

Feature Engineering vs. Deep Learning for Automated Coin Grading: A Comparative Study on Saint-Gaudens Double Eagle Gold Coins

Tanmay Dogra, Eric Ngo, Mohammad Alam, Jean-Paul Talavera, Asim Dahal

Department of Electrical and Computer Engineering

Virginia Tech

Blacksburg, VA, USA

tanmaydogra@vt.edu

Abstract—We challenge the common belief that deep learning always trumps older techniques, using the example of grading Saint-Gaudens Double Eagle gold coins automatically. In our work, we put a feature-based Artificial Neural Network built around 192 custom features pulled from Sobel edge detection and HSV color analysis up against a hybrid Convolutional Neural Network that blends in EfficientNetV2, plus a straightforward Support Vector Machine as the control. Testing 1,785 coins graded by experts, the ANN nailed 86% exact matches and hit 98% when allowing a 3-grade leeway. On the flip side, CNN and SVM mostly just guessed the most common grade, scraping by with 31% and 30% exact hits. Sure, the CNN looked good on broader tolerance metrics, but that is because of some averaging trick in regression that hides how it totally flops at picking out specific grades. All told, when you are stuck with under 2,000 examples and lopsided classes, baking in real coin-expert knowledge through feature design beats out those inscrutable, all-in-one deep learning setups. This rings true for other niche quality checks where data's thin and know-how matters more than raw compute.

Index Terms—Automated coin grading, Artificial Neural Networks, Convolutional Neural Networks, Feature extraction, Sobel edge detection, Transfer learning, Class imbalance, Numismatics

I. INTRODUCTION

The Saint-Gaudens Double Eagle, minted from 1907 to 1933, represents one of the most sought-after collectible coins in American numismatics [1]. Professional graders evaluate these gold coins on the 70-point Sheldon scale [2], and their assessments directly impact market value. Single-grade differences can represent thousands of dollars. However, traditional manual grading suffers from subjectivity, inter-grader variability, and significant time delays. These limitations drive the need for automated grading systems.

Recent advances in computer vision and machine learning allow for automated approaches to numismatic analysis [3], [4]. Two fundamental paradigms exist: feature-based approaches that rely on domain-expert-designed feature extractors, and end-to-end deep learning approaches that learn features directly from raw images. Although CNNs have achieved success in various image classification tasks [5], their effectiveness in specialized domains with limited labeled data is still an open question.

We address this gap by developing and comparing feature-based and hybrid deep learning approaches for automated Saint-Gaudens coin grading. Our primary contributions include:

- A comprehensive feature extraction pipeline using Sobel edge detection with wedge-based spatial analysis, K -means color clustering in HSV space, and perceptually-weighted brightness computation, yielding 192 features
- A comparative evaluation of an ANN using extracted features, a hybrid CNN with EfficientNetV2 backbone processing raw images concatenated with engineered features, and an SVM baseline
- Evidence that feature engineering with SMOTE outperforms end-to-end deep learning when training data is limited and class-imbalanced
- Analysis of why domain-specific features provide better performance, with implications for similar quality assessment applications

II. RELATED WORK

1) *Automated Coin Analysis*: Computer vision has been used for coin evaluation for about 20 years. Early work used histogram matching to grade Lincoln cents, achieving results close to human graders [6]. Later examples include CNNs for identifying Philippine coins [3] and Siamese networks for checking both sides of coins [4]. Most studies focus on identifying coin types rather than assigning detailed grades on the Sheldon scale. Detailed grading requires detecting small signs of wear, scratches, and shine, which needs either custom features or large datasets.

2) *The Sheldon Grading Scale*: William Herbert Sheldon created the Sheldon scale in 1949, and the American Numismatic Association updated it in the 1970s [7]. It uses a 70-point system to rate coin condition. For uncirculated coins (grades 60-70), the main factors are surface condition, how sharply the details were struck, shine, and overall look. Grading services like PCGS and NGC follow this scale, but graders agree only about 85-90% of the time within one grade [8]. This sets a practical limit for automated systems.

