Detecting Knot by Mosaic Image

Touseef Haider[1]

[1]Department of Mathematics and Computer Sciences, Rutgers - Newark, Newark, USA

October 30, 2025

Abstract

Recently machine learning model has been used to recognize knot type. We apply convolutional neural networks to detect knot types in rigid diagrams created with Mosaics (introduced by Lomonaco and Kauffman).

Keywords: Knot theory, Mosacis tile, Neural Network, CNN, Few-Shot Learning, Machine Learning, Computer Vision.

MSC (2020): 57K10; 57M25; 68T07.

1 Introduction

A central problem in knot theory is determining whether two knots are equivalent. The known algorithm for solving this problem operates with tower of exponential complexity relative to the crossing number of the knots [CL14]. Recent advances have produced more practical algorithms, including a polynomial time algorithm for alternating knots and links by author and Tsvietkova [HT25]. However, the development of an efficient algorithm for general knot recognition remains unresolved.

The application of artificial intelligence (AI) to knot theory has increased significantly, with AI-driven methods providing new insights and experminetal approaches to longstanding problems. For instance, there is work on the unknot decision problem [GHRS21], the slice ribbon conjecture [GHMR23], and the study of the properties of the Jones polynomial using dimensionality reduction and topological data analysis [LHS22]. Kauffman et al. [KRT22] investigate the use of neural networks to learn Dynnikov moves for knot recognition. Furthermore, deep learning approaches have been employed to study algebraic and geometric structures [HZDM20], as well as the relationship between knot invariants through topological data analysis tools [DGS21].

Hughes [HES] developed generative adversarial networks (GANs) that take a Jones polynomial as input and output corresponding knot. Reinforcement learning methods, such as Q-Learning and Deep Q-Learning, have been applied to the problem of untangling braids [VL+22]. Additional approaches, including linear regression and deep neural network have been utilized to identify relationships among hyperbolic knot invariants [Grü22].

Davis et al. [DVB+21, DJLT24] relate the signature, slope, volume, and injectivity radius of hyperbolic knots an insight that was reveal by neural network model. Jejjala et al. [JKP19] used deep learning to estimate the hyperbolic volume of knots. Craven [CJK21] approximated the hyperbolic volume of a knot using a single evaluation of its Jones polynomial at a root of

unity. More recently, Craven and Jejjala [CHJK22] explored the relationships between over 40 different knot invariants. More recently the colored jones polynomial has been used for volume conjecture using the neural network [HJR+25].

In this study, we use computer vision techniques to recognize knots from their mosaic diagrams. Convolutional Neural Networks (CNNs) are effective in image pattern recognition and can be trained to interpret knot diagrams and predict their properties. Dranowski et al. [DKT25] recently proposed a strategy combining CNN-based image analysis with traditional knot invariants. Their methodology use neural network first to extract the knot's features from a picture, and then use classic topological invariants to verify or refine the classification. This two-step "recognize-then-compute" approach integrate the strengths of both domains: CNN handles noisy visual data, while mathematical invariants provide rigorous classification.

To enhance CNN based knot recoginin, a more structured representation of knot diagram is utilized in this study. Knot mosaics, introduced by Lomonaco and Kauffman [LJK08], provide a systematic method for encoding knot diagrams on a grid. Knot mosaics were originally part of quantum knot system, which maps knot diagrams to states in a Hilbert space. Kauffman and others explored application of quantum algorithms to compute Jones polynomial [LJK08]. There have been long list of papers by Lomonaco and Kauffman about use of mosaic in quantum related area but it is beyond the scope of this study.

This work focuses on square mosaic grids on the plane. Although, there are work about hexagonal mosaic tiles [HLL19] and square mosaics on the torus [OHL+17]. The planar square mosaic is selected because it is more developed and have tabulations available.

