

Physics-Informed Machine Learning for Predicting Rutting Resistance of Asphalt Mixtures Using Aggregate Shape Parameters

Bakiman Sihotang

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

bakiman.sihotang@uab.edu

Eliana Cafarchia

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

elianac@binghamton.edu

Abstract

Rutting remains one of the most consequential distress modes in asphalt pavements, affecting safety, ride quality, drainage, and life-cycle cost. Conventional empirical and mechanistic models can predict rutting resistance, but they often struggle to represent the coupled influence of aggregate morphology, binder rheology, mixture gradation, compaction, traffic loading, climate, and construction variability. This paper examines physics-informed machine learning as a system-level strategy for predicting rutting resistance of asphalt mixtures using aggregate shape parameters. The central argument is that aggregate shape is not merely a descriptive input but a meso-structural control variable that conditions contact fabric, interlock, stress distribution, and permanent deformation accumulation. Physics-informed machine learning offers a way to embed mechanical constraints, viscoelastic behavior, conservation principles, and domain knowledge into data-driven models, thereby improving extrapolation, uncertainty awareness, and engineering credibility. The paper discusses the architecture of such systems, including imaging and characterization pipelines, laboratory and field data integration, model training and validation, deployment in pavement management, and governance arrangements. It also considers sustainability, equity, robustness, and policy implications, especially where predictive tools influence material specifications, procurement, and maintenance prioritization. The analysis concludes that physics-informed machine learning can support more resilient and sustainable asphalt infrastructure, but only if technical performance is matched by transparent validation, institutional accountability, and careful attention to the socio-technical conditions of deployment.

Keywords

physics-informed machine learning; asphalt mixtures; rutting resistance; aggregate shape; pavement engineering; infrastructure governance; sustainability.

1. Introduction

Rutting in asphalt pavements is a cumulative permanent deformation phenomenon that emerges from the interaction of material composition, environmental exposure, traffic loading, and structural support conditions. Its consequences extend beyond surface irregularity, because rutting can trap water, accelerate moisture damage, increase hydroplaning risk, raise rolling resistance, and shorten rehabilitation cycles. For these reasons, rutting resistance has

become a central criterion in asphalt mixture design, quality assurance, and pavement management. At the same time, the properties that govern rutting resistance are difficult to represent with simple empirical relationships. Aggregate shape, gradation, binder stiffness, aging, compaction energy, and field conditions all interact in ways that are nonlinear, scale-dependent, and sensitive to local variability. This complexity has encouraged the use of machine learning, but purely data-driven approaches can produce models that fit historical observations while failing to respect mechanical principles or extrapolate reliably to new materials and climates.

Physics-informed machine learning has emerged as a response to this tension between predictive flexibility and physical credibility [1]. By embedding governing constraints, symmetries, conservation laws, constitutive relationships, and domain knowledge into learning architectures, physics-informed models seek to combine the adaptability of data-driven methods with the interpretability and robustness of mechanics-based modeling [2]. In civil and environmental systems, such integration is especially valuable because data are often sparse, noisy, heterogeneous, and generated under conditions that do not cover the full operational envelope [3]. Pavement engineering is a clear example. A model trained only on laboratory rutting tests may not capture field aging, moisture gradients, or traffic wander. A model trained only on field performance data may lack the resolution needed to distinguish the influence of aggregate angularity from binder modification or compaction quality.

Aggregate shape parameters are particularly important in this context. Coarse aggregate morphology affects particle packing, contact area, interlock, and resistance to rearrangement under load [4]. Surface texture and angularity influence friction and bonding with asphalt mastic, while flat and elongated particles can create preferred orientation and localized stress concentrations [5]. These morphological features also influence dynamic modulus and permanent deformation behavior, although the magnitude and direction of their effects depend on gradation, binder properties, and loading conditions [6]. Imaging and computational approaches have improved the ability to quantify aggregate shape and simulate mixture meso-structure, but translating those measurements into reliable rutting predictions remains challenging [7]. Recent image-recognition research has further demonstrated how coarse aggregate morphology can be linked to asphalt mixture meso-structure, providing new opportunities for data-rich characterization [8].

