





AI IN

BY CODY A. PENNETTI, PH.D., P.E., AND MICHAEL D. PORTER, PH.D.

Artificial intelligence offers unprecedented possibilities for transportation engineering. But engineers must also understand the technology's limitations.

PHOTOGRAPH COURTESY OF ISTOCK.COM/HALLOJULIE

Artificial intelligence is set to significantly transform the field of transportation engineering, offering extraordinary opportunities to enhance efficiency, safety, and sustainability. Advanced AI algorithms can optimize traffic flow, monitor infrastructure systems in real time, and swiftly respond to or predict incidents, streamlining transportation management and operations.

The integration of AI into transportation systems requires a rigorous and nuanced

understanding of its exciting possibilities and inherent limitations. Remember, though, that AI is not a panacea; it is a tool that depends on the quality and comprehensiveness of the data it processes. As such, effective implementation of AI relies on the expertise of the users to comprehend and interpret its outputs.

But first, civil engineers must understand the foundational mechanics of AI to responsibly implement these systems.

AI can continuously monitor the “health” of transportation systems, recommend actions, or even automate functions.

IMAGE COURTESY OF ISTOCK.COM/PHUTTAPHAT TIPSANA

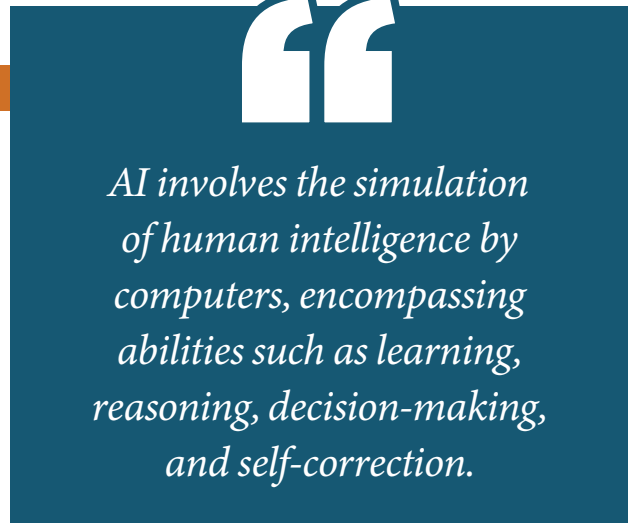


AI POTENTIAL

To understand the effective application of AI to transportation systems, we can think of a wearable health monitor that tracks various metrics such as heart rate, blood pressure, movement, and oxygen levels. While the raw data generated by wearable fitness trackers can be overwhelming and not immediately informative, an AI system trained on comprehensive historical health data can uncover subtle patterns and provide personalized insights, even in real time. This AI training process requires healthcare professionals to guide the development and interpretation of these systems.

Similarly, AI has the potential to continuously monitor the “health” of transportation systems, recommending actions or even automating functions to optimize the system. To succeed, however, this effort requires a combination of appropriate data, well-designed algorithms, and experts who understand the nuances of transportation infrastructure.

For civil engineers, this means learning about AI systems and using the extensive data collected from road sensors, bridges, vehicles, cameras, public transportation networks, and commuters to train such systems on how to operate in a practical, safe, and ethical manner. These AI systems can then suggest



long-term planning solutions, adjust traffic signals to prevent congestion, identify effective safety improvements, monitor bridge deterioration, and immediately dispatch emergency services when hazards are detected or predicted.

In addition to discussing the development of AI tools and how they function, this article will highlight the work of several transportation agencies — in New York, Delaware, and Texas — that have already demonstrated the potential of AI to improve operations.

AI EVOLUTION

While we often think of AI as a type of software application, it is more accurately seen as a field of data science that encompasses various subfields. Broadly, AI involves the simulation of human intelligence by computers, encompassing abilities such as learning, reasoning, decision-making, and self-correction.

Although various forms of AI have existed for decades, earlier systems relied on curated rules — set by human experts — that operated more like a complex series of if-then statements. Recent advances have focused attention on systems that learn from data to identify patterns and improve their performance over time, leading to more complex and adaptable applications.

