Selected — best directional alignment
Jensen-Shannon	96.94%	Theoretically optimal for distributions
Euclidean	96.61%	Baseline geometric distance

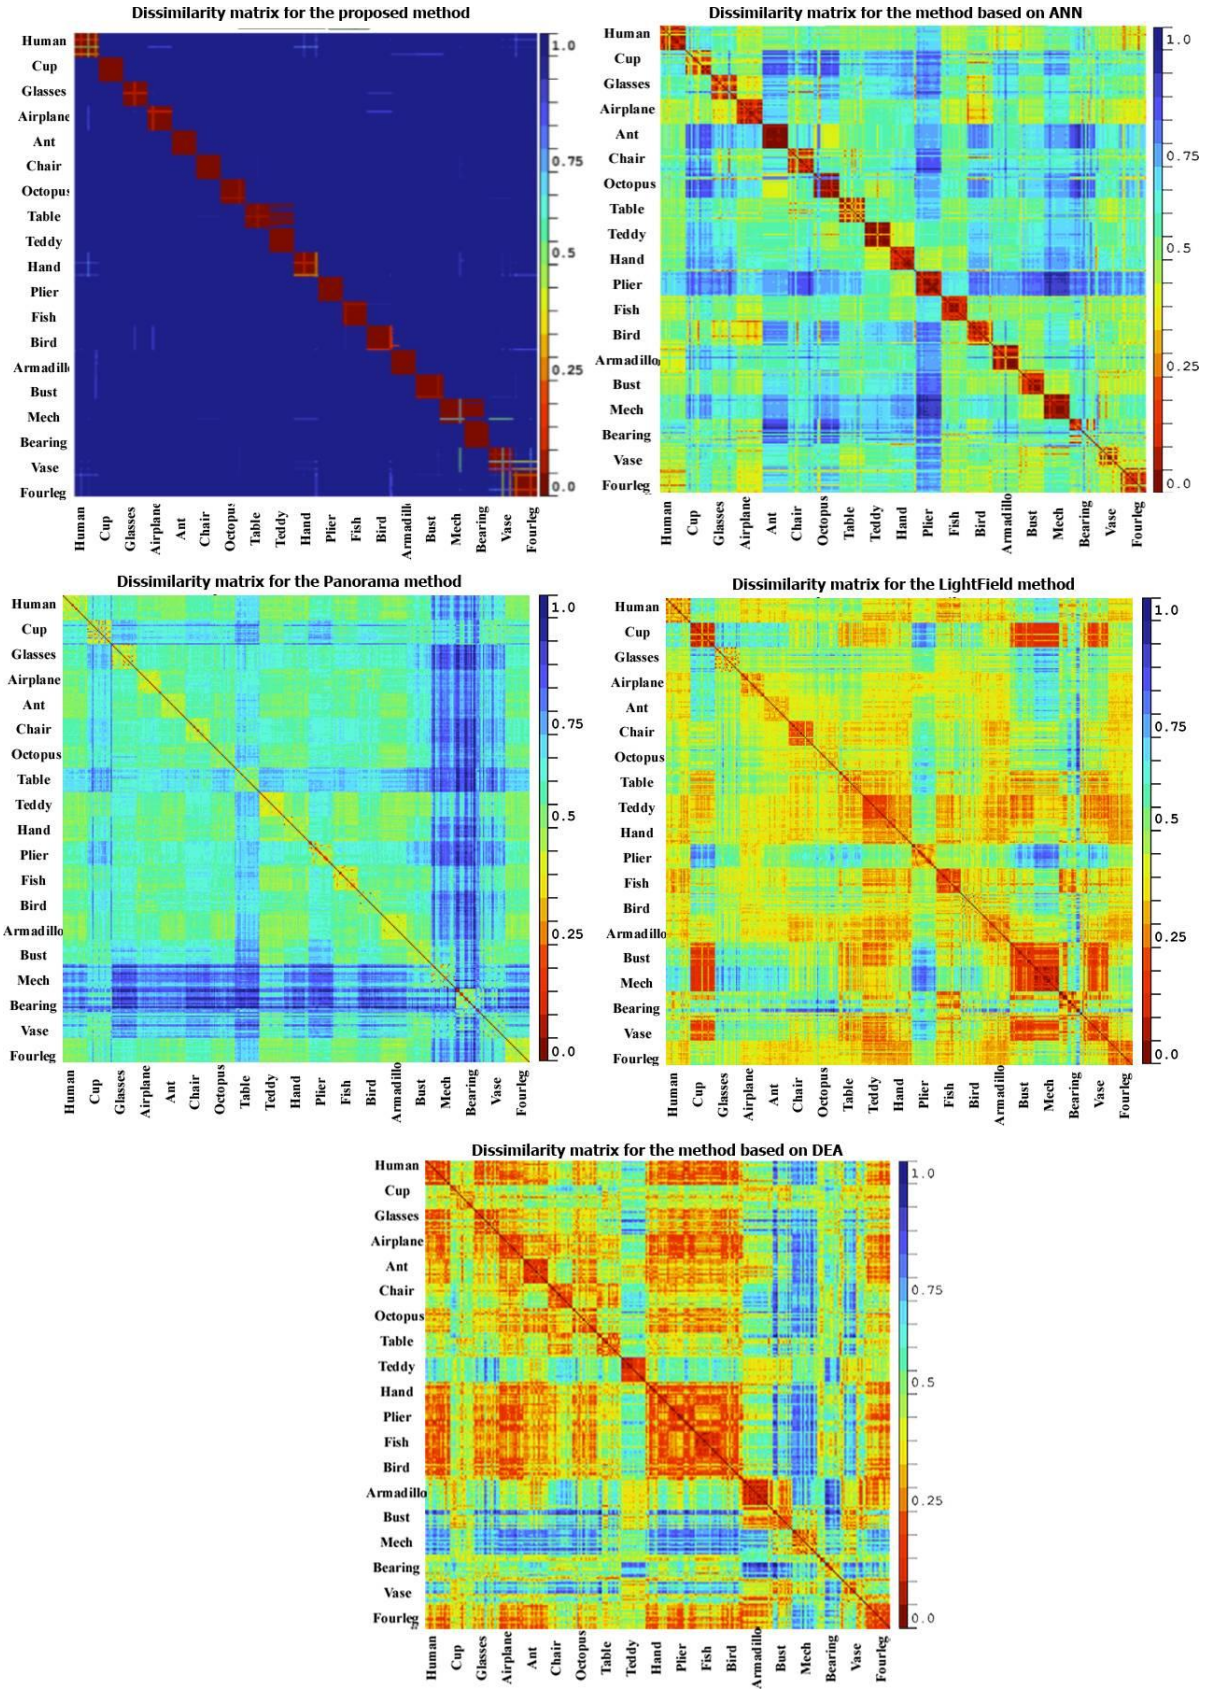

Fig. 4. Dissimilarity matrices computed for the proposed method and four baseline approaches on the PSB dataset. Well-defined dark diagonal blocks reflect strong intra-class cohesion, while uniformly high off-diagonal values confirm effective inter-class separation.

6) Discussion with deep learning approaches: Recent advances in 3D shape retrieval have been largely driven by deep learning approaches operating on meshes, point clouds, or multi-view representations, such as MeshCNN, SpiralNet++, and PointNet-based models. These methods have demonstrated strong performance by learning hierarchical and highly discriminative features directly from raw geometric data.

However, such approaches typically require large-scale labeled datasets, significant computational resources, and often lack interpretability due to their black-box nature. In contrast, the proposed method relies on handcrafted geometric features combined with an Explainable Boosting Classifier, offering a lightweight and interpretable alternative.

Although direct experimental comparison with these models is beyond the scope of this work, the proposed approach achieves competitive retrieval performance (97.23% mAP) while maintaining low computational complexity and full transparency, making it particularly suitable for applications where interpretability and efficiency are critical.

C. Sensitivity Analysis and Justification of the Probability Space

