

A Possible Human-Centered Embedding Space Search in Degenerate Clifford Algebras

Isaac Roberts^{1*}, Louis Mozart Kamdem Teyou², Alexander Schulz¹,
N'Dah Jean Kouagou², Axel-Cyrille Ngonga Ngomo², and Barbara Hammer^{1 †‡}

1- Bielefeld University – Faculty of Technology
Inspiration 1, 33619 Bielefeld – Germany

2- Heinz Nixdorf Institute-Paderborn University –
Faculty of Computer Science,
Electrical Engineering and Mathematics
Fürstenallee 11, 33102 Paderborn – Germany

Abstract. Recent knowledge graph embedding (KGE) models increasingly exploit algebraic structures to encode relational semantics. Clifford-based models, in particular, offer strong expressiveness and geometric interpretability. In this work, we analyze the representations and decision boundaries of such models using an embedding-based reasoner as a classification function. To interpret Clifford-based geometric effects, we adapt DeepView, a visualization framework that approximates decision functions of deep classification models. This study provides one of the first systematic visual analyses of Clifford-based KGE models, helping bridge algebraic representation learning and interpretability.

1 Introduction

Knowledge graphs (KGs) provide a structured representation of real-world data by encoding entities and their relationships as a set of triples. Their graph-based organization enables symbolic reasoning, but large-scale inference over KGs is often computationally demanding. To address this challenge, knowledge graph embedding (KGE) models project entities and relations into algebraic vector spaces where relational patterns can be captured through mathematical operations, enabling classical machine learning inference over the KG.

Over the years, a wide variety of KGE models have been proposed. Bilinear models such as DistMult and ComplEx capture interactions through tensor or complex-valued operations, while geometric models increasingly rely on richer algebraic structures such as Quaternions, Octonions and Clifford algebras. DeCaL [1], based on Clifford algebras, unifies and generalizes many of these approaches: several widely used models—including DistMult, ComplEx, and Keci—emerge as special cases of its multivector formulation.

*Corresponding author: iroberts@techfak.uni-bielefeld.de

†This project has received funding from the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101073307, from the German Federal Ministry of Research, Technology and Space (BMFTR) within the project KI-OWL under the grant no 01IS24057B, and from the Ministry of Culture and Science of North Rhine-Westphalia (MKW NRW) within the project "WHALE" (LFN 1-04).

‡Experiment code: https://github.com/Louis-Mozart/UPB_UBI

Despite their effectiveness, KGE models are typically black boxes. Their decisions arise from high-dimensional algebraic transformations whose geometric meaning is not directly observable. Thus, it remains unclear how these models encode semantic classes and how latent structures influence reasoning behavior. This lack of interpretability is particularly problematic when embeddings are used for downstream reasoning tasks, where trust and transparency are essential.

To address this gap, we analyze the representations learned by DeCaL through the lens of Embedding-Based Reasoning (EBR). We treat EBR [2] as a classifier over ontology classes and visualize its embedding space using discriminative dimensionality reduction (DR) [3]. This enables systematic human inspection of emergent geometric structure in Clifford-based embeddings. By coupling these qualitative observations with targeted experiments, we demonstrate that certain geometric patterns correspond to measurable differences in predictive performance and generalization. Our study provides initial evidence that visual inspection can guide the search over embedding spaces, laying groundwork for a principled, human-centered approach to embedding-space selection.

2 Background

2.1 Embedding-based Reasoning over Knowledge Bases

A knowledge base (KB) consists of a TBox that describes the schema and an ABox containing assertions about individuals. EBR retrieves instances of a class expression by operating entirely in an embedding space. It encodes individuals and concept constructors using a KGE model ϕ , which maps each entity or constructor to a d -dimensional vector and then utilizes a scoring function f to measure how plausible it is that an embedding satisfies a class expression. EBR recursively composes the embeddings of class constructors (e.g., conjunction, existential restriction) and evaluates the plausibility $f(\phi(x), C)$ that an individual x belongs to a class expression C . Individuals are finally ranked or classified based on the scores, permitting EBR to perform instance retrieval even when the KB is incomplete or inconsistent.

