

AI-Driven Optimization of Recycled Aggregate Gradation for Sustainable Asphalt Mixture Design

Ira A. Shroff

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
iramail@colostate.edu

Abstract

This paper examines artificial intelligence driven optimization of recycled aggregate gradation as a system-level strategy for sustainable asphalt mixture design. Recycled aggregates, including reclaimed asphalt pavement and recycled concrete aggregates, offer substantial environmental and economic benefits, yet their morphological variability, absorption characteristics, and heterogeneous surface textures complicate conventional volumetric design. Artificial intelligence can support gradation optimization by integrating aggregate imaging, mixture performance data, construction records, and lifecycle assessment into adaptive decision architectures. The analysis considers not only predictive accuracy but also structural trade-offs among rutting resistance, fatigue durability, moisture susceptibility, workability, and long-term pavement resilience. It argues that the technical optimization problem is inseparable from governance, data infrastructure, procurement policy, and fairness across suppliers and agencies. A central claim is that AI-driven gradation design should be understood as a socio-technical infrastructure rather than a standalone computational tool. The paper reviews relevant advances in machine learning, computer vision, pavement performance modeling, and infrastructure governance. It then discusses architecture choices, deployment challenges, robustness requirements, and policy implications for public agencies. Cross-domain comparisons with supply chain optimization, commons governance, and case-study evaluation are used to clarify how optimization models can remain accountable under uncertainty. The conclusion emphasizes that sustainable asphalt mixture design requires hybrid intelligence in which data-driven optimization is bounded by material science, local conditions, institutional capacity, and transparent validation. Such an approach can improve recycled aggregate utilization while protecting pavement performance and public trust.

Keywords

recycled aggregate; asphalt mixture design; artificial intelligence; gradation optimization; sustainable infrastructure; pavement performance; data governance; socio-technical systems.

1. Introduction

Asphalt pavements form the circulatory system of modern mobility, but their material demands impose significant environmental burdens through aggregate extraction, bitumen production, transport, and construction. The use of recycled solid waste materials in asphalt pavements has therefore become a major area of research and practice, motivated by landfill diversion, resource conservation, and cost reduction [1]. Life-cycle assessment of pavements has shown that material selection and mixture design decisions can meaningfully alter energy use, greenhouse gas emissions, and resource depletion, although the magnitude depends on transport distances, plant efficiency, maintenance cycles, and functional equivalence [2]. Sustainable pavement systems require integrated consideration of materials, structural design, construction quality, use phase, and end-of-life recovery rather than isolated component

improvements [3]. Within this broader context, recycled aggregate gradation is a particularly consequential design variable because it influences the skeleton of the mixture, the distribution of voids, binder demand, compactability, and the mechanisms by which distress develops. Traditional mixture design methods, including those codified in established handbooks and specifications, provide essential guidance but often struggle to capture the combinatorial complexity introduced by recycled materials [4].

Artificial intelligence has emerged as a potential means of managing that complexity. Machine learning can identify nonlinear relationships among aggregate properties, gradation parameters, binder content, compaction effort, and laboratory or field performance. Computer vision can characterize aggregate shape, angularity, texture, and size distribution with increasing automation. Optimization algorithms can search large design spaces for mixtures that satisfy multiple performance and sustainability objectives. Yet the application of AI to recycled aggregate gradation cannot be treated as a purely mathematical exercise. The optimization problem is embedded in infrastructure systems that include quarries, recycling facilities, asphalt plants, testing laboratories, transportation agencies, contractors, and communities. It also involves trade-offs that are technical, economic, environmental, and political. A mixture that maximizes recycled content may reduce virgin material consumption but increase moisture sensitivity or variability. A gradation that improves rutting resistance may reduce fatigue tolerance. An AI model that performs well on historical data may fail when feedstock changes, climate conditions shift, or maintenance regimes evolve.

This paper argues that AI-driven optimization of recycled aggregate gradation should be designed as a socio-technical capability rather than a singular software intervention. The technical architecture matters, but so do data provenance, model validation, procurement rules, supplier incentives, and mechanisms for contesting automated recommendations. The discussion is organized around several themes. First, recycled aggregate gradation is framed as a complex systems problem involving material heterogeneity and institutional fragmentation. Second, AI architectures for gradation optimization are examined in relation to sensing, representation, learning, and multi-objective search. Third, the paper analyzes structural trade-offs in mixture design and lifecycle sustainability. Fourth, governance and deployment are considered through the lens of data infrastructure and commons management. Fifth, robustness, fairness, and policy implications are explored, especially for public agencies responsible for long-lived assets. Finally, case illustrations and cross-domain comparisons are used to identify forward-looking design principles.