3) *Feature Engineering vs. Deep Learning*: In image tasks, manual features compete with features learned by models. Deep learning works well with large datasets, but custom features perform better when data is limited or the task is specific [9]. Tools like Sobel filters detect edges [10], which help measure textures. HSV color space is used for color analysis [11].

III. METHODOLOGY

A. Dataset Description

We compiled a dataset of 1,785 Saint-Gaudens Double Eagle gold coins from the David Lawrence Rare Coins (DLRC) database. PCGS or NGC professionally graded each coin, with grades ranging from MS-50 to MS-68 on the Sheldon scale. We excluded coins graded below 50 due to insufficient representation and our focus on Mint State specimens.

The dataset has severe class imbalance typical of real-world numismatic collections. Grade MS-64 is the largest class with 165 samples in the test set, while grades MS-50, MS-51, and MS-68 contain fewer than 15 samples each. This imbalance presents significant challenges for machine learning, particularly for methods sensitive to class distribution.

We captured high-resolution images of both obverse (Liberty) and reverse (Eagle) sides for each coin under standardized lighting conditions. We split the dataset 70%/30% for training and testing (1,248 training coins and 537 test coins), with stratification to maintain grade proportions across splits.

B. Approach Overview

We implemented and compared two different paradigms:

Feature-Based Approach: We extracted domain-specific features encoding numismatic expertise from coin images, then fed them to shallow machine learning models (ANN, SVM) for grade prediction.

Hybrid CNN Approach: We processed raw coin images with a deep CNN using transfer learning while also feeding the engineered features, allowing the network to leverage both learned and domain-specific representations.

C. Feature-Based Approach

1) *Condition Feature Extraction*: Coin condition assessment depends on detecting surface wear, which appears as smoothing of fine details and reduction of edge sharpness. We used the Sobel edge detection operator [10] to quantify edge characteristics, as it provides continuous gradient values suited for measuring wear progression.

For an input grayscale image I , the Sobel operator computes horizontal and vertical gradient components:

$$G_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} * I, \quad G_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} * I \quad (1)$$

We compute the gradient magnitude as:

$$G = \sqrt{G_x^2 + G_y^2} \quad (2)$$

We applied Gaussian blurring with $\sigma = 1.5$ before edge detection to reduce noise while preserving significant edges. Background removal isolated the coin from its holder, and prong removal eliminated mounting artifacts that could confound analysis.

2) *Wedge-Based Spatial Analysis*: To capture localized wear patterns and increase feature dimensionality, we divided each coin into angular wedges using circular masking. We define the wedge mask M_k for the k -th slice spanning angles $[\theta_k, \theta_{k+1}]$ as:

$$M_k(x, y) = \begin{cases} 1 & \text{if } \theta_k \leq \arctan\left(\frac{y-c_y}{x-c_x}\right) < \theta_{k+1} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where (c_x, c_y) is the coin center. We evaluated both 4-slice and 8-slice configurations, with 8 slices providing finer spatial resolution.

For each wedge region, we extracted statistical features from the gradient components:

- Minimum, maximum, mean, and median of G_x
- Minimum, maximum, mean, and median of G_y
- Minimum, maximum, mean, and median of G (magnitude)

This yields 12 features per wedge times 8 wedges times 2 coin sides, giving us 192 edge-based features. We also masked the eagle region on the reverse side for focused analysis, as it contains the highest detail density.

3) *Color Feature Extraction*: Gold coin coloration varies based on alloy composition, environmental exposure, and surface oxidation, all of which influence perceived grade. We performed color analysis in HSV (Hue, Saturation, Value) color space, which separates chromatic content from intensity. This is better for numismatic analysis where we must normalize lighting variations.

We applied K -means clustering with $K = 5$ to the HSV values of coin pixels, grouping coins into five color categories:

- 1) Rustic Gold: Lower saturation, darker tones indicating age patina
- 2) Golden Bronze: Moderate saturation with warm undertones
- 3) Autumn Gold: Rich golden hues with slight oxidation
- 4) Golden Sand: Lighter, more yellow-dominant coloration
- 5) Sunlit Gold: High saturation, bright golden appearance

The cluster assignment serves as a categorical feature encoding the coin's dominant color profile.

