



## Artificial and Human Intelligence with Digital Twins

### Authors:

**Michael Thomas**

Senior Systems Architect

SAS

[michael.thomas@sas.com](mailto:michael.thomas@sas.com)

**Brad Klenz**

Distinguished Systems Architect

SAS

[brad.klenz@sas.com](mailto:brad.klenz@sas.com)

**Prairie Rose Goodwin**

Senior Product Developer

SAS

[PrairieRose.Goodwin@sas.com](mailto:PrairieRose.Goodwin@sas.com)

Human intelligence has been creating and maintaining complex systems since the beginnings of civilizations. One long-lived example is the Qanat in Gonabad, Iran, which was created 2,700 years ago and still provides drinking and agricultural water for around 40,000 people today.<sup>1</sup> In modern times, digital twins have emerged to aid operations of complex systems, as well as improve design and production. In parallel, both Artificial Intelligence (AI) and eXtended Reality (XR)—including Augmented Reality (AR) and Virtual Reality (VR)—have emerged as tools that can help manage the operations of complex systems. This article explores how digital twins can be enhanced with AI and how emerging UI technologies like XR

can improve people's abilities to manage complex systems via digital twins.

Digital twins can marry human and artificial intelligence to produce something far greater by creating a usable representation of complex systems. End users do not need to worry about the formulas that go into machine learning, predictive modeling and artificially intelligent systems but can nonetheless capitalize on their power as an extension of their own knowledge and abilities. Digital twins combined with AR, VR and related technologies provide a framework to overlay intelligent decision making into day-to-day operations, as shown in Figure 1.

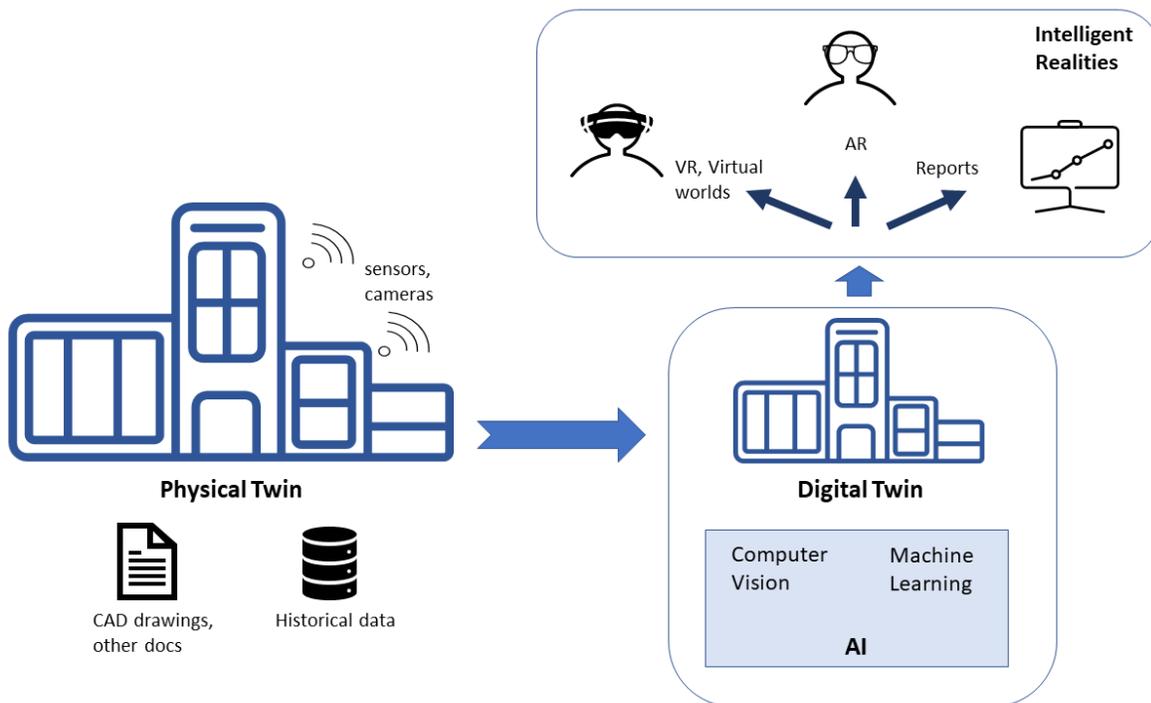


Figure 1: A digital twin enhanced with AI and Intelligent Realities

<sup>1</sup> P. English, "The Origin and Spread of Qanats in the Old World". Proceedings of the American Philosophical Society. [JSTOR](#), 112 (3): 170–181. [JSTOR 986162](#). June 21, 1968.