2. System Context: Recycled Aggregate Gradation as a Socio-Technical Problem

Recycled aggregates differ from virgin aggregates in ways that matter for gradation optimization. Reclaimed asphalt pavement contains aged binder, while recycled concrete aggregate may include adhered mortar, cement paste, and higher porosity. These materials can have lower specific gravity, higher water absorption, and more irregular shape compared with many virgin aggregates. Their properties vary by source, processing method, stockpile management, and seasonal conditions. The resulting heterogeneity means that a gradation curve that performs well with one recycled feedstock may not perform well with another. Conventional volumetric design often compensates through binder adjustment, blending rules, or conservative limits, but these responses can reduce the environmental benefit of recycling or increase cost. AI offers the possibility of adaptive blending and gradation selection that responds to measured feedstock characteristics rather than relying solely on prescriptive envelopes.

Aggregate morphology is central to this adaptive logic. Recent image recognition research on coarse aggregate morphology shows that shape, angularity, and surface texture influence the meso-structure of asphalt mixtures and therefore affect mechanical response [5]. Image analysis of asphalt concrete has long demonstrated that internal structure, including aggregate orientation, contact patterns, and void distribution, is not fully captured by bulk volumetric properties [6]. This creates an opportunity for AI systems that combine imaging data with mixture performance data to recommend gradations that produce desirable meso-structural arrangements. However, it also creates a representational challenge. A neural network may learn statistical associations between image features and performance without understanding the physical mechanisms that produce those associations. For infrastructure decisions, such models must be interpreted alongside mixture design principles and established construction knowledge [7].

The institutional context further complicates optimization. Public agencies rely on specifications, test methods, and quality assurance procedures to ensure that mixtures meet minimum performance requirements. Superpave volumetric requirements, for example, provide a structured framework for balancing binder content, air voids, voids in mineral aggregate, and voids filled with asphalt [8]. These requirements evolved to manage risk across diverse materials and climates. Introducing AI-driven gradation optimization into this environment means engaging with specifications that are not merely technical but also legal and administrative. A model recommendation may need to be translated into a submittal, verified through laboratory testing, and accepted by an agency engineer. If the model cannot explain why a gradation is expected to work, it may be rejected even if it is statistically promising. Thus, deployment depends on boundary objects that allow data scientists, material engineers, contractors, and regulators to coordinate.

The system also has temporal and spatial dimensions. Pavement performance unfolds over years and is influenced by traffic loading, climate, drainage, subgrade conditions, and maintenance history. Modeling asphalt concrete behavior requires attention to viscoelasticity, damage accumulation, and environmental aging [9]. Recycled aggregates interact with binder chemistry and modification, affecting rheology, adhesion, and moisture damage resistance [10]. An AI optimization system that ignores these interactions may overfit to short-term laboratory results. A system that attempts to model everything may become too complex to validate. The practical challenge is to design an optimization architecture that is sensitive to the most consequential variables while remaining robust to missing data, measurement error, and local variation. This is not a reason to abandon AI, but it is a reason to treat AI as one component within a larger infrastructure of evidence.

3. AI Architectures for Gradation Optimization

AI architectures for recycled aggregate gradation can be understood in terms of sensing, representation, prediction, and search. Sensing involves capturing aggregate properties through mechanical testing, image analysis, laser scanning, or plant sensors. Representation involves converting those measurements into features that machine learning models can use. Prediction involves estimating performance outcomes such as stiffness, rutting depth, fatigue life, moisture susceptibility, or workability. Search involves identifying gradations and blending proportions that satisfy multiple objectives under constraints. Each layer introduces choices that affect the system's behavior. For example, a convolutional neural network can learn features directly from aggregate images, reducing reliance on handcrafted descriptors. Deep learning has transformed image recognition and other domains by enabling hierarchical

feature learning from large datasets [11]. General deep learning principles emphasize the importance of data volume, architecture design, regularization, and validation [12]. Machine learning more broadly offers tools for prediction, clustering, and optimization across scientific and engineering domains [13].