4) *Brightness Feature Extraction*: Luster, the quality of light reflection from a coin's surface, is a critical grading criterion. We developed a perceptually-weighted brightness measure:

$$B = \sqrt{S \times V} \cdot \exp\left(-\frac{(H - H_0)^2}{2\sigma_H^2}\right) \quad (4)$$

where H , S , V represent the mean hue, saturation, and value of the coin image, H_0 is the target gold hue angle, and σ_H controls the hue weighting bandwidth. This formulation

emphasizes brightness contributions from gold-colored regions while reducing weight for off-hue areas.

We discretized brightness values into five categories: Dim, Soft, Bright, Vivid, and Brilliant, matching the terminology used by professional numismatists.

5) *Feature Vector Construction*: We constructed the complete feature vector for each coin from:

- 192 edge gradient features (96 obverse + 96 reverse)
- 6 mean HSV values (3 obverse + 3 reverse)
- 1 color cluster assignment
- 2 brightness levels (obverse + reverse)
- 1 grading service indicator (PCGS/NGC)

Total dimensionality: 202 features per coin. We standardized features to zero mean and unit variance using training set statistics.

6) *ANN Architecture*: We implemented a feedforward neural network using Keras with the following architecture:

- Input Layer: 202 neurons (feature dimensionality)
- Hidden Layer 1: 128 neurons, ReLU activation
- Hidden Layer 2: 64 neurons, ReLU activation
- Output Layer: 13 neurons (grades 51 to 68, excluding gaps), softmax activation

We compiled the network with the Adam optimizer (learning rate 0.001), categorical cross-entropy loss, and trained for 97 epochs with batch size 32 and 10% validation split.

7) *Class Imbalance Handling*: To address severe class imbalance, we applied Synthetic Minority Over-sampling Technique (SMOTE) [12] before training. SMOTE generates synthetic samples for minority classes by interpolating between existing samples and their k -nearest neighbors. We adaptively set the number of neighbors k to $\min(5, n_{\min} - 1)$ where $n_{\min}$ is the smallest class size. SMOTE was applied exclusively to the training set to prevent data leakage.

We also added Gaussian noise with $\sigma = 0.01$ to standardized features as data augmentation, improving model robustness to input variations.

8) *SVM Baseline*: We trained a Support Vector Machine with radial basis function (RBF) kernel on identical features as a baseline comparison. We set SVM hyperparameters to default values ($C = 1.0$, $\gamma = \text{'scale'}$), with training completing in approximately 6 seconds.

D. Hybrid CNN Approach

1) *CNN Architecture*: We implemented a hybrid CNN architecture combining transfer learning with explicit features:

Image Branch:

- Input: $224 \times 224 \times 3$ RGB images
- Data Augmentation: Random horizontal/vertical flips, rotation ($\pm 20\%$), zoom ($\pm 10\%$), contrast ($\pm 20\%$), brightness ($\pm 20\%$), translation ($\pm 10\%$)
- Backbone: EfficientNetV2-B0 [5] pre-trained on ImageNet, initially frozen
- Global Average Pooling
- Dropout (0.2)
- Dense layer: 128 neurons, ReLU

Feature Branch:

- Input: Extracted features (same as ANN)
- Dense layer: 128 neurons, ReLU
- Batch Normalization
- Dense layer: 64 neurons, ReLU

Combined:

- Concatenation of image and feature branches
- Dense layer: 64 neurons, ReLU
- Batch Normalization
- Dropout (0.2)
- Output: 1 neuron, linear activation (regression)

We trained the model as a regression task predicting continuous grade values, then rounded to integers for classification metrics.