1) Impact of probability calibration: A core assumption of our approach is that the proximity of output probability vectors correctly maps to the geometric and semantic similarity of 3D shapes. To empirically validate this assumption and assess the robustness of the EBC descriptor, we conducted a comprehensive sensitivity analysis evaluating divergence choices, class count, data imbalance, and class structure.

Impact of the Distance Metric: We evaluated the retrieval performance using Cosine Similarity (97.23% mAP), Jensen-Shannon (96.94% mAP), and Euclidean (96.61% mAP). While Jensen-Shannon is theoretically optimal for comparing probability distributions, Cosine Similarity achieves the highest performance, demonstrating that the angular direction of the probability vector is more discriminative than probability-specific metrics. This result validates that the EBC generates highly separable semantic representations where directional similarity outperforms distributional divergence measures.

Sensitivity to Class Count (Semantic Spanning): We evaluated the retrieval task using subsets of 5, 10, and 19 classes. The performance remains consistently high across all configurations (96.13%, 100.0%, and 97.23%, respectively, all measured with Cosine distance). This demonstrates that the probability space dynamically adjusts its capacity; even with a small number of base classes, the descriptor remains rich enough to separate shapes effectively.

Robustness to Class Imbalance: To test stability, we intentionally degraded the training set by keeping only 15% of the samples for half of the classes (severe imbalance). Despite this extreme condition, the overall retrieval mAP only decreased to 79.79%. This minimal degradation highlights the strong resilience of the EBC model to imperfect or skewed datasets.