One of the most significant of these applications is known as machine learning, which is currently used by most AI systems. Such systems learn from historical data and apply statistical models and algorithms to identify patterns and improve their performance over time without significant human intervention. This shift has enabled more complex and adaptable AI applications.

Some of the most familiar AI applications, such as ChatGPT and Gemini, are powered by large language models trained using machine learning techniques. These models fall under the umbrella of generative AI, which means they are capable of handling sophisticated tasks such as writing text in response to conversational prompts.

HOW MACHINES LEARN

As the foundational technology behind many advanced AI applications, machine learning focuses on developing models capable of leveraging data patterns to perform various tasks. The training process involves providing the model with examples and adjusting parameters to minimize errors and enhance performance. To achieve optimal results, successful machine learning applications often require balancing model

complexity in response to the dataset size and available computational resources.

Engineers must recognize that the effectiveness of an AI system hinges on the quality of the training data and the selected AI learning process. The data must be comprehensive, representative, and ethically sourced. Neglecting diverse data inputs can

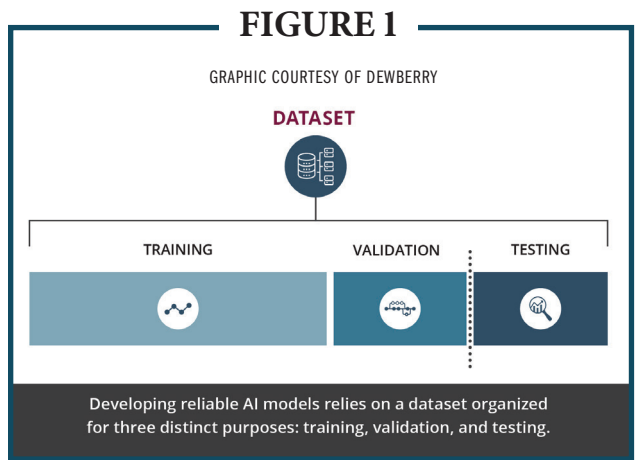
lead to biased AI outputs. For instance, an AI system trained only on data from high-traffic urban areas will likely fail to address the needs of rural transportation networks.

To build a reliable AI system for transportation, the data should be carefully organized for three distinct purposes: training, validation, and testing (see Figure 1).

◆ **Training:** This is the core dataset that the AI model learns from. It should be extensive, diverse, and representative of the real-world scenario the model will be used for. XGBoost, a machine learning algorithm, applies recursive binary partitioning and training data to build a collection of decision trees that establish classifications based on similarities. The model refines the decision trees iteratively to improve its accuracy at predicting or classifying new data, such as which visual characteristics on bridges might indicate deterioration.

◆ **Validation:** This subset serves as a quality check during training. It is a separate set of the data used to fine-tune the model's hyperparameters (settings that control its learning process). This helps prevent overfitting, which happens when the model performs well on training data but poorly on unseen data. For instance, a traffic congestion prediction model might overfit if it uses only data from similar road segments, such as city streets. Validation data ensure that the model can generalize its learnings to different situations and datasets.

◆ **Testing:** Finally, new data are used to evaluate the final AI model. This step is crucial to assess how well the model performs in real-world scenarios, such as real-time traffic forecasting. Testing confirms the model's accuracy and reliability and helps identify any potential issues before deployment.



Given high-quality training data, AI models can leverage various learning paradigms to extract knowledge and generate outputs that guide actions. In recent decades, three categories have emerged as AI's core machine learning paradigms: supervised learning, unsupervised learning, and reinforcement learning. Each offers unique benefits and challenges based on the potential applications.

Supervised learning is an approach for making accurate predictions for a particular task, but it requires labeled data, which means data in which both the inputs and the desired outputs are known. Unsupervised learning is excellent for discovering hidden patterns in data, but it needs expert interpretation to make the results actionable.

Reinforcement learning is great at optimizing complex operations, such as a city-wide traffic signal system, but it requires accurate simulation data and a clear reward structure, such as earning points for the number of vehicles that move through a signalized intersection simulation. The AI system learns which actions result in the best score.