In a (square) knot mosaic, a continuous knot diagram is decomposed into an $n \times n$ grid of mosaic tiles. There are exactly 11 types of mosaic tiles (figure 1), which are sufficient to construct any knot diagram. These tiles include straight segments, curved segments, crossings, and empty space, arranged to ensure that connected strands continue from one tile to the next. Every knot (indeed every link) can be represented on a mosaic board using some configuration of these 11 tiles, providing a universal discrete encoding of knots. Mosaic knot theory has been proven equivalent to standard tame knot theory (the Lomonaco–Kauffman conjecture, proved by Kuriya & Shehab [KS14]). That is any ambient isotopy of knots can be realized through a sequence of moves on the mosaic, such as adding, removing, or altering tiles, which are analogous to Reidemeister moves. If two knots are topologically equivalent, their mosaic diagrams are connected by a series of tile transformations, and conversely, any valid tile moves that transform one mosaic into another correspond to an isotopy of the underlying knots. This ensures that the mosaic representation preserves all topological information of the knot.

Figure 1: Mosaic Tiles

This study employs a Convolutional Neural Network (CNN) approach to address the knot identification problem. Compare to work of Dranowski et al. [DKT25], who did not use mosaic tiles and included a verification step based on computed invariants such as the Jones polynomial, this study utilizes mosaic images of knots. These images provide a more structured and less noisy representation. Although a verification step is not included, the CNN model demonstrates efficient knot identification.

Despite the structured nature of mosaic images, there are limitations to their use. The number of distinct mosaics increases rapidly with board size, making comprehensive tables for higher crossing knots challenging. For example there are over a billion knot mosaics on 6×6 board [OHLL15]). Therefore, this study is restricted to the available tabulation of up to 10-crossing prime knots available in [LLPP18, Hea25]. In practice, a CNN-based pipeline using mosaics represents a knot as a matrix of discrete tile identifiers, which are then used to generate mosaic images. During training, the network learns to recognize specific configurations of tiles corresponding to known knot. For instance, particular arrangements of crossing and arc tiles may indicate a trefoil knot or a figure-eight knot. Because the input is limited to valid mosaic patterns, the CNN's classification task is simplified. However, with only 249 knots in the dataset, the image database is limited, and a standard machine learning model would not be able to learn each knot type effectively. To address this, a variation of CNN suitable for problems with limited examples per class is used, incorporation rotated images of each knot at 90, 180 and 270 degrees. Section 2 provides a brief explanation of this model and the preliminary concepts of mosaic knot theory.

Section 3 describes the data processing, CNN with prototypical network model training and evaluation, and the results. Section 4 discusses potential improvements and future research directions.

2 Prerequisite

Mosaic Number of a knot L is the size of the smallest square grid needed to represent the knot. It is defined as the minimum integer n such that the knot can be drawn on an $n \times n$ matrix of tiles.

Mosaic number of a knot L is related to the crossing number; an upper bound on mosaic number in terms of crossing number is given in [LHLO14]. They proved that for any nontrivial knot (or non-split link) except the Hopf link, the mosaic number is at most one more than the crossing number. Thus, mosaics do not drastically increase the complexity of knot representations.

The tile number of a knot is the minimum number of non-blank tiles required to form a mosaic diagram of the knot. Although the total number of tiles in an $n \times n$ grid is n^2, many tiles may remain blank (tile T_0). Heap & Knowles [HK18] demonstrate that numerous knots can be represented efficiently, using significantly fewer tiles than the full n^2 grid.

2.1 Prototypical Networks for few-shot classification

In neural networks, the well-known image classification model is the convolutional neural network (CNN), which was used by Dranowski et al. [DKT25] for knot identification. Standard CNNs require large numbers of labeled examples. In contrast, metric learning approaches trained episodically can leverage structural inductive biases to generalize from limited data. This study evaluates prototypical networks trained episodically on rotated images generated from standard knot mosaic images.

Episodic prototypical networks [SSZ17] is used to this knot identification problem because this architecture is well-suited when there are many distinct classes but few examples per class. For training we uses rotated variants of the available examples to form the training pool while evaluation is performed on held-out standard mosaic images. We compare three episodic configurations of training (20-way, 10-way, 5-way).

3 Experiment

3.1 Data

- The original mosaic tabulation images exhibit inconsistencies in format and background color. Therefore, the mosaic images were recreated using the tiles shown in the Figure 1 to ensure consistency for training purposes.
- Rotated version of each knot image were generated at $90, 180$ and 270 degree. These rotations yield valid diagrams of the same knot, as planar rotation does not alter knot topology. The resulting images comprised the training pool.
- Mirror images were excluded from the dataset because not all knots are amphicheiral.
- The dataset contains 249 unique prime knots with up to 10 crossings.
- Afer generating rotated images, the training pool consisted of 747 rotated samples, while the validation set included 249 original images.