The practical value of predicting rutting resistance is not limited to laboratory curiosity. Rutting predictions inform asphalt mix design, binder selection, aggregate sourcing, quality control, warranty provisions, and maintenance scheduling. Standard rutting tests, such as wheel-tracking methods, provide essential performance indicators, but they are expensive and time-consuming and cannot cover every material combination or climate scenario [9]. Mechanistic-empirical design frameworks offer a broader systems perspective by linking material response, structural response, and performance over time, yet they still depend on calibration factors and simplified assumptions [10]. A physics-informed machine learning approach can complement these established tools by learning from available data while respecting the physical mechanisms that make rutting resistance meaningful. The remainder of this paper develops this argument by examining aggregate shape parameters, the physics-informed modeling paradigm, system architecture, validation and robustness, deployment and governance, sustainability and equity, and future research directions.

2. Aggregate Shape Parameters and Rutting Mechanisms

Aggregate shape parameters describe geometric and surface features that influence how particles pack, interlock, and transmit stress. Angularity, sphericity, roundness, surface texture, form, and flatness or elongation are commonly used descriptors. These parameters are not independent. A particle may be angular but relatively equant, or rounded but elongated; each combination produces different contact mechanics and mixture behavior. In asphalt mixtures, coarse aggregates form a skeleton that carries a significant portion of load, especially in stone mastic and open-graded mixtures. When the skeleton is composed of angular and textured particles, interlock can increase resistance to permanent deformation. When particles are rounded or smooth, the skeleton may rearrange more easily under repeated loading, particularly at high temperatures and low loading frequencies [4]. The meso-structural arrangement of particles, mastic, and air voids therefore becomes a critical link between aggregate shape and rutting resistance.

Surface texture and angularity also affect the binder-aggregate interface. A rougher surface can provide greater mechanical adhesion and increase the effective contact area between aggregate and mastic. This can improve resistance to shear displacement, but the benefit depends on binder properties, coating thickness, and moisture conditions. If the mastic is too soft or if moisture weakens adhesion, the advantages of texture may be diminished. Similarly, flat and elongated particles can increase interlock in some gradations but also create anisotropic packing and stress concentrations that promote dilation, particle rearrangement, or fracture. These trade-offs explain why aggregate shape parameters must be considered within a system that includes binder rheology, volumetric properties, compaction, and environmental aging [5]. Binder modification can improve stiffness and elasticity, but it cannot fully compensate for a poorly interlocked aggregate skeleton [11]. Rheological properties of modified binders, including polymer-modified systems, influence rutting resistance by controlling deformation recovery and viscous flow [12].

Rutting resistance is also shaped by moisture damage and aging. Moisture can reduce adhesion between aggregate and mastic, leading to stripping, loss of cohesion, and accelerated permanent deformation [17]. Aging increases binder stiffness over time, which may improve rutting resistance at moderate temperatures but can also increase brittleness and cracking susceptibility. These coupled processes mean that a predictive model must represent time, temperature, moisture, and loading frequency, not just aggregate shape. Nevertheless, aggregate shape remains a high-value input because it is measurable, increasingly automatable through image recognition, and directly tied to mixture meso-structure [8]. Computational studies of particle surface characteristics using optical and X-ray computed tomography have shown that shape descriptors can be quantified with high resolution [5]. Virtual testing approaches that generate random aggregate structures have also made it possible to simulate how morphological variability affects mechanical response [7]. These developments provide a foundation for physics-informed models that treat aggregate shape as a first-class variable rather than a background nuisance.

3. Physics-Informed Machine Learning as a Modeling Paradigm

Physics-informed machine learning refers to a family of methods that integrate scientific knowledge into data-driven models. The integration can occur through loss functions that penalize violations of governing equations, architecture design that enforces known symmetries, constraints on parameters, hybrid models that couple mechanistic submodels with learned corrections, or Bayesian priors that encode physical expectations [1]. In the case of rutting prediction, relevant physics includes viscoelasticity, viscoplasticity, conservation of

mass and momentum, contact mechanics, damage evolution, and moisture transport. A purely data-driven model may learn correlations between aggregate shape and rutting without representing why those correlations exist. A physics-informed model can use that knowledge to regularize learning, especially when data are limited or biased toward particular materials and conditions [2]. This is important because pavement data are rarely balanced across aggregate types, climates, traffic levels, and maintenance histories.