2) *Training Procedure*: Training proceeded in two phases:

Phase 1: Head Training (20 epochs):

- EfficientNetV2-B0 backbone frozen
- Adam optimizer, learning rate = 0.001
- Mean Squared Error (MSE) loss
- Early stopping (patience=10) and model checkpointing

Phase 2: Fine-tuning (20 epochs):

- Backbone unfrozen for end-to-end training
- Adam optimizer, learning rate = 10^{-5}
- Continue from Phase 1 weights

Total training time: approximately 2.5 minutes on GPU. Model size: 71 MB.

IV. RESULTS

A. Overall Performance Comparison

Table I presents the comparative performance of all three models across multiple metrics (70/30 stratified split, 537 test coins).

TABLE I
MODEL PERFORMANCE COMPARISON

Metric	ANN (5-fold CV)	Hybrid CNN	SVM
Exact Accuracy	86%	31%	30%
Accuracy (± 1)	88%	88%	31%
Accuracy (± 2)	94%	94%	38%
Accuracy (± 3)	98%	98%	43%
Weighted Precision	0.85	0.09	0.09
Weighted Recall	0.86	0.31	0.30
Weighted F1-Score	0.85	0.14	0.14
Training Time	4 min	2.5 min	6 sec
Inference Time (E2E)	~ 1.8 s	~ 80 ms	< 1 ms
Model Size	< 1 MB	71 MB	–

The feature-based ANN substantially outperforms both the hybrid CNN and SVM on exact accuracy and classification metrics. Notably, both the CNN and SVM exhibit nearly identical collapse behavior (weighted precision of 0.09, weighted recall ≈ 0.30), suggesting that neither transfer learning nor kernel methods alone can overcome severe class imbalance without explicit rebalancing strategies like SMOTE. While the CNN achieves comparable tolerance-based accuracy through regression averaging, its weighted F1-score of 0.14 indicates severe classification failures across individual grades.

B. Per-Grade Analysis

Table II details per-grade classification performance for the CNN, revealing systematic failures on minority classes.

TABLE II
CNN PER-GRADE CLASSIFICATION REPORT

Grade	Precision	Recall	F1	Support
50.0	0.00	0.00	0.00	1
55.0	0.00	0.00	0.00	2
57.0	0.00	0.00	0.00	3
58.0	0.00	0.00	0.00	12
60.0	0.00	0.00	0.00	3
61.0	0.00	0.00	0.00	12
62.0	0.00	0.00	0.00	44
63.0	0.00	0.00	0.00	108
64.0	0.31	1.00	0.47	165
65.0	0.00	0.00	0.00	127
66.0	0.00	0.00	0.00	52
67.0	0.00	0.00	0.00	7
68.0	0.00	0.00	0.00	1
Accuracy			0.31	537
Weighted Avg	0.09	0.31	0.14	537

The CNN shows catastrophic collapse to predicting grade 64 for nearly all inputs, achieving 100% recall for grade 64 but 0% recall for all other grades. This is a classic failure mode under class imbalance: the model learns to predict the majority class to minimize overall loss.

C. Tolerance-Based Accuracy Analysis

The ANN’s prediction distribution relative to professional grades shows concentration within the ± 3 tolerance band, confirming the 98% accuracy at this threshold. The ANN shows a slight conservative bias, tending to predict grades lower than professional services. This is actually a desirable property that protects buyers from overpaying.

D. Inference Time Breakdown

The ANN’s approximately 1.8 second inference time comprises:

- Image preprocessing: 200 ms
- Sobel edge detection: 400 ms
- Wedge masking and feature extraction: 800 ms
- HSV and brightness computation: 300 ms
- ANN forward pass: 50 ms
- Total: approximately 1.75 seconds

The CNN’s end-to-end inference time of ~ 80 ms includes image resizing/normalization and forward pass, both suitable for web deployment (< 2 s total with UI overhead). The CNN offers substantial speed for batch processing.

V. DISCUSSION

A. Why Feature Engineering Outperforms Deep Learning

We attribute the performance gap between the ANN and CNN/SVM to several factors in this problem domain:

Limited Training Data: With only 1,785 coins (approximately 1,248 in training), the dataset is not enough to train the millions of parameters in EfficientNetV2-B0. Deep CNNs

require 10,000+ samples to learn robust visual features from scratch, even with transfer learning from ImageNet [13].