3.2 Model and training

We adopt a prototypical network framework. The embedding network is a lightweight convolutional feature extractor followed by a projection head producing 128-dimensional embeddings. During each episode, the sampler constructs N-way K-shot tasks by selecting K support examples and Q query examples for each of the N classes. Class prototypes are computed as the mean of support embeddings for each class, and queries are claasified based on the nearest prototype in Euclidean space. Training optimizes the cross-entropy loss between softmax-normalized negative squared distances and the query labels.

The Adam optimizer was used with a learning rate of $1e^{-3}$. The number of episodes per epoch and total epochs were selected to ensure stab per-epoch evaluation curves, as detailed in the experiments section.

3.3 Experimental protocol

We evaluate combinations of episode width $N \in \{5, 10, 20\}$ and support size $K \in \{2, 3\}$ with a consistent query count $Q = 3$ per class. These settings are referred to as "3-shot" and "2-shot" respectively. Each run was trained for 10 epochs with 200 episodes per epoch.

For each configuration, class prototypes were constructed from the rotated training pool after each epoch. Top-1, top-2 and top-3 accuracy were evaluated on the held-out set. These metrics were logged per epoch and used to generate comparison plots.

The evaluation metrics used throughout the paper are defined as follows. Top-1 accuracy is the fraction of query examples for which the nearest prototype (the single predicted class) equals the true class. Top-2 and top-3 accuracy indicate whether the true class appears among the two or three nearest prototypes respectively. Because we hold out a single canonical example per class, the top-2 metric offers an informative complement to top-1 while top-3 highlights whether the correct class remains near the top of the similarity ranking.

3.4 Results

Table 2 presents the final-epoch evaluation performance, including top-1, top-2, and top-3 accuracy, for each (N, K) combination. Figure 2a, 2b shows the top-1 accuracy trajectories, Figure 2c, 2d highlights top-2 accuracy, and Figure 2e 2f shows top-3 accuracy curves by support size. These curves illustrate how quickly each setting reaches high accuracy and how performance

shifts as the number of support examples decreases.

For comparison, classical baseline models were also trained on the rotated training pool, with canonical originals reserved for evaluation. Table 1 shows their performance. Both k-nearest neighbors (k-NN) and radial basis function support vector machine (RBF SVM) performed poorly under the limited-data set, underscoring the effectiveness of the episodic metric-learning approach.

Table 1: Baseline classifiers trained on rotated samples and evaluated on canonical originals.

Model	Top-1	Top-2	Top-3
k-NN ($k = 3$, distance weights)	0.0482	0.0683	0.0843
SVM (RBF kernel, $C = 1$)	0.0040	0.0080	0.0120

Table 2: Final epoch evaluation on held-out canonical examples (Q=3 queries per class during evaluation).

Shot	Configuration	Eval top-1	Eval top-2	Eval top-3
3-shot	20-way	0.9639	0.9960	0.9960
	10-way	0.7992	0.9157	0.9518
	5-way	0.6667	0.8795	0.9116
2-shot	20-way	0.9679	0.9960	0.9960
	10-way	0.9237	0.9839	0.9960
	5-way	0.8353	0.9558	0.9759

The prototypical training demonstrates strong generalization to held-out canonical examples when the training pool contains rotated variants. With three support examples, the 20-way configuration attains top-1 of 0.96 and ranks the correct class within the top two 99.6% of the time. The 5-way configuration maintains high top-2 (0.88) despite lower top-1 score, indicating that most misclassifications are ranked second. When reducing to two support examples, performance improves for broader episodes: both 20-way and 10-way setting exceed 0.98 top-2 accuracy and approach perfect top-3 results, while the 5-way configuration recovers to 0.96 top-2 and 0.98 top-3. In contrast, classical classifiers trained on the same rotated pool (Table 1) remain below 9% top-3 accuracy, highlighting the limitations of non-episodic methods under limited canonical supervision. The observed variance in 3-shot 5-way episodes indicates that episode composition still matters, suggesting the need for extended training schedules or curriculum strategies to tabilize low-shot performance.