One advantage of physics-informed machine learning is improved extrapolation. If a model learns that stiffer binders reduce permanent deformation at high temperature, or that angular particles increase interlock, it may generalize better to unseen mixtures than a model that simply memorizes training examples. Another advantage is uncertainty awareness. Physics-informed models can be designed to distinguish between aleatoric uncertainty, which arises from natural variability, and epistemic uncertainty, which arises from limited knowledge. This distinction is valuable for infrastructure decisions because it helps engineers and policymakers understand whether a prediction is uncertain because of inherent material randomness or because the model is operating outside its training domain [3]. In pavement engineering, such distinctions can influence how conservative a design should be, how much testing is required, and how maintenance budgets should be allocated.

Hybrid modeling is particularly promising for rutting resistance. A mechanistic component can represent the viscoelastic response of the binder and mastic, while a machine learning component can learn corrections related to aggregate morphology, compaction, or aging that are difficult to capture analytically. Neural network models have been used to predict dynamic modulus and rutting potential from mix properties, demonstrating the feasibility of data-driven pavement prediction [13], [14]. Soft computing and intelligent systems have also been applied across pavement engineering, including distress prediction and material characterization [15]. However, these earlier models often lacked explicit physical constraints. Physics-informed approaches can build on them by incorporating constitutive relationships and domain knowledge, thereby reducing the risk of physically impossible predictions. For example, a model should not predict that a mixture becomes less resistant to rutting as aggregate angularity increases when all other variables are held constant, unless there is a specified mechanism such as particle fracture or excessive void content. Embedding such monotonic or mechanistic expectations can improve trust and support regulatory acceptance.

4. System Architecture and Data Infrastructure

A physics-informed machine learning system for rutting prediction is not merely a model; it is a socio-technical architecture that includes data acquisition, preprocessing, feature engineering, physics encoding, model training, validation, deployment, and feedback. The first layer is aggregate characterization. Imaging systems, laser scanners, X-ray computed tomography, and computer vision pipelines can measure particle shape, angularity, texture, and size distribution [5], [8]. These measurements must be linked to mixture design data, including gradation, binder grade, binder content, air voids, compaction method, and additive type. The second layer is laboratory performance data. Wheel-tracking tests, repeated load tests, and dynamic modulus tests provide controlled observations of rutting resistance and viscoelastic response [9]. The third layer is field performance data, including rut depth, traffic loading, climate, maintenance history, and subsurface conditions. Field data are essential for validating whether laboratory-based predictions hold under real-world conditions [10].

The integration of these layers raises significant infrastructure challenges. Data may be stored in incompatible formats, collected under different standards, or governed by proprietary

constraints. Aggregate imaging data can be large, requiring storage and computational resources. Field data may be sparse and irregular, with missing values and inconsistent measurement intervals. A successful system therefore requires a data governance framework that addresses provenance, quality, privacy, security, and access. It also requires interoperability standards so that models can be updated and audited across agencies, laboratories, and contractors. In this sense, physics-informed machine learning is as much an institutional project as a technical one. The value of the model depends on the reliability of the data supply chain and the willingness of stakeholders to share and maintain data over time.

Deployment architecture must also account for different user needs. A mix designer may need rapid predictions for alternative aggregate sources. A state transportation agency may need network-level forecasts for maintenance planning. A contractor may need quality control feedback during production. These use cases require different levels of granularity, latency, and explainability. A cloud-based platform can support large-scale training and periodic updates, while edge computing can support real-time quality control at the plant or paving site. Digital twin representations of pavements can integrate physics-informed predictions with sensor data, inspection records, and traffic models. However, such integration increases dependence on cyber-physical infrastructure and raises questions about failure modes, cybersecurity, and liability. If a model recommends a mix design that underperforms, responsibility must be clearly assigned among material suppliers, designers, agencies, and software providers. Governance arrangements must therefore be designed alongside the model itself.