Since EBR relies on a knowledge graph embedding model ϕ to represent individuals and class constructors, the choice of its space directly shapes the geometry of the decision boundaries of f . In this work, we instantiate ϕ with DeCaL, a KGE model defined in degenerate Clifford algebras $Cl_{p,q,r}(\mathbb{R})$. These algebras generalize many familiar number systems depending on the signature parameters p , q , and r . This makes DeCaL an ideal embedding model for studying how algebraic structure influences KGE reasoning behavior.

2.2 Visualizing Neural Network Classifier: Deepview

DeepView [3] is a framework to visualize the classification function of a deep neural network classifier together with data. It consists of four core steps which include (i) projecting the data to two dimensions using a discriminative DR method, (ii) sampling a grid in the 2D space and mapping it to the original input

space, (iii) applying the classifier to the projected samples to obtain the predicted class label and a certainty estimate using entropy and finally (iv) visualizing the certainties and predicted classes in the background of the 2D scatter plot to obtain an approximation of the decision function. DeepView typically employs a discriminative variant of UMAP [4], which combines an unsupervised distance (i.e. Euclidean) with a discriminative distance metric that emphasizes directions important to the classifier; however, in this work, we reduce the dimensions of our data preserving only the Euclidean distance, since the model has been trained in a supervised fashion [5] and the inner product induced by Clifford Algebras is isomorphic to computing the Euclidean distance in \mathbb{R} [1]. To measure projection faithfulness, Q_{data} and Q_{-data} are introduced [3] for the positions of the data points and for areas without data, respectively.

3 Methodology

Our objective is to visualize positive and negative instances of a *class expression* C in a two-dimensional projection that preserves the inner-product structure induced by different Clifford algebras. We utilize the plots to make testable observations about the properties of the embedding spaces.

Visualization Setup. To this end, we apply ϕ , trained on a complete KB, to its instances, obtaining d -dimensional vectors embedded in a Clifford algebra, controlled by parameters p , q , and r . We interpret f as a classifier that provides a probability for each embedding. Then we project the d -dimensional vectors to 2D and visualize how f separates instances of C in the projected space.

Visualization Orientation. Figure 1 shows visualizations of class expressions across several Clifford algebras, each denoted by its corresponding pqr parameters. In these plots, the dots represent an entity in the Family [6] dataset with its corresponding label depicted with its respective color. Negative inputs are shown in blue, positive inputs in green, and uncertain regions (if any) appear in white. The background color represents f 's decision function i.e., what the f predicts in this region. From Figure 1, we derive several key observations.

Observation 1: Distinctive Predictive Regions. For $pqr = 000$ and $pqr = 111$, we observe little, if any, green color in the background, indicating no positive prediction region. For 000 , the white regions around the positive samples suggest uncertainty in the classifier's decision. On the contrary, $pqr = 100$, we observe smaller blue regions in the background. By contrast, $pqr = 001$ and $pqr = 010$ exhibit more structured and distinct predictive regions. Sometimes, f 's decision boundary is linearly separable, such as in $pqr = 001$ on the "Male" class.

Observation 2: Clustering of Similarly Labeled Data Points. Ideally, inputs with the same label cluster together due to shared similar attributes. For $pqr = 000$ and $pqr = 111$, the positive examples appear dispersed, while $pqr = 001$, 010 , and 100 demonstrate clearer clustering behavior among both sets of inputs.

Desirable Embedding Properties. From these observations, we identify three key properties: 1. ϕ should embed similarly-labeled inputs closely, 2. f should form distinguishable regions corresponding to negatively and positively predicted