In pavement engineering, deep convolutional neural networks have been used with transfer learning for distress detection, demonstrating the value of visual data for infrastructure assessment [14]. Neural network applications in pavement engineering have also been surveyed across prediction, classification, and optimization tasks [15]. These advances suggest that AI can support gradation optimization not only by predicting performance but also by automating quality control. For instance, a plant-mounted imaging system could estimate aggregate gradation and shape in real time, allowing adjustments before a mixture leaves the facility. A laboratory-based system could combine image-derived morphology with compaction and performance tests to recommend blend proportions. A cloud-based platform could aggregate data across projects and learn from regional differences. However, these architectures differ in latency, cost, reliability, and governance. Real-time plant control requires low latency and high robustness. Long-term mixture optimization can tolerate more computation but demands stronger validation and traceability.

A critical architectural question is whether optimization should be centralized or federated. Centralized models can benefit from large pooled datasets and consistent preprocessing, but they may struggle with local material conditions and create dependencies on a single vendor or agency. Federated learning could allow plants and agencies to train shared models without moving raw data, preserving confidentiality and local autonomy. Hybrid architectures might combine a global prior with local fine-tuning. Such designs align with infrastructure studies that emphasize the importance of installed base, local adaptation, and boundary-spanning work [17]. They also raise questions about model ownership, update rights, and accountability. If a model recommends a gradation that later underperforms, who is responsible? The supplier, the agency, the software vendor, or the engineer who accepted the recommendation? These questions must be addressed before deployment, not after failure.

Another architectural dimension is multi-objective optimization. Recycled aggregate gradation typically involves competing objectives: maximize recycled content, minimize cost, maintain workability, resist moisture damage, achieve adequate rutting resistance, and satisfy fatigue requirements. A single scalar objective can hide value judgments. A Pareto-based approach can present trade-off frontiers rather than one optimal solution. Decision-makers can then choose a mixture based on policy priorities, risk tolerance, and local experience. This approach also creates a role for human judgment. Instead of automating decisions, AI can structure them by making trade-offs explicit. This is particularly important in public infrastructure, where environmental benefits, user costs, and agency budgets are distributed unevenly across stakeholders. The architecture should therefore support transparency, scenario analysis, and sensitivity reporting.

4. Trade-offs in Mixture Design, Performance, and Lifecycle Sustainability

The most important trade-offs in recycled aggregate gradation involve performance, durability, and sustainability. Increasing recycled aggregate content can reduce virgin material extraction and landfill disposal, but it can also increase variability, absorption, and binder demand. A gradation that includes a high proportion of recycled material may have a different void structure and aggregate interlock than a conventional mixture. If the recycled aggregates are angular and well-graded, they may improve rutting resistance through better stone-on-

stone contact. If they are flaky, elongated, or coated with aged binder, they may reduce workability and increase compaction difficulty. These effects are not uniform across climates or traffic levels. A mixture designed for a low-volume road may tolerate higher recycled content than a mixture designed for a high-volume interstate subjected to heavy truck loads and extreme temperatures.

Lifecycle sustainability adds another layer. A mixture with higher recycled content may have lower upstream impacts, but if it requires more binder, additional processing, or more frequent maintenance, the net benefit may be reduced. Life-cycle assessment of pavements has shown that use phase and maintenance assumptions can dominate results, making it necessary to define functional units carefully [2]. Sustainable pavement systems require balancing environmental, economic, and social performance across the entire lifecycle [3]. AI optimization can help by evaluating many scenarios, but it cannot resolve normative questions about how to weigh different impacts. For example, should a design that reduces carbon emissions but increases cost be preferred? Should a design that extends pavement life but requires more energy-intensive recycling be considered sustainable? These choices require public deliberation and policy guidance.

There is also a structural trade-off between optimization and robustness. A model that is finely tuned to historical data may recommend gradations that perform well under past conditions but fail when feedstock or climate changes. Robust optimization seeks solutions that perform acceptably across a range of plausible conditions rather than optimally under one prediction. This is especially relevant for recycled aggregates because their properties can shift with source, season, and processing. Robustness can be improved through conservative constraints, ensemble modeling, uncertainty quantification, and regular retraining. However, each of these measures adds cost and complexity. Agencies must decide how much uncertainty is acceptable and how to monitor performance over time. The optimization system should therefore include feedback loops that compare predicted and observed performance, allowing models to be updated as evidence accumulates.

Workability and constructability are often underemphasized in optimization research but are critical in practice. A gradation that looks excellent in a laboratory may segregate during transport, compact poorly, or require excessive roller passes. These outcomes affect density, permeability, and long-term durability. AI models trained only on laboratory data may miss plant and field constraints. Including construction data in the training process can improve realism, but it also introduces noise and confounding factors. The system must distinguish between mixture properties and construction quality. If a pavement fails because it was under-compacted, an AI model might incorrectly attribute the failure to gradation. Causal inference, experimental design, and field validation are therefore essential complements to predictive modeling. The goal is not merely to predict but to support decisions that are technically sound and operationally feasible.