5. Model Training, Validation, and Robustness

Training a physics-informed model for rutting resistance requires careful attention to data splitting, feature selection, and validation strategy. Random splits can overestimate performance because data from the same project, aggregate source, or climate zone may appear in both training and testing sets. Grouped or spatial validation is often more appropriate. Temporal validation is also important because aging and traffic loading change over time. A model trained on early-life performance may not predict long-term rutting accurately. Cross-validation should therefore reflect the operational conditions in which the model will be used. Physics-informed constraints can help, but they do not eliminate the need for rigorous empirical validation [2]. The model must be tested against independent laboratory tests and field sections, including mixtures with aggregate shapes, binder grades, and climates that were not represented in training.

Robustness is another critical dimension. Pavement materials are variable, and small changes in gradation, moisture content, or compaction can alter performance. A robust model should not produce dramatically different predictions for minor input perturbations unless those perturbations correspond to real physical thresholds. Sensitivity analysis can reveal whether the model is relying on spurious correlations. Adversarial testing can examine how the model responds to extreme but plausible inputs, such as highly elongated particles, very high air voids, or aged binder. Uncertainty quantification can communicate confidence intervals to decision-makers. In safety-critical infrastructure applications, a model that is accurate on average but unreliable in edge cases may be less useful than a simpler model with known limitations. Physics-informed machine learning can improve robustness by constraining predictions to physically plausible regions, but it must be paired with transparent documentation of assumptions and failure modes [3].

Transfer learning and continual learning are also relevant. Aggregate sources, construction practices, and climate conditions vary across regions. A model trained in one region may not transfer directly to another. Transfer learning can adapt a pre-trained model using a smaller amount of local data, while continual learning can update the model as new performance data arrive. However, continual learning introduces risks of catastrophic forgetting and drift. If the model is updated too aggressively, it may lose previously learned physical relationships. If it is updated too slowly, it may fail to reflect changing materials or climate. A governance process is needed to decide when updates are warranted, how they are validated, and how version changes are communicated to users. These are not purely technical questions; they involve institutional trust, regulatory acceptance, and professional responsibility.

6. Deployment, Governance, and Policy Implications

The deployment of physics-informed machine learning in pavement engineering intersects with standards, procurement, and public accountability. Rutting resistance is often specified through standard tests and performance criteria [9]. If a machine learning model is used to supplement or replace some testing, agencies must define how predictions are validated and how errors are handled. A model that predicts low rutting risk but is later contradicted by field performance could lead to premature failure, safety concerns, and costly remediation. Therefore, deployment should be staged, beginning with decision support rather than autonomous decision-making. The model can rank alternatives, flag high-risk mixtures, or optimize testing schedules, while human engineers retain authority over final decisions. This human-in-the-loop approach aligns with broader principles of responsible artificial intelligence in infrastructure.

Policy implications extend to data sharing and transparency. Public agencies fund much of the data collection for pavement performance, and there is a strong argument that such data should be accessible for research and validation, subject to privacy and security safeguards. Proprietary aggregate imaging algorithms and contractor data may complicate openness, but standards for model documentation, versioning, and auditability can help. Model cards, datasheets, and validation reports can make assumptions explicit. Independent review panels can assess whether a model is fit for regulatory use. Procurement rules can require vendors to provide explainability, uncertainty estimates, and performance guarantees. These measures are not obstacles to innovation; they are conditions for trustworthy adoption. Without them, physics-informed machine learning may remain confined to pilot projects or be used in ways that are difficult to defend publicly.

Governance also involves equity and distributional effects. Pavement condition affects access to jobs, healthcare, education, and emergency services. If predictive models prioritize maintenance in ways that systematically disadvantage rural, low-income, or minority communities, they can reinforce existing inequities. Conversely, well-designed models can support more equitable allocation by identifying sections with high social and economic consequences of failure. This requires data on traffic, population, accessibility, and environmental burden, not just material properties. It also requires participatory processes that allow affected communities to understand how maintenance decisions are made. Physics-informed machine learning cannot resolve these normative questions on its own, but it can make trade-offs more visible if it is designed with transparency and public accountability in mind.

