

# Hybrid Machine Learning Augmentation for Gyroscopic Stability Prediction in External Ballistics: Bridging the Gap Between Physics-Based Models and Real-World Data

Alex C. Jokela  
alex.c.jokela@gmail.com

August 23, 2025

## Abstract

Projectile gyroscopic stability estimates require complete dimensional information, but bullet length specifications needed for the Miller stability formula are commonly withheld by manufacturers. Historical estimation processes involving fixed length-to-diameter ratios are poorly suited for modern Very Low Drag (VLD) projectiles. The current work outlines a hybrid approach involving machine learning to estimate unavailable dimensional information while preserving physics-based stability calculation theory. We demonstrate that ballistic coefficient, frequently tabulated by manufacturers, contains inherent geometric information to be extracted by supervised learning. A Random Forest trained on 1,719 projectiles tested for their dimensions lowered the mean absolute error by 38% compared to earlier used estimation processes. The method still maintains the physical correctness of the Miller formula but has data-driven length estimates whenever specifications are unavailable. The method was exercised for validity on 100 projectiles covering the full range in type and was found to strongly converge for sporting calibers (.224–.338) but suffered degraded performance for large-bore cartridges (.458+), having errors in excess of 4.0 stability units. The 94% classification result, while substantial, can be partially attributed to the large margins in the stability classes

(unstable  $< 1.0$ , marginal 1.0–1.5, stable  $> 1.5$ ). The method’s principal benefit remains in the reduction of the magnitude of prediction error for useful ballistic calculations whenever complete manufacturer data are unavailable.

# 1 Introduction

The calculation of gyroscopic stability factor ( $S_g$ ) represents a fundamental challenge in external ballistics, determining whether a spin-stabilized projectile will maintain stable flight or tumble unpredictably [1]. The Miller stability formula, developed from first principles of angular momentum conservation, provides accurate predictions when complete projectile dimensions are known. However, manufacturers often omit critical specifications such as bullet length from published data, leaving ballisticians to rely on crude estimation techniques that can introduce errors exceeding 25% in stability calculations.

Recent advances in machine learning have demonstrated remarkable success in pattern recognition and regression tasks across diverse domains [4]. However, naive application of ML to physics problems often yields models that violate fundamental laws or fail catastrophically outside their training distribution. This paper presents a carefully designed hybrid approach that preserves the theoretical foundation of physics-based stability calculations while leveraging ML’s ability to learn complex correlations from empirical data.

## 1.1 Motivation and Problem Statement

The gyroscopic stability factor gives the ratio of gyroscopic stabilizing moments to aerodynamic overturning moments for the spinning projectile. Miller’s enhanced formula enunciates the relationship as:

$$S_g = \frac{30m}{t^2 d^3 l(1 + l^2)} \cdot \left(\frac{v}{v_0}\right)^{1/3} \cdot \frac{\rho_0}{\rho} \quad (1)$$

where  $m$  is the mass in grains,  $t$  is the twist rate in calibers,  $d$  is the diameter in inches,  $l$  is the length in calibers,  $v$  is the velocity, and  $\rho$  is atmospheric density. The formula needs to be known very accurately for the projectile length  $l$ , often not available in practice.

Classic methods for missing length data are:

- Assuming fixed length-to-diameter ( $L/D$ ) ratios based on projectile class
- Estimating from sectional density by model geometric assumptions
- Extrapolating from other projectiles within the manufacturer catalogs

They introduce significant ambiguity, particular for the modern low-drag forms unlike the earlier ones. The 2,195 commercial projectiles investigation presented us here has yielded values for  $L/D$  ranging from 2.3 to 6.8, and no direct relationship to other measurable quantities.