The operations of a physical twin can be digitized by sensors, cameras and other such devices, but those digital streams are not the only sources of data that can feed the digital twin. In addition to streaming data, accumulated historical data can inform a digital twin. Relevant data could include data not generated from the asset itself, such as weather and business cycle data. Also, CAD drawings and other documentation can help the digital twin provide context. AI and other analytical models can take raw data and process it into forms that help humans understand the system. Additionally, AI can make intelligent choices of content on behalf of the user. Such guidance could be very welcome to users because user input facilities are very different from the typical keyboard and mouse. Finally, as displayed in the upper right corner of Figure 1, humans can perceive the system as an intelligent reality—a technologically enhanced reality that can aid their cognition and judgement.<sup>2</sup>

With the blueprint in Figure 1 as a basis, this article explores how to create digital twins that utilize AI and reality technologies to achieve operational benefits. Any number of operations could be enhanced with the techniques described here. For example, the paper “Augmented Reality (AR) Predictive Maintenance System with Artificial Intelligence (AI) for Industrial Mobile Robot”

details how a machine learning model can be used to classify the state of a robot motor which can then be presented to factory personnel with AR.<sup>3</sup> This article applies the blueprint concepts to facilities management after first exploring each concept in depth. While the various streams of data reach their conclusions in human perception, the starting point of a digital twin for a user is how it is perceived. Thus, the starting point for this exploration are user interfaces for digital twins, followed by a discussion of AI.

### Human Reality of Digital Twins

Humans have a long history of interfacing with data through a canon of data visualization, starting with William Playfair’s inventions of line, bar and pie charts in the late 1700’s. Digital twins can present data in such familiar forms, but the traditions of the late eighteenth century should not unduly restrain the expressive power of a digital twin.

When using mobile technologies such as tablets, smart phones and AR headsets, the digital reality is overlaid directly on the physical reality into a single view, as shown in Figure 2. AR headsets may be the obvious choice for this use case, but it is not the only one. Traditional interfaces rendering 3D models also allow workers to take advantage of digital twins.

---

<sup>2</sup> M. Thomas, “Intelligent Realities For Workers Using Augmented Reality, Virtual Reality and Beyond.” The Industrial Internet Consortium Journal of Innovation, Mar 2019. Available: <https://www.iiconsortium.org/news/joi-articles/2019-March-Intelligent-Realities-For-Workers-Using-Augmented-Reality-Virtual-Reality-and-Beyond.pdf>

<sup>3</sup> Y. Tay et al., “Augmented Reality (AR) Predictive Maintenance System with Artificial Intelligence (AI) for Industrial Mobile Robot.” SAS Global Forum 2019. Available: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3628-2019.pdf>

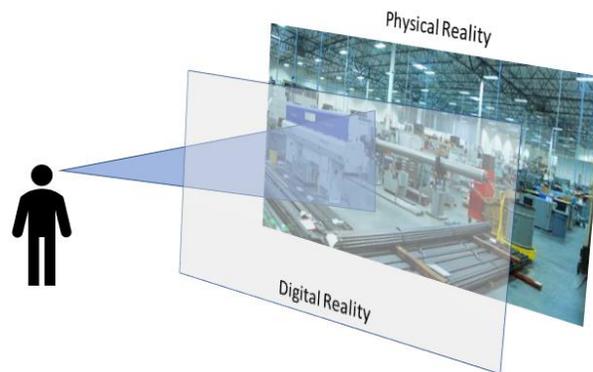


Figure 2: Digital reality produced by digital twin overlays the physical twin

The first step in considering the creation of intelligent realities for digital twins is understanding data visualization options across the User Interface (UI) spectrum. Next, a reporting integration approach is considered which can operationalize analytics and AI without requiring a new hardware paradigm, like an AR headset. AR headsets have the potential to greatly benefit operations but only if applications are successfully designed for usability, which is the next consideration. This section concludes by outlining how to build a digital twin interface for remote experts.

### Visualizing Digital Twin Output Across The UI Spectrum

In Cap Gemini’s “Augmented and Virtual Reality in Operations” report, Jan Pflueger from Audi’s Center of Competence for AR/VR encouraged a business-first approach for reality projects. “First, focus on your use case and not on the technology itself. After you identify your use case, focus on your information handling and data so you can deliver the right information to the technology.”<sup>4</sup>

To that end, the following table considers five different technological approaches for rendering digital twins and their respective capabilities.