SUPERVISED: LEARNING BY EXAMPLE

For a supervised learning scenario, a dataset might include images of bridge cracks with varying severity (inputs) and corresponding human expert-generated labels indicating the bridge's structural health (outputs). The AI learns to associate the crack characteristics with the bridge's condition based on the patterns found in these labeled data (see Figure 2).

Suppose you want to predict the structural condition of a bridge based on visual data collected for the structure. The inputs could include bridge images categorized according to the condition assessment (such as good, fair, poor) based on engineering assessments by professionals. The AI system would then be trained using these labeled data, in which the bridge's condition is known for each image. The AI learns

the relationships between the visual features and the bridge's structural integrity using image recognition models such as convolutional neural networks.

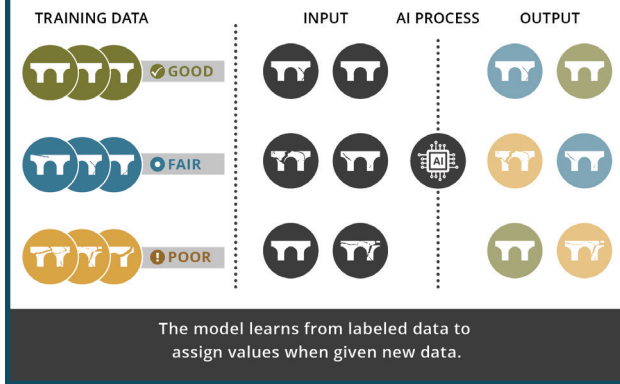
Given a new set of bridge images, the AI could predict the severity of the damage and classify the bridge's condition, which will be the outputs. The system has now learned to assess a bridge's condition by analyzing images of cracks and deterioration. In the future, it can then identify features in the cracks that indicate potential structural problems.

To realize the benefits of supervised learning, AI models require a large dataset of labeled data (such as images of bridges and their structural health). A diverse training set will also

FIGURE 2

GRAPHIC COURTESY OF DEWBERRY

SUPERVISED LEARNING



improve the application. For example, the training data could include images with different lighting conditions, crack types, and materials. Although these training data can sometimes be expensive and time-consuming to collect, once the model is properly trained, it can automate a well-defined task to expedite routine operations.

Several transportation agencies have already implemented AI methods to evaluate traffic camera footage using supervised learning techniques. In New York City, for example, computer vision, a form of AI, was used with traffic cameras facing crosswalks or sidewalks “to develop a scalable data acquisi-

tion framework to collect large-scale pedestrian data automatically and continuously from sampled streaming footage,” according to the case study, “Leveraging Existing Infrastructure and Computer Vision for Pedestrian Detection.”

The case study focused on social distancing during the COVID-19 pandemic and was conducted by researchers at Connected Cities for Smart Mobility towards Accessible and Reliable Transportation, a U.S. Department of Transportation university transportation center at New York University’s Tandon School of Engineering. The study is highlighted on the U.S. DOT’s website for the Intelligent Transportation Systems Joint Program Office.

The technology in the case study was able to approximate the distance between pedestrians, which could then be “scaled up to calculate distances under different environmental conditions (e.g., cameras with different angles),” per the case study.

Because the system can “monitor pedestrian density and distribution as well as temporal variations in behavior,” the information collected can help urban planners and engineers better analyze new pedestrian and mobility patterns, the case study noted. The technology can also be applied to other situations, including the “detection of parking and bus lane occupancy, detection of on-street illegal and double parking, and usage of pedestrian density information at bus stops to assess transit demand.”



AI-powered computer vision combined with traffic cameras can collect large-scale pedestrian data.