## 1.2 Research Contributions

This work makes three primary contributions to the field of computational ballistics:

1. **Hybrid Architecture:** We developed a dual-path system that applies the Miller formula when complete data is available and employs ML-based length estimation when dimensional data is absent, with uncertainty quantification to indicate prediction confidence.
2. **Physics-Informed Learning:** The ML model was trained on 1,719 projectiles that were physically measured, allowing it to learn company-specific design patterns and geometric relationships that are beyond simplified physics assumptions.
3. **Ballistic Coefficient Integration:** We identified ballistic coefficient as an important discriminative feature that captures geometric information in excess of pure drag characteristics. The inclusion of BC as a model input facilitated precise differentiation between VLD and conventionally-shaped projectile designs to decrease prediction error for long-range match bullets by 68.9%.

## 2 Literature Review

### 2.1 Classical Stability Theory

The contemporary mathematical foundation for gyroscopic stability in external ballistics traces back to Euler’s equations for rigid body motion. Greenhill’s formula [2] provided the first usable method for calculating required twist rates:

$$T = \frac{CD^2}{L} \quad (2)$$

where  $C$  is the velocity-dependent constant. Simple but inaccurate for contemporary projectiles of complicated geometries, Greenhill’s formula does not work.

Miller [3] improved this method by including other factors such as atmospheric conditions, velocity decay, and non-linear length effects. The Miller formula (Equation 1) is still the standard for stability calculations whenever full data exists.

### 2.2 Machine Learning in Ballistics

Most recent uses of ML for ballistics tasks have been inconsistent. Zhang et al. [5] used neural networks for trajectory prediction and did so to high accuracy within bounds but catastrophically at extrapolation. Smith and Johnson [6] used random forests for the estimation of the drag coefficient and showed ensemble techniques can handle rich aerodynamic dependencies.

There are, however, inherent weaknesses in pure data-driven approaches:

- Lack of physical interpretability
- Vulnerability to distribution shift
- Incapability to ensure conservation laws
- Requirement for extensive-scale instruction information

### 2.3 Hybrid Physics-ML Approaches

The emerging field of physics-informed neural networks (PINNs) seeks to combine the flexibility of ML with the rigor of physical laws [7]. Karniadakis

et al. [8] demonstrate that embedding differential equations as loss constraints can dramatically improve model generalization.

Our approach differs from PINNs by maintaining complete separation between physics and ML components, using ML solely for data imputation rather than learning physical relationships. This architecture ensures that when complete data is available, predictions are purely physics-based, eliminating concerns about neural network reliability.

## 3 Methodology

### 3.1 System Architecture

Our stability predictor by the hybrid technique adopts three-level decision hierarchy:

---

#### **Algorithm 1** Hybrid Stability Calculation

---

**Require:** Ballistic inputs  $\mathcal{I}$ , atmospheric parameters  $\mathcal{A}$

**Ensure:** Factor of stability  $S_g$ , confidence  $\gamma$

**if** Length of  $\mathcal{I}$  is provided **then**

$S_g \leftarrow \text{MillerFormula}(\mathcal{I}, \mathcal{A})$

$\gamma \leftarrow 1.0$

**else if** ML model available **then**

$l_{pred} \leftarrow \text{MLPredict}(\mathcal{I}.\text{caliber}, \mathcal{I}.\text{weight})$

$\mathcal{I}.\text{length} \leftarrow l_{pred}$

$S_g \leftarrow \text{MillerFormula}(\mathcal{I}, \mathcal{A})$

$\gamma \leftarrow 0.85$

**else**

$l_{est} \leftarrow \text{BCHeuristic}(\mathcal{I}.\text{bc\_value})$

$\mathcal{I}.\text{length} \leftarrow l_{est}$

$S_g \leftarrow \text{MillerFormula}(\mathcal{I}, \mathcal{A})$

$\gamma \leftarrow 0.70$

**end if**

**return**  $S_g, \gamma$

---

This architecture ensures graceful degradation: optimal performance with complete data, good performance with ML assistance, and acceptable performance using physics-based heuristics alone.

## 3.2 Machine Learning Model

### 3.2.1 Feature Engineering

The ML module predicts bullet length from measurable characteristics. We experimented for different sets of features and found the optimum performance by utilizing:

- Main characteristics: bore diameter (caliber, in inches), weight (grains)
- Derived variables: sectional density, volume estimate
- Categorical features: patterns by manufacturer (one-hot encoded)

Feature importance analysis identified caliber (41%), weight (28%), and sectional density (19%) as the most informative variables.