---

<sup>4</sup>“Augmented and Virtual Reality in Operations: Investment Guide.” Capgemini Research Institute, 2018. Available: <https://www.capgemini.com/wp-content/uploads/2018/09/AR-VR-in-Operations1.pdf>

	Traditional Desktop	Smart phone or tablet	Monocle AR	Stereoscopic AR, including mixed reality (MR) devices	Fully Immersive VR
2D data visualization	Yes	Yes	Yes	Yes	Yes
3D data visualization	Yes	Yes	Yes	Yes	Yes
Augmenting a digital 3D model of a physical asset	Yes	Yes	Not practical	Yes, and model appears as 3D	Yes, and model appears as 3D
Augmenting a physical asset with charts, heat map overlay, etc. in physical space	Only if PC is in the physical space with the asset	Yes	Yes	Yes	No, but it could augment a virtual rendering
Heads up and hands free in a physical environment	No	No	Yes	Yes	No
3D rendering of a remote asset	Yes	Yes	Not practical or typical	Yes	Yes
Eye and hand tracking	Possible, but not typical	Not practical or typical	Not practical or typical	Yes for some products	Yes for some products

Figure 3: Five technological approaches for rendering digital twins and their capabilities

Within each class of device, capabilities vary, and the variance may significantly affect a product’s viabilities for different use cases. This is especially true for AR headsets. Display resolution, field-of-view and computational power differ from product to product. In addition, design decisions about whether to put battery and compute units on the headset or on a separate tethered module can affect comfort and practicality. One practical concern for AR headsets is how they integrate with work clothing and

uniforms such as those required for clean room and food processing operations.

#### Reporting with a Digital Twin Context

Given an interactive visual analytics application, intelligent reality reports can be created with integrated 3D models like the one shown in Figure 4. In this example, Autodesk Forge is integrated into the reporting interface of SAS Visual Analytics. The digital twin presents a custom visualization that can interact with other

objects in the report, including showing data in a table or graph.

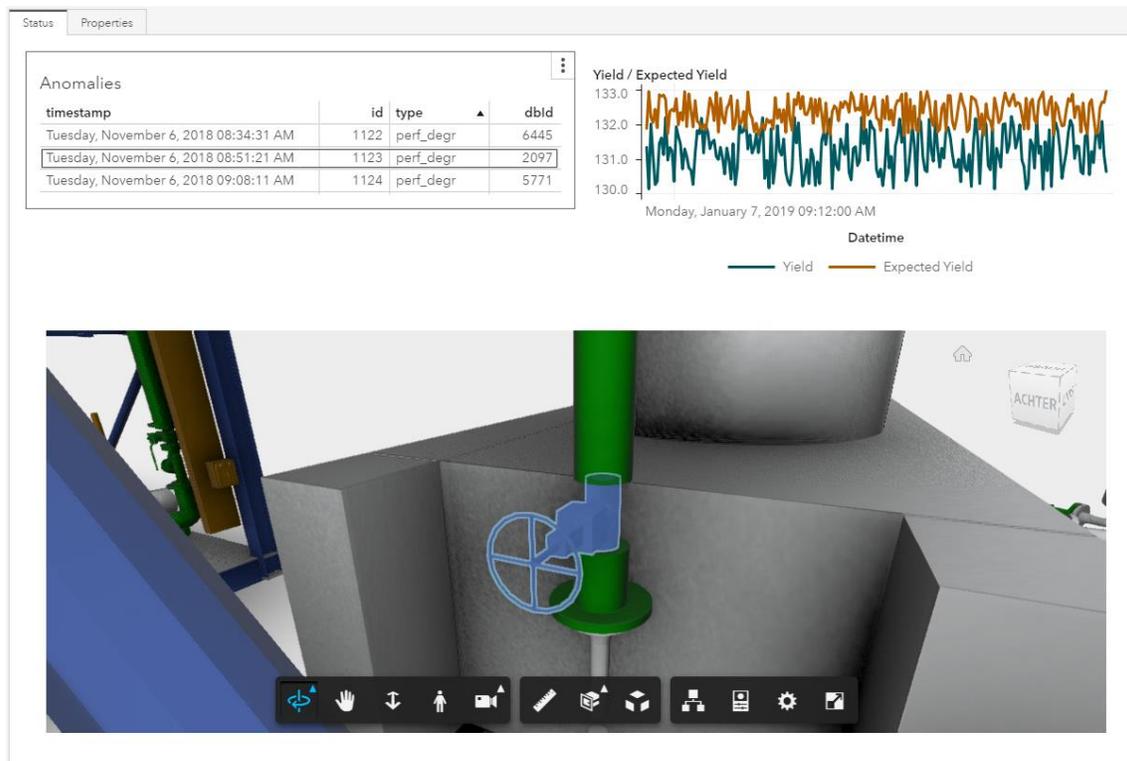


Figure 4: SAS Visual Analytics and Autodesk Forge

This visualization approach adheres to long standing data presentation traditions without requiring new hardware beyond a regular desktop set up. The user interface is presented on a typical computer with a mouse and keyboard, and users need little additional training in order to utilize the power of the digital twin.

