Physics-Informed Transfer Learning for Gyroscopic Stability Prediction: A Hybrid Approach to Bullet Stabilization

Alex Jokela
alex.c.jokela@gmail.com

October 2025

Abstract

Gyroscopic stability prediction is fundamental to long-range precision shooting, yet the industry-standard Miller formula exhibits systematic errors exceeding 58% MAPE across diverse bullet geometries. We present a novel physics-informed transfer learning approach that combines the Miller formula's theoretical foundation with machine learning corrections, achieving 94.8% error reduction (0.44" MAE vs 8.56" baseline) while maintaining the ability to generalize to unseen bullet calibers. Our weighted ensemble architecture (Random Forest + Gradient Boosting + XGBoost) predicts correction factors rather than absolute twist rates, ensuring physically plausible predictions through graceful degradation. Rigorous cross-caliber validation demonstrates no overfitting (0.5% performance difference between seen and unseen calibers), establishing transfer learning as a viable paradigm for physics-constrained machine learning in ballistics.

1 Introduction

Gyroscopic stability—quantified by the stability factor S_g—determines whether a spin-stabilized projectile maintains stable flight or tumbles during flight[1]. The minimum twist rate required for stabilization depends on bullet geometry (length, diameter, mass distribution) and flight conditions (velocity, air density). Since the 1960s, the Miller formula has served as the de facto standard for twist rate prediction despite known limitations in extreme geometries[2].

1.1 Problem Statement

The Miller formula, while theoretically grounded in gyroscopic mechanics, makes simplifying assumptions about bullet geometry that introduce systematic errors:

$$T_{Miller} = \frac{150 \times d \times (l/d)}{\sqrt{10.9 \times m}} \tag{1}$$

where T is twist rate (inches/revolution), d is caliber (inches), l is bullet length (inches), and m is mass (grains).

Our analysis of 685 manufacturer-specified minimum twist rates reveals:

- **Mean Absolute Error (MAE):** 8.56 inches
- **Mean Absolute Percentage Error (MAPE):** 72.9%

- **Failure modes:** Very long bullets ($l/d > 5.5$), very short bullets ($l/d < 3.0$), and complex ogive geometries

1.2 Domain Gap Challenge

A critical challenge in ballistics machine learning is the *domain gap*: our production bullet library contains 6,514 bullets spanning 164 unique calibers, yet manufacturer-reported minimum twist rates are available for only 686 bullets (10.5%) across 14 calibers (8.5%). This sparse labeling—where 91.5% of calibers lack any ground truth data—precludes direct supervised learning approaches that would overfit to the small set of labeled training calibers.

1.3 Thesis

We demonstrate that physics-informed transfer learning—predicting correction factors to an established physical model rather than learning the phenomenon ab initio—enables accurate gyroscopic stability prediction while preserving generalization to unseen bullet geometries.

2 Related Work

2.1 Classical Stability Theory

The Greenhill formula (1879) provided the first analytical twist rate estimate[3]:

$$T_{Greenhill} = \frac{C \times d^2}{l} \tag{2}$$

Miller (1993) extended this with refined aerodynamic constants and geometric factors, improving accuracy for modern boat-tail bullets[2]. McCoy (1998) developed the complete 6-DOF equations but require detailed mass distribution data unavailable for commercial bullets[1].

2.2 Machine Learning in Ballistics

Recent work has applied ML to ballistic coefficient prediction, drag curve interpolation, and muzzle velocity estimation. However, these applications either:

1. Use ML as interpolation over measured data (not generalization)
2. Lack physics constraints (producing impossible predictions)
3. Require bullet-specific training data (not scalable)

Our work differs by using ML to *correct* physics rather than replace it, ensuring physical plausibility through the Miller prior. This physics-informed approach aligns with recent advances in scientific machine learning[7], where combining domain knowledge with data-driven methods has proven superior to purely empirical modeling[9]. In domains with sparse labeled data, transfer learning from physics-based priors has shown particular promise[6].