5. Governance, Data Infrastructure, and Deployment

Governance is a central concern for AI-driven gradation optimization because the system depends on data that are often fragmented, proprietary, or inconsistent. Aggregate producers may not want to share stockpile data. Contractors may view performance data as commercially sensitive. Agencies may lack the capacity to manage complex data pipelines. These conditions resemble common-pool resource problems, where collective action is needed to sustain a shared resource [16]. In this case, the shared resource is not only the data but also the trust required for collaboration. If participants believe that data will be used

punitively, they may withhold information or manipulate measurements. Governance arrangements must therefore balance transparency with confidentiality, and accountability with learning.

Infrastructure studies emphasize that large information systems succeed or fail through their relationship with existing practices, standards, and institutions [17]. An AI platform for gradation optimization cannot simply be installed; it must be integrated into procurement, testing, reporting, and quality assurance workflows. This requires attention to data standards, metadata, versioning, and interoperability. The data revolution has made vast quantities of information available, but it has also revealed the importance of data infrastructures that make information usable and trustworthy [18]. For asphalt mixtures, essential metadata include aggregate source, processing method, moisture condition, binder grade, additive type, mixing temperature, compaction method, and testing protocol. Without such context, machine learning models may learn spurious correlations. With it, they can support meaningful comparisons and continuous improvement.

Deployment should be incremental and layered. A first stage might involve decision support, where AI recommendations are reviewed by engineers but not automatically implemented. A second stage might involve plant-level quality control, where models flag deviations from target gradation. A third stage might involve adaptive optimization, where the system adjusts blend proportions within pre-approved limits. Each stage requires different validation, cybersecurity, and liability arrangements. Public agencies should define performance thresholds, audit trails, and human override procedures. They should also require model documentation, including training data scope, known limitations, and uncertainty estimates. These requirements are not bureaucratic obstacles; they are mechanisms for maintaining public accountability in infrastructure decisions that affect safety and mobility.

Procurement policy can accelerate or inhibit adoption. If contracts reward lowest initial cost, contractors may have little incentive to invest in AI or higher recycled content. If contracts include lifecycle performance incentives, pay factors for sustainability, or shared savings from reduced virgin material use, behavior may change. Standards can also create demand for better data. For example, requiring digital gradation records or image-based aggregate characterization can generate the data needed for AI models. However, mandates must be accompanied by training, technical assistance, and realistic transition periods. Small producers and local agencies may lack the resources to comply immediately. A just transition requires capacity building rather than simply imposing new requirements on the weakest participants.

6. Robustness, Fairness, and Policy Implications

Robustness in AI-driven gradation optimization has technical, organizational, and political dimensions. Technically, models must handle missing data, sensor drift, outliers, and distributional shift. Organizationally, they must remain useful when key personnel leave, vendors change, or budgets fluctuate. Politically, they must withstand scrutiny from stakeholders who may question automated recommendations. These dimensions interact. A technically robust model that is not trusted will not be used. A trusted model that is technically fragile may produce poor decisions. Robustness therefore requires redundancy, transparency, and institutional learning. It also requires humility: the system should know when it does not know and should defer to human judgment or additional testing when uncertainty is high.

Fairness is a less obvious but important consideration. AI models trained on data from well-resourced agencies and large contractors may perform poorly for smaller suppliers or rural jurisdictions. If optimization favors certain aggregate sources or processing methods, it may concentrate market power. If recycled aggregate specifications are tightened without support for smaller plants, some producers may be excluded. Fairness also concerns environmental justice. Recycling facilities and asphalt plants can impose local air quality and noise burdens, while the benefits of sustainable pavements are distributed across road users. AI optimization that focuses narrowly on carbon or cost may ignore these distributional effects. A fair system would include stakeholder engagement, accessible reporting, and mechanisms for affected communities to participate in decisions about siting and operations.

Policy implications follow from these concerns. Governments should fund open benchmark datasets, shared testing protocols, and independent validation. They should support pre-competitive research on aggregate characterization and mixture performance. They should develop standards for model transparency, uncertainty reporting, and auditability. They should also protect against lock-in by requiring interoperability and data portability. In addition, policy should recognize that optimization is not a substitute for public investment in maintenance. A pavement that is well-designed but under-maintained will still deteriorate. Sustainable asphalt mixture design must be part of a broader asset management strategy that includes timely preservation, drainage, and rehabilitation.