### Usability & Augmented Reality

When moving beyond the desktop into AR headsets, application designers face a new set of usability challenges. Usability is the cornerstone for any technology to be a tool rather than a hindrance. While AR is a new interaction paradigm, the long-standing standards of usability still apply. These should guide efforts to integrate AR with a digital twin. A good interface is efficient, learnable, memorable, error-infrequent and

pleasant to use.<sup>5</sup> Leveraging domain knowledge with the advantages of AR and digital twins makes this space uniquely situated to maximize usability in many of these categories.

### Learnability

Users on AR platforms show significant improvements with minimal instruction over a short period of time.<sup>6</sup> Some users find that AR is difficult initially because they cannot rely on their intrinsic knowledge to operate the system. This setback is temporary, however, and users improve rapidly. Learnability varies significantly based on the target audience. Tools made for experts have a higher learning curve but are more powerful overall, and expert efficiency should justify the extended training period.

### Efficiency

AR has the effect of “embodiment” when the technology fades away and becomes an extension of our senses.<sup>7</sup> When this happens, the technology extends our sensory, cognitive and motor limitations, so we spend fewer cognitive resources thinking about the interface and more on the task at hand. The information in the world

combined with information in the program improves the efficiency of a task through embodied cognition.

### Memorability

Users are more engaged with the content when accessing it in an AR system because of the novelty. Engagement is a key factor in thinking critically and remembering details. When a user comes back to the interface after a period of inactivity, they will be more likely to remember both the content and the actions.

### Low error rate

As with any application, the error rate is primarily affected by the design of the interface. A good interface designer will be able to create a user experience well within the limits of human factors, and this applies for AR. While the interaction paradigm is different from point and click, system designers have considered the kinds of inputs that can be reliably recognized and limit the amount of irrecoverable errors during use.

---

<sup>5</sup> Holzinger, A., 2005. Usability engineering methods for software developers. *Communications of the ACM*, 48(1), pp.71-74. Available at

[https://www.researchgate.net/profile/Andreas\\_Holzinger/publication/220422205\\_Usability\\_Engineering\\_Methods\\_For\\_Software\\_Developers/links/5460e30c0cf27487b4526442/Usability-Engineering-Methods-For-Software-Developers.pdf](https://www.researchgate.net/profile/Andreas_Holzinger/publication/220422205_Usability_Engineering_Methods_For_Software_Developers/links/5460e30c0cf27487b4526442/Usability-Engineering-Methods-For-Software-Developers.pdf)

<sup>6</sup> Avery, B., Piekarski, W., Warren, J. and Thomas, B.H., 2006, January. Evaluation of user satisfaction and learnability for outdoor augmented reality gaming. In *Proceedings of the 7th Australasian User interface conference-Volume 50* (pp. 17-24). Australian Computer Society, Inc.. Available at <https://pdfs.semanticscholar.org/f1f3/6cc2d4fe33d9bee7ed838b9455dc0661cfff.pdf>

<sup>7</sup> Tussyadiah, I.P., Jung, T.H. and tom Dieck, M.C., 2018. Embodiment of wearable augmented reality technology in tourism experiences. *Journal of Travel research*, 57(5), pp.597-611. Available at <https://e-space.mmu.ac.uk/618427/1/Embodiment%20of%20Wearable%20Augmented%20Reality%20Technology%20in%20Tourism%20Experiences-Authors%27%20Copy.pdf>

### Satisfaction

Multiple studies have indicated that users prefer AR over traditional interaction paradigms. Satisfaction is the culmination of learnability, efficiency, memorability and the ability to use the system without catastrophic errors. When the other categories of usability are well-balanced, the user will be satisfied with their experience.

### Remote Experts and Distant Twins

An AR device can help a field worker by overlaying their reality with digital twin output. The same AR device can also serve as

a platform for remote expert assistance. The simplest way is to transmit the video feed from the AR device's video cameras to the remote expert; however, the video feed alone would not be able to go beyond what the field worker can physically access and does not contain live IoT sensor data.

Instead of relying completely on video, the remote expert could view the system as a virtual world. A VR or MR headset could be used, but a traditional flat screen would also work well. The picture below shows a digital twin created with the Unity gaming engine.



Figure 5: Digital twin of a retail store

In this kind of VR system, the expert can change to a vantage that the field worker cannot see. For example, they could view the system from any angle, go through locked doors or even go inside of components.

Unlike the field worker's digital reality, the expert's reality must be created. It could be created by a 3D artist or with digital artifacts such as CAD drawings—or some combination of the two. An artist would have full control, but utilizing CAD drawings would be more scalable.

While it is technically possible to convert CAD drawings directly for use in gaming engines, CAD drawings tend to be too detailed for real-time rendering in the gaming engine style. CAD is purposed towards creating models that can be handed to manufacturing or building, while gaming engines pursue photo realism, believable lighting and low latency response to changes in camera position.