PHOTOGRAPH COURTESY OF GETTYIMAGES.COM/DRAZEN_

UNSUPERVISED: PROMOTING DISCOVERY

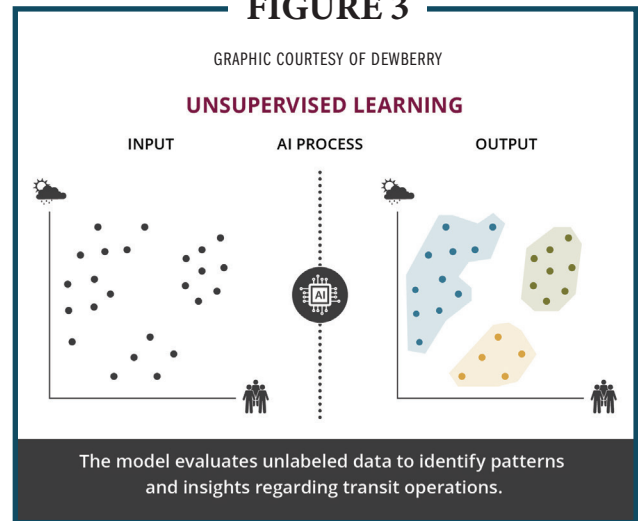
The unsupervised learning approach involves training AI on data without any specific output targets, meaning the data are unlabeled, without input-output pairings. The AI independently seeks patterns and relationships within the data that will help users better understand the data and data-generating process. For example, an unsupervised learning algorithm applied to public transit arrival data may identify many useful patterns without being told specifically what to look for (see Figure 3).

In this scenario, the unlabeled data inputs include scheduled arrival and actual arrival times for various transit stop locations, as well as the departure time, date, and passenger volume. The AI system analyzes these data to find inherent patterns and relationships. It may begin to associate peak usage times and anomalies in performance or cluster different transit stops based on ridership demand patterns.

The system's outputs would then cluster the data into different groups and identify unusual performance conditions that indicate the system's operational health. The AI could identify certain peak times and days that will help transportation engineers set schedule strategies and ideal maintenance times and identify precursors to potential technical issues.

Since the input data are unlabeled, the patterns identified by the AI may sometimes only be meaningful or valuable with additional expert interpretation. Clustering does not imply that the AI assigns a meaningful label; it only estimates that there are shared characteristics in that set of observations that could be meaningful to an informed engineer. Still, the

FIGURE 3



AI output could bring attention to complex relationships that inform transportation operations.

The Delaware Department of Transportation is developing an AI-based transportation operations and management system using unsupervised learning to improve how traffic flows are predicted. The system identifies anomalies and inefficiencies, generating, evaluating, and executing response solutions to existing and predicted traffic congestion.

“It will be a system that goes beyond notifying technicians and the traveling public of congestion to proactively make changes that enhance mobility while continuing to evolve over time,” according to DelDOT’s January 2024 final report on the project.



AI systems might try different phasing policies for traffic signals to minimize congestion.

PHOTOGRAPH COURTESY OF ISTOCK.COM/ANYABERKUT

Reinforcement learning involves training AI through rewards and penalties. In football, imagine the “reward” as the points earned for scoring a touchdown. Given a set of scenarios, the AI then tries many strategies before it learns which plays (actions) are most effective at scoring points (rewards) in a given scenario. The AI learns by trying different actions and receiving feedback based on the outcomes.

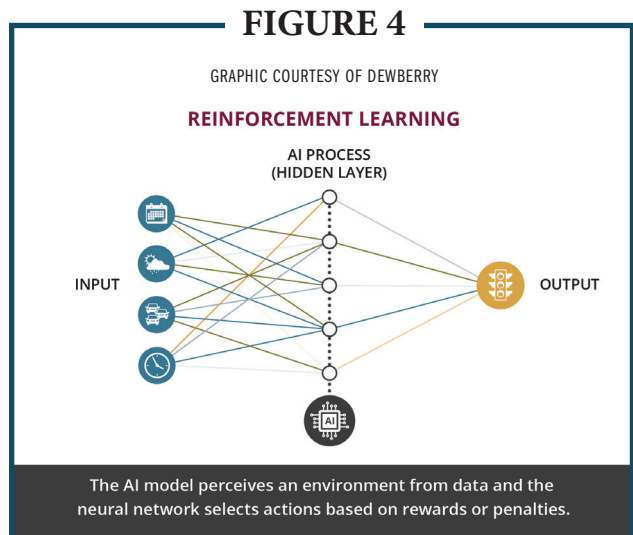
For example, an AI system for traffic signal operations might try different phasing policies to determine which are most effective at minimizing congestion. Reinforcement learning assigns rewards for the actions that result in better outcomes (less congestion) and penalties for poorer outcomes (such as crashes). (See Figure 4.)