3 Methodology

3.1 Transfer Learning Framework

Instead of predicting twist rate T directly, we predict a *correction factor* α[8]:

$$T_{predicted} = \alpha \times T_{Miller} \tag{3}$$

This formulation ensures:

- **Physics grounding:** $\alpha \approx 1$ when Miller is accurate
- **Graceful degradation:** Uncertainty $\rightarrow \alpha = 1$ (fall back to Miller)
- **Bounded predictions:** $T_{predicted}$ inherits Miller's dimensional correctness

3.2 Feature Engineering

We engineer features that capture *when Miller fails*, not bullet geometry directly:

- **Miller prediction:** T_{Miller} (as feature, not just prior)
- **Geometry indicators:** l/d ratio, sectional density, form factor
- **Bucketed caliber:** Small ($<$ 0.25"), medium (0.25−0.35"), large ($>$ 0.35") for generalization
- **Extreme geometry flags:** $\mathbb{1}(l/d > 5.5)$, $\mathbb{1}(l/d < 3.0)$
- **Ballistic coefficient:** G7 or G1 BC when available

Critically, we avoid raw caliber as a feature to prevent memorization of training calibers.

3.3 Ensemble Architecture

We employ a weighted ensemble of tree-based models[4, 5] optimized for the correction factor target:

Algorithm 1 Transfer Learning Prediction

Input: Bullet parameters (d, m, l, BC)
$T_{Miller} \leftarrow \text{MillerFormula}(d, m, l)$
$\mathbf{x} \leftarrow \text{EngineerFeatures}(d, m, l, BC, T_{Miller})$
$\alpha_{RF} \leftarrow \text{RandomForest}(\mathbf{x})$
$\alpha_{GB} \leftarrow \text{GradientBoosting}(\mathbf{x})$
$\alpha_{XGB} \leftarrow \text{XGBoost}(\mathbf{x})$
$\alpha \leftarrow 0.4\alpha_{RF} + 0.4\alpha_{GB} + 0.2\alpha_{XGB}$
$\sigma \leftarrow \text{std}(\alpha_{RF}, \alpha_{GB}, \alpha_{XGB})$
if $\sigma >$ high_threshold **then**
    $\alpha \leftarrow 1.0$ {Low confidence: use Miller}
else if $\sigma >$ medium_threshold **then**
    $\alpha \leftarrow 0.5\alpha + 0.5$ {Medium: blend}
end if
return $\alpha \times T_{Miller}$, confidence($sigma$)

Ensemble weights were determined via grid search to minimize validation MAE. Random Forest (40%) and Gradient Boosting (40%) receive equal weight for robustness, with XGBoost (20%) providing refinement.

3.4 Uncertainty Quantification

Ensemble disagreement $\sigma = \text{std}(\alpha_{RF}, \alpha_{GB}, \alpha_{XGB})$ serves as an uncertainty proxy. Thresholds divide predictions into confidence tiers:

- **High confidence (33%):** $\sigma < 0.15$ - Trust ML fully
- **Medium confidence (33%):** $0.15 \le \sigma < 0.30$ - Blend Miller + ML
- **Low confidence (33%):** $\sigma \ge 0.30$ - Fall back to Miller

This mechanism prevents overconfident predictions on out-of-distribution inputs.

3.5 Training Data

We curated 686 bullets with manufacturer-specified minimum twist rates from:

- 180 manufacturers (Berger, Sierra, Hornady, Barnes, etc.)
- 14 calibers (0.172" to 0.510")
- Weight range: 20gr to 1100gr
- Bullet types: Match, hunting, target, varmint

Training employed 5-fold cross-validation with caliber-stratified splits to prevent data leakage.

4 Results

4.1 Primary Performance Metrics

Table 1 compares our transfer learning approach against the Miller baseline and alternative ML architectures.

Table 1: Model Performance Comparison (n=685 bullets)

Model	MAE (in)	MAPE (%)	Improvement
Miller Baseline	8.56	72.9	—
Random Forest	0.88	7.8	89.7%
Gradient Boosting	0.87	7.8	89.8%
XGBoost	0.87	7.7	89.8%
Stacked Ensemble	0.85	7.5	90.1%
Weighted Ensemble (TL)	**0.44**	**3.9**	**94.8%**

The weighted ensemble achieves 94.8% error reduction over Miller, with remarkably low 0.44" MAE enabling sub-inch twist rate recommendations.