Tools exist to optimize CAD drawings for virtual engines. In the blog post “Design, build and operate faster with the PiXYZ Plugin for AEC,” Unity evangelist Kieran Colenutt outlines a way to perform a conversion from CAD to model formats supported by Unity.<sup>8</sup>

### **FROM IOT SENSORS TO ARTIFICIAL INTELLIGENCE**

---

With IoT, data is collected from sensors on a device, on neighboring devices, the

environment around a device and whatever interacts with the device. The speed is real time, and connectivity often allows us to span distances instantly. Advances in streaming analytics now enable us to process this real-time data using machine learning and artificial intelligence.

While very simple systems can be twinned from raw data readings, AI and other analytical techniques are necessary to make a human-consumable digital twin of complex systems. Consider a vehicle. If you see it as only an object on a map, then a digital twin can be very simple and easy to digest. There are only two variables, latitude and longitude, and the variables are easily understood by humans. But if you wish to twin the operations of the vehicle, you are now dealing with hundreds of megabytes of data per second and thousands of variables. While all that data is important for the operations of the vehicle, that much raw data would overwhelm the ability of a human to make sense of it. AI synthesizes the data so that the digital twin can present it in a human consumable format. Conversely, AI enhances the digital twin experience by providing additional information about the environment not otherwise available to the user.

Underneath the umbrella term of AI are several specific categories of machine learning. These are presented in this section as a toolbox for your digital twin. First, the general architectural practices of AI are

---

<sup>8</sup> K. Colenutt, “Design, build and operate faster with the PiXYZ Plugin for AEC,” Unity Blog, Jan 2019. Available: <https://blogs.unity3d.com/2019/01/30/design-build-and-operate-faster-with-the-pixyz-plugin-for-aec/>

presented, and then specific deep learning techniques are reviewed.

### Common Practices of Creating Artificial Intelligence

AI must become intelligent somewhere, and it is usually not “on the job.” Deep learning models are trained on large databases and are almost always done offline. It is not unusual to take hours or days to train a model. Once the model is trained, the application of the model through inferencing is less compute-intensive but still requires more compute resources than is typical for digital twin applications.

For some applications, near real-time or slightly delayed results are sufficient. For example, in the computer vision defect detection described below, it might be acceptable to hold a production batch while the defect detection is performed. In other cases, real-time inferencing is needed. Inferencing can be done in the cloud or data center where sufficient resources are readily available. For edge inferencing, edge

gateways are now becoming available with sufficient compute power, but you must plan for this specialized need.<sup>9 10</sup>

### Recurrent neural networks

Recurrent neural networks (RNNs) are a special class of deep learning neural networks designed for sequence or temporal data. Within IoT and digital twins, there are many examples of such sequence and temporal data. Many sensors are collecting data over time. The sequence or pattern of the measurements over time can be used to understand interesting characteristics of the digital twin asset. One example is measuring energy circuits in a smart building or power grid. The pattern of the energy use on a circuit can capture the start or end of an asset operation such as a motor start, which signals an operation change in the digital twin asset. Another use of RNNs is for forecasting unusual time series data. An example is forecasting the energy output from a solar farm, shown in Figure 6.

---

<sup>9</sup> Klenz, Brad. 2018. “How to Use Streaming Analytics to Create a Real-Time Digital Twin.” *Proceedings of the SAS Global Forum 2018 Conference*. Cary, NC: SAS Institute Inc. Available: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2018/2004-2018.pdf>

<sup>10</sup> Williams, David. 2019. “NVIDIA Graphics Processing Units Accelerating SAS® Analytics.” *Proceedings of the SAS Global Forum 2019 Conference*. Cary, NC: SAS Institute Inc. Available: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3618-2019.pdf>

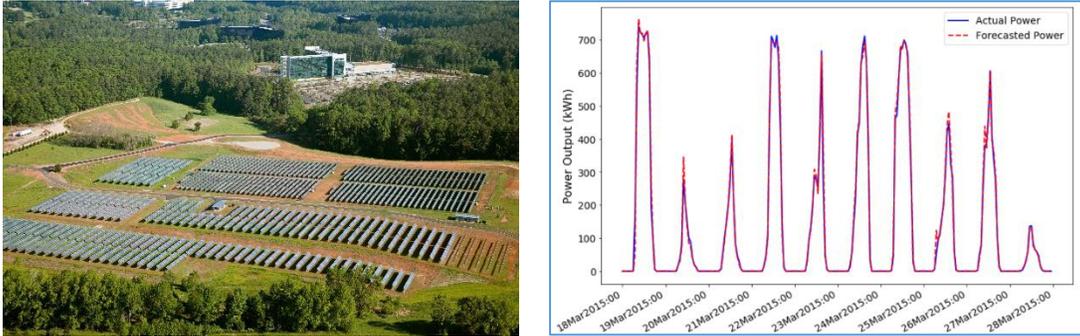


Figure 6: Solar Farm and Power Output Chart

In this case, there is a cyclical component that could be forecasted using traditional methods, but there is a less well modeled component of weather and cloud cover. With the large amount of data available from the solar farm and nearby solar farms, a deep learning RNN can capture the more sporadic aspects of the energy output.<sup>11</sup>