7. Sustainability, Equity, and Life-Cycle Considerations

Sustainability in asphalt pavements involves reducing greenhouse gas emissions, conserving virgin aggregates, increasing the use of recycled materials, and extending service life. Rutting resistance interacts with all of these goals. Mixtures that resist rutting may require higher-quality aggregates, polymer-modified binders, or more energy-intensive production. Conversely, premature rutting leads to frequent rehabilitation, which consumes additional materials, energy, and funds. A physics-informed model can help evaluate these trade-offs by predicting how different aggregate shapes and recycled material contents affect performance. For example, recycled asphalt pavement and recycled asphalt shingles can alter binder stiffness and aggregate gradation, potentially affecting rutting resistance [12]. If models can quantify these effects reliably, agencies can optimize for both durability and environmental performance.

However, sustainability assessment must consider the full life cycle. A mixture that performs well in the laboratory but requires rare aggregate sources transported over long distances may have a larger carbon footprint than a locally sourced mixture with slightly lower rutting resistance. Similarly, a model that recommends more frequent testing or more conservative designs may improve reliability but increase costs and emissions. Life-cycle assessment can provide a broader perspective, but it requires data on material extraction, production, transportation, construction, maintenance, and end-of-life. Integrating such data with physics-informed rutting prediction is challenging but valuable. It moves the discussion from single-parameter optimization toward multi-objective decision-making that balances performance, cost, environmental impact, and social equity.

Equity considerations also arise in maintenance prioritization. Pavements in wealthier areas may receive more attention because of political influence or higher traffic volumes. A data-driven model trained on historical maintenance patterns may reproduce those biases. To avoid this, model development should include fairness audits and consider distributive justice. For example, predictions could be evaluated by region, income level, and demographic composition to ensure that high-risk sections are not systematically overlooked. This does not mean that all sections should receive equal treatment; engineering risk and traffic loading matter. It means that the criteria for prioritization should be explicit and justifiable. Physics-informed machine learning can support such transparency by linking predictions to physical mechanisms and documented assumptions, but only if governance structures are designed to ask those questions.

8. Limitations and Future Research Directions

Despite its promise, physics-informed machine learning for rutting prediction faces several limitations. First, the physics of asphalt mixtures is extraordinarily complex. Viscoelasticity, viscoplasticity, damage, moisture diffusion, aging, and particle rearrangement interact across multiple scales. No model can capture all of these processes perfectly. Simplifications are necessary, but they introduce model error that must be acknowledged. Second, data are often scarce and uneven. High-quality aggregate imaging data may exist for some projects but not others. Field performance data may be limited to certain climate zones or traffic levels. This limits the ability to train and validate complex models. Third, explainability remains a challenge. Even physics-informed models can be difficult to interpret when they combine many features and nonlinear transformations. Engineers and regulators need to understand why a prediction was made, not just that it is accurate on average [3].

Future research should focus on multi-scale integration, uncertainty quantification, and human-centered design. Multi-scale models could link particle-level shape descriptors to

mixture-level performance and pavement-level response. Uncertainty quantification could provide probabilistic predictions that are more useful for risk-based decision-making. Human-centered design could ensure that models fit existing workflows and decision cultures rather than forcing users to adapt to opaque systems. Collaborative data infrastructures, shared benchmarks, and open validation datasets could accelerate progress. There is also a need for longitudinal studies that track mixtures from construction through many years of service. Such studies would provide the empirical grounding needed to assess whether physics-informed models truly improve long-term predictions. Finally, research should examine the institutional conditions under which these models are trusted, contested, and revised. Technical accuracy alone will not guarantee adoption if governance and accountability are weak.

9. Conclusion

Physics-informed machine learning offers a promising path for predicting rutting resistance of asphalt mixtures using aggregate shape parameters. By combining data-driven learning with mechanical constraints and domain knowledge, such models can improve extrapolation, robustness, and engineering credibility. Aggregate shape is a crucial input because it influences meso-structure, interlock, stress distribution, and permanent deformation. However, aggregate shape alone is not sufficient. Binder rheology, gradation, aging, moisture, compaction, and traffic loading must be integrated into a system-level framework. The architecture of such a framework includes imaging and characterization, laboratory and field data integration, physics-informed model training, rigorous validation, deployment in decision support, and governance that addresses transparency, liability, and equity. Sustainability and life-cycle considerations add further dimensions, requiring trade-offs among performance, cost, emissions, and social impact. The conclusion is not that physics-informed machine learning will replace traditional engineering judgment, but that it can augment that judgment when designed and governed responsibly. Its value will depend less on algorithmic novelty than on data quality, validation discipline, institutional trust, and a commitment to public accountability. If these conditions are met, physics-informed machine learning can contribute to more resilient, sustainable, and equitable asphalt infrastructure.