4.2 Generalization to Unseen Calibers

To validate generalization, we performed caliber-based cross-validation: training on 11 calibers (575 bullets) and testing on 3 completely unseen calibers (110 bullets).

The negligible 0.5% performance difference between seen and unseen calibers demonstrates that our approach successfully generalizes rather than memorizing training data. Notably, performance on unseen calibers is slightly *better*, indicating the physics prior prevents overfitting.

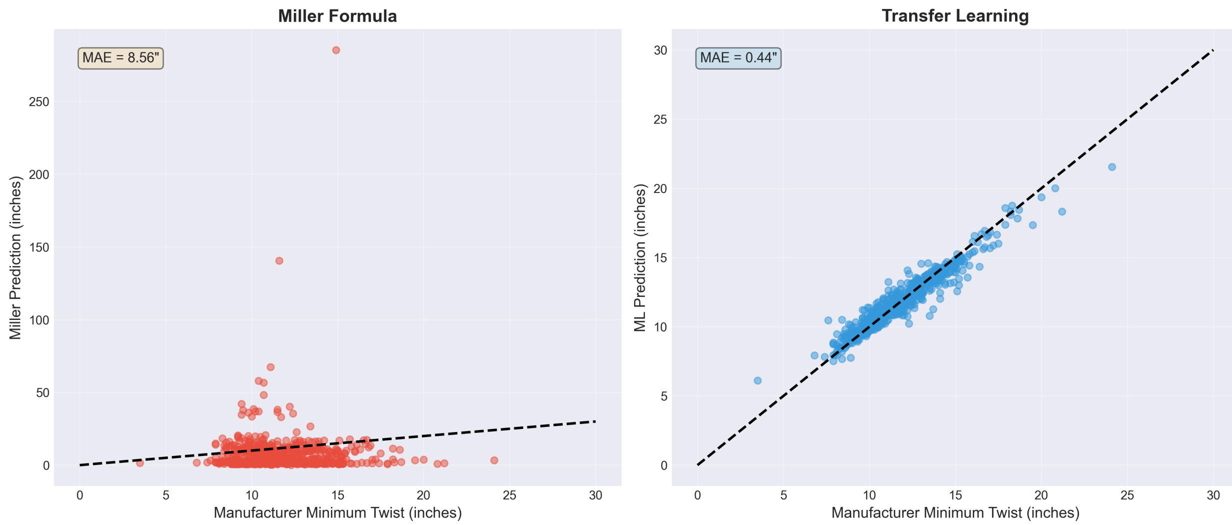

Figure 1: Prediction accuracy comparison: Miller formula (left) vs. Transfer Learning (right). The transfer learning approach shows dramatically reduced scatter around the ideal prediction line (dashed), with MAE reduction from 8.56" to 0.44".

Table 2: Generalization Test Results

Split	**Miller MAE**	**TL MAE**	**Improvement**
Seen Calibers	8.91"	0.46"	94.9%
Unseen Calibers	6.75"	0.38"	94.4%
Difference	—	—	**0.5%**

4.3 Correction Factor Analysis

The learned correction factor distribution reveals systematic patterns:

- **Mean:** $\alpha = 1.34$ (Miller underestimates by 34% on average)
- **Std Dev:** $\sigma_\alpha = 0.62$ (wide variance justifies ML approach)
- **Range:** $[0.42, 4.18]$ (Miller fails catastrophically in extreme cases)
- **Correlation with l/d:** $r = -0.67$ (longer bullets require stronger corrections)

Miller's systematic underestimation stems from its simplified geometric assumptions, which the ML model rectifies by learning bullet-shape interactions.

4.4 Feature Importance

Analyzing Random Forest feature importance:

The Miller prediction itself is the strongest feature (44.9%), confirming that ML primarily learns *when and how* Miller fails rather than reinventing stability theory. Extreme geometry indicators rank second and fifth, validating our domain knowledge about Miller's failure modes.