The process for training an RNN is different if you are working with sequence data versus working with temporal data. The process for training the RNN with sequence data is as follows:

- Break the data into segments of sequential measurements. The length of the segment is determined by the time interval of the data and the expected duration of the precursor to an event. For the energy circuit example in smart buildings, the data is collected at 5-second

intervals, and we use the previous minute of data.

- Create a target variable for the events of interest and use it to label the sequences where the event occurs. For our example, we are using motor starts and identifying weak motor starts indicating capacitor failure.
- Train the RNN. Bidirectional model fitting is not needed in this case because measurement data is always moving forward in time.

The trained model can then be deployed for inferencing. In most cases, the model inferencing function will be sufficiently fast to be used on the real-time measurement stream, either in the cloud, server or edge device.

The second type of RNN is used to forecast. The example in this case is to forecast the energy output of a solar farm for short time periods in the future (1 hour). The key in this

<sup>11</sup> Kahler, Susan. 2018. "Using Deep Learning to Forecast Solar Energy." Available: <https://blogs.sas.com/content/subconsciousmusings/2018/07/05/deep-learning-forecasts-solar-power/>

case is to create a set of lagged variables for the predictors and the response variable. The response variable is the energy produced.

To train this RNN, take the historical input database and create lagged variables for the predictors and response variable. The number of lags is determined by the time interval of the measurement data and the expected correlation of previous measurements on the forecast time horizon. For the solar farm example, we are producing one-hour-ahead forecasts, and the data over the last few hours is sufficient to capture the primary effects for the forecast. Note that there are a large variety of conditions possible throughout the year and previously observed weather, even though the forecast horizon is fairly short. Since we have a large amount of historical data of the various conditions, the use of an RNN is appropriate for this problem.

Since training and evaluating the RNN model is dependent on the sequence, partitioning the data requires more care than typical random partitioning. In this case, we need to preserve the sequence of the data for use in the model creation steps (training, validation, test). The easiest way to do this is to partition the data based on the time variable. Use the earliest historical data for the training data set. Then use the next time partition for the validation data set. Finally, use the most recent data for the testing data set. This is sufficient if the performance of the asset has been consistent over the historical data sample. If there have been periods of degraded performance, it is best to eliminate that data from the data sets used to create the model.

You can use RNNs for one-step-ahead forecasting where the forecast interval matches, or is less than, the desired forecast interval. This will yield the most accurate forecast. In some cases, you might need a multistep forecast to project future time periods based on the near term forecast estimates. These forecasts are typically less accurate but can be tested to determine if they have sufficient accuracy.

### Reinforcement learning

Reinforcement learning (RL) is a subfield of machine learning and deals with sequential decision-making in a stochastic environment. In any RL problem, there is at least one agent and an environment. The agent observes the state of the environment and takes and executes a decision. The environment returns a reward and a new state in response to the action. With the new state, the agent takes and executes another action, the environment returns a reward and new state and this procedure continues iteratively. RL algorithms are designed to train an agent through this interaction with the environment, and the goal is maximizing the summation of rewards.

RL has recently received much attention due to its successes in computer games and

robotic applications.<sup>12 13</sup> Beside the simple RL applications, there are still few real-world applications of RL to increase efficiency. We studied and did some research to extend an RL algorithm for controlling the heating, ventilation and air conditioning (HVAC) systems. HVAC includes all the components that are supposed to maintain a certain comfort level in the building. Buildings consume 30% to 40% of all consumed energy in the world, so any improvement could result in huge savings in energy consumption and carbon dioxide release.<sup>14</sup> Advances in new technologies in recent years have improved the efficiency of most components in the HVAC systems. Nevertheless, there are still several directions to reduce the energy consumption by controlling different decisions on these systems.

We considered a multizone HVAC system and selected the amount of air flow as the main control decision. Using the obtained data from an SAS building in Cary, NC, we trained an environment and used it to train an RL algorithm when there are 10 zones in the system with a set-point of 72 with  $\pm 3$  allowance. The figure below shows the results of 50 cases with different initial temperatures. The upper figure is the temperature and the lower figure is taken actions over 150 minutes in which every three minutes a decision is taken. We compared this result to the commonly used rule-based algorithm (in which the system is turned on/off at 69/75), and RL obtained 47% improvement on combination of obtained comfort and energy consumption.