Domain Shift from ImageNet: The visual characteristics of coin surfaces (metallic reflections, subtle wear patterns, fine engraving details) differ substantially from natural images in ImageNet. Transfer learning works best when source and target domains share visual similarities, which does not hold here.

Severe Class Imbalance: The CNN and SVM collapse to predicting grade 64 reflects optimization dynamics under imbalanced data. MSE loss (CNN) and default SVM behavior minimize error on the majority class while ignoring minorities. The ANN’s use of SMOTE-augmented data and categorical cross-entropy loss better addresses this challenge.

Domain Knowledge in Features: The hand-crafted features encode specific numismatic expertise. Wedge-based edge analysis isolates high-detail regions (Liberty’s face, Eagle’s feathers) where wear is most diagnostic. Color clustering captures patina patterns specific to gold coins. Brightness computation models human perception of luster. A CNN must rediscover these patterns from data alone, which is not feasible with limited samples.

B. Implications for Similar Applications

Our findings suggest guidelines for choosing between feature engineering and deep learning:

Favor Feature Engineering When:

- Training data is limited (less than 5,000 samples)
- Strong domain expertise is available
- Class imbalance is severe
- Model interpretability is important
- Highest accuracy is required

Favor Deep Learning When:

- Large datasets are available (more than 10,000 samples)
- Domain knowledge is limited or hard to formalize
- Real-time inference is critical
- Source and target domains are similar

Similar quality assessment applications in manufacturing inspection, medical imaging with rare conditions, and specialized artifact analysis may benefit from feature-based approaches despite the prevailing assumption that deep learning is always better.

C. Inference Speed vs. Accuracy Trade-off

The CNN’s speed advantage presents interesting deployment scenarios. For single coin grading, the ANN’s 1.8-second inference is acceptable for web services meeting the 10-second requirement. Batch processing 10,000 coins would require 5.5 hours with the ANN versus approximately 13 minutes with the CNN. A hybrid approach using the CNN for initial screening and the ANN for final grading could balance speed and accuracy.

D. Conservative Grading Bias

The ANN’s tendency to undergrade compared to professional services may reflect genuine model conservatism or

systematic bias in the training data. From a practical standpoint, conservative grading protects buyers from overpaying, builds user trust in the automated system, and may counteract potential grade inflation by professional services. Quantifying this bias against a held-out set graded by multiple services would provide valuable calibration data.

VI. LIMITATIONS

Several limitations constrain the generalizability of our findings:

Dataset Constraints:

- Limited to 1,785 coins from a single source (DLRC)
- Severe class imbalance with some grades having fewer than 5 samples
- Restricted to grades 50 to 68, lower grades excluded
- Single coin type (Saint-Gaudens Double Eagle)

Methodological Limitations:

- ANN results from 5-fold CV; CNN/SVM from single split (no cross-validation)
- Professional grades treated as ground truth despite known inter-grader variability
- No statistical significance testing between models
- Limited hyperparameter optimization

Technical Constraints:

- 2D images only, no 3D surface topology analysis
- Standardized imaging conditions may not reflect real-world user photos
- No comparison with human graders on identical test sets

VII. FUTURE WORK

Several directions could extend this research. Collecting 10,000+ coins with balanced grade distribution would test whether CNN performance improves with scale. Evaluating focal loss, class weights, and advanced augmentation techniques could improve CNN training under class imbalance. Combining CNN feature extraction with hand-crafted features in ensemble models is another promising direction. Testing pre-trained models on metal surface or numismatic datasets, if available, could reduce the domain shift problem.

Grading identical coins with multiple professional graders would establish inter-rater reliability benchmarks. Adding structured light scanning or photometric stereo would enable 3D surface topology features. Evaluating models on other gold coin types (Liberty Head, Indian Head) would test generalization. Identifying which extracted features contribute most to ANN accuracy through feature importance analysis would provide insights into the grading process. Parallelizing feature extraction could reduce inference latency for real-time optimization. Developing confidence intervals for grade predictions would enable uncertainty quantification.