In a traffic signal timing scenario, the inputs might involve traffic data such as vehicle counts, wait times at intersections, and traffic flow rates. The system could set an operational parameter, such as minimizing wait times for the traffic signal while adhering to acceptable signal phase transitions. The AI system would then learn through a reward and penalty system by trying different traffic signal timings and observing the outcomes. It adjusts its actions to maximize rewards based on performance metrics (such as short delays per vehicle and short queue length) and minimize penalties.

The system’s outputs might iteratively adjust the signal timings to improve performance, receiving rewards for improvements causing shorter delay times and penalties for longer delay times. Over time, the AI would learn the most effective traffic signal patterns for different times of day and traffic volumes.

Engineers need to know, however, that reinforcement learning requires considerable realistic simulation data — perhaps terabytes of text data and hundreds of billions of written words — and an expectation of continuous learning for real-time applications, which can be computationally intensive and complex to manage. The system also needs a clear and consistent reward structure, such as points for each vehicle moved through an intersection, to learn effectively. Otherwise, the AI will identify ways to exploit a poorly defined training structure. The system might “cheat,” for example, by repeating the same task multiple times to maximize the number of reward points earned.

The Texas Department of Transportation is evaluating the use of reinforcement learning algorithms to dynamically



Engineers must recognize that the effectiveness of an AI system hinges on the quality of the training data and the selected AI learning process.

adjust signal timing, especially to respond to unexpected traffic incidents. To train the signal timing model, the AI relies on a reward function measured by the queue length, learning to find which traffic signal actions reduce the queue length for all network links.

Using simulation data, the AI model can outperform traditional pre-timed traffic signals during both the expected traffic operations and random disrupted conditions. “The incident generation module creates a realistic learning environment for (reinforcement learning) agents, resulting in improved system performance and reduced congestion,” according to the project’s September 2023 final report.

CHATTING WITH AI

Generative AI — including generative pre-trained transformers, or GPTs — is one of the most familiar, accessible, and valuable AI tools for a variety of applications. But engineers need to remember that a GPT operates on the data it was trained on, which may not always be disclosed to the users or readily interpretable by them. So, engineers must understand that making inquiries or seeking applications that differ from the intended use of the GPT system could result in inaccurate outputs — called hallucinations, gibberish, garbage, and other equally undesirable terms.

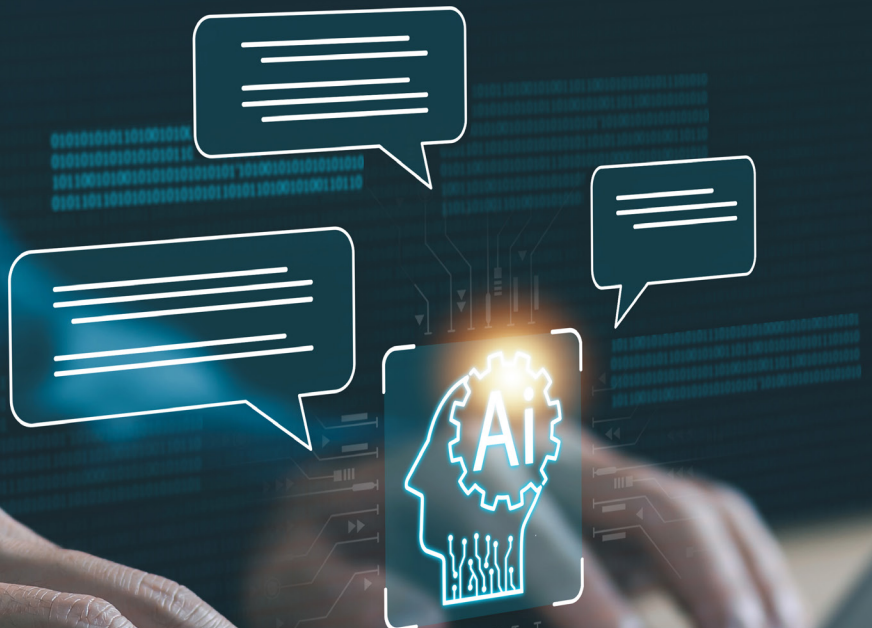
Initially, GPT models undergo unsupervised learning, absorbing language structure and context from vast textual data without specific guidance. This complex training process enables the model to identify and replicate complex language patterns autonomously.