---

<sup>12</sup> Hao, Karen. 2019. "The Rise of Reinforcement Learning" in "We Analyzed 16,625 Papers to Figure Out Where AI Is Headed Next." Available: <https://www.technologyreview.com/s/612768/we-analyzed-16625-papers-to-figure-out-where-ai-is-headed-next/>

<sup>13</sup> Burda, Yuri. 2018. "Reinforcement Learning with Prediction-Based Rewards." Available: <https://blog.openai.com/reinforcement-learning-with-prediction-based-rewards/>

<sup>14</sup> U.S. Department of Energy, 2008, "Energy Efficiency Trends in Residential and Commercial Buildings" Available: [https://www1.eere.energy.gov/buildings/publications/pdfs/corporate/bt\\_stateindustry.pdf](https://www1.eere.energy.gov/buildings/publications/pdfs/corporate/bt_stateindustry.pdf)

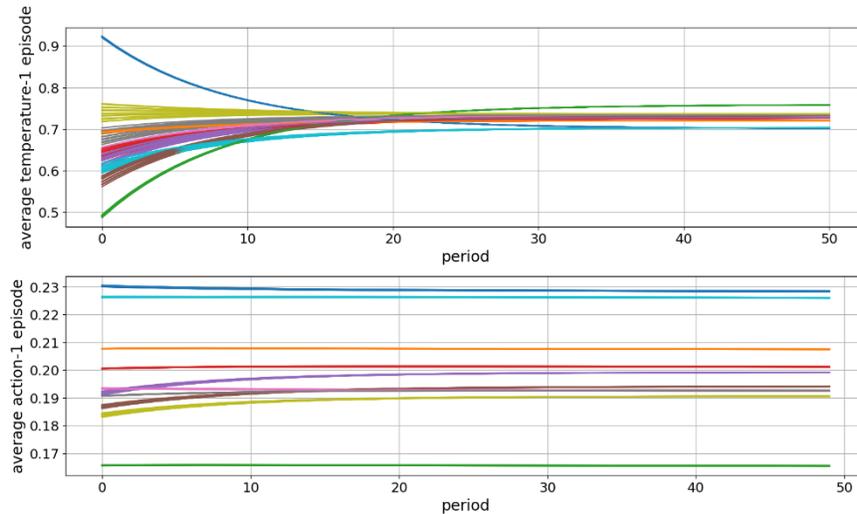


Figure 7: The Average Temperature and Average Action of the RL Algorithm

### Hyperparameter tuning

For all deep learning methods, hyperparameter tuning is an important step. Hyperparameter settings are often dependent on the domain knowledge of the application. Research into the specific application can yield a set of parameter settings to be tested. In some cases, a set of parameter settings has been established as best practices. In other cases, research will be needed to determine the best settings.

One feature in SAS® Visual Data Mining and Machine Learning is hyperparameter autotune. This feature will take a range of potential parameter settings and perform an optimal search for the best performing

settings. This will greatly help cases where research is needed on the parameter settings.<sup>15 16</sup>

### Computer vision

Computer vision is a powerful tool that has caught the attention of many with its ability to recognize faces and objects within a scene. For digital twins, it can add important information about the quality of the things being monitored. A task that requires visual inspection could be enhanced with an AR interface to a digital twin. For example, computer vision can detect defects by comparing thousands of images for anomalies that may not be as easily detected by a human. Moreover, specialized cameras, such as infrared, allow for even further

---

<sup>15</sup> Koch, Patrick, et al. 2018. "Autotune: A Derivative-Free Optimization Framework for Hyperparameter Tuning." Available: <https://www.kdd.org/kdd2018/accepted-papers/view/autotune-a-derivative-free-optimization-framework-for-hyperparameter-tuning>

<sup>16</sup> Koch, Patrick, Brett Wujek, and Oleg Golovidov. 2018. "Managing the Expense of Hyperparameter Autotuning." Proceedings of the SAS Global Forum 2018 Conference. Cary, NC: SAS Institute Inc. Available: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2018/1941-2018.pdf>

analyses by combining multiple streams of information. A process for implementing a computer vision model is as follows.

If possible, fix the camera to a stable mount point so that all images will be taken from the same angle and with the same proportions. This vastly simplifies the model training as compared to general object recognition models, which must capture objects from many angles. The fixed camera location also simplifies the process of determining the location of defects on the piece.

Another option is to initially create a model that finds easily identified features on the piece. For a power substation, you could have general instructions on how to point the camera at a transformer in the substation. An object recognition model could identify the bushings on the top of the transformer. This would provide reference points to scale the images with images captured at similar angles, similar to how facial recognition models determine various key points on a face.<sup>17 18</sup>

You will use the images to create a classification model using Convolutional neural networks (CNN). Depending on how well labeled your data is, you can create models of various complexity.