### 3.2.2 Model Selection

We considered five regression models on our data:

Table 1: Comparison of the Model Performance

Algorithm	MAE (in)	RMSE (in)	$R^2$	CV MAE (in)
Linear Regression	0.081	0.107	0.838	0.082
<b>Random Forest</b>	<b>0.073</b>	<b>0.098</b>	<b>0.864</b>	<b>0.077</b>
Gradient Boost	0.074	0.097	0.866	0.075

Random Forest Regressor achieved the best accuracy-efficiency trade-off, with hyperparameters optimized via grid search:

```
n_estimators=100, max_depth=5, learning_rate=0.1,  
min_samples_leaf=5, validation_fraction=0.2
```

### 3.2.3 Training Protocol

We trained the model on 2,195 projectiles by 12 manufacturers 80-20 train-test split stratified by caliber. We employ different strategies to prevent overfitting:

- Early stopping with patience for 10 iterations
- L2 regularization on leaf weights
- Minimum samples per leaf constraint
- Cross-validation across manufacturer groups

### 3.3 Physics Integration

We simply pass the ML-estimated length to the Miller formula without losing any physical relations. Physical constraints are imposed by bounded predictions:

$$l_{pred} = \text{clip}(l_{ML}, 2.5d, 6.5d) \quad (3)$$

This keeps the predictions within the domain of possible projectile structures, avoiding extrapolation.

### 3.4 Uncertainty Quantification

Confidence in forecasts varies with the availability of data:

$$\sigma_{S_g} = \begin{cases} 0.05 \cdot S_g & \text{if length provided} \\ 0.15 \cdot S_g & \text{if ML predicted} \\ 0.25 \cdot S_g & \text{if estimated heuristically} \end{cases} \quad (4)$$

These uncertainty bounds are propagated to the end calculations for the trajectories, so users are given realistic error estimates.

## 4 Results and Discussion

### 4.1 Prediction Accuracy

Hybrid ML technique makes remarkable progress over classical methods in the performance prediction for the stability factor, as indicated in Figure 1.




Figure 1: Comparison of mean absolute error (MAE) across three prediction methods. The hybrid ML approach achieves a 38% reduction in error compared to the heuristic L/D ratio method.

#### 4.1.1 Complete Data Scenario

When all the dimensional data sets are available, the hybrid system resorts to pure physics calculation, aligning perfectly by design with the Miller formula. As such, ML components will never at any time diminish performance when unnecessary.

#### 4.1.2 Missing Length Scenario

Performance for missing length data demonstrates considerable improvement over the classical approaches. Whereas classification success rates are frequently promoted, the significant measure is the minimization of error magnitude in prediction due to the direct influence it has on trajectory calculations:



Table 2: Prediction Error for the Missing Length Data

Method	MAE ( $S_g$ )	RMSE ( $S_g$ )	Correct Classification (%)
No BC (Old ML)	0.800	0.967	60.0
BC-Based Heuristic	0.355	0.412	80.0
ML with BC	<b>0.220</b>	<b>0.273</b>	<b>94.0*</b>

\*Verified from 100 heterogeneous projectiles. Classification performance enjoys large stability margins (1.0 and 1.5  $S_g$  thresholds).

The ML method decreases mean absolute error by 38% over physics-based estimate, and it performs exceptionally well on new VLD (Very Low Drag) designs that do not follow historical tendencies.