D. Robustness and Error Characterization

Statistical Stability across Multiple Splits: To ensure the generality of our method and verify that our near-perfect results are not sensitive to a specific data split, we conducted 5 independent runs using different random seeds for the stratified train/test splits. Table IV reports the mean and standard deviation (dispersion) for all key metrics. The extremely low standard deviations (e.g., 0.38% for mAP) demonstrate that the EBC descriptor is highly robust and consistently extracts discriminative semantic information regardless of the specific training subset. Note: The difference between the reference-split mAP (97.23%) and the mean mAP over 5 random splits (98.74% ± 0.38%) reflects normal variability across train/test partitions and confirms the stability of the proposed method.

TABLE IV. MEAN AND STANDARD DEVIATION OF RETRIEVAL METRICS COMPUTED OVER 5 INDEPENDENT RUNS WITH DIFFERENT RANDOM STRATIFIED SPLITS. THE LOW STANDARD DEVIATIONS CONFIRM THE STABILITY AND GENERALIZABILITY OF THE PROPOSED APPROACH

Metric	Mean ± Standard Deviation
mAP(%)	98.74±0.0.38
NN(%)	97.90±0.0.52
FT(%)	97.90±0.0.52
ST(%)	49.50±0.0.34
DCG(%)	99.05±0.0.28

V. CONCLUSION

In this study, we presented a hybrid 3D shape retrieval method combining handcrafted geometric descriptors with the Explainable Boosting Classifier (EBC). Experiments on the Princeton Shape Benchmark demonstrate a mean Average Precision of 97.23%, surpassing state-of-the-art methods including CatBoost, Panorama, and LightField. Despite these results, the evaluation remains limited to a moderate-scale dataset, and the method relies on supervised training, which requires labeled data. Future work will focus on validating the approach on larger benchmarks such as SHREC and ModelNet40 and on exploring hybrid combinations with deep learning features to further improve retrieval performance and scalability.

REFERENCES

[1] J. W. H. Tangelder and R. C. Veltkamp, "A survey of content based 3D shape retrieval methods," Proc. - Shape Model. Int. SMI 2004, 2004, doi: 10.1109/SMI.2004.1314502.

[2] B. Bustos, D. Keim, D. Saupe, and T. Schreck, "Content-based 3D object retrieval," IEEE Comput. Graph. Appl., vol. 27, no. 4, pp. 22–27, Jul. 2007, doi: 10.1109/MCG.2007.80.

[3] Y. Yang, H. Lin, and Y. Zhang, "Content-based 3-D model retrieval: A survey," IEEE Trans. Syst. Man Cybern. Part C Appl. Rev., vol. 37, no. 6, pp. 1081–1098, 2007, doi: 10.1109/TSMCC.2007.905756.

[4] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Matching 3D models with shape distributions," 2001, doi: 10.1109/SMA.2001.923386.

[5] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz, "Rotation invariant spherical harmonic representation of 3D shape descriptors," 2003.

[6] Z. Wu et al., "3D ShapeNets: A deep representation for volumetric shapes," Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit., vol. 07-12-June-2015, pp. 1912–1920, Oct. 2015, doi: 10.1109/CVPR.2015.7298801.

[7] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view Convolutional Neural Networks for 3D Shape Recognition," in 2015 IEEE International Conference on Computer Vision (ICCV), 2015, pp. 945–953, doi: 10.1109/ICCV.2015.114.

[8] D. Maturana and S. Scherer, "VoxNet: A 3D Convolutional Neural Network for real-time object recognition," IEEE Int. Conf. Intell. Robot. Syst., vol. 2015-December, pp. 922–928, Dec. 2015, doi: 10.1109/IROS.2015.7353481.

[9] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad, "Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission," Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., vol. 2015-August, pp. 1721–1730, Aug. 2015, doi: 10.1145/2783258.2788613/SUPPL_FILE/P1721.MP4.

[10] L. Shapira, A. Shamir, and D. Cohen-Or, "Consistent mesh partitioning and skeletonisation using the shape diameter function," Vis. Comput., vol. 24, no. 4, pp. 249–259, 2008, doi: 10.1007/s00371-007-0197-5.

[11] J. J. Koenderink and A. J. van Doorn, "Surface shape and curvature scales," Image Vis. Comput., 1992, doi: 10.1016/0262-8856(92)90076-F.

[12] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung, "On Visual Similarity Based 3D Model Retrieval," Comput. Graph. Forum, 2003.

[13] M. Bouksim, F. R. Zakani, K. Arhid, M. Aboulfatah, and T. Gadi, "New Approach for 3D Mesh Retrieval Using Data Envelopment Analysis," Int. J. Intell. Eng. Syst., vol. 11, no. 1, pp. 1–10, 2018, doi: 10.22266/ijies2018.0131.01.

[14] B. I. van Blokland and T. Theoharis, "Partial 3D object retrieval using local binary QUICCI descriptors and dissimilarity tree indexing," Comput. Graph., vol. 100, pp. 32–42, Nov. 2021, doi: 10.1016/J.CAG.2021.07.018.

[15] K. Arhid, F. R. Zakani, B. Sirbal, M. Bouksim, M. Aboulfatah, and T. Gadi, "A novel approach for partial shape matching and similarity based on data envelopment analysis," Comput. Opt., vol. 43, no. 2, pp. 316–323, Mar. 2019, doi: 10.18287/2412-6179-2019-43-2-316-323.

[16] A. Kanezaki, Y. Matsushita, and Y. Nishida, "RotationNet: Joint Object Categorization and Pose Estimation Using Multiviews From Unsupervised Viewpoints." pp. 5010–5019, 2018.

[17] Y. Feng, Z. Zhang, X. Zhao, R. Ji, and Y. Gao, "GVCNN: Group-View Convolutional Neural Networks for 3D Shape Recognition," Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit., pp. 264–272, Dec. 2018, doi: 10.1109/CVPR.2018.00035.

[18] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," Proc. - 30th IEEE Conf. Comput. Vis. Pattern Recognition, CVPR 2017, vol. 2017-January, pp. 77–85, Nov. 2017, doi: 10.1109/CVPR.2017.16.

[19] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in Advances in Neural Information Processing Systems, 2017, vol. 2017-Decem.

[20] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph Cnn for learning on point clouds," ACM Trans. Graph., vol. 38, no. 5, Oct. 2019, doi: 10.1145/3326362;CSUBTYPE:STRING:JOURNAL;ISSUE:ISSUE:DOI.

[21] HanockaRana, HertzAmir, FishNoa, GiryesRaja, FleishmanShachar, and Cohen-OrDaniel, "MeshCNN: A network with an edge," ACM Trans. Graph., vol. 38, no. 4, Jul. 2019, doi: 10.1145/3306346.3322959/SUPPL_FILE/REPOSITORY.ZIP.

[22] Y. Feng, Y. Feng, H. You, X. Zhao, and Y. Gao, "MeshNet: Mesh Neural Network for 3D Shape Representation," Proc. AAAI Conf. Artif. Intell., vol. 33, no. 01, pp. 8279–8286, Jul. 2019, doi: 10.1609/AAAI.V33I01.33018279.

[23] S. Gong, L. Chen, M. Bronstein, and S. Zafeiriou, "SpiralNet++: A Fast and Highly Efficient Mesh Convolution Operator," Proc. - 2019 Int. Conf. Comput. Vis. Work. ICCVW 2019, pp. 4141–4148, Nov. 2019, doi: 10.1109/ICCVW.2019.00509.

[24] M. Bouksim, K. Arhid, F. R. Zakani, M. Aboulfatah, and T. Gadi, "New approach for 3D mesh retrieval using artificial neural network and histogram of features," Sci. Vis., vol. 10, no. 2, pp. 84–94, 2018, doi: 10.26583/SV.10.2.07.

[25] M. Bouksim, M. A. Madani, F. R. Zakani, K. Arhid, T. Gadi, and M. Aboulfatah, "Catboost3D: A Novel CatBoost-based Approach for Efficient Classification of 3D Models," 1st Int. Conf. Comput. Internet Things Microw. Syst. ICCIMS 2024, 2024, doi: 10.1109/ICCIMS61672.2024.10690498.

[26] M. Bouksim, F. Rafii Zakani, K. Arhid, A. Dahbi, T. Gadi, and M. Aboulfatah, "New 3D Shape Descriptor Extraction using CatBoost Classifier for Accurate 3D Model Retrieval," Int. J. Adv. Comput. Sci. Appl., vol. 15, no. 5, pp. 1063–1071, Jun. 2024, doi: 10.14569/IJACSA.2024.01505107.

[27] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, "Shape distributions," ACM Trans. Graph., vol. 21, no. 4, pp. 807–832, 2002, doi: 10.1145/571647.571648.

[28] J. Lin, "Divergence Measures Based on the Shannon Entropy," IEEE Trans. Inf. Theory, vol. 37, no. 1, pp. 145–151, 1991, doi: 10.1109/18.61115.

[29] S. Kullback and R. A. Leibler, "On Information and Sufficiency," https://doi.org/10.1214/aoms/1177729694, vol. 22, no. 1, pp. 79–86, Mar. 1951, doi: 10.1214/AOMS/1177729694.