References

1. 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.
2. 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.
3. Willard, J., Jia, X., Xu, S., Steinbach, M., & Kumar, V. (2022). Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4), 1–37.
4. Pan, T., Tutumluer, E., & Carpenter, S. H. (2006). Effect of coarse aggregate morphology on permanent deformation behavior of hot mix asphalt. *Journal of Transportation Engineering*, 132(7), 580–589.
5. Masad, E., Saadeh, S., Al-Rousan, T., Garboczi, E., & Little, D. (2007). Computations of particle surface characteristics using optical and X-ray CT images. *Computational Materials Science*, 38(3), 495–509.

6. Wang, L., Wang, H., Haberl, A., & Oeser, M. (2015). Effect of aggregate morphology on the dynamic modulus of asphalt mixtures. *Construction and Building Materials*, 93, 719–727.
7. Chen, J., Li, H., Wang, L., & Wu, S. (2018). Virtual testing of asphalt mixture with two-dimensional and three-dimensional random aggregate structures. *International Journal of Pavement Engineering*, 19(10), 915–927.
8. L.W.Shi, Z.Y.Hu, H.M.Huang, et al. Effect of coarse aggregate morphology on asphalt mixture meso-structure using image recognition technology, *Measurement*, 2026, 290: 122757.
9. AASHTO. (2020). Standard method of test for determining the rutting susceptibility of asphalt paving mixtures using the Hamburg wheel-tracking device (AASHTO T 324). American Association of State Highway and Transportation Officials.
10. National Cooperative Highway Research Program. (2004). Guide for mechanistic-empirical design of new and rehabilitated pavement structures (NCHRP 1-37A). Transportation Research Board.
11. Bahia, H. U., Hanson, D. I., Zeng, M., Zhai, H., Khatri, M. A., & Anderson, R. M. (2001). Characterization of modified asphalt binders in Superpave mix design (NCHRP Report 459). Transportation Research Board.
12. Airey, G. D. (2003). Rheological properties of styrene butadiene styrene polymer modified road bitumens. *Fuel*, 82(14), 1709–1719.
13. Tarefder, R. A., & Zaman, M. (2010). Neural network modeling of rutting potential of asphalt pavements. *Journal of Materials in Civil Engineering*, 22(11), 1104–1112.
14. Kutay, M. E., Arambula, E., Gibson, N., & Youtcheff, J. (2011). Use of artificial neural networks to predict dynamic modulus of asphalt mixtures. *Transportation Research Record*, 2210(1), 93–100.
15. Gopalakrishnan, K., Ceylan, H., & Attoh-Okine, N. (2017). Intelligent and soft computing systems for pavement engineering. Springer.
16. Kim, Y. R., & Little, D. N. (2004). Linear viscoelastic analysis of asphalt mastics. *Journal of Materials in Civil Engineering*, 16(2), 122–132.
17. Little, D. N., & Jones, D. R. (2003). Chemical and mechanical processes of moisture damage in hot-mix asphalt pavements. *Transportation Research Record*, 1829(1), 20–28.
18. Molenaar, A. A. A. (2000). Rutting of asphalt pavements. In *Proceedings of the 2nd Eurasphalt & Eurobitume Congress* (pp. 1–10). European Asphalt Pavement Association.
19. Wang, H., Huang, X., & Chen, J. (2020). Review of aggregate shape characterization and its application in asphalt mixtures. *Construction and Building Materials*, 237, 117600.
20. Zhang, Y., & Luo, R. (2020). An overview of machine learning in asphalt pavement engineering. *Journal of Traffic and Transportation Engineering (English Edition)*, 7(4), 491–507.