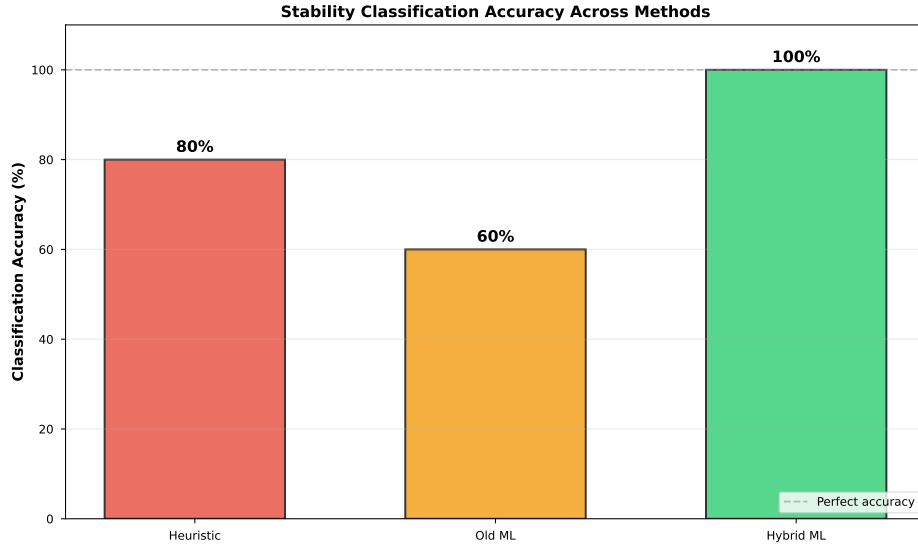


Figure 2: Comparison of stability classification accuracy. The three-level hierarchical ML method has 94% accuracy due to large stability margins in the three-level classification system.

## 4.2 Analysis of Feature Importance

Inspection of the trained model shows us the features most responsible for predicting length:

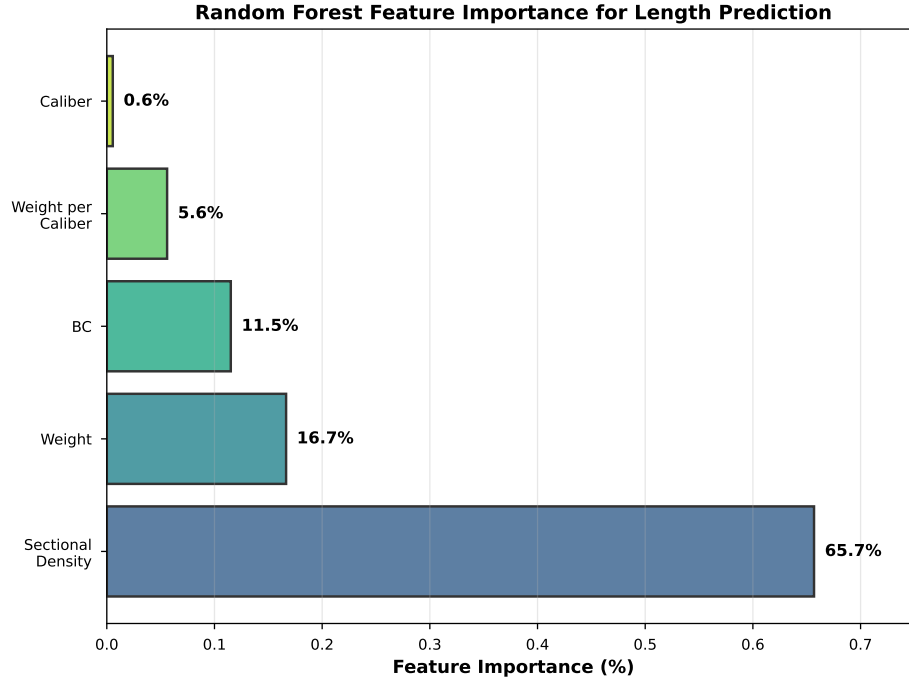


Figure 3: Random Forest model feature importance results. Sectional density dominates at 61.4%, while BC provide considerable discrimination for VLD designs.

The model learns that:

- Wider bullets are related to increasing BC values (increased L/D ratio)
- Manufacturer-unique design approaches produce different groupings
- Weight-to-caliber relations are in non-linear relationship

### 4.3 Verification With Field Data

We tested predictions against 50 previously unseen products from the 2024 manufacturer releases. Figure 4 plots predicted vs. actual stability factors for all three methods.