If you primarily have a collection of known good images, you can create a binary

classification model that identifies images with a high likelihood of known good or suspected anomaly images. The power transformer is an example of this.

If you have images that have been labeled with known defect types, you can create a more complex classification model that identifies the various defects. The discrete parts are an example of this. There might be previous images labeled with an incorrect bearing insertion and other images labeled with incorrect part milling.

If you have good location identification, you can also break down the images and find the portions of the image with defects. The semiconductor wafer is an example here. You can quantify the expected yield based on the proportion of the wafer with defects.

When you have trained the model, you can determine at what latency you can infer and test new images being captured and if you need to stream image-by-image and get immediate results. Alternatively, you might be able to capture a batch of images and process in batch. Also determine if the inferencing can be done in the cloud or server or if an edge gateway is needed.

### **APPLICATIONS TO SMART FACILITIES**

---

Smart facilities offer a perfect example of how a digital twin can offer features that cannot be accomplished in the physical

---

<sup>17</sup>Long, Xindian. 2018. "Understanding Object Detection in Deep Learning." Available: <https://blogs.sas.com/content/subconsciousmusings/2018/11/19/understanding-object-detection-in-deep-learning/>

<sup>18</sup> Long, Xindian. 2019. "Exploring Computer Vision in Deep Learning: Object Detection and Semantic Segmentation." Proceedings of the SAS Global Forum 2019 Conference. Cary, NC: SAS Institute Inc. Available: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3317-2019.pdf>

world. Although buildings are becoming smarter and smarter, we do not yet have walls that turn transparent on command. If a building administrator wants to look through walls, AR and digital twins provide that ability. In this case, the information pipeline goes through the following steps.

### Raw Data

An air-handler can have hundreds of sensors that monitor things like duct pressure, valve-positions, outside air temperature and power draw. Traditionally, these systems are left alone until they break down or need maintenance; however, that approach does not consider how efficiently the system is running and offers little insight into problems or places for improvements. The data streams generated from the sensors can improve maintenance. Other sources of raw data include thermometers, motion detectors and signal monitors for Wi-Fi and other wireless networking.

### AI and Analytical Models

Models can be trained to twin the system and alert administrators when the digital model does not match the physical performance. This approach minimizes downtime and helps pinpoint issues. The end user does not need to know all the math behind the model, but the twin can be crucial to separating important information from the noise. For example, the model can detect when the power draw of the air-handler is more than expected given the outside weather.

### Always Aware Building Management Via Digital Twin and AR

With the raw data processed into intelligence, many options are possible for rendering the digital twin output. Most importantly, the physical and digital realities can be combined for a facility manager such that they can be situationally aware. When AR headsets are utilized, the digital twin can be merged into the physicality of the facility as a manager moves around the facility. Just as a manager can take note of physical flaws and issues, they can use AR to see into walls, make invisible Wi-Fi coverage visible and see temperature differences throughout the facility. While an alert-driven approach based on defined thresholds would remain important, a situationally aware approach lessens the chances that alerts would be a surprise to a manager. Also, a manager can use intuition and judgement to prioritize issues that may fall in blind spots of defined rules.

A situationally aware approach is possible due to the advent of lightweight AR headsets with strong battery life. Without modern AR headsets, the digital reality of the facility is only visible at a desktop or perhaps on a tablet; but even with mobile use of a tablet, the manager would still have to operate the computing device as they move about. This approach is not heads up and hands free. With an AR headset, the digital output is ambient as the manager moves about.

Since a digital twin of a facility can produce a lot of output, it is likely that it could produce many visual representations in the same physical space. A digital twin AR application that attempts to render all possible

information for that space would not be usable. An AI agent can select the most pertinent digital twin output based on several criteria such as the manager's role, newness of information, situational urgency and the manager's history of interest.

### CONCLUSIONS

---

When properly architected and integrated under the Intelligent Realities umbrella, IoT, AI and UI technologies can open new possibilities, and digital twin provides a

usable representation to consume the massive amounts of information inherent in such an architecture. This article worked through various UI options for interacting with digital twins. AR and VR are included, but more traditional options like tablets and desktop computers should also be considered.

- Return to [IIC Journal of Innovation landing page](#) for more articles and past editions

The views expressed in the *IIC Journal of Innovation* are the contributing authors' views and do not necessarily represent the views of their respective employers nor those of the Industrial Internet Consortium.

© 2019 The Industrial Internet Consortium logo is a registered trademark of Object Management Group®. Other logos, products and company names referenced in this publication are property of their respective companies.