A primary limitation of this approach is its dependence on rotated/augmented variants to expand the training pool, while canonical examples per class remain limited. Future work could evaluate robustness to more realistic covariate shifts and explore alternative prototype estimators (e.g., median or robust mean) when rotated variants are noisy.

4 Future Direction

The experiment considered only $90°, 180°$ and $270°$ planar rotations to the mosaic images, which are topologically trivial transformations. The knot equivalence problem requires knots to be ambient isotopic, which is more complex than simple rotation. The next phase of this study is to extended the experiment by first generating augmentation images of each class by implementing the mosaic tile equivalent Reidmeister moves. Training a network on this a dataset is expected to

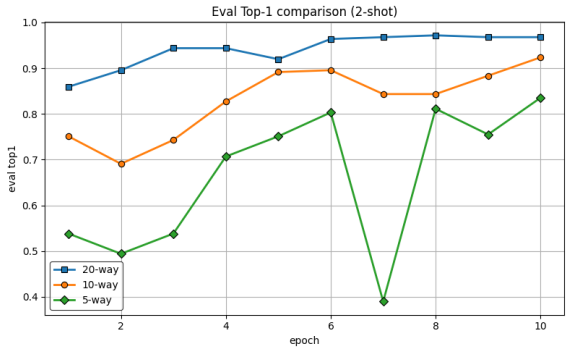

(a) Top-1, 2-shot.

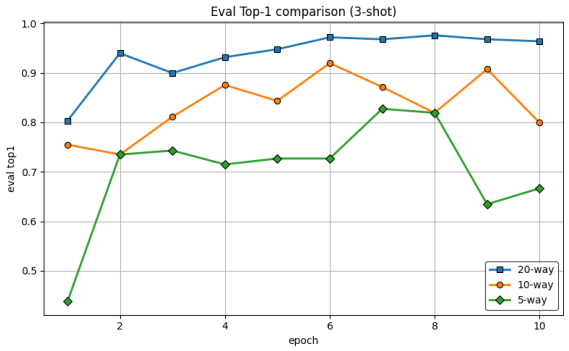

(b) Top-1, 3-shot.

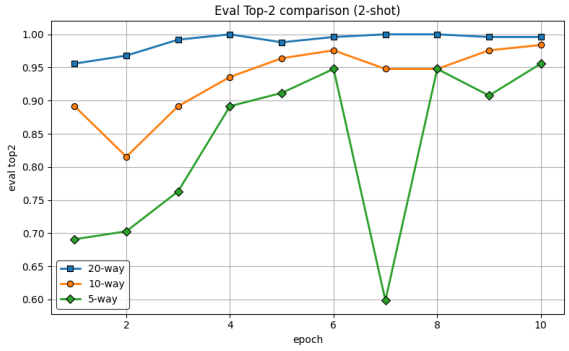

(c) Top-2, 2-shot.

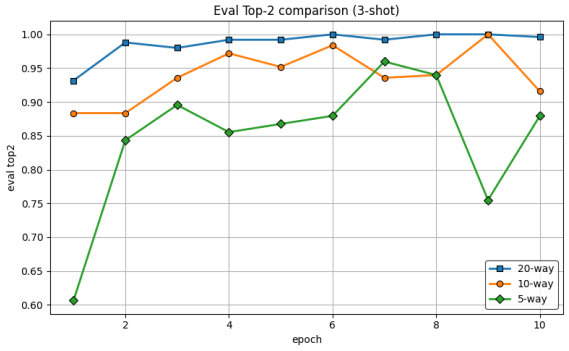

(d) Top-2, 3-shot.

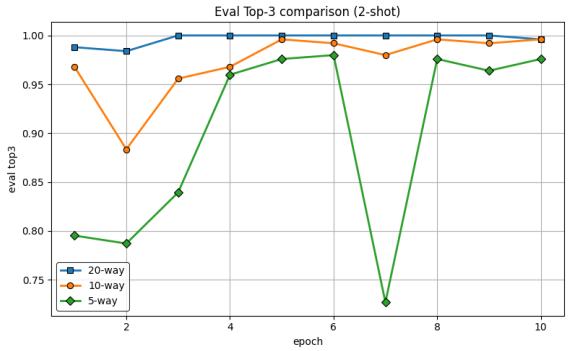

(e) Top-3, 2-shot.