Following this, supervised fine-tuning adjusts the model's responses for specific tasks, enhancing relevance and effectiveness. For example, ChatGPT-3, released to the public in June 2020, uses a sophisticated type of deep neural network known as a transformer, which contains more than 175 billion tunable parameters. These parameters are fine-tuned through training to help the system accurately predict the next word in a sentence. By using hundreds of terabytes of text available on the internet, GPT systems can develop an incredible understanding of grammar and facts about the world, and even some reasoning abilities.

Because the quality of these training data is crucial, engineers and other decision-makers must rigorously audit GPT systems to detect biases and ensure diverse data representation. Such scrutiny is necessary to prevent misinformed decisions and to ensure that the system serves the public good, adhering to current ethical and societal standards.

To successfully implement AI, civil engineers must gain skills in data science and AI model interpretation.

PHOTOGRAPH COURTESY OF ISTOCK.COM/SUCHAT LONGTHARA




TRANSPORTATION APPLICATIONS

For civil engineering applications, GPT systems can analyze datasets such as community surveys, public meeting transcripts, transportation performance data, and project records. By leveraging unsupervised and supervised learning, the model can uncover complex relationships between factors such as public opinions, transportation conditions, and project outcomes. The resulting outputs could inform policy decisions and offer insights through conversational interfaces, enabling stakeholders to quickly access critical information about transportation projects.

There is considerable potential for GPTs tailored to civil engineering applications. Unlike standard AI models that rely on broad internet searches, custom GPTs can be trained on curated datasets defined by the user, allowing them to focus on relevant projects and policies. This customized approach empowers stakeholders to then “converse” with the transportation system, using natural language to rapidly extract insights from these tailored datasets, such as:

- ◆ What are the community’s greatest concerns about the new project?
- ◆ How do different demographic groups perceive transportation services?



The potential of AI offers an exciting future, one fueled by the unsatiated curiosity of the engineers.

- ◆ What are the common causes of maintenance issues in specific types of infrastructure?
- ◆ Are there any patterns in crash severity, locations, or timing?
- ◆ What construction practices have resulted in the lowest financial and safety risks?

While AI can provide valuable insights and automate routine tasks, its integration into civil engineering requires careful consideration, ongoing learning, and robust collaboration between engineers and data scientists to maximize benefits and minimize risks. AI does not replace engineering judgment.

PREPARING FOR THE AI DISRUPTION

Over the years, the civil engineering field has dealt with various disruptive technologies, including computer-aided design, geographic information systems, building information modeling, and virtual design and construction, each bringing about significant changes to project planning and operations. AI, however, represents a disruption that extends beyond core engineering production processes and services.

Unlike adopting new software, using AI represents a tectonic shift comparable to the introduction of personal computers or the internet, which were already well-established when many in the current engineering workforce began their careers. Compounding this unfamiliar experience, the pace of commercialization of AI systems is unprecedented. Although the core machine learning paradigms have existed for decades, AI’s recent rapid acceleration has caused impulsive responses — some overly trusting, others excessively cautious.

To successfully implement AI, civil engineers must gain skills in data science and AI model interpretation, while organizations need to foster a culture of innovation that embraces continuous learning and experimentation with AI technologies.

In this way, the civil engineering industry can navigate this technological shift and leverage AI to create smarter, more efficient, and more sustainable transportation systems.

The optimal approach to integrate AI within civil engineering firms remains an open question. The answer will likely vary depending on the size, resources, and specific needs of each firm. But the potential of AI offers an exciting future, one fueled by the unsatiated curiosity of the engineers. **CE**



CODY A. PENNETTI, PH.D., P.E.
senior associate/Dewberry



MICHAEL D. PORTER, PH.D.
associate professor/University of Virginia