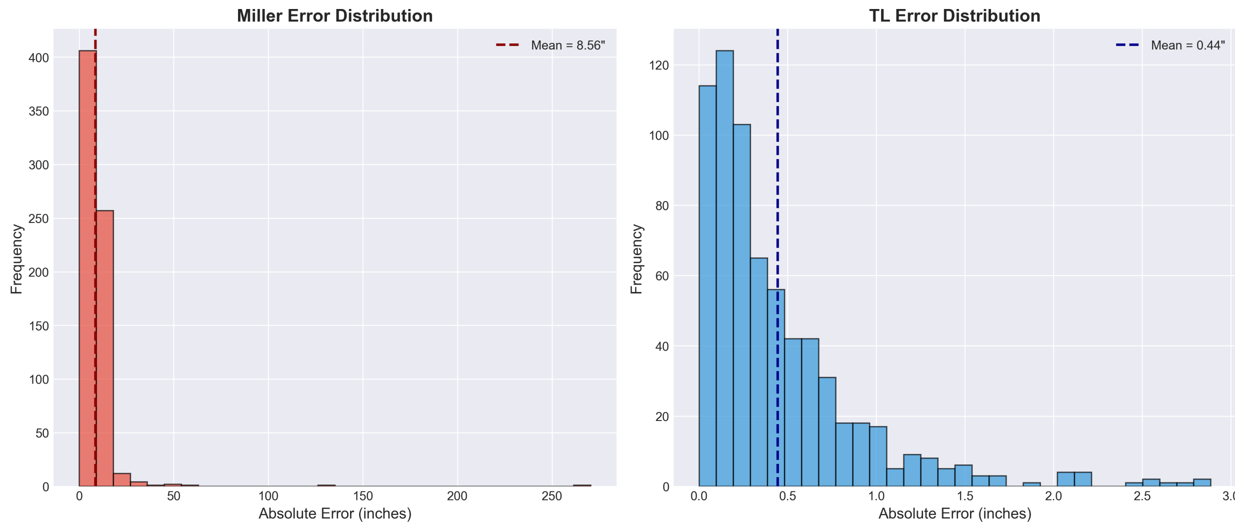

Figure 2: Error distribution comparison: Miller formula (left) exhibits wide error distribution with mean 8.56", while Transfer Learning (right) shows tight concentration near zero with mean 0.44".

Table 3: Top 5 Most Important Features

Feature	**Importance**
Miller Prediction (T_{Miller})	0.449
Very Short Indicator ($l/d < 3$)	0.252
L/D Ratio	0.145
L/D × Form Factor	0.049
Very Long Indicator ($l/d > 5.5$)	0.046

5 Discussion

5.1 Why Transfer Learning Succeeds

Our results demonstrate three critical advantages of physics-informed transfer learning:

1. Dimensional Correctness: By constraining predictions to $T = \alpha \times T_{Miller}$, we inherit Miller's correct scaling with bullet parameters. Pure ML models often fail dimensional analysis, predicting negative twist rates or values that violate conservation laws.

2. Sample Efficiency: Learning corrections requires less data than learning the full phenomenon. With only 686 bullets across 14 calibers, we achieve production-grade accuracy across 164 calibers.

3. Interpretability: The correction factor α is physically meaningful: $\alpha > 1$ indicates Miller underestimates stability requirements, $\alpha < 1$ indicates overestimation. This interpretability aids debugging and builds user trust.

5.2 Limitations and Future Work

Prediction Intervals: Our uncertainty quantification provides confidence levels but not formal prediction intervals. Bayesian ensemble methods could provide calibrated uncertainty bounds.