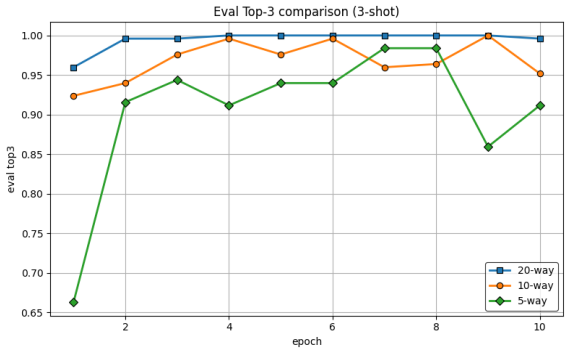

(f) Top-3, 3-shot.

Figure 2: Per-epoch evaluation curves grouped by support size for top-1, top-2, and top-3 metrics.

encourage the model to learn topological feature rather than relying solely on geometric rotations.

The dataset used in this study was limited in size. A future direction involves creating mosaic images for knots and links with higher crossing numbers. The number of possible $n \times n$ knot mosaics increases rapidly with n, making exhaustive tabulation infeasible for larger crossing number. Work by Hong et al. [HOLL14, LHLO14] and Lee et al. [LLPP18] demonstrates the exponential grow in the number of possible knot mosaic.

A critical sub-problem in this direction is the conversion of hand-drawn or computer-generated knot diagrams into mosaic representations. Integrating the approach of Dranowski et al. [DKT25] and also incorporate the verification step based on knot invariants would strengthen knot detection and mitigate errors introduced by convolutional neural networks.

Remark: All the data and CNN architect is available in the GitHub repository : https://github.com/thaidermath/knot_mosaic_cnn

References

[CHJK22] Jessica Craven, Mark Hughes, Vishnu Jejjala, and Arjun Kar, *Illuminating new and known relations between knot invariants*, arXiv preprint arXiv:2211.01404 (2022).

[CJK21] Jessica Craven, Vishnu Jejjala, and Arjun Kar, *Disentangling a deep learned volume formula*, Journal of High Energy Physics **2021** (2021), no. 6, 1–40.

[CL14] Alexander Coward and Marc Lackenby, *An upper bound on reidemeister moves*, American Journal of Mathematics **136** (2014), no. 4, 1023–1066.

[DGS21] Pawel Dlotko, Davide Gurnari, and Radmila Sazdanovic, *Knot invariants and their relations: a topological perspective.*, CoRR (2021).

[DJLT24] Alex Davies, András Juhász, Marc Lackenby, and Nenad Tomašev, *The signature and cusp geometry of hyperbolic knots*, Geometry & Topology **28** (2024), no. 5, 2313–2343.

[DKT25] Anne Dranowski, Yura Kabkov, and Daniel Tubbenhauer, *On knot detection via picture recognition*, 2025.

[DVB+21] Alex Davies, Petar Veličković, Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomašev, Richard Tanburn, Peter Battaglia, Charles Blundell, András Juhász, et al., *Advancing mathematics by guiding human intuition with ai*, Nature **600** (2021), no. 7887, 70–74.

[GHMR23] Sergei Gukov, James Halverson, Ciprian Manolescu, and Fabian Ruehle, *Searching for ribbons with machine learning*, Machine Learning: Science and Technology (2023).

[GHRS21] Sergei Gukov, James Halverson, Fabian Ruehle, and Piotr Sułkowski, *Learning to unknot*, Machine Learning: Science and Technology **2** (2021), no. 2, 025035.

[Grü22] Daniel Grünbaum, *Narrowing the gap between combinatorial and hyperbolic knot invariants via deep learning*, Journal of Knot Theory and Its Ramifications **31** (2022), no. 01, 2250003.