Planning theory reminds us that infrastructure decisions are often wicked problems, meaning they involve contested goals, incomplete information, and irreversible consequences [19]. AI can help organize information, but it cannot eliminate disagreement about values. The choice of whether to prioritize recycled content, cost, durability, or emissions reduction is political and ethical. A legitimate optimization process should make these choices explicit and subject them to deliberation. It should also allow for course correction. If monitoring reveals unexpected performance, the system should support learning rather than blame. This requires institutions that reward transparency and continuous improvement, not just compliance with predetermined specifications.

7. Case Illustrations and Cross-Domain Comparisons

Several case illustrations can clarify how AI-driven gradation optimization might work in practice. In one scenario, a state transportation agency partners with regional asphalt plants to create a shared data platform. Aggregate images are captured at each plant, and gradation is measured automatically. Mixture performance data from laboratory tests and field monitoring are uploaded to the platform. A machine learning model predicts rutting and moisture susceptibility for candidate gradations. Engineers use the model to explore blends that increase recycled content while maintaining performance. The model is retrained quarterly with new data. An audit trail records recommendations, decisions, and outcomes. Over time, the agency identifies which aggregate sources and gradations perform best under different climate and traffic conditions. This scenario requires governance agreements, data standards, and trust among partners.

In another scenario, a small contractor lacks advanced imaging equipment. The contractor uses a mobile application that photographs aggregate piles and sends images to a cloud service. The service returns an estimated gradation and suggested blend adjustments. The contractor can accept, modify, or reject the recommendation. Because the contractor's data are limited, the model relies on a regional prior. This raises concerns about accuracy and dependency on external vendors. A public agency could mitigate these risks by providing a

shared service or by certifying third-party tools. It could also require that contractors retain the ability to export their data and switch providers. Such policies support fairness and resilience in the market.

Cross-domain comparisons offer additional insights. Supply chain optimization has long dealt with multi-objective trade-offs, supplier variability, and demand uncertainty. AI systems in logistics use real-time data to reroute shipments and manage inventories, but they also require exception handling and human oversight. Similarly, gradation optimization must handle feedstock variability and construction disruptions. Commons governance research shows that shared resources can be managed sustainably when rules are participatory, monitoring is transparent, and sanctions are graduated [16]. This suggests that data-sharing consortia for asphalt materials should be designed with clear rules, inclusive membership, and dispute resolution mechanisms. Infrastructure studies show that systems succeed when they align with existing routines and power structures [17]. Therefore, deployment strategies should begin with pilot projects, co-design with end users, and iterative refinement.

Case-study methodology is relevant here because infrastructure innovations are context-dependent. A single case cannot establish universal best practice, but careful case selection and process tracing can reveal causal mechanisms and implementation challenges [20]. Comparative case studies across regions with different climates, materials, and governance arrangements can identify which factors are essential and which are contingent. For AI-driven gradation optimization, useful cases might include a high-recycled-content highway project, a rural low-volume road program, and a dense urban corridor with strict environmental regulations. Each case would examine data availability, model performance, decision processes, and outcomes. The goal is not to produce a universal algorithm but to build a repertoire of validated practices that can be adapted locally.

8. Discussion and Forward-Looking Perspectives

The future of AI-driven gradation optimization depends on how well technical advances are integrated with institutional design. Several promising directions merit attention. First, digital twins of asphalt plants and pavements could provide continuous feedback between design assumptions and observed performance. A digital twin would combine sensor data, material models, and maintenance records to simulate how gradation choices affect long-term behavior. This could support adaptive specifications that evolve as evidence accumulates. Second, advances in computer vision and multimodal learning could integrate aggregate images, spectral data, and mechanical tests into richer representations of material quality. Third, reinforcement learning could optimize blending decisions over time, although its use in safety-critical infrastructure would require strict constraints and validation.

However, these opportunities come with risks. More complex models can be harder to explain and validate. Greater data collection can raise privacy and security concerns, even if the data are not personal. Automation can deskill the workforce if engineers and technicians are not involved in model development and interpretation. Vendor lock-in can stifle innovation and increase costs. To avoid these pitfalls, organizations should invest in data literacy, model governance, and independent evaluation. They should treat AI as a tool for augmenting expertise rather than replacing it. They should also maintain the ability to operate without AI if necessary, ensuring that basic testing and design capabilities remain strong.