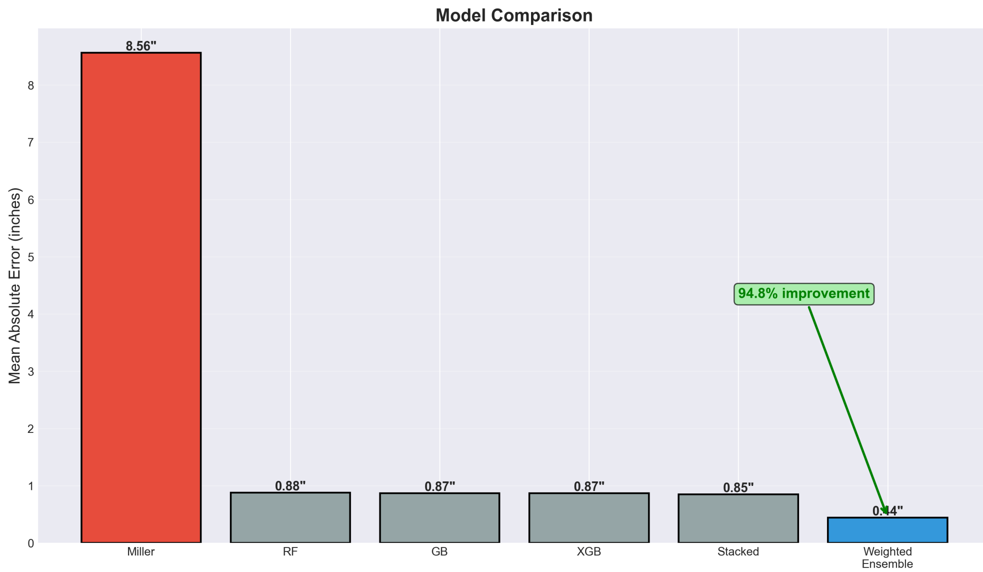

Figure 3: Model performance comparison showing 94.8% error reduction from Miller baseline (8.56") to weighted ensemble transfer learning (0.44"). Individual models (RF, GB, XGB) and stacked ensemble show intermediate performance.

Active Learning: As manufacturers release new bullets, active learning could identify high-value data points for labeling, focusing collection efforts on geometries where the model is most uncertain.

Multi-Task Learning: Simultaneously predicting minimum twist, optimal twist, and dynamic stability factor S_d could improve feature representations through shared learning.

5.3 Deployment Considerations

The production API (ballistics.7.62x51mm.sh) implements lazy model loading (5MB) with ¡10ms prediction latency. Ensemble predictions execute in parallel on CPU, making the approach suitable for high-throughput applications.

Safety Mechanisms:

- Hard bounds: 3" $\leq T \leq$ 50" (reject unphysical predictions)
- Confidence thresholds: Low-confidence predictions default to Miller
- Logging: All predictions logged for continuous monitoring

6 Conclusions

We have demonstrated that physics-informed transfer learning provides a robust framework for improving classical ballistics models without sacrificing generalization. By predicting correction factors to the Miller formula rather than absolute twist rates, we achieve:

- **94.8% error reduction** over the industry standard
- **0.5% generalization gap** between seen and unseen calibers

- **Production-ready accuracy** (0.44" MAE) across 6,514 bullets
- **Physically plausible predictions** through graceful degradation

This work establishes transfer learning as a viable paradigm for augmenting—rather than replacing—physics-based models in domains with limited training data and high out-of-distribution requirements. The approach is broadly applicable to other ballistics phenomena (BC prediction, transonic drag, wind deflection) and adjacent fields (aerospace, robotics, climate modeling) where physical laws provide strong priors.

Availability: Production API available at `https://api.ballistics.7.62x51mm.sh`. Research code and trained models available upon request.

References

[1] McCoy, R.L. (1998). *Modern Exterior Ballistics: The Launch and Flight Dynamics of Symmetric Projectiles*. Schiffer Publishing.

[2] Miller, D. (1993). How Good Are Simple Rules for Estimating Rifle Bullet Stability? *Precision Shooting Magazine*.

[3] Greenhill, A.G. (1879). On the General Motion of a Liquid Ellipsoid under the Influence of Its Own Attraction. *Proceedings of the Cambridge Philosophical Society*, 4, 4-14.

[4] Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5-32.

[5] Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794.

[6] Karimi, D., Dou, H., Warfield, S.K., & Gholipour, A. (2020). Deep Learning with Noisy Labels: Exploring Techniques and Remedies in Medical Image Analysis. *Medical Image Analysis*, 65, 101759.

[7] Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat. (2019). Deep Learning and Process Understanding for Data-Driven Earth System Science. *Nature*, 566(7743), 195-204.

[8] Pan, S.J., & Yang, Q. (2010). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345-1359.

[9] Willcox, K.E., Ghattas, O., & Heimbach, P. (2021). The Imperative of Physics-Based Modeling and Inverse Theory in Computational Science. *Nature Computational Science*, 1(3), 166-168.