VIII. CONCLUSION

We show that feature-based machine learning significantly outperforms deep learning for automated Saint-Gaudens coin grading when training data is limited. Our feature-based ANN,

using Sobel edge detection, wedge-based spatial analysis, HSV color clustering, and perceptually-weighted brightness features, achieves 98% accuracy within ± 3 grades on the Sheldon scale. This surpasses the CNN's and SVM's 31%/30% exact accuracy, which collapses to majority-class prediction due to severe class imbalance.

These results challenge the assumption that end-to-end deep learning always beats traditional feature engineering. In specialized domains with constrained data, encoding expert knowledge through hand-crafted features provides better performance, interpretability, and reliability. The 1.8-second inference time meets practical deployment requirements for web-based grading services.

Our findings have broader implications for quality assessment tasks in manufacturing, medicine, and cultural heritage preservation, where domain expertise is available but labeled data is scarce. Future work should investigate hybrid architectures that combine the strengths of learned and engineered features.

ACKNOWLEDGMENTS

This work was sponsored by the Bradley Department of Electrical and Computer Engineering at Virginia Tech. We thank Dr. Luke Lester for his guidance and support, Jianzhu Chen for technical assistance, Dr. Scot Ransbottom and Dr. Creed Jones for their subject matter expertise, and David Lawrence Rare Coins (DLRC) for providing access to the coin database.

AUTHOR CONTRIBUTIONS

Tanmay Dogra designed and implemented all machine learning models (ANN, CNN, SVM), performed statistical analysis, and wrote the manuscript. Eric Ngo developed the color clustering subsystem using K-Means and HSV analysis. Mohammad Alam developed the brightness classification subsystem and feature extraction methodology. Jean-Paul Talavera designed and implemented the web interface for the grading system. Asim Dahal contributed to quality assurance and model validation. All authors reviewed and approved the final manuscript.

REFERENCES

- [1] J. W. Dannreuther, *The Official Guide to Coin Grading and Counterfeit Detection*, Professional Coin Grading Service. New York: House of Collectibles, 2004.
- [2] W. H. Sheldon, *Early American Cents, 1793–1814: A Quantitative Scale for Condition*. New York: Harper & Brothers, 1949.
- [3] Santos et al., "Utilization of Convolutional Neural Networks in Coin Grading for Philippine Peso Coins," *Technologique: A Global Journal on Technological Developments and Scientific Innovations*, 2024.
- [4] S. Korchagin and T. Getahun, "Automatic Coin Grading: Model Based on Siamese Neural Network," *Semantic Scholar*, 2023.
- [5] M. Tan and Q. V. Le, "EfficientNetV2: Smaller Models and Faster Training," in *Proc. International Conference on Machine Learning (ICML)*, 2021, pp. 10096–10106.
- [6] "Computer-based Objective Interactive Numismatic System," *Ph.D. Dissertation*, Pace University, 2004.
- [7] "Sheldon Coin Grading Scale," Wikipedia. [Online]. Available: https://en.wikipedia.org/wiki/Sheldon_coin_grading_scale
- [8] "Grading St. Gaudens \$20 (1907-1933)," NGC. [Online]. Available: <https://www.ngccoin.com/coin-grading-guide/grading-st-gaudens-20/>

- [9] M. Blagus and L. Lusa, "SMOTE for High-Dimensional Class-Imbalanced Data," *BMC Bioinformatics*, vol. 14, no. 106, 2013.
- [10] "Sobel Edge Detector," University of Edinburgh, HIPR2. [Online]. Available: <https://homepages.inf.ed.ac.uk/rbf/HIPR2/sobel.htm>
- [11] "Interactive Color Image Segmentation using HSV Color Space," *Science and Technology Journal*, vol. 7, no. 1, 2019.
- [12] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.
- [13] F. Zhuang et al., "A Comprehensive Survey on Transfer Learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2021.