From a sustainability perspective, the ultimate test of AI-driven gradation optimization is whether it reduces environmental burdens without shifting them elsewhere. This requires

lifecycle thinking that includes material extraction, processing, transport, construction, maintenance, and end-of-life. It also requires social and economic assessment. A mixture that is environmentally beneficial but unaffordable for local agencies may not be sustainable in practice. A mixture that is cost-effective but harmful to nearby communities may not be just. AI can help evaluate these dimensions, but only if the objectives and constraints reflect them. The optimization process should therefore be participatory and transparent, with stakeholders involved in defining what counts as success.

A forward-looking research agenda should address several gaps. There is a need for open datasets that link aggregate properties, gradation, mixture design, construction, and performance across diverse conditions. There is a need for validation studies that compare AI recommendations with conventional design methods over multiple project cycles. There is a need for governance models that specify data ownership, liability, and access. There is a need for fairness metrics that assess how optimization affects different suppliers, regions, and communities. There is also a need for educational programs that prepare engineers to work with AI systems critically and ethically. Addressing these gaps will require collaboration among universities, agencies, industry, and communities.

9. Conclusion

AI-driven optimization of recycled aggregate gradation has the potential to improve sustainable asphalt mixture design by integrating material characterization, performance prediction, and multi-objective search. It can help agencies and producers use recycled materials more effectively while managing risk. However, the value of AI depends on how it is embedded in technical and institutional systems. The optimization problem is not only about finding a better gradation curve; it is about building an infrastructure of data, trust, governance, and accountability. The analysis in this paper emphasizes that recycled aggregate gradation involves structural trade-offs among performance, durability, cost, constructability, and environmental impact. It also involves fairness, robustness, and policy choices that cannot be resolved by algorithms alone. The most promising path is hybrid intelligence, where AI augments human expertise, makes trade-offs explicit, and supports continuous learning. Public agencies should proceed incrementally, with pilot projects, transparent validation, and inclusive governance. If designed well, AI-driven gradation optimization can contribute to a more circular, resilient, and equitable pavement system. If designed poorly, it may automate bias, obscure uncertainty, and undermine trust. The difference lies in treating optimization as a socio-technical capability rather than a purely computational solution.

References

1. Huang, Y., Bird, R., & Heidrich, O. (2007). A review of the use of recycled solid waste materials in asphalt pavements. *Resources, Conservation and Recycling*, 52(1), 58–73.
2. Santero, N. J., Masanet, E., & Horvath, A. (2011). Life-cycle assessment of pavements. Part I: Critical review. *Resources, Conservation and Recycling*, 55(9–10), 801–809.
3. Van Dam, T. J., Harvey, J. T., Muench, S. T., Smith, K. D., Snyder, M. B., Al-Qadi, I. L., ... Kendall, A. (2015). Towards sustainable pavement systems: A reference document (FHWA-HIF-15-002). Federal Highway Administration.
4. Read, J., & Whiteoak, D. (2003). *The Shell Bitumen Handbook* (5th ed.). Thomas Telford.

5. 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.
6. Masad, E., Muhunthan, B., Shashidhar, N., & Harman, T. (1999). Internal structure characterization of asphalt concrete using image analysis. *Journal of Computing in Civil Engineering*, 13(2), 88–95.
7. Roberts, F. L., Kandhal, P. S., Brown, E. R., Lee, D.-Y., & Kennedy, T. W. (1996). *Hot mix asphalt materials, mixture design, and construction* (2nd ed.). NAPA Research and Education Foundation.
8. Christensen, D. W., & Bonaquist, R. F. (2006). Volumetric requirements for Superpave mix design (NCHRP Report 567). Transportation Research Board.
9. Kim, Y. R. (2009). *Modeling of asphalt concrete*. McGraw-Hill.
10. Airey, G. D. (2003). Rheological properties of styrene butadiene styrene polymer modified road bitumens. *Fuel*, 82(14), 1709–1719.
11. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
12. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
13. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260.
14. Gopalakrishnan, K., Khaitan, S. K., Choudhary, A., & Agrawal, A. (2017). Deep convolutional neural networks with transfer learning for computer vision-based data-driven pavement distress detection. *Construction and Building Materials*, 157, 322–330.
15. Ceylan, H., Bayrak, M. B., & Gopalakrishnan, K. (2009). Neural networks applications in pavement engineering: A recent survey. *International Journal of Pavement Research and Technology*, 2(6), 249–259.
16. Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
17. Star, S. L., & Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1), 111–134.
18. Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. Sage.
19. Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4(2), 155–169.
20. Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12(2), 219–245.