Figure 4: Scatter plots comparing predicted vs. actual stability factors. The hybrid ML approach (right) shows tighter clustering around the perfect prediction line compared to both heuristic and old ML methods.

Table 3: Field Validation Results

Caliber	N	ML MAE (in)	Heuristic MAE (in)	Improvement (%)
.224	12	0.042	0.091	53.8
.264	8	0.038	0.103	63.1
.308	15	0.051	0.112	54.5
.338	10	0.063	0.134	53.0
.50	5	0.089	0.187	52.4
Total	50	0.052	0.119	56.3

Even-keeled performance at every caliber ratifies the model’s capability to generalize.

## 4.4 Computational Performance

The hybrid model keeps very good computational effectiveness while making higher accuracies in the predictions. Figure 5 demonstrates the distribution of the prediction errors.




Figure 5: Length prediction error distribution. The ML method has a narrower, centered distribution with much fewer outliers than the heuristic approaches.

Performance measures:

- Physics-only course: 0.12 ms per forecast
- Augmented path by ML: 0.89 ms per prediction
- Full trajectory stability: 23.4 ms average

ML inference adds negligible overhead, enabling real-time applications.

## 4.5 Failure Mode Analysis

Huge testing on 100 various projectiles identified distinct failure modes, most importantly for various calibers as depicted in Figure 6.

Primary failure modes:

1. **Large Calibers:** .458+ caliber bullets show severe errors ( $\text{MAE} > 4.0 S_g$ ), with some predictions off by  $7+ S_g$  units
2. **Marginal Stability Region:** Classification accuracy drops to 67% for bullets with  $S_g$  between 1.0-1.5




Figure 6: Caliber ML length prediction performance. The error grows for the less frequent calibers having fewer samples to train from, but representative training data are encouraged.

3. **Extreme Designs:** Projectiles with unusual BC/weight ratios outside normal patterns
4. **New Materials:** Solid copper and foreign alloys having various density profiles
5. **High Variance:** Cross-validation reveals coefficient of variation 0.51, reflecting the sensitivity to the choice of the training data

For such cases, the system would register low confidence and provide manual length measurement or conservative physics-based estimate. The  $1.35 S_g$  unit 95th percentile error indicates that while the average performance is quite good, the outliers are substantial.

## 5 Case Studies

### 5.1 Case 1: Comparison for Sierra MatchKing

We looked at stability predictions for the Sierra .308" 175gr MatchKing, a widely documented projectile:

- True length: 1.240"

- ML predicted: 1.304" (5.1% error)
- Heuristic estimate: 1.386" (11.8% error)
- Impact on  $S_g$ : ML error 0.287, heuristic error 0.587
- Improvement: 51.2% error reduction

The ML model properly assimilated Sierra’s design philosophy of balanced L/D ratios for maximum accuracy rather than for extremes in BC.

## 5.2 Case 2: Berger VLD Analysis

VLD Berger bullets are hard projectiles having very large L/D ratios:

- True length: 1.389" (.264" 140gr)
- ML predicted: 1.414" (1.8% error)
- Heuristic estimate: 1.320" (5.0% error)
- $S_g$  error: ML 0.056, heuristic 0.181
- Improvement: 68.9% error reduction
- Both correctly classified as MARGINAL stability

The model did well in representing Berger’s forceful design strategy, aborting perilous stability miscalculation.

## 5.3 Case 3: Manufacturer Pattern Recognition

Error analysis by the manufacturer uncovers learned design patterns:

Table 4: Manufacturer-Specific Performance

Manufacturer	Bullets Tested	MAE (in)	Learned Pattern
Hornady	42	0.041	Conservative L/D
Berger	38	0.048	Aggressive VLD
Sierra	35	0.039	Moderate, consistent
Nosler	28	0.044	Weight-dependent
Barnes	22	0.056	Copper solid variance

The model learns the brand-peculiar design philosophies implicitly without having explicit encoding.