[Hea25] Aaron et. al. Heap, *Knot mosaic space*, https://www.geneseo.edu/knotmosaics, 2025, Created by Aaron Heap and his students; Copyright © 2025 SUNY Geneseo.

[HES] Mark Hughes, Amy Eubanks, and Jared Slone, *Using generative adversarial networks to produce knots with specified invariants*, 2021 Fall Western Sectional Meeting, AMS.

[HJR+25] Mark Hughes, Vishnu Jejjala, P. Ramadevi, Pratik Roy, and Vivek Kumar Singh, *Colored jones polynomials and the volume conjecture*, 2025.

[HK18] Aaron Heap and Douglas Knowles, *Tile number and space-efficient knot mosaics*, Journal of Knot Theory and Its Ramifications **27** (2018), no. 06, 1850041.

[HLL19] Hugh Howards, Jiong Li, and Xiaotian Liu, *An infinite family of knots whose hexagonal mosaic number is only realized in non-reduced projections*, arXiv preprint arXiv:1912.03697 (2019).

[HOLL14] Kyungpyo Hong, Seungsang Oh, Ho Lee, and Hwa Jeong Lee, *Upper bound on the total number of knot n-mosaics*, Journal of Knot Theory and Its Ramifications **23** (2014), no. 13, 1450065.

[HT25] Touseef Haider and Anastasiia Tsvietkova, *Polynomial algorithm for alternating link equivalence*, 2025.

[HZDM20] Mustafa Hajij, Ghada Zamzmi, Matthew Dawson, and Greg Muller, *Algebraically-informed deep networks (aidn): A deep learning approach to represent algebraic structures*, arXiv preprint arXiv:2012.01141 (2020).

[JKP19] Vishnu Jejjala, Arjun Kar, and Onkar Parrikar, *Deep learning the hyperbolic volume of a knot*, Physics Letters B **799** (2019), 135033.

[KRT22] LH Kauffman, NE Russkikh, and IA Taimanov, *Rectangular knot diagrams classification with deep learning*, Journal of Knot Theory and Its Ramifications **31** (2022), no. 11, 2250067.

[KS14] Takahito Kuriya and Omar Shehab, *The lomonaco–kauffman conjecture*, Journal of Knot Theory and Its Ramifications **23** (2014), no. 01, 1450003.

[LHLO14] Hwa Jeong Lee, Kyungpyo Hong, Ho Lee, and Seungsang Oh, *Mosaic number of knots*, Journal of Knot Theory and Its Ramifications **23** (2014), no. 13, 1450069.

[LHS22] Jesse SF Levitt, Mustafa Hajij, and Radmila Sazdanovic, *Big data approaches to knot theory: understanding the structure of the jones polynomial*, Journal of Knot Theory and Its Ramifications **31** (2022), no. 13, 2250095.

[LJK08] Samuel J Lomonaco Jr and Louis H Kauffman, *A quantum manual for computing the jones polynomial*, Quantum Information and Computation VI, vol. 6976, SPIE, 2008, pp. 73–76.

[LLPP18] Hwa Jeong Lee, Lewis Ludwig, Joseph Paat, and Amanda Peiffer, *Knot mosaic tabulation*, Involve, a Journal of Mathematics **11** (2018), no. 1, 13–26.

[OHL+17] Seungsang Oh, Kyungpyo Hong, Ho Lee, Hwa Jeong Lee, and Mi Jeong Yeon, *Period and toroidal knot mosaics*, Journal of Knot Theory and Its Ramifications **26** (2017), no. 05, 1750031.

[OHLL15] Seungsang Oh, Kyungpyo Hong, Ho Lee, and Hwa Jeong Lee, *Quantum knots and the number of knot mosaics*, Quantum Information Processing **14** (2015), no. 3, 801–811.

[SSZ17] Jake Snell, Kevin Swersky, and Richard Zemel, *Prototypical networks for few-shot learning*, Proceedings of the 31st International Conference on Neural Information Processing Systems (Red Hook, NY, USA), NIPS'17, Curran Associates Inc., 2017, p. 4080–4090.

[VL+22] Alexei Vernitski, Alexei Lisitsa, et al., *Reinforcement learning algorithms for the untangling of braids*, The International FLAIRS Conference Proceedings, vol. 35, 2022.