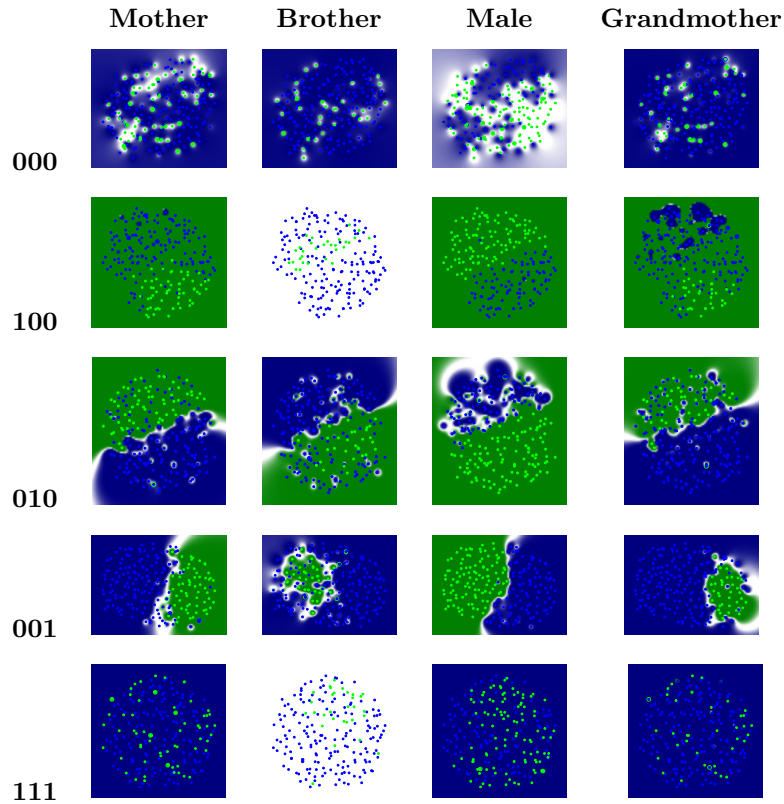


Fig. 1: Deepviews of Clifford embeddings using the Grandmother, Brother, Male, and Mother class expressions from the Family dataset. The average Q_{data} is $> 99\%$ and Q_{-data} is 89.71% over the pqr signatures and class expressions.

inputs, and 3. closely embedded inputs should be predicted similarly. Considering the class expressions visualized in Figure 1, the embeddings from $pqr = 001$ and $pqr = 010$ yield representations that most exhibit the desired properties.

4 Experiment Evaluation

In the previous section, we have identified promising pqr configurations based on DeepView visualizations. Since the projections provide an approximation of the classifier, we validate these observations in the following experiments.

Classifier Robustness. We randomly remove (set to negative) a percentage of positive assertions *from a single class* in the Family dataset [6]. DeCaL models are trained transductively using $pqr \in \{000, 100, 010, 001, 111\}$. Performance is evaluated over the entire dataset using Recall, Precision, Accuracy, and Jaccard score. The experiment is repeated 10 times for each class, with removal

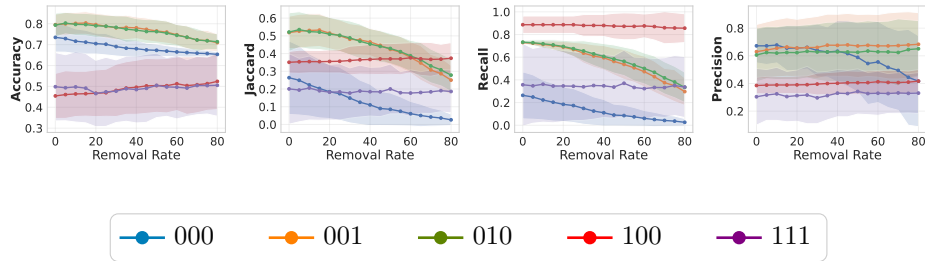


Fig. 2: We report each classifier’s performance averaged over the classes with predictive metrics(↑) for each pqr across incompleteness levels.

percentages ranging from 0 to 80 in increments of 5.

Structural Robustness. To evaluate embedding structure, positive assertions are randomly removed *across all classes*, and a DeCaL model is trained on the modified KB. Since entities may belong to multiple classes, this is treated as a multi-label classification task. After training, predictions from f are used to train a multi-label kNN (mlkNN) classifier [7]. Performance is measured using Hamming loss over the entire dataset and over entities with removed assertions.

5 Discussion

The DeepView-based visualizations shown in Figure 1 indicate that embeddings produced with $pqr = 001$ and $pqr = 010$ exhibit clearer clustering of similarly labeled entities and more structured decision regions compared to other configurations. Our analysis validated this with two experiments.