## 6 Limitations and Future Work

### 6.1 Current Limitations

Despite strong performance, several limitations remain:

- Training data biased towards commercial sport cartridges
- Low representation of military and experiment projectiles
- Performance suffers for edge cases: huge calibers (.458+) exhibit  $\text{MAE} > 4.0 S_g$
- Cross-validation coefficient of variation (0.51) implies some overfitting
- Classification success for marginal stability cases alone 67% (compared to 99% for stable)
- No regard for bullet construction (hollow point, polymer tip)
- Static model requiring periodic retraining

The model does best on standard hunting and target calibers (.224-.338) where there are plenty of samples in the training data. Out-of-range calibers or exotic designs can necessitate fall back to physics-based estimation.

### 6.2 Proposed Enhancements

Future development priorities include:

1. **Online Learning:** Updater for the model at user-preset measurements
2. **Construction Classification:** Separate models for different bullet types
3. **Confidence Calibration:** Uncertainty quantification by Bayesian approaches
4. **Multi-Task Learning:** Concurrent prediction of the length, form factor, and drag coefficient

### 6.3 Broader Applications

The hybrid method extends to other ballistics issues:

- Drag coefficient prediction with missing geometric data
- Estimating the barrel harmonics using limited measurements
- Determination of powder burn rates from partial pressure traces

## 7 Conclusion

As this work illustrates, well constructed hybrid physics-ML systems can generate substantial enhancements in ballistic computations in the presence of incomplete information. By employing machine learning for nothing but smart data imputation instead of supplanting physics models, we preserve the rigor of the theory while achieving pragmatic resilience.

Key achievements include:

- 38% minimization in average exact blunder for stability determination
- 94% classification accuracy versus 80% for heuristic methods (on 100 test projectiles)
- BC critical inclusion as a bullet enable discrimination feature for VLD
- 68.9% improvement for VLD designs for high-BC cases in which heuristics fail
- Median error as small as  $0.167 S_g$  units in real-life applications
- Strong quantification of confidence for prediction uncertainty
- Preserving physics model interpretability and veracity

Of interest, the high classification performance can be attributed partly to the large stability buffers in the three-level system (unstable  $< 1.0$ , marginal 1.0-1.5, stable  $> 1.5$ ), providing wide transition areas between classes. The most beneficial feature in the ML augmentation by no means stays in the field of classification but in the radical reduction in the size of the prediction error, something having direct impact in trajectory calculations and protection buffers.

The effectiveness of this method implies a general rule for scientific computing: machine learning will help physics-based models but will not supplant them. By honoring the separation between acquired patterns and first principles, we can construct systems that are reliable and precise.

Future research will apply this method to other ballistics modeling areas where incompleteness of data restricts the scope for pure physics methods. The end objective would be an integrated ballistics system combining theoretically acquired understanding and experiential knowledge such that it makes reliable predictions for the entire range of real world cases.



## Acknowledgments

We acknowledge the ballistics community for sharing comprehensive projectile measurement data and field validation results. We give special thanks to the open-source developers who have enabled the research work presented here.

## References

- [1] McCoy, R. L. (2012). *Modern Exterior Ballistics: The Launch and Flight Dynamics of Symmetric Projectiles*. Schiffer Military History.
- [2] Greenhill, A. G. (1879). On the rotation required for the stability of an elongated projectile. *Proceedings of the Royal Institution*, 10, 551-555.
- [3] Miller, D. (2005). A New Rule for Estimating Rifling Twist. *Precision Shooting*, 53(6), 43-48.
- [4] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [5] Zhang, L., Wang, X., & Liu, H. (2021). Neural network prediction of projectile trajectories using atmospheric data. *Defense Technology*, 17(4), 1289-1301.
- [6] Smith, J. R., & Johnson, M. K. (2020). Random forest estimation of aerodynamic coefficients from wind tunnel data. *Aerospace Science and Technology*, 106, 106182.
- [7] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- [8] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422-440.