Classifier robustness experiments highlight that these configurations $pqr = 001$ and $pqr = 010$ maintain the highest accuracy across increasing levels of missing positive assertions, whereas $pqr = 100$ displays a distinct trade-off between recall and precision, indicating many false positives. This behavior is consistent with the visualizations, exhibiting extensive positive prediction regions.

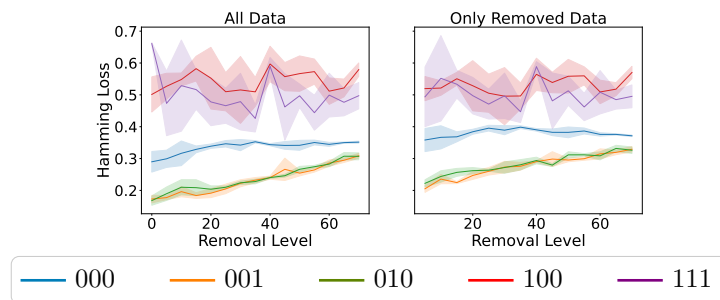


Fig. 3: We report the Average Hamming loss (↓) of the mlkNN when $k = 5$ for pqr across removal levels. Shaded areas represent the standard deviation.

Structural robustness results reinforce these findings. The lower Hamming loss achieved by $pqr = 001$ and $pqr = 010$ suggests that their embedding spaces preserve neighborhood structure in a way that is consistent with semantic class membership. This indicates that, even when explicit assertions are removed during training, the learned embeddings retain sufficient relational information to recover missing labels through nearest-neighbor reasoning.

The combined qualitative and quantitative results demonstrate that the choice of Clifford algebra parameters has a pronounced effect on both the geometry of the embedding space and the robustness of downstream predictions. The visualization analysis and robustness experiments provide mutually reinforcing evidence regarding the relative behavior of different pqr configurations.

Conclusion We visually inspected the learned representations of DeCaL with Embedding-based Reasoning. This analysis enabled key observations that affect downstream performance to be extracted and further tested. Namely, we notice that ϕ should map similarly-labeled inputs closely and that f should have distinctive predictive regions. Thus, we measure their alignment by removing information from the knowledge base, focusing on the classifier’s robustness and the embedding’s structural robustness. Through our experiments, we confirm that $pqr = 001$ and 010 provide the best embeddings for the Family dataset. In further work, we aim to expand our findings to larger datasets and eventually automate the process of visual embedding selection. Additionally, a user study of the embedding space selection process is necessary to verify the potential of our human-centered approach.

References

- [1] Louis Mozart Kamdem Teyou, Caglar Demir, and Axel-Cyrille Ngonga Ngomo. Embedding knowledge graphs in degenerate clifford algebras. In *ECAI 2024*, pages 1293–1300. IOS Press, 2024.
- [2] Louis Mozart Kamdem Teyou, Luke Friedrichs, N’Dah Jean Kouagou, Caglar Demir, Yasir Mahmood, Stefan Heindorf, and Axel-Cyrille Ngonga Ngomo. Neural reasoning for robust instance retrieval in shoiq. In *K-CAP 2025*, pages 61–68, 2025.
- [3] Alexander Schulz, Fabian Hinder, and Barbara Hammer. Deepview: Visualizing classification boundaries of deep neural networks as scatter plots using discriminative dimensionality reduction. In Christian Bessiere, editor, *IJCAI*, pages 2305–2311, 7 2020. Main track.
- [4] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction, 2020.
- [5] Isaac Roberts, Alexander Schulz, Sarah Schroeder, Luca Hermes, and Barbara Hammer. Targeted visualization of the backbone of encoder llms. In *ECML Workshop AIMLAI*, In Press 2024.
- [6] Caglar Demir and Axel-Cyrille Ngonga Ngomo. Neuro-symbolic class expression learning. In *IJCAI*, pages 3624–3632, 2023.
- [7] Min-Ling Zhang and Zhi-Hua Zhou. Ml-knn: A lazy learning approach to multi-label learning. *Pattern recognition*, 40(7):2